Controlled Dynamics of Laser-Cooled Ions in a Penning Trap

Eoin Seymour Gwyn Phillips

December 2004

Thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy of the University of London
and for the Diploma of Imperial College

Abstract

I report on a scheme for improved control over the motion of small numbers of laser-cooled ions in a Penning trap. This work is part of a larger project to assess the suitability of the Penning trap for quantum information processing (QIP). Laser cooling of ions in this type of trap is complicated by large Zeeman shifts of the cooling- and repumping-transition wavelengths from those in zero magnetic field. Lasers can not, therefore, be tuned to the unshifted atomic spectral lines using an external reference such as a hollow-cathode lamp. In order to address this difficulty I built a wavemeter with a precision of one part in 10^7 for direct tuning of the lasers to the transition wavelengths. It was also necessary to improve the laser-frequency stabilities by building reference cavities to which they could be locked. In evaluating the suitability of the Penning trap for QIP, one must address the poor localisation of the ions due to the difficulty in cooling one of the radial motions: the magnetron motion. I have implemented a technique called axialisation that reduces this motion's amplitude by coupling it to the other radial motion: the modified cyclotron motion, which is efficiently laser-cooled. This was done with small clouds of magnesium ions and later with small clouds of calcium ions. The damping rate of both radial motions was measured using near-resonant motional excitation and an RF–photon correlation technique. The coupling rate could be inferred from a shift of the motional frequency in a classical analogue of a quantum avoided crossing. Increases in damping rates of the magnetron motion of more than an order of magnitude were observed for both calcium and magnesium ions and coupling rates of order 1 kHz were measured.

Do mo thuismitheoirí

Contents

CONTENTS

List of Figures

List of Tables

Acknowledgements

Firstly, I would like to thank my supervisors, Dr Danny Segal and Prof. Richard Thompson for all of their support, encouragement and friendship. For financial support, I am very grateful to the European Commission for a generous bursary and travel grant awarded through the QUEST Research and Training Network.

I have been very fortunate to have had colleagues in the lab who have not only been excellent physicists but have also been excellent fun to work with. Thanks to Hebe, Jake, Arno, Tim, Rafael, and Hamid. A special thank you to Kingston, Sebastián, Mike (you really should have learned your lesson after a few years in our group!), and Bakry. Most of all to Rich, fellow zombie in my final year.

Excellent technical assistance was made available in the department. I have benefited greatly from the help and experience of Paul, Martin (×2) and Simon in the optics workshop, Brian in the main workshop, and Bandu in the field of electronics. Karen and Perminder always seemed to know how to fix the finance problems.

Beyond the lab in Imperial College, I have enjoyed the company and assistance of our collaborators in the National Physical Laboratory. Thanks to Peter, Val, and especially Sarika. Thanks also to Dr Stephen Webster, Dr Alastair Sinclair, and Dr Mark Oxborrow.

I spent a very happy three months in Innsbruck in the group of Prof. Rainer Blatt. I was given a great welcome by everyone there. I am especially grateful to Dr Christoph Becher, my supervisor there and Carlos Russo, my co-conspirator who dreamt only of cavity modes. And for showing me the dazzling lights of Innsbruck night-life after a long day in the lab, thanks to Alex, Flic, Christian, and Herbert. I promise I'll learn more German next time!

Two great and very different groups I've been involved with in Imperial were the Chaplaincy and the Physics Society. I've received a lot of support

ranging from the moral to the culinary from the chaplains and fellow students. Thanks especially to Alan, Andrew, Charles, Geoff and Ivor. Thanks also to chaplaincy assistants and other students who made it such a lively, compassionate and thought-provoking group. From the Physics Society there have been dozens I've worked with on the committee. Most of all thanks to my preceding and following chairmen, Libby, Seb, James (how did we manage to be the only club in the black?!), and Brett.

For many coffee- and alcohol-fuelled discussions thanks to John, Alex and Gillian, fellow wanderers on the fun-filled path of PhDs.

Outside of Imperial I've had great support from many friends, amongst them John and PP, neither of whom smoke a pipe yet. I've also had frequent visits from de lads (and lasses) back home: Jon, Stu, Simon, Tim, Eleanor, Gráinne, Catherine, Aisling, Sinéad, Laura and many others. Thanks especially to Jon for all the pies. Thanks also to Stu for showing us what to do in a tight corner.

I was given wonderful hospitality and encouragement while writing up by Linda, David and Jonathan Lawrence for which I'm very thankful.

To my own parents and to my sister, thank you! For all the financial help, words of encouragement and everything else you've done for me, I'm eternally grateful.

And finally, but certainly not least, thank you to Catherine for being there throughout my final year. For making the tough times bearable and every other time wonderful.

Chapter 1

Introduction

The idea that the storage, processing, and retrieval of information is governed by the laws of physics has profound implications both philosophically and technically. It is not only computer designers and other engineers who need to be aware of this, as it affects everyone's daily life.

All speech and cognitive process is governed by some simple laws. Normally, when we think of something governed by laws we see restrictions placed in our way. In the past two decades, however, there has been a realisation that instead of restricting what we can do, the application of some of the laws of physics can lead to radically new ways of thinking about information. The new field that has grown up based on this realisation is called *quantum information theory.*

The marriage of quantum physics and information theory has allowed theorists to devise ways of performing calculations that are simply impossible to perform with any existing computer, no matter how fast. The key to the speed-up is quantum parallelism: Nature allows us to encode many possible starting points in the same memory and process them simultaneously. With some ingenuity of our own we can apply some operations to the system so that we get the answer we desire at the end. Where a normal ("classical") computer would have to look at many individual cases sequentially, a quantum computer could examine all the possibilities simultaneously. The problems that could theoretically be solved with such a computer are small in number, to date. However, those that are already known to be solvable are of sufficient interest to inspire a vast world-wide effort to build a quantum computer and to encourage many theorists to look for solutions to other problems.

As with any technology it is up to us as responsible human beings as to how we should put it to use. Some of the calculations that will become possible if we can build a "quantum computer" will allow us to break secret codes to malicious ends. On the other hand, the speed of database searches could be dramatically increased. Such a breakthrough would allow us to solve complex problems in the life sciences such as searching the genome for sequences of interest, or in financial services, finding patterns that could predict market movements.

The proposal in 1995 that charged atoms in an electromagnetic "cage" could be used as a prototype quantum computer spurred a flurry of interest in the atomic physics community. The proposal required these charged atoms, or ions, to be cooled to temperatures near absolute zero using laser interactions. The ions could then hold the information to be processed in their internal energy levels.

Ion traps were already widely used for the delicate manipulation of individual atoms, the ultimate building block of matter. Later in 1995, the same year as it was proposed, a basic demonstration of a "quantum gate" was realised. This "gate" took information as an input, processed it and presented the output. The output was as predicted by the original proposal and work has continued at an astounding pace in many groups around the world since then to further develop the idea and build a scalable quantum processor.

The particular type of trap that was originally proposed for use in a quantum computer was a linear radiofrequency (RF) trap. Such a trap has a rapidly-alternating voltage, trapping the ions first in one direction and then in another in quick succession. It has several disadvantages, however. The rapidly-varying voltage can heat the ions, potentially destroying the fragile states that record the information. It also requires the electrodes to be small so that the voltage can be alternated sufficiently rapidly. Despite these problems, this design of trap has proved remarkably successful in many experiments. Another type of trap exists that may be even better-suited to quantum information processing. This trap is called the Penning trap and is based on a static electric and a static magnetic field. The absence of alternating voltages means that the heating of the ions should be much smaller.

The Ion Trapping Group in Imperial College has a strong background in the manipulation of ions in Penning traps. A project was initiated several years ago to investigate the performance of a Penning trap for quantum computing.

One of the first obstacles to be overcome in the use of a Penning trap is

laser-cooling the trapped ions. This is considerably more complicated than in the RF trap due to the presence of a strong magnetic field. On the most basic level, more lasers are normally required to cool the ions than would be needed in an RF trap due to the splitting of the ion's energy levels by the magnetic field. A slightly more subtle problem is that one of the three characteristic motions of the ion in a Penning trap is very difficult to cool. This motion is called the magnetron motion and it has an overall negative energy, meaning that energy has to be added to it in order to reduce its "temperature" and the radius of the motion.

Cooling of calcium ions in a Penning trap was demonstrated by the group in Imperial. Although the cooling was efficient overall, the radius of the magnetron motion was still relatively large. A technique for damping the magnetron motion, which had first been developed for mass spectroscopy experiments with Penning traps, was adapted for use in the presence of laser-cooling. Successful experiments were performed with trapped magnesium ions, in parallel with early calcium-ion experiments. The technique, known as axialisation, couples the magnetron motion with one of the other motions, the so-called modified cyclotron motion. Laser-cooling the modified cyclotron motion reduces its radius and, through the coupling, the magnetron radius too.

In this thesis, I describe how we can control the dynamics of small clouds of ions in a Penning trap using this axialisation technique. I provide some background on the field of ion traps and laser-cooling in Chapter 2. I then outline some of the theoretical aspects of quantum information processing in Chapter 3 with particular reference to schemes for working with ion traps.

A sophisticated experimental set-up has been built up over the course of the project to investigate QIP in Penning traps. I describe this set-up in Chapter 4 and a major addition to the set-up, a "wavemeter" in Chapter 5.

The theory of axialisation is covered in Chapter 6 along with a detailed description of ion motion in a Penning trap. The application of axialisation to small clouds of magnesium and calcium ions is presented in Chapters 7 and 8 respectively.

Some related experiments I worked on in the University of Innsbruck are discussed in Chapter 9. These experiments also involved calcium ions, but this time in a radiofrequency trap. A pair of highly-reflective mirrors around the trap will allow cavity quantum electrodynamics (QED) methods to be applied to the quantum computing system.

Finally, in Chapter 10, I summarise the work performed and discuss the possible future directions it might take.

Most of the work described in this thesis was undertaken as part of a team. Where possible I have explicitly mentioned in the text where a part of an experiment was developed by somebody else. The results in Chapter 7 were published in a refereed article [1] and a paper describing the work in Chapter 8 is in preparation [2]. Interim results have been presented in several talks and seminars including the EU Young European Physicists meeting, 2003 (organised by the QUEST Research and Training Network). A poster was also presented at the Trapped Charged Particles and Fundamental Interactions (TCPFI) conference, 2002.

Chapter 2

Ion trapping and laser-cooling

The isolation of single ions with minimal perturbation from their environment is desirable from both an aesthetic and a practical point of view in Physics. In the former case it allows physicists to achieve almost the ultimate level of control over the subject of their study. From the practical point of view, such trapped particles can be used for a wide variety of applications ranging from frequency standards to quantum information processing as outlined in Chapter 1.

Control over charged particles is most obviously exerted by the application of external electric and magnetic fields, with which they interact via the Lorentz force,

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \,, \tag{2.1}$$

where q is the charge on the particle, $\mathbf{E}$ and $\mathbf{B}$ the electric and magnetic fields[1] respectively, and $\mathbf{v}$ the particle's velocity.

Maxwell's Laws indicate that either an electrostatic field or a magnetostatic field alone cannot trap a charged particle:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon} \,, \tag{2.2}$$

$$\nabla \cdot \mathbf{B} = 0 \,. \tag{2.3}$$

This statement was originally formulated in Earnshaw's Theorem. If a single static field alone is insufficient to trap a charged particle, two possibilities suggest themselves: either vary one of the fields sufficiently rapidly that trapping occurs

[1] Throughout this thesis, the magnetic flux density, $\mathbf{B}$, will be referred to as the magnetic field (normally represented by $\mathbf{H}$. The two are related by the magnetic permeability: $\mathbf{B} = \mu\mathbf{H}$.). While technically inaccurate, this is a common usage.

alternately in different directions, or else combine suitable electrostatic and magnetostatic fields such that the magnetostatic field aids confinement in the direction(s) in which the electrostatic field does not trap. A broad-ranging overview of how such fields are generated and used has been provided by Ghosh [3].

2.1 RF traps

Wolfgang Paul was awarded the Nobel prize in 1989 for his part in developing the ion trap that now bears his name. The other recipients who shared the prize in that year were Norman Ramsey and Hans Dehmelt. Ramsey and Dehmelt were recognised for their work in separated oscillatory fields and the development of the Penning trap respectively.

Paul's Nobel-prize acceptance speech [4] gives a very succinct account of the RF trap and its development from the linear quadrupole mass filter. A more detailed treatment may be found in Ref. [3].

The trapping of any particle requires the application of a force that varies with the distance from a desired trapping location,

$$\mathbf{F} = -c\mathbf{r} \ . \tag{2.4}$$

Thus the particle is elastically bound to a point. This linear force is achieved if the particle moves in a parabolic potential. An example of such a potential is,

$$\Phi = \frac{\Phi_0}{2r_0^2}(\alpha x^2 + \beta y^2 + \gamma z^2) \ , \tag{2.5}$$

being an electric quadrupole potential where r_0 is a constant dependent on the separation of the electrodes generating the field. A three-dimensional electric quadrupole field subject to the Laplace condition ($\nabla^2\Phi = 0$) will have coefficients α, β, and γ adding to zero. One possible solution of the equation has $\alpha = \beta = 1$, $\gamma = -2$ and thus the form,

$$\Phi = \frac{\phi_0(r^2 - 2z^2)}{r_0^2 + 2z_0^2} \ , \tag{2.6}$$

with $r^2 = x^2 + y^2$ and $2z_0^2 = r_0^2$ (to express the potential in cylindrical coordinates).

In order to achieve the potential described by Equation 2.6 it should ideally be applied with electrodes with the same shape as the desired lines of equipotential. This is done with a ring electrode with a hyperbolic cross-section and two end-cap electrodes shaped as hyperbolae of revolution above and below it (Fig. 2.1). In practice, hyperbolic electrodes are difficult to manufacture and approximations to the shape are normally made. The deviation from the true quadrupole can be obtained computationally using a software package such as SIMION [5].

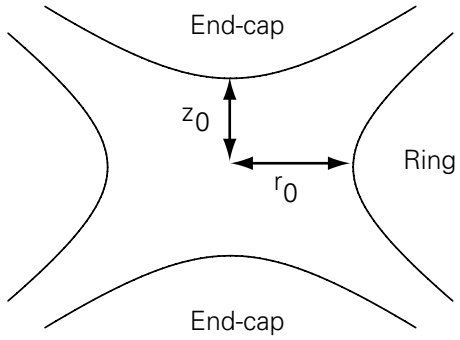

Figure 2.1: A cross-section through the electrode structure of a Paul trap. The electrodes have a hyperbolic cross-section to match the equipotentials of the quadrupole field.

In order to confine ions in such a trap, an RF voltage must be applied to the electrodes [6] (possibly in conjunction with a DC offset). A simplified description of the trapping field is that the RF traps alternately in the axial and radial directions (Fig. 2.2). The frequency of the RF must be high (typically of order tens of megahertz) to afford stable trapping. A DC offset added to the RF serves to confine the ions more tightly in one direction (axial or radial) at the expense of the other. Typically the confinement is less strong in the radial direction with RF alone so the ring is positively biased with respect to the end caps to counteract this.

The motion of ions in such a trap can be described by Mathieu equations [3],

$$\frac{\mathrm{d}^2 r}{\mathrm{d}\zeta^2} + (a_r - 2q_r \cos 2\zeta) r = 0 \,, \tag{2.7}$$

$$\frac{\mathrm{d}^2 z}{\mathrm{d}\zeta^2} + (a_z - 2q_z \cos 2\zeta) z = 0 \,, \tag{2.8}$$

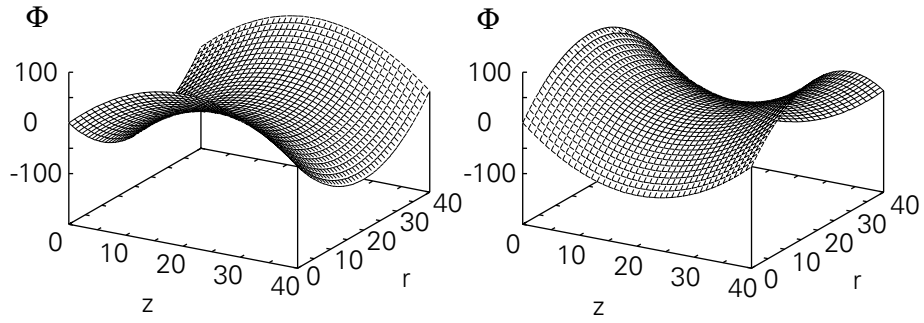

Figure 2.2: The potential in a Paul trap at the two phases of the RF drive. Ions are trapped alternately in one direction and an orthogonal direction. The scales on all axes are in aribtrary units.

with a, q, and ζ defined as follows,

$$\begin{aligned} a_z &= -2a_r = -\frac{8eU}{mr_0^2\Omega^2} , \\ q_z &= -2q_r = -\frac{4eV}{mr_0^2\Omega^2} , \\ \zeta &= \frac{\Omega t}{2} . \end{aligned} \tag{2.9}$$

U is the DC voltage applied, V is the amplitude of the RF voltage, m is the mass of the ion, the subscripts r, z refer to the radial and axial directions respectively and Ω is the frequency of the RF. a and q essentially indicate the trap parameters due to the DC and RF fields respectively.

The solutions to the Mathieu equations show that certain combinations of DC and RF voltages give rise to stable motions of the ions while others are unstable. A section of the stability diagram has been reproduced in Fig. 2.3.

Ions trapped stably in a Paul trap exhibit two motional frequencies: a secular frequency, $\beta_{r,z}\Omega/2$, and a micromotional freqency at the driving frequency, Ω. The value of β will lie between zero and one in the lowest trapping regime and is a function of a and q (see Ref. [3]).

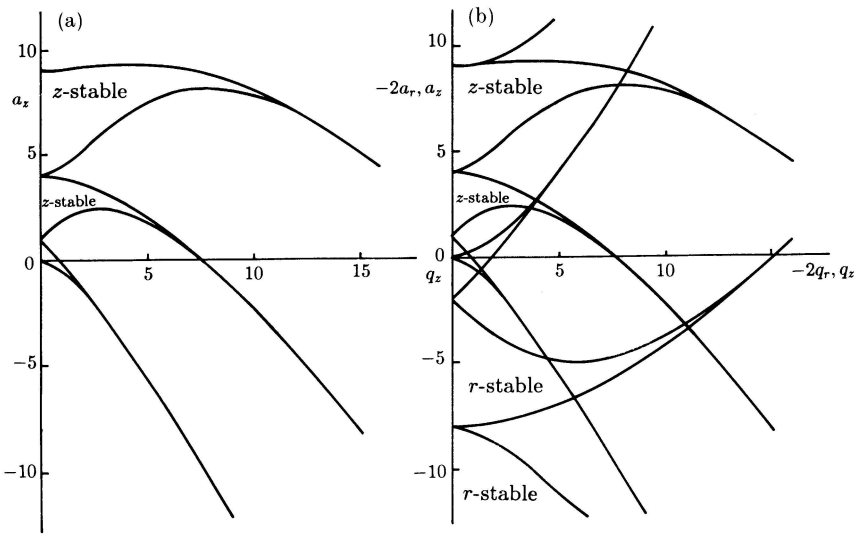

Figure 2.3: A section of the stability diagram for a three-dimensional oscillating quadrupole field (from Ref. [3]). *a* shows the stable regions for the axial motion while *b* shows the stable regions for the radial motion superimposed on those of the axial motion. Parameters must be chosen such that both radial and axial motions are stable in order to trap ions.

2.2 The Penning trap

The second approach to creating an ion trap is to combine both electrostatic and magnetostatic fields in a device called the Penning trap. There are two motivating reasons for doing this: firstly, the micromotion that is present in the Paul trap is absent in the static fields of the Penning trap; secondly, it is somtimes of interest to study an ion's properties in a high magnetic field.

The idea for the Penning trap came originally from a proposal by Penning to extend the path length of electrons between a pair of electrodes [7]. The first description of the use of hyperbolic electrodes along with a magnetic field to stably trap ions occurs in Ref. [8].

The hyperbolic electrodes of the Penning trap are in the same configuration as in the Paul trap. A static three-dimensional quadrupole is set up with the end-cap electrodes at a positive voltage with respect to the ring,

$$U(r,z) = -\frac{U_0}{r_0^2 + 2z_0^2}(r^2 - 2z^2) \ , \tag{2.10}$$

similar to that in the Paul trap, but in this case static. A saddle point in the potential exists in the centre of the trap and ions are trapped along the axis of the trap (the z-direction). The electrostatic field also pushes the ions outward in the radial direction. To counteract this, an axial magnetic field is applied. The ions move in a cyclotron motion perpendicular to the magnetic field and are thus confined in the radial plane.

Ion dynamics in a Penning trap are described in detail in Refs. [9, 10, 11, 12, 3]. The characteristic motions of ions in a Penning trap are the axial simple-harmonic motion and the radial cyclotron motion. The cyclotron motional frequency is modified from that which would exist in the absence of the electric field. The electric field cancels part of the Lorentz force due to the magnetic field and reduces the frequency of the cyclotron motion. The ions' acceleration along the electric field lines creates a velocity component along those lines and thus the second term in the Lorentz equation (2.1) gives rise to an $\mathbf{E}\times\mathbf{B}$ motion, called the magnetron motion.

The equation of motion in the axial direction is simply,

$$\ddot{z} = -\frac{4eU_0}{mR^2}z \ , \tag{2.11}$$

with $R = r_0^2 + 2z_0^2$. This is the familiar equation for simple harmonic motion with an angular frequency, $(4eU_0/mR^2)^{1/2}$.

The equations of motion in the radial directions are,

$$\ddot{x} = \frac{2eU_0}{mR^2}x + \frac{eB}{m}\dot{y} \ , \tag{2.12}$$

$$\ddot{y} = \frac{2eU_0}{mR^2}y - \frac{eB}{m}\dot{x} \ . \tag{2.13}$$

The first term is the repulsive electric field while the second is the effect of the magnetic field. The coupling of the two equations is due to this magnetic field but can be removed by moving to a rotating frame [12].

The appropriate choice of a rotating frame to remove the coupling between the x- and y-motions is one rotating at half the true cyclotron frequency ($\omega_c/2 =$

$eB/2m$). In this frame Equations 2.12 and 2.13 are [12],

$$\ddot{x} = \left(\frac{\omega_z^2}{2} - \frac{\omega_c^2}{4}\right) x \,, \tag{2.14}$$

$$\ddot{y} = \left(\frac{\omega_z^2}{2} - \frac{\omega_c^2}{4}\right) y \,, \tag{2.15}$$

The solution to this equation is simple harmonic motion in both the x- and y-directions with frequency, ω_1, given by,

$$\omega_1^2 = \frac{\omega_c^2}{4} - \frac{\omega_z^2}{2} \,. \tag{2.16}$$

It is clearly seen in Equation 2.16 how the magnetic field creates a potential to counteract the electrostatic potential ($-\omega_z^2/2$). The two solutions to this equation are rotation in opposite directions at the frequency $(\omega_c^2/4 - \omega_z^2/2)^{1/2}$. The motion in the positive direction is the modified cyclotron motion while that in the negative direction is the magnetron motion.

The frequencies of the three motions in the lab frame are,

$$\omega_z = \sqrt{\frac{4eU_0}{mR^2}} \,, \tag{2.17}$$

$$\omega_c' = \frac{\omega_c}{2} + \omega_1 \,, \tag{2.18}$$

$$\omega_m = \frac{\omega_c}{2} - \omega_1 \,. \tag{2.19}$$

An example of the variation with the trap voltage can be seen in Fig. 2.4.

In order for the motion to be bound, ω_1 must be real, giving the trapping condition,

$$\omega_c^2 - 2\omega_z^2 > 0 \,. \tag{2.20}$$

Some interesting relationships betewen the motional frequencies exist:

$$\omega_c' + \omega_m = \omega_c \,, \tag{2.21}$$

$$\omega_c' \omega_m = \frac{\omega_z^2}{2} \,, \tag{2.22}$$

$$\omega_c'^2 + \omega_m^2 + \omega_z^2 = \omega_c^2 \,. \tag{2.23}$$

While the axial motion is simple harmonic motion in a potential well the ra-

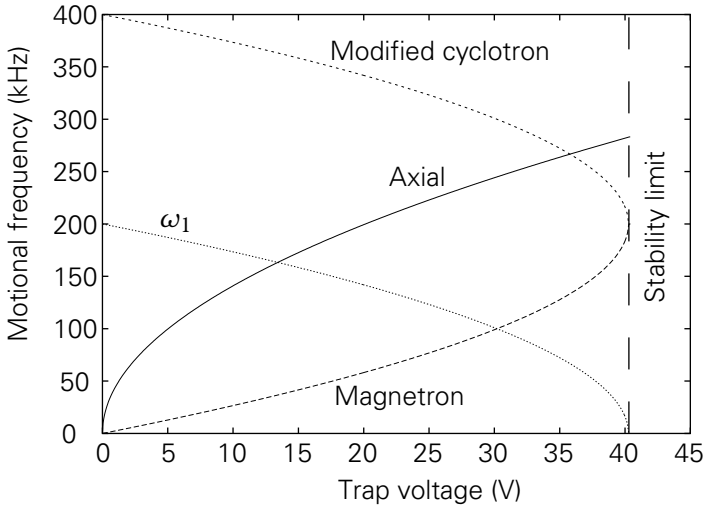

Figure 2.4: The variation of the motional frequencies with the trap voltage for a given magnetic field, ion species and electrode geometry. The two radial motions become degenerate at the stability limit.

dial motions are more complicated. The two radial motions describe an epicycle (Fig. 2.5) with the ion orbiting a field line in a fast modified cyclotron motion and precessing slowly around the axis of the trap in the magnetron motion. The magnetron motion is also rather unusual in that it is dominated by the negative potential energy and thus has an overall negative energy. This has the consequence that in order to reduce its radius, energy must be added rather than removed. This point will prove crucial in considering how to cool ions in a Penning trap. The energies of all three motions are derived by Schweikhard and Marshall [13],

$$E_{c'} = \frac{m}{2}(\omega_c'^2 - \omega_c'\omega_m)r_{c'}^2 \,, \tag{2.24}$$

$$E_m = \frac{m}{2}(\omega_m^2 - \omega_c'\omega_m)r_m^2 \,, \tag{2.25}$$

$$E_z = \frac{m}{2}\omega_z^2 z_{\max}^2 \,. \tag{2.26}$$

It can be seen that the trapping condition (Equation 2.20) requires that the total magnetron energy (Equation 2.25) be negative.

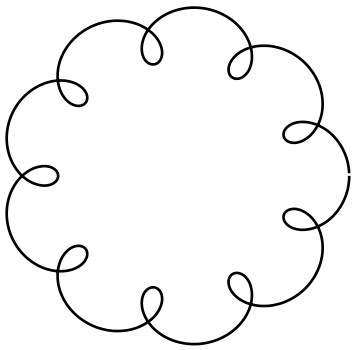

Figure 2.5: A typical ion motion in a Penning trap in the radial plane. The relative amplitudes of the two motions depend on the initial trajectory of the ion when it entered the trap. Here, the modified-cyclotron amplitude is one-fifth of the magnetron amplitude. The modified-cyclotron frequency was set to be ten times the magnetron frequency.

2.3 The combined trap

A third main category of ion trap is the combined trap. Essentially, this is a Paul trap with an axial magnetic field superimposed upon it. This magnetic field increases the radial confinement of the ions and thus extends the stable regions of operation of the trap. All three motional frequencies are altered by the addition of the magnetic field. Further details of such traps may be found in Ref. [3].

2.4 Buffer-gas cooling

Ions entering a trap typically have very high velocities (of order $1000\,\mathrm{ms}^{-1}$ in our experiments) and should be cooled if they are to remain in the trapping potential for a useful length of time. Many techniques for reducing ions' motional amplitudes in a trap have been developed. I will outline several of these here.

Perhaps the simplest cooling technique conceptually is the use of a neutral "buffer" gas. If an ion's temperature is 10^4 K on entering the trap (equivalent to a velocity of $1000\,\mathrm{ms}^{-1}$), then it may be easily understood that filling the trap with a gas at room temperature will reduce the ion's temperature rapidly through collisions. This is regularly done with helium or another inert gas at a pressure of about 10^{-4} mbar in RF traps used in mass spectrometry. The resolution of

the spectrum is improved through the increase in the measurement time.

Although buffer-gas cooling has been in use in RF traps for some time, it is only in the past 15 years that it has been possible to use it in Penning traps. The reason for this is that all of the characteristic motions of an ion in a Paul trap are in potential wells; collisions with background gas will reduce the ion's energy and thus reduce its radius. In the Penning trap, however, only two of the motions are in potential wells, while the magnetron motion is around a potential "hill". Thus, buffer-gas cooling in a Penning trap would serve to reduce the axial and modfied cyclotron motions but rapidly increase the radius of the magnetron motion; the ion would then be lost from the trap.

To get around the problem of the increasing magnetron radius, a method was developed to couple the magnetron and modified cyclotron motions [14]. With this coupling, the reduction in the modified cyclotron radius due to the buffer-gas cooling also led to a reduction in the magnetron radius. The technique is called "axialisation" and is described in Chapter 6.

2.5 Resistive-cooling

The idea of coupling two oscillators and using a damping of one to damp the other indirectly was not new. It had been used in ion traps for some time, especially in the Penning trap. A moving ion in a trap will induce small currents in the electrodes as it approaches or moves away from them. An analysis of these induced currents shows that they are at the same frequency as the ion's motion [10]. Thus, if the electrodes are connected to an external, tuned, circuit the motion can be coupled to that circuit if the tuning matches the ion's motional frequency. The circuit and the ion's motion are the two coupled oscillators and it is the circuit that is to be damped. The damping occurs through heat dissipation from the components of the electronic circuit and can be enchanced by cooling the circuit [15].

The method of resistive cooling can be improved by introducing an active feedback to the electrodes. Here, a small perturbation to the trapping voltage is applied in anti-phase to the electronically-detected ionic motion [3].

2.6 Light pressure and Doppler-cooling

The concept of light pressure dates back to J. C. Maxwell in the 19^{th} century. Maxwell realised that there is a momentum associated with electromagnetic radiation and that this can be transferred to matter in an interaction. This idea was radically revised when Planck's quantisation was proposed and it was realised that the quanta (photons) had a momentum related to their frequency. The study of the absorption and emission of photons by matter gave birth to the field of Quantum Optics.

In the 1930s, the first experiments on the recoil of atoms after absorption of light were performed by Frisch [16]. It was not until 1975, however, when the laser was widely available to physics laboratories, that the reduction of atoms' kinetic energy by radiation pressure was proposed. Two papers appeared in that year, the first published by Hänsch and Schawlow[2] [17] and the second by Wineland and Dehmelt [18].

Since 1978, when laser-cooling was first demonstrated [19, 20, 21], it has become a standard technique in atomic and molecular physics laboratories. The theory of laser-cooling has been covered in many books and review articles including Refs. [3, 22]. The original article by Hänsch [17] remains remarkably accessible, however, and this short account is based on that article.

If light is near-resonant with a sharp absorption line of an atom there will be a large scattering cross-section for interaction between that light and the atom. If the light is monochromatic (e.g. from a laser), and detuned to the red from resonance, then only atoms travelling in the opposite direction to the propagation of the light will see the light Doppler-shifted on-resonance. Atoms travelling in the same direction as the photons will see them Doppler-shifted further away from resonance.

Momentum transfer from the photon to the atom takes place in the absorption process. If the light is detuned as described such that atoms will only absorb a photon when they are travelling in opposite directions the momentum transfer will always serve to slow the atoms down. If such detuned light is applied isotropically to an atom (or at least along three non-parallel axes) the atom's translational motion will be slowed in all directions.

Re-emission of a photon from the excited atom will take place and this will,

[2]There is a link between this paper and Imperial College: Schawlow was visiting this institution when he wrote it

of course, also change the atom's momentum. This time the momentum change will be in the opposite direction to that of the photon-emission. The direction of emission, however, is random. The momentum kicks from many emission processes describe a random walk in momentum space and this stochastic process will have no net effect on the atom's momentum when averaged over a large number of emissions.

A lower bound on the temperature that may be reached by this method of cooling is imposed by the natural linewidth of the transition. When the Doppler width is comparable to the natural linewidth, the uncertainty due to the homogeneous broadening means that there is no longer selection of the absorbed-photon direction. At this stage of the cooling there are two competing processes. The momentum kicks after emission of a photon must now be considered as a heating process since the individual kicks are no longer negligible in comparison to the momentum of the atom. The cooling process is, as before, the Doppler-selected absorption of photons in the opposite direction to the atom's motion. The two processes will enter equilibrium when the atom's "temperature" is at the "Doppler limit". To calculate the Doppler limit, the natural linewidth is compared to the thermal energy. The following result follows,

$$T_{\text{Doppler}} \sim \frac{\hbar\gamma}{2k_B} \,, \tag{2.27}$$

where γ is the natural linewidth of the cooling transition, $\hbar$ is the reduced Planck's constant and k_B is Boltzmann's constant. Typical Doppler limits are tens to hundreds of micro-Kelvin.

The number of absorption and emission processes required to cool an atom from a typical starting temperature of 10^4 K to its Doppler limit is generally very large. The reduction in speed due to a single absorption process is $h\nu/mc$ where m is the mass of the ion and ν is the frequency of the light. For a typical starting energy of 1 eV, of order 1×10^5 photon interactions will be required to bring the atom to rest. Thus, although a small natural linewidth will result in a low Doppler limit, it will not allow a high enough interaction rate to efficiently cool the ion. For this reason, dipole-allowed transitions are used with linewidths of tens of megahertz. A second consequence of the large number of interactions required is that "re-pumper" lasers may be required. If there are long-lived levels to which the upper level of the cooling transition can decay other than the ground level, it will be necessary to re-pump population out of those levels

to shorter-lived levels in order to maintain a sufficiently high interaction rate.

As well as choosing an appropriate cooling transition, it is also necessary to use an appropriate laser power. It might naively be thought that the higher the power the greater the cooling rate. This is not the case since stimulated emission begins to reduce the cooling rate. If an atom absorbs a photon and is then stimulated to emit by a second photon arriving from the same direction, there will be no net change in the atom's momentum. A parameter I (not the laser intensity) can be defined to express the ideal laser power [3],

$$I = \left(\frac{\mu E}{\hbar} / \gamma \right)^2 , \tag{2.28}$$

where μ is the electric dipole moment of the atom and E is the electric field due to the laser. I is the square of the ratio of the induced flipping rate to the spontaneous flipping rate. It is found that for the maximum cooling rate $I \approx 1$.

2.7 Doppler-cooling in the Penning trap

Laser-cooling of ions in a Penning trap is more complicated than that of unbound atoms or of ions in a Paul trap. As well as the Zeeman splitting of the energy levels, which can require a larger number of lasers to close the cooling cycle than in a Paul trap, the unbound magnetron motion makes the choice of cooling parameters rather difficult. Despite this, laser-cooling was first demonstrated in a Penning trap [19] and several species of ion have since been successfully cooled [23, 24].

The theory of laser-cooling in Penning traps has been covered in Refs. [25, 26, 27] amongst others. It is found that the following factors affect the cooling rates: the laser detuning and spatial beam offset, the laser power, the intensity gradient of the laser beam and its orientation. In addition, the ion's oscillation frequencies and amplitudes will also determine the cooling rate.

Since the energy associated with the magnetron motion is negative, it is necessary to add energy to it in order to reduce the radius and velocity of its motion. The modified cyclotron motion is bound, however, and one must remove energy from it to cool it. To achieve both goals simultaneously is difficult and requires careful attention to be paid to the laser detuning and spatial offset [26, 27].

In order to add energy to the magnetron motion, the beam must be radially off-set away from the trap axis such that the photons are travelling in the same direction as the magnetron motion (see Fig. 2.6). If the laser is detuned to the red, the ion will absorb a photon when the modified cyclotron motion is moving it towards the laser and it will thus lose energy. Simultaneously, the magnetron motion will be moving the ion away from the laser during the absorption, and energy will therefore be added to the magnetron motion.

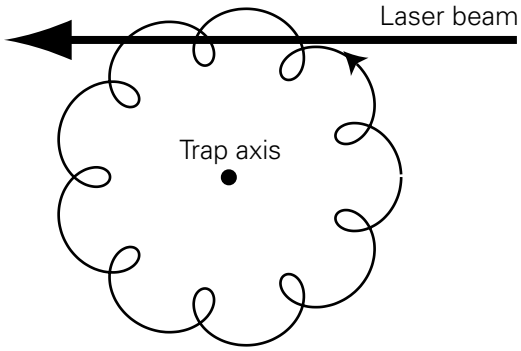

Figure 2.6: The cooling beam must be off-set with respect to the trap axis in order to cool the magnetron motion and modified cyclotron motions simultaneously.

Horvath and Thompson [26] described how the difference between the modified cyclotron and magnetron frequencies determines the degree to which the laser detuning and spatial offset can be varied while maintaining cooling of both the motions. This is seen in Fig. 2.7 where the extent of the region where both motions are cooled is determined by the difference in the two motional frequencies. The area in which the modified cyclotron motion is cooled is larger than that in which the magnetron motion is cooled, due to its larger frequency.

Thompson and Papadimitriou [27] created a simple analytical model to calculate the cooling rate as determined by the laser frequency and spatial offset. The results are found to be affected by the saturation parameter (laser intensity). Some results are reproduced in Figs. 2.8 and 2.9. It can be seen that the area in which both motions can be cooled simultaneously (where the shade is light at the same position on both the left- and right-hand plot) is very small. In most experiments, one motion will be cooled (usually the modified cyclotron motion) while the damping of the other is very low.

The importance of the the intensity gradient of the laser beam is discussed

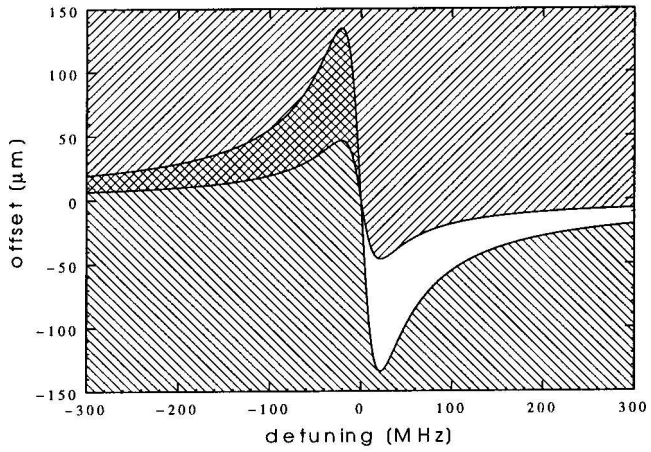

Figure 2.7: Cooling of the magnetron and modified cyclotron motions. In the top half of the plot, the magnetron motion is cooled (positive hatching) and in the bottom half, the modified cyclotron is cooled (negative hatching). The extent of the overlap is determined by the difference of the two motional frequencies. From Ref. [26].

by Itano and Wineland [25]. Magnetron cooling is achieved by having a steep intensity gradient in the beam. They also emphasize the importance of cooling all three motions of the ion since recoil from cooling two of them will heat the third. This requirement can be achieved by tilting the beam at an angle to the radial plane of the trap such that a component of the photon-momentum is in the axial direction.

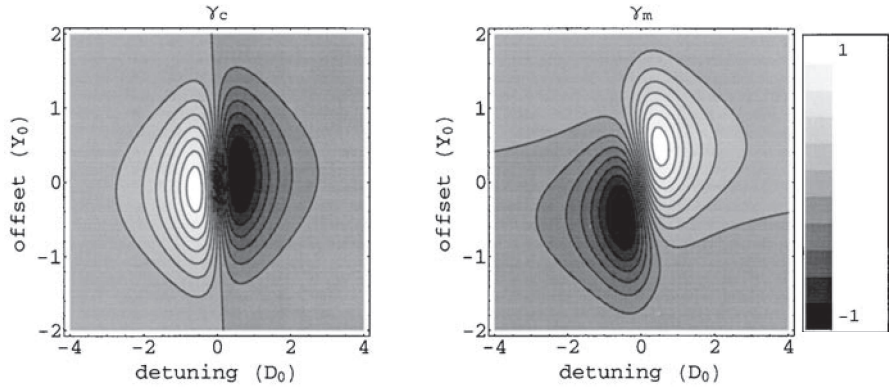

Figure 2.8: Modified cyclotron (left) and magnetron (right) cooling rates. Lighter colours mean strong cooling while dark means strong heating. This pair of plots were calculated for a low laser saturation parameter. There is no combination of frequency and spatial offset here for which both motions are cooled simultaneously. As expected, detuning the laser to the red cools the modified cyclotron motion while heating the magnetron motion. From Ref. [27].

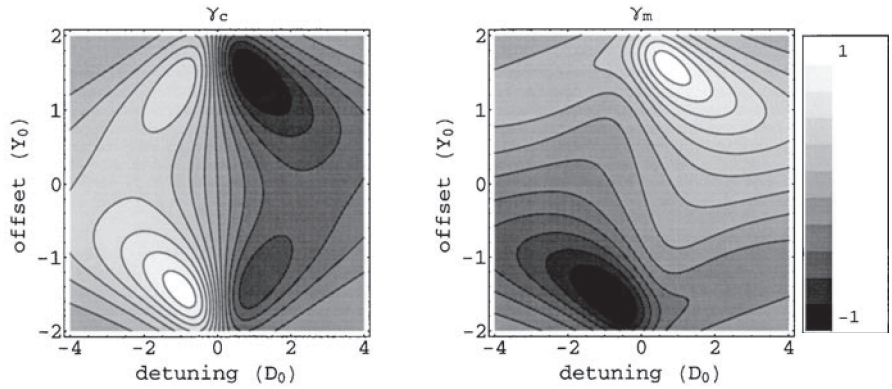

Figure 2.9: Modified cyclotron (left) and magnetron (right) cooling rates. These plots were calculated with a high laser saturation parameter. In both these plots and those in Fig. 2.8 it can be seen that the modified cyclotron motion can only be cooled with a laser detuned to the red. For the higher laser powers in this figure a small amount of magnetron damping can be achieved with the laser detuned to the red. This overlap with the modified cylotron damping region is small and both damping rates are very low. From Ref. [27].

2.8 Resolved-sideband cooling

An atom in periodic motion will have a spectrum with "sidebands" on either side of the normal spectral line, or "carrier". This spectrum can easily be derived by substituting the ion's equation of motion into the equation of the electric field due to the absorbed or emitted radiation,

$$E = E_0 \sin(kx - \omega_0 t) \ , \tag{2.29}$$

with

$$x = x_0 \sin(\Omega_x t) \ , \tag{2.30}$$

giving,

$$E = E_0 \sin(k x_0 \sin \Omega_x t - \omega_0 t) \ , \tag{2.31}$$

which can be rewritten in terms of Bessel functions as follows,

$$E = E_0 \sum_{n=0}^{\infty} J_n(k x_0) \sin\left((n\Omega_x - \omega_0)\, t\right) \ . \tag{2.32}$$

Taking the Fourier transform of Equation 2.32 to give an expression in the frequency domain and squaring to give the intensity yields,

$$I(\omega) \propto \sum_{n=-\infty}^{+\infty} \frac{|J_n(\beta)|^2}{(\omega - (\omega_0 - n\Omega_x))^2 + (\gamma/2)^2} \ , \tag{2.33}$$

where the linewidth of the transition, γ, has been introduced, and with $\beta = kx_0$ being the "modulation depth". This modulation depth determines the amplitude of the sidebands. The carrier will always be present and the sidebands will always have frequency separations from this carrier which are integer multiples of the motional frequency (see Fig. 2.10).

Advantage may be taken of the sidebands on a trapped ion's spectrum to reduce the ion's motional amplitudes beyond what is possible with the Doppler cooling described earlier. The ion's energy levels may be represented as shown in Fig. 2.11. If the sidebands can be resolved, i.e. the motional frequency is greater than the cooling transition's linewidth, then transitions may be made between resolved motional states. The laser is simply detuned from the normal transition frequency (the carrier) by the motional frequency. Detuning to the red will cause a reduction of one in the motional quantum number.

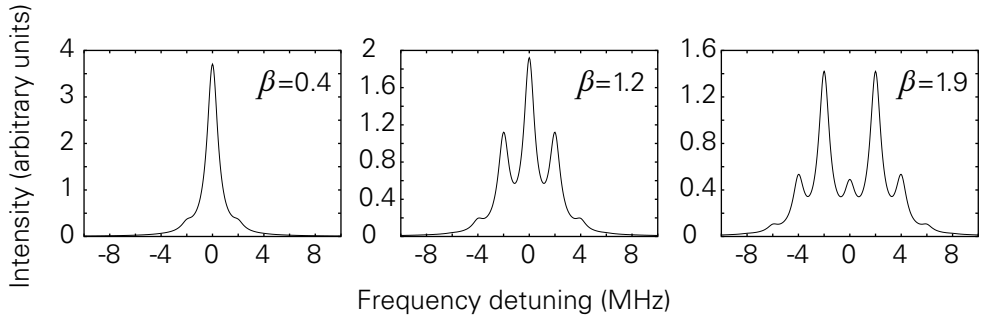

Figure 2.10: Sidebands on either side of a carrier, as calculated from Equation 2.32. The amplitude of the sidebands is determined by the modulation depth, β (the amplitude of the ion's motion in this case). The calculations were performed for a linewidth of 1 MHz and a motional frequency of 2 MHz.

The major difficulty that arises is that a narrow transition must be used such that the motional sidebands can be resolved. This requirement can be eased slightly by ensuring the motional frequencies are large e.g. through the use of a large magnetic field in a Penning trap. The use of a narrow transition also means that the decay back to the ground state will be slow and the overall cooling rate will be low. In order to speed this process up, a repumper laser may be employed to pump population out of the upper level of the sideband-cooling transition to another, shorter-lived, state.

Once the motional quantum number is zero no further absorption can take place on the red sideband. Near the ground state of motion the average motional quantum number can be measured by comparing the relative strengths of the red and blue sidebands. The ratio of the two sidebands is $\tilde{n}/(\tilde{n}+1)$.

2.8.1 The Lamb–Dicke parameter

In the context of the discussion of motional sidebands, an atom is said to be in the Lamb–Dicke regime when the second-order sidebands have a negligible height; that is, the atom's amplitude of motion is small in comparison to the wavelength of the light and the modulation depth is consequently low. When this condition holds, the first-order Doppler shift is eliminated and only the second-order Doppler shift remains.

The Lamb–Dicke parameter, η, relates the extent of the ion's wavepacket in the ground state of motion to the wavelength of light used to probe it. The

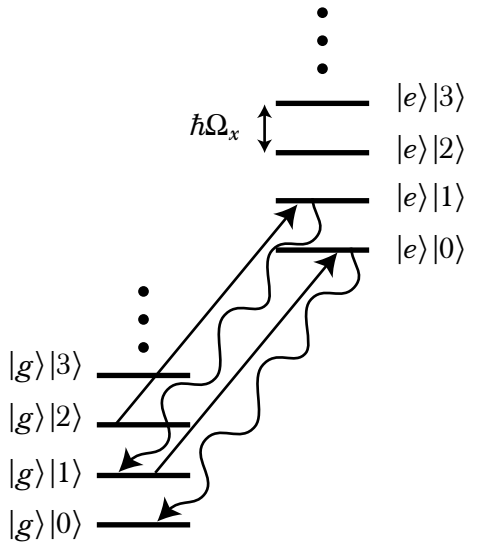

Figure 2.11: An ion's motional quantum states superimposed on its ground and excited energy levels (not to scale). Detuning the laser to the red by an amount equal to the motional frequency causes a transition to the next lowest motional state ($|g\rangle|n\rangle \rightarrow |e\rangle|n-1\rangle$). Decay back to the ground state via spontaneous emission tends, on average, to cause no change to motional quantum number.

Lamb–Dicke parameter is defined,

$$\eta = (\mathbf{k} \cdot \hat{\mathbf{z}}) z_0 \ , \tag{2.34}$$

where $\mathbf{k}$ is the propagation vector of the light ($|\mathbf{k}| = 2\pi/\lambda$), $\hat{\mathbf{z}}$ is the unit vector in the direction in which η is to be measured, and z_0 is the extent of the ion's ground-state wavepacket in that direction. Thus, an ion may be tightly confined along one axis but more weakly confined along the others and the Lamb–Dicke parameters will be different along each axis. Due to the scalar product in the definition, the direction of propagation of the probe light will also be important.

The Lamb–Dicke limit applies when the ion's spatial extent in its motional ground state is less than $\lambda/2\pi$. That is,

$$\eta \ll 1 \ . \tag{2.35}$$

A definition of the Lamb–Dicke parameter equivalent to Equation 2.34 is

given by,

$$\eta = \sqrt{\frac{\hbar^2 k^2}{2m}/(h\nu)}\ , \tag{2.36}$$

with ν being the oscillation frequency of the trap.

Different interpretations of the Lamb–Dicke parameter exist. Lindberg [28] describes it as the speed of cooling versus the speed of internal processes. In its form shown in Equation 2.36 it is clearly the square root of the ratio of the recoil energy to the vibrational quantum energy. Similarly Wunderlich [29] describes how it controls the coupling between internal and motional states of the ion.

This last interpretation is an important consideration for some quantum-computing schemes. As will be described in Chapter 3, many ion-trap schemes require the motional state for inter-ion communication. An incoming photon must have sufficient momentum to excite a phonon of the motional mode. Thus a very small η will mean that the photon's energy is small in comparison with the motional energy and no excitation of the motional mode will be possible. It is for this reason that many schemes are designed to work with optical radiation, with a high photon energy, rather than microwaves for example.

2.9 Raman- and EIT-cooling

Another method of cooling to the ground state of motion is to narrow the linewidth artificially by using a Raman transition. The transition is a three-level transition with two lasers (Fig. 2.12). The two lasers are each detuned from resonance with their particular half of the transition by a small amount, Δ. This detuning effectively narrows the resonance and, if sufficiently narrow, allows the motional sidebands to resolved. Even if this limit does not hold, a narrower linewidth will allow a lower Doppler limit to be achieved.

Electromagnetically induced transparency (EIT) is a quantum interference effect where a dark resonance is formed by the interference of two transitions. A bright resonance occurs higher in frequency than this dark resonance. If the separation of the two resonances can be set such that it is equal to the motional frequency, absorption will take place only on the red sideband and will be suppressed on the carrier.

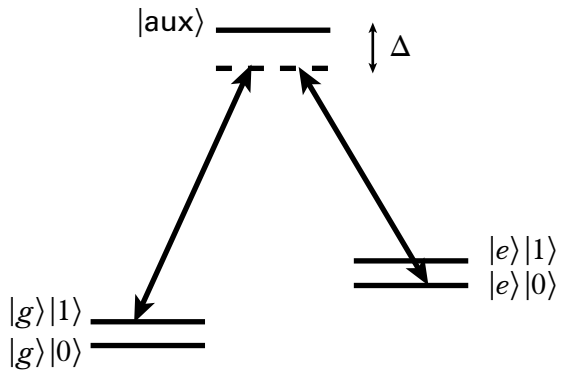

Figure 2.12: The Raman cooling scheme (not to scale). Two beams are used, each detuned from resonance with its respective transition by the same amount, Δ. Only two of the motional states are shown for $|g\rangle$ and $|e\rangle$ but these extend, of course, to infinity. Motional states have been omitted on the auxiliary state since their separation is very small in comparison with Δ.

Chapter 3

Quantum Information Processing

The idea of quantum information has captured the imagination of researchers across the world in the fields of mathematics, computer science, and physics. We live in a knowledge-driven society and proposals that promise to revolutionise the way we process information and acquire new knowledge are inevitably likely to attract the attention of both academia and industry alike.

In this chapter I will briefly outline the concept of quantum information processing (QIP) and the ways in which it might be applied to real-world problems. Many thorough treatments of the subject have now been published including Refs. [30, 31].

It will be seen that ions in a trap are considered to be one of the most attractive systems in which to investigate QIP. It is for this reason that ongoing work in the group in Imperial College aims to assess the feasibility of QIP in a Penning trap. An understanding of some of the requirements imposed by existing QIP proposals is necessary then for this assessment.

3.1 Information theory

Computer science and information theory have been most heavily influenced by two men: Alan Turing and Claude Shannon. Turing developed an abstract computing device, a Turing machine, whose state and output stream could be changed based on the order and content of its input stream. This provided the theoretical framework for the investigation of the capabilities and limitations

of classical computing devices[1] and is also credited with inspiring the development of the first digital computer by the British intelligence service during the Second World War. Shannon's contribution formalised the theory of information transfer. He calculated ultimate limits on transfer rates on both ideal and noisy transmission channels. His theorems also linked the abstract ideas of information theory with physical laws.

3.1.1 Classical information

At a very basic level, information can be represented in numerical form and manipulated mathematically. This applies both to numerical data and to the spoken or written word. Encoding methods are used to represent the data to be analysed or transmitted in a form that is understandable by the machine employed for this purpose. In designing such machines it is often found desirable to work in a binary (base 2) number system with switches, transistors, or other devices representing '1' with an 'on' state and '0' with an 'off' state. For example, to encode a word in a binary form, the individual letters might be converted to integers indicating their position in the alphabet ('1' for 'a', '2' for 'b', etc.) and the integers converted to binary ('00001' for 'a', '00010' for 'b', etc.). These numbers could then be represented internally by a row of five switches in the appropriate on- or off-states. The binary system of 0 and 1 is sometimes described as low and high or as false and true. Each '0' or '1' is called a ***b***inary dig***it*** or bit.

Processing of information is performed by gates. A gate takes one or more inputs and produces outputs based on those inputs. Many gates in common use today take two inputs and produce one output. The combination of the inputs to produce an output is based on Boolean algebra. An example of a gate which takes just one input and produces one output is the NOT gate. This gate simply inverts the input (Table 3.1). A simple two-input gate is the AND gate. The output of the AND gate is true if and only if both inputs are true.

A gate related to the AND gate is the NAND gate. It is equivalent to an AND gate followed by a NOT gate. The importance of the NAND gate is that it is easily implemented using transistors and that all possible logical operations can be performed using suitably interconnected NAND gates alone.

[1] A classical computer is one governed by the laws of classical physics. By contrast a quantum computer can only be described by the rules of quantum mechanics.

A	NOT A
0	1
1	0

A	B	A AND B
0	0	0
0	1	0
1	0	0
1	1	1

A	B	A NAND B
0	0	1
0	1	1
1	0	1
1	1	0

Table 3.1: Truth tables for some common gates

3.1.2 Quantum information

The importance of Shannon's association of information theory with physics should not be underestimated. All information must be represented somehow by physical means, from electrical impulses in the brain and vibrations of air (e.g. speech) to light pulses in optical fibres. The laws of physics will restrict the ways that information may be represented and the rate at which it can be transmitted. Rather than viewing this solely as a restriction, however, some saw this as an opportunity. In the early 1980s Benioff [32] and Feynman [33] separately proposed using quantum systems for computation. Using a quantum system would allow the solution of problems currently considered intractable on classical computers.

The primary feature of a quantum system that makes it so attractive for computation is the existence of superpositions of states. If we represent false and true by the states $|0\rangle$ and $|1\rangle$, it is also possible to construct a state,

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle \ . \tag{3.1}$$

The probabilities of finding the system in one of the two states, $|0\rangle$ or $|1\rangle$, are $|\alpha|^2$ and $|\beta|^2$ respectively. Until the system is measured, however, it remains in some superposition of the two states. A quantum system which can be placed in such a superposition of two states is called a ***qu***antum ***bit*** or qubit.

If an initial superposition state, $|\psi\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, is formed there is an equal likelihood of finding the system in either $|0\rangle$ or $|1\rangle$ upon measurement. Any operation on this state (such as addition, subtraction, etc.) will result in a superposition of the possible answers based on the two different inputs, $|0\rangle$ and $|1\rangle$. A quantum parallelism is in operation here such that the values of a function are being simultaneously calculated for several different input values. Of course, upon readout, only one of the possible answers can be obtained and this is no better than a classical computer. In order to overcome this restriction, quantum

algorithms have been developed whereby quantum interference effects cause one, desired, answer to be more likely to be read out at the end than the others.

3.2 Quantum algorithms

The first quantum algorithm to be developed was by Deutsch [34]. Deutsch proposed a quantum analogue of the Turing machine. He also provided an algorithm which would run more quickly on such a quantum computer than the best algorithm for the same problem could run on a classical computer. Algorithms (quantum or otherwise) are classified by the way in which their running times scale with the number of inputs. Fast algorithms may execute in polynomial time; that is, for N inputs the running time is proportional to N^2, say. Other algorithms may scale rather more badly and have execution times proportional to 2^N for example.

Deutsch's algorithm was developed further in collaboration with Jozsa and later by other researchers. It concerns whether or not a function is balanced or constant. In a simplified version of the algorithm it can be compared to determining if a coin is fair (one "head" and one "tail") or biased (two "heads" or two "tails"). In this simplified version, a classical algorithm would have to look at each side of the coin in turn (i.e. two operations) to determine if it was fair or not. However, the Deutsch–Jozsa algorithm could determine whether it was fair in a single operation.

The speed-up of the Deutsch-Jozsa algorithm over the equivalent classical algorithm is considerable. The algorithm does not perform a useful task, however. Other quantum algorithms were developed following Deutsch's original paper and their potential usefulness is what has prompted the considerable interest from funding bodies in the development of QIP theory and experiments.

The two algorithms which have attracted the most attention are Shor's algorithm for the factorisation of large numbers and Grover's algorithm for searching an unsorted data-set. Shor's algorithm is of use primarily because the security of most encryption technology in use today is based on the difficulty of factorising large numbers by classical computers. The Fast Fourier Transform (FFT) method of factorisation requires $n2^n$ steps on a classical computer for an n-bit number. Shor's algorithm can perform the same task with n^2 steps. This is an exponential speed-up over the classical algorithm.

A	B	A XOR B	A	B	A'	B'
0	0	0	0	0	0	0
0	1	1	0	1	0	1
1	0	1	1	0	1	1
1	1	0	1	1	1	0

Table 3.2: The truth tables for the XOR and c-NOT gates. For the c-NOT gate it is seen that the control channel (A) remains unchanged (A'). The target (B) is flipped if and only if the control is true. It should be noted that the output of the target (B') is identical to the output of the XOR gate.

Grover's algorithm demonstrates a much more modest improvement over the classical equivalent. Classically, an unsorted list can be searched in N steps where N is the number of records. To find the same information with Grover's algorithm requires just $\sqrt{N}$ steps. The importance of Grover's algorithm is not the degree of speed-up but rather the importance of the problem it solves. Databases are at the core of most IT utilisation in modern commerce and an improvement in search times would be tremendously useful. The usefulness is moderated slightly by the need to load the data into the quantum computer before searching. This will, of course, take N steps.

3.2.1 The c-NOT gate

Just as classical algorithms are built up from a long series of logical operations performed by classical logic gates (Table 3.1), quantum algorithms are built up from a series of operations performed by quantum logic gates. The crucial difference between the two types of gate is that quantum gates must be reversible. In fact, it was this property and the associated lack of energy dissipation that first led to the investigation of quantum computing [32]. This reversibility requires that there be the same number of outputs as inputs.

An important example of a quantum logic gate is the controlled-NOT (c-NOT) gate. It is closely related to the eXclusive OR (XOR) gate (Table 3.2). The XOR gate has two inputs and one output. The output of the gate is true if one and only one of the inputs is true. The c-NOT gate has two inputs and two outputs. The first input is the control while the second is the target. The first output is always the control, unchanged. If the control is false the second output is the target, also unchanged. If the control is true, however, then the second output is the target subjected to the NOT operation (i.e. it is inverted).

The importance of the c-NOT gate is similar to that of the NAND gate in classical computing. Any arbitrary quantum algorithm can be decomposed into a series of c-NOT gates and single qubit rotations [35] just as any classical algorithm can be decomposed into a series of NAND gates. This property is referred to as universality. In fact it turns out that almost any pair of quantum gates is universal. The reason the c-NOT gate has acquired particular prominence is that it is relatively easy to imagine some physical processes where one qubit is flipped based on the state of another.

3.3 Quantum error correction

An important point which has not been mentioned so far is that quantum states are rather "fragile". It is very difficult to put a system into an arbitrary superposition state and have it remain in such a state for a long enough period to apply several quantum gates. The reason for this fragility is that the qubit will interact with the environment and the state of the qubit when measured will be different than that expected. In order to measure the expected outcome, one would also have to somehow make a measurement of the environment with which the qubit has interacted. This is, of course, impossible to do in most cases. The process of loss of quantum information to the environment is called decoherence.

The concept of environmental noise affecting stored or transmitted information is not new. Shannon's theorem deals with this very issue and predicts how well classical information can be transmitted given a certain level of environmental noise. Shannon also describes an important method for overcoming the noise limitations of a system. This method requires the use of additional bits for encoding information that can later be used to reconstruct the original information in the event of some of the bits changing due to environmental noise.

The direct application of such an error correction scheme is impossible with qubits. The reason for this is that one would need an exact knowledge of the state of the qubit (the values of α and β in equation 3.1). This is an impossibility because the measurement would cause the superposition state to collapse to one of the two basis states. A more general statement can be made that an arbitrary, unknown, superposition state cannot be copied perfectly. This is called the no-cloning theorem.

Perhaps surprisingly, it is found that there are ways of correcting errors in qubits despite the limitation of the no-cloning theorem. Such quantum error

correction schemes were first proposed by Calderbank, Shor *et al.* [36], and independently by Steane [37].

Quantum error correction (QEC) requires the encoding of a single logical qubit, $\alpha|0_L\rangle + \beta|1_L\rangle$, in a number of physical qubits, say $\alpha|000\rangle + \beta|111\rangle$. Although it is not possible to measure the logical qubit during a computation, one of the physical qubits may be measured without destroying the coherence. A suitable choice of encoding means that the measured state of one of the ancillary qubits gives us some information about the error that has been introduced to the logical qubit without giving any information about the actual state of that logical qubit. The error can then be corrected with an appropriate choice of quantum gates and the computation may continue.

The fundamental importance of QEC is that it makes it far less likely that a fundamental physical obstacle will be encountered in the future that makes large-scale quantum computing impossible. Were it not for QEC the demands placed on the system used for the computation would be phenomenally difficult to achieve. Decoherence rates far lower than anything currently known would be required.

3.4 The DiVincenzo criteria

Even though quantum error correction makes the choice of a quantum system for computation much easier, there are still some other requirements to be considered. These requirements were originally gathered together by DiVincenzo [38] and have become known as the DiVincenzo criteria. They are listed below with a brief explanation of how trapped ions fulfil or in some cases partially fulfil them. Needless to say, there are other systems which also fulfil many of the criteria, including ones based on nuclear magnetic resonance (NMR), neutral atoms in optical lattices and charge pairs in superconductors. At present, there is no known system which rigorously satisfies all the criteria.

One needs to be able to:

1. Have well-characterised qubits in a scalable system.
   The qubit representation in a trapped-ion scheme is a pair of electronic energy levels between which coherent transitions may be driven. This representaion is well-characterised and the system is scalable through the addition of extra ions to the trap. Long strings of hundreds of ions may be

formed and segmented traps may be used to trap even largers number of ions in many micro-traps.

2. Store a superposition state for a time much greater than the typical gate time.
   Either a long-lived meta-stable state of the ion may be used (e.g. the $^2D_{5/2}$ state in $^{40}Ca^+$ with a lifetime greater than 1 s) or a pair of hyperfine levels of the ground state (e.g. those of $^2S_{1/2}$ in $^{43}Ca^+$). By comparison, typical gate times are of order tens of microseconds.

3. Initialise the qubit register.
   For either of the two qubits (meta-stable states or hyperfine levels) described above, the state can be initialised faithfully with optical radiation.

4. Perform a universal set of gates.
   Two-ion gates have already been demonstrated by two different groups [39, 40]. These gates will be reviewed in Section 3.7.

5. Measure the final result.
   An electron-shelving technique, where one of the electronic states of the ion is insensitive to the detection radiation, may be used to detect the qubit state with greater than 99% fidelity.

Although other systems also partially fulfil these criteria, trapped ions are currently favoured by many due to the ease of control with optical radiation, and the existence of proposals that will allow the current, small, demonstrations to be scaled to much larger devices.

3.5 A proposal for a quantum gate: the Cirac–Zoller scheme

The original suggestion for using trapped ions for QIP came in a paper by Cirac and Zoller in 1995 [41]. They proposed that a pair of electronic energy levels, $|g\rangle$ and $|e\rangle$, could represent the two logical states, $|0\rangle$ and $|1\rangle$. Interactions between qubits, as required for multi-qubit quantum gates, could be facilitated by a common mode of motion of the ions. This mode is referred to as a "bus" qubit.

The main attraction of the Cirac–Zoller scheme is that the decoherence rates of trapped ions should be small in comparison with other proposed systems since the ions are well isolated from the environment. The main source of decoherence the authors foresaw was the spontaneous decay of the electronic state and the damping of the motional state (e.g. by residual gas in the vacuum chamber). The choice of a suitably long-lived state and a low enough background pressure could make either of these sources sufficiently low to carry out many quantum gates before decoherence became significant.

Key requirements of the Cirac–Zoller scheme are the laser cooling and the tight confinement of the ions. The ions must be laser cooled to the ground state of motion using a technique such as resolved-sideband cooling. Initial occupation of the motional ground state creates the conditionality of the gate as will be seen. Once in the motional ground state, it is important the confinement of the ions by the trap is sufficently strong that the Lamb–Dicke limit holds (see Chapter 2) and the wavelength of the addressing laser is consequently large in comparison with the amplitude of the ion's motion. Of the two conditions it is that of cooling to the ground state that is the most demanding experimentally.

The scheme is outlined in Ref. [41] with the two qubit states represented by one of the $S_{1/2}$ and one of the $D_{5/2}$ states of a barium ion. This is a dipole-forbidden transition and is similar in character to the same transition in calcium ions. A simplified level diagram is shown in Fig. 3.1. Selection of a particular transition between the $S_{1/2}$ levels and the $D_{5/2}$ levels is performed by a combination of polarisation selection and wavelength detuning. Typically, a small magnetic field ($\sim 3\,\mathrm{G}$) will be applied to an ion trap and this is sufficient to remove the degeneracy of the levels by way of the Zeeman effect.

In order to describe the gate we consider the two ions' electronic states and their common motional state. The motional state will be represented by $|n_b\rangle$ where n is the motional quantum number. If both electronic and the one motional state are all in the ground state these can be represented, $|g_1\rangle|g_2\rangle|0_b\rangle$ where the numerical subscript is the ion number. It is assumed that lasers can be used to address individual ions. With inter-ion distances on the order of $15\,\mu$m in a weak trap and a tightly focused beam this is achievable. The addressing laser can be tuned to the carrier or one of the motional sidebands. As was seen in the discussion of sideband cooling (Chapter 2), tuning the laser to the red motional sideband will result in the removal of a motional quantum as well as excitation of the electronic state. If the motional quantum number is al-

ready zero this will have no effect on either the electronic state or the motional quantum state since $a|0_b\rangle = |0_b\rangle$ where a is the annihilation operator. With the laser tuned to the carrier frequency, excitation of the electronic state occurs with no effect on the motional state.

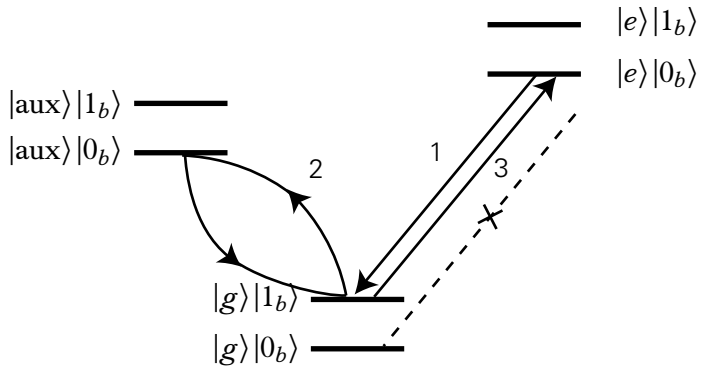

Figure 3.1: The level diagram of the Cirac–Zoller scheme. Both ions start in state $|g\rangle|0_b\rangle$, $|e\rangle|0_b\rangle$ or some superposition of the two. If the control qubit is in state $|g\rangle|0_b\rangle$ no transition can take place in step 1. If it is in state $|e\rangle|0_b\rangle$ the labelled transition to $|g\rangle|1_b\rangle$ takes place. If this latter transition takes place, one phonon is introduced into the ions' common mode of motion and the second ion can make a transition to the auxiliary state if it is in state $|e\rangle$. A 2π pulse on this transition makes the transition to the auxiliary state and back again, with a phase of -1 introduced into the state. Finally, step 3 removes the phonon from the motional mode and the first ion is returned to its original state.

The gate is a three-step process. The first step entangles the first ion's state with the bus qubit; the second performs a conditional operation on the second ion dependent on the bus qubit's state; the third maps the bus qubit back onto the first ion.

In the first step, the laser is tuned to the red sideband and applied to ion 1. If the ion is in state $|g_1\rangle|0_b\rangle$ there will be no change. However, if the ion is in state $|e_1\rangle|0_b\rangle$ then the state can be changed to $|g_1\rangle|1_b\rangle$ (Fig. 3.1). For transfer to this state the length of the pulse must be chosen carefully such that it is a so-called π-pulse. That is, the pulse length is Ω/π where Ω is the Rabi frequency for the transition. It should be clear now why sideband-cooling to the motional ground state is required: if the motional state were $|1_b\rangle$ to start with, a transition from $|g_1\rangle|1_b\rangle$ to $|e_1\rangle|0_b\rangle$ could take place and the conditionality would be lost.

The second step is a 2π-pulse applied to the second ion on the red sideband

Initial state		Step 1		Step 2		Step 3
$\|g_1\rangle\|g_2\rangle\|0_b\rangle$	$\rightarrow$	$\|g_1\rangle\|g_2\rangle\|0_b\rangle$	$\rightarrow$	$\|g_1\rangle\|g_2\rangle\|0_b\rangle$	$\rightarrow$	$\|g_1\rangle\|g_2\rangle\|0_b\rangle$
$\|g_1\rangle\|e_2\rangle\|0_b\rangle$	$\rightarrow$	$\|g_1\rangle\|e_2\rangle\|0_b\rangle$	$\rightarrow$	$\|g_1\rangle\|e_2\rangle\|0_b\rangle$	$\rightarrow$	$\|g_1\rangle\|e_2\rangle\|0_b\rangle$
$\|e_1\rangle\|g_2\rangle\|0_b\rangle$	$\rightarrow$	$-i\|g_1\rangle\|g_2\rangle\|1_b\rangle$	$\rightarrow$	$i\|g_1\rangle\|g_2\rangle\|1_b\rangle$	$\rightarrow$	$\|e_1\rangle\|g_2\rangle\|0_b\rangle$
$\|e_1\rangle\|e_2\rangle\|0_b\rangle$	$\rightarrow$	$-i\|g_1\rangle\|e_2\rangle\|1_b\rangle$	$\rightarrow$	$-i\|g_1\rangle\|e_2\rangle\|1_b\rangle$	$\rightarrow$	$-\|e_1\rangle\|e_2\rangle\|0_b\rangle$

Table 3.3: The intermediate steps in the Cirac–Zoller c-NOT scheme

of the transition from $|g\rangle$ to an auxiliary state. The purpose of this is to introduce a phase of -1 if the state of the system is $|g_2\rangle|1_b\rangle$. If the initial state of the first ion had been $|g_1\rangle$ or the initial state of the second ion had been $|e_2\rangle$ then no transition would have taken place.

The final step is to return the first ion to its original state. This is done with another π-pulse on the red sideband of the $|g\rangle$–$|e\rangle$ transition i.e. a repeat of step 1. The total process is shown in Table 3.3.

This three-step procedure does not produce a c-NOT gate as such. However, a Hadamard transformation, defined,

$$H|0\rangle \rightarrow \frac{1}{\sqrt{2}}\left(|0\rangle + |1\rangle\right) \ , \quad H|1\rangle \rightarrow \frac{1}{\sqrt{2}}\left(|0\rangle - |1\rangle\right) \ , \tag{3.2}$$

applied to the second ion before and after the three-step sequence will convert it to a c-NOT gate.

The two important features of the Cirac–Zoller proposal to note are that it requires individual addressing of ions and sideband-cooling to the ground state of motion. The first requirement means that the ions must be well-localised and well-spaced in the trap.

Many other schemes have been proposed for use in an ion trap quantum computer. Some have less stringent cooling requirements than the Cirac–Zoller proposal. For example, the Mølmer–Sørensen proposal [42] is described as a "hot" gate and uses the common motion of the ions only as a virtual level, much like a Raman transition puts a negligible population in the uppermost of the three electronic levels involved. The Mølmer–Sørensen proposal also has the advantage that individual ion addressing is not required during the gate procedure. While this eases the localisation requirements of the scheme, it is still necessary to somehow individually initialise and read-out the individual qubits respectively at the beginning and end of a computation. This initialisation and read-out will still require individual-ion addressing.

3.6 QIP in a magnetic field

Although the Cirac–Zoller scheme and most other ion-trap QIP proposals are designed for linear RF traps, the same principles can be applied to ions in a Peninng trap. For example, with tight radial confinement from a magnetic field and weak axial confinement from an electric field, a string of ions may crystallise along the trap axis. This ion string is analagous to that formed in a linear RF trap and the Cirac–Zoller scheme could in principle be applied to it.

There have been several proposals, however, that make use of the presence of a magnetic field. These are clearly of interest to those working with Penning traps and linear combined traps. Two of the proposals are briefly presented here.

3.6.1 Spin-spin coupling

Mintert and Wunderlich proposed [43] to address individual ions in a trap, not in spatial coordinates as in the the Cirac–Zoller and most other schemes, but rather in frequency space. A magnetic field gradient applied along the trap axis will cause a Zeeman splitting varying along the axis. Thus, the transition energy between the two qubit levels will be different for each ion, dependent on its position along the magnetic field gradient. Their choice of qubit levels are hyperfine states of an ion, separated by an energy equivalent to radiation in the microwave regime.

As well as allowing the addressing of ions in frequency space, this magnetic field gradient also couples internal and external degrees of freedom of the ions. As seen in Fig. 3.2, the ion's position changes after being excited from the lower state to the excited state. The displacement occurs since the $|0\rangle$ and $|1\rangle$ states' energies vary in opposite directions in the magnetic field gradient. When the ion changes from one state to the other, the equilibrium position changes by an amount [44],

$$\mathrm{d}z = \frac{\hbar\,\partial_z\omega}{m\nu^2}\ , \tag{3.3}$$

where $\partial_z\omega$ is the change in the transition frequency between the two qubit levels per unit length along the trap axis. The frequency of the common motional mode of the ions is given by ν while m is the mass of the ion.

Thus, all the necessary components of an ion-trap quantum computer are in place: individual ion addressing is possible for initialisation and read-out and a

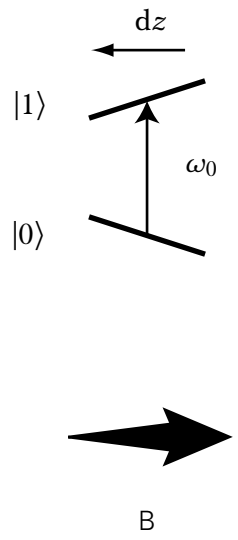

Figure 3.2: The effect of the magnetic field gradient on the qubit states. The Zeeman shift affects the splitting of the states as a function of the ion's spatial position. After a transition from one state to another the position of minimum potential energy changes and the ion moves by a distance dz along the trap axis to this new equilibrium position.

method of achieving ion–ion interactions is proposed via their common motion. It was soon realised, however, that hidden in the proposal was an even more powerful method of QIP with ions.

Although the original proposal in Ref. [43] treated the spin–spin coupling between ions in the trap as a negligible perturbation to the scheme, it was later realised [44, 29] that this coupling could form the basis of a novel form of quantum gate. Essentially the ion string could be thought of as a molecule in an NMR experiment. The major advantages over NMR experiments are that the couplings between ions can be controlled and scale well with the length of the "molecule", and that the computation is based on a single quantum system rather than a statistical analysis of an ensemble.

Mc Hugh and Twamley [45] extended the idea of spin–spin coupling between ions in a single linear trap to a system of ions in an array of micro-traps. They showed that such a system was analagous to an ion crystal.

3.6.2 QIP with electrons in a Penning trap

The group of P. Tombesi has published several papers [46, 47] describing a proposal to perform QIP with electrons in a Penning trap. The principle feature of the design is that individual electrons are stored in a linear array of Penning traps. A single electron bound to a Penning trap is a well-studied system and indeed has its own name, "geonium", coined by Dehmelt. The proposal in Ref. [47] is to use the three motional degrees of freedom of the bound electron as qubits.

A linear array of traps is proposed since the electron dynamics become chaotic after the introduction of more than one electron into the trap. Electron–electron interactions are mediated by the Coulomb force. The interactions are turned on and off by tuning the axial frequencies of the electrons using the end-cap voltage. When the frequencies of adjacent electrons are nearly equal, a strong interaction between the two is predicted.

The ideas behind QIP with electrons are not directly applicable to ions in a Penning trap. The proposals do have a feature in common with the proposal [45] described in the previous section, namely the use of an array of traps. Such an array of trapping potentials has already been proposed [48] for RF traps. There is also a similar proposal for an array of Penning traps [49].

3.7 Experimental realisations of QIP

A simplified version of the Cirac–Zoller scheme was demonstrated by the group of D. Wineland in NIST [50] in the same year as the original proposal [41] was published. Instead of using a pair of ions as qubits, a single ion was used and its motional mode performed the function of the control qubit. A single $^9Be^+$ ion was sideband-cooled to the ground state of motion with Raman beams near 313 nm. The internal qubit states were a pair of hyperfine ground states, $|F = 2, m_F = 2\rangle$ and $|F = 1, m_F = 1\rangle$ separated by 1.25 GHz and represented as $|\downarrow\rangle$ and $|\uparrow\rangle$. This internal qubit was flipped conditional on the external qubit being in the $|1\rangle$ state (one motional quantum). A probability of finding the qubits in the desired states at the end of the gate greater than 80% was recorded and the gate speed was about 20 kHz.

The next experimental stage was a demonstration [51], also by the group of Wineland, of two- and four-ion entanglement with a scheme similar to that pro-

posed by Sørensen and Mølmer [52]. The desired outcome of the entanglement operations is a state (for two ions), $|\psi\rangle = (|\uparrow\uparrow\rangle - i\,|\downarrow\downarrow\rangle)/\sqrt{2}$. In measurements on the entangled state, the probabilities of finding the states $|\uparrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$ were respectively 0.43 and 0.46 (compared to the ideal case of 0.5 for each).

The first demonstration of the Cirac–Zoller c-NOT gate [50] suffered from a lack of scalability since the motional mode was used as a control qubit instead of a bus qubit. In order to conclusively demonstrate the scalability of the approach, it was necessary to demonstrate the operation of a multi-ion gate. This was done, eight years after the original proposal and experiment, simultaneously by the groups of Wineland [40] and of R. Blatt (in the University of Innsbruck) [39].

Leibfried *et al.*, in the group of Wineland, used the same system of Be^+ ions they had used in their previous demonstrations [50, 51]. They did not use the scheme proposed by Cirac and Zoller [41]. Instead they demonstrated a phase gate [53], logically equivalent to the c-NOT gate. Such a gate relies on a state-dependent displacement force: the ions move only if their spins are opposite to one another. It is based on a similar concept to the entanglement scheme previously demonstrated [51]. A fidelity of 97% was reported. While still far below the required fidelity for scalable operation, even with the assistance of error correction, this is still a very impressive feat.

The c-NOT gate demonstrated in the group of Blatt [54, 39] closely followed the original proposal in Ref. [41]. Important refinements included the use of a composite-pulse sequence, similar to that used in NMR experiments. This corrects for unwanted evolution of the qubit state during the gate. The ion used was $^{40}Ca^+$, with the qubit represented by Zeeman sub-states of the $^2S_{1/2}$ and long-lived $^2D_{5/2}$ levels. Despite the many advances displayed in achieving their result, the fidelity of the gate of 71% was rather lower than that achieved in NIST. The main source of error was frequency-noise in the laser at 729 nm. Although there are many advances that can be made in improving the parameters used for this experiment, it is generally believed that for scalable quantum processors it is important to move to a Raman transition between hyperfine levels as a qubit basis. The importance of laser-frequency stability is much less in such a system and the lifetime of the upper level of the qubit is generally many orders of magnitude greater. Groups are already working on $^{43}Ca^+$, which has a similar level structure to $^{40}Ca^+$ but has hyperfine splitting of the $^2S_{1/2}$ state due to the odd number of nucleons.

There have been some other, notable, experiments reported more recently.

Again, it is the groups of Blatt and Wineland that have led the way in the field. The most recent results concern the creation of three-qubit entangled states [55, 56] and of the teleportation of a qubit from one ion to another [57, 58]. Similar techniques to those used in the entangled-states and teleportation experiments allowed the implementation of a three-qubit quantum error correction scheme by the group in NIST [59]. This was the first time quantum error correction was demonstrated with qubits represented by individual quantum systems rather than an ensemble as used in NMR experiments [60, 61].

3.8 Technical limitations

Although some wonderful experimental results have already emerged, as described in the previous section, there are still some major hurdles to be overcome before a large-scale ion-trap quantum computer can be built. Many of these were described in an article by Wineland *et al.* [62].

An important issue is the addressing of individual ions. Even if it is possible to perform a gate such as that proposed by Mølmer and Sørensen [42], which requires even illumination of the ions, it is still necessary to initialise and read out the final state of each ion separately. The technique employed until now by the group of Blatt has been to use a very tightly-focused beam with a small waist (about 2.5μm compared with an inter-ion separation of 5μm). A more promising method in the longer term is one proposed and implemented by the group of Wineland [48, 58] where a segmented trap is used to separate the ions and re-combine them in a "computation" region. A high level of coherence of the qubit state was reported after transport of the ion by Rowe *et al.* [63].

Scaling ion-trap quantum computation is hindered by the introduction of many more motional modes as more ions are placed in the trap. Resolving these modes when addressing a given sideband becomes difficult if there are too many in close proximity. A segmented trap can be used to address this problem too. With separate storage and computation regions, just two ions (for example) can be placed in the computational region during a gate process.

Important sources of decoherence of trapped-ion qubit states include electric-field and magnetic-field fluctuations [54, 39]. Interestingly, micromotion of the ions in an RF trap is no longer considered to be an important direct source of decoherence. Electric-field fluctuations can move the ions out of the potential minimum of the trap, however, and introduce some heating

through micro-motion. Such sources of electric fields include noise on the electrode voltages, AC Stark shifts and patch potentials on electrodes. Some progress has been made towards eliminating the latter problem through the use of photoionisation schemes [64] and the consequent lower atomic fluxes and the absence of electron beams during the loading process.

Great lengths must be taken to cancel out external magnetic-field fluctuations near the ions. Typically, a system consisting of a control circuit connected to a Gaussmeter and some large field-cancellation coils is used. Despite these precautions, stray magnetic fields still contribute greatly to the errors in quantum gate processes [39].

It is to be expected that neither electric- nor magnetic-field fluctuations will be as problematic in a Penning trap as they have been found to be in RF traps. Larger electrodes, further away from the ions, can be used since there is no need to apply a high-frequency alternating voltage. This should greatly reduce the problems of voltage-source noise and patch potentials. The use of a superconducting magnet to generate a large, stable, magnetic field should be beneficial in reducing the effect of external magnetic-field fluctuations. The achievable magnetic-field stability is of order one part in 10^8 in a typical superconducting magnet and external magnetic fields are excluded through the Meissner effect.

3.9 Discussion

Ions in a trap have already proved to be a promising avenue for the realisation of scalable quantum information processing. Refinements to existing schemes and entirely new schemes continue to be dreamt up and implementation of some of these schemes is progressing rapidly. Many of the serious sources of decoherence in the system are being addressed by groups working in the field.

The environment of the Penning trap continues to appeal as an improvement over existing RF-trap implementations. Some of the most serious sources of decoherence in RF traps such as patch potentials and magnetic field instability have been seen to be much less problematic in Penning traps. The greatest experimental hurdles still to be addressed before the Penning trap can seriously be considered as a contender are the sideband-cooling of the trapped ions and their localisation. The former is a work in progress in this group while the latter point is addressed by the technique of axialisation, described in this thesis.

Chapter 4

Experimental set-up

Doppler cooling of Ca^+ ions in a Penning trap was demonstrated for the first time by the group in Imperial College [24]. It is described in some detail by Koo [65] and Sudbery [66]. Since those early experiments, we have moved to using a different trap, with a ring electrode divided into quadrants. This has allowed the application of an axialisation drive. Both the earlier trap and the split-ring trap can be used as a Paul, Penning, or combined trap since they share the same electrode geometry. In fact, the split-ring trap can also be run as a linear RF trap or linear combined trap. These configurations are described in Chapter 2. The linear RF trap is described in considerable detail by Ghosh [3]. In this chapter I will restrict the description to that of the Penning trap since this is the configuration we have used predominantly.

4.1 Trap

The trap itself is the same as that used for experiments with magnesium ions (see Chapter 7). It was designed by M. van Eijkelenborg [67] and built in-house from beryllium copper. The original technical drawings by van Eijkelenborg may be found in Ref. [68] and are reproduced in Appendix B. The most important specifications of the trap are the radius of the ring electrode and the separation of the end-cap electrodes. Respectively, r_0 is equal to 5 mm and $2z_0$ is 7 mm.

The trap consists of two end-cap holders, a frame for the ring segments, two end caps and four ring segments (see Fig. 4.1). A lens holder fits on top of the frame for the ring and has tapped holes for supporting the atomic-beam ovens pointing in towards the trap. Filaments fit inside the end caps themselves. Holes

in the end caps allow the electron-beam into the centre of the trap. There are holes in all four sides of the ring-frame and between the ring segments themselves. Those in the vertical sides are input and output apertures for the laser beams. The hole in the top allows the imaging of the ions while that in the bottom reduces laser scatter into the imaging system.

The end-cap holders and the frame for the ring are screwed to a hollow cylindrical stainless steel pedestal. The end-cap holders are electrically isolated from the pedestal by Macor spacers. Where electric isolation is required between trap components, ceramic washers and nuts are used. Although very brittle and difficult to work with, they have the advantage of being ultra-high-vacuum (UHV) compatible. Electrical connections are made to the feedthrough by Kapton (polyimide)-coated wires. In order to carry the high currents used to power the filaments and ovens, thick wires rated to 5 A are used.

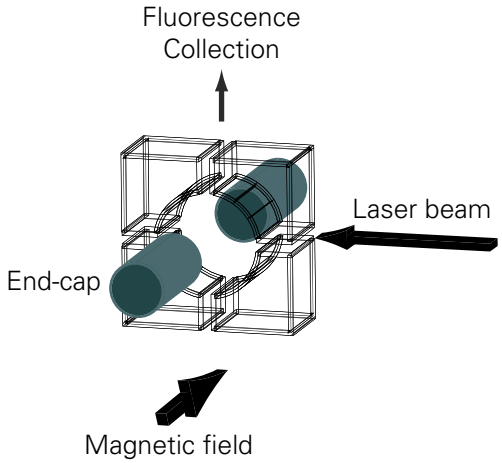

Figure 4.1: The split-ring trap. The magnetic field direction is along the axis of the trap, defined by the two end caps.

One draw-back to the trap design is the difficulty in aligning the end caps. It is necessary to have the same end-cap–ring separation for both end caps. Without this, the ions will be trapped off-centre and may not be easily imaged through the hole in the ring electrode. The angle of the end caps is also difficult to adjust. These difficulties were not due to a bad design as such, but rather to a reliance on the manufacture reaching certain tolerances. These tolerances were not met when the trap was fabricated and later adjustments would have either required the re-machining of the trap or changes to be made to allow

greater freedom of movement of the components. We chose the latter route when changes became necessary as described below.

There were three failed attempts at trapping calcium ions. After each failed attempt it was necessary to open the vacuum system, inspect the trap looking for faults, attach new ovens and filaments and bake out the entire vacuum system again. This process is very time-consuming, taking at least two weeks to complete.

The cause of the first failure to trap ions was a bad connection to an oven and misalignment of the end-cap electrodes. The second attempt to trap ions was rather longer, taking around six months. It coincided with the move to temporary lab space and required a very thorough investigation of all the laser systems, electrical systems, etc. After this attempt, the main focus of investigation was the pair of atomic-beam ovens (see Section 4.1.1).

After the third, and much shorter, attempt the vacuum system was opened again and the alignment of the end caps examined. It was considered that the misalignment may have been sufficiently severe to either prevent trapping or else for the ions to have been trapped off-centre. The end-cap holders were re-machined in-house to allow a greater lateral movement and to facilitate easier adjustment of the angle. In addition, shields were designed by R. Hendricks to fit onto either side of the ring-electrode frame in order to reduce laser scatter and improve the signal-to-noise ratio of the fluorescence-detection system. Calcium ions were succesfully trapped on the first attempt. We conclude that the most likely cause of failure in the attempt immediately previous to this, and perhaps in the other unsuccesfful attempts, was the misalignment of the electrodes.

4.1.1 Atomic-beam ovens

Initial attempts to work with calcium immediately followed the magnesium experiments described in Chapter 7. Although attempts were made over a long period to trap calcium ions, none were observed. One possible reason for this was that an insufficient flux of calcium atoms was being emitted into the trap during the loading procedure.

The temperature at which an atomic-beam oven must operate is determined by the vapour pressure of the atomic species. For example, the vapour pressure of calcium is 10^{-10} mbar at a temperature of 230 °C. The maximum rate of

sublimation is given by [69],

$$\Gamma = 2.63 \times 10^{24} \frac{PA}{(MT)^{1/2}} \ , \tag{4.1}$$

where P is the vapour pressure of the atom, A is the surface area of the emitter, M is the atomic weight, and T is the temperature of the sample. The vapour pressure increases with temperature and curves showing the relationship can be found in Ref. [69].

The temerature reached after a given time depends greatly on the oven design (its length, mass, material) and on the way in which it is mounted (e.g. the lens-holder on which it is mounted will act as a heat sink), and the current passed through it. Although data were avilable for previous calcium ovens these were of a different design and mounted in a different trap. Currents and loading times for magnesium ovens were available for the split-ring trap but the vapour pressure of magnesium is rather different (about 10^{-10} mbar at 130 °C). In addition, the magnesium ovens were a slightly different size than the calcium ovens which complicated calculations further.

The oven design used with the split-ring trap is shown in Fig. 4.2. It consists of a tantalum tube (I.D. 0.86 mm, O.D. 1.02 mm, Goodfellow 713-389-90). A tantalum wire (ϕ 0.25 mm, Advent TA550515) is spot-welded at several points along the length of the oven. A pin is used to make a hole in the middle of the oven on the side opposite the spot-welded wire; this is done after the calcium has been placed in the oven. A current passed through the tantalum wire will heat the tube, primarily through heating of the wire and conduction to the tube. In order to reduce the possibility of heat-induced oxidation of the calcium, the wire was spot-welded to the tube before the calcium was added. One end of the tube was crimped; the other was left open and small particles of calcium (purity 99.0%, natural abundance, Goodfellow 445-429-08) were placed in the tube with tweezers.

Tests were carried out in a small vacuum chamber with a viewport. The ovens were tested, both hanging in free-space and mounted to heat sinks. Two different lengths of tube were tested: 8 mm and 10 mm. 8-mm ovens required lower currents as the mass to be heated was lower, and the total resistance was higher. The test procedure for the ovens involved running them at a low current (~2 A) for long times (up to one hour) in order to complete the "de-gassing". During this time the pressure in the test chamber would rise from approximately

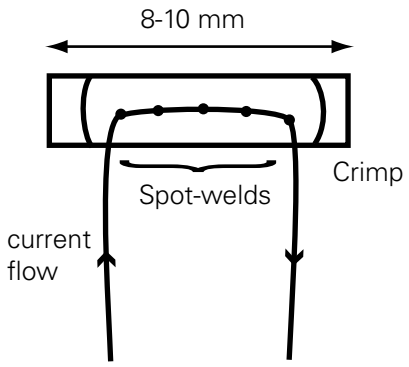

Figure 4.2: The atomic-beam oven design. A tantalum wire is spot-welded along the length of a tantalum tube. The ends of the tube are crimped shut with a small pair of pliers. A hole is punched in the middle of the length of the oven, opposite the spot-welds. The heating of the oven is primarily through the resistive heating of the wire and conduction to the oven.

10^{-7} mbar to 10^{-5} mbar and fall back to the original value slowly. At low currents the wire and oven did not change colour. Above 3 A a slight glow was observed. At high currents ($\sim$3.45 A) the wire and oven glowed bright orange and a black spot formed on a piece of copper which had been placed opposite the hole in the oven. The spot changed to a silver colour after a few seconds of deposition. The calcium deposit disappeared over the course of about 30 sec after air was allowed into the test chamber.

The effect of mounting the oven on a copper-block heat sink was to raise the current required to cause deposition to 3.8 A for a similar oven. Although this is only a 10% change it is very significant as the heat generated rises quadratically with current.

Although the easiest measurement to make when testing an oven is to look for the current required for deposition, the oven will typically be run at a much lower current in ion-trapping experiments. Deposition of the element on the trap electrodes would not only deplete the oven very rapidly but it would also give rise to patch potentials on the electrodes due to the different work functions of the calcium and the beryllium copper. Typically the ovens are run at 1.4 A when loading small clouds of ions. This current is lower even than that used for de-gassing.

Another interesting point to note regarding the ovens is that in tests they were still found to work after being exposed to air for 12 hours. It had been

feared that a calcium compound (calcium oxide or calcium hydroxide) might form on the outside of the calcium granules, thus raising the temperature required to generate a beam of calcium from the oven. This was clearly not the case. It was found, however, in dissecting ovens from an early, failed, trapping attempt that the calcium inside the oven had turned white. This may have been due to the ovens being exposed to air for longer than 12 hours. Alternatively, it may have been due to spot-welding the ovens after they had been filled with calcium, a method which was later changed as described above.

4.1.2 Vacuum system

Most ion traps are used under vacuum conditions, with the exception of large-scale dust-particle traps used mainly for demonstration purposes. The vacuum requirements for Penning traps are even more stringent than those for RF traps due to the unstable nature of the magnetron motion. Stray background gas particles can cause an ion to rapidly exit the trap through collisions. Although buffer gas cooling (with a pressure of order 10^{-4} mbar) is used in Fourier-Transform Ion Cyclotron Resonance (FT-ICR) experiments with Penning traps, this is only in conjunction with an axialising field. Where buffer-gas cooling is not being used the trap is normally operated at ultra-high vacuum (UHV) conditions with a background pressure of 10^{-9} mbar or less.

In order to attain UHV conditions a very clean vacuum vessel is required. Dirt and especially oily residue can outgas for a long time (as long as years). For this reason, great care is taken in cleaning the vacuum vessel and all components that will go inside it.

For cleaning metals a liquid known as "Copper Brite" (75% concentrated nitric acid, 23% concentrated sulphuric acid, 2% concentrated hydrochloric acid) is used, diluted to one part in four with distilled water[1]. It is similar to a commercially-available cleaning agent. When used for cleaning copper trap components it reacts with the top layer of the metal or, more usually, with the metal-oxide layer. The reaction is very dependent on temperature and the best results were obtained by heating the Copper Brite solution in a water bath before use. When hot, dipping a trap component in the solution for about 10 sec was sufficient to leave it clean and shiny.

[1]The popular rhyme "Here lies Jill, still and placid: added the water to the acid" is a useful reminder of the *wrong* way in which to dilute Copper Brite. Adding acid to water avoids rapid local boiling and splashing.

On one occasion, on which the Kapton coating on the wires to the filaments had been overheated in operation, a black coating had formed on the electrodes. The overheating of the Kapton was obvious since it had discoloured and partially detached from the wire. The coating on the electrodes could not be removed with Copper Brite. On the assumption that it was indeed a Kapton coating, the electrodes were soaked in a strong sodium hydroxide solution to soften the Kapton and were then rinsed in water before being dipped in Copper Brite again. This procedure was successful in removing the coating and leaving the electrodes clean.

After cleaning with Copper Brite, metal components were cleaned in acetone in a warm ultrasonic bath; this removed any greasy residue. Ceramic components were also cleaned in the same ultrasonic bath. Finally, Kapton-coated wires and all other items (such as the stainless-steel pedestal on which the trap sits) were wiped with methanol.

The inside of the vacuum chamber itself was cleaned with Copper Brite where deposits had formed from previous use of atomic-beam ovens. This was followed by wiping with acetone and methanol.

The assembled vacuum chamber is depicted in Fig. 4.3. It consists of three four-way crosses. On the first four-way cross are the laser-input and fluorescence-output windows along with the electrical feed-through. The trap sits in the centre of this cross. The second four-way cross has connections to an ion gague and a Brewster-cut exit window. The plane of the third cross is at right angles to plane of the other two crosses. The third cross has a valve for connecting the system to a turbo-molecular pump, a leak-valve with a blanking plate attached and the $33\,\mathrm{dm}^3\,\mathrm{sec}^{-1}$ ion pump (Varian VacIon Plus 40 Triode). The leak valve can be used for letting in small amounts of gas to the chamber when performing ion-detection experiments with resonant electronic circuits [65].

The pumping system was based around a Pfeiffer TPU-062 turbo-molecular pump backed by a low-capacity rotary pump. The turbo-molecular pump could pump out the vacuum chamber to about 1×10^{-9} mbar on its own. In order to achieve such a low pressure in a reasonable time-scale the temperature of the whole vacuum chamber was raised to over $200\,^{\circ}\mathrm{C}$ for a period of several days. The heating was done by fin heaters and the entire system was enclosed by an aluminium surround. At this high temperature the vapour pressure of any remaining contaminants was higher than the background pressure in the system and they quickly entered the gas phase and could be pumped out. Thus,

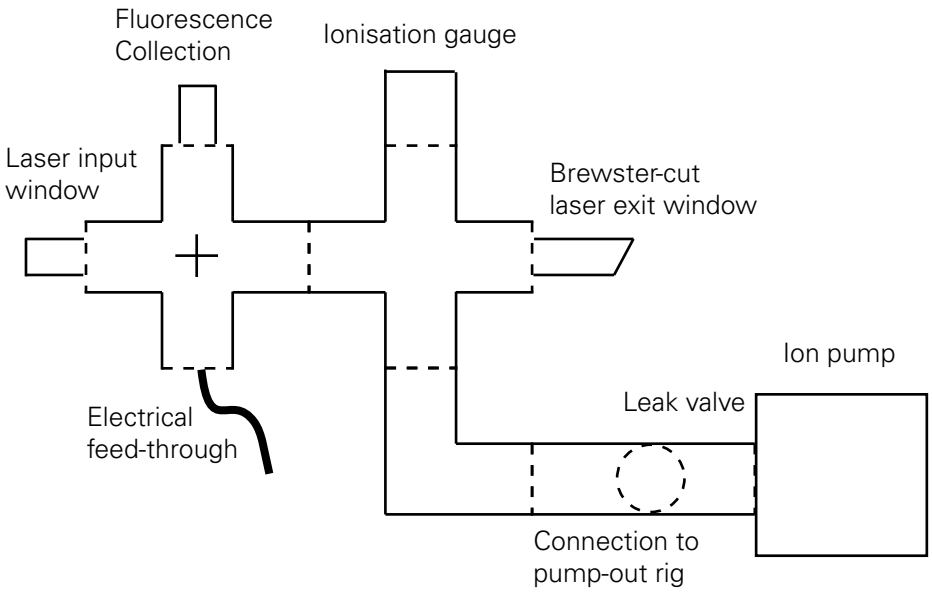

Figure 4.3: Schematic of the vacuum chamber. The three four-way crosses are shown (the right-most is perpendicular to the others with the arms into and out of the page respectively. A small cross denotes the position of the trap; the magnet pole pieces are tightened against the outer wall of the chamber at this point.

the system's pressure was reduced from typically 5×10^{-7} mbar before baking to about 1×10^{-9} mbar afterwards. This reduction would have taken months to achieve had the system not been heated. After the rate of pressure change had dropped almost to zero, the temperature was slowly reduced to prevent temperature-induced stress on the glass windows.

The ion pump was engaged during the cooling period. This reduced the pressure of the system by a further order of magnitude to about 1×10^{-10} mbar. The pressure could be reduced, in principle, by another order of magnitude again through the use of a titanium sublimation pump. Such a pump works by evaporating a film of a titanium onto a surface. This surface adsorbs light atoms and molecules such as helium and hydrogen which are not as effectively pumped by an ion pump. It has not been found necessary to achieve a higher vacuum to date but the inclusion of such a pump may be considered when setting up vacuum systems for new traps.

Cyclotron	ω_c	$2\pi \times 376.1\,\mathrm{kHz}$
Axial	ω_z	$2\pi \times 140.5\,\mathrm{kHz}$
Magnetron	ω_m	$2\pi \times 28.39\,\mathrm{kHz}$
Modified cyclotron	ω_c'	$2\pi \times 347.7\,\mathrm{kHz}$

Table 4.1: The motional frequencies of Ca^+ for the typical operating parameters of a magnetic field of 0.98 T and 4 V trap bias. All frequencies are quoted as angular frequencies.

4.1.3 Magnet

The vacuum system sits between the pole pieces of a water-cooled electromagnet (Oxford Instruments 4-inch Electromagnet, Newport Series N100). The pole pieces are retractable and when the vacuum chamber is in place they are screwed in so that they touch the outside of the chamber. The magnet is driven by a supply (Oxford Instruments SCT stabilised power supply) with current-stabilised output capable of supplying up to 16 A. At 15 A the magnetic field generated is 0.98 T. The magnetic field has been found to be stable to around one part in 10^4 over the course of several hours.

With a magnetic field of 0.98 T and a typical trap bias of 4 V, the theoretical motional frequencies were calculated. These are presented in Table 4.1.3.

4.2 Laser-cooling of Ca^+

Calcium ions have been laser-cooled in several other groups in RF traps to date [70, 71, 72]. There are several additional challenges associated with cooling calcium in a Penning trap as opposed to an RF trap. Although calcium is a group-II element like beryllium and magnesium, which have also been cooled in a Penning trap [23, 73], it has D states between the upper and lower levels of the cooling transition and the ions must be repumped from these states. A simplified energy-level diagram is shown in Fig. 4.4

4.2.1 The cooling cycle

Many successive interactions between an atom and a cooling laser are required to cool from typical starting temperatures to the Doppler limit. Even after this initial cooling, there will inevitably be heating effects on the atom and a high

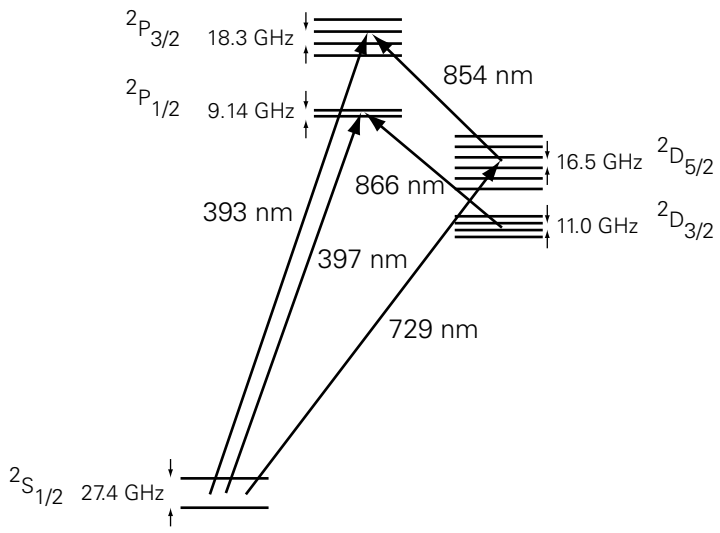

Figure 4.4: Simplified level scheme of Ca^+ (not to scale). Zeeman splittings are shown for a magnetic field of 0.98 T. Lasers are available in our experimental set-up for all wavelengths except for 393 nm.

rate of interaction must be maintained if the atom is to be kept near the Doppler limit.

In order to achieve these high cooling rates, a cooling transition with a low lifetime is chosen. In the case of Ca^+ ions this transition is the $^2S_{1/2}$–$^2P_{1/2}$ dipole transition. The Zeeman splitting of the lower level at 0.98 T is 27.4 GHz. There are four transitions between the lower and upper levels and in order to close the cooling cycle two of these must be addressed. We chose to address the π transitions since these are the most closely spaced of the four. This allowed a greater ease of tuning when the lasers had to be tuned from the zero-field transition to the Zeeman-shifted transition at 0.98 T (see Section 4.5).

Decay from the $^2P_{1/2}$ level to the $^2D_{3/2}$ level is possible. There are six possible transitions between these two levels. Four of them must be addressed in order to close the cycle. It is most convenient to choose the two pairs of σ transitions as will be discussed in the next section.

Finally, a small amount of population will be found in $^2P_{3/2}$ level which can decay to the metastable $^2D_{5/2}$ level. A single laser on the transition between these two levels is sufficient to repump any ions that end up in the $^2D_{5/2}$ level.

4.3 Lasers

One of the main reasons for choosing Ca^+ ions for laser-cooling is that the relevant transitions can be addressed by diode and solid state lasers. This is greatly preferable to using dye lasers such as that described in Chapter 7 since they are much easier to run. A schematic of the laser systems used is shown in Fig. 4.5.

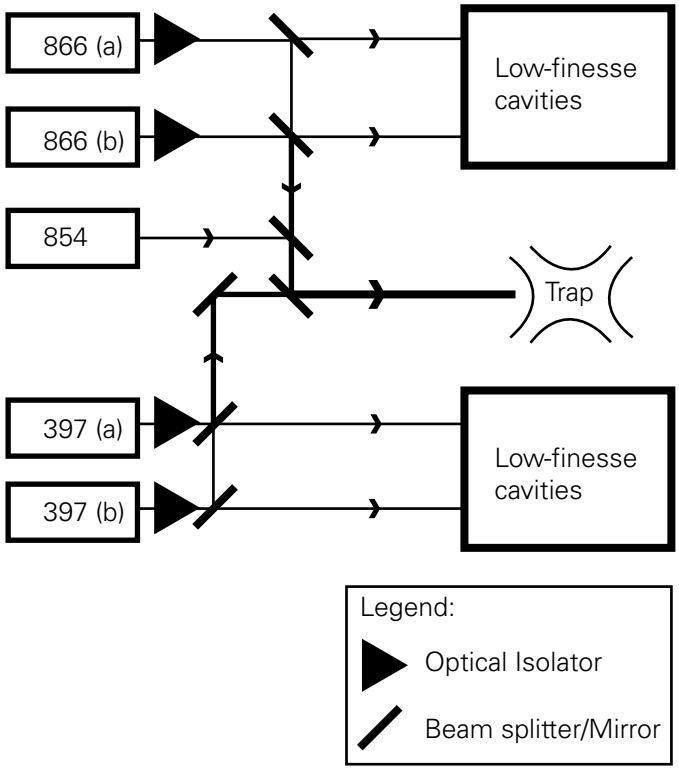

Figure 4.5: Schematic of the laser system used for cooling Ca^+ in a Penning trap. Two cooling lasers (397 (a) and 397 (b)) are used to address the two Zeeman-split $^2S_{1/2}$ states. Two re-pumper lasers at 866 nm are used (866 (a) and 866 (b)); each of these has sidebands imposed to repump from the four Zeeman-split $^2D_{3/2}$ states.

4.3.1 Extended-cavity diode lasers

Laser diodes are not used in a 'bare' configuration for laser spectroscopy or cooling experiments since their bandwidth is too large and the frequency stability

is too low. The output frequency is strongly affected by the ambient temperature through change to the free-running wavelength and the length of the diode. Although active stabilisation can be used to control the frequency stability, it is necessary to first confer a greater inherent stability with passive techniques.

One of these passive techniques is to place the diode in an external cavity. The cavity can then be temperature-stabilised, providing a better level of frequency stability. The cavity also serves to reduce the linewidth of the laser since the width of the mode of the external cavity is rather smaller than that of the diode itself as shown in Fig. 4.6.

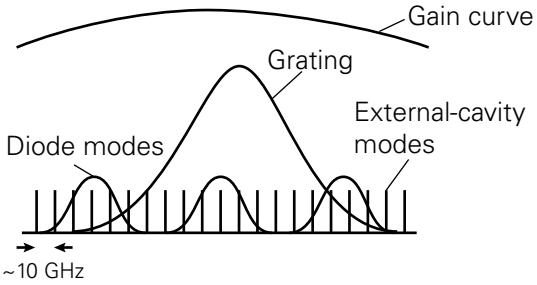

Figure 4.6: The gain curve and modes present in an ECDL (after Toptica DL100 data sheet). The grating selects one of the diode-cavity modes. Then, within a single diode-cavity mode, a single extended-cavity mode will dominate. Thus, the normally broadband output of the diode is dominated by the narrow-bandwidth extended-cavity mode.

A diode in a cavity is referred to as an external-cavity diode laser (ECDL). Two popular optical layouts for ECDLs are the Littman and Littrow configurations. They are similar but the Littrow has a better inherent stability against acoustic noise due to fewer components. The Littrow configuration is shown in Fig. 4.7. Several papers have been published describing in technical detail particular implementations of the Littrow configuration. The most widely-used in the field of laser-cooling are the Boshier design [74] and the Hänsch design [75]. The commercial diode lasers we use follow the Hänsch design while the lasers which were built in-house follow the Boshier design. There is little to choose between them in terms of inherent stability against thermal and acoustic noise.

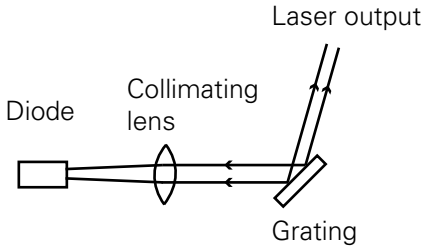

Figure 4.7: The Littrow configuration of an ECDL (after Toptical DL100 data sheet). The first order light reflected from the grating is directed back into the diode while the zeroth order forms the laser output.

4.3.2 Cooling lasers

The lasers addressing the cooling transition are TUI Optics (now Toptica) DL100 laser heads with TUI Optics control electronics. The control electronics are designed in a modular way and we have the temperature control, current control, piezo scan-control, and PID (Proportional-Integral-Differential) feedback modules.

The diodes themselves are Nichia NLHV3000E and NLHV500A, both with a free-running wavelength near 397 nm. The diodes have different intensities (30 mW and 5 mW respectively) and the higher-power diode does not have an internal photodiode for intensity-stabilisation. In other respects they are identical. In normal operation their output intensities were found to be typically 9.3 mW and 2.9 mW respectively. Neutral density filters were used to even out the intensities before combination such that the final powers entering the trap were approximately 400 μW and 600 μW respectively. Control over the intensity of the combined beams was provided by a wave-plate ($\lambda/2$) and a polarising beam splitter. Rotating the wave-plate allowed the smooth reduction of laser power, necessary for the loading of small clouds of ions.

The linewidth of the lasers was measured using a heterodyne technique by K. Koo and J. Sudbery [65, 66]. The two lasers were tuned close in frequency to one another and the beams overlapped onto a high-speed photodiode. A fast Fourier transform (FFT) was taken of the signal and the difference frequency between the two beams was seen as a prominent peak. The width of that peak was a measure of (twice) the linewidth of the lasers. The linewidth was found to be 2 MHz.

As with any laser, there is a certain amount of amplified spontaneous emission (ASE). This is light close in wavelength to the lasing wavelength but which is not coherent. It arises from amplification of a photon on a single pass through the laser cavity. The ASE from the cooling lasers contains some light at 393 nm which can excite the ions from $^2S_{1/2}$ to the $^2P_{3/2}$ state. From there it can decay to the long-lived $^2D_{5/2}$ state. It is for this reason that a repumper at 854 nm is required. This becomes much less problematic at the low power level normally used for small numbers of ions as the ASE is almost eliminated by the strong attenuation.

4.3.3 Re-pumping lasers

The infrared lasers were all built in-house following the Boshier design [74]. They are described in detail by Sudbery [66]. Diodes were not available at exactly 866 nm or 854 nm so diodes with a free-running wavelengths close to the desired wavelengths had to be used and tuned by controlling their temperature. The feedback provided by the external cavity also helps greatly in "pulling" the diodes from their free-running wavelength to the desired wavelength for the experiment. The greater the feedback the further from the free-running wavelength the diode can be tuned. A large amount of feedback into the diode has the disadvantage of reducing the output power of the laser, however.

The 854-nm laser used a SDL 5411 diode with a free-running wavelength of 852 nm and a typical operating power of 13 mW. It was found that with small clouds of ions, so little population entered the $D_{5/2}$ state that it was not necessary to tune the laser using the piezo. Simply setting the current to within 10% of the optimal value (known from measuring the wavelength on the wavemeter) was sufficient to saturate the transition. No difference was observed whether the laser ran multi-mode or single-mode.

The 866-nm lasers used SDL 5411 and LD1377 diodes respectively with typical output powers of 7.15 mW and 3.33 mW. Their nominal free-running output powers are both approximately 100 mW indicating that a large percentage of the available power must be fed back in order to achieve the desired wavelengths.

Previous experiments had used four separate 866-nm lasers to address the $D_{3/2}$–$P_{1/2}$ transition. This method, as well as being time-consuming to set-up and keep tuned to resonance, would have the disadvantage of requiring four sep-

arate locking cavities in order to stabilise the lasers. Mixing the $D_{3/2}$ sub-levels using a single 866-nm laser and microwaves was attempted but was unsuccessful [65].

A different, elegant, method of reducing the number of 866-nm lasers required has been found to work very well [65, 24]. The principle is to modulate the input current to the laser diode at high-frequency [76, 77]. This will result in a central peak or "carrier" and a "sideband" on either side of the carrier, separated from the carrier by the modulation frequency. Thus, the laser can be tuned to half-way between a pair of transitions and modulated at half the frequency separation of those two transitions. The sidebands will then be resonant with the transitions and the carrier will not be involved in repumping.

The frequency of the carrier should not, in principle, be different from the output of the unmodulated laser; the only effect should be a decrease in its intensity. In fact, the modulation perturbs the normal running of the laser and the wavelength is found to be different when measured with the wavemeter. The wavelength measured is found to be stable for reasonably high modulation powers, however.

The sidebands are due to both amplitude- and frequency-modulation. The number of sidebands due to frequency modulation may, of course, be greater than two if the modulation depth is particularly high as with the sidebands on a moving atom's spectrum (see Chapter 2). In practice, it would be difficult to achieve such a high modulation depth and we have not observed it.

The use of the current-modulation technique restricts the choice of transition. It is much easier to address the two pairs of σ transitions (Fig. 4.8) as they are closer in frequency than any other choice of four frequencies that would close the cooling cycle.

The modulation is provided by an Agilent 83650L function generator at about 15 dBm. The output is split by an RF splitter and sent to the two laser diodes via semi-rigid cable, a stub tuner for impedance matching, a bias-tee for the addition of the normal current drive, and a length of rigid cable soldered directly to the diode-holder pins. It has been found that certain diodes work much better with this technique than others. It is thought that this may be to do with impedance-matching.

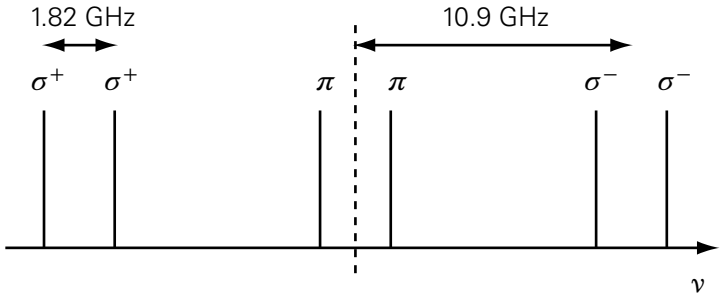

Figure 4.8: The transitions from $^2D_{3/2}$ to $^2P_{1/2}$ at 0.98 T. The 866-nm lasers are tuned half way between each pair of σ transitions and sidebands are applied at half the separation of the transitions. Thus the sidebands are resonant with the transitions while the carrier does no work.

4.3.4 Quadrupole-transition laser

In the experiments performed to-date in the group of R. Blatt [39], the $^2S_{1/2}$–$^2D_{5/2}$ quadrupole transition has been used to represent the qubit. It is also a suitable transition to use for sideband-cooling. The lifetime of the upper level has been the subject of several experiments [78, 71, 79] and found to be 1.1 s. Since the transition is forbidden, a narrow-linewidth, high-power laser is required to address it. A Ti:Sapphire laser is used for this purpose.

At Imperial College, the Ti:Sapphire laser is a Coherent MBR-100 with a quoted linewidth of 100 kHz. It is pumped by a Coherent Verdi V5 with approximately 5 W of power at 532 nm. The output of the MBR-100 is about 300 mW at 729 nm. The linewidth of the Ti:Sapphire laser is found to be closer to 1 MHz than the quoted linewidth of 100 kHz. The quoted linewidth was inferred from an examination of the error signal with the laser locked to the the (relatively) unstable Coherent reference cavity.

In order to resolve the sidebands on such a narrow transition it is necessary to achieve both a very stable frequency and a very narrow linewidth. Work was begun on a Pound–Drever–Hall lock (see Chapter 9) for the laser by J. Sudbery [66] and is currently being completed by H. Ohadi. Moving away from the poorly-understood loop filter with bypass topology [80] to a more conventional loop filter based on a high-speed PID (Proportional Integral Differential, see Section 4.4.3) circuit appears to have brought the lock much closer to completion. The loop filter controls the response and the bandwidth of the feedback loop and the need for bypass topologies is less with the availability of faster

op-amps.

I note here that the locking scheme is very similar to that on which I worked in Innsbruck and which is described in Chapter 9. The main difference is the servo mechanism. The MBR-100 does not have room for an intra-cavity electro-optic modulator (EOM). As a result, the feedback is to an external double-pass acousto-optic modulator (AOM). The main disadvantage to this technique is that the bandwidth is limited by the propagation time of the sound waves in the AOM crystal. The laser itself should be more inherently stable than the CR-899 described in Chapter 9 due to its monolithic design.

4.4 Locking cavities

Although ions had been trapped and laser-cooled in a Penning trap previously, it was necessary to lock the lasers' frequencies in order to trap small clouds and, ultimately, a single ion. The cooling lasers' frequency drifts were measured to be 300 MHz per hour. In order to compensate for this slow but relatively large-scale drift it was decided to implement a side-of-fringe lock. This type of lock, along with some background theory of laser locking, is described in more detail in Chapter 9. I will concentrate on the practical aspects of the cavity design and building here.

Cavities formed by a pair of curved mirrors and a spacer form excellent frequency discriminators. The transmission of the cavity varies with frequency, with peaks seen at regular intervals where the length of the cavity matches an integer number of half-wavelengths of light. This interval is known as the free spectral range (FSR) of the cavity and is given by,

$$\mathrm{FSR} = \frac{c}{2L} \,, \tag{4.2}$$

where L is the length of the cavity. The peaks are not delta functions and their width is determined primarily by the reflectivity of the mirrors. The width of the peak relative to the separation of the peaks (the free spectral range) is called the finesse. The finesse is related to the reflectivity, r, of the mirrors,

$$F = \frac{\pi\sqrt{r}}{1-r} \,. \tag{4.3}$$

A high finesse (narrow peak) gives rise to a large slope on the side of the

peak. If this side of the peak is used as a frequency reference for a laser the high slope gives a large gain in feeding back to the laser frequency controls and a tight lock. Conversely, a low finesse will give a shallower slope but will allow further excursions in frequency while still giving a usable signal to return the laser to the desired lock point, i.e. it has a bigger "capture range". The lock will not be as tight in frequency terms, however.

A high finesse was not required for this laser lock as the goal was to prevent long-term frequency excursions rather than to remove high-frequency laser noise. Mirrors were chosen to give a finesse of about 60. They were custom-made by CVI Technical Optics, part number PR1-866-95-0537-0.30CC (for the 866-nm laser cavities) and PR1-400-95-0537-0.30CC (for the 397-nm laser cavities). The radius of curvature was 30 cm.

The cavity spacer was constructed from standard quality Zerodur (supplied by HV Skan Ltd.), a glass with a very low coefficient of expansion. Although it doesn't have quite as good a thermal stability as ULE (ultra-low expansivity glass) it is considerably cheaper and its performance is certainly adequate for this application given the thermal stabilities of the other components present in the design. The Zerodur was cut in-house according to my design (Fig. 4.9).

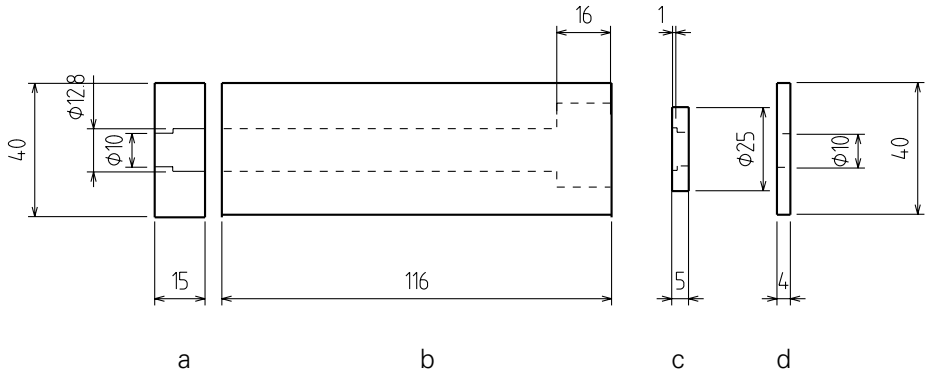

Figure 4.9: Re-entrant cavity design. The mirrors are mounted in the recesses in parts *a* and *c*. Part *c* is connected to *d* with three piezos. Part *d* is connected to *b* with another three piezos. *a* is glued to *b* with Torr Seal. The set of two cavities for the 397-nm lasers differs slightly from those used for the 866-nm lasers in having a smaller outer dimension (40 mm square as shown above compared to 50 mm square).

It was desired that the cavities should be tunable, i.e. their length could be altered to change the resonant frequency. A laser locked to the cavity would

follow this change in resonant frequency due to the electronic feedback to the laser control. Tunability of a cavity is normally achieved with piezo-electric transducers (PZTs or "piezos") which have a range of motion of about ten microns. Changing the PZT's length by half the wavelength of the light will tune the cavity over a full free spectral range.

Unfortunately PZTs do not have particularly good temperature stability. For this reason a re-entrant design was employed: there are two sets of counteracting PZTs. Expansion in one set, due to a rise in temperature, causes the end-plate to move away from the main block slightly. The second set of PZTs, connected to the mirror holder, should expand by the same amount, however, and there should be no net movement of the mirror.

The length of the cavity was originally designed to allow for the locking of two lasers to the one cavity. This can be done by rotating the polarisation of one of the lasers such that the two beams have orthogonal polarisations. The polarisations can be separated again at the far side of the cavity and detected on separate photodiodes. If this scheme were used it would require that the lasers were locked to different longitudinal modes of the cavity separated by the frequency difference of the two cooling transitions. Thus, it was necessary to choose a long enough cavity that the difference in longitudinal mode number was quite a large integer. This would give the resolution of locking point required to allow for the operating magnetic field to be changed at a later date. The splitting of the two cooling transitions at 1 T is 18 GHz while the splitting at 2.5 T (the maximum field of the superconducting magnet available for future experiments) is 45 GHz. A cavity length of 10 cm gives a FSR of 1.5 GHz and a difference in longitudinal mode number of 12 and 30 respectively for the two fields quoted.

The scheme with orthogonal polarisations was successfully demonstrated with the 397-nm lasers. A separate cavity for the second laser was soon constructed, however, as the scheme required a greater magnet stability than was available from the electromagnet. Although greater field stability would be possible with a superconducting magnet in future there would be little scope for altering the field by a small amount such that the Zeeman splitting was a multiple of the cavity mode-spacing.

The PZTs (PI open-loop LVPZT 15 μm, current model equivalent P-802.10) were glued to the Zerodur using Varian Torr Seal, a vacuum-compatible resin sealant. Slight imbalances in the amount of Torr Seal used at different joins will

be one of the dominant causes of thermal instability of the final cavity. Great care must be taken to use as small an amount of the epoxy as possible and to ensure that the same thickness is applied at each join.

4.4.1 Temperature stabilisation

The coefficient of expansion of Zerodur is very low. Even if it were zero, however, it would still be necessary to keep the cavities at a stable temperature. This is because of the rather higher expansion coefficients of epoxy and the mirror substrates. To stabilise the temperature, the cavities are placed in a large diecast aluminium box with a 13-mm layer of expanded Neoprene insulation on each side. Heating wire was wrapped around the cavity and the box was then covered with solid Neoprene (13 mm thick).

A thermistor under the outer layer of Neoprene was used for temperature measurements. It was connected to one arm of a Wheatstone bridge and the voltage across the bridge was connected to a differential amplifier to give a sensor signal. The sensor signal was connected to a PI (proportional–integral) circuit and then to the servo mechanism, a high-current transistor controlling the current to the heater wire around the diecast box. Although the servo can only heat, and not cool, it is found to be very effective at keeping the temperature constant to about 0.1 °C.

The sensor and PI circuits were designed by R. Rowley in the National Physical Laboratory and are described in Ref. [66].

4.4.2 Mode matching

A cavity can support both longitudinal modes and transverse modes of the electromagnetic field internally. This discretisation appears because the phase of the field after a round trip of the cavity must match the field after just one traversal or else destructive interference will take place. The transverse modes are given names based on the number of nodes in the mode in each of the two transverse axes. Thus TEM_{00} has no node in either axis and has a simple Gaussian profile, while TEM_{01} has a node in the y-axis and two bright spots are seen if it is viewed on a CCD camera.

The transverse modes differ in frequency from one another. It is desirable that only one of them shoud be visible in each free spectral range so that there is no danger of the laser losing the lock on one peak and locking to another

instead. It also reduces the amount of power required for the lock since the power of the laser is not distributed between several transverse modes.

In order to concentrate the power of the laser in one mode it is necessary to match both its waist size and position to the natural waist of the cavity. Our cavities are symmetric and thus the longitudinal position of the mode will be at the mid-point of the cavity. The width of this mode can be calculated from [81],

$$w_0 = \sqrt{\frac{\lambda}{2\pi}}\,[L(2R-L)]^{1/4}\ , \tag{4.4}$$

where λ is the wavelength of the light, L is the length of the cavity, and R is the radius of curvature of the mirrors. The waists will be 118 μm and 176 μm for the 397-nm and 866-nm laser cavities respectively.

Once the position and size of the cavity mode are known, lenses can be placed in the laser beam before the cavity such that the beam waist overlaps the cavity waist in the longitudinal direction and so that the waist is closely matched in size as well. If the latter condition is not met, light will exist outside the TEM_{00} mode and will be coupled into other transverse modes.

To calculate the focal lengths and correct positions of the lenses, use was made of the "ABCD" matrices for optical interfaces and the ABCD law of Gaussian beams. This law [81] says that the ABCD matrices originally derived from the field of geometric optics can be used to describe the propagation of a Gaussian beam through an optical system. The field of a laser beam is described by [81],

$$E_0(\mathbf{r}) = A\exp\left(ik(x^2+y^2)/2q(z)\right)\exp\left(ip(z)\right)\ , \tag{4.5}$$

where A is the amplitude, z is the propagation direction. $q(z)$ is a function determined by the intensity profile of the beam. For a Gaussian profile $q(z)$ is defined,

$$q(z) = \frac{-i\pi w(z)^2}{\lambda}\ , \tag{4.6}$$

where w is the beam spot-size. p is defined in terms of q:

$$p(z) = i\ln\frac{q(0)+z}{q(0)}\ . \tag{4.7}$$

Just as matrices can be used to show how a beam's height and angle change in geometrical optics (see, for example, Ref. [82]), those same matrices can be used

to calculate the change in q. Thus, we can calculate the beam spot-size at any point in the system.

The relevant matrices can be found in many textbooks such as Ref. [83] and I wrote a short program to calculate the spot-size resulting from the most basic lens system as two parameters were varied. That lens system is shown in Fig. 4.10 along with the relevant matrices. The program was written in Octave [84], a freely-licensed package similar in syntax and functionality to Matlab. The programme is reproduced in Appendix A.

$$\begin{pmatrix} 1 & \mathrm{d}_1 \\ 0 & 1 \end{pmatrix} \qquad \begin{pmatrix} 1 & \mathrm{d}_2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & \mathrm{d}_3 \\ 0 & 1 \end{pmatrix} \qquad \begin{pmatrix} 1 & L/2 \\ 0 & 1 \end{pmatrix}$$

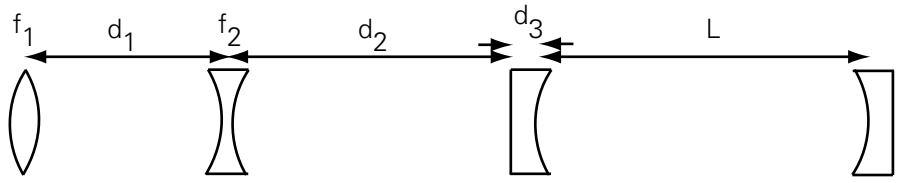

$$\begin{pmatrix} 1 & 0 \\ \frac{-1}{\mathrm{f}_1} & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & 0 \\ \frac{-1}{\mathrm{f}_2} & 1 \end{pmatrix} \qquad \begin{pmatrix} 1 & 0 \\ 0 & \frac{n_a}{n_s} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \frac{n_s - n_a}{n_a R} & \frac{n_s}{n_a} \end{pmatrix}$$

Figure 4.10: The lens system modelled by the program in Appendix A. The matrices for each element of the system are shown below that element. f_1, f_2 are the focal lengths of the two lenses and will be positive and negative values respectively to reflect the lens-types. n_a and n_s are the refractive indices of air and the mirror substrate respectively. d_3 is the thickness of the cavity mirror, R is its radius of curvature and other variables are as indicated in the figure.

The output of the program is the beam spot-size in the centre of the cavity given a variation of the distance between the lens pair and the first mirror of the cavity on one axis and the separation of the lenses themselves on the other axis. The three-dimensional plot shows a valley, the line along the bottom of which gives the parameters at which the beam waist coincides with the cavity waist. As previously described, matching the longitudinal positions of the two waists is only one part of the mode-matching process. The other part is to ensure that the waist sizes are the same. To see this, a plane is plotted showing the cavity waist. The intersection of the bottom of the "valley" and the plane gives the ideal mode-matching parameters (see Fig. 4.11).

When viewed from above (Fig. 4.12), this plot gives a very good intuitive feel to the user of the system as to how easy or difficult it will be to implement

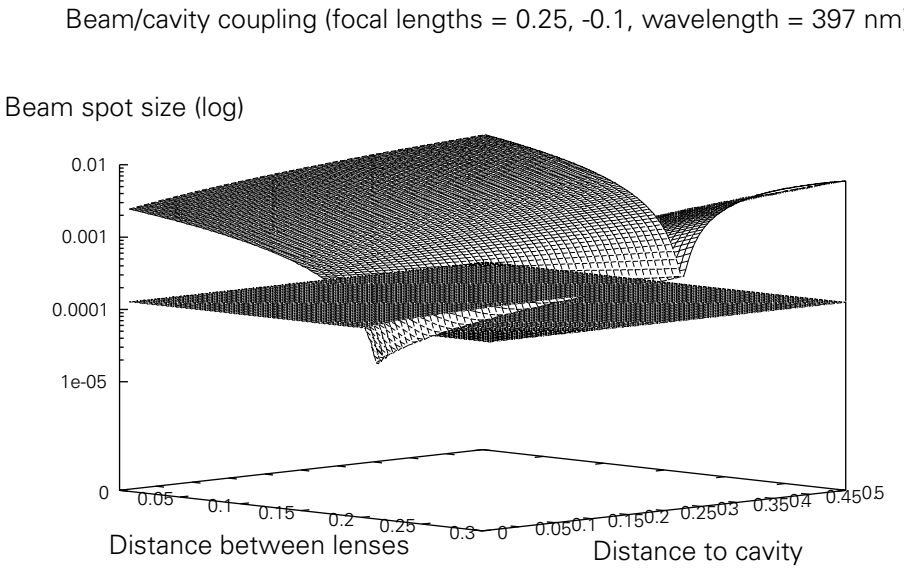

Figure 4.11: Output of mode-matching program (all distances are in metres). When viewed from the side, the beam spot-size at the centre of the cavity can easily be compared with the cavity waist size.

good mode-matching. A roughly triangular shape is seen in the plot which shows the parameters for which the beam waist is smaller than the cavity-waist. If this tapers slowly to a point, it indicates that the mode-matching will not be overly-sensitive to the values of either parameter.

As well as getting the correct waist size at the correct longitudinal point, it is also important that the transverse alignment be good. This is achieved by scanning either the laser or the cavity PZT and then walking the beam across the cavity transverse axes with a pair of steering mirrors until the TEM_{00} mode is seen. This mode can then be optimised in height to the detriment of the other transverse modes by fine adjustments on the two steering mirrors.

4.4.3 Locking the lasers

Stabilising a system's property to a desired value is known as Control Engineering. The property to be stabilised might be temperature (as in the case of the locking cavities), laser frequency (as will be described here and in Chapter 9) or some other property. The desired value is known as the set-point. The difference between the measured value and the set-point is the error and the

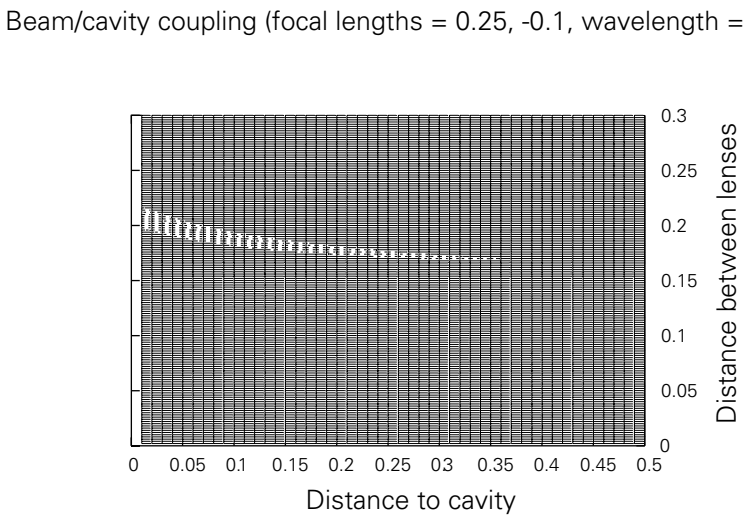

Figure 4.12: Output of mode-matching program. When viewed from the top, the user can get an intuitive grasp of how difficult the system will be to mode-match for the given parameters. A sharp dependence on one or both lens separations may give excellent mode-matching but be very difficult to implement.

mechanism to correct the value (e.g. the heater/cooler in the case of temperature control) is the servo mechanism.

It might be thought that simply feeding the error signal directly back to the servo mechanism (perhaps with some gain) would be sufficient to stabilise the system. This method is known as proportional feedback (proportional to the error signal) and on its own will generally bring the system closer to the desired value. It is found, however, that the system will oscillate around the set-point with such feedback.

In order to force the system to settle to a stable value, some memory of the recent behaviour of the system is required. This is introduced using an integrator circuit in addition to the proportional circuit. An additional benefit of including an integrator circuit is that the resistor and capacitor values used can be chosen to give a frequency cut-off below the first natural oscillation of the servo mechanism. Were signals higher than this frequency to be fed back, they would cause the system to move in the wrong direction (see the discussion of forced, damped harmonic oscillators in Chapter 6).

Finally, although a combination of proportional and integrator circuits can

cause the system to settle to a stable value, there may still be a difference between the stable value and the desired set-point. This difference can be reduced by introducing a differentiating circuit.

Typically a circuit with three separate op-amps is used, one to control each of the three values (proportional, integrating, and differentiating) above. The three signals are then combined and an overall gain is added. The choice of appropriate time constants and gains is an involved process. Some mechanisms exist for determining the ideal values theoretically, by measuring the system's response to a step function. In practice, it is often possible to follow a simplified method of adjusting each property in turn until stable operation is achieved. This latter method was used in the locking schemes described in this thesis.

All of the lasers were locked with PID circuits. In the case of the commercial lasers, the supplied PID modules were used. For the infrared lasers, PID circuits were designed and built by R. Hendricks based on the design in Ref. [85].

The power of the beams picked off from the main beam for the lock was as low as $10\,\mu$W for one of the blue lasers and $180\,\mu$W for the other. A neutral-density filter was used on the output of the latter cavity so as not to saturate the photodiode.

The powers of the infrared laser beams going to the cavities were $220\,\mu$W and $131\,\mu$W respectively.

4.5 Laser tuning

Several diagnostics were available to assist in tuning the lasers to the correct wavelength. The most versatile of these is the wavemeter described in Chapter 5. I describe the others here briefly.

For tuning the blue lasers to resonance with the cooling transition, it was useful to have a separate, reliable, source of calcium ions other than the trap itself. This was provided by a home-made hollow cathode lamp with calcium granules in the cathode. The laser beam was directed through the cathode and absorption spectroscopy was performed using a photodiode at the far side of the lamp. When the laser was tuned on-resonance a dip appeared in the output of the photodiode. This diagnostic was only useful, of course, when loading at zero magnetic field in a Paul trap and when calibrating other diagnostics.

The second diagnostic for the blue lasers was a Fabry-Perot cavity with a free spectral range of $6.3\,$GHz. Shining the laser through and observing the output

while the laser frequency was scanned allowed the detection of mode-hops.

A spectrum analyser (Tec Optics SA-7.5, 7.5 GHz FSR) was used with the infrared lasers. It allowed the measurement of the sideband amplitude with respect to the carrier. It was also useful as an indicator that a red laser had changed from single-mode to multi-mode operation.

The two wavemeters, one commercial and the other built in the group are described in Chapter 5.

4.6 Fluorescence Detection

The ions were imaged through a hole in the centre of the ring electrode. A lens on the top of trap assisted in the collection of light and two further lenses outside the vacuum chamber focused the light onto the entrance window of a photomultiplier tube. The lens system was symmetrical about its centre and is depicted in Fig. 4.13.

An iris is installed above the topmost lens and a pair of filters (Comar band-pass filter, 395GB25 and a Comar low-pass filter, 435IK25) reduce the stray light entering the PMT.

The PMT was a Thorn EMI phototube (Model 9893 QB 350). It was run at 2.3 kV and its output was connected to a discriminator (Electron Tubes AD2). This was then connected to an ECL-to-TTL converter (Electron Tubes ET1) and sent to a multi-channel scaler or MCS card (EG&G Ortec MCS-Plus) in a PC. With the fluorescence viewed on the MCS, the normal time bin was 10 ms. When looking for quantum jumps in small clouds of ions this bin-time was often increased to 20–50 ms. The dark count rate was quoted as 15 counts s^{-1}. The background count rate could be reduced as far as 100 counts s^{-1} when working with small clouds. This was done by excluding as much room-light as possible and by working with low laser powers to reduce scatter from trap electrodes.

A mirror was installed on a Newport Flipper mount after the middle lens to allow re-direction of light towards a camera. The camera was an Andor DH534-18S-03 intensified CCD. The CCD array is 1024×1024 pixels in size, with the side of each pixel being 13 μm. The intensifier gives a maximum gain of about 4320 per photoelectron. The controller card has a 1-MHz, 16-bit A/D converter allowing for rapid image acquisition.

The last lens before the camera was chosen to give a slightly higher magnification than for the PMT. A pair of filters identical to those used for the PMT

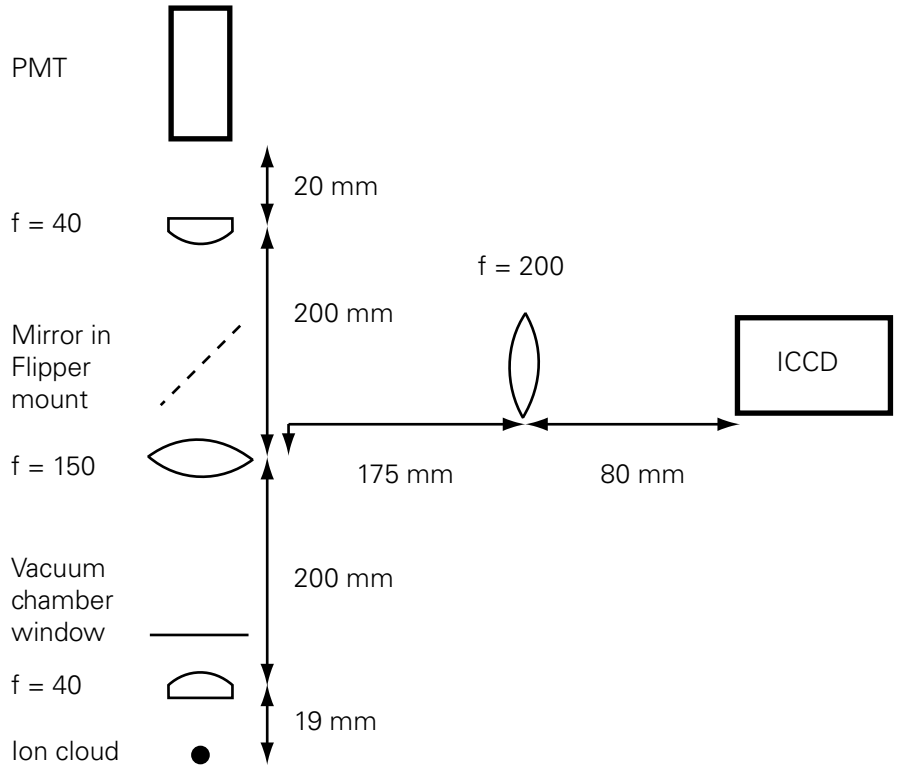

Figure 4.13: The imaging system. A pin-hole aperture (0.5 mm) and filters (Comar 395GB25 bandpass and 435IK25 low-pass with transmissions of 80% and 70% respectively at 397 nm) were placed in front of the PMT. An identical pair of filters (but no pin-hole) were placed in front of the camera. The magnification of the image on the PMT was approximately 1.0 while that of the image on the camera was 1.73.

were inserted immediately before the camera. The magnification of the system was calculated to be 1.73.

Different operating modes were used for different purposes in the course of experiments. "Real-time" acquisitions (single frames taken at regular intervals) of 0.1 s or longer exposure were helpful in identifying changes in the cloud due to the tuning of lasers or changes in axialisation parameters. "Accumulations" (a certain number of frames taken consecutively and added together) allowed for clearer pictures of ions; longer effective exposure times could be attained without risk of saturating the CCD. Finally, "Kinetic Series" were useful in identifying sub-Hertz cloud dynamics. A "Kinetic Series" is a sequence of frames taken in rapid succession by the camera. The exposure time and time between

shots can be controlled in software. As with all of the acquisition types, an important limiting factor on the acquistion rate was the data transfer from the CCD chip to the PC (1 MHz).

4.7 Loading the trap

Ions were formed close to the trap centre by crossed atomic and electron beams. The electrons were accelerated by putting a bias voltage between the filament and the opposite end cap. A voltage was chosen to give the electron an energy close to the first ionisation potential of neutral calcium atoms (6.11 eV). A larger bias (> 11.87 eV) could potentially result in doubly-ionised calcium ions being formed which would not be laser-cooled. A lower bias would mean the cross-section for ionisation would be too low and a very large atomic-beam flux and a long loading time woud be required to load ions into the trap. In fact, bias voltages as high as 13 V were used. No noticeable decrease in the fluorescence level was observed for a lower voltage of 10 V.

Early experiments used a relatively high oven current (of order 3 A), activated for 32 s. During the last 12 s of this period, the electron filament would be turned on at a current of 4.5 A. This was effective in loading large clouds and the number of ions loaded could be reduced by lowering either or both of the oven and filament currents.

The first draw-back to the short loading procedure was that if it was desired to load again afterwards, it was necessary to allow the oven to cool off first. This could take as long as ten minutes. In attempting to load very small clouds, it is often necessary to load ion clouds several times in a relatively short period to gauge the cloud size from different oven and filament currents. The other draw-back is that the atomic flux is likely to be less consistent from shot to shot. A better technique for this type of work, then, is to run the oven at a lower current for a longer time such that a uniform flux of calcium atoms has been achieved when the filament is turned on. The improved technique involved running the oven at a current of 1.45 A for a period of 2 min before turning the filament on for 20 s and then turning both currents off. The cloud size could be controlled via the filament current, ranging typically from 4.5 A for a large cloud to just over 4 A for single ions in a Paul trap.

While loading ions, one of the cooling lasers is scanned in frequency from a few gigahertz below resonance up to resonance. Thus, the very energetic ions

are resonant with the laser when it is far-detuned and are continuously addressed as their velocity falls with the cooling action. The fluorescence has the shape of half-peaks on the MCS output. These can be fitted to a Voigt function to estimate the temperature of the ions [65].

Chapter 5

Wavemeter

5.1 Introduction

It is clear from the description of laser-cooling in Chapter 2 that the lasers must be tuned to the correct wavelengths with some precision. The cooling transition's natural width, for example, is just 22 MHz. In practice, of course, this will initially be Doppler-broadened. Tuning further than a few megahertz from the transition, however, will result in either no signal being observed (detuned to the red) or rapid heating of the ions out of the trap (detuned to the blue). Similarly, it is important to hit the repumping transitions; this becomes even more critical when addressing the transitions with sidebands and consequently with lower power.

Tuning was originally performed with a Burleigh Wavemeter (WA-20). This measures the wavelength with an accuracy of six significant figures. The sensitivity of the photodiode on which the "unknown" laser is incident is much greater in the infra-red than in the UV. Substitution of an op-amp with a larger gain [66] helped in reading the blue laser wavelengths but the resolution was still too low to tune these lasers to the cooling transition: 0.001 nm resolution when measuring a wavelength of roughly 397 nm corresponds to a frequency resolution of 2 GHz.

The procedure developed in the absence of a more flexible wavelength measure at 397 nm was to tune the lasers using the atomic transition itself. This was done with a hollow-cathode lamp containing small fragments of calcium. The plasma containing calcium ions, generated when a large DC voltage was passed through the lamp, absorbed light from the 397 nm laser when it hit resonance

with the $^2S_{1/2}$–$^2P_{1/2}$ transition.

The main drawback with this method, of course, is that there is no magnetic field imposed on the hollow-cathode lamp. The lasers, when tuned, are resonant with the transition with no Zeeman shift. This method, therefore, only allows the correct transition frequency for laser-cooling of ions when the trap is in a Paul-trap configuration. The following procedure may be used to tune the lasers for the correct Penning trap wavelengths. Since the geometry of the Paul and Penning traps is identical, it is straightforward to set up our trap with suitable RF voltages and DC offsets to trap calcium ions in a Paul trap. When ion signal has been detected and optimised in the Paul trap, the magnetic field is turned on to a low value. The lasers are tuned such that the signal level is regained. The magnetic field is increased by another increment and the lasers are tuned again. This is repeated until the full magnetic field (usually around 0.8 T) has been reached. Operating the trap with an RF voltage applied along with an axial magnetic field gives rise to what is known as a "combined trap" [3].

As well as a wish to avoid the time-consuming nature of the tuning method just described, there was another pressing need for the development of a more precise wavemeter. The laser near 729 nm will address a quadrupole transition. This, by its nature, is only weakly allowed and, with a very narrow width, is hard to find when tuning the laser. In order to have a good chance of tuning the laser to the transition it is necessary to start close to it and search over a small range. A wavemeter measuring the 729-nm laser wavelength with an accuaracy of at least one part in 10^7 is thus required.

5.2 Basic wavemeter design

Strategies for measuring CW laser wavelengths include the use of atomic or molecular spectra (see Section 7.2.2 in Chapter 7), the use of high-resolution spectrometers, and various interferometers. In most situations the most convenient method is the use of a Michelson interferometer. Wavemeters using this type of interferometer have been in use almost as long as the laser [86]. A review of wavemeters in general is presented by Snyder and Hänsch [87]. A recent and easily-built design is presented by Fox *et al.* [88].

The principle behind the Michelson interferometer is very simple. As the relative path-length in the interferometer changes, fringes are observed at the output. This is due to constructive and destructive interference and the

fringes observed display a $\cos^2$ pattern provided the intensity in the two arms is matched; if this condition is not met, the contrast of the fringes is reduced. A complete period of this signal is seen when the relative path length changes by the wavelength of the incident light.

The number of fringes passing detectors at the output of the interferometer from the two lasers is recorded as the length of the interferometer is changed. If two lasers are coincident on the interferometer, two fringe patterns will be observed. The number of fringes on each is inversely proportional to the respective wavelength. Thus, the expression

$$\frac{\lambda_U}{\lambda_R} = \frac{N_R}{N_U} \,, \tag{5.1}$$

is obtained. λ_U and λ_R are the wavelengths of the unknown and reference wavelengths respectively while N_U and N_R are the number of fringes observed at the two wavelengths.

If one wavelength, λ_R, is well-known then the other can be easily inferred. Equation 5.1 can then be used to calculate the wavelength of the second laser. It should be noted that the change in path-length need not be known.

To conveniently change the relative path lengths in the interferometer, one generally mounts a retroreflector (cube-corner) on a cart. This cart may then be moved by means such as motors and wires [89], solenoids and an air-track [88], or even on an adapted toy train [90].

Although the use of an air-track might seem quite primitive, it is well-proven and has been used in commercial [91] as well as lab-built wavemeters. The advantage is that the motion is smooth and can easily extend over a long distance (being nearly frictionless). The primary disadvantage is that the wavemeter can not be evacuated to eliminate dispersion errors and variations in refractive index due to air currents. Dispersion effects can be accounted for to a great precision but air currents still have the potential to be problematic. To reduce their effect air is forced out of the bottom of the cart instead of through holes in a track (as would normally be seen in school experiments on momentum transfer). This should give rise to a laminar flow of air away from the beam paths of the lasers.

To further reduce errors due to changes in refractive index, the beams can be made to counterpropagate. The use of a retroreflector introduces a spatial offset to the return beam in the moving arm (see Fig. 5.2). Thus, instead of the secondary output of the reference laser returning to the laser itself, it leaves

the wavemeter spatially off-set. The unknown beam can be sent in along the exit path of the reference beam and the two beams exactly counterpropagate throughout the interferometer. This means that the two beams experience the same changes in refractive index due to air currents.

5.3 Resolution

It is clear from Equation 5.1 that the easiest way to measure the unknown wavelength more precisely is to increase the number of fringes counted. Assuming that the reference wavelength is well-enough known then, say, 10^6 fringes may be counted on the reference laser output and the unknown wavelength may then be stated to a precision of one part in 10^6. Measuring 10^6 fringes of a 633-nm reference laser corresponds to a cart displacement of 16 cm in the the optical set-up depicted in Fig. 5.2. The two limits on increasing the precision in this way are the readout time, which would become proportionally longer, and the coherence length of the two lasers, which must be greater than the length of the longer arm in the interferometer. Of greater relevance in our experiment is the readout time: we wish to use the wavemeter to tune lasers in real-time. One can not, of course, just increase the velocity of the cart and its distance of travel: the detection and processing electronics will have a finite bandwidth. A realistic limit is therefore a fringe rate of about 1 MHz. Higher-speed electronics are possible but require a great deal of attention to issues such as impedance-matching, and commercial design becomes the only viable option. The issue of the coherence length of the lasers is unimportant in our set-up since all of the lasers have linewidths of 3 MHz or less, giving coherence lengths of 100 m or more.

Two main approaches have been adopted in previous designs to improving the resolution of the wavemeter without increasing read-out time. The first is frequency multiplication with a phase-locked loop (PLL). The second is matching the phase of the fringes when starting and stopping counting. From a theoretical perspective these two strategies seem to be almost identical. Experimentally, however, they are rather different.

A phase-locked loop has as its output a signal with a fixed frequency multiple of its input:

$$\sin(\omega t) \longrightarrow \sin(k\omega t)\ , \tag{5.2}$$

where k is an integer constant. The finer comb of the PLL allows the effective interpolation of the signal from the unknown laser to a fraction of the optical wavelength. For example, if k were set to 100 and 10^6 fringes are measured, the use of a PLL would make this equivalent to measuring 10^8 fringes. The two major disadvantages are that the scheme requires a very stable frequency and that PLLs are rather difficult to work with, especially when they are required to work at different frequencies. The first problem equates to a requirement of a stable cart velocity. This is overcome in some schemes by using a very long track and making the reading in the central region where the velocity should be at its steadiest, or through the use of some motional feedback to stabilise the velocity. The second problem would affect us quite badly since we desire to measure wavelengths ranging from 400 nm to 900 nm. The frequencies of the fringes consequently vary by a factor of 2 approximately.

Phase-matching of fringes is arguably the easier method to implement. The cart velocity need only be stable over a short time-scale while a phase match is sought. The principle of the method is quite simple. In the normal fringe-counting method a certain integer number of fringes (say, 10^6) is counted in one channel and the number of fringes in the other channel is read out. It is very unlikely that an integer number of fringes has been observed in the second channel. However, we have no way of knowing what fraction of a fringe is left over since a simple zero-crossing technique has been used to count the fringes. If we knew the fraction left over to n significant figures, we would improve the resolution by a factor of 10^n. Instead of trying to measure the fraction of a fringe, an alternative method is to search for phase coincidences in the two channels. Such phase coincidences are illustrated in Fig. 5.1 and will occur (within an arbitrary resolution) at regular time intervals dependent on the two frequencies. If the counting is started and stopped at a phase coincidence, an integer number of fringes is counted in both channels. The improvement in resolution is given by the precision with which we make a phase-match. For example, phase-matching to 1/10 fringe improves the precision by a factor of 10.

The method of phase-matching of fringes is clearly described by Kahane *et al.* [92]. A circuit is presented based on the use of a high-speed AND gate to detect a phase-coincidence. The zero-crossing detector outputs are sent to the AND gate and if the phases match then they will both be incident on the AND gate at the same time. Of course, this simplified scheme would require

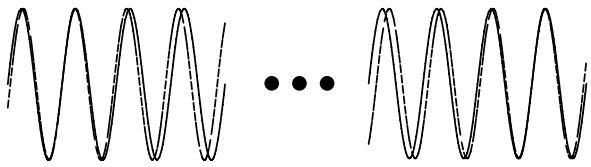

Figure 5.1: Counting of fringes begins when the phase of two signals matches to a given precision (say, 1/100 fringe). When a given number of fringes has been counted on one channel, counting continues until a second phase-match is observed. Thus an integer number of fringes is counted on both channels

the phases to be matched much more closely than to 1/100 of a fringe. To allow some control over the tolerance over which a pair of fringes will be considered to be phase-coincident, the pulses are sent through a cascade of NOR gates. This broadens them slightly in the time-domain. Thus, there is a greater chance that they will overlap in time at the inputs to the AND gate. The number of NOR gates required depends on the choice of IC and the rate at which the fringes pass the detectors.

Work was begun on building a circuit based on that in Ref. [92]. This was abandoned after the offer of a complete circuit by Rainer Blatt's group in the University of Innsbruck. The circuit they had designed and built was based on the same principles as those in Ref. [92]. Thus it would fit very well with the optical and mechanical set-up we had already put in place. The circuit is described in the Diplom thesis of Blasbichler [93]. The main improvement over the original design is the use of a microprocessor to perform the counting and mathematical analysis.

5.4 Errors

Although a desired precision may be achieved, the accuracy of the instrument may not be sufficient. The total error in a wavelength measurement is given by,

$$\left(\frac{\Delta\lambda}{\lambda}\right)^2 = \left(\frac{\Delta\lambda_{\text{align}}}{\lambda}\right)^2 + \left(\frac{\Delta\lambda_{\text{HeNe}}}{\lambda}\right)^2 + \left(\frac{\Delta\lambda_{\text{count}}}{\lambda}\right)^2 , \tag{5.3}$$

along with an additional systematic error,

$$\frac{\Delta\lambda_{\text{WC}}}{\lambda} . \tag{5.4}$$

λ is the true wavelength being measured, and the various deviations are due to:

$\Delta\lambda_{\mathrm{WC}}$ wavefront curvature

$\Delta\lambda_{\mathrm{align}}$ beam misalignment

$\Delta\lambda_{\mathrm{HeNe}}$ reference laser uncertainty

$\Delta\lambda_{\mathrm{count}}$ count resolution

The first of these is described by,

$$\left(\frac{\Delta\lambda_{\mathrm{WC}}}{\lambda}\right)^2 = \frac{\lambda^2}{4\pi^2 w_0^2} = \frac{\Delta\theta^2}{4}\ , \tag{5.5}$$

where w_0 is the waist of the laser beam and $\Delta\theta$ is the divergence of the beam. Essentially there is a variation of phase across the diameter of the beam due to the spherical nature of the wavefronts. It can be minimised by increasing w_0, by expanding the beam size with a pair of lenses. This gives an error comparable to the counting error of one part in 10^8 when,

$$\begin{aligned} \frac{\lambda}{2\pi w_0} &= 10^{-8} \\ \Rightarrow w_0 &= \frac{10^8\lambda}{2\pi}\ . \end{aligned} \tag{5.6}$$

This corresponds to a waist of $113\,\mu$m evaluated for a wavelength of 633 nm. It is not difficult to achieve a beam waist bigger than this in the rare case that a laser might have a larger divergence, through the use of a beam expander.

The alignment error is due to the optical path lengths being different for the two lasers:

$$\frac{\Delta\lambda_{\mathrm{align}}}{\lambda} = \frac{\Delta L}{L} \approx \frac{1}{2}\left(\frac{\Delta x}{L}\right)^2\ . \tag{5.7}$$

Assuming that the unknown and reference lasers are aligned to within 0.5 mm over a distance of 0.5 m then this error will be less than one part in 10^8.

The uncertainty in the HeNe laser wavelength is specified as 1.3×10^{-15} m. Thus its contribution will be of order $1.3\times10^{-15}/633\times10^{-9} = 2.0\times10^{-9}$.

5.5 Optical and Mechanical Design

The decision was made to build the optical set-up on an optical breadboard. This allows some freedom of movement later if it becomes necessary to measure wavelengths on another optical table. Since the intention was to mount the breadboard on the optical table itself, a solid aluminium breadboard (Thorlabs MB3045/M) was chosen over one with a honey-comb interior. If the wavemeter is being run on a side bench in a laboratory, without vibration suppresion, a breadboard with a greater level of damping would be appropriate.

The choice of mirrors was difficult as they had to be highly reflective across a very wide range of wavelengths (397–866 nm). We chose ion-plated silver (Comar[1] 25MX06) as the cut-off in reflectivity of the reflective coating at the lower end of the spectrum was sufficiently below 400 nm to allow good reflection of the 397-nm beams. Ion-plated silver also has the benefit of being quite a tough coating, resistant to scratches. The first mirror after the HeNe laser need only be reflective at 633 nm. A cheaper mirror (Comar 25MX01) was used here.

The choice of beam splitter was also complicated by the wide range of wavelengths. Although it is not necessary to have a 50/50 split between the reflected and transmitted beam, deviation from this will reduce the contrast of the fringes. It is also desirable to have the same reflection–transmission ratio across the whole range of wavelengths to be measured. Our choice was to use a TiO_2 coating (Comar BT) on a glass wedge (Comar 25GQ00).

Cube corners work by total internal reflection and thus work consistently well over a very wide range of wavelengths. Those chosen have a 25 mm diameter front face (Comar 180KQ25).

A neutral density filter is placed immediately after the HeNe laser. This is to avoid saturation of the photodiodes. A transmission of 1% is sufficiently high for read-out.

Finally, since the precision is limited by diffraction effects, it is beneficial to use a beam expander on one or both laser beams. It is impractical to use one on the diode laser beams, but since these already have a large spot size this is not problematic. A beam expander with a magnification factor of four (Comar 01TE04) is thus placed after the HeNe laser.

The optics were mounted in standard mirror mounts (Comar 300BR00) with the exception of the beam-splitter, the bottom half of which was glued to a stan-

[1]Comar Instruments, Cambdrige CB1 7UH, UK, Phone: +44 1223 245470

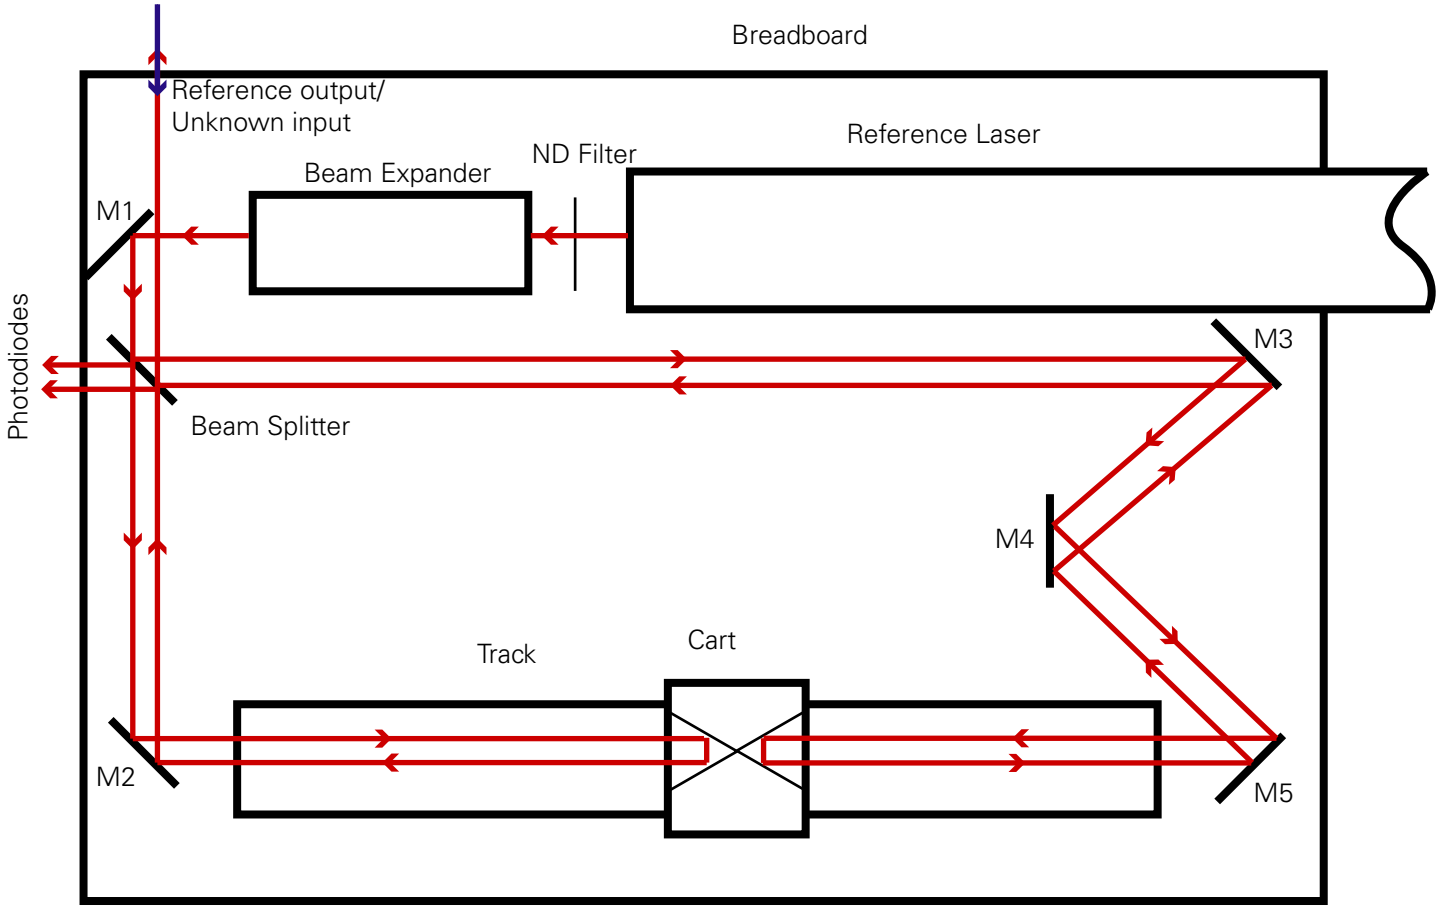

Figure 5.2: The optical layout of the wavemeter

dard mirror mount. This method of mounting the beam-splitter was required in order to maximise the available area available to entry and exit beams. Most of the mirror mounts were attached to 45-mm posts. Posts of diameter 14 mm were chosen since it was felt that the greater the diameter the better the stability against acoustic noise. Several posts had to be custom-made in our workshop in order to achieve greater flexibility in positioning and in order to achieve the required beam height. Specifically the beam splitter was mounted on a 45-mm pedastal post; this gives a lower beam height than a normal 45-mm post since the thickness of the bottom of a post-holder is absent.

5.5.1 Stabilised reference laser

The stabilised laser is a commercial system from Melles Griot (05-STP-903), 1.5 mW output power. It is designed such that only two longitudinal modes can exist within the gain curve, with a mode separation of 640 MHz. Adjacent longitudinal modes in a gas laser are linearly polarised perpendicular to one another. It is this property which is used in stabilising the laser. The laser output is incident on a polarising beam splitter to split the two longitudinal modes into different optical channels and the intensity of each is measured. If the intensity of the two modes is equal, then the modes are symmetrically placed around the centre of the gain curve (see Figure 5.3). An error signal is derived from the difference of the two polarisation intensities. The servo mechanism of the feedback loop is a heater around the tube of the laser. The heating of the tube lengthens it and thus causes the longitudinal mode frequencies to change (the change in the free spectral range can be ignored in this system). A description of a similar method (with a different form of feedback) can be found in Ref. [94].

It is normally desired to have single-mode output from the laser[2]. The Melles Griot system allows just one of the longitudinal modes to exit the laser. The quoted extinction ratio is better than 1000:1 for the unwanted mode. The nominal frequency of the laser is given as 473.61254 THz (corresponding to a vacuum wavelength of 632.99963 nm). The Burleigh wavemeter measure the (vacuum) wavelength as 632.9915 nm.

[2]In fact the Blasbichler design [93] makes use of the two orthogonally polarised modes. The two modes are split at the end of the interferometer by a polarising beam splitter and are incident on two separate photodiodes. Fringes are counted from both polarisations, along with those from the unknown laser. Simple mathematical manipulation is performed inside the microprocessor to obtain the unknown wavelength based on the average number of counts from the different modes.

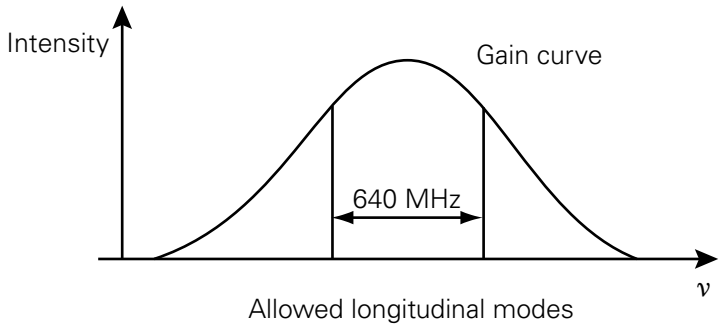

Figure 5.3: The HeNe gain curve with the allowed longitudinal modes (after Melles Griot manual). The intensities of the two modes are kept the same with a heating servo system, thus keeping them equally spaced around the centre of the gain curve and stable in frequency.

One of the most important figures of merit for the HeNe laser is, of course, the frequency stability. This is quoted to be ± 0.3 MHz over the course of a minute and ± 0.8 MHz over the course of an hour. This is certainly better than the desired resolution of the wavemeter.

5.5.2 Cart

The cart is based on the design in [88]. It was constructed from brass according to our own modified design (see Figure 5.4). Its weight is approximately 500 g. A compromise must be sought between stable motion of a heavier cart and the greater air-flow required to float a heavier cart. Higher air flow-rates could result in more turbulence in the beam path, much greater acoustic noise levels during operation, and greater expense in provision of the compressed air supply.

A pipe was welded to the top of the cart to allow a rubber or silicone tube to be attached. The diagonal internal pipes were drilled from the bottom of the side of the cart and plugged with a suitable length of brass.

Alignment of the cube corners is not critical although mechanical stability is required. Two small PTFE dowels are used at the bottom of the cube-corner receptacle. The cube corner is held in place by a nylon screw at the top of the receptacle.

The track on which the cart runs can be made from one of many different materials. Some examples in use in other designs include granite and aluminium. The chief concerns are that the track should be level along its length

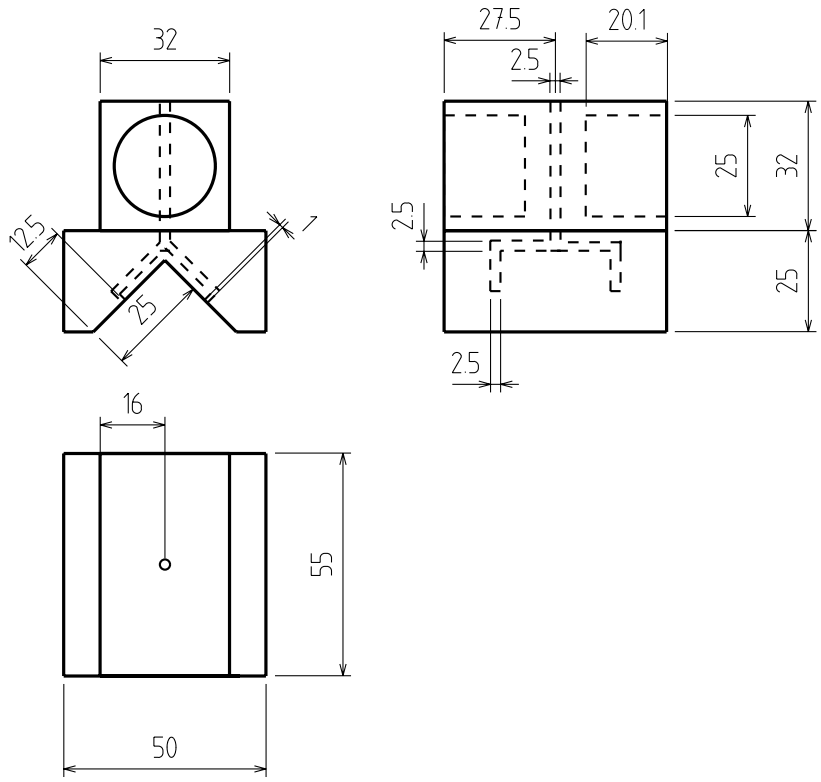

Figure 5.4: Technical drawings for the cart. Not shown are the teflon dowels and M2 nylon screw at either end for holding the cube corners in place. Additionally, a metal pipe was welded to the top to allow the attachment of the air supply.

and that it should be smooth. Since glass is very easily polished to a very flat surface and does not sag greatly, it was decided to make the bar from glass. A bar, 25 mm side and 335 mm in length was cut and polished. It is rests on a pair of heavy v-blocks. It was initially considered that it might be necessary to clamp it or stick it to the v-blocks; this has not proved to be the case.

The wavemeter is designed to read in only one direction of travel of the cart and to display the wavelength while the cart returns to the other end. At one end the cart is pushed by a pair of solenoids (BLP Components 45B Series, miniature PCB-mount, 24Vdc) connected in parallel. A voltage of around 22 V is required to push the cart along the track. A pair of springs is used at the other end to return the cart. There is very little friction observed but the torque exerted by the air hose on the cart can be significant in slowing or accelerating

the cart. This torque is easily counteracted by twisting the connection of the hose to the valve further along the hose.

Measurements were made of the flow rate and pressure required to float the cart. Different pressures were required with different sources of compressed air. In the first instance compressed air piped into the lab was used and a supply pressure of 0.6 bar was needed. During the time in our temporary labspace there was concern that we would not have such a supply and we evaluated the possibility of using an industrial compressor (ShushAir 2-9, flow rate $30.8\,\mathrm{dm^3\,min^{-1}}$). The compressor's capacity was slightly too low to supply the cart. Another industrial compressor (Schneider Profilmaster Silent Typ 200-6 ölfrei) floated the cart at a pressure of 0.2 bar but was too loud for use in a lab. Finally, after problems emerged with the compressed air supply in our refurbished labs, we used cylinders of gas. Both compressed air and oxygen-free nitrogen (OFN) were used; the OFN is a cleaner gas than compressed air and thus preferable. The pressure required was 0.5 bar and a single cylinder was found to last more than a week of moderate use (around two hours flow a day).

5.6 Electronics

We were fortunate enough to be able to use a read-out box designed by the Ion Trapping group in the University of Innsbruck and built by a third party. The box is described in [93].

Some electronic design was necessary for auxiliary components. Most important of these is the photodiode circuit. This is based on the AEPX65 photodiode and OPA655P op-amp in a current-to-voltage configuration. The photodiode's response across the spectrum from 400 nm to 900 nm is good. The biggest challenge is to achieve a high enough bandwidth with sufficient gain. The bandwidth required can be calculated from the number of fringes passing the photodiode and the velocity of the cart. Of order 10^6 fringes pass in 10 s yielding a rate of 100 kHz; the bandwidth must therefore be close to 1 MHz in order to sample this. Although the op-amp chosen is designed to have a high bandwidth, in order to achieve this, care must be taken in the construction of the circuit. Specifically, it was found that moving from regular stripboard to a ground-plane board improved the bandwidth by an order of magnitude or more. It was also important to filter all power inputs via capacitors to ground.

The output of the photodiodes in normal operation is of about 1 V peak-to-peak for typical laser powers (of order 3 mW). The minimum power required to read the wavelength of a blue laser was 0.23 mW while that for the red was 0.5 mW. The required power could be lowered at a later date if desired by including an extra gain stage after the first OPA655P. In fact, an identical op-amp in an inverting-gain configuration would work well (and has been used in a similar set-up elsewhere). The amount of additional gain that could be achieved while retaining a sufficiently high bandwidth would have to experimentally determined.

5.7 Alignment

Due to the complexity of the optical set-up it can be hard to align the system such that fringes are observed. A quick way of achieving the correct alignment is outlined here (mirror labels are shown in Fig. 5.2):

1. Roughly line up the HeNe beam such that it hits all mirrors at a constant beam height (87 mm in our case)

2. With the cart removed, the beams in the two arms of the interferometer should overlap. Mirrors M2 and M5 may be adjusted as necessary to overlap the beams

3. Check that the combined beams overlap on the photodiode. Note that they will fall on the right-hand photodiode at this point but will fall on the left-hand photodiode when the cart is in place

4. With the cart in place and running up and down the track, the track itself may be adjusted until the fringes observed are roughly optimised

5. Block the longer (right-hand) arm of the interferometer. With the cart in place but the air supply turned off, the signal on the photodiode should be optimised iteratively. Firstly, with the cart nearest M5, mirror M2 should be used to optimise the signal. Then with the cart near M2, mirror M1 should be adjusted. The iterations of the two adjustments should be continued until a constant signal level is seen along the length of the track

6. The air supply should be turned on again and the right-hand arm of the interferometer unblocked.

Theoretical Prediction	Wavemeter reading	Burleigh WA-20 reading (vacuum)	
396.84717	396.84723(8)	397.959	*Paul trap*
866.21448	866.21569(8)	866.455	
396.85198	396.85216(8)		*Penning trap*
396.84237	396.84237(8)		
866.24195	866.24315(8)		
866.18701	866.18755(14)		

Table 5.1: Paul-trap and Penning-trap (0.98 T) wavelengths. Comparison of the new wavemeter with theoretical predictions and and older commercial device. Errors in the wavemeter reading are statistically calculated from a Gaussian fit to the frequency of discrete wavelengths (50 readings in total). Readings were not available from the commercial wavemeter for the transitions at 0.98 T. The readings of the infrared lasers for the Penning trap transitions are of the carrier wavelengths. Sidebands on each of these carrier wavelengths address the transitions as described in Section 4.3.3 of Chapter 4.

7. A second iterative procedure should now be followed. Firstly trundle the cart up and down along a distance of about 5 cm at the end of the track nearest M2. Meanwhile optimise the size of the fringes using M5. Next trundle the cart up and down a short distance at the other end of the track and optimise the fringe contrast using M3. The two steps should be iterated until the fringes have the same contrast along the whole track

8. Coupling of the unknown laser into the interferometer can be achieved with two external mirrors only. The internal mirrors should not need to be adjusted.

5.8 Calibration

The wavemeter was calibrated by measuring the wavelengths of the lasers tuned to the correct wavelengths for laser-cooling, i.e. by comparison with the atomic spectral lines. The lasers were locked to their reference cavities during the measurements. The theoretical predictions and experimental measurements are given in Table 5.1. An example of plots of wavelength measurements is given in Fig. 5.5.

The theoretical predictions have two corrections built in. The first is that the wavemeter electronics have the HeNe laser wavelength hard-coded as

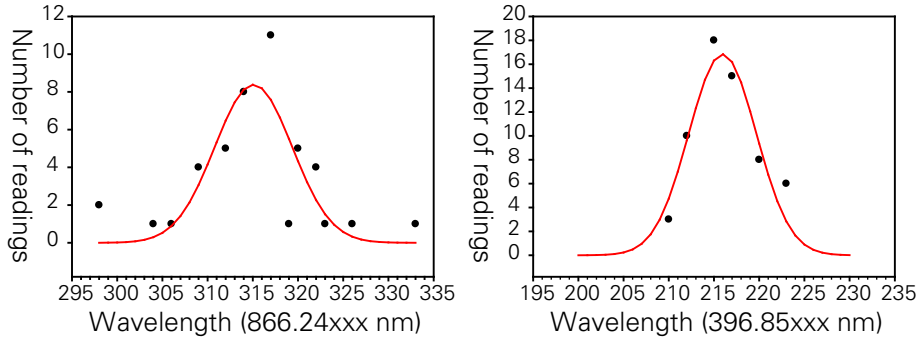

Figure 5.5: Graphs of the occurrences of discrete wavelengths with Gaussian fits. The errors in Table 5.1 were calculated from the width of the Gaussian fits. *a* is the measurement of one of the 866-nm lasers and *b* is the measurement of one of the 397-nm lasers when tuned to the relevant transitions for laser-cooling at 0.98 T. The lasers were locked to the reference cavities during the readings and the readings were taken over a short time-scale (several minutes). Laser-frequency drift should not, therefore, contribute to the width of the Gaussian peaks.

632.81654 (close to the wavelength in air) while the actual (vacuum) wavelength is 632.99096 nm, calculated from the quoted frequency. This second value was confirmed with the commercial wavemeter within its precision (632.991 nm). The second correction is for variation in refractive index at different wavelengths (Fig. 5.6). This variation is also temperature dependent [95] and an ambient temperature of 20 °C is assumed. This is reasonable as the laboratory temperature is stable to within 2 °C.

The results in Table 5.1 show that the precision of the wavemeter is high (better than one part in 10^7). The accuracy and repeatability of the device are also high. The wavelength readings vary little in the fourth decimal place from day to day. This also indicates a good level of temperature stability within the laboratory as a change of even 1 °C would cause a change in the fourth decimal place of the wavelength readings.

A close examination of Table 5.1 shows that the wavemeter readings in the ultraviolet (UV) more closely match theoretical predictions than in the infrared (IR). The reason for this difference in accuracy being higher in the UV than in the IR is unknown but may be due to the fact that for the shorter wavelength a smaller length of the track which must be traversed for a measurement of the desired number of fringes. The repeatability of readings right across the

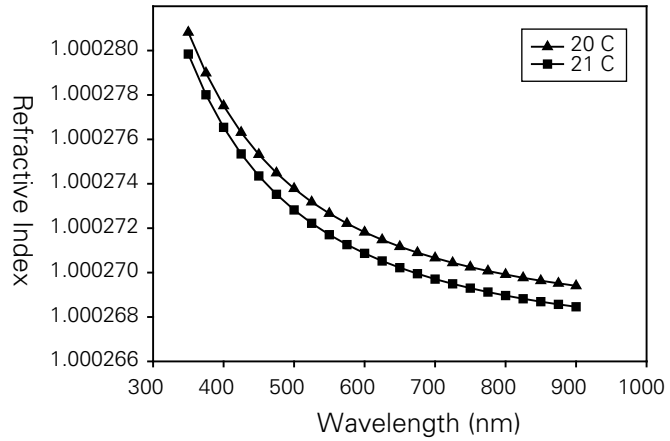

Figure 5.6: The variation of the refractive index of air with wavelength calculated using Ref. [96]. A change in the air temperature of 1 °C is enough to cause a change in the fourth decimal place of the wavemeter reading (i.e. 10^{-4} nm).

spectrum means that an offset for wavelengths in the infrared is not problematic, in any case.

One feature which has been omitted from the graphs in Fig. 5.5 is the occurrence of readings rather lower that the others observed. The typical difference between the "true" readings and these anomalous readings is approximately five parts in the third decimal place. After calculation, it is apparent that these other readings are "off-by-one" errors where the wavemeter has read one too few fringes of the unknown laser. The frequency of occurrence of these errors depends very greatly on the motion of the cart. A careful balancing of the torque exerted by the air tube is required to minimise the number of such errors. In typical use the occurrence has been 6% or less. Since such readings are well understood and easily seen by inspection to be very different from the "true" readings this does not cause a problem with the wavemeter functionality.

5.9 Discussion

A new wavemeter has been designed and built. With the help of electronics supplied by another group, wavelength readings have been consistently made with a precision of greater than one part in 10^7. The design was based around a cart mounted with retroreflectors floating on a cushion of air providing a very smooth motion. The primary disadvantage with this design is that the waveme-

ter cannot be operated in vacuum and as such is subject to variations in the refractive index of air both through air currents and temperature changes. During more than one year's use it has been found that neither effect is large enough in our laboratory environment to cause difficulty.

The immediate advantage of having such an accurate wavelength-measuring device has been the ability to load calcium ions in a Penning trap directly. Previously, it had been necessary to tune the 397-nm lasers to the cooling transition using another source of calcium ions (a hollow cathode lamp) and to increase the magnetic field in increments while following the change in the transition wavelengths with the laser tunings. This laborious process has now been eliminated.

Looking to the future, the new wavemeter will become invaluable when the Ti:Sapphire laser is ready for use. The narrow transition at 729 nm will be difficult to locate but this process should be made much easier with such a precise knowledge of the laser's wavelength.

Chapter 6

Axialisation of ions in a Penning trap

Trapping single Ca^+ ions in a Penning trap is an essential step towards demonstrating that the required level of control over the trapped ions can be achieved in order to implement quantum computing schemes. In order to achieve this goal, the trap must be very stable. The instability of the magnetron motion in a Penning trap can make the trapping of small numbers of ions very difficult, however.

One way to improve the stability of the Penning trap is to use an azimuthal RF field in a manner known as azimuthal quadrupolar excitation [14]. Given the correct geometry and frequency, this field can couple the two radial motions. Thus the normally errant magnetron motion can be efficiently cooled through coupling to the laser-cooled modified cyclotron motion. The consequent reduction of the magnetron radius results in the ions being tightly confined around the axis of the trap, hence the other name for the scheme: axialisation. This tight confinement is also necessary for QIP schemes that need individual ion addressing. The theoretical basis for axialisation is described in this chapter while its implementation with both Mg^+ and Ca^+ ions is discussed in Chapters 7 and 8 respectively.

6.1 Plasma confinement

While our work focuses on the trapping of very small number of ions in a Penning trap, there are others seeking to trap large clouds instead [97, 98], with sizes

ranging over many orders of magnitude from about 10^3 to 10^9 ions or more. In such large clouds the ions no longer move in a similar manner to a single ion in a Penning trap. Instead, the space charge effects become important and the ion cloud must be thought of as a non-neutral plasma.

The radial motion of a plasma in a Penning trap is, of course, still unstable. In order to increase trapping times it is desired to compress the plasma in the radial plane. Two methods have been developed to do this. The first [97] is to apply a torque to the plasma with a laser beam perpendicular to the trap axis, near-resonant with the cooling transition, and off-set from the trap axis in a manner similar to that described for conventional laser cooling in a Penning trap (Chapter 2). This method is suitable for medium-sized plasmas (10^3–10^6 ions). For larger clouds (up to 10^9 ions), a technique known as the "rotating wall" [98] may be applied. Here, an RF voltage is applied to a segmented ring electrode. The frequency is chosen to be close to the plasma rotation frequency (analagous to the magnetron frequency). Turning the frequency up forces the plasma to rotate faster and a consequent compression of the cloud is achieved.

Neither of the techniques described can be applied to small clouds of ions in a Penning trap since the behaviour of such small clouds is not plasma-like. It is for this reason that we investigate the technique of axialisation for such small clouds.

6.2 FT-ICR spectroscopy

One of the major uses of Penning traps in modern science is the measurement of molecular masses to a very high precision. Some of the techniques developed by the mass spectroscopy community can be applied to laser-cooled ions in Penning traps. For this reason I present some of the salient features of mass spectroscopy.

The measurement of molecular masses can be performed using one of several methods including the ejection of ions followed by time-of-flight measurements as they pass through a magnetic field, the ejection of ions based on their motional frequency (and hence their mass), and the detection of their motional frequencies inside the trap. It is this last method that is of most interest here.

If a signal close to a motional frequency is applied to the ions, that motion may be excited. This excitation can be detected by measuring the induced currents in the trap electrodes. This is similar to the concept of resistive cooling

of ions where the motional frequency is coupled to a tuned circuit through the induced currents in trap electrodes.

The excitation frequency may be slowly scanned across a frequency range of interest and the resultant excitation detected with a narrow-bandwidth detector. This will give a measure of the excitation amplitude as a function of the excitation frequency. Comisarow and Marshall developed a major improvement to this motional detection scheme [99] by applying a strong excitation close to the resonant frequency (normally the modified cyclotron frequency) and then performing a Fourier analysis of the electrode signal (Fig. 6.1). This allowed the motional frequencies to be determined with great accuracy in a fraction of the time needed for the "slow-scanning" method. It is Comsiarow and Marshall's method that is still used in many experiments around the world. It is known as Fourier Transform Ion Cyclotron Resonance (FT-ICR) spectroscopy.

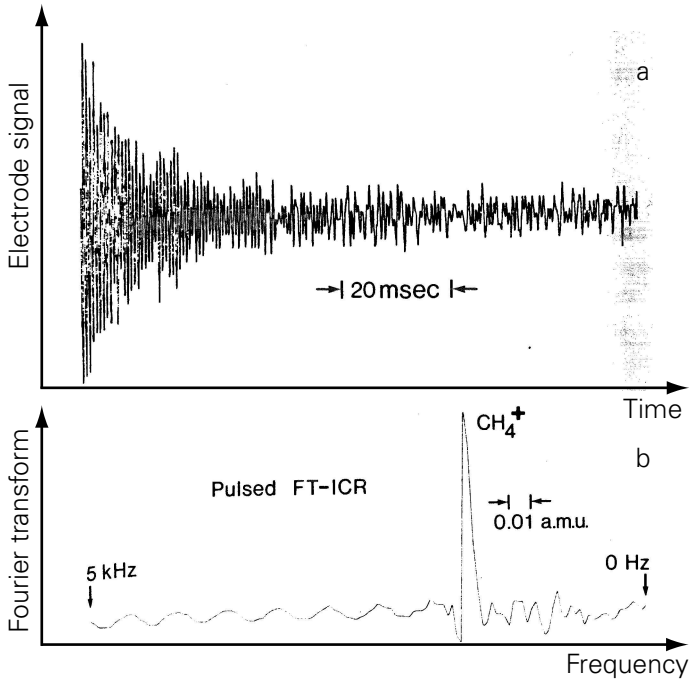

Figure 6.1: An example of FT-ICR spectroscopy. *a* the "ring-down" signal detected in the electrodes and *b* the Fourier transform of the first 256 channels of that signal showing the peak corresponding to CH_4^+. The frequency in *b* is offset from the resonant frequency due to a mixing process. Axes have been added to both plots. From Ref. [99]

6.3 Excitation frequencies and geometries

The three modes of motion in an ideal Penning trap (axial, modified cyclotron, and magnetron) can be treated as distinct, uncoupled motions. They may be individually excited with fields of an appropriate geometry, oscillating at the motional frequency. These geometries are described in detail by Schweikhard and Marshall [13].

6.3.1 Dipolar excitation

One of the geometries of interest to us is the "standard" excitation mode or radial dipolar excitation (Fig. 6.2). This allows the excitation of one or other radial mode. The dipolar geometry is easily achieved by dividing the ring electrode in two and applying opposite phases of the excitation signal to each half. If the ring electrode has been divided into quadrants, an opposing pair of quadrants may be connected in the same way. This is a better approximation to the dipole geometry and in FT-ICR experiments the other pair of ring quadrants may be used to detect the ion signal.

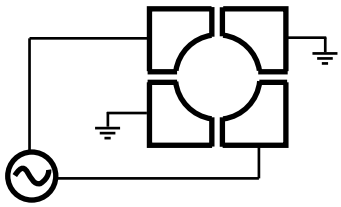

Figure 6.2: Radial dipolar excitation geometry. The signal is connected to an opposing pair of ring segments.

6.3.2 Parametric excitation

In order to address the axial mode, parametric excitation may be used. Parametric means that the excitation has the same geometry as the trapping potential. The excitation signal is applied between the pair of end-caps and the ring (Fig. 6.3).

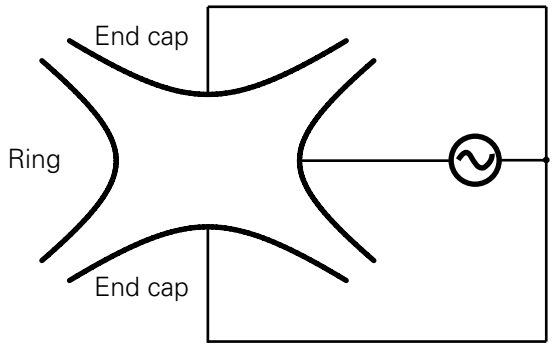

Figure 6.3: Parametric excitation geometry. This is the same as the trapping geometry. A signal at $2\omega_z$ in this geometry will excite the axial motion.

6.3.3 Azimuthal quadrupolar excitation

Finally, and of most interest to us, is an xy-excitation. Azimuthal quadrupolar excitation requires the ring electrode to be divided into quadrants. A signal is applied to one pair of opposing quadrants with the same signal in anti-phase applied to the other pair of quadrants (Fig. 6.4).

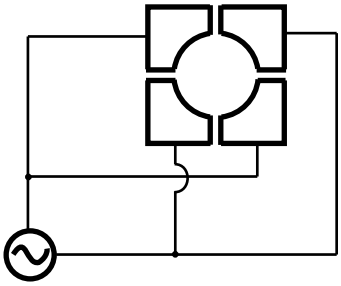

Figure 6.4: Geometry for azimuthal quadrupolar excitation. The ring is divided into quadrants. One pair of opposing quadrants is connected to one phase of the excitation signal while the other is connected to the other phase.

Bollen *et al.* [10] show that signals at three frequencies can be detected (and hence efficiently excited) for a ring divided into an even number of segments and where the number of segments is four or more. They also calculate expressions for the electric field from such a ring electrode, divided into quadrants. This geometry is very similar to that shown in Fig. 6.5 and used in the experiments described in Chapters 7 and 8. Abbreviated to first order, this expression

is,

$$\begin{aligned} E_x &= -2.0V_0\frac{y}{r_0^2}\,, \\ E_y &= -2.0V_0\frac{x}{r_0^2}\,, \end{aligned} \tag{6.1}$$

where $\pm V_0$ is the voltage applied to the segments and r_0 is the inner radius of the ring electrode. Thus, the field closely matches a true quadrupole.

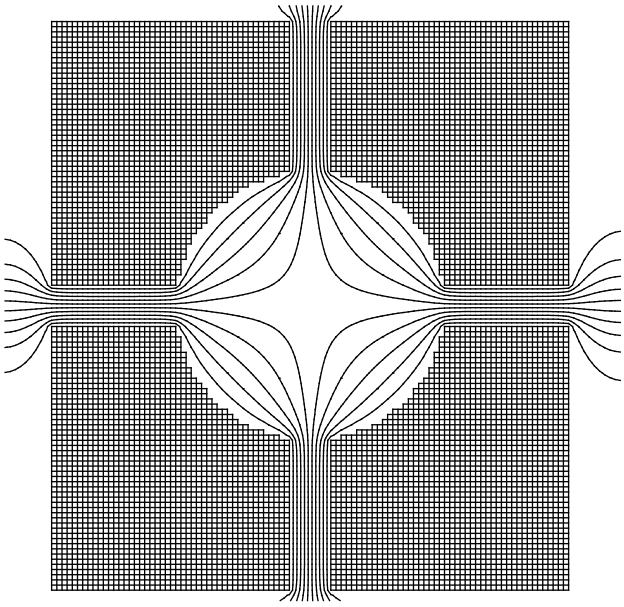

Figure 6.5: Equipotential lines of the quadrupole field generated by the split-ring electrode used in the experiments described in Chapters 7 and 8. Although the geometry is very different from the ideal hyperbolic ring segments, the field in the centre of trap very closely matches a pure quadrupole field. The inner radius of the ring is 5 mm. Figure produced by R. Castrejon-Pita using SimIon.

The motion of an ion in the radial plane can be represented as follows,

$$\begin{aligned} x &= \rho_c'\cos\omega_c' t + \rho_m\cos\omega_m t\,, \\ y &= \rho_c'\sin\omega_c' t + \rho_m\sin\omega_m t\,, \\ \dot{x} &= -\rho_c'\omega_c'\sin\omega_c' t - \rho_m\omega_m\sin\omega_m t\,, \\ \dot{y} &= \rho_c'\omega_c'\cos\omega_c' t + \rho_m\omega_m\cos\omega_m t\,, \end{aligned} \tag{6.2}$$

where ρ'_c and ρ_m are the radii of the modified cyclotron and magnetron motions respectively.

The Lorentz force due to the field of Equation 6.1 is,

$$\begin{aligned} q\mathbf{v}\cdot\mathbf{E} &= 2.0\frac{q}{r_0^2}\left[\rho_c'^2\omega'_c\cos 2\omega'_c t + \rho_m^2\omega_m\cos 2\omega_m t\right. \\ &\quad \left. +\rho'_c\rho_m(\omega'_c+\omega_m)\cos(\omega'_c+\omega_m)t\right] . \end{aligned} \tag{6.3}$$

Thus, the magnetron and modified cyclotron motions can be independently excited with an azimuthal quadrupole field at twice their respective motional frequencies ($2\omega'_c$ and $2\omega_m$ terms). The ion also has a resonance at the sum frequency, $\omega'_c+\omega_m \equiv \omega_c$.

While the resonances at twice the motional frequencies will clearly detect or excite the respective radial motions, it is not immediately clear what the effect of the resonance at the sum frequency will be. It is found that an applied field at this frequency couples the two motions as shall be seen in the next section.

6.4 Coupling motional modes

We are not limited to the simple excitation of particular modes of motion; it is also possible to couple any pair of modes. This is done by applying a signal at the sum or difference frequency (dependent on the pair of modes to be coupled) in a particular geometry.

6.4.1 Coupling frequencies

Normally, two harmonic oscillators can be coupled (so that energy may be transferred between the oscillators) by applying a force oscillating with a frequency close to the difference of the oscillators' natural frequencies. For example, the axial and cyclotron motions may be coupled by a field oscillating at their difference frequency in just this way. However, the negative energy of the magnetron motion complicates matters. In order to couple another motion to it, the sum frequency must be used instead of the difference frequency. Although unfamiliar from our everyday experience, an intuitive understanding of this may be gained by considering a quantum mechanical model of this classical system (Section 6.5.3).

6.4.2 Coupling to the axial mode

The generation of a coupling between two motional modes was first proposed by Wineland and Dehmelt [15] who suggested the coupling could be used to damp a normally undamped motion. The example they gave was of a radial motion coupled to the axial motion, where the axial motion was damped through resistive cooling. Dehmelt's group demonstrated such a coupling between the magnetron motion and the axial motion a few years later [100]. They observed a factor-of-14 reduction in the magnetron radius. The coupling was achieved using an inhomogeneous field at the sum frequency of the two motions.

The geometry to achieve the coupling between the axial and magnetron modes is discussed in more detail by Guan *et al.* [101]. It requires an xz-symmetry where z is the trap axis. Thus, both a segmentation of the end-caps and a double segmentation of the ring electrode is required (Fig. 6.6). In practice, approximations to this geometry are often made, including the original demonstration of the coupling by van Dyck *et al.* [100]. Other interesting examples of the coupling of these two modes include Refs. [102, 103].

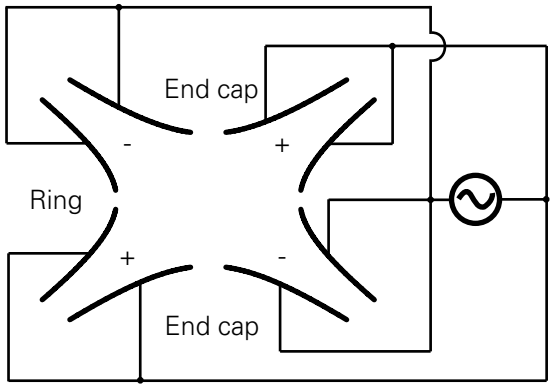

Figure 6.6: The geometry for coupling the axial motion to one of the radial motions. A quadrupole is created with an xz-symmetry axis. In order to do this the end caps must be segmented and the ring must be divided both into quadrants radially and in half axially. An opposing pair of quadrants of the ring is connected as shown while the other pair of quadrants is connected to earth (not shown).

6.5 Coupling the magnetron and modified cyclotron motions

Just as the magnetron motion can be coupled to the axial motion, with an appropriate change of geometry and driving frequency, it can also be coupled to the modified cyclotron motion. This was first proposed by Gräff *et al.* [104]. Such a coupling is desirable since the damping rate of the modified cyclotron motion in the presence of laser-cooling is almost always much greater than the damping rate of the magnetron motion. Thus, coupling the two would lead to a much larger damping of the magnetron motion. A similar argument applies to the case of buffer-gas cooling except that in that case it is always the modified cyclotron motion which is damped and the magnetron motion is heated (when the motions are uncoupled).

The necessary geometry for the coupling is the azimuthal quadrupolar excitational geometry described in Section 6.3.3. As described in that section, an ion will respond to a quadrupole field oscillating at the true cyclotron frequency, ω_c and it is found that this is the correct frequency to couple the two radial modes. Since it is this coupling which leads to axialisation, it will be described in more detail here.

Bollen *et al.* [10] treat the interaction of an oscillating azimuthal quadrupole field with a single ion in detail. They follow the formalism developed by Brown and Gabrielse [11] by introducing velocity vectors, $\mathbf{V}^+$ and $\mathbf{V}^-$ for the modified cyclotron and magnetron motions respectively. They show that a coupling between the two motions is introduced by the oscillating field and that if the field oscillates at a frequency close to the true cyclotron frequency ($\omega_c = \omega_c' + \omega_m$) a slow variation of the two velocity amplitudes is introduced. The coupling constant, k, is found to be,

$$k = \frac{qV_0}{mr_0^2}\,, \tag{6.4}$$

where V_0 is the peak-to-peak amplitude of the oscillating voltage applied to the ring segments and r_0 is the inner radius of the ring. This assumes a ring geometry similar to the one used in our experiments.

The change in either motion is opposite to that of the other, i.e. as the magnetron motion decreases in amplitude, the modified cyclotron amplitude increases. This interconversion is analogous to Rabi flopping in a driven atomic system. The time taken for the conversion of one mode to the other (i.e. the

length of a π-pulse) is given by,

$$T = \frac{\pi}{k} \,, \tag{6.5}$$

assuming the driving field is resonant with the sum frequency of the motions.

Clearly, if one of the two motions is damped, i.e. its motional amplitude is reduced, this interconversion will cause a damping of the other motion. Since the modified cyclotron motion is relatively easily damped, either by buffer gas cooling or by laser-cooling, this indicates an efficient way of damping the magnetron motion.

6.5.1 The rotating frame

A clearer picture of the radial motions in the Penning trap may be obtained by moving from the laboratory frame to one rotating at half the cyclotron frequency ($\omega_c/2$) [12]. In this frame it is found that the force due to the magnetic field is exactly cancelled by the Coriolis force. First, I consider the treatment of the radial motion of the ions in the laboratory frame. I then compare the treatment in the rotating frame. It will be seen that the motion of the ions can be more clearly expressed in this rotating frame and that the coupling field's action is more easily understood.

In the laboratory frame, the standard equations of motion in the Penning trap apply (see Chapter 2):

$$\ddot{x} = \frac{2eU}{mR^2}x + \frac{qB}{m}\dot{y} \,, \tag{6.6}$$

$$\ddot{y} = \frac{2eU}{mR^2}y - \frac{qB}{m}\dot{x} \,, \tag{6.7}$$

$$\ddot{z} = -\frac{4eU}{mR^2} \,, \tag{6.8}$$

where U is the DC voltage applied between the end caps and the ring electrode and $R = r_0^2 + 2z_0^2$ is a geometric factor depending on the ring-electrode radius (r_0) and half the separation of the end caps (z_0).

Equations 6.6 and 6.7 may be rewritten,

$$\ddot{x} - \beta\dot{y} - \alpha x = 0 \,, \tag{6.9}$$

$$\ddot{y} + \beta\dot{x} - \alpha y = 0 \,, \tag{6.10}$$

where $\alpha = 2qU/mR^2$ and $\beta = qB/m$. Noting that Equation 6.8 describes simple harmonic motion with an angular frequency, $\omega_z^2 = 4qU/mR^2$, and that the true cyclotron frequency is $\omega_c = qB/m$, we may write $\alpha = \omega_z^2/2$ and $\beta = \omega_c$. The two radial parameters, x and y, may be combined,

$$u = x + iy \ , \tag{6.11}$$

yielding a single equation in u,

$$\ddot{u} + i\beta\dot{u} - \alpha u = 0 \ . \tag{6.12}$$

Writing,

$$u = u_0 \exp(i\omega t) \ , \tag{6.13}$$

and substituting into Equation 6.12 yields a quadratic equation in ω,

$$\omega^2 - \beta\omega + \alpha = 0 \ , \tag{6.14}$$

whose solutions are,

$$\omega = \frac{\beta}{2} \pm \frac{\sqrt{\beta^2 - 4\alpha}}{2} \ , \tag{6.15}$$

$$= \omega_c/2 \pm \omega_1 \ , \tag{6.16}$$

where $\omega_1^2 = \omega_c^2/4 - \omega_z^2/2$. Thus the two solutions are the modified cyclotron and magnetron frequencies respectively.

Now, in a rotating frame a new coordinate suggests itself,

$$v = u \exp(i\omega_R t) \ , \tag{6.17}$$

where ω_R is the angular rotation frequency of the frame. Differentiating this expression twice and substituting into Equation 6.12 gives,

$$\ddot{v}e^{-i\omega_R t} + i(\beta - 2\omega_R)\dot{v}e^{-i\omega_R t} + (\beta\omega_R - \omega_R^2 - \alpha)ve^{-i\omega_R t} = 0 \ . \tag{6.18}$$

Now, setting the rotation frequency to be $\omega_R = \beta/2 = \omega_c/2$ greatly simplifies this expression,

$$\ddot{v} + \left(\frac{\beta^2}{4} - \alpha\right) v = 0 \ , \tag{6.19}$$

which is simple harmonic motion in the complex plane with $\omega^2 = \beta^2/4 - \alpha$. The motional frequencies are in fact,

$$\omega = \pm\sqrt{\beta^2/4 - \alpha} = \pm\sqrt{\omega_c^2/4 - \omega_z^2/2}\ , \tag{6.20}$$

where a negative frequency indicates rotation in the opposite direction to that of the frame. These frequencies, when converted back to the laboratory frame, are once again the modified cyclotron and magnetron frequencies respectively. However, in the rotating frame they are essentially the same motion but in opposite directions. A hint that they should become degnerate in this frame is seen in Fig. 2.4.

Some examples comparing motion in the lab frame with that in the rotating frame are given in Fig. 6.7. The simplest case of pure magnetron or pure modified cyclotron motion is shown in Fig. 6.7 *a* and *b*. It is important to note that there is no difference in the shape of the orbit for either motion; the only difference between the two is that they move in opposite directions in the rotating frame.

A superposition of the two motions with equal amplitudes will give rise to a linear oscillation in the rotating frame (Fig. 6.7 *c* and *d*). This is analagous to the polarisation of light where linearly polarised light can be decomposed into two circular polarisations.

In the normal case where the magnetron motion has a much larger radius than the modified cyclotron motion (Fig. 6.7 *e* and *f*), the motion in the rotating frame is elliptical and the direction is counter to the frame's rotation (i.e. in the direction of pure magnetron motion). The semimajor and semiminor axes are $(\rho_m + \rho_c')$ and $(\rho_m - \rho_c')$ respectively.

The same shape of orbit as in Fig. 6.7 *e* but in the opposite sense (i.e. in the direction of the modified cyclotron motion), gives a very different result in the lab frame (Fig. 6.7 *g* and *h*).

6.5.2 The azimuthal quadrupole field in the rotating frame

It has been seen that the description of the motion of the ions can be greatly simplified by moving to a rotating frame. This description is even more helpful when a quadrupole field oscillating at ω_c is introduced. The effect of this is to change the normal centre-symmetric parabolic potential in the rotating frame to one that is compressed in one direction (Fig. 6.8).

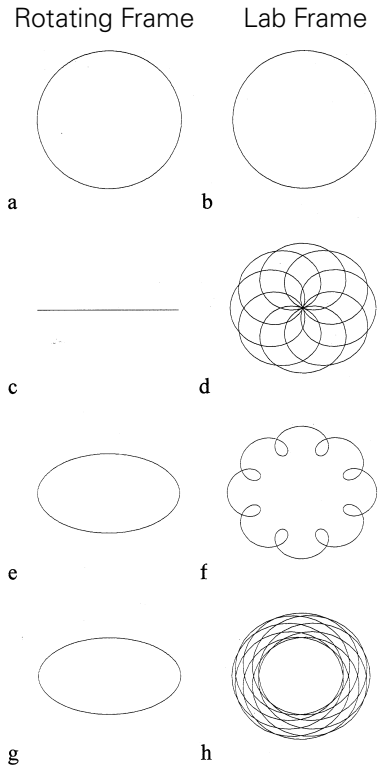

Figure 6.7: Comparison of motion in the rotating frame (left) and the lab frame (right). *a* and *b*: pure magnetron or pure modified cyclotron motion; *c* and *d* magnetron and modified cyclotron motions with equal amplitudes; *e* and *f* large magnetron amplitude, small modified cyclotron amplitude; *g* and *h* small magnetron amplitude, larged modified cyclotron amplitude. Note that the direction of rotation in *a*, *e*, and *g* distinguishes magnetron from modified cyclotron motion in the rotating frame. From Ref. [12].

In the absence of the azimuthal quadrupole, two normal modes of motion exist but they are degenerate. The normal modes can be expressed either as linear oscillations or as circular oscillations in a manner analgous to the description of polarised light. For example, one of the linear normal modes can be decomposed into the two circular normal modes. When the quadrupole field is introduced, the two modes become nondegenerate, with directions along each of the major and minor axes of the potential.

The effect of the change in the potential is to place the motions in a superposition of the two new normal modes. The motions interconvert, as seen in Fig. 6.9. Circular motion in the compressed potential changes to elliptical mo-

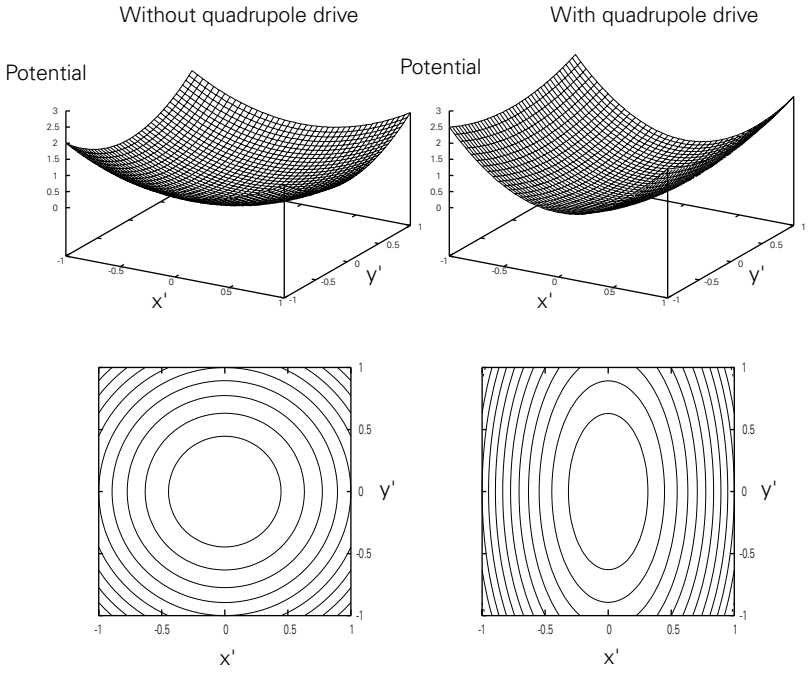

Figure 6.8: The effect of applying the azimuthal quadrupole field, as seen in the rotating frame. The normally centre-symmetric field becomes compressed in one direction. Two normal modes exist in this new potential, along the major and minor axes.

tion at a rate dependent on the eccentricity of the potential. The eccentricity itself is determined by the amplitude of the quadrupole field. The elliptical motion continues to change shape and eventually reaches the limit of a straight line, at 45° to the two normal modes.

The linear motion seen in Fig. 6.9 begins to become circular again, but the motion is now in the opposite direction than before. This means that where the initial motion was, say, entirely of magnetron character it now acquires more of a modified cyclotron character. Ultimately the motion becomes circular again i.e. pure magnetron in character.

The process of interconversion continues and this time, as the motion becomes more ellipsoidal, it does so with a major axis perpendicular to that seen in the previous transition from circular to linear. The behaviour described is shown in the laboratory frame in Fig. 6.10.

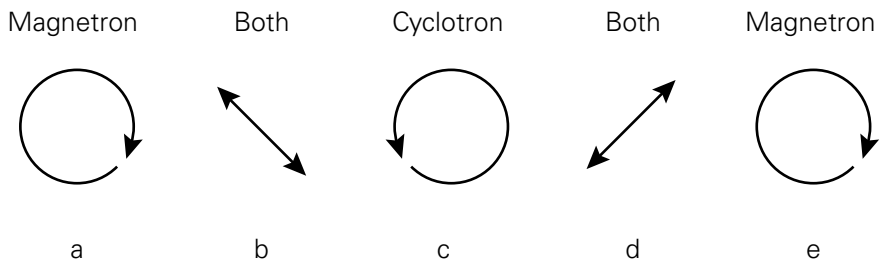

Figure 6.9: Interconversion of the two radial motions in the rotating frame. In the compressed potential, circular motion in the clockwise direction (*a*) becomes gradually more elliptical until it is reaches the limit of being linear (*b*). It then begins to become more circular again but this time circulating in the opposite direction (*c*). Having become purely cyclotron in character its ellipticity increases again but this time with the major axis orthogonal to that in *b*. Having gone through a stage of displaying both motional characteristics equally again (*d*) it becomes purely magnetron in character once again (*e*).

6.5.3 A quantum mechanical picture

In the experiments described in this thesis, the coupling between the two radial motions can be considered in an entirely classical framework. The motional quantum state is high and in the early stages of laser-cooling motional sidebands cannot be resolved on any of the optical transitions. Nevertheless, it is instructive to consider a third explanation of the axialisation process based on a quantum mechanical approach. Indeed, this was the first one to be published, put forward by Dehmelt *et al.* [100].

Dehmelt's description is in terms of motional-sideband cooling. In his experiment the axial motion was coupled to the magnetron motion. This is entirely analagous to coupling the modified cyclotron motion to the magnetron motion.

An energy level diagram is shown in Fig. 6.11. The modified cyclotron motion imposes motional sidebands on the optical transition between the ground state and an excited state (not shown). The magnetron motion, in turn, imposes sidebands on these sidebands. It is important to note that while the sidebands due to the modified cyclotron motion are in a ladder going upwards in energy (i.e. the energy associated with n_c is greater than that associated with $n_c - 1$), the sidebands associated with the magnetron motion go in the opposite direction. This is a simple consequence of the magnetron motion's total energy being negative.

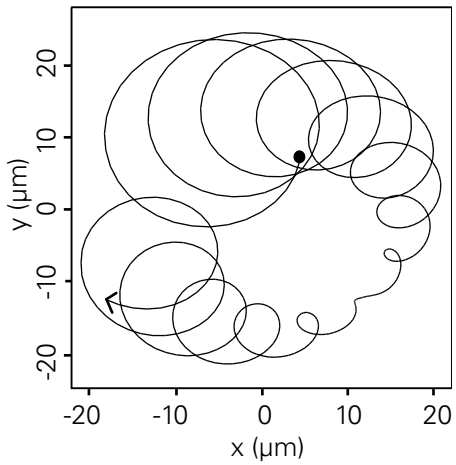

Figure 6.10: Interconversion of the two radial motions in the laboratory frame. The motion changes from pure cyclotron to pure magnetron and back again. From Ref. [105]. The program simulating the system is described in the same reference.

An RF field applied on the blue motional sideband on the modified cyclotron frequency couples (n_c, n_m) to $(n_c + 1, n_m - 1)$. The rate of change between these states is proportional to axialisation-drive amplitude. It is important to note that it is the blue sideband which results in a reduction of the magnetron motional quantum number. This is the clearest way to see why radiation at $\omega_c' + \omega_m$ must be used instead of radiation at the difference frequency.

Damping of the modified cyclotron motion introduces an asymmetry into the coupling process just described. Laser-cooling induces transitions between the ground state and the excited state with a reduction in n_c, i.e. the laser is red-detuned. Even when the sidebands are not resolved this statement holds. When the ion decays from the excited state there may be a change in the phonon number of the motional modes in either direction. On average, however, there will be no net change. Thus, there will be downward trend in the modified cyclotron quantum number.

In summary, then, there is a symmetric coupling between (n_c, n_m) and $(n_c + 1, n_m - 1)$ which will have no net effect on the two quantum numbers when averaged over time. The addition of damping of the cyclotron motion causes a decrease of n_c and a consequent decrease of n_m.

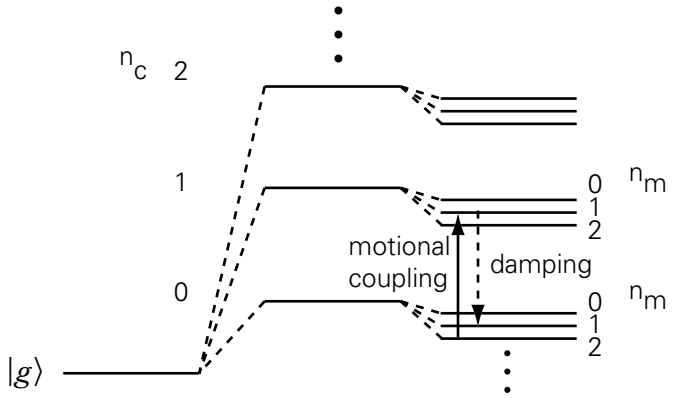

Figure 6.11: Motional-sideband cooling. Here the motional levels are the modified cyclotron (n_c) and the magnetron (n_m) although any pair of motions may be coupled with an appropriate field. Note that the modified cyclotron levels increase in energy while the magnetron levels decrease. The coupling links (n_c, n_m) with ($n_c + 1$, $n_m - 1$). In the absence of damping the motional coupling would be a balanced two-way process. With damping, however, the direction shown predominates and there is a net damping effect of the magnetron motion.

6.6 Measuring the damping rate of the motions

In determining the effectiveness of axialisation it is desirable to measure the damping rate of the ion motions. One method of measuring damping rates and motional frequencies of each of the characteristic motions of the ions is to use an RF–photon correlation technique [67]. The principle behind the method is explained here and the changes introduced by the coupling of the motions described.

6.6.1 Uncoupled motion

Once again, it is appropriate to treat the ion's motion as a damped harmonic oscillator with a natural frequency equal to its motional frequency. The damping, as before, arises from the laser-cooling.

As described in Section 6.3 an additional driving force or excitation can be applied to the ion, close to its motional frequency. In this case the system can be treated as a classical forced, damped harmonic oscillator (see for example Ref. [106], p. 216). As with any similar system the phase between the drive and the motion changes as the drive passes through the natural resonance of the

system.

The system is modelled using the inhomogeneous equation:

$$m\ddot{x} + \gamma\dot{x} + \kappa x = F_0 \cos\omega_f t \ , \tag{6.21}$$

where m is the mass of the particle, x is the projection of the radial motion along one of the radial axes, γ is the damping rate of the particle's motion, κ is the restoring force (the trapping potential), F_0 is the amplitude of the extra applied force, ω_f the angular frequency of that force, and t is time.

A solution of the form

$$x = a\cos\omega_f t + b\sin\omega_f t \ , \tag{6.22}$$

is sought. Differentiating this proposed solution, inserting into Equation 6.21 and equating cos and sin terms yields expressions for a and b:

$$a = \frac{F_0(\kappa - m\,\omega_f^2)}{(\kappa - m\omega_f^2)^2 + \gamma^2\omega_f^2} \ , \tag{6.23}$$

$$b = \frac{F_0\gamma\omega_f}{(\kappa - m\omega_f^2)^2 + \gamma^2\omega_f^2} \ . \tag{6.24}$$

The restoring-force constant, κ, can be expressed in terms of the natural frequency of the particle: $\kappa = m\,\omega_0^2$. Thus the solution can be written,

$$x = \frac{F_0}{m^2(\omega_0^2 - \omega^2)^2 + \gamma^2\omega^2}\left[m(\omega_0^2 - \omega^2)\cos\omega t + \gamma\omega\sin\omega t\right] \ . \tag{6.25}$$

Making use of the trigonometrical identity,

$$c_1\cos\theta + c_2\sin\theta = \sqrt{c_1^2 + c_2^2}\cos\left(\theta - \arctan\left(\frac{c_2}{c_1}\right)\right) \ , \tag{6.26}$$

Equation 6.25 has the form,

$$x = C\cos(\omega_f\, t - \eta) \ . \tag{6.27}$$

That is, the motion is at the frequency of the external force but with a phase difference between the motion and the force. This phase can be evaluated by

substituting a and b for c_1 and c_2 in Equation 6.26:

$$\tan\eta = \frac{b}{a} = \frac{\gamma\,\omega_f}{m(\omega_0^2 - \omega_f^2)}\ . \tag{6.28}$$

Thus, the phase changes as the driving force passes through the natural motional frequency of the ion (ω_0), and the width of the change is determined by γ, the damping rate (Fig. 6.12). The phase change is centred around the natural frequency of the oscillator.

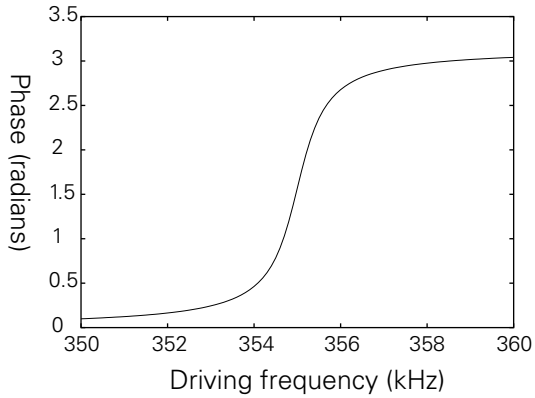

Figure 6.12: The phase change as the driving frequency passes through resonance with the natural frequency of the oscillator. The natural frequency was set to be 355 kHz and the damping was set at 1 kHz. The width of the arctan curve is the damping rate of 1 kHz, with $m = 1$.

6.6.2 Coupled oscillators

The system of two coupled harmonic oscillators is clearly more appropriate when an axialising drive is applied to the ion. The two oscillators are the modified cyclotron and magnetron motions. While a detailed theory describing the effect of coupling the two radial motions in the presence of damping is currently being developed in the group, a simpler model is presented here. Initial results of the more sophisticated model agree with this simpler theory.

Both of the radial motions are considered to be damped by the laser-cooling. Where there is a heating effect of one or both motions this will be manifest in a negative damping rate of the respective motion. A negative damping rate

could be due, for example, to the laser beam direction opposing the magnetron motion. We will consider only positive damping rates in this section.

As will be seen, the motional frequency and damping rate of the oscillators can be measured by applying an external force close to one of the motional frequencies and observing the relative phase between the external force and the motion. The equations of motion can be written,

$$\ddot{x_1} + \gamma_1 \dot{x_1} + \omega_1^2 x_1 - k x_2 = (F/m) e^{i\omega t} , \tag{6.29}$$

$$\ddot{x_2} + \gamma_2 \dot{x_2} + \omega_2^2 x_2 - k x_1 = 0 . \tag{6.30}$$

where γ_1 and γ_2 are the damping rates of the two motions, ω_1, ω_2 are their natural frequencies, k is the coupling constant, and F and ω are the amplitude and frequency of the excitation. x_1 and x_2 are respectively the coordinates of the two motions.

The solutions will have the form,

$$x_1 = x_{10} \exp i(\omega t + \phi) , \tag{6.31}$$

$$x_2 = x_{20} \exp i(\omega t + \theta) . \tag{6.32}$$

The amplitudes of the motion are x_{10} and x_{20} respectively while their phases with respect to the excitation are ϕ and θ. Differentiating and substituting back into the equations of motion yields the solution,

$$x_{10} e^{i\phi} = F \times \left\{ (\omega_1^2 - \omega^2) + i\gamma_1 \omega - \frac{k^2}{(\omega_2^2 - \omega^2) + i\gamma_2 \omega} \right\}^{-1} . \tag{6.33}$$

The amplitude of the first mode (x_1) is given by the modulus of the reciprocal of the denominator in Equation 6.33. The phase of the motion is given by the arg of the denominator. These two characteristics of the motion have been plotted in Fig. 6.13.

It is seen that not one but two phase changes are observed. The stronger peak in Fig. 6.13 is associated with the phase change corresponding to the damping rate of the motion being driven. The second, much weaker peak is associated with the other oscillator. The separation of the peaks is found to be determined by the coupling rate.

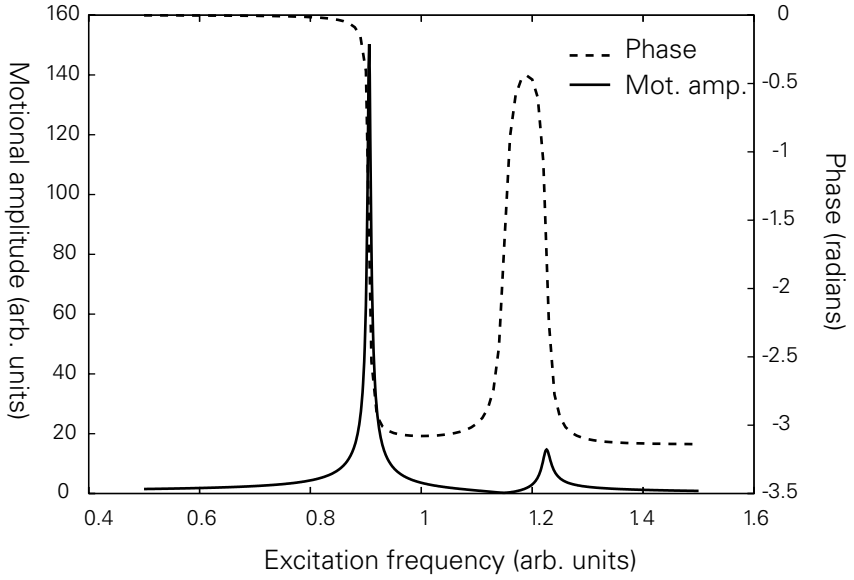

Figure 6.13: The modes of motion and the phase changes associated with them when one of the modes is driven by an external force. The two peaks arise due to the coupling of the driven motion with another motion. The two phase changes will have widths corresponding to the damping rates of the two motions. The parameters used were natural frequencies of 1.0 and 1.15 and damping rates of 2×10^{-3} and 2×10^{-2} for these two motions respectively. The coupling rate was 0.3. The frequency units are consistent but arbitrary.

6.7 Coupling and avoided crossings

If, instead of solving for x_{10} as in the previous section, we solve for ω, the normal modes of the system, we get the expression,

$$\omega^4 - (\omega_1^2 + \omega_2^2 + \gamma_1\gamma_2)\omega^2 + (\omega_1^2\omega_2^2 - k^2) = 0 \ , \tag{6.34}$$

having set the phases, ϕ and θ, to zero and in the limit that F goes to zero since we are interested in the normal modes rather than the motion induced by the external excitation. This is a quadratic equation in ω^2 with the solutions,

$$\omega^2 = 1/2 \left[(\omega_1^2 + \omega_2^2 + \gamma_1\gamma_2) \pm \sqrt{(\omega_1^2 + \omega_2^2 + \gamma_1\gamma_2) - 4(\omega_1^2\omega_2^2 - k^2)} \right] \tag{6.35}$$

That is, the separation of the two peaks is dependent on the natural frequencies of the motions and the coupling strength, k. This results in what has been

termed a "classical avoided crossing" [102] (see Fig. 6.14).

Avoided crossings can be observed in two ways: either by changing one frequency such that it crosses the other, or by changing the frequency of the coupling such that it passes through resonance with the two modes. The classical avoided crossing described by Equation 6.35 occurs as one motional frequency crosses the other. This is, in fact, equivalent to the coupling frequency passing through resonance with the two modes (i.e. the sum or difference of the two motional frequencies as appropriate).

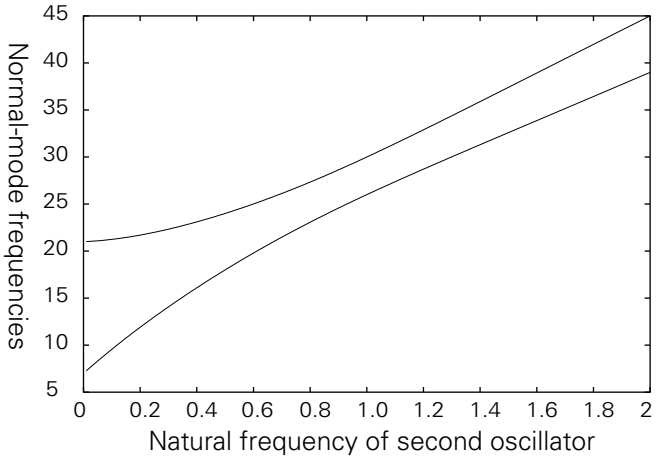

Figure 6.14: The Classical avoided crossing. As the natural frequency of one of the oscillator approaches that of the other, the normal modes deviate from those expected to produce a classical avoided crossing. The first oscillator's mode is 1.0 (arbitrary units) and the second oscillator's natural frequency is varied.

6.7.1 Dressed states

Avoided crossings also occur in other branches of physics, most notably in the dressed-state picture of atom–photon dynamics. Where a pair of atomic energy levels is coupled by a light field consisting of photons with an energy close to the separation of the energy levels, the system may be described in terms of "dressed" states [107].

As seen in Fig. 6.15, the photons couple the ground state and excited state and the two levels become almost degenerate. Instead of representing the system by the atomic states $|g\rangle$ and $|e\rangle$, and the states of the single-mode field, $|n\rangle$ where n is the number of photons in the field, the system may be represented

by dressed states. The new basis states are represented $|+,n\rangle$ and $|-,n\rangle$ and their separation is determined by the Rabi frequency of the transition. The Rabi frequency in turn is determined by the intensity of the light field (the coupling strength) and the detuning of the photons from resonance with the atomic transition.

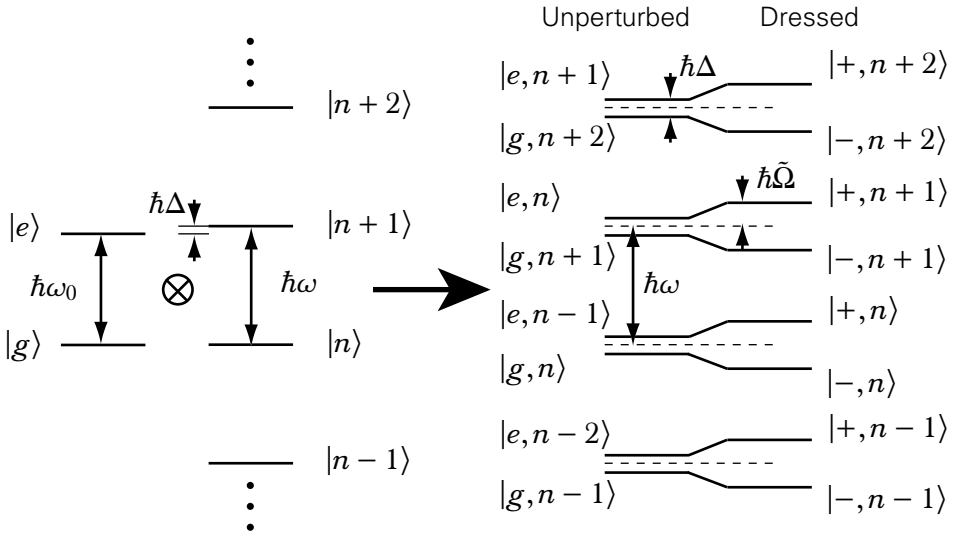

Figure 6.15: Dressed-state picture of photon–atom interactions. The state $|g,n\rangle$ is degenerate with $|e,n-1\rangle$ if the photons in the single-mode field are resonant with the atomic transition. In the more general picture depicted here the photons have a slightly higher energy by an amount $\hbar\Delta = \hbar\omega - \hbar\omega_0$. After Ref. [107].

If the frequency of the photons coupling the two states is increased from below resonance, the energy of the lower state, $|-,n\rangle$ approaches that of the higher one, $|+,n\rangle$, as seen in Fig. 6.16.

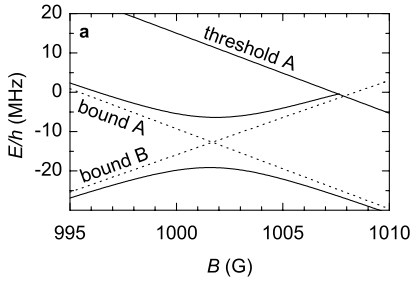

Figure 6.16: A quantum avoided crossing. In this case the two states are those of a free atom and an atom bound in a molecule. The coupling is achieved by bringing the states into resonance with each other with a magnetic field. *Threshold A* is the dissociation threshold of the molecule and is not a feature of avoided crossings in general. The plot is a theoretical plot from Ref. [108].

6.8 Axialisation in the presence of laser-cooling

Some important differences exist between axialisation with buffer-gas cooling and axialisation with laser-cooling. The first is that with buffer-gas cooling there is little control over the damping rate for the modified cyclotron motion and the heating rate for the magnetron motion. By contrast, the laser frequency and spatial position both affect the damping or heating rates for the two motions. Indeed, it is possible to choose whether one or other motion is damped or heated by appropriate choice of the two parameters. The damping rates for the motions are γ_c and γ_m respectively and negative values of these damping rates indicate heating of the motion.

The second difference is that there is a spatial dependence of the damping rates. That is, with the laser parameters fixed, the ions will only be heated or cooled in a restricted volume of the trap through which the beam passes. The cooling and damping rates associated with buffer-gas cooling are homogeneous throughout the trapping volume.

Expressing the axialisation effect in terms of the radii of the two motions, the following coupled equations are obtained [1],

$$\dot{\rho_c'} = k\rho_m - \gamma_c\rho_c' \,, \qquad \dot{\rho_m} = -k\rho_c' - \gamma_m\rho_m \,, \tag{6.36}$$

where, once again, ρ_c' and ρ_m are the radii of the modified cyclotron and mangetron motions respectively, γ_c and γ_m are their respective damping rates, and k is the coupling rate between the two motions. Assuming both ρ_c' and ρ_m evolve as $\exp(-\Gamma t)$, the expression for this coefficient, Γ is,

$$\Gamma = \left[-(\gamma_c + \gamma_m) \pm \sqrt{(\gamma_c + \gamma_m)^2 - 4(\gamma_c\gamma_m - k^2)}\right]/2 \tag{6.37}$$

This expression indicates that the types of behaviour possible are more varied with laser-cooling than with buffer-gas cooling. The two standard cases are that of no damping and that of strong modified cyclotron damping and weak magnetron heating. These are the same as with buffer-gas cooling. Without any damping at all, $\gamma_c = \gamma_m = 0$ and $\Gamma = \pm k$. This results in an oscillation between pure modified cyclotron and pure magnetron motion. In the latter case, $\gamma_c \gg 0$ and $\gamma_m < 0$. With $k^2 > -\gamma_c\gamma_m$, Γ is large and the two radii tend towards zero. This describes axialisation and is illustrated (in the lab frame) in Fig. 6.17.

This gives an important relation on how strong the coupling must be in order to achieve axialisation.

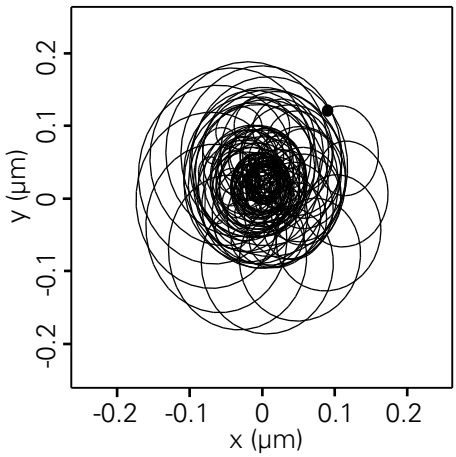

Figure 6.17: When the modified cyclotron motion is laser-cooled and there is a slight heating of the magnetron motion (as is the normal situation), along with a laser beam through the centre of the trap, the modified cyclotron radius is reduced. Through the coupling between the two motions the magnetron radius is also reduced and the ion spirals in towards the axis of the trap. From Ref. [105].

The behaviour unique to a laser-cooled system is where the laser-beam is off-set from the centre of the trap slightly. The radii decrease as in normal axialisation but eventually become smaller than the laser's off-set. The ions start to heat as they spend less time in the laser beam and the radii increase again. An equilibrium between these cooling and heating effects is set up such that the condition $k^2 = -\gamma_c\gamma_m$ is fulfilled. Such stable orbits have been observed in experiments with magnesium ions [105].

Chapter 7

Axialisation of magnesium ions

7.1 Introduction

Initial work in the Imperial College group on axialisation was performed with magnesium ions by H. Powell [109, 105]. This was the first time azimuthal quadrupolar excitation was demonstrated in the presence of laser cooling. It was decided to continue the work on magnesium in order to measure the enhanced cooling rate of the magnetron motion in the presence of the radial quadrupole field [1].

Magnesium is a good choice for the study of ion dynamics in the presence of laser cooling due to its simple level structure. The major drawback with the use of magnesium is that the cooling transition is currently only accessible with a frequency-doubled dye laser, which is rather harder to use than semiconductor and solid state lasers.

7.2 Doppler-cooling magnesium

Magnesium has been used in several ion-trapping experiments elsewhere[110, 111, 112, 113]. Since no re-pumping lasers are required the set-up for laser cooling is quite straightforward. Difficulties encountered include the rather low atomic mass, however: trap parameters suitable for the trapping of magnesium can also trap lighter atomic or molecular impurity ions.

7.2.1 Energy level structure

The relevant energy levels are shown in Figure 7.1. The $^2S_{1/2}$–$^2P_{3/2}$ transition near 280 nm is 43 MHz in width (giving a Doppler limit of 1 mK). The absence of D states between the lower and upper levels of the cooling cycle means that repumping lasers are not required. It also means, of course, that there are no long-lived energy levels that can be used for a qubit. As experimental work in QIP moves towards the use of hyperfine levels for qubits, magnesium may, however, come into its own again. As can be seen in Ref. [114], the properties of the ion are suitable for such work.

In Ref. [114], Wineland assesses the suitability of various ions with hyperfine structure for QIP. One of his criteria for suitability is the ratio of the spontaneous decay rate of the relevant transition to the hyperfine splitting. Although there is negligible decoherence due to decay of the qubit states, some decoherence can be introduced during operations through decay of the upper level of the Raman transition. Although $^{25}Mg^+$ has a relatively high spontaneous decay rate of 43 MHz, the hyperfine splitting is sufficiently large at 1.79 GHz to make the probability of spontaneous emission nearly an order of magnitude smaller than that of $^9Be^+$ (which has already been used in QIP experiments [40]). Although other ions have probabilities an order of magnitude lower than this, the simplicity of laser cooling magnesium is in its favour. New semiconductor sources at 280 nm have recently been developed, albeit not in the form of laser diodes. Were a didoe laser available at 280 nm laser cooling of magnesium would be very straightforward indeed.

7.2.2 Laser system

The cooling laser was a modified dye laser (Coherent CR699-21) pumped by an argon-ion laser (Coherent Innova 400). The argon-ion laser was operated at 514.5 nm and was capable of delivering up to 10 W of output power. It was normally run at 6 W in the experiments described below. A home-made mount was included in the cavity of the dye laser to hold a frequency-doubling crystal (β-barium borate or BBO) at the auxiliary waist of the ring-cavity beam. A dichroic mirror allowed the doubled light to exit the laser while the the fundamental beam was reflected back around the ring.

The ring dye laser contained many intra-cavity components (see Figure 7.2) in order to ensure single-mode operation with ease of tunability. The inclusion

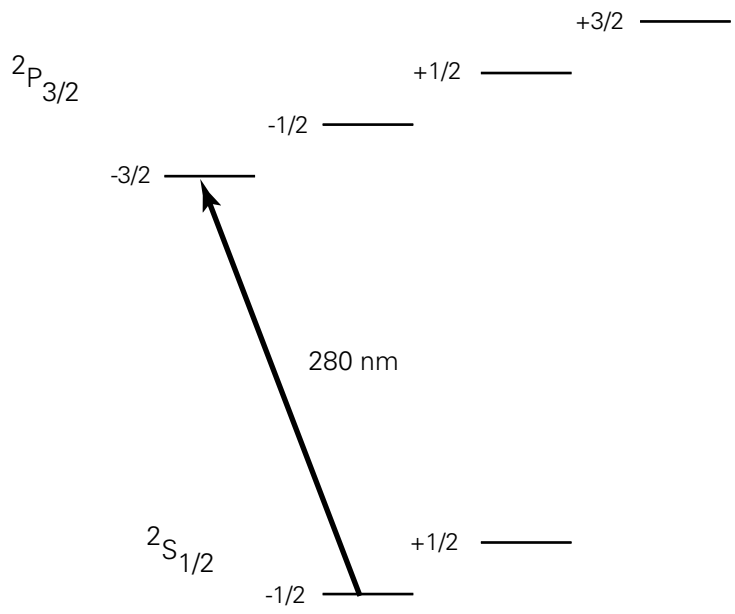

Figure 7.1: Zeeman-split energy levels of $^{24}Mg^+$ with cooling transition labelled (not to scale)

of such a large number of components makes lasing harder to achieve since there is an inevitable loss associated with each one. The main components are outlined here:

Dye jet The dye (Rhodamine 6G) is pumped by an external pumping module to the laser head. The dye passes through a jet assembly which causes it to pass through free space as a well-collimated, flat jet to a collector a few centimetres away. The angle and shape of the jet can be adjusted with 12 machine screws. The angle with respect to the pump beam is set to be Brewster's angle to ensure that as much of the pump light as possible is absorbed by the jet, and in order to minimise stray reflections, which would be potentially dangerous. The collected dye flows back to the pumping module and is recirculated.

Mirrors In the factory-shipped laser there are four highly reflective mirrors plus an output coupler. The output coupler was replaced with a high reflector when the direction of travel of the beam was changed. The first mirror focuses the pump beam onto the dye jet while the other four form the ring around which the beam passes. There are two waists in the beam, the first

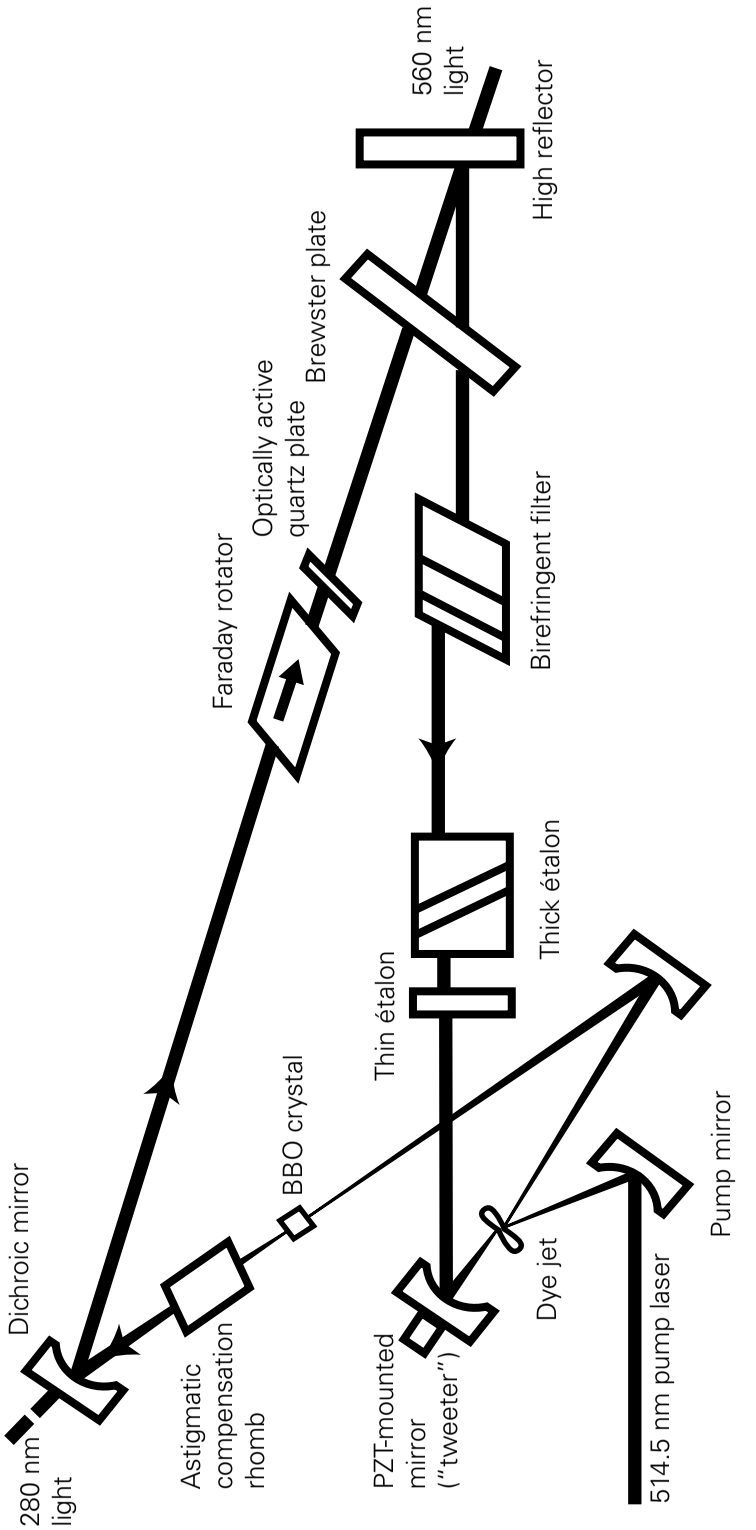

Figure 7.2: Simplified schematic of the Coherent CR699 dye laser.

occurring in the dye jet as just mentioned, and the other occurring between the first high reflector and the dichroic mirror. It is in this auxiliary waist that the frequency-doubling crystal is placed.

Astigmatic compensation rhomb The astigmatism for which this compensates arises from the reflections from the mirrors in the cavity. The rhomb was previously modified to compensate for the additional astigmatism introduced by the Brewster-cut frequency-doubling crystal [105].

Optical diode An optical diode based on the Faraday effect introduces significant losses in one direction of travel. The polarisation of the light is rotated by a small amount by a strong magnetic field in a crystal. The direction and amount of rotation are independent of the direction of propagation. A wave-plate rotates the polarisation as well, but this time the direction is dependent on the direction of propagation. For the combined system of the Faraday crystal and wave-plate, the rotations will add in one direction and in the other they will subtract. The additional rotation in one direction causes losses through reflection off Brewster-angled components. Without the diode in place the lasing would take place in both directions around the ring cavity and the power would be split between the two directions. The direction of travel was, in fact, reversed from the standard operation of this laser. This modification was made in order to obtain maximum output power of the frequency-doubled light through the dichroic mirror.

Birefringent tuning filter This allows very large-scale tuning of the laser. The output colour of the laser can easily be seen to change as the filter is turned through large angles. The filter is formed from three crystalline quartz plates and limits the output bandwidth to a 20 GHz section of the dye's gain curve.

Etalon One of the most lossy intra-cavity elements, this ensures single-mode operation through a combination of two low-finesse etalons. The thin etalon has a free spectral range (FSR) of 225 GHz while the thick etalon has an FSR of 10 GHz. The angle of the former with respect to the beam is adjusted mechanically. The latter etalon is formed from a pair of Littrow prisms which can be moved by changing the voltage on their PZT mounts. This allows frequency tuning in everyday operation while keeping losses low.

PZT-mounted mirror The "folding" mirror is mounted on a piezoelectric transducer (PZT) with a large bandwidth (greater than 1 kHz). This mirror is used for locking the cavity and for small-scale tuning of the laser. The high bandwidth of the the servo means that it is sometimes referred to as the "tweeter", a reference to hi-fi technology.

Vertex-mounted Brewster-angle plate Rotating the plate changes the optical path length of the cavity as the beam passes through a longer or shorter section of the glass plate. Changing the effective cavity length continuously tunes the laser output frequency. The plate provides low-frequency but large-scale corrections to the laser frequency when the laser is locked. Again, in reference to hi-fi terminology it is referred to as a "woofer".

The laser is frequency-stabilised through the use of a temperature-stable reference cavity. A small portion of the laser output is picked off by a beam splitter. A photodiode on part of this pick-off beam monitors the laser intensity while the remainder is passed through the reference cavity. A side-of-fringe lock is implemented by the control electronics, compensated for intensity fluctuations through the first photodiode signal. Feedback is to two components: the folding mirror PZT mount, and the Brewster plate near the output coupler. The quoted jitter is less than 1 MHz RMS with a frequency drift less than 100 MHz/hour.

The reference cavity also contains a Brewster plate, with a galvo-mount. Rotating the plate changes the path length within the cavity (in the same way as in the laser cavity). Thus the fringes of the reference cavity move in frequency space and the lock point corresponds to a different laser frequency. This allows the laser frequency to be scanned through feedback to the laser-cavity elements as the plate is rotated at low frequency.

The main output of the fundamental beam (at 560 nm) was passed through an iodine cell and the fluorescence monitored on a photodiode. By scanning the laser via the reference-cavity galvo plate, the iodine spectrum could be observed on an oscilloscope. This gave an aboslute frequency reference which could be used to tune the laser to the cooling transition (or more exactly to half the frequency of the cooling transition).

Typical output powers of the laser (with a 6-W pump beam) were 300 mW at 560 nm and 50 μW at 280 nm. The maximum conversion efficiency of the frequency-doubling crystal was about 0.1%. For this reason, and given the huge loss involved in having an intra-cavity doubling crystal, it would have been

preferable to have had an external frequency-doubling cavity. The quoted output power for the laser with Rhodamine 6G dye is 800 mW single-mode with a 6-W pump beam. With this input power to a commercial frequency-doubling cavity such as the Coherent MBD-200, one would expect an output power of roughly 40 mW.

7.3 Axialisation

Single ions and small clouds of magnesium had been routinely trapped and axialised [109, 105]. Our goal in this experiment was, therefore, not to demonstrate the feasibility of axialisation but rather to go some way towards characterising it.

The trap used is described in Chapter 4. Two difference between this and subsequent axialisation experiments with calcium are the drive electronics for applying the radial quadrupole field and the method of motional excitation.

The axialisation-drive electronics were those built by Powell and described in Ref. [105]. A gain of 0.4 (i.e. a 60% reduction) is incorporated in the electronics. Where axialisation drive amplitudes have been quoted in this chapter, they are without the gain of 0.4. The method of excitation of motional frequencies was to apply an RF voltage to the disused atomic beam oven. Although we aim to drive radial motions and would thus prefer to have a radial dipole driving these motions, it was found that driving them through the oven was effective. There was a sufficiently large component of the generated field in the correct geometry to drive the motions and to see signal in RF–photon correlation traces. Radial dipole excitation was later performed with axialised clouds of calcium ions (see Chapter 8).

The axialising drive and motional-excitation signal were provided by two Hewlett Packard function generators (models HP8643A and HP3325B, respectively). One important point to note is that the frequency scan functions on both devices provide only pseudo-phase-continuous scans; the signal switches off and back on again at each frequency step. This was circumvented in later experiments with calcium (Chapter 8) using a more restricted but truly phase-continuous scan available on the HP8643A function generator.

7.3.1 Loading technique

In order to load the ions the dye laser was typically set to scan from 1 GHz in the fundamental (i.e. 2 GHz at 280 nm) from below resonance up to resonance, with a scan-time of 1 s. This "sweeps in" ions with large initial velocities and large Doppler shifts with respect to the cooling laser. The ions were loaded by heating an oven with a magnesium sample of natural abundance to generate an atomic beam and ionising the atoms in the centre of the trap with a crossed electron beam. Typical heating currents (of order 3 A) for the oven are lower than for calcium due to the higher vapour pressure of magnesium.

The number of ions entering the trap was controlled by the oven current. In all experiments the activation times for the oven and filament were fixed and the only parameter varied was the current passed through the oven. The oven was run for 20 sec on its own and then for a further 10 sec in unison with the filament producing electrons. Both devices were turned off simultaneously.

Fluorescence signal from the ions was normally seen towards the end of the loading process. The shape of the peaks observed on the MCS is the Doppler-broadened transition linewidth.

As discussed in Chapter 2 the laser beam position is vital in achieving efficient cooling of ions in a Penning trap. The beam was offset upwards of the centre of the trap by about 100–200 μm such that the **k**-vector of the photons was parallel to the magnetron motion. After loading, the beam was moved down the trap to maximise the fluorescence signal. Moving the beam too far down caused the signal to rapidly die away as the magnetron motion was heated. It was normally possible to recover the signal after this by rapidly moving back up and slowly back down towards the centre again. Too slow a response at this stage meant that the ions were heated out of the trap before the beam was moved back up.

7.3.2 Turning on the drive

The magnetic field was generally well known (to within 1%) since the laser had previously been tuned to a given point on the iodine spectrum corresponding to the transition wavelength in magnesium at a given magnetic field. In addition, the magnet was operated near its maximum current; at low currents the magnetic field increases roughly linearly with current but near the top of the range the curve flattens off considerably. Thus, a small change in current within the

time-scale of the experiment or from day to day did not yield a greatly different magnetic field. Given this knowledge it was possible to assume the same cyclotron frequency at the beginning of each experiment rather than scan the axialisation drive frequency looking for a clear response in ion signal.

Working at a field of 0.98 T the true cyclotron frequency was $\omega_c = 2\pi \times 620\,\text{kHz}$. A signal of this frequency and of amplitude 1 V peak-to-peak was applied to the circuit providing the phase inversion. The outputs of the circuit provide both the in-phase and the out-of-phase signals with a gain of 0.4 (i.e. 400 mV peak-to-peak in this case) and these were applied to trap. Each pair of opposing ring segments was connected to one of the phases.

Before axialising the cloud, the laser scan had been turned off and the laser frequency manually tuned to within about 100 MHz of resonance. This would cause the fluorescence signal to rise to approximately 2/3 of its maximum. Going closer to resonance than this was difficult due to the danger of going over resonance and rapidly heating the ions out of the trap.

Turning on the drive gave rise to an immediate increase in fluorescence. It was also found that it was possible to increase the laser frequency to be much closer to resonance and indeed to go over resonance without catastrophic effects: it was possible to simply reduce the frequency again and restore the original fluorescence level. This was a clear indication that the magnetron motion was being efficiently cooled. It was also possible to move the laser beam closer to the centre of the trap than before, typically by an amount comparable to the beam waist or more. In such a position the magnetron motion would previously have been heated.

7.4 RF–photon correlation

In order to measure the improvement in the magnetron damping rate afforded by the axialisation drive, the method of RF–photon correlation described in Chapter 6 was employed. This requires the measurement of the phase difference between the excitation and the motion. This is possible by measuring the time between a given phase of the excitation (say, 0 rad) and the arrival of the next photon at the photomultiplier. If these times are passed to a multi-channel analyser (MCA), an exponentially decaying signal is obtained with modulation at the motional frequency. The exponential decay is a result of the high probability of the arrival of one photon shortly after another. The phase of the

modulation at the start of the trace changes linearly with the phase difference between the drive and the motion.

The reason the motion of the ion manifests itself as a modulation of the MCA signal is that the ion moves in and out of the beam. Each time it passes through the beam it can absorb and re-emit a photon, and the time between each pass through the beam is half its motional period. This description is generally true for the magnetron motion but may not apply to the modified cyclotron motion (or indeed to the magnetron motion itself) if the beam waist is larger than the magnetron radius. In this case, the modulation is due to the ion coming in and out of resonance with the beam thanks to the Doppler shift (see Fig. 7.3). At certain points on its trajectory the component of its momentum in the direction of travel of the photons is such that the Doppler-shifted resonance matches the detuning of the laser and a photon is absorbed.

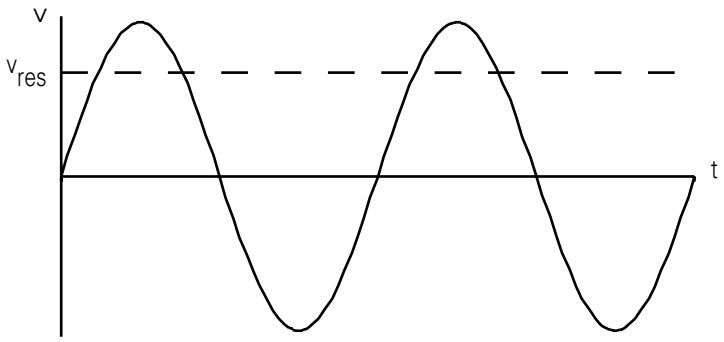

Figure 7.3: As the ion's velocity varies in its orbit, it passes through the velocity equivalent to resonance with the detuned laser beam.

In oder to measure the time between the given phase of the RF and the next photon arrival, a Time-to-Amplitude Converter (TAC) is used. The TAC is triggered by the phase of the motional drive signal and stopped by the next photon detected. Fitting a modulated exponential curve to the MCA trace allows the phase of the modulation (i.e. the phase difference between the excitation and the motion) to be determined.

It is usually desirable to have a small cloud of ions when making these measurements since the overall fluorescence level is higher. This gives rise to a shorter data acquistion time and makes the experiment less susceptible to effects such as poor laser frequency- and intensity-stability. It was found that clouds of about 10–20 ions allowed rapid acquisition of data (less than a minute)

while still preserving a sufficient modulation that an accurate fit of phase and frequency could later be performed.

Traces were taken for discrete drive frequencies around the motional frequency of interest. A sample for a small cloud of ions is show in Fig. 7.4. The calculated phases were then plotted against drive frequency to give curves such as those in Fig. 7.5. The width of the arctan function on the phase plot is a direct measure of the cooling rate.

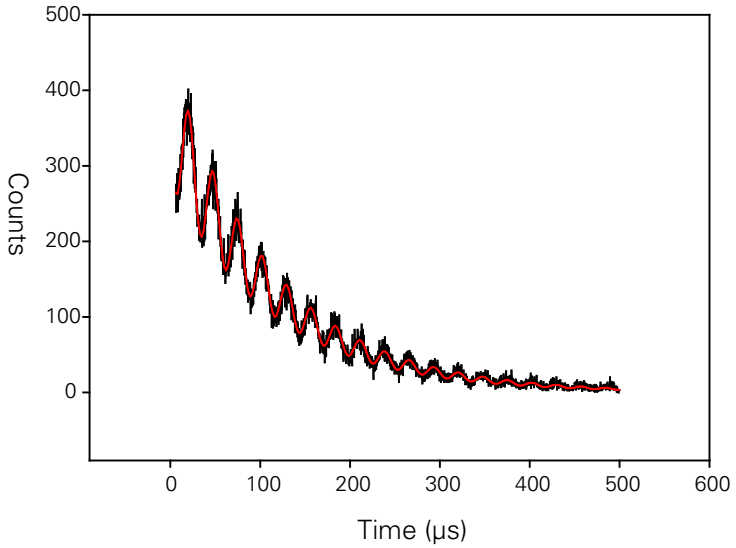

Figure 7.4: A typical MCA plot for a small cloud of ions. The oscillation frequency superimposed on the exponential is at the frequency of the driving radiation (36.6 kHz). The initial phase is calculated from the fitted function.

7.5 Measured cooling rates

Previous RF–photon correlation experiments with beryllium and magnesium ions [67] indicated a very low cooling rate (<100 Hz) for the magnetron motion. At the same time, the modified cyclotron cooling rate was of order several kilohertz. This fits with a computational model [27] which predicts that simultaneous weak cooling of the magnetron motion and strong cooling of the modified cyclotron motion is possible with appropriate choice of laser frequency and spatial offset.

Experiments to measure damping rates prior to this work, described in Ref. [67], were performed without an axialisation drive. It is to be expected that the magnetron cooling rate will increase through coupling to the laser-cooled modified cyclotron motion.

Small clouds of magnesium ions were loaded. The number of ions can be estimated by knowing the count rate per ion. In the case of magnesium, single ions had previously been trapped and the count rate measured to be about 5 kHz. For very small numbers of ions the number can be inferred from the number of levels in quantum jumps in the measured fluorescence. For a single ion there will be two levels (shelved and fluorescing), for two ions there will be three levels (both shelved, one shelved, both fluorescing), and so on. Beyond about four ions quantum jumps are harder to distinguish. The count rate also fails to rise exactly linearly since space charge means that the cloud no longer spends as much time within the laser beam.

Typical count rates from the PMT were 100 kHz without axialisation indicating a cloud size of roughly 20 ions. The motional drive frequency was applied to the unused atomic beam oven and that frequency was also connected to the TAC (EG&G Ortec 566 TAC). Using this signal to trigger the TAC, and detected photon signals to stop the TAC as described above, the times were fed to the multi-channel buffer card (EG&G Ortec 916 MCB) in a PC and analysed with the MCA software (EG&G Ortec Maestro II, A64-BI). A modulated exponential curve rapidly became visible on-screen and was saved to disc after about one minute.

MCA traces were taken at 100 Hz steps for the 2 kHz region around the motional frequency. Some additional traces were taken at 10 Hz steps very close to the motional frequency when the cloud was axialised. It was considerably more difficult to take traces without axialisation (as one would expect) since the cloud was easily heated out of the trap by the motional excitation. In both cases of axialised and unaxialised clouds it was necessary to load a new cloud regularly due to laser mode-hops and other laser instabilities. This should not impact the quality of the results.

The following function was fitted to the MCA traces using SigmaPlot:

$$f(\Delta t) = \exp(-a\,\Delta t)\left[b + |c|\sin\left(\omega\,\Delta t - \phi\right)\right] \ . \tag{7.1}$$

a, b, c, and ϕ are all fit parameters. The first three are related to the average

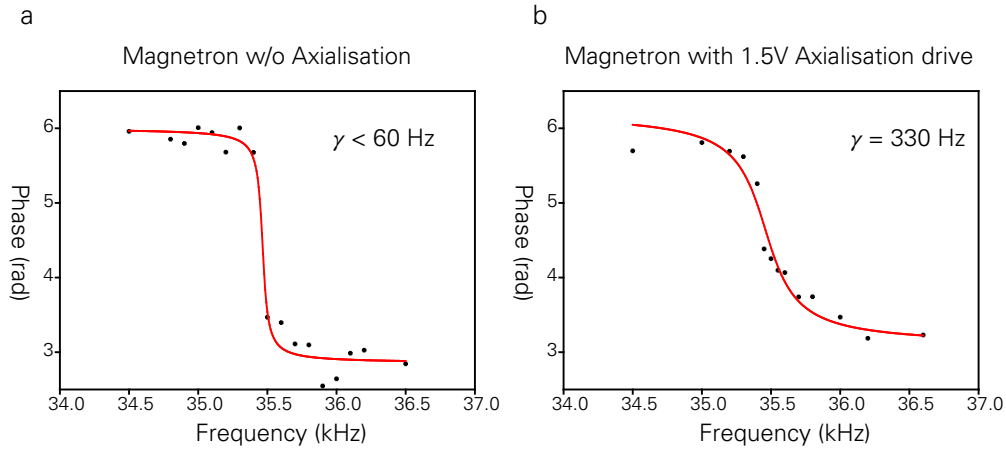

Figure 7.5: The phase change with change in the driving frequency near the magnetron frequency. In *a* the ions were unaxialised and γ, the cooling rate is low and the transition is narrow. The phase change is broader in *b* where the ions were axialised.

photon count rate, the modulation depth, and the acquisition time. ϕ is the phase of the motion with respect to the drive, the parameter of interest to us here.

ϕ was plotted against the driving frequency and is shown in Fig. 7.5. The width of the change in phase when the cloud is unaxialised is less than 0.06 kHz, as seen before [67]. When the axialisation drive is on, the width of the change of phase is considerably broadened to 0.33 kHz. Equivalent plots for drive frequencies close to the modified cyclotron frequency show no increase (or decrease) in the width of the phase change.

Very few datasets could be taken before technical difficulties arose with the laser. A general trend of increasing magnetron damping with axialisation voltage could be seen, however (Fig. 7.6).

These results demonstrate both efficient cooling of the modified cyclotron motion and coupling between the modified cyclotron and magnetron motions.

Previous work [109, 105] demonstrated the reduction in the magnetron radius with an imaging technique and assigned a consequent upper limit on the ion temperature of approximately 10 mK. This value approaches the Doppler limit and shows that the problems of localisation of the ion and of Doppler-cooling can be overcome.

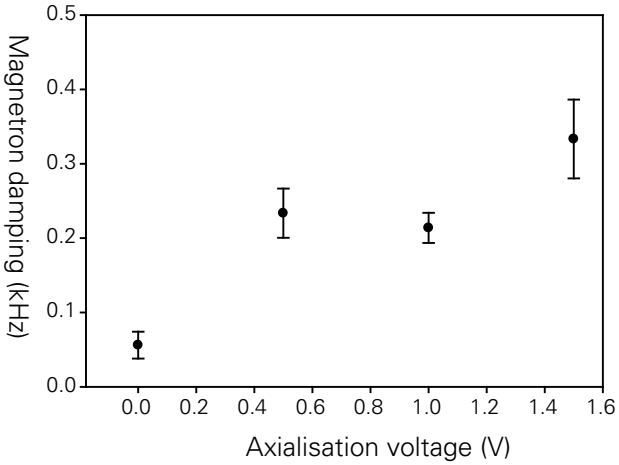

Figure 7.6: Magnetron damping with axialisation amplitude. The damping rate increases with the axialisation-drive amplitude. The first pair and second pair of data points were taken on different days with different clouds of ions. The points should be seen as a lower limit of the achievable damping rate since it is difficult to ensure that the cooling is always optimal.

Chapter 8

Axialisation of calcium ions

One of the main goals of the group over the past two years has been the axialisation of single ions of calcium. As previously described, this is part of an effort to investigate the feasibility of quantum information processing with ions in a Penning trap. Single Ca^+ ions have been routinely loaded with the trap in a Paul-trap configuration but the smallest number of ions loaded in a Penning trap in earlier experiments was estimated to be two or three [65]. Single ions of Mg^+ have been consistently loaded in a Penning trap in the past [105] and one of the major differences between the trapping of the two ion species is the complexity of the laser-cooling system required. It was hoped, therefore, that improvements to the efficiency of the laser-cooling would help in respect of loading single calcium ions and a large number of changes were made to the experimental set-up to this end.

One of the major improvements to the laser-cooling to be implemented was the axialisation of the ions. This has allowed the more-consistent loading of very small clouds of Ca^+ ions. We have yet to succeed, however, in trapping a single ion. Nonetheless, the ability to work with such small clouds of ions has allowed us to infer that the maximum count-rate per ion is much lower than for ions trapped in the Paul-trap configuration. It is believed that the reason for this low count rate may be the shelving of ions in coherent trapped states (lambda systems). Simulations are currently being developed by the group to test this theory and to seek methods to overcome the limitations.

8.1 Axialisation-drive circuit

A new circuit was built to provide the correct phases of signal to the ring-electrode segments. The new design allows the combination of an oscillating quadrupole field with an oscillating dipole field, enabling a small radial excitation voltage (dipole) to be applied at the same time as the axialisation drive (quadrupole).

The main requirements of the design were that the phase should be accurate in both the zero- and π-phase outputs, and that the circuit would be able to handle intermediate frequencies (up to 1 MHz). To these ends, buffers were included in the circuit to ensure that the signal path-lengths were similar, and relatively high-performance op-amps (model TL074ACN) were used. Care was taken in the construction of the circuit by building it on a ground-plane board, shielding it well from external sources of RF, and by decoupling power inputs to chips to the ground plane via capacitors.

The circuit is shown in Fig. 8.1. The two inputs are from function generators tuned to the axialisation frequency and the desired motional excitational frequency respectively. The three outputs provide the two phases of the axialisation signal with the excitation signal's two phases superimposed on one.

Later tests on the circuit showed that a variable DC offset existed on the axialisation outputs. This was a few millivolts in magnitude and did not increase with the input signal's amplitude (Fig. 8.2). The gain was also measured for both channels and found to be linear. A slight difference in gain (approximately 9%) was found between the channels which is probably due to the use of imprecise resistor values. This difference in gain should not affect the functioning of the axialisation drive and will not introduce any additional DC offset.

8.2 Imaging of clouds and qualitative results

Having trapped some ions, the clearest way to demonstrate axialisation is to image them in real-time and turn on the drive. With the intensified CCD (ICCD) camera on maximum gain, exposure times as low as 100 ms gave clear images of the ion cloud's extent. The geometry of the trap and the fluorescence-collection optics is such that the images are side views of the cloud, i.e. taken in the radial plane. The laser beams passing through the trap follow a horizontal line in the images. Both the camera and optics are described in Chapter 4.

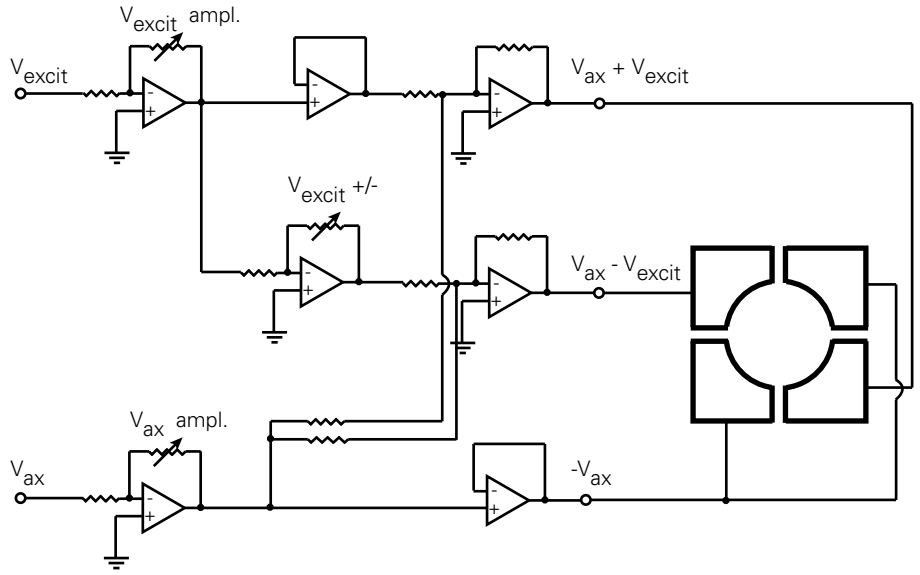

Figure 8.1: The axialisation-drive circuit and the connections to the ring electrode. Control over the axialisation-drive and excitation signal gains (both between 0 and 1) is provided by a pair of variable resistors (V_{ax} ampl. and V_{excit} ampl. respectively). There is also the facility to change the relative amplitudes of the excitation drive phases (V_{excit} +/-) if necessary. All fixed resistances are the same and are equal to the maximum value of the variable resistors (10 kΩ here).

The approximate value of the axialisation-drive frequency was known from a previous calibration of the magnetic field generated by a given current. This gave a starting value for the axialisation drive frequency, calculated from the formula,

$$\omega_c = \frac{qB}{m} , \tag{8.1}$$

where q is the charge on the ion, B is the magnetic field and m is the mass of the ion.

With the axialisation-drive circuit connected to the ring electrode as shown in Fig. 8.1, the function generator (V_{ax}) was turned on. The frequency was varied while observing the cloud with the camera. With the frequency close to the true cyclotron frequency, a clear reduction in the cloud's width was seen. This change took place on a time-scale that was fast compared to the refresh-rate of the camera (here set to 10 Hz). Both blue lasers were kept static in frequency while taking images in order to get a uniform, high, fluorescence rate. The

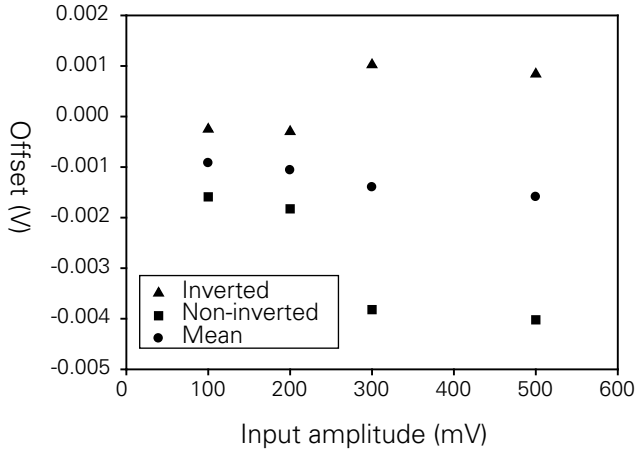

Figure 8.2: Variation of the DC offset of the outputs of the circuit in Fig. 8.1 with input amplitude. It can be seen that the average offset is small and does not vary greatly with the input voltage. This is important as a DC offset could change the motional frequencies of the ions by altering the effective trap bias.

lasers could be tuned higher in frequency once the axialisation drive had been applied and the beam could be moved further towards the centre of the trap than before without heating the ions. This behaviour was the same as that seen with magnesium ions (Chapter 7). The immediate increase in fluorescence that had been seen with magnesium ions when the axialisation drive was turned on was not observed with calcium ions. This may have been due to the fact that the typical cloud size was larger and that space-charge effects may have prevented the cloud from being reduced in size to the same extent.

The frequency-sensitivity of the axialisation drive was found to be high. A detuning of tens of Hertz had a noticeable effect on the cloud size. This sensitivity was dependent on beam position as before. The drive frequency could initially be set to an approximate value with the beam position optimised for an unaxialised cloud. The beam could then be moved closer to the centre and the axialisation frequency re-optimised with greater sensitivity. It was also found that the optimal frequency remained the same over a period of many hours. This indicated that the magnetic field stability was at least one part in 10^4 on this time-scale.

In early experiments where this sensitivity of the axialisation-drive frequency was investigated, a Hewlett Packard 3325B function generator was used,

with a quoted accuracy of 1.9 Hz at our operating frequency. We changed to using a Hewlett Packard 8643A function generator in later experiments, when attempting phase-continuous scans of frequency. Although the quoted frequency accuracy of the 8643A was better (0.14 Hz), a malfunction was discovered after some weeks of use that meant that the output frequency differed by as much as 5% from the value displayed on-screen. The output frequency was found to remain unchanged when the desired frequency was varied over quite a large range (of order 500 Hz). This behaviour was entirely reproducible and gave rise to an irregular stepping effect on a plot of output frequency versus programmed frequency. Although this malfunction required us to re-evaluate some later data, the sensitivity of the axialisation drive had already been established using the HP 3325B, a stable and accurate function generator.

Images of the ion clouds were taken using the Andor intensified CCD (ICCD) camera described in Chapter 4. The side of each pixel in the camera is 13 μm in length. The image intensifier has the disadvantage of reducing the effective resolution through aberrations; the smallest resolvable feature size was measured before shipping as 22 μm. With the lens-system shown in Chapter 4, a magnification of 1.73 ± 0.05 is achieved. Scale bars in the images are based on the actual pixel size and take into account the magnification factor.

The focus error was estimated using a ray-tracing program to model the imaging system. A pessimistic adjustment-precision of the final lens in the imaging system of 18 mm was assumed. If the lens were moved by that amount it would result in the focal point of the imaging system moving by 60 μm. The image was calculated to occupy an area with a diameter of 14 mm on the final lens. Given a distance from the lens to the ICCD camera of 80 mm (see Fig. 4.13 in Chapter 4), this gives a focus error of,

$$\frac{14\,\mathrm{mm}}{80\,\mathrm{mm}} \times 60\,\mu\mathrm{m} = 10.5\,\mu\mathrm{m}\ . \tag{8.2}$$

The blue laser spot-size (w_0) is estimated to be roughly 50 μm at the trap centre.

Images of a cloud before and after axialisation are shown in Fig. 8.3. The scale of the reduction of the cloud width depends on several factors:

a. Laser-cooling rate, itself dependent on blue laser detunings and red laser frequencies along with the beam position

b. Axialisation frequency

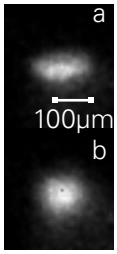

Figure 8.3: The effect of turning on the axialisation drive. The cloud is unaxialised in *a*. An immediate reduction in the width of the cloud is observed *b*.

c. Axialisation-drive amplitude

d. Cloud size

Unfortunately, most of these, with the exceptions of (b) and (c) are hard to quantify exactly. However, we have made some attempts to demonstrate qualitatively the influence of these factors.

8.2.1 Laser-cooling rate

For beam positions close to the centre of the trap, the magnetron motion can be heated easily if the cooling rate drops due to a laser being detuned too far from resonance. This is shown for an axialised cloud in Fig. 8.4 where the only parameter changed between the two shots is the frequency offset of one of the blue lasers. Were the cloud unaxialised the ions would have been heated by having the beam so low in the trap and would not be visible at all in Fig. 8.4 *b*.

8.2.2 Axialisation-drive amplitude

For the images in Fig. 8.5 the trap bias and laser parameters were kept constant while the axialisation amplitude was changed. Two points appear in the graph at a drive amplitude of 250 mV: a second image was taken at this drive voltage after all the other points had been recorded. The consistency of the cloud size between these two points suggests that there really was a lower limit of the cloud radius that could be achieved. The aspect ratios were calculated from full-width-half-maxima measured on a single row or column of pixels in the image. A more detailed analysis would require the integration of the data in horizontal and vertical directions.

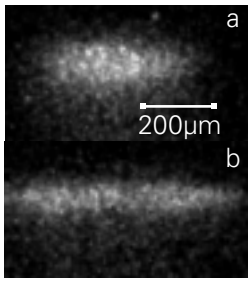

Figure 8.4: The effect of laser-cooling rate on cloud shape. In *a* the laser-cooling rate is high. One of the blue lasers has been detuned in *b* to show how the heating of the magnetron motion by one of the lasers is not entirely compensated for by the axialisation drive.

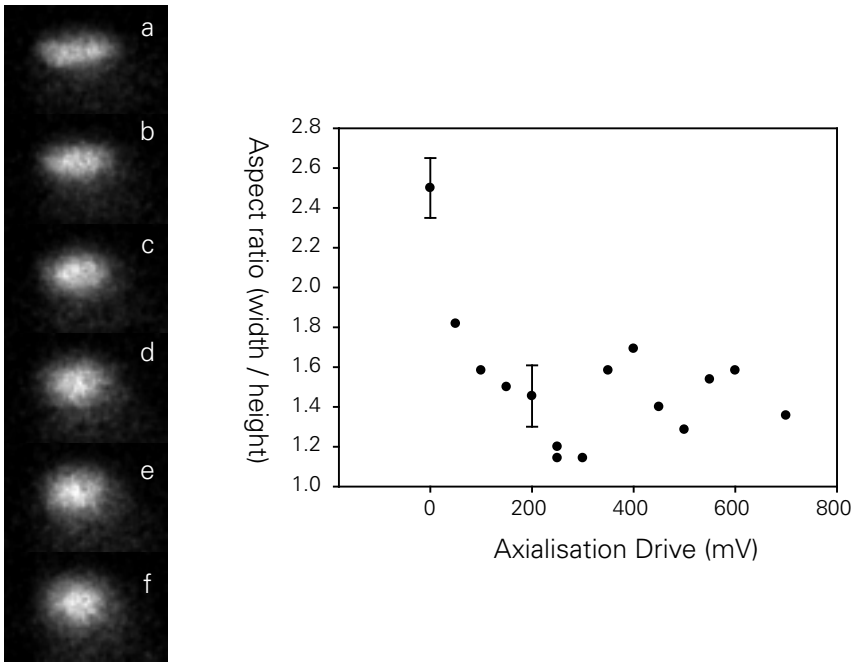

Figure 8.5: The effect of increasing the axialisation-drive amplitude on the width of the cloud. *a* is an unaxialised cloud while *b* through *f* are axialised with a drive amplitude of 50 mV and increasing in 50 mV steps. The aspect ratios for images *a* through *f* and images taken with higher drive amplitudes are shown in the accompanying graph, The ratios were crudely calculated but emphasise the fact that the cloud size shrinks in the radial direction with even a small drive but that there appears to be a limit to how far it can be reduced. The representative error bars are calculated from the focus error as caculated in Equation 8.2.

8.2.3 Cloud size

There are insufficient data to draw strong conclusions about the effect of ion number on the ultimate radius of an axialised cloud but clearly space charge will ultimately impose a limit on how small the cloud can be. Experience suggests that with smaller clouds a larger axialisation drive can be applied before a deterioration in signal level is observed. This is because small clouds tend to be located closer to the centre of the trap and therefore spend more time within the laser beam. The effect of space charge on the ultimate cloud radius was also observed in experiments with magnesium ions [105].

A second but less important factor is that the azimuthal field deviates further from the ideal quadrupole field away from the trap centre. It is possible that larger clouds, located further from the trap centre, will experience some degree of RF heating due to the wrong geometry of the field.

As the experiment progresses towards the trapping of single ions in future, more data should become available on the effect of cloud size.

8.2.4 Trap bias

As would be expected, a reduction in the trap bias causes an increased confinement in the radial plane as the ions are less tightly confined along the axis. This is clearly demonstrated in the series of images in Fig. 8.6, taken with a relatively weak axialisation drive of 40 mV amplitude. The blue laser detunings were adjusted between shots in order to re-optimise the cooling and reduce the cloud size.

8.2.5 Observation of motional excitation on the MCS

One of the most unusual results is a modulation of the ion fluorescence at frequencies a little less than 1 Hz when certain combinations of red- and blue-laser detunings are in place. Similar behaviour had previously been observed in other ion species without axialisation by the group manifested as a photon count-rate oscillation measured using a PMT. Oscillations of the fluorescence level were also observed by Powell [105] in her experiments with axialised magnesium ions.

On this occasion, there appeared to be a dependence of the frequency and depth of the modulation on the amplitude of the axialisation drive. Previous

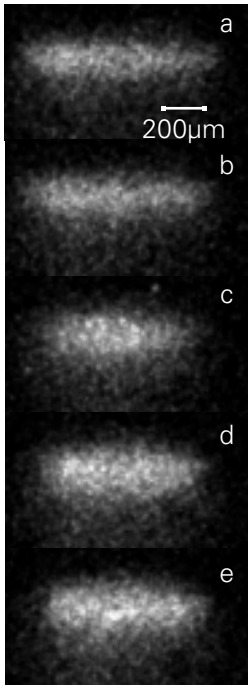

Figure 8.6: The effect of reducing the trap bias on the radius of the cloud. The bias starts at 10 V in *a* and decreases in 1 V decrements to 6 V in *e*. Weaker confinement in the axial direction allows the ions to be "squashed" in towards the axis in the radial plane by both the magnetic field and the axialisation drive.

observations of such a modulation with axialised ions suggested that these parameters of the modulation were dependent on the frequency of the drive. The sensitivity of the phenomenon to laser frequencies meant that it was difficult to set up and maintain the modulation for long enough to acquire data to quantify any dependence in these experiments with calcium ions, however. Figure 8.7 is a sample MCS trace showing the modulation. The modulation depth was, on occasion, as large as 80–100%.

Despite observing such a phenomenon before, this is the first time we have able to image the cloud during these oscillations of fluorescence intensity. Previous images of magnesium ions have shown a superposition of two cloud states [105]. Figure 8.8 is a series of frames taken by the camera during such oscillations. Although the oscillations are not as clear when shown as stills, the images are convincing evidence that the entire cloud changes shape during the fluorescence modulation: there is a definite "breathing" in and out in the radial

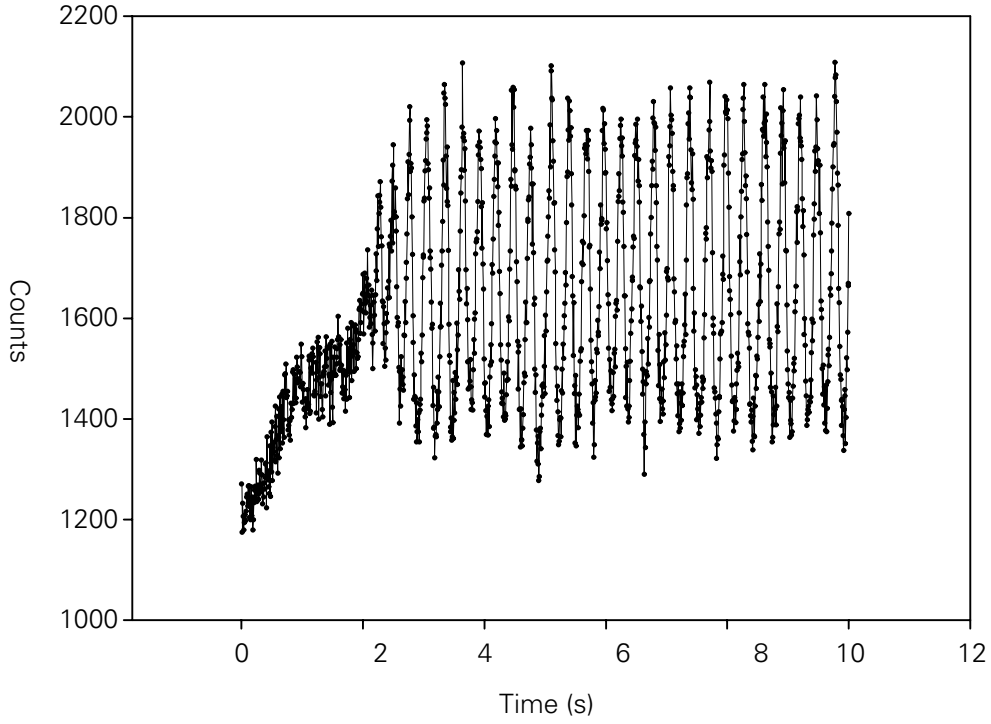

Figure 8.7: Modulation of the fluorescence with specific laser detuning combinations. There appeared to be a dependence of the frequency and the depth of the modulation on the amplitude of the axialisation drive. The modulation occurs because of an iris in front of the PMT which prevents it from imaging the entire trap volume. Expansion of the cloud outside of the volume imaged will reduce the count rate measured from the PMT while compression into that volume will increase the count rate.

plane. Each exposure was 50 ms in length and the frames were taken at a rate of 6.9 Hz. The graduated scale has been changed from the normal range (of zero counts to maximum counts) to one which emphasises the change in cloud size, which is more apparent in moving images than in stills. The full width at half maximum (FWHM) for the integrated images was calculated for integration in both vertical and horizontal directions. The aspect ratio of the cloud was calculated from these FWHM data and is displayed on the right-hand side. This is intended to guide the eye rather than to quantitatively indicate the actual cloud dynamics. The modulation depth of the fluorescence could not be measured simultaneously with the acquisition of the images for technical reasons. It is likely, however, that the depth was lower during the acquisition of the images

in Fig. 8.8 than in Fig. 8.7.

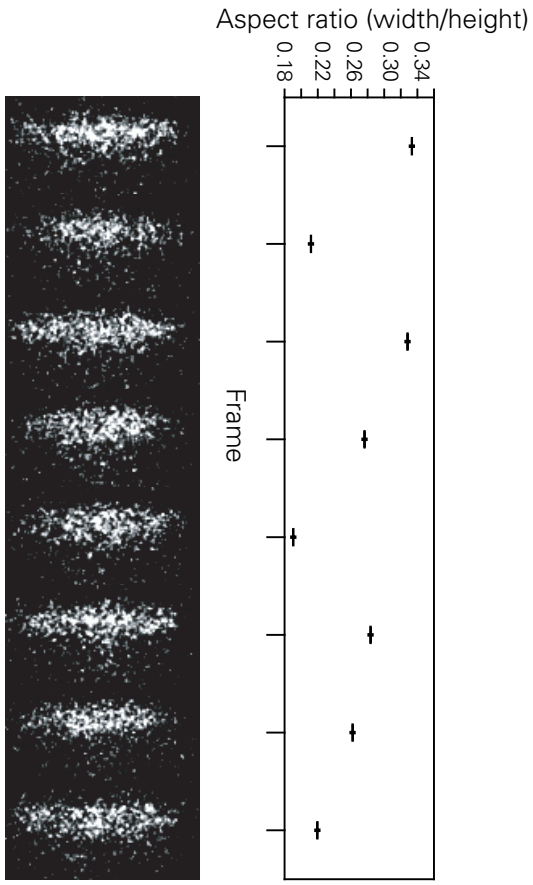

Figure 8.8: Successive frames taken in a “kinetic series” showing a dynamic change in cloud shape. The aspect ratio on the right is intended to guide the eye rather than to indicate the motional frequency.

8.3 Quantitative results

The primary method for measuring the effect of the axialisation in a more quantitative way was the technique of RF–photon correlation, as with magnesium ions. As described in Chapter 6, this allows the measurement of the damping rate of the motion of the ions. The main difference between experiments described here and those performed with magnesium is that the motional excitation is performed with a dipole signal across the ring electrode rather than with an oven acting as an antenna. The cloud should be much more sensitive to strict dipolar excitation.

The amplitude of the excitation signal was typically between a few millivolts and a few tenths of millivolts, with the smaller values required when the excitation frequency was close to the motional frequency. The required amplitude was determined empirically so as to give a sufficiently large modulation depth on the MCA trace without exciting the ions out of the trap.

8.3.1 Magnetron-motion damping rate

When the laser-cooling rate is high for the modified cyclotron motion and the axialisation drive is on, one expects an increased magnetron damping rate compared with the unaxialised case. This was certainly borne out by the experimental evidence but many subtleties arose when a detailed study was made.

The unaxialised magnetron damping rate was measured with several different clouds. It was typically of order 3 Hz but was measured on one occasion as being 30 Hz (Fig. 8.9). The rarity of such a high damping rate shows the difficulty in achieving the laser parameters necessary for good laser-cooling of this motion without axialisation.

The damping rate was typically higher with the axialisation drive on, with damping rates of up to 60 Hz seen. There appears to be a roughly linear increase in the damping rate with the amplitude of the axialisation drive (Fig. 8.10). It should be emphasised that the error bars represent only the goodness of fit in the phase plot. It might be expected that error bars indicating the possible extremes of the cooling rate would be imposed, determined by measurable systematic effects such as laser-frequency drift. Estimating such extremes is very difficult, however. In particular the upper limit of the cooling rate can usually be revised upwards by ever-more careful adjustment of laser frequencies and

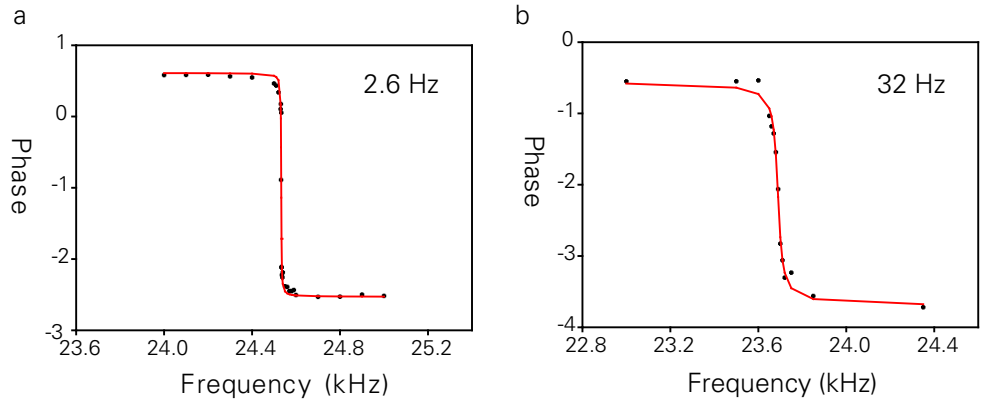

Figure 8.9: Damping rate for the magnetron motion when unaxialised. *a* shows a typcial damping rate, with the much high rate of 32 Hz (*b*) being observed on only one occasion.

beam position. In the absence of a reliable measure of such an error estimate the goodness-of-fit errors have been used.

The error bars generally increase with the damping rate. The reason for this is that when the cooling rate is high the cooling laser frequencies are close to one linewidth away from resonance. The line-shape of the cooling transition is steepest at this point and a small change in laser frequency will cause a large change in the cooling rate. Such changes during the acquisition of data, which typically took half an hour, will give rise to a large uncertainty in the cooling rate. This point demonstrates the need for an improved laser lock in future experiments. Suggestions for improvements to the locking set-up are discussed in Section 8.5.1.

Many measurements of the magnetron damping rate were made with different clouds of ions on different days. Two examples of such data sets are shown in Figs. 8.10 and 8.11. Although every effort was made to keep the laser-cooling rate high by optimising the signal with the laser frequencies between each data point, it was still impossible to know if the maximum cooling rate was achieved. This makes it very difficult to interpret some data sets such as that shown in Fig. 8.11. It appears to show that there is no significant increase in damping rate with axialisation drive. This could be due to poor laser-frequency or beam-position optimisation at higher drive amplitudes, however. All of these rates should therefore be interpreted as lower limits.

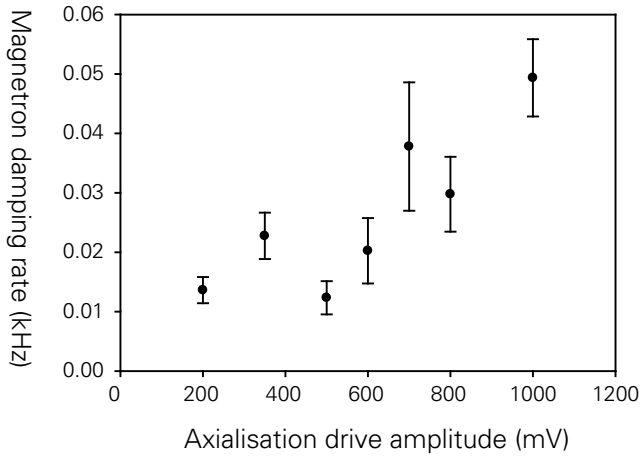

Figure 8.10: Variation of damping rate with axialisation drive. In general, the rate increases as the drive is made stronger. The error bars are derived from the fitting process in the phase plot.

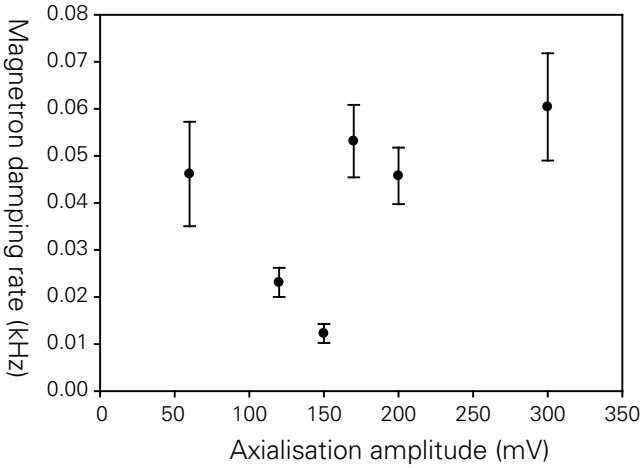

Figure 8.11: A relatively small variation of the magnetron damping rate with axialisation-drive amplitude. The rate at 50 mV is notable for being so large.

It was observed in most data sets that the damping rate increased with the axialisation-drive amplitude and I believe that counter-examples such Fig. 8.11 are the result of sub-optimal laser-cooling at higher drive amplitudes. This point can only be resolved when better laser stability has been achieved and greater reproducibility of damping rates has been demonstrated.

8.3.2 Modified-cyclotron-motion damping rate

It was expected that there would be little effect on the modified-cyclotron damping rate by the axialisation drive. The increase in the magnetron damping rate with axialisation is a small fraction of the typical unaxialised damping rate of the modified cyclotron motion and the coupling of this damping to the magnetron motion should have little effect on the modified-cyclotron damping rate. Indeed, it might be expected that the rate would rise as the ions were better localised and spent more time in the laser beam. To investigate this, measurements of the damping rates of both motions were made alternately by R. Hendricks.

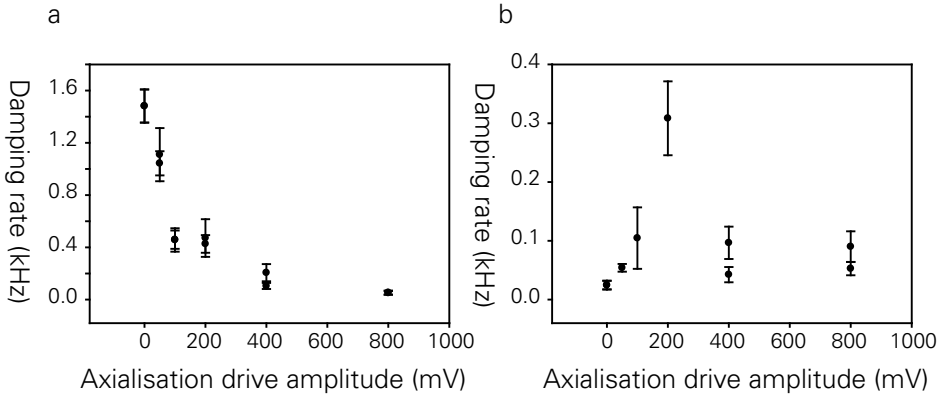

Figure 8.12: The damping rate of *a* the modified cyclotron and *b* the magnetron motions. The anomalously high magnetron damping at 200 mV drive is likely to be exaggerated by the small number of points available for the fit.

Intriguingly, these preliminary data show that the damping rate for the modified cyclotron motion drops rapidly as the axialisation-drive amplitude is increased (Fig. 8.12). It is believed that this is due to RF-heating of the modified cyclotron motion. This heating would be an off-resonant excitation of the motion by the axialising field at the true cyclotron frequency. The separation of the true cyclotron frequency from the modified cyclotron frequency is typically

small, being the magnetron frequency. Such off-resonant heating of the modified cyclotron motion would be greater in the experiments with calcium ions than with magnesium since the magnetron motional frequency was smaller. This theory will be tested in future by examining the effect on the modified-cyclotron damping rate of a quadrupole drive off-resonant with the true cyclotron frequency (i.e. in the absence of axialisation).

Further hints to the dynamics behind the axialisation process are provided by Fig. 8.12. It is seen that there is a maximum damping rate of the magnetron motion near 200 mV which then falls off again. The damping rate fitted from the phase change plot at this drive amplitude is likely to be rather higher than it ought to be due to the small number of points available for the fit. Nevertheless, in this and in previous experiments it was found that the axialisation seemed to be most effective at a small (∼50–200 mV) drive amplitude. It is likely that Fig. 8.12 holds the key to this behaviour. Initially, the axialisation drive provides some damping to the magnetron motion through the coupling. For low drive amplitudes the modified-cyclotron damping rate remains high and coupling this damping to the magnetron motion effectively damps that motion too. As the drive amplitude increases, the RF heating of the modified cyclotron motion increases and its damping rate falls rapidly. Although the coupling between the motions may be stronger, the damping rate of the modified-cyclotron motion has fallen so far that at a certain stage no further benefit to the magnetron damping rate can be obtained from a larger coupling rate between the two motions. The key to axialisation appears to be achieving a careful balance between strong coupling between the two motions and low heating of the modified cyclotron motion.

8.3.3 Shifts in motional frequencies with axialisation

It was observed by Powell [105] that there was no perceptible shift of the magnetron frequency with the amplitude of the axialisation drive in her discussion of stable orbits (Chapter 6). She sought to determine whether ions in a stable orbit under certain conditions of axialisation and laser-cooling were distributed around a ring or moved as a ball. If they were distributed around a ring, space charge effects would cause an increase of the rotational frequency. She compared frequencies measured using a Fourier transform of photon–photon correlation data to the theoretical prediction based on the trap parameters. The

frequency increase she sought would have changed the rotational frequency from the normal magnetron frequency by around a factor of two and no such shift was observed. However, in the work with Ca^+ a more subtle shift has been consistently observed. We believe that this effect is due to the motional coupling in a manner similar to that described by other authors [102, 103] and which will be discussed in detail later in this section.

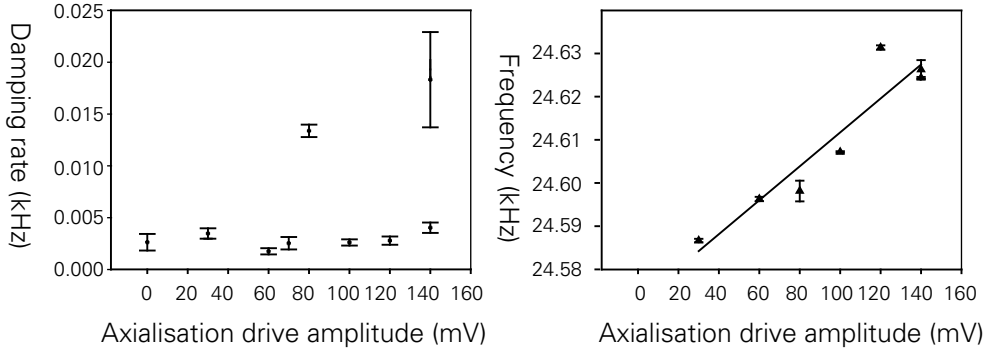

Figure 8.13: The change in the magnetron damping rate and the magnetron frequency with axialisation-drive amplitude. Although the damping rate is generally low and there is significant variation, the frequency increases linearly with the drive amplitude. In particular, two data sets were taken at 140 mV drive amplitude, one with a low and one with a high laser-cooling rate. The damping rates varied dramatically while the two motional frequencies were almost identical. Error bars for both frequencies and damping rates are derived from the fitting process and the linear regression for the frequencies is plotted.

The frequency shift is typically of order a few kilohertz with an axialisation drive of 1 V. It is found to be linear with the axialisation drive and to have the same amplitude and sign for both motional frequencies. An important point to make is that it was clear from the outset that the shift was independent of the damping rate and depended only on the amplitude of the drive. This is shown in Fig. 8.13 where both damping rates (dots) and frequency shifts (triangles) have been plotted. Although the damping rate varies considerably and with no clear trend, there is a clear linear increase in the magnetron frequency. Two measurements of the damping rate were made at an axialisation drive of 140 mV, one with a higher fluorescence rate and consequently higher laser-cooling rate than the other. The measured damping rates of the magnetron motion differ by 15 Hz while the magnetron frequency is almost unchanged (about 2 Hz differ-

ence between the two data sets). To ensure that this shift was not due to slowly varying experimental parameters, another set of data was taken over a shorter time-scale. The results are shown in Fig. 8.14. The linearity of the frequency shift is better because drifts of the magnetic and electric fields were smaller.

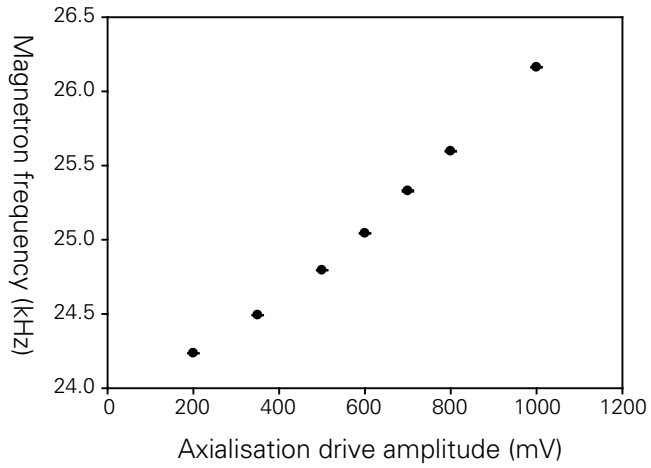

Figure 8.14: In a collection of data sets taken more rapidly, the frequency shift is more linear as other effects such as magnetic field drift are reduced.

In order to conclude that these shifts are due to a motional coupling similar to that described in Refs. [102, 103] it was first necessary to discount the possibility that they had been caused by one of the following:

1. A DC offset is present in the axialisation drive, changing the ring electrode from ground potential and thus affecting the overall trap bias.

2. Drift of the magnetic field could appear to give a change related to another parameter varied slowly with time (in this case the axialisation-drive amplitude).

3. The trap can be thought of as a linear combined trap with the ring segments treated as the rods in a normal RF trap. The addition of an RF voltage will cause frequency shifts that vary with both the applied frequency and amplitude.

The first possibility, a DC offset, would cause changes to all three motional

frequencies. The axial frequency changes with the square root of the trap bias,

$$\omega_z = \sqrt{\frac{4eU}{m(2z_0^2 + R_0^2)}} \ , \tag{8.3}$$

and the magnetron and modified cyclotron frequencies change in opposite directions as this changes such that they always add to form the true cyclotron frequency,

$$\omega_c' = \frac{1}{2}\left(\omega_c + \sqrt{\omega_c^2 - 2\omega_z^2}\right) \ , \tag{8.4}$$

$$\omega_m = \frac{1}{2}\left(\omega_c - \sqrt{\omega_c^2 - 2\omega_z^2}\right) \ . \tag{8.5}$$

Thus, a change in trap bias could not explain frequency shifts of the motions that go in the same direction. Measurements on the circuit described in Section 8.1 also indicated that any DC offset applied to the ring was very small indeed.

The variation of magnetic field (point 2) has already been seen to be very small (less than one part in 10^4). Care was taken, too, to take readings at high and low axialisation amplitudes alternately to eliminate any slow drift of parameters from the results.

Finally, to address point 3, the equations of motion in a linear combined trap were examined. Huang [115] gives the frequencies of the equivalents of the modified cyclotron and magnetron motions in such a trap as,

$$\omega_\pm = \frac{qB_0}{2m}\left[1 \pm \sqrt{1 + \frac{4m}{qB_0^2}\left(\frac{qV_{\mathrm{RF}}^2}{2m\Omega^2 R_0^4} - \frac{U}{a^2}\right)}\right] \ , \tag{8.6}$$

where V_{RF} is the amplitude of the RF voltage, and Ω is the frequency. It can thus be seen that a quadratic relation between the motional frequencies and the amplitude of the RF exists. The results obtained for the frequency shifts all show a linear relationship. Again, the two frequencies change in opposite directions such that they still add to the true cyclotron frequency, in contrast to the experimental results. Finally, but most importantly, the frequency shift was found to be a resonant effect: detuning the axialisation drive far away from resonance with the true cyclotron frequency resulted in no shift of the motional frequencies.

A literature survey was conducted to see if there were any other reports of shifts of this kind. To our knowledge there are only two references to similar frequency shifts in the literature [102, 103]. Although the motions of interest, and consequently the driving frequencies, were different in these papers the fundamental principle should remain the same.

In the work of Cornell *et al.* [102], the modified cyclotron motion was coupled to the axial motion. The axial motion could then be detected and the modified cyclotron frequency calculated from the phase of the axial motion. This method required the use of π pulses of the coupling field rather than a continous coupling. These π pulses converted all of the modified cyclotron motion into axial motion.

Häffner *et al.* [103] coupled the magnetron motion to the axial motion. Their work is more similar to ours in that the coupling was done for the purposes of damping the magnetron motion (the axial motion was resistively damped) and the coupling was continuously on rather than pulsed. In both cases the coupling field was applied with a non-ideal geometry since the required xz-geometry is difficult to achieve (see Chapter 6).

The shift arises from the coupling of the two harmonic oscillators. The relevant equations describing this system are set out in Chapter 6. Analysis of the system shows that the motional frequencies of the coupled oscillators are shifted by an amount dependent on both the axialisation drive amplitude and frequency. Cornell *et al.* describe the result as a "classical avoided crossing" and draw the analogy with a dressed-state system: one more familiar to many in the field of quantum optics than the classical system. The offset of the frequency of the motional resonance from the uncoupled frequency is given by,

$$\epsilon = \frac{\delta}{2} \pm 1/2 \left(\delta^2 + |k|^2\right)^{1/2} , \tag{8.7}$$

where δ is the detuning of the axialisation drive from resonance and k is the coupling rate of the two motions (proportional to the axialisation-drive amplitude). For small detunings, δ, the frequency shift ϵ is directly proportional to the axialisation-drive amplitude, as observed in the experiment.

An important consequence of these shifts and the avoided crossing is that for axialisation-drive frequencies close to resonance with the true cyclotron frequency, two motional frequencies may be measured. As was seen in Chapter 6 one of the two peaks will be rather weaker than the other and with the axi-

alisation drive detuned too far from resonance the second peak will not be observed at all. Although two phase transitions very close to one another had been observed in early experiments, these were interpreted at the time as motional frequency changes due to magnetic-field drift during the timescale of the measurement. When it was understood that these shifts were probably described by Equation 8.7 it was then possible to look at the behaviour of the shift with the axialisation drive parameters.

8.3.4 Coupling rate of the motions

In order to verify that Equation 8.7 really did describe our system accurately, we sought to measure the predicted avoided crossing. We changed the axialisation-drive frequency in small steps, going from below to above resonance. For each frequency, we measured the motional frequency using the RF–photon correlation method. The results are shown in Fig. 8.15 and are in excellent agreement with theory. They indicate a coupling rate of about 900 Hz for the applied drive of 250 mV. The frequency error bars in Fig. 8.15 and the following figures in this section are calculated from the accuracy of the function generator which was used (Thurlby Thandar TG550).

The important difference between the result presented here and those obtained by Cornell *et al.* and by Häffner *et al.* is the method used to measure the motional frequencies. Here, they were measured using a phase technique while in Refs. [102] and [103] they were measured using a Fourier-transform technique. It is also of note that the coupling rate achieved in our experiments was nearly two orders of magnitude bigger than that achieved in the experiments coupling a radial motion to the axial motion. This is, to a large degree, due to the more accurate geometry used to apply the coupling drive in our experiments; the xz-geometry required for coupling to the axial motion is difficult to achieve experimentally. Another factor is that the coupling of motions in the other two experiments was applied in order to measure a motional frequency rather than to damp one of them. Such a high coupling rate was therefore not necessary.

This result may help to explain why the system was found to be unstable when the drive frequency was on resonance. This phenomenon was previously observed with magnesium ions [105] and was also found in these experiments with calcium. Despite the apparent instability of the system, the measured

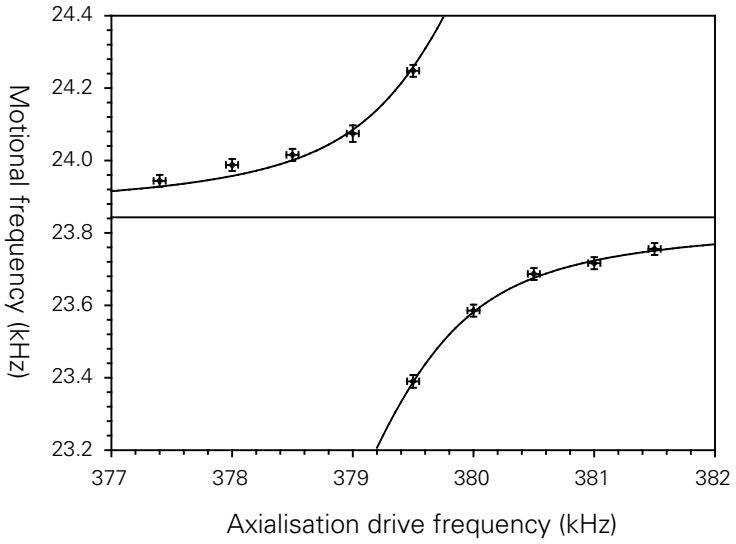

Figure 8.15: The classical avoided crossing. The axialisation-drive frequency was changed in steps and the motional frequencies measured using the RF–photon correlation technique. Near resonance with the true cyclotron frequency, two motional frequencies are clearly visible. The fitted curve indicated a true cyclotron frequency of 379.54 kHz and a coupling rate of 870 Hz.

damping rate was found to be high, close to resonance (Fig. 8.16).

A series of phase measurements with a high resolution of excitation frequency was made, while keeping both the axialisation-drive frequency and amplitude constant (Fig. 8.17). The shape is very similar to that predicted by the simple model of coupled classical oscillators in Chapter 6.

Finally, in order to confirm that the same frequency shift applied to the modes associated with both radial motions, RF–photon correlation measurements close to the two motional frequencies were made by R. Hendricks. The axialisation-drive frequency was kept constant while the drive amplitude was increased. Equation 8.7 predicts that the separation of the two peaks near each motional frequency will increase with drive amplitude (i.e. as k increases). The results confirm the shift of frequency is the same for each of the two motions (Fig. 8.18).

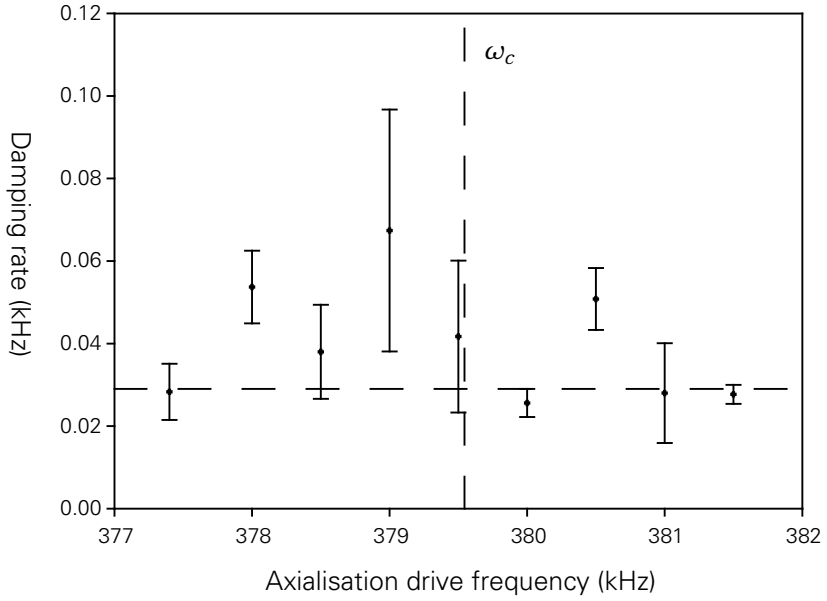

Figure 8.16: The variation of the magnetron damping rate with the axialisation-drive frequency. The vertical reference shows the true cyclotron frequency as determined from the fit in Fig. 8.15. The horizontal reference is the maximum unaxialised damping rate achieved.

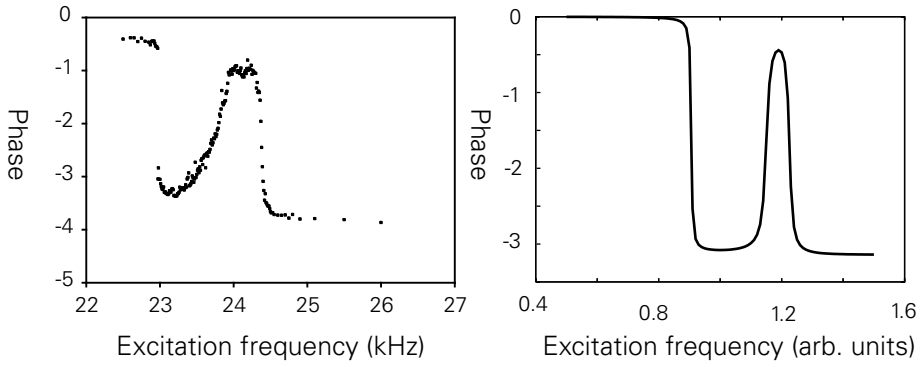

Figure 8.17: Phase measurements with a fixed axialisation-drive frequency and amplitude. This matches theoretical phase plots for the simple model of coupled oscillators well (reproduced from Fig. 6.13 in Chapter 6).

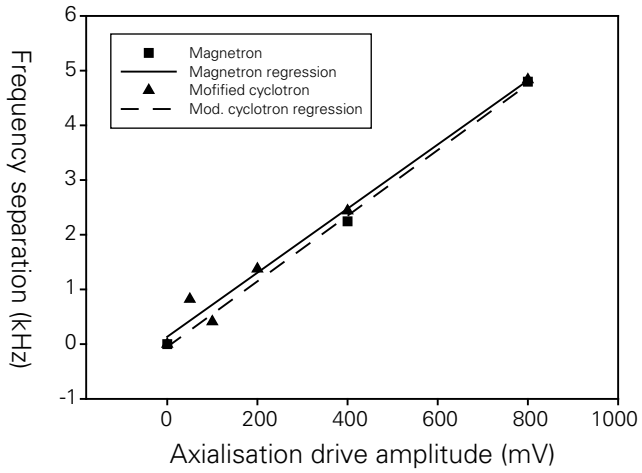

Figure 8.18: The separation of the frequencies of the two motional modes measured from the phase changes in RF–photon correlation experiments. The pair of modes near the magnetron frequency have almost the same separation as those near the modified cyclotron frequency, as expected.

8.4 Loading with axialisation

One of the desired outcomes of using the axialisation technique is that it should become easier to load small clouds of ions in the Penning trap. In some early experiments all of the lasers were left static in frequency and close to the trap centre while loading. After a couple of minutes a small cloud of ions suddenly appeared in the trap (as imaged with the ICCD). This was certainly a great improvement over the unaxialised case where it is necessary to scan one of the cooling lasers in frequency or spatially in order to cool the ion cloud and begin to see fluorescence.

It was found that allowing full laser power into the trap during the loading process resulted in the immediate appearance of ions on the ICCD image. The laser power can then be reduced with neutral density filters or a combination of a wave-plate and a polarising beam-splitter (Chapter 4). This process may not be effective for very small clouds, however, where the initial oversaturation of the cooling transition may prevent the ion being cooled.

A new technique has been developed recently by A. Abdulla and R. Hendricks where the penultimate mirror before the trap is mounted with actuators (Newport NSA12). These actuators can be computer-controlled and allow a sweep of the beams in toward the centre of the trap. This appears to be a convenient and reproducible way of performing the initial cooling on a cloud of ions after loading. It has the advantage, too, that it works at a low laser power which should be advantageous in loading single ions.

More work is necessary before conclusive statements can be made about the ease of loading in the presence of an axialising drive. Preliminary indications are that it has become a more reliable process.

8.5 Discussion

The axialisation of calcium ions has been conclusively demonstrated. The method has been characterised using information gleaned from RF–photon correlations. This has allowed the damping rate of both radial motions to be determined. It has also allowed a measurement of the coupling rate of the motions.

Lower damping rates of both radial motions were achieved than with magnesium ions. This is in large part due to the width of the cooling transition and the presence of other levels amongst which the population can be spread. It is

clear from many of the results presented here, however, that a severe limitation on both the characterisation of the technique and of its implementation is the stability of the laser systems.

It is expected from theory and shown in both images of the ions and the damping rates measured that the lasers need to be tuned close to resonance with the relevant transitions to achieve axialisation, i.e. the closer confinement of the ions around the axis rather than just the coupling of the two motions. With significant drift rates, and most likely higher-frequency laser-frequency noise, it is impossible to achieve consistently high damping rates. This problem was not so apparent with magnesium ions where only one cooling laser was required.

Laser instability has partially obscured what is probably a more fundamental and less-easily tackled problem in attempting to trap single ions of calcium. It has been observed that the count rate per ion is very much lower for the Penning trap than for the Paul trap. The trap used for the experiments described in this thesis has been operated in both configurations, and a count rate per ion of 6 kHz was observed in the Paul trap and less than 1 kHz in the Penning trap. In another trap used in the group, experiments with calcium ions indicated a count rate per ion of 2 kHz in the Paul trap and 0.6 kHz in the Penning trap [65].

Although the instability of the lasers would certainly make it difficult to achieve the maximum count rate, and especially so in the Penning trap where more lasers are required to be kept on-resonance, it is highly unlikely that this is the only factor contributing to the lower count rate in the Penning trap. A possible explanation is the existence of "trapped states" [116]. Trapped states are closely related to electromagnetically induced transparency. A pair of lasers coupling a three-level Λ system can create destructive interference on one of the transitions for certain detunings of the lasers. This will suppress this transition and reduce the observed fluorescence. In the context of laser-cooling it will also have the effect of reducing the damping rate.

8.5.1 Future directions

The fundamental limitation discussed previously of the laser-frequency instability will need to be addressed before single ions can be routinely cooled and axialised. Firstly the long-term drift will need to be removed by better temperature-stabilising the locking cavities. Experiments without any temperature stabilisation of the cavities have shown that they have a significant drift (of

order 100 MHz per hour). Without any lock in place at all, the inherent drift of the lasers of about 300 MHz per hour makes many experiments very difficult indeed. Recent measurements of the drift measured with the wavemeter indicate that the drift of the temperature-stabilised cavities is between 15 and 75 MHz per hour. The drift was found to be roughly linear in time. It is thought that the reason for this poor stability is the use of resistors with relatively high thermal coefficients in the Wheatstone bridge used to measure the cavity temperature in the locking circuit. Steps are already being taken to address this.

The removal of serious long-term drift from the laser frequency is likely to reveal frequency instability on a shorter time-scale. Firstly, there is no decoupling of intensity noise in the laser lock. Thus any intensity fluctuations will cause frequency fluctuations as the servo seeks to compensate for the intensity change. This can be addressed at a very basic level by introducing a second photodiode into the locking scheme which monitors the intensity of the beam. Secondly, the finesse of the cavities is very low (~ 60). This finesse was chosen in the design phase since the primary function of the cavities was to remove large-scale frequency drift over a long time period. The slope is not high enough, however, to remove higher-frequency noise. It remains to be seen if this will prove a serious obstacle in laser-cooling single ions of calcium. Higher-finesse cavities and perhaps a different locking scheme such as the Pound–Drever–Hall scheme like that being built for the Ti:Sapphire laser might be considered in future and would allow the removal of higher-frequency noise.

It will also be necessary to assess the degree to which trapped states are likely to be reducing the signal level and address this problem if it exists. Trapped states are certainly much more likely to be present with the large number of lasers needed for cooling in the Penning trap compared with the Paul trap. The method of generation of the repumping light, using pairs of sidebands generated from just two 866-nm diode lasers means that the level of the coherence in the system will be very high. This coherence will also contribute to probability of trapped states existing. Thus, one possible method of increasing the fluorescence level is to reduce the coherence of the light. This could be done by the addition of a light-emitting diode (LED) near the correct wavelength, although care would need to be taken that it did not excite the 850-nm transition too strongly and cause shelving in the $D_{5/2}$ level. The coherence could also be reduced by using higher-bandwidth lasers on the repumping transitions. Reducing the coherence on the cooling transition, by contrast, would be undesirable

as it would reduce the minimum temperature achievable.

An open question regarding the dynamics of the axialisation itself is why the system is less stable when the drive frequency is greater than the true cyclotron frequency. In principle, the only difference between being above or below resonance is that the motional frequencies will be shifted downwards instead of upwards. This in itself does not appear to explain the instability.

In general, we found the magnetic field provided by the electromagnet to be stable (to approximately one part in 10^4). Evidence for the stability of the field came in the form of the stability of the optimal axialisation frequency over a long period of time. Further evidence comes from the fact that when tuning the cooling lasers for maximum fluorescence it was normally found to be necessary to tune them in the same direction by approximately the same amount. Had the drift been in the magnetic field, the Zeeman splitting would have separated the cooling transitions ($^2S_{1/2}$–$^2P_{1/2}$) further and the lasers would have had to have been tuned in opposite directions. Instead, the drift was a temperature-related drift of the locking cavities.

Although we found the magnetic field to be relatively stable, it would be desirable to have an even more stable magnet. As well as providing an even more stable cyclotron frequency, the extra stability will become particularly important when the sidebands of the 729-nm transition are probed. Very small magnetic-field fluctuations would change the motional frequencies and thus change the sideband frequency. For these reasons a trap is being modified for trapping calcium ions in the bore of a superconducting magnet.

The results presented in this chapter will be submitted for publication in a journal [2]. They also form the subject of a poster which will be presented at a major international conference (Physics 2005, April 2005, Warwick, UK).

Chapter 9

Cavity QED experiments with trapped ions

9.1 Introduction

Quantum Electrodynamics (QED) has proved to be the most precise theory developed in physics to date. Many of its predictions have been tested experimentally and agreement between experiment and theory found to a very high precision.

The theory concerns the interaction of atoms with their electromagnetic environment. In free space this environment is treated as an infinite number of quantised modes of vibration of the electromagnetic field.

The subset of QED theory describing the interaction of an atom with a cavity formed by two mirrors has been the focus of a great deal of study for some time. The cavity restricts the number of modes which must be considered and greatly simplifies calculations. It is only in recent years, that the experimental conditions to realise the requirements imposed by theoretical study have been achieved. The primary condition is the manufacture and precise alignment of very highly-reflective mirrors. This was first achieved in the microwave regime where superconducting materials provided the very high reflectivity required. New dielectric coatings have allowed the manufacture of mirrors in the optical regime with similarly high reflectivity. The secondary condition, of ensuring a long interaction time between the atom and the cavity, has long been achievable with trapped ions, which can be held stationary for some time. Recently, sufficiently long interaction times have also been achieved with neutral atoms using

a variety of techniques including dropping atoms with a low velocity through a cavity from a trap [117] and creating dipole traps within the cavity [118].

The research programme of the group of Professor Rainer Blatt in Innsbruck, Austria includes the practical realisation of a trapped ion–cavity system [78, 119, 120]. This is part of a larger project with the goal of building an ion-trap quantum computer with a linear RF trap and calcium ions [39, 57]. In this respect their goals are very similar to those of the group in Imperial College. The complete refurbishment of our laboratories in Imperial College was scheduled to take place during the summer of 2003. The opportunity therefore arose for me to spend some months in the group in Innsbruck while this refurbishment work took place. While there I worked on a laser-locking scheme and assisted in building a cavity to be used for Cavity QED (CQED) experiments.

The cavity being built will form an integral part of a source of single photons on demand, a very desirable commodity in tests of quantum mechanics and for quantum cryptographic protocols. In the tests of quantum mechanics, such a source will allow loophole-free tests of Bell's inequality [121]. Single-photon sources in quantum cryptographic systems will allow higher data transfer rates using an entirely secure protocol [122].

9.2 Single-photon sources

Sources of single photons on-demand are required for some proposals for quantum cryptography and entanglement experiments [123, 31]. They are also needed for quantum computing proposals involving linear optical components [124]. Early experiments [125] producing single photons relied on cascade spontaneous-emission events or on parametric down-conversion. In both cases the probability of single-photon generation is low but success is indicated by a 'heralding' photon e.g. one of the two emitted photons in a down-conversion process indicates the other one has also been emitted simultaneously. These experiments suffered from the inability to control the time of emission. Later experiments employing quantum dots have allowed control over the time of emission but not over its direction. Work on quantum dots has progressed significantly over the years and shows some promise.

Another approach is to use an atom strongly coupled to a cavity. This has the advantage that the direction of the emitted photon is well-defined as there is a strong preference for emission into a cavity mode. Control of the cavity

length and of external laser parameters provides excellent control over the time of emission (essentially limited by the cavity life-time). The significant level of existing expertise with the laser systems required for such a set-up, and the interest in the basic physics associated with atom–cavity systems means that this is a very attractive method to investigate.

9.3 Quantum communication and state transfer

In a more complex variation of the atom–cavity experiment, it is proposed by Cirac *et al.* [126] and by Law and Kimble [127] to use the atom–cavity coupling to transfer internal electronic states of a trapped ion or atom to another ion or atom via a photon. This has been referred to as a 'flying qubit' and may allow for a great increase in complexity over current quantum computing experiments [39, 57, 40, 48]. Using a photon to transfer quantum states between individual trapped ions is much faster than physical shuttling of ions around a trap [63]. It should also allow for greater integrity of the transferred state. Indeed, an extension to the scheme [128] was published shortly after Cirac *et al.*'s original proposal that promises perfect fidelity over a "noisy" channel.

The proposals and experiments to be reviewed here have a common feature, namely that they employ a Λ (lambda) level scheme and a Raman transition between two of the levels (Fig. 9.1). The atom is held inside a cavity formed by two highly-reflective mirrors. The length of the cavity is set such that it is nearly resonant with one of the transitions in the Λ level scheme, forming half of a Raman transition. The other half of the transition is provided by a pump laser. The detuning from resonance with their respective halves of the transition is the same in the case of both the cavity and the laser.

Although Raman spectroscopy was developed some time ago, an important variation was recently introduced, namely stimulated Raman adiabatic passage (STIRAP) [129]. In its standard implementation it uses two coherent pulses of light to completely transfer population between two quantum states with negligible population transfer to a third, auxiliary state ($|2\rangle$ in Figure 9.1). The important features of the scheme are that the two laser pulses are detuned from resonance with their respective transitions and that the pulse sequence is counter-intuitive: that is, the Stokes pulse arrives before the Pump pulse (see Figure 9.1 for the definition of these) but with a temporal overlap of the two.

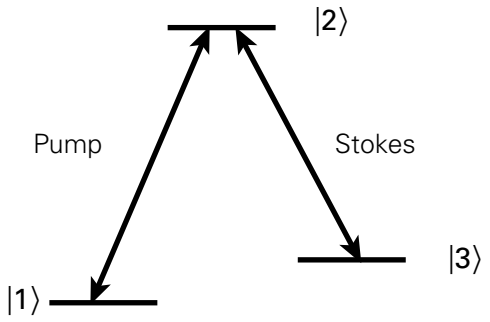

Figure 9.1: The "lambda" three-state level scheme. The name arises from the Greek capital letter, Λ, which is the same shape as the two transitions in the scheme. In the terminology of STIRAP, the left-hand transition is the "pump" while the right-hand is the "Stokes" transition.

The STIRAP process is used in the schemes to be described in order to ensure that the state transfer is coherent. For an atom in an arbitrary state (even a superposition of eigenstates), the state should be mapped faithfully to the emitted photon.

9.3.1 Cirac proposal

The proposal [126] by Cirac *et al.* is for a quantum "network", with nodes consisting of atoms in high-Q (i.e. high-finesse) cavities. The nodes store quantum information in the energy levels of atoms and process this locally using appropriate operations e.g. Ref. [41]. Photons are then used to connect these nodes and transfer quantum information between them. It should be noted that the statement in Ref. [126] that "fast and internal-state-preserving transportation of atoms or ions seems to be technically intractable" has been brought into question by recent results in the group of Wineland [63]. These results do not, however, obviate the need to explore the possibility of using photons for quantum information transfer. Large-scale quantum information processors would require such an elegant system for speedy operation, and quantum communication between remote sites obviously does not allow for the shuttling of ions between nodes.

The goal in such a quantum network is to start with an electronic qubit stored in ion 1,

$$\alpha\,|g\rangle_1 + \beta\,|e\rangle_1 \ , \tag{9.1}$$

and to "map" that qubit onto the cavity mode of cavity 1,

$$|0\rangle_1 \longrightarrow \alpha\,|1\rangle_1 + \beta\,|0\rangle_1 \ . \tag{9.2}$$

That is, the cavity initially has no photons in it and a photon is emitted into the cavity mode if and only if the ion is in state $|g\rangle$. That is not to say that the ground state emits a photon, but rather a process takes place that results in photon emission if and only if the atom starts in the ground state. This counterintuitive situation is a result of the choice of emission process that ensures coherence. The resulting photonic state is then transmitted along a fibre to cavity 2 where the reverse process takes place,

$$|0\rangle_2 \longrightarrow \alpha\,|1\rangle_2 + \beta\,|0\rangle_2 \ , \tag{9.3}$$

$$|g\rangle_2 \longrightarrow \alpha\,|g\rangle_2 + \beta\,|e\rangle_2 \ . \tag{9.4}$$

In summary, then we have the process,

$$(\alpha\,|g\rangle_1 + \beta\,|e\rangle_1)\,|g\rangle_2 \otimes |0\rangle_1\,|0\rangle_2\,|\mathrm{vac}\rangle \longrightarrow |g\rangle_1\,(\alpha\,|g\rangle_2 + \beta\,|e\rangle_2) \otimes |0\rangle_1\,|0\rangle_2\,|\mathrm{vac}\rangle \ , \tag{9.5}$$

where $|\mathrm{vac}\rangle$ is the vacuum mode of the fibre between the cavities.

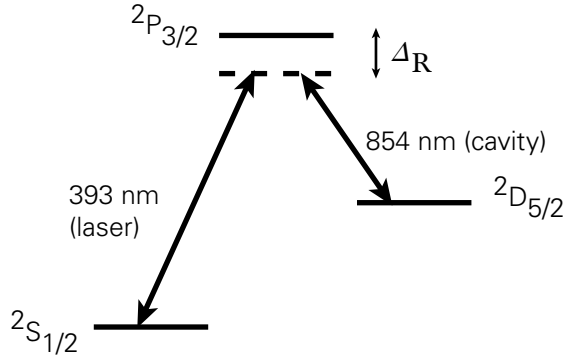

Figure 9.2: Coherent mapping of the electronic qubit onto the cavity mode. The two halves of the STIRAP process are provided by a laser detuned from 393 nm and the cavity mode detuned from 854 nm. The two pathways are detuned by the same amount from resonance, Δ_R.

The atom's internal state is mapped onto the cavity mode using a laser pulse slightly detuned from the 393 nm transition as shown in Figure 9.2. If the ion is

in the ground state, $^2S_{1/2}$ ($|g\rangle$), it will be transferred to the excited state, $^2D_{5/2}$ ($|e\rangle$), with the emission of a photon into the cavity mode. The photon in the cavity mode leaks out of one of the cavity mirrors (one is made more lossy than the other to make one side the preferred channel). If, in contrast, the atom is in the excited state to start with, no photon is emitted into the cavity mode. Since the initial state of the atom has been mapped onto the photon, feeding this photon along a transmission line (e.g. an optical fibre) and into another, identical, cavity causes atom 1's state to be mapped onto cavity 2. The mode of cavity 2 may then be mapped onto the atom in cavity 2.

There will be inevitable inaccuracies inherent in the system. These can include photon absorption in the mirrors and fibre optic transmission line. The authors propose using quantum error correction [37, 36] to overcome these problems manifest in the physical implementation.

Van Enk [128, 130] shows how the use of an extra atom in each cavity allows a perfect transmission, by detecting when the transmission has failed and retrying until it succeeds. The scheme requires initial entanglement of the pair of atoms in the sending node, measurements on the backup atom after attempting to send the information and further local operations to ensure that the state has been succesfully transmitted.

9.3.2 Experimental progress

The group of G. Rempe has done some detailed modelling and further development [131] of the scheme described by Cirac *et al.* Refs. [126, 128]. Their scheme is based around rubidium-85 atoms [132] and in their theoretical paper [131] they derive the pulse shape required for adiabatic transfer of population from the ground to the excited state.

Initial experiments [132] involved large numbers of atoms passing through the cavity. Raman scattering into the cavity mode was demonstrated. A pump laser was directed transverse to the cavity axis and the atoms emitted photons on the anti-Stokes side of the Raman transition into the cavity mode. The counter-intuitive pulse sequence of the STIRAP scheme was achieved by displacing the cw pump probe slightly below the cavity mode but with some spatial overlap of the two. Thus the normal time sequence of the pulses was converted to a spatial sequence due to the atoms' motion.

Later experiments by the same group reduced the atomic flux such that

only one atom was present in the cavity at any given time. The pump laser was changed from being continuously on to being ramped up in power in short pulses. The pulses of the pump laser were interspersed with pulses of a repumping laser that put the atoms back into the ground state after they had emitted a photon. Thus, several sequential photons were emitted by a single atom in the cavity. The major technical hurdle still to be overcome by the group is to trap an atom within the cavity so that photons can be generated on-demand instead of being reliant on the atom flux pattern.

H. Kimble's group in Caltech has achieved the trapping of an atom within the cavity mode using a Far Off-Resonant Trap (FORT) [118], otherwise known as a dipole trap. Since then they have produced single photons [133] and counted the number of atoms trapped [134].

Another group pursuing the goal of a single-atom laser is that of W. Lange and H. Walther. Like the Innsbruck group, their system is that of a single calcium ion in an RF trap. The Pump and Stokes lines of the Raman transition are different from those in the Innsbruck experiment, however. The 397-nm cooling transition is used instead of the 393 nm $^2S_{1/2}$–$^2P_{3/2}$ transition and output is at 866 nm. Their achievements to-date include using a single ion as a probe of the internal electromagnetic field of the cavity [135]. They recently published results demonstrating the production of single photons [136, 137]. With a low ion–cavity coupling constant, however, the photon was emitted into the cavity mode just 8% of the time with the other 92% of emitted photons emitted via spontaneous decay into free space.

9.4 Laser and cavity locking

Although CW lasers are almost monochromatic in comparison with other light sources, they do still have a finite linewidth. The centre frequency is also found to drift with time. Even in the 1960s, shortly after their invention, efforts were made to "lock" the lasers such that these drifts were reduced. As the field of spectroscopy has developed, the need for better laser-frequency stability and narrower linewidths has increased. The narrower the linewidth and the greater the frequency stability, the greater is the resolution of spectrsocopy that can be performed.

A second field has also driven the development of laser and cavity locking. It is that of gravitational-wave detection. Einstein's theory of general relativity

predicts the existence of gravitational waves. These weak ripples in space can be detected by the lengthening and contraction of a body. The change in length is extremely small, however, and a very precise measure of a body's length had to be developed. Interferometry is an obvious candidate for this and is used in the American LIGO project, the European VIRGO observatory, and the proposed space-based detector LISA, all based on ultra-stable optical cavities. The cavity is stabilised by monitoring the reflected light from a stable laser incident on it. In most gravitational observatories, the cavity is several kilometres in length.

As has been hinted at in the previous two paragraphs, the locking of a cavity's length and a laser's frequency are closely related techniques. Either the laser or the cavity can be seen as a reference frequency and the interaction with the other can be used to detect fluctuations in frequency. I shall primarily describe the locking of a laser here and show at the end of the description how the technique is easily modified to lock a cavity.

9.4.1 Sources of noise

There are two broad categories of noise in a laser's output: "fundamental" noise and "technical" noise [138]. These categories apply both to the laser's intensity and to its frequency. We shall be concerned here only with its frequency.

Fundamental noise, as its name suggests, can not be removed from a laser source. It is due primarily to shot-noise and to spontaneous emission. Shot noise arises from the quantum nature of light: photon arrival times are inherently random. The power spectral density is given by:

$$S(f) = 2\bar{I}h\nu \ , \tag{9.6}$$

where $\bar{I}$ is the average intensity. This indicates that the noise is "white", i.e. it is independent of the Fourier frequency.

Spontaneous emission will inevitably occur in any lasing medium to some extent. In contrast to the stimulated emission, the spontaneous emission has a random phase with respect to other photons. This phase difference gives rise to a small frequency change as can be seen in the equation for the instantaneous frequency,

$$\omega(t) = \frac{\mathrm{d}\phi}{\mathrm{d}t} \ . \tag{9.7}$$

The Schawlow–Townes limit is the ultimate limit on the linewidth of a laser

source.

$$\Delta\nu = \frac{h\nu_0}{2\pi}\, c^2\, \frac{(1-R)^2}{PL^2}\ . \tag{9.8}$$

The removal of all technical noise will leave a laser operating at the Schawlow–Townes limit. The parameters in the equation are ν_0, the centre frequency; R, the reflectivity of the mirrors of the laser cavity; P, the output power; L, the laser-cavity length. Control over the length and reflectivity is generally possible in the design of a system. It can be seen that a diode laser with a cavity of order a few millimetres in length will have a high limit while adding a (longer) external cavity around the diode (see Chapter 4) will reduce the linewidth limit. The reflectivity of the external mirrors will be much higher than the low reflectivity of the diode faces and will help to reduce the linewidth still further by reducing $(1-R)$ in Equation 9.8.

Sources of technical noise are more varied than those of fundamental noise. They are typically environmental and those that receive the most attention are acoustic and thermal noise. Acoustic noise is often high in a laboratory given the amount of equipment running at any given time. Extensive measures are normally in place to minimise its effect on the optical systems. The measures include optical tables with high levels of damping, rigid supports for optical mounts, and the enclosure of the entire table with plastic curtains. Whilst these measures are very effective, it is impossible to completely eliminate acoustic noise. Thermal drift affects laser systems through the expansion and contraction of cavities. The resonant frequency depends on the cavity length so these fluctuations in length are mapped onto the output frequency of the laser.

9.4.2 Measuring frequency fluctuations

The first step in locking a laser system is to detect the frequency fluctuations so that they can be corrected for. This requires the comparison of the laser frequency with a frequency reference at least as stable as the ultimate desired laser stability. An error signal is derived from this comparison: a voltage proportional to the deviation from the desired frequency. The most commonly-used references are atomic or molecular spectral lines (especially in laser-cooling, see Chapter 4) and Fabry–Perot cavities.

The use of a spectral line (especially an atomic line) places some constraints on the tunability of the laser: there may be no line near the desired wavelength for example. The locking scheme may be complicated by the requirement of a

technique such as saturated absorption spectroscopy [139] in order to lock to a line within a Doppler profile. Finally, a problem inherent in all schemes locking to spectral lines is the lack of control over the slope of the error signal since this is related to the lineshape of the relevant atomic or molecular transition.

With the advent of modern dielectric mirror coatings with extremely high reflectivities, coupled with the use of new cavity-spacer materials such as Ultra-Low Expansivity (ULE) glass, cavities can be constructed with very high finesse ($\sim 250,000$) and with almost negligible thermal drift. These provide excellent frequency references.

In the most basic configuration, the light transmitted through the cavity is measured using a photodiode. Peaks of transmission are seen at regular intervals as the laser frequency is scanned. The interval is equal to the free spectral range (FSR) of the cavity where the cavity length is an integer number of half-wavelengths of the light.

$$\mathrm{FSR} = \frac{c}{2L} \ . \tag{9.9}$$

The shape of the peaks is known as an Airy function (Equation 9.13). As the reflectivity of the mirrors increases the peaks become sharper while the free spectral range remains constant. Thus, a measure of the slope or narrowness of the peak is given by the "finesse" of the cavity:

$$F = \frac{\mathrm{FSR}}{\gamma} \ , \tag{9.10}$$

where γ is the full width at half maximum of the peaks. Another measure frequently used is the "coefficient of finesse", F. It is defined,

$$\mathsf{F} = \left(\frac{2r}{1-r^2}\right)^2 \ , \tag{9.11}$$

with r being the reflectivity of the mirrors, and is related to the finesse by the relation,

$$F = \frac{\pi\sqrt{\mathsf{F}}}{2} \ . \tag{9.12}$$

The Airy function in terms of the coefficient of finesse is,

$$\mathcal{A} = \left[1 + \mathsf{F}\sin^2(\delta/2)\right]^{-1} \ , \tag{9.13}$$

where δ is the phase difference between the first incident beam and the first "round-trip" beam. For example, if the cavity length is an integer number of half-wavelengths, the phase difference after a round trip will be $2n\pi$ with n an integer. The peaks for a range of coefficients of finesse are shown in Figure 9.3.

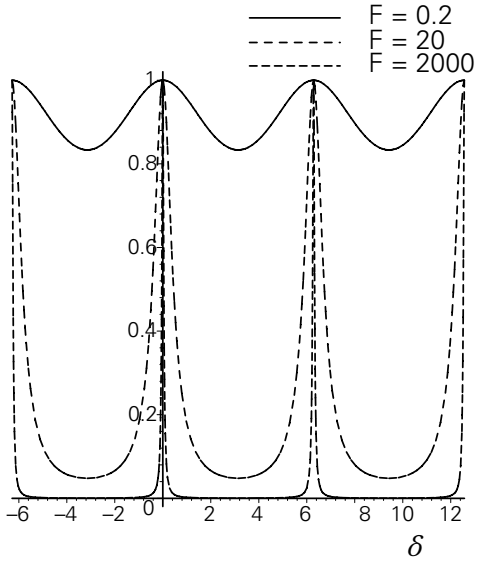

Figure 9.3: Transmission (Airy) functions for cavities with different coefficients of finesse, F.

9.4.3 Locking to the side of a fringe

The most simple locking scheme conceptually is a side-of-fringe lock to a transmission peak. The laser frequency is initially chosen such that it closely matches one of the transmission peaks of the cavity. The laser is tuned such that the transmitted intensity sits roughly half-way up one of the transmission peaks, since this is the point of greatest slope. A small change in frequency will result in a move up or down the peak depending on the direction of the change. Thus, an error signal with an approximately linear relation to frequency drift is obtained. It would not have been possible to obtain information about the direction of the frequency change had the lock been to the peak of the fringe since the intensity would vary in the same manner in either direction.

The biggest problem with side-of-fringe locks is that the error signal is not decoupled from intensity fluctuations. A change in the intensity of light trans-

mitted through the cavity could result either from a frequency change or from a change in the intensity of the laser. The most common way to combat this is to make a separate measurement of the intensity of the laser on another photodiode. An electronic subtraction of the two measurements, independent intensity and transmitted intensity, can be made to partially overcome the problem. It is not an ideal solution as it requires a careful balancing of the two photodiode signals and is dependent on the linear response of the photodiodes. Nevertheless, the scheme has been successfully implemented in many commercial systems such as the Coherent CR-699 dye laser (see Chapter 7) and even without the intensity-fluctuation compensation as described in Chapter 4.

Decoupling intensity and frequency fluctuations can be achieved using other locking schemes. One, commonly used in frequency-doubling enhancement cavities, is the Hänsch–Couillaud lock [140] which analyses the intensity of the two orthogonal polarisations of the light reflected from the cavity. It is a relatively inexpensive scheme to build and has a better performance than side-of-fringe locks. It still requires the use of two photo-diodes, however, and thus the problem of matching photodiodes' frequency response and drift is not alleviated. The best performance to-date is achieved using a scheme known as the Pound–Drever–Hall lock, a locking scheme which has the advantage of requiring just one photodiode.

9.4.4 The Pound–Drever–Hall lock

Pound [141] described a scheme for stabilising a microwave source based on measuring the output of an impedance-matching tank circuit with a high-Q cavity. It provided the inspiration for a later scheme devised by Drever and Hall *et al.* to stabilise a laser system [80]. The key feature of both schemes is the frequency modulation of the source and the detection of the radiation reflected from a high-Q cavity.

The Pound–Drever–Hall scheme is described in overview by Hamilton [138] and in more detail by Black [142]. The original paper by Drever *et al.* [80] is very accessible with the exception of the bypass-topology circuit which is unfamiliar to most physicists. Pound's 1946 paper is primarily of historic interest given that it refers to the microwave régime and uses now-defunct technology such as valves.

The scheme measures the reflected light from the cavity. When the laser is

resonant with the cavity this signal should be at a minimum, i.e. as much light as possible is transmitted through the cavity[1]. This has the great advantage that fluctuations in laser intensity are irrelevant to the lock as the aim is to keep the reflected intensity at a minimum rather than at some arbitrary value.

In the side-of-fringe scheme, the direction of the frequency shift can be inferred from the rise or fall of the intensity. This is not true, however, if we lock to the centre of a fringe as the change in intensity is the same in either direction. This can be overcome by dithering the laser frequency slightly. The reflected intensity will thus vary and if the centre-frequency drifts above or below resonance the intensity variation will be in-phase or out-of-phase (respectively) with the dither. A phase-sensitive detection scheme can be used to generate a useful error signal.

If the dither frequency is so large that the field inside the cavity does not have time to respond to the change, then a new régime is entered. The lock still functions as required, and in fact has some rather desirable properties such as the ability to reduce the linewidth of the laser.

With a large frequency-modulation (i.e. dither), sidebands appear on the laser output. This is analogous to the sidebands on a moving particle's spectrum, see Chapter 2. The electric field of the laser beam can be described,

$$E_{\text{inc}} = E_0 e^{i(\omega t + \beta \sin \Omega t)} \ , \tag{9.14}$$

where E_0 is the initial amplitude of the field before the addition of sidebands, ω is the centre frequency of the laser, β is the modulation depth (which controls the intensity of the sidebands), and Ω is the modulation frequency. This can be expressed in terms of Bessel functions as seen in Chapter 2,

$$E_{\text{inc}} \approx E_0 \left[J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega+\Omega)t} - J_1(\beta) e^{i(\omega-\Omega)t} \right] \ , \tag{9.15}$$

neglecting Bessel functions of order greater than one, i.e. only the first-order sidebands are treated here. This is valid if the modulation depth is not very large.

The system can be considered in terms of a forced, damped harmonic oscil-

[1]The reflected light will only fall to zero if the mode-matching and impedance-matching of the beam to the cavity is perfect. Even if there is some reflected light when on-resonance, however, the only effect on the Pound–Drever–Hall scheme is to change the slope of the error signal rather than the frequency-offset of the lock-point.

lator (see Chapter 7). The cavity field is the oscillator itself while the incident beam is driving the field. The reflected beam from the cavity consists of the prompt reflected beam and a leakage beam due to the reflectivity of the coupling mirror of the cavity being less than one. A phase difference will exist between the incident beam and the cavity field if the beam is not on resonance. As the frequency of the beam is scanned through resonance with the cavity, the phase will undergo a π-shift. This is a crucial point in understanding the Pound–Drever–Hall lock: we can detect which side of the cavity resonance the laser is on (in theory) by measuring the phase of the reflected beam. In practice, it is very difficult to directly measure the electric field of a laser beam (see, however, Ref. [143]).

Measurement of the relative phase of two laser beams is much more straightforward than of absolute phase of a single beam. A reference phase is required and this is the function that the sidebands perform. Since they are derived from the same source as the carrier, there will be a constant phase relationship between each sideband and the carrier. Since they are so far away from resonance with the cavity, their phase in the reflected beam is unaffected by frequency drift (by contrast with the carrier). The sidebands in the reflected beam interfere with the carrier (which consists of both a prompt reflection and cavity leakage) to form a beat pattern with the same frequency as the modulation. The phase of this beat signal can then be measured. It can be shown [142] that the power of the reflected beam is given by:

$$\begin{aligned} P_{\text{ref}} &= P_c|G(\omega)|^2 + P_s\left\{|G(\omega+\Omega)|^2 + |G(\omega-\Omega)|^2\right\} \\ &\quad +2\sqrt{P_cP_s}\Big\{\text{Re}\left[G(\omega)G^*(\omega+\Omega) - G^*(\omega)G(\omega-\Omega)\right]\cos\Omega t \\ &\quad +\text{Im}\left[G(\omega)G^*(\omega+\Omega) - G^*(\omega)G(\omega-\Omega)\right]\sin\Omega t\Big\} \\ &\quad +(2\Omega \text{ terms}) \,. \end{aligned} \tag{9.16}$$

2Ω terms are omitted as they are due to interference of sidebands with each other. G is the reflection coefficient, the ratio of the reflected and incident electric fields. The symbol G is used rather than the more common F to avoid confusion with the coefficient of finesse. It is defined,

$$G(\omega) = E_{\text{ref}}/E_{\text{inc}} = \frac{r\left(\exp\left(i\frac{\omega}{\Delta\nu_{\text{fsr}}}\right) - 1\right)}{1 - r^2\exp\left(i\frac{\omega}{\Delta\nu_{\text{fsr}}}\right)} \,. \tag{9.17}$$

E_{inc} and E_{ref} are the incident and reflected electric fields, ω is the angular frequency of the light, and r is the reflectivity of the mirrors.

Of the terms in Equation 9.17, those oscillating at frequency Ω are of interest as they include the phase of the reflected carrier. Whether it is the $\cos\Omega t$ or the $\sin\Omega t$ term that is observed depends on Ω, the modulation frequency. For the low-frequency régime where the cavity field has time to respond, $\Omega \ll \Delta\nu_{\text{FSR}}/F$ and the cosine term is observed while the sine term vanishes. In the high-frequency régime, $\Omega \gg \Delta\nu_{\text{FSR}}/F$ and it is the sine term that is observed.

The selection of terms oscillating at frequency Ω can be done with an electronic mixer with inputs from the photodiode detecting the reflected light and the oscillator at the modulation frequency, Ω. This gives rise to both 2Ω terms and a dc signal at the output of the mixer, of which the latter is isolated using a low-pass filter. The resulting error signal is

$$\epsilon = -2\sqrt{P_c P_s}\,\text{Im}\left\{G(\omega)G^*(\omega+\Omega) - G^*(\omega)G(\omega-\Omega)\right\} \ . \tag{9.18}$$

At resonance, the frequency ω is a multiple of the cavity's free spectral range and thus for a small deviation, $\delta\omega$,

$$\frac{\omega}{\Delta\nu_{\text{FSR}}} = 2\pi N + \frac{\delta\omega}{\Delta\nu_{\text{FSR}}} \ , \tag{9.19}$$

where N is a large integer. Assuming a high cavity finesse, the reflection coefficient can then be approximated by,

$$G \approx \frac{i}{\pi}\frac{\delta\omega}{\gamma} \ , \tag{9.20}$$

expressing the frequency deviation as a fraction of the cavity linewidth. This yields a new approximation for the error signal in the high-finesse and small frequency-deviation régime:

$$\epsilon \approx -\frac{4}{\pi}\sqrt{P_c P_s}\,\frac{\delta\omega}{\gamma} \ . \tag{9.21}$$

A linear relationship with $\delta\omega$ means that it is straightforward to feed this error signal back to the laser. A small change in the laser frequency results in a proportional small change in the error signal.

A schematic of the locking system is shown in Figure 9.4. An EOM provides

frequency-modulation (really phase-modulation, but this amounts to the same thing here). A method of sending the reflected light along a different path than it came in on is required. This is done by introducing a polarising beam-splitter followed by a quarter-wave plate. The first passage of the beam through the wave plate introduces a quarter-wave rotation while the second passage, after reflection from the cavity, introduces an additive rotation. The full rotation of $\lambda/2$ means that the beam reflected from the cavity is then reflected from the polarising beam splitter.

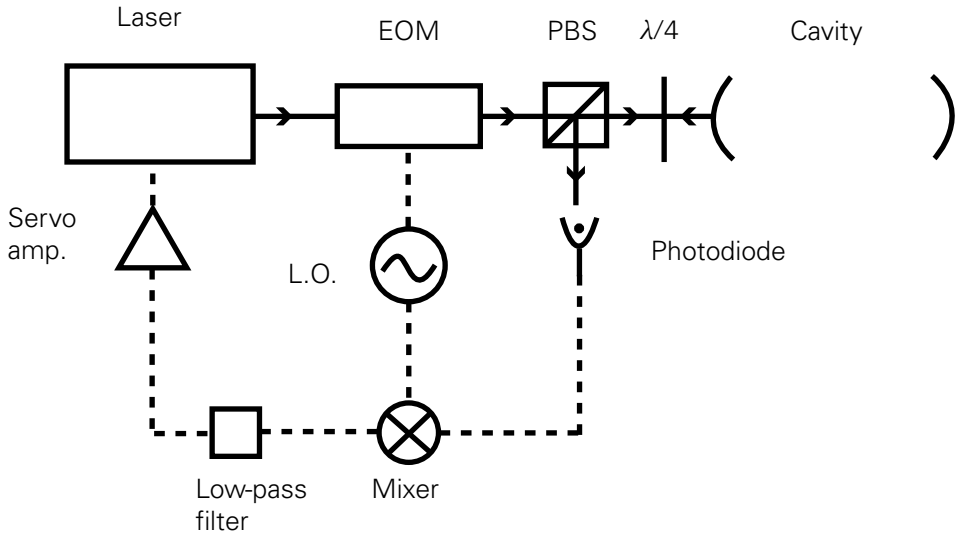

Figure 9.4: Schematic of the Pound–Drever–Hall lock. Solid lines indicate the laser beam with the arrows showing the direction of propagation. Dashed lines indicate signal paths in the scheme. Some components necessary in practice, such as a Faraday optical-isolator, are omitted for clarity.

The error signal in Equation 9.21 is seen after the low-pass filter. It is fed to some standard control electronics such as a PID regulator before being fed-back to the laser frequency control. The original Drever–Hall paper suggested the use of an amplifier with bypass topology instead of a PID regulator. This is unnecessary today with major advances in the speed of operational amplifiers.

9.4.5 Cavity Stabilisation

What has preceded mainly dealt with the stabilisation of a laser frequency. The idea of stabilising a cavity to a laser is very closely-related to this. In the simple

side-of-fringe lock, it was seen that a frequency excursion of the laser would cause the transmitted intensity to move up or down the Airy fringe. A change in cavity length will cause the fringe itself to move in frequency space. The result is the same, however: a change in the transmitted intensity of the laser. Given a sufficiently stable laser, the error signal can be fed back via some servo mechanism to the cavity. This might be a heater that heats the cavity to lengthen it or turns off to let it contract again. In a higher-bandwidth system the servo mechanism will typically be a PZT connecting one of the mirrors to the cavity spacer.

9.5 Experimental Setup

The experiment being built in Innsbruck places an ion in the centre of a high-finesse cavity. The ion will be laser-cooled and addressed with a laser at 729 nm in order to implement qubit rotations. Another laser, at 393 nm, will be used as the pump laser in the STIRAP process while the cavity will be tuned to the Stokes transition near 854 nm.

Two Pound–Drever–Hall locks are required for the planned experiments as shown in Figure 9.5. The pump laser is stabilised to an ultra-stable reference cavity. A second laser is locked to the same reference cavity and the cavity around the ion is stabilised to this second laser. This is referred to as a "transfer lock", where the stability of one cavity is conferred onto another using a "transfer laser".

9.5.1 Ion-trap cavity

The system of interest, an atom located in the mode of a Fabry–Pérot cavity formed by two concave, highly-reflective mirrors, is a difficult one to achieve experimentally. Stability of the ion trap with respect to the cavity must be ensured, along with the ability to precisely position the trapped ion in the centre of the cavity mode. Additionally, the finesse of the cavity must be very high so that the ion is in the strong-coupling régime required for many CQED experiments. The coupling is also affected by the volume of the mode. The smaller the mode, the higher the intensity of light per unit volume and the greater the coupling

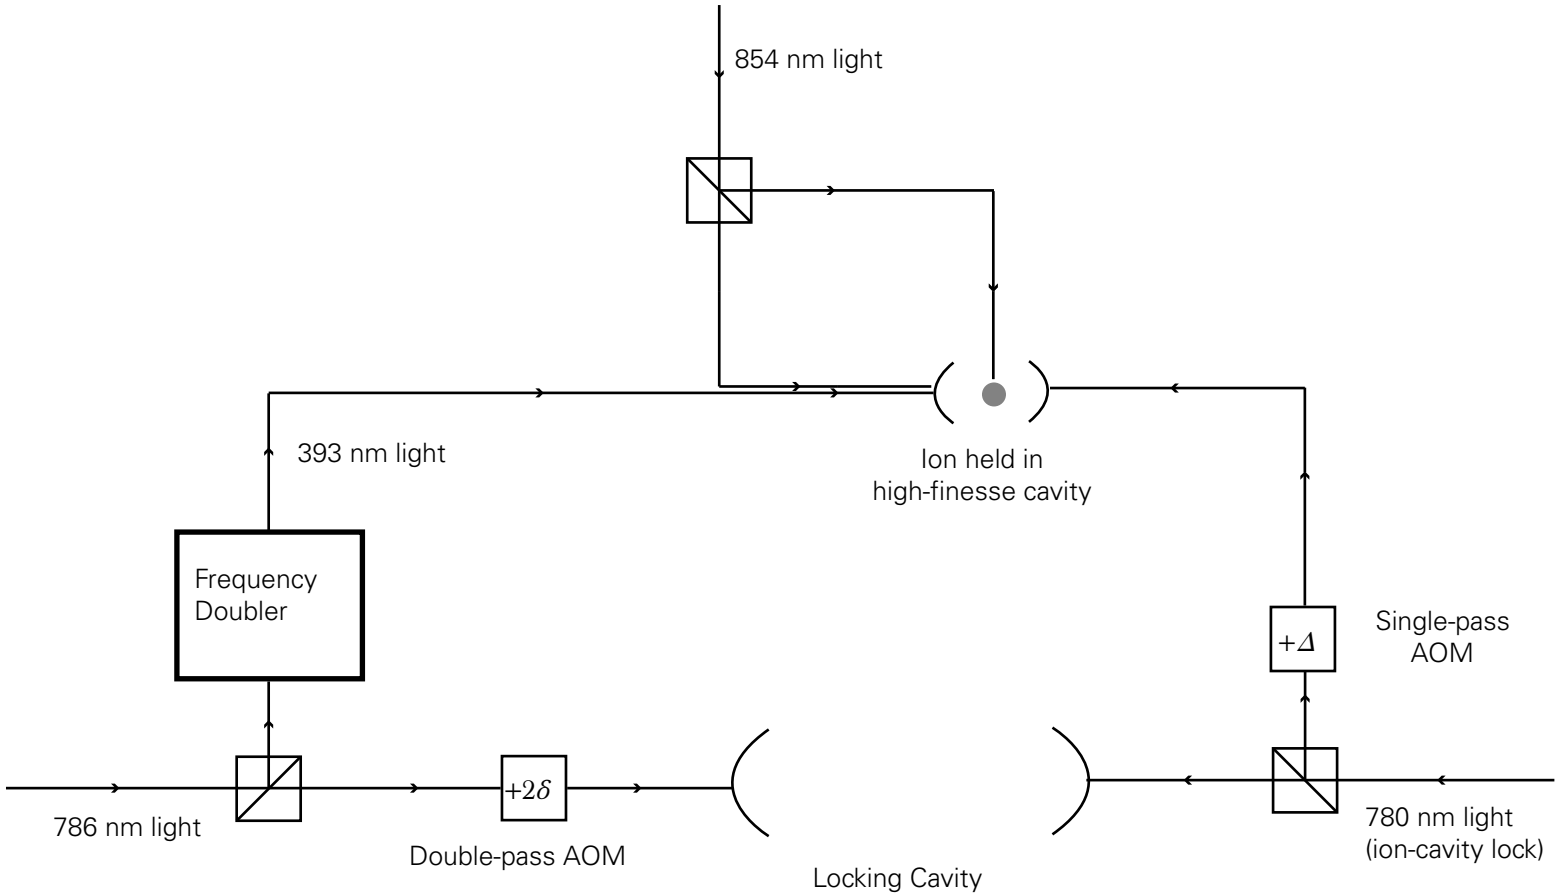

Figure 9.5: A schematic of the dual Pound–Drever–Hall lock used to stabilise both sides of the STIRAP scheme. The use of the same cavity to lock both sides reduces the effect of the locking-cavity drift on the experiment: both the laser and cavity frequency will drift in the same direction as the cavity. Thus Δ_R will not vary as dramatically as it might with separate locking cavities. The fact that one laser is frequency-doubled means that any drift will be double that of the other, however.

strength, g, between the ion and the cavity [120]:

$$g = \sqrt{\frac{2c\gamma\lambda^2}{\pi^2 L w_0^2}}\ , \tag{9.22}$$

where γ is the spontaneous decay lifetime of the transition coupled to the cavity, λ is the wavelength of that transition, L is the length of the cavity, and w_0 is the size of the waist of the cavity mode.

The mode size of the cavity is determined by the radius of curvature of the mirrors and their separation (see Fig. 9.6). The ideal mirror separation to achieve the smallest mode volume is such that the cavity is "concentric", i.e. the centres of curvature of the two mirrors overlap. For mirrors with equal radii of curvature the mirror separation is twice the radius of curvature. As can be seen from Fig. 9.6 the mode size reduces rapidly near concentricity.

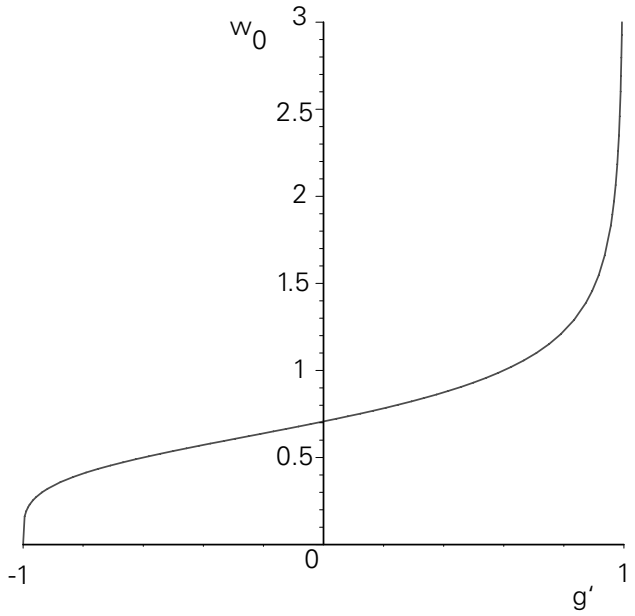

Figure 9.6: Variation of cavity waist with $g' = 1 - L/R$. The mirrors are confocal at $g' = 0$ and are concentric at $g' = -1$.

Concentric cavities are unfortunately extremely unstable. Slight lateral movements of the mirrors can cause the mode to disappear entirely. Similarly, any increase in the length of the cavity will cause the mode to disappear from the centre of the cavity. Alignment of a near-concentric cavity requires a very

stable mechanical system.

The mirrors used were very high-reflectivity (at 854 nm) concave mirrors with a radius of curvature of 10 mm, manufactured by Research Electro-Optics (REO). The transmissivities of the mirrors were 20 ppm and 4 ppm respectively. The mirror with the larger transmissivity acts as the output coupler as there is a greater probability of a photon emitted into the cavity mode passing through that mirror.

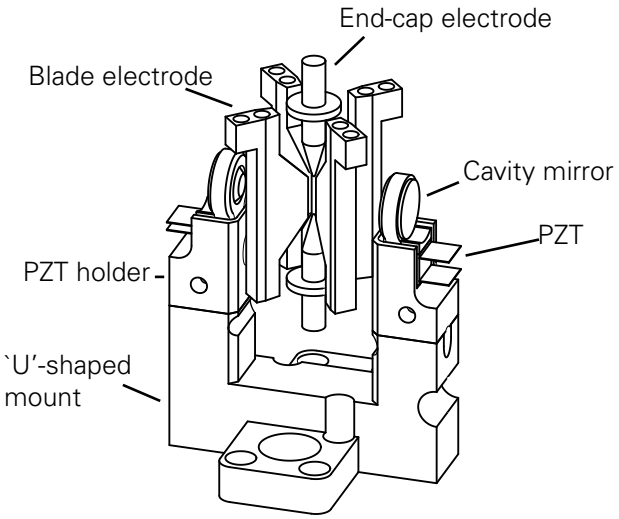

Figure 9.7: The trap and cavity. The two mirrors are mounted on piezo stacks and then glued to a 'U'-shaped holder. They surround the trap (the long vertical blades and the two points in the centre). Figure courtesy of C Russo.

The entire mechanical design was done by Russo. The mirrors were mounted in stainless steel holders which were then attached to piezoelectric transducers (PZTs). The PZTs allow longitudinal adjustment of the mirror position; this adjustment will provide a means of locking the length of the cavity using a Pound–Drever–Hall scheme. Different PZTs were used at each end of the cavity such that one offered a larger translation while the other offered a larger bandwidth. Finally, the PZTs were to be attached to a 'U'-shaped piece of stainless steel (see Fig. 9.7) such that the mirrors would be approximately 40 mm apart (i.e. producing a near-concentric cavity). This work has been completed since the end of my time working in the group. Epoxy was used to hold the structure together. Epoxies have come to the market recently that have very low out-gassing rates and are suitable for use in UHV applications. In attaching

the mirrors to the steel holders, and in attaching the PZTs to the 'U', an ultraviolet-cure epoxy (Dymax OP-66-LS) was used. This allowed a fast cure *in situ*, necessary since the alignment of the pieces was crucial. After the initial cure with a high-power UV lamp, the pieces were baked at 100 °C for 45 minutes to ensure that epoxy away from the surface (and hence not exposed to the light) was properly cured. Kapton- (polyimide) coated wires were attached to the PZT electrodes with a conductive epoxy (Polytec/Epotec H20-E). The mirror holders were glued using a standard heat-cure epoxy. Where heat-cure was required *in situ*, this was done with a heat gun. A thermocouple was used to measure the temperature and ensure that the PZTs were not heated above the Curie temperature ($\sim$ 120 °C) as this could damage them by depolarising the crystal structure. Once depolarised, the piezoelectric effect of different parts of the crystal would act in random directions and no macroscopic motion would be observed.

A preliminary test of the cavity lock has been conducted and the PZTs' bandwidth has shown to be ample. The finesse of the cavity has been measured to be 90,000.

9.6 Ti:sapphire laser lock

The other locking system on which I worked was that of a Ti:sapphire laser (Coherent 899) to which an intra-cavity EOM had been added (see Figure 9.8). The Ti:sapphire laser was pumped by a Coherent Verdi V5 with output of 5 W at 532 nm. The output power of the Ti:sapphire at 786 nm is approximately 400 mW with all optical elements apart from the EOM in place inside the cavity. The addition of the EOM introduced negligible losses after it had been well aligned.

The output of the Ti:sapphire laser was intensity-stabilised by an AOM (Isomet 1205-603F). This was done by controlling the amount of power picked off in the first order by a feedback circuit connected to a photodiode. Only the zeroth-order spot was allowed into the frequency doubler and the locking optics.

The light at 786 nm is frequency-doubled to produce light at 393 nm. Since frequency-doubling is a non-linear process, an increase in power of the fundamental frequency (here, equivalent to 786 nm) gives rise to a quadratic increase in the power of the frequency-doubled output. This increase in power of the fundamental is achieved by placing the frequency-doubling crystal in a

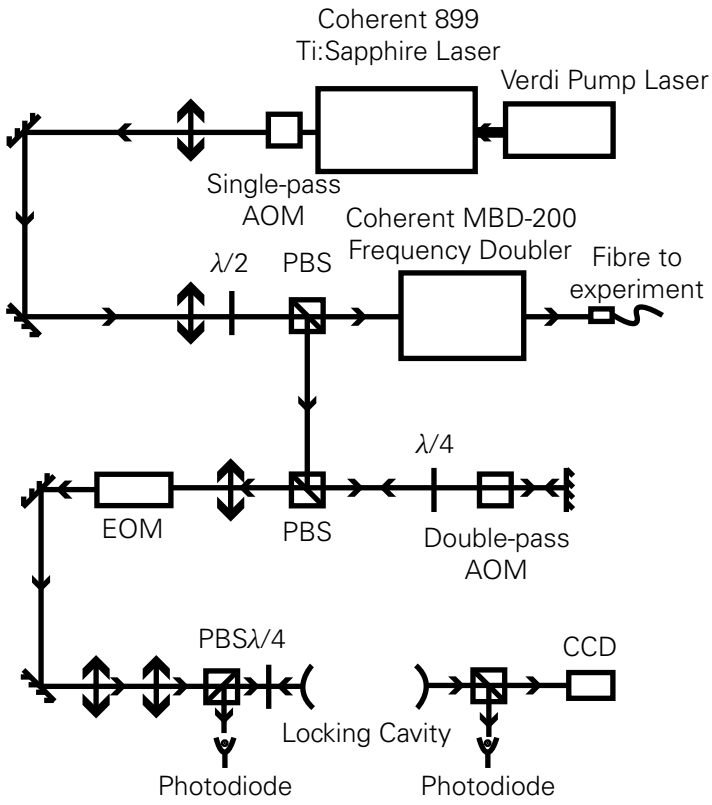

Figure 9.8: The Pound–Drever–Hall locking set-up for the 786 nm laser. The optical set-up has been simplified with the exclusion of steering mirrors and some other components.

stable cavity which, when locked to the fundamental-beam frequency, causes a build-up of power inside the cavity. The intensity-stabilised Ti:sapphire output is frequency-doubled in a commercial doubling-enhancement cavity (Coherent MBD-200) to produce light at 393 nm (close to the $S_{1/2} - P_{3/2}$ transition of the Ca^+ ion). The cavity contains a lithium borate (LBO) crystal which has good non-linear properties. Conversion efficiencies of up to 20% were measured.

The Ti:sapphire laser was frequency-stabilised using a Pound–Drever–Hall lock. Here, the cavity to which the laser was locked was formed by two highly reflective (at 729 nm) mirrors separated by a ULE spacer. The mirrors had been bought with the intention of locking a 729 nm laser but the reflectivity at 786 nm was found to be sufficiently high to give a high finesse ($\sim 16,000$) at this wavelength also. The mirrors were identical and had a curvature of 50 cm.

Their separation was 15 cm. The waist of the mode at the centre of the cavity is given by [83]:

$$w_0 = \sqrt{\frac{\lambda}{2\pi}} \left(2RL - L^2\right)^{1/4} \ , \tag{9.23}$$

yielding a waist of 212 μm for the given cavity.

In all schemes of locking a laser to a cavity, it is important to ensure that the "mode-matching" is good, i.e. that as much power as possible goes into the TEM_{00} mode. This ensures that the lock is to the correct peak and reduces the likelihood of capture of an adjacent feature in the spectrum. It also means that almost all of the available power is being used since none is lost to other modes, allowing a larger proportion to be used for the experiment proper.

Mode-matching imposes two requirements, namely that the spatial position and the waist of the beam match those of the cavity's mode. Lenses are used to achieve the correct beam waist, in the correct longitudinal position. A pair of mirrors are used to achieve the correct transverse position of the beam. Due to the sensitivity of the system, it is normal for the lenses to be mounted on translation stages to allow fine adjustment. High-quality mirror mounts (Radiant Dyes) are used to steer the beam into the cavity.

Rough mode-matching was achieved using a "waist meter" (Thorlabs WM-100 Omega meter). This device measures the spot size of the beam (not necessarily the waist, unless the correct longitudinal position is chosen). Measurement in both transverse directions was important to ensure that no astigmatism had been introduced by non-perpendicular alignment of mode-matching lenses.

The EOM before the locking cavity was driven by a circuit built from Mini-Circuits components (the "local oscillator" in the PDH scheme), and amplified using an amplifier chip from a photocopier[2] (chip model number Xerox 1165756). The driving frequency was 17 MHz.

When the frequency of the local oscillator closely matched that of the induction coil attached to the EOM, and when the amplification was sufficiently large, sidebands could be observed on the laser's spectrum when in scanning mode (Figure 9.9). As expected these were separated from the carrier by about 17 MHz. This separation was calculated from the knowledge of the laser scan range.

[2]Acousto-optic modulators are commonly used to scan laser beams in photocopiers and RF power amplifiers are included to drive them.

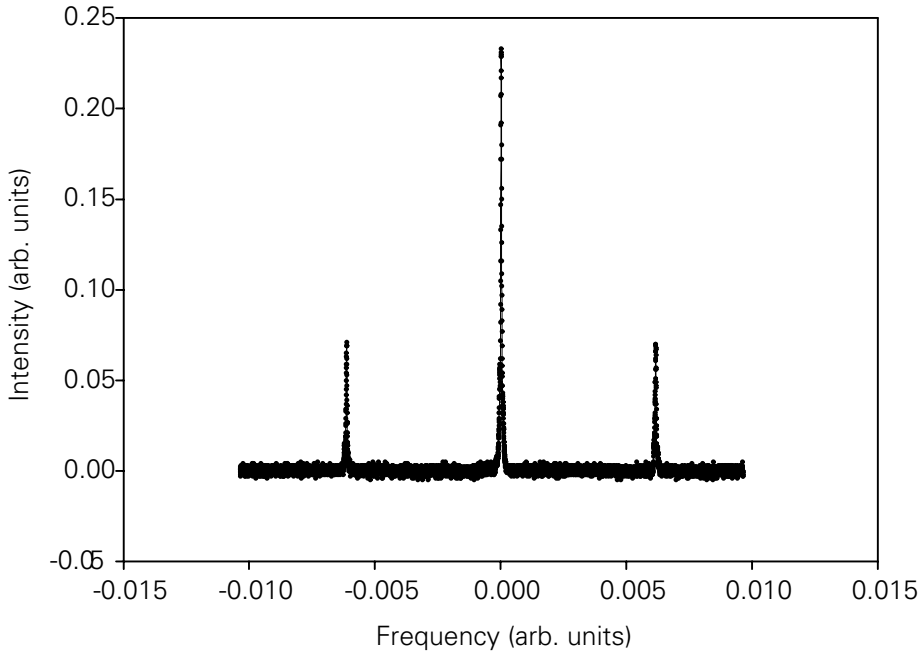

Figure 9.9: Cavity transmission measured on a photodiode. Although the frequency scale is uncalibrated, the sidebands are known to be separated from the carrier by the modulation frequency of 17 MHz.

The signal from the photodiode measuring the reflected light from the cavity was fed through a series of electronic devices (all manufactured by Mini-Circuits) in order to obtain an error signal. Firstly it was sent to a mixer (ZRPD-1) and then to a low-pass filter (BLP-19). The other input to the mixer came from the local oscillator supplying the EOM imposing the sidebands. A sample of this signal with the laser scanning can be seen in Fig. 9.10.

The signal was fed back to the laser via the Coherent-laser control box. Normally this box takes as its input the error signal from the reference cavity attached to the laser, a low-finesse (~100), temperature-stabilised cavity. Feeding back the PDH error signal instead of the laser's built-in error signal should provide a better lock. There was not time to measure whether or not this was the case. It was possible, however, to observe the transmitted mode on a CCD camera and to see the TEM_{00} mode stable on it. The error signal with the laser locked via the PDH error signal is shown in Fig. 9.11. This error signal shows how poor this preliminary lock was as the noise is substantial.

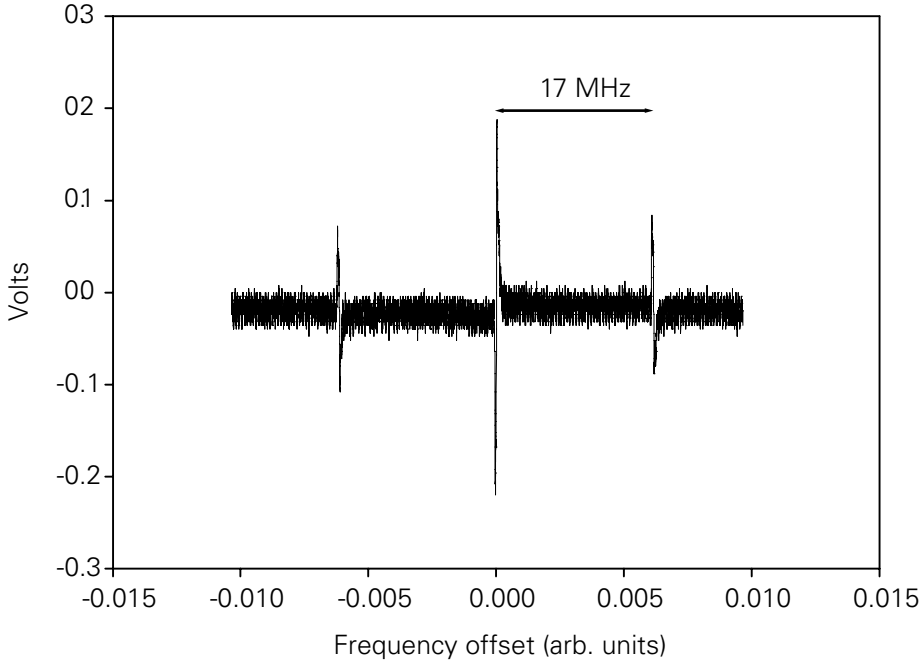

Figure 9.10: The Pound–Drever–Hall error signal with the laser scanning over the cavity resonance

9.7 Future Directions

Since I left the group, the cavity and trap have been completed and put in vacuum. Work on the laser lock has continued but has yet to be completed. It is hoped that in the next twelve months single-photon generation will be demonstrated.

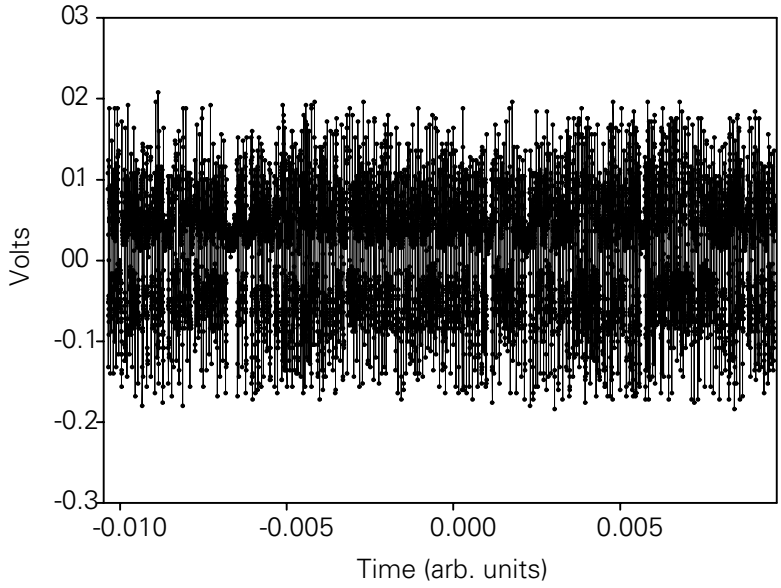

Figure 9.11: The error signal when the laser was locked. Ideally there would be no deviation from zero. The large noise indicates that the laser was badly locked.

Chapter 10

Discussion

The goal of this work has been the development of a system to allow the investigation of quantum information processing with calcium ions in a Penning trap. I have implemented experimental developments to this end. The most important of these has been the first demonstration of the axialisation of laser-cooled calcium ions.

Doppler-cooling of calcium ions in a Penning trap had already become an established technique in our group by the time I started. Major improvements to the system were required, however, including new wavelength measurement apparatus. and laser-locking cavities. The Michelson interferometer built here, combined with analysis electronics designed by the University of Innsbruck, allows the measurement of wavelengths to one part in 10^7. This has allowed the direct loading of ions in the Penning trap, replacing a cumbersome procedure of tuning lasers to a hollow-cathode lamp and gradually increasing the magnetic field while tuning the lasers to stay on resonance.

Locking the lasers with a side-of-fringe lock to cavities has greatly eased the running of the experiment. Large, independent, frequency drifts of the different diode lasers have been reduced. While this development has certainly allowed much more complex experiments than before, the temperature stability of the cavities appears now to be a limiting factor. Even without the design of new cavities, several steps can be taken to improve matters. The first step is to improve the temperature-stabilisation circuit. The addition of photodiodes to monitor laser intensity fluctuations and decouple these from frequency fluctuations would also be beneficial in reducing laser frequency noise. To improve the frequency-stability much further would require the design of new cavities

with higher-reflectivity mirrors. A locking scheme which is inherently insensitive to intensity fluctuations such as the Pound–Drever–Hall scheme might then be considered.

In order to gain experience of the technique of axialisation, work was performed first on laser-cooled magnesium ions. The simpler level structure of magnesium means that laser-cooling can be achieved with just one laser. This simplicity makes it a useful system for testing techniques. Magnesium has two main disadvantages for our purposes, however: firstly the laser system required is rather difficult to use and secondly there is no metastable level for use as an upper qubit state.

The first disadvantage, namely the laser system, might be remedied in the future in one of several ways. H. Powell looked at possibility of generating the light at 280 nm with a third-harmonic generation system [105], thus avoiding the need for a dye laser. Alternatively, it may be that narrow-band diode lasers become available near that wavelength in the future. Finally, it would be possible to continue using a dye laser as before but with removal of the BBO crystal from the laser cavity and the use, instead, of an external frequency-doubling cavity. This would greatly ease the use of the dye laser and would produce more power at the second harmonic. Pumping the dye laser with a solid state laser such as an ELS VersaDisk disc laser would also be beneficial as the beam-pointing stability of the argon-ion laser was found to be one of the major hindrances in the reliable operation of the dye laser.

Although no metastable levels in Mg^+ are available for the qubit transition, the possibility of using the hyperfine structure of $^{25}Mg^+$ exists. Most groups at the forefront of QIP with trapped ions are already using species with hyperfine structure or are moving towards such a species. Examples of these include $^{9}Be^+$ in the group of D. Wineland and $^{43}Ca^+$ in the groups of R. Blatt and A. Steane.

A simple yet effective technique to measure the damping rate of the ions' motions was employed. This RF–photon correlation showed that the damping rate of the magnetron motion was increased by more than a factor of five by the axialisation of magnesium ions. The same technique applied to calcium ions showed that the unaxialised damping rate was very much lower (about 3 Hz) than that of magnesium (around 60 Hz). A difference between the two is certainly not unexpected since the cooling transition is broader in magnesium than in calcium and the presence of metastable D states in calcium will reduce the cooling rate further. The scale of the difference was larger than expected,

however. This may indicate the presence of some laser-induced "trapped" states in the calcium cooling scheme. The axialised damping rate of the magnetron motion was also correspondingly lower in calcium than in magnesium.

An important observation made during the measurement of the damping rates was that there was a frequency shift of the motional frequency as the axialisation-drive amplitude was changed. This linear shift was found to be due to the coupling of the two motions and the magnitude of the shift was an indication of the coupling rate between the two motions. Similar shifts had been observed in experiments by Cornell *et al.* [102] and Häffner *et al.* [103]. Those experiments measured a coupling rate between the axial and one of the radial motions and the coupling rates were found to be tens of Hertz. By contrast, the coupling rate achieved in our axialisation experiments was of order 1 kHz. The lower coupling rate in Ref. [102] was partly by design and in both cases was largely due to the non-ideal excitation geometry. Such a high coupling rate as we have demonstrated is almost certainly unnecessary but it indicates that efficient coupling is taking place.

Although firm conclusions can not be reached at this time, preliminary indications are that the optimal axialisation-drive amplitude is relatively small. Low aspect ratios of the cloud as measured from ICCD images along with high damping rates measured by RF–photon correlation suggest that a drive amplitude of around 50 to 200 mV may be optimal for the sizes of cloud loaded (about 20 ions). With more stable lasers it will be possible to verify that the damping rate is indeed highest at these low drive amplitudes.

The last set of experiments reported in the thesis are some which were conducted in the University of Innsbruck. The experience gained there of Pound–Drever–Hall locking has proved very useful as such a lock continues to be developed in the group in Imperial. Suggestions that have arisen for improving the lock here include a low-bandwidth feedback from the external lock to the piezo scan control of the laser itself to provide a better level of frequency-stability than provided by the laser itself.

10.1 Future Directions

The two remaining intermediate goals in the project to measure decoherence in the Penning trap are the trapping of a single ion of calcium and probing the 729-nm ($S_{1/2}$–$D_{5/2}$) transition of the single ion with light from an ultra-

stable Ti:Sapphire laser. Once the spectrum of sidebands on this transition has been measured, the use of the same transition for sideband cooling may be considered.

It was projected that an axialising field should greatly ease the loading of very small clouds and individual ions in the Penning trap: the damping of the unbound magnetron motion should afford a much greater level of stability and allow the positioning of the laser beam closer to the trap centre to yield higher damping rates of the modified cyclotron motion. Early experiments loading small clouds with and without the presence of the axialisation drive indicated that a great improvement could be obtained. This, coupled with the new, automated, spatial scanning of the beam using actuators on a mirror mount, should make the loading process much more reliable.

Single calcium ions have already been routinely trapped in a Paul trap configuration. Signal levels from the ion were around 4 kHz. Analysis of quantum jumps seen on the fluorescence output suggests that the smallest clouds loaded thus far in the Penning trap consisted of as few as two or three ions. The signal level per ion for such small clouds in the Penning trap was much lower, however, than for the single ion in a Paul trap. This suggests that there may be a serious obstacle to overcome before high signal levels in the Penning trap can be achieved. Such an obstacle may be the existence of trapped states arising from the coherent interaction of some of the laser beams with the atomic levels. Such a problem would not have been encountered before as the only other ion species to have been cooled in a Penning trap do not have as complicated a level structure. Possible ways of removing this hurdle are to use some less-coherent light, most likely on the 866-nm transition. Such light might be generated by a broadband diode laser, or even by an LED in addition to the existing repumper lasers.

Assuming these technical hurdles can be overcome and single ions can be trapped routinely, the next stage will be to probe the 729-nm transition. The narrow linewidth of the transition (< 1 Hz) means that substantial power and a low laser linewidth will be required. Both of these requirements will be satisfied by the Ti:Sapphire system which is currently being stabilised using a Pound–Drever–Hall lock. The new wavemeter's accuracy will be essential in tuning the laser to the atomic transition. Once tuned to the transition, the laser can be used to perform spectroscopy and the motional sidebands should be apparent. Such sidebands will be significantly more difficult to resolve where the magnetic

field is provided by an electromagnet than by a superconducting magnet. The much-improved stability (three orders of magnitude) of the superconducting magnet will confer greater stability on the motional frequencies and thus make the sidebands easier to resolve. A trap is currently being prepared for use in a superconducting magnet and experiments with this trap should begin soon.

The work described in this thesis has been presented at several conferences. A poster describing the magnesium axialisation results was shown at the Trapped Charged Particles and Fundamental Interactions Conference 2002 and a refereed article on the subject was published [1]. I have made several oral presentations, including one at a conference organised by the EU research network, QUEST, for young researchers, as well as seminars both in Britain and Ireland. The calcium axialisation results will be presented in poster format at Physics 2005, a major international conference. A paper describing those results is also in preparation [2].

Appendix A

Octave code for cavity mode matching

```
% coupling.m
%
% Octave script to work out the appropriate focal lengths of a pair
% of lenses for the Fabry Perot stabilising cavity.
%
% The focal lengths are chosen and the beam waist is calculated for various
% lens-lens and lens-cavity separations
%
% Eoin Phillips, Ion Trapping Group, Imperial College, London
% 25th June 2002
% 5th September 2002
% 2nd June 2003
% 27th June 2003
%

clear;
clg;

% Define constants
L = 0.15;          % Length of cavity
f1 = 0.3;          % Focal length of positive lens
f2 = -0.2;         % Focal length of negative lens
```

```
d1 = 0.05;          % Distance between two lenses
d2 = 0.1;           % Distance between second lens and cavity
d1max = 0.3;        % Max separation of lenses in metres
d2max = 0.5;        % Max separation of lenses and cavity in metres
d1step = 0.002;
d2step = 0.01;
nz = 1.5;           % Refractive index of mirror
na = 1.0;           % Refractive index of air
d3 = 0.01;          % Thickness of mirror
R = 0.5;            % Radius of curvature of mirror
lambda = 786e-9;    % Wavelength of laser


% Positive lens
M1 = [1 0; -1/f1 1];

% Inter-lens separation
M2 = [1 d1; 0 1];

% Negative lens
M3 = [1 0; -1/f2 1];

% Lens -> mirror - to be redefined in calculation
M4 = [1 d2; 0 1];

% Air -> mirror
M5 = [1 0; 0 na/nz];

% Traverse mirror
M6 = [1 d3; 0 1];

% Mirror -> cavity
M7 = [1 0; (nz - na)/(na * R) nz/na];

% Traverse half of cavity
M8 = [1 L/2; 0 1];
```

```
                                                                    60

% Minimum spot size of beam in cavity, see Milonni & Eberly, p. 498
wcav = sqrt(lambda / (2 * pi)) * (2*R*L - L^2)^0.25;

% Waist of laser beam
wbeam = 1.5e-3;
qbeam = -i * pi * wbeam^2 / lambda;



                                                                    70
% Now calculate q over a range of lens-lens and lens-cavity separations
k = 1;

for d1 = d1step:d1step:d1max
   M2 = [1 d1; 0 1];
   l = 1;
   for d2 = d2step:d2step:d2max
      M4 = [1 d2; 0 1];
      M = M8 * M7 * M6 * M5 * M4 * M3 * M2 * M1;
      qnew = M * [qbeam; 1];                                        80
      qstorinv(k,l) = qnew(2) / qnew(1); % i.e. 1/q                 qstorinv
      d2stor(l) = d2;
      l++;
   end
   d1stor(k) = d1;
   k++;
end

qstorinv = imag(qstorinv);
qstor = qstorinv .^(-1);                                            90

wstor = sqrt(lambda / pi * qstor);
```

```
for kiter = 1:k-1
        for liter = 1:l-1
                wplot(kiter,liter) = wcav;
        end
end

clearplot;
ylabel("Distance between lenses");
xlabel("Distance to cavity");
zlabel("Waist size");
mesh (d2stor, d1stor, wstor);
hold on;
mesh (d2stor, d1stor, wplot);
gset view 0,0;
gset logscale z;
replot;
```

Appendix B

Split-ring trap

The following figures are the original technical designs for the split-ring trap by Martijn van Eijkelenborg. Some modifications have been made to the trap since it was first designed. Most notably, one of the two holes in the sides of the ring-electrode frame has been enlarged to reduce laser scatter and a blind hole has been drilled through the bottom of the ring electrode (i.e. cutting an arc in the sides of the lower ring segments and a hole through the bottom of the frame).

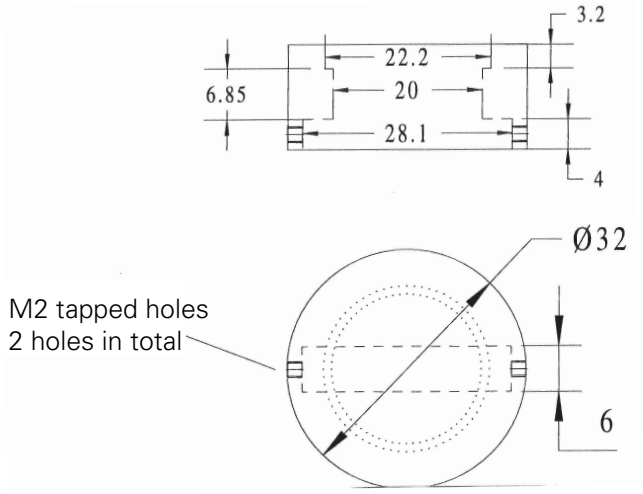

Figure B.1: The lens holder that sits on top of the ring-electrode frame. Four tapped M2 holes through the ring hold ceramic screws that support the oven.

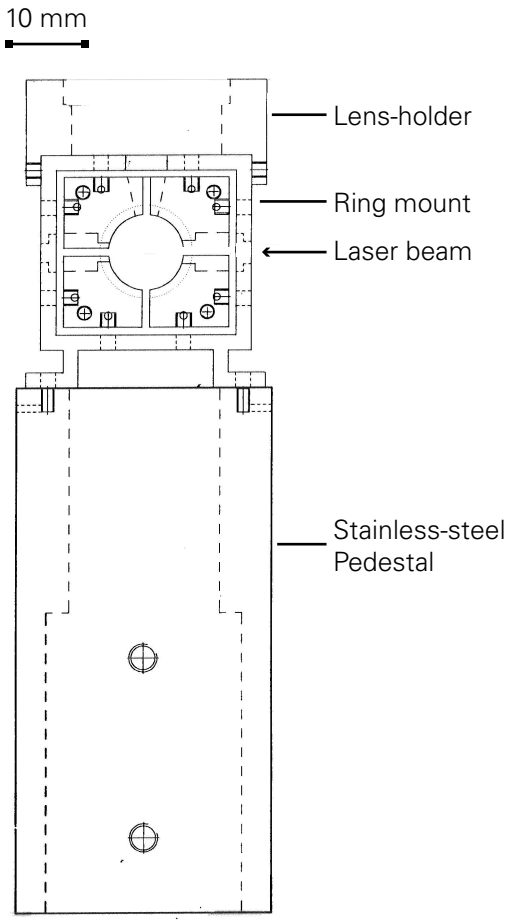

Figure B.2: The ring electrode in its holder, mounted on the stainless-steel pedestal.

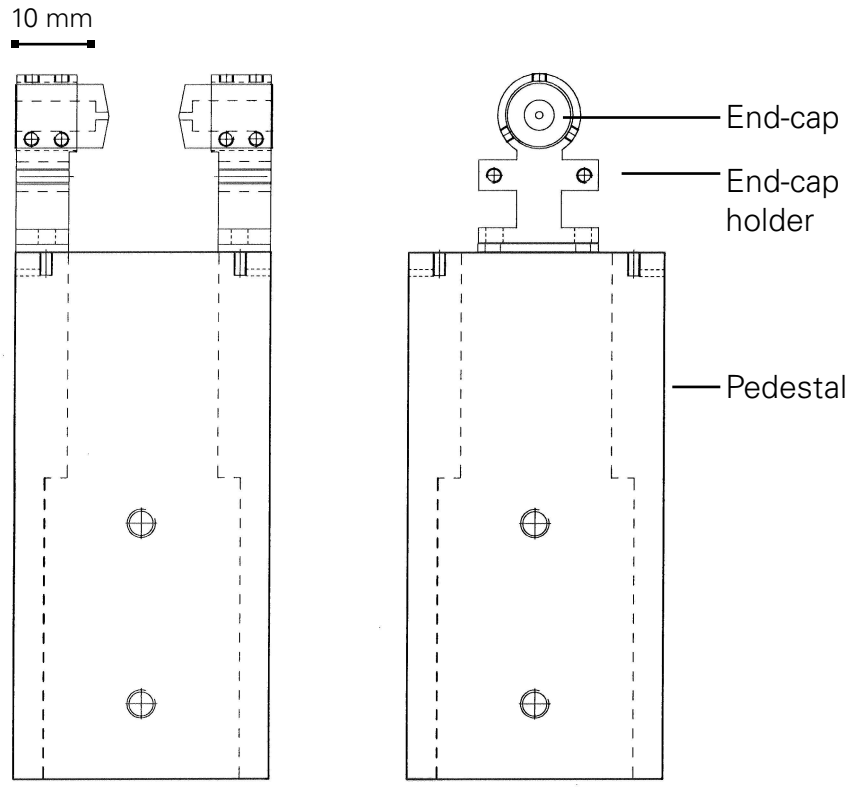

Figure B.3: The end-cap electrodes in their holders, mounted on the stainless-steel pedestal.

Bibliography

[1] H. F. Powell, S. R. de Echaniz, E. S. Phillips, D. M. Segal, R. C. Thompson, Improvement of laser cooling of ions in a Penning trap by use of the axialization technique, Journal Of Physics B 36 (2003) 961–970.

[2] E. S. Phillips, R. J. Hendricks, A. M. Abdulla, D. M. Segal, R. C. Thompson, in preparation (2005) .

[3] P. K. Ghosh, Ion Traps, Clarendon Press, Oxford, 1995.

[4] W. Paul, Electromagnetic traps for charged and neutral particles, Reviews Of Modern Physics 62 (1990) 531–540.

[5] Scientific Instrument Services Inc., SIMION 3D Ion and Electron Optics Software, http://www.sisweb.com/simion.htm.

[6] W. Paul, O. Osberghaus, E. Fischer, An ion cage, Forschungsberichte des Wirtschafts- und Verkehrsministeriums Nordrhein-Westfalen (1958) 42 pp.

[7] F. M. Penning, Glow discharge at low pressure between coaxial cylinders in an axial magnetic field, Physica 3 (1936) 873–894.

[8] J. R. Pierce, Theory and design of electron beams, Bell Telephone Laboratories series, Van Nostrand, New York, 1949.

[9] M. Kretzschmar, Single-particle motion in a Penning trap - description in the classical canonical formalism, Physica Scripta 46 (1992) 544–554.

[10] G. Bollen, R. B. Moore, G. Savard, H. Stolzenberg, The accuracy of heavy-ion mass measurements using time of flight-ion cyclotron-resonance in a Penning trap, Journal Of Applied Physics 68 (1990) 4355–4374.

[11] L. S. Brown, G. Gabrielse, Geonium theory - physics of a single electron or ion in a Penning trap, Reviews Of Modern Physics 58 (1986) 233–311.

[12] R. C. Thompson, D. C. Wilson, The motion of small numbers of ions in a Penning trap, Zeitschrift für Physik D 42 (4) (1997) 271–277.

[13] L. Schweikhard, A. G. Marshall, Excitation modes for Fourier transform-ion cyclotron-resonance mass-spectrometry, Journal Of The American Society For Mass Spectrometry 4 (1993) 433–452.

[14] G. Savard, S. Becker, G. Bollen, H. J. Kluge, R. B. Moore, T. Otto, L. Schweikhard, H. Stolzenberg, U. Wiess, A new cooling technique for heavy-ions in a Penning trap, Physics Letters A 158 (1991) 247–252.

[15] D. Wineland, H. Dehmelt, Line shifts and widths of axial, cyclotron and g-2 resonances in tailored, stored electron (ion) cloud, International Journal Of Mass Spectrometry And Ion Processes 16 (1975) 338–342.

[16] R. Frisch, Experimental demonstration of Einstein's radiation recoil, Zeitschrift für Physik 86 (1933) 42.

[17] T. W. Hänsch, A. L. Schawlow, Cooling of gases by laser radiation, Optics Communications 13 (1) (1975) 68–69.

[18] D. J. Wineland, H. Dehmelt, Proposed $10^{14}\Delta\nu < \nu$ laser fluorescence spectroscopy on Tl^+ mono-ion oscillator III (side band cooling), Bulletin of the American Physical Society 20 (1975) 637.

[19] D. J. Wineland, R. E. Drullinger, F. L. Walls, Radiation-pressure cooling of bound resonant absorbers, Physical Review Letters 40 (1978) 1639–1642.

[20] H. Walther, K. W. Rothe (Eds.), Laser Spectroscopy, Vol. IV, Springer, Berlin, 1979, Ch. Laser cooling of ions bound to a Penning trap, pp. 66–72.

[21] W. Neuhauser, M. Hohenstatt, P. Toschek, H. Dehmelt, Optical-sideband cooling of visible atom cloud confined in parabolic well, Physical Review Letters 41 (1978) 233–236.

[22] C. S. Adams, E. Riis, Laser cooling and trapping of neutral atoms, Progress in Quantum Electronics 21 (1997) 1–79.

[23] J. J. Bollinger, J. M. Kriesel, T. B. Mitchell, L. B. King, M. J. Jensen, W. M. Itano, D. H. E. Dubin, Laser-cooled ion plasmas in Penning traps, Journal Of Physics B 36 (2003) 499–510.

[24] K. Koo, J. Sudbery, D. M. Segal, R. C. Thompson, Doppler cooling of Ca^+ ions in a Penning trap, Physical Review A 69 (2004) art. no.–043402.

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

[26] G. Z. K. Horvath, R. C. Thompson, Laser cooling of ions stored in a Penning trap: A phase-space picture, Physical Review A 59 (1999) 4530–4546.

[27] R. C. Thompson, J. Papadimitriou, Simple model for the laser cooling of an ion in a Penning trap, Journal Of Physics B 33 (2000) 3393–3405.

[28] M. Lindberg, Steady-state of a laser-cooled trapped ion in the Lamb-Dicke limit, Journal Of Physics B 17 (1984) 2129–2139.

[29] H. Figger, D. Meschede, C. Zimmermann (Eds.), Laser Physics at the Limits, Springer, 2002, Ch. Conditional spin resonance with trapped ions, pp. 261–274.

[30] M. A. Nielsen, I. L. Chuang, Quantum Computation and Quantum Information, Cambridge University Press, Cambridge, 2001.

[31] D. Bouwmeester, A. Ekert, A. Zeilinger (Eds.), The Physics of Quantum Information, Springer, Berlin, 2000.

[32] P. Benioff, Quantum-mechanical models of Turing-machines that dissipate no energy, Physical Review Letters 48 (1982) 1581–1585.

[33] R. P. Feynman, Simulating physics with computers, International Journal Of Theoretical Physics 21 (1982) 467–488.

[34] D. Deutsch, Quantum theory, the Church-Turing principle and the universal quantum computer, Proceedings of the Royal Society of London A 400 (1818) (1985) 97–117.

[35] T. Sleator, H. Weinfurter, Realizable universal quantum logic gates, Physical Review Letters 74 (1995) 4087–4090.

[36] A. R. Calderbank, E. M. Rains, P. W. Shor, N. J. A. Sloane, Quantum-error correction and orthogonal geometry, Physical Review Letters 78 (1997) 405–408.

[37] A. Steane, Multiple-particle interference and quantum error correction, Proc. R. Soc. London Ser. A-Math. Phys. Eng. Sci. 452 (1996) 2551–2577.

[38] D. P. DiVincenzo, The physical implementation of quantum computation, Fortschritte Der Physik-Progress Of Physics 48 (2000) 771–783.

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

[40] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband, D. J. Wineland, Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate, Nature 422 (2003) 412–415.

[41] J. I. Cirac, P. Zoller, Quantum computations with cold trapped ions, Physical Review Letters 74 (1995) 4091–4094.

[42] A. Sørensen, K. Mølmer, Quantum computation with ions in thermal motion, Physical Review Letters 82 (9) (1999) 1971–1974.

[43] F. Mintert, C. Wunderlich, Ion-trap quantum logic using long-wavelength radiation, Physical Review Letters 87 (2001) 257904.

[44] C. Wunderlich, C. Balzer, T. Hannemann, F. Mintert, W. Neuhauser, D. Reiss, P. E. Toschek, Spin resonance with trapped ions, Journal Of Physics B 36 (2003) 1063–1072.

[45] D. McHugh, J. Twamley, A quantum computer using a trapped-ion spin molecule and microwave radiation, arXiv pre-print quant-ph/0310015.

[46] S. Mancini, A. M. Martins, P. Tombesi, Quantum logic with a single trapped electron, Physical Review A 61 (2000) 012303.

[47] G. Ciaramicoli, I. Marzoli, P. Tombesi, Scalable quantum processor with trapped electrons, Physical Review Letters 91 (2003) 017901.

[48] D. Kielpinski, C. Monroe, D. J. Wineland, Architecture for a large-scale ion-trap quantum computer, Nature 417 (2002) 709–711.

[49] S. Stahl, et al., A planar Penning trap, in preparation .

[50] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, D. J. Wineland, Demonstration of a fundamental quantum logic gate, Physical Review Letters 75 (1995) 4714–4717.

[51] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, I. C. Monroe, Experimental entanglement of four particles, Nature 404 (2000) 256–259.

[52] K. Molmer, A. Sorensen, Multiparticle entanglement of hot trapped ions, Physical Review Letters 82 (1999) 1835–1838.

[53] A. Sorensen, K. Molmer, Entanglement and quantum computation with ions in thermal motion, Physical Review A 6202 (2000) 022311.

[54] F. Schmidt-Kaler, H. Haffner, S. Gulde, M. Riebe, G. P. T. Lancaster, T. Deuschle, C. Becher, W. Hansel, J. Eschner, C. F. Roos, R. Blatt, How to realize a universal quantum gate with trapped ions, Applied Physics B 77 (2003) 789–796.

[55] D. Leibfried, M. D. Barrett, T. Schaetz, J. Britton, J. Chiaverini, W. M. Itano, J. D. Jost, C. Langer, D. J. Wineland, Toward Heisenberg-limited spectroscopy with multiparticle entangled states, Science 304 (2004) 1476–1478.

[56] C. F. Roos, M. Riebe, H. Haffner, W. Hansel, J. Benhelm, G. P. T. Lancaster, C. Becher, F. Schmidt-Kaler, R. Blatt, Control and measurement of three-qubit entangled states, Science 304 (2004) 1478–1480.

[57] M. Riebe, H. Haffner, C. F. Roos, W. Hansel, J. Benhelm, G. P. T. Lancaster, T. W. Korber, C. Becher, F. Schmidt-Kaler, D. F. V. James, R. Blatt, Deterministic quantum teleportation with atoms, Nature 429 (2004) 734–737.

[58] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, D. J. Wineland, Deterministic quantum teleportation of atomic qubits, Nature 429 (2004) 737–739.

[59] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, D. J. Wineland, Realization of quantum error correction, Nature 432 (2004) 602–605.

[60] D. G. Cory, M. D. Price, W. Maas, E. Knill, R. Laflamme, W. H. Zurek, T. F. Havel, S. S. Somaroo, Experimental quantum error correction, Physical Review Letters 81 (1998) 2152–2155.

[61] D. Leung, L. Vandersypen, X. L. Zhou, M. Sherwood, C. Yannoni, M. Kubinec, I. Chuang, Experimental realization of a two-bit phase damping quantum code, Physical Review A 60 (1999) 1924–1943.

[62] D. J. Wineland, C. Monroe, W. M. Itano, B. E. King, D. Leibfried, D. M. Meekhof, C. Myatt, C. Wood, Experimental primer on the trapped ion quantum computer, Fortschritte Der Physik-Progress Of Physics 46 (1998) 363–390.

[63] M. A. Rowe, A. Ben-Kish, B. Demarco, D. Leibfried, V. Meyer, J. Beall, J. Britton, J. Hughes, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband, D. J. Wineland, Transport of quantum states and separation of ions in a dual rf ion trap, Quantum Information & Computation 2 (2002) 257–271.

[64] N. Kjaergaard, L. Hornekaer, A. M. Thommesen, Z. Videsen, M. Drewsen, Isotope selective loading of an ion trap using resonance-enhanced two-photon ionization, Applied Physics B 71 (2000) 207–210.

[65] J. L. K. Koo, Laser cooling and trapping of Ca^+ ions in a Penning trap, Ph.D. thesis, Imperial College London (2003).

[66] J. Sudbery, Studies of laser cooled calcium ions in the Penning and combined traps, Ph.D. thesis, Imperial College London (2003).

[67] M. van Eijkelenborg, K. Dholakia, M. Storkey, D. Segal, R. Thompson, A driven, trapped, laser cooled ion cloud: a forced damped oscillator, Optics Communications 159 (1999) 169–176.

[68] M. E. M. Storkey, Studies of laser cooled trapped ions, Ph.D. thesis, Imperial College London (2001).

[69] J. F. O'Hanlon, A User's Guide to Vacuum Technology, Wiley Interscience, New York, 1980.

[70] G. P. T. Lancaster, H. Haffner, M. A. Wilson, C. Becher, J. Eschner, F. Schmidt-Kaler, R. Blatt, Doppler cooling a single Ca^+ ion with a violet extended-cavity diode laser, Applied Physics B 76 (2003) 805–808.

[71] P. Staanum, I. S. Jensen, R. G. Martinussen, D. Voigt, M. Drewsen, Lifetime measurement of the metastable $^2D_{5/2}$ state in the $^{40}Ca^+$ ion using the shelving technique on a few-ion string, Physical Review A 69 (2004) 32503–1–9.

[72] D. M. Lucas, C. J. S. Donald, J. P. Home, M. J. McDonnell, A. Ramos, D. N. Stacey, J. P. Stacey, A. M. Steane, S. C. Wester, Oxford ion-trap quantum computing project, Philosophical Transactions Of The Royal Society Of London Series A 361 (2003) 1401–1408.

[73] H. F. Powell, M. A. van Eijkelenborg, W. Irvine, D. M. Segal, R. C. Thompson, Quantum jumps in singly ionized magnesium, Journal Of Physics B 35 (2002) 205–216.

[74] A. S. Arnold, J. S. Wilson, M. G. Boshier, A simple extended-cavity diode laser, Review Of Scientific Instruments 69 (1998) 1236–1239.

[75] L. Ricci, M. Weidemuller, T. Esslinger, A. Hemmerich, C. Zimmermann, V. Vuletic, W. Konig, T. W. Hansch, A compact grating-stabilized diode-laser system for atomic physics, Optics Communications 117 (1995) 541–549.

[76] C. J. Myatt, N. R. Newbury, C. E. Wieman, Simplified atom trap by using direct microwave modulation of a diode-laser, Optics Letters 18 (1993) 649–651.

[77] M. J. Snadden, R. B. M. Clarke, E. Riis, Injection-locking technique for heterodyne optical phase locking of a diode laser, Optics Letters 22 (1997) 892–894.

[78] A. Kreuter, C. Becher, G. P. T. Lancaster, A. B. Mundt, C. Russo, H. Haffner, C. Roos, J. Eschner, F. Schmidt-Kaler, R. Blatt, Spontaneous emission lifetime of a single trapped Ca^+ ion in a high finesse cavity, Physical Review Letters 92 (2004) art. no.–203002.

[79] P. A. Barton, et al., Measurement of the lifetime of the $3d\,^2D_{5/2}$ state in $^{40}Ca^+$, Physical Review A 62 (2000) 32503–1–10.

[80] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, H. Ward, Laser phase and frequency stabilization using an optical resonator, Applied Physics B 31 (1983) 97–105.

[81] P. W. Milonni, J. H. Eberly, Lasers, Wiley-Interscience, New York, 1988.

[82] E. Hecht, Optics, 3rd Edition, Addison-Wesley, Reading, Massachusetts, 1998.

[83] A. Yariv, Optical Electronics in Modern Communications, Oxford University Press, 1997.

[84] The Octave programming language, http://www.octave.org/.

[85] K. B. MacAdam, A. Steinbach, C. Wieman, A narrow-band tunable diode-laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb, American Journal Of Physics 60 (1992) 1098–1111.

[86] V. Daneu, et al., Accurate laser wavelength measurements in the infrared and far infrared using a Michelson interferometer, Physics Letters 29A (6) (1969) 319–320.

[87] F. P. Schäfer (Ed.), Dye Lasers, 3rd Edition, Vol. 1 of Topics in Applied Physics, Springer-Verlag, New York, 1990, Ch. 6, pp. 201–219.

[88] P. J. Fox, R. E. Scholten, M. R. Walkiewicz, R. E. Drullinger, A reliable, compact, and low-cost Michelson wavemeter for laser wavelength measurement, American Journal Of Physics 67 (7) (1999) 624–630.

[89] S. M. Jaffee, W. M. Yen, A precision drive for a Michelson wavemeter, Review Of Scientific Instruments 64 (6) (1993) 1459–1462.

[90] X. Hui-Rong, S. V. Benson, T. W. Hänsch, A 'toy-train' wavemeter, Laser Focus 17 (1981) 54–58.

[91] National Institute of Standards and Technology, Boulder, CO, USA, NIST wavemeter manual.

[92] A. Kahane, et al., Vernier fringe-counting device for laser wavelength measurements, Review Of Scientific Instruments 54 (9) (1983) 1138–1142.

[93] G. Blasbichler, Ein Lambdameter mit 10^{-7} Meßunsicherheit, Master's thesis, University of Innsbruck (2000).

[94] R. Balhorn, H. Kunzmann, F. Lebowsky, Frequency stabilization of internal-mirror helium-neon lasers, Applied Optics 11 (1972) 742–744.

[95] P. E. Ciddor, Refractive index of air: New equations for the visible and near infrared, Applied Optics 35 (1996) 1566–1573.

[96] T. Doiron, http://emtoolbox.nist.gov/wavelength/ciddor.asp.

[97] D. J. Heinzen, J. J. Bollinger, F. L. Moore, W. M. Itano, D. J. Wineland, Rotational equilibria and low-order modes of a nonneutral ion plasma, Physical Review Letters 66 (1991) 2080–2083.

[98] X. P. Huang, F. Anderegg, E. M. Hollmann, C. F. Driscoll, T. M. O'Neil, Steady-state confinement of non-neutral plasmas by rotating electric fields, Physical Review Letters 78 (1997) 875–878.

[99] M. B. Comisarow, A. G. Marshall, Fourier transform ion cyclotron resonance spectroscopy, Chemical Physics Letters 25 (1974) 282–283.

[100] A. Perlmutter, L. F. Scott (Eds.), New Frontiers in High-Energy Physics, Vol. 14 of Studies in the natural sciences, Plenum Press, New York, 1978, Ch. Electron magnetic moment from geonium spectra, pp. 159–181.

[101] S. H. Guan, X. Z. Xiang, A. G. Marshall, Axial and radial ion cloud compression - coupling of magnetron and cyclotron motion to axial motion in a segmented cubic Fourier-transform ion-cyclotron resonance ion trap, International Journal Of Mass Spectrometry And Ion Processes 124 (1993) 53–67.

[102] E. A. Cornell, R. M. Weisskoff, K. R. Boyce, D. E. Pritchard, Mode-coupling in a Penning trap: π-pulses and a classical avoided crossing, Physical Review A 41 (1990) 312–315.

[103] H. Haffner, T. Beier, S. Djekic, N. Hermanspahn, H. J. Kluge, W. Quint, S. Stahl, J. Verdu, T. Valenzuela, G. Werth, Double Penning trap technique

for precise g factor determinations in highly charged ions, European Physical Journal D 22 (2003) 163–182.

[104] G. Graff, H. Kalinowsky, J. Traut, A direct determination of the proton electron mass ratio, Zeitschrift für Physik A 297 (1980) 35–9.

[105] H. F. Powell, Quantum optics of a single ion, Ph.D. thesis, Imperial College London (2002).

[106] M. Alonso, E. J. Finn, Physics, Addison Wesley Longman Limited, Harlow, England, 1996.

[107] B. W. Shore, The Theory of Coherent Atomic Excitation, Vol. 1, Simple Atoms and Fields, Wiley-Interscience, New York, 2000.

[108] S. Durr, T. Volz, A. Marte, G. Rempe, Observation of molecules produced from a bose-einstein condensate, Physical Review Letters 92 (2004) art. no.–020406.

[109] H. F. Powell, D. M. Segal, R. C. Thompson, Axialization of laser cooled magnesium ions in a Penning trap, Physical Review Letters 89 (2002) 93003.

[110] F. Diedrich, E. Peik, J. M. Chen, W. Quint, H. Walther, Observation of a phase-transition of stored laser-cooled ions, Physical Review Letters 59 (1987) 2931–2934.

[111] M. D. Barrett, B. DeMarco, T. Schaetz, V. Meyer, D. Leibfried, J. Britton, J. Chiaverini, W. M. Itano, B. Jelenkovic, J. D. Jost, C. Langer, T. Rosenband, D. J. Wineland, Sympathetic cooling of $^{9}Be^{+}$ and $^{24}Mg^{+}$ for quantum logic, Physical Review A 68 (2003) art. no.–042302.

[112] D. N. Madsen, S. Balslev, M. Drewsen, N. Kjaergaard, Z. Videsen, J. W. Thomsen, Measurements on photo-ionization of 3s3p p-1(1) magnesium atoms, Journal Of Physics B 33 (2000) 4981–4988.

[113] T. Hasegawa, T. Shimizu, Resonant oscillation modes of sympathetically cooled ions in a radio-frequency trap, Physical Review A 66 (2002) art. no.–063404.

[114] D. J. Wineland, M. Barrett, J. Britton, J. Chiaverini, B. DeMarco, W. M. Itano, B. Jelenkovic, C. Langer, D. Leibfried, V. Meyer, T. Rosenband, T. Schatz, Quantum information processing with trapped ions, Philosophical Transactions Of The Royal Society London Series A-Mathematical Physical Engineering Sciences 361 (2003) 1349–1361.

[115] Y. L. Huang, G. Z. Li, S. H. Guan, A. G. Marshall, A combined linear ion trap for mass spectrometry, Journal Of The American Society For Mass Spectrometry 8 (1997) 962–969.

[116] J. P. Marangos, Topical review: electromagnetically induced transparency, Journal Of Modern Optics 45 (1998) 471–503.

[117] A. Kuhn, M. Hennrich, G. Rempe, Deterministic single-photon source for distributed quantum networking, Physical Review Letters 89 (2002) art. no.–067901.

[118] J. McKeever, A. Boca, A. D. Boozer, J. R. Buck, H. J. Kimble, Experimental realization of a one-atom laser in the regime of strong coupling, Nature 425 (2003) 268–271.

[119] A. B. Mundt, A. Kreuter, C. Russo, C. Becher, D. Leibfried, J. Eschner, F. Schmidt-Kaler, R. Blatt, Coherent coupling of a single $^{40}Ca^{+}$ ion to a high-finesse optical cavity, Applied Physics B 76 (2003) 117–124.

[120] C. Maurer, C. Becher, C. Russo, J. Eschner, R. Blatt, A single-photon source based on a single Ca^{+} ion, New Journal Of Physics 6.

[121] C. Simon, W. T. M. Irvine, Robust long-distance entanglement and a loophole-free Bell test with ions and photons, Physical Review Letters 91 (2003) art. no.–110405.

[122] N. Lutkenhaus, Security against individual attacks for realistic quantum key distribution, Physical Review A 6105 (2000) art. no.–052304.

[123] C. H. Bennett, G. Brassard, Quantum cryptography: Public key distribution and coin tossing, in: Proceedings of IEEE International Conference on Computers, Systems and Signal Processing, IEEE, New York, 1984, pp. 175–9.

[124] E. Knill, R. Laflamme, G. J. Milburn, A scheme for efficient quantum computation with linear optics, Nature 409 (2001) 46–52.

[125] J. Kim, O. Benson, H. Kan, Y. Yamamoto, A single-photon turnstile device, Nature 397 (1999) 500–503.

[126] J. I. Cirac, P. Zoller, H. J. Kimble, H. Mabuchi, Quantum state transfer and entanglement distribution among distant nodes in a quantum network, Physical Review Letters 78 (1997) 3221–3224.

[127] C. K. Law, H. J. Kimble, Deterministic generation of a bit-stream of single-photon pulses, Journal Of Modern Optics 44 (1997) 2067–2074.

[128] S. J. vanEnk, J. I. Cirac, P. Zoller, Ideal quantum communication over noisy channels: A quantum optical implementation, Physical Review Letters 78 (1997) 4293–4296.

[129] K. Bergmann, H. Theuer, B. W. Shore, Coherent population transfer among quantum states of atoms and molecules, Reviews Of Modern Physics 70 (1998) 1003–1025.

[130] S. J. VanEnk, J. I. Cirac, P. Zoller, H. J. Kimble, H. Mabuchi, Quantum state transfer in a quantum network: a quantum-optical implementation, Journal Of Modern Optics 44 (1997) 1727–1736.

[131] A. Kuhn, M. Hennrich, T. Bondo, G. Rempe, Controlled generation of single photons from a strongly coupled atom-cavity system, Applied Physics B 69 (1999) 373–377.

[132] M. Hennrich, T. Legero, A. Kuhn, G. Rempe, Vacuum-stimulated Raman scattering based on adiabatic passage in a high-finesse optical cavity, Physical Review Letters 85 (2000) 4872–4875.

[133] J. McKeever, J. R. Buck, A. D. Boozer, H. J. Kimble, Determination of the number of atoms trapped in an optical cavity, Physical Review Letters 93 (14) (2004) art. no.–143601.

[134] J. McKeever, A. Boca, A. D. Boozer, R. Miller, J. R. Buck, A. Kuzmich, H. J. Kimble, Deterministic generation of single photons from one atom trapped in a cavity, Science 303 (2004) 1992–1994.

[135] G. R. Guthohrlein, M. Keller, K. Hayasaka, W. Lange, H. Walther, A single ion as a nanoscopic probe of an optical field, Nature 414 (2001) 49–51.

[136] M. Keller, B. Lange, K. Hayasaka, W. Lange, H. Walther, A calcium ion in a cavity as a controlled single-photon source, New Journal Of Physics 6.

[137] M. Keller, B. Lange, K. Hayasaka, W. Lange, H. Walther, Continuous generation of single photons with controlled waveform in an ion-trap cavity system, Nature 431 (2004) 1075–1078.

[138] M. W. Hamilton, An introduction to stabilized lasers, Contemporary Physics 30 (1989) 21–33.

[139] K. Razdan, D. A. Van Baak, Demonstrating optical saturation and velocity selection in rubidium vapor, American Journal Of Physics 67 (1999) 832–836.

[140] T. W. Hansch, B. Couillaud, Laser frequency stabilization by polarization spectroscopy of a reflecting reference cavity, Optics Communications 35 (1980) 441–444.

[141] R. V. Pound, Electronic frequency stabilization of microwave oscillators, Review Of Scientific Instruments 17 (11) (1946) 490–505.

[142] E. D. Black, An introduction to Pound-Drever-Hall laser frequency stabilization, American Journal Of Physics 69 (2001) 79–87.

[143] E. Goulielmakis, M. Uiberacker, R. Kienberger, A. Baltuska, V. Yakovlev, A. Scrinzi, T. Westerwalbesloh, U. Kleineberg, U. Heinzmann, M. Drescher, F. Krausz, Direct measurement of light waves, Science 305 (2004) 1267–1269.