Imperial College
London
100 years of living science
100

University of London
Imperial College of Science, Technology and Medicine
Department of Physics

Single Ca$^+$ Ions in a Penning Trap for Applications in Quantum Information Processing

Hamid Ohadi

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics of the University of London and
the Diploma of Imperial College, January 2008

Abstract

I report on work aimed towards realising Quantum Information Processing with single $^{40}\mathrm{Ca}^{+}$ ions in a Penning trap, where trapping is achieved using static electric and magnetic fields. Unlike RF traps, in the Penning trap there is no heating due to micromotion. Therefore the qubit should in principle have a longer decoherence time. However in the Penning trap, the internal levels of the ion shift and split due to the magnetic field. As a result, one needs to implement additional lasers to address the split levels. In the experiment reported here two cooling lasers at 397nm and five repumping lasers at 866nm and 850nm are used. An additional Ti:Sapphire laser has been developed to address the qubit transition at 729nm. Because the natural linewidth of this transition is extremely narrow, the laser needs to be very stable in frequency. For this reason the laser is locked to a stable, high finesse cavity ($\mathcal{F} \sim 250,000$) using the Pound-Drever-Hall technique. This brings the stability of the laser from $\sim$ 3MHz down to 3kHz sufficiently low for the experiments we intend to do in the near future, for example sideband cooling. Previous attempts in our group to trap a single ion in the Penning trap failed due to an unexpectedly low fluorescence count rate per ion. This led us to simulate the cooling dynamics to investigate the possibility that “trapped states” might be the cause. Although there were trapped states in our scheme we were able to show that it was possible to avoid them by appropriately tuning the laser frequencies. Further experimental work showed that the low count rate was due to low laser power in the 866nm repumping lasers. Once this had been addressed we were able to trap and Doppler cool a single $^{40}\mathrm{Ca}^{+}$ ion in a Penning trap for the first time. The ion is further cooled using axialisation.

Acknowledgements

I wish to thank my supervisors Dr Danny Segal and Prof. Richard Thompson. They are the most friendly enjoyable supervisors. Not only they helped me with the research but they also provided wisdom, encouragement and friendship. I would like to thank all my colleagues all these years: Bakry, Danyal, Eoin, Rich, Rafael and Dan. They made a sad atmosphere of our basement lab into a pleasurable and fun place to work.

For technical assistance I would like to thank Brian Willey from the mechanical workshop and Bandu Ratnasekara from the Electronics section. Their offices were always open to my continuous requests.

Thanks to my flatmates Ehsan and Omid for being such a great friends. My special thank to Parizad who made the last year of my research the most enjoyable year of my life.

Finally but not the least I would like to thank my parents who made this research possible.

Contents

List of Tables

List of Figures

CHAPTER 1

Introduction

Information processing is fundamentally a physical process and physical laws determine how information can be stored or processed. Information "technology" has changed as progressively faster realisations have become available; from gears to relays, from valves to transistors and to integrated circuits. The modern information processing revolution, that started back in the 1940s by building the first computers, was realised by encoding and controlling information using the classical laws of Physics. According to Moore's law [1], the number of transistors that can be commercially placed in an integrated circuit doubles approximately every two years as the transistors get smaller. However there is a fundamental barrier for how small the size of a transistor could be i.e. a few tens of atoms. Classical computer processors will reach their lower limit of miniaturisation during the first few decades of the 21st century. Quantum mechanical effects will become dominant in the operation of logic gates as they become smaller and smaller. Logic gates fundamentally rely on classical Physics. As the scale of the components gets to the limit of a few atoms they become noisier, the error rate increases and eventually they become entirely inaccurate. Quantum computation (QC) is the theory of calculations at this scale. QC looks to turn these limitations into advantages by exploiting Quantum Mechanics.

Quantum Computation and Quantum Information form the interface between Computer Science and Quantum Mechanics. They constitute the study of information processing tasks (normally considered to be the province of computer scientists) that can be accomplished using quantum mechanical systems. The potential power of Quantum Mechanics to perform computation was first discovered by Richard Feynman. He noticed that it seems to be impossible to simulate the evolution of a quantum mechanical system on a classical computer in an efficient way [2]. Instead of viewing this issue as an obstacle he pointed out that we can use quantum systems to perform complex calculations reversibly.

However, implementing Quantum Information Processing (QIP) is not as straightforward as it may at first seem. Quantum mechanical systems are much harder to control due to the constant danger of interaction with the environment, irreversibly dissipating the information to the environment (a process which is known as *decoherence*). One may ask, 'if quantum technology is so difficult to deal with why not stay within the boundaries set by classical Physics? Why not make a network of classical computers that work in parallel instead? Why not build processors that have multiple cores instead of having just a single core?' The answer is that QIP has much more to offer than just increasing the clock-speed or parallelism of conventional microprocessors. It turns out that QIP has awesome potential capabilities that have no classical analogue. Quantum computers are theoretically *much* faster and more efficient at specific important calculations. So much so that certain tasks that are effectively impossible to perform on a classical computer are tractable using a QC. This is due to the fact that quantum mechanical systems can be in a *superposition of states* which allows one to perform parallel operations in a natural way. Moreover there are protocols in quantum cryptography that instead of relying on numerical complexities (as classical cryptography does) are based on physical laws to assure their utmost security. The basic principles of QIP are introduced in chapter 2.

There are seven necessary (and sufficient) requirements, which are called the DiVincenzo criteria [3], that any candidate for QIP needs to fulfill. Fulfilment of all these requirements at the same time is not an experimentally easy task. The physical system which stores the information (qubit) needs to be isolated from the environment otherwise the information will be lost by the interaction with the environment. On the other hand the system also cannot be completely isolated because one needs to process the information and eventually interact with the system. Processing the information introduces inevitable decoherence into the system. The candidate qubit needs to be relatively robust against decoherence caused by the environment or the observer. Furthermore to process the information "quantum logic gates" need to be realised and ideally, to do sensible calculations, a large scale quantum computer must be built. There need to be protocols for expanding the system to large numbers of qubits.

There are few candidates which are suitable for the realisation of QIP and *ion trapping* seems to be one of the most promising. In chapter 3 we discuss how ion trapping fulfils all the requirements of a serious candidate. The chapter ends with the current status, the obstacles met and future prospects of ion trap QIP.

There are essentially two kinds of ion traps: the *RF (Paul) trap* and the *Penning trap*. The geometry of both traps is similar: they consist of two opposing endcap electrodes plus a ring electrode in between. In the RF trap ion confinement is achieved using a time-varying potential at radio frequency which is applied to the ring or the endcap electrodes. The motion of an ion in the trap is the superposition of a slow secular motion and a fast "micromotion" at the RF drive frequency. The micromotion, which is unwanted, is greater as the distance from the trap centre increases. The majority of ion trapping groups in the world have decided to implement the RF trap and two leading groups in ion trapping QIP have already realised the essential building blocks of QIP [4, 5]. As explained in chapter 4, laser cooling in the RF trap is, in principle, much easier than the Penning trap. However the RF trap has also disadvantages which result from the micromotion and the relatively strong effects of "patch potentials" formed on the electrodes (see chapter 4). These problems are worse in smaller traps.

In the ion trapping group at Imperial College we are trying to realise QIP in the Penning trap with $^{40}\mathrm{Ca}^{+}$ ions. In the Penning trap instead of using an RF potential a static magnetic field is imposed along the axis which forces the ion to 'drift' around the axis in a "magnetron" motion. Therefore, in principle, the ion should not suffer from the heating and decoherence caused by the micromotion and patch potentials. In chapter 4 the basic principles of ion trapping and laser cooling in the Penning trap are described. The chapter finishes by introducing 'Pad traps', a new Penning trap design suitable for QIP with many ions, which we have recently proposed [6].

The status of the experiment when I started my PhD is explained in chapter 5. When I first joined the group, before performing any experiment with trapped ions, I worked on improving the performance of a Ti:Sapphire laser. The laser is needed to drive the qubit transition at 729nm. The laser frequency was stabilised to a linewidth of ∼1.7kHz, using the Pound-Drever-Hall locking technique, to an external high finesse cavity. The Ti:Sapphire laser stabilisation is explained in detail in chapter 6.

Up to this point trapping and laser cooling of clouds of Ca^{+} ions in a Penning trap had been achieved in our group. We were also able to work with single ions when operating our trap as an RF trap. However all previous attempts to trap single Ca^+ ions in the Penning trap were unsuccessful due to a low signal count rate per ion. In chapter 7, I investigated the possibility that the inefficient laser cooling of Ca^+ ions in the Penning trap could be due to “coherent population trapping”. It was found that an ideal excited state population of 0.25 could be achieved just as in the case of the RF trap, ruling this out as a possible cause for the low signal rate per ion.

Certain improvements to the experiment resulted in trapping and laser cooling of single Ca^+ ions in the Penning trap for the first time in the last months of my research. We also managed to see two-ion crystals along the axis and in the radial plane [7]. Improvements to the experiment that led to the observation of individual ions are explained in chapter 8.

Although ion trap QIP is well established, the use of ions in the Penning trap for this purpose is much less advanced. However, the eventual benefits make the pursuit of this goal highly worthwhile. Chapter 9 draws conclusions from the work presented here and proposes suggestions for further work.

CHAPTER 2

Introduction to QIP

2.1 Introduction

The foundations of quantum computation were laid down by David Deutsch in a paper in which he described a *universal quantum computer* [8]. Since then a lot of research has been done by theoretical scientists to find other interesting things that quantum computers can do. Shor's algorithms [9] for finding the prime factors of a large integer and the "discrete logarithm", Grover's search algorithm [10] and simulating quantum mechanical systems are among those areas which classical computers have no efficient role. Similarly Quantum Mechanics has great potential in cryptography which is not available in classical protocols. The two most common types of cryptography are "private key" and "public key". With "quantum cryptography" it is possible to distribute private keys with complete security [11]. Moreover quantum computers can break public keys because current classical cryptography algorithms rely on factoring large integers or discrete logarithm problems which are only difficult to solve for "classical" computers.

Controlling a quantum mechanical system experimentally in such a way that the information is well preserved and operations can be performed without introducing major errors is a very difficult task. By increasing the number of bits (qubits) and quantum operations that are performed, more sources of decoherence will be introduced into the system, the errors pile up and the computation goes off track. Some physicists are pessimistic about the future of QIP and argue that we are indeed quite far from building a large-scale quantum computer that is able to do sensible calculations [12]. Others take a different view contending that whether we can ever build a large-sale quantum computer is not really an issue at this stage. The point is to make the first steps and instead of simply understanding quantum mechanical phenomena to try to control and employ them for computing purposes. Moreover QIP has stimulated the discovery of new areas of Physics and a new understanding of Quantum Mechanics.

2.2 Basics of Quantum Computation

2.2.1 The Quantum Bit (Qubit)

The qubit, the analogue to a classical bit, is the basic unit of quantum information. It is the general quantum state of a physical system which has two possible eigenstates of a measurement operator represented by $|0\rangle$ and $|1\rangle$. Such a system is capable of storing binary information. However unlike its classical counterpart a qubit can exist in a *superposition* of these two states and form another state

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} , \tag{2.1}$$

where α and β are probability amplitudes which can be complex numbers. When this qubit is measured we will find either the $|0\rangle$ state, with the probability $|\alpha|^2$ or $|1\rangle$, with the probability $|\beta|^2$ (assuming the state $|\psi\rangle$ is normalised). Equation 2.1 represents a *superposition* state. A superposition state is different from just a mixture of states. A mixed state applies to an ensemble prepared by statistically combining two different states [*]. The superposition state can always be expressed as a pure state in a well chosen basis while for a mixed state this is not possible. Many unique features of QIP such as entanglement, as we discuss later on, are the result of the underlying principle

---

[*] The distinction can still be made with a single particle repeatedly prepared and then measured. The preparation can be coherent leading to a superposition state or the preparation can itself be a measurement leading to a mixed state when a histogram of results is obtained.

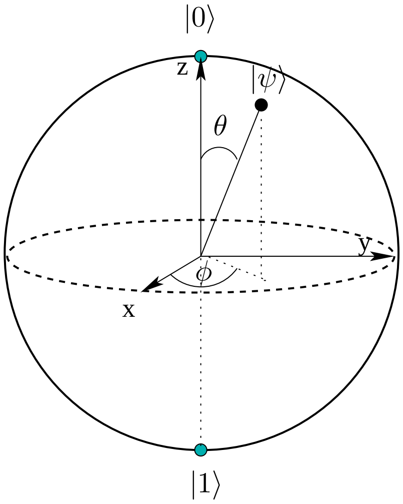

Figure 2.1. Bloch sphere representation of a qubit.

of superposition.

The physical object carrying the information can be any two-level quantum system such as electronic or nuclear spin, a pair of atomic energy levels, photon polarisation, the state of a Josephson junction or a quantum dot etc. Different candidates suitable for QIP are described in section 2.3.

To be able to develop an independent mathematical model for QC a qubit is defined as:

> A quantum system whose state lies in a two dimensional Hilbert space. In other words a qubit is a normalised vector in the two dimensional complex vector space spanned by two orthogonal vectors $|0\rangle$ and $|1\rangle$.

Because the state is normalised we have $|\alpha|^2 + |\beta|^2 = 1$ and the state can be rewritten as

$$|\Psi\rangle = e^{i\gamma}(\cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle).$$

Since $e^{i\gamma}$ is just an arbitrary phase factor we can ignore it and rewrite the state as

$$|\Psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle.$$

Here ϕ and θ are real numbers which define a point on a sphere called the *Bloch sphere* as shown in figure 2.1. The radius of the Bloch sphere is equal to 1. North pole of the Bloch sphere represents state $|0\rangle$ and the south pole represents $|1\rangle$.

A qubit can be any point on the surface of the Bloch sphere. It would therefore seem that it can hold infinite amount of information. However this is actually not the case because a qubit collapses to either $|0\rangle$ or $|1\rangle$ after being physically measured. An n-qubit register $\mathcal{Q}$ is generally defined to be a quantum state in $\otimes^n \mathcal{H}_i$ Hilbert space where $\mathcal{H}_i$ is the Hilbert space of a single qubit and $\otimes^n$ represents the product of the n Hilbert spaces.

The number of states that we could produce by n qubits is 2^n. If we then perform an operation (function) on these qubits the final answer will be the superposition of the values obtained by performing the operation on each initial state of the qubits. Therefore we could carry out 2^n operations by only applying one operation. This a phenomenon called *quantum parallelism* which gives the real power of QC. Taking a deeper look reveals that it is not as simple as that because after we make a measurement the superposition collapses to one state and we are left with an answer which is, at best, probabilistic. Nevertheless thanks to *quantum interference* it is possible to carefully arrange a set of quantum operations such that only the answer to the function interferes constructively [13].

Another major distinction from classical information that results from the superposition principle is *entanglement*. Suppose we have two qubits. Each one takes the values $|0\rangle$ or $|1\rangle$ and thus has a vector space spanned by vectors $\{|0\rangle, |1\rangle\}$. Now the Hilbert space of two qubits together is simply the multiplication of each individual qubit Hilbert space ie. $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. The state in general can be any superposition of these states

$$|\Psi\rangle = \alpha_{00}|00\rangle + \alpha_{01}|01\rangle + \alpha_{10}|10\rangle + \alpha_{11}|11\rangle,$$

where, as before, if the state is normalised $|\alpha_{ij}|^2$ is the probability of being in state $|ij\rangle$ where i and j are 0 or 1. Of particular interest are the states

$$\frac{|00\rangle + |11\rangle}{2}$$
$$\frac{|00\rangle - |11\rangle}{2}$$
$$\frac{|01\rangle + |10\rangle}{2}$$
$$\frac{|01\rangle - |10\rangle}{2}.$$

In these very important cases the two qubits are entangled in the sense that the

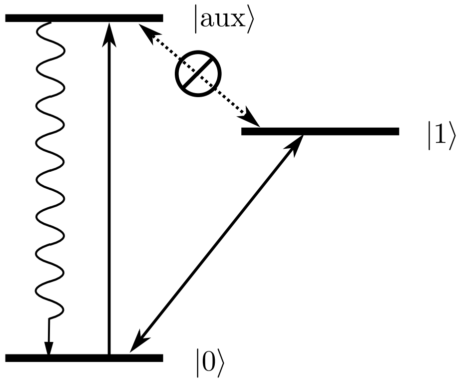

Figure 2.2. A laser driving the dipole transition $|0\rangle \rightarrow |\mathrm{aux}\rangle$ is illuminated at the ion. The ion scatters lots of photons if it is in the ground state $|g\rangle$ whereas it is dark if it is $|1\rangle$.

quantum state of one qubit should be described with reference to another. A measurement on one qubit determines the outcome of the measurement on the other qubit with a probability of 1. These particular entangled states are called the Bell states or EPR (Einstein-Podolsky-Rosen) states [14]. Entanglement is the key feature of quantum information processing and is used in every case where there is a great advantage in using QIP.

EPR states are some special cases of entangled states. A more rigorous definition of an entangled state is as follows [15]:

> Supposing q_1 and q_2 are two qubits with corresponding $\mathcal{H}_1$ and $\mathcal{H}_2$ Hilbert spaces, a state $|\psi\rangle$ in the Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$ is entangled if it cannot be written as
>
> $$|\psi_1\rangle \otimes |\psi_2\rangle.$$

2.2.2 Quantum Measurement

Unlike classical computation in which reading a bit does not affect its state, in QC we are faced with the problem that measurement of the quantum system generally affects the state of the qubit. Measurement causes our quantum mechanical system to jump into an eigenstate of the dynamical variable that is being measured.

If a normalised quantum state $|\psi\rangle$ is measured with respect to an observable $\mathcal{O}$ then the quantum state collapses to

$$\frac{\mathcal{O}|\psi\rangle}{|\mathcal{O}\psi|} ,$$

with the probability $|\mathcal{O}\psi|^2$. Suppose we operate a measurement such as $|0\rangle\langle 0|$ on a qubit with a state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$. This measurement determines whether the qubit state is $|0\rangle$ or not. The result of measurement will be 0 with the probability of $|\alpha|^2$ and the state after the measurement collapses to $|0\rangle$ with a phase factor. One experimental realisation of such a measurement is determining the atomic state of an ion in a trap. Suppose the simple case where there is only a ground state $|0\rangle$, an excited state $|1\rangle$ and an auxiliary state $|\mathrm{aux}\rangle$. The $|0\rangle \rightarrow |\mathrm{aux}\rangle$ is a fast dipole transition whereas $|\mathrm{aux}\rangle$ is not coupled to $|1\rangle$ directly (see figure 2.2). Shining a laser which is in resonance with the dipole transition scatters lots of photons provided that the ion is in the state $|0\rangle$ whereas it does not affect the ion if it is $|1\rangle$. Suppose that the ion is prepared in a coherent state of $(|0\rangle + |1\rangle)/\sqrt{2}$. A laser driving the dipole transition illuminates the ion. With the probability of 1/2, photons are scattered from the ion and the state of the ion collapses to $|0\rangle$. Now if the state of the ion is measured again, with the probability of 1, state $|0\rangle$ will be measured. This state detection technique is called *quantum jumps* and is explained more in section 3.2.

2.2.3 Quantum Gates

In a QC information is stored in qubits. In order to process the information, fundamental quantum gates similar to logic gates in classical computers need to be implemented. Quantum gates enable us to do calculations with qubits. A quantum gate is essentially an operator that converts a state in a vector space to another state in the same space. The difference with classical computation is that this operator cannot be just any operator because the state after the operation needs to be normalised again. It can be shown easily that to fulfil this condition the operator needs to be unitary i.e.

$$UU^\dagger = U^\dagger U = 1,$$

where $U^\dagger$ is the adjoint of U (transposing and then complex conjugating). This implies that any quantum gate needs to have the same number of outputs as it has for the input. As a result most classical logic gates like OR, AND, NAND,

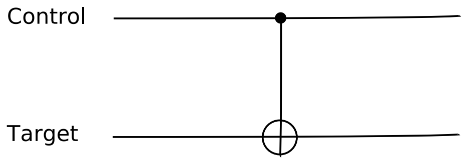

Figure 2.3. Two qubit C-NOT gate.

XOR, FANOUT etc. cannot be built in QC.

Another difference is that in QC a qubit cannot be duplicated. This is known as the *no-cloning theorem* [16]. A simple proof for this theorem is that if it were possible to clone a qubit state, one could clone an unknown system into as many copies as one needed and measure the dynamical variables that would not commute and hence break the uncertainty principle. For example one could measure the position of a particle and the momentum on a clone of that particle at the same time with infinite precision which is forbidden by the uncertainty principle. Moreover classical error correction is based on generating redundant information by cloning it in a certain way so that if the information is somehow damaged it could be detected and possibly recovered. Due to the no-cloning theorem, classical error correction techniques are not applicable. On the other hand it is essential to quantum cryptography as it prevents eavesdroppers from making a copy of the encryption key.

Similar to qubits which can be shown in vector form, quantum gates can be shown in matrix form. Some special single qubit gates we will use later are as follows. The X gate

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

which, like a classical NOT gate, flips the qubit. The Z gate

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

leaves $|0\rangle$ unchanged and flips the sign of $|1\rangle$ to give $-|1\rangle$ and finally the Y gate

$$Y = \begin{bmatrix} 0 & -\mathrm{i} \\ \mathrm{i} & 0 \end{bmatrix}.$$

X, Y and Z gates could be viewed as rotations by π along the x, y and z axis of the Bloch sphere respectively. X, Y and Z, which are in fact the famous Pauli matrices, are the generators of any rotation of the qubit about the x,

y and z axis and together with the identity I form an orthonormal basis for all one-qubit gates. Another useful class are the Hadamard, Phase and $\Pi/8$ gates. The Hadamard gate

$$H = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

converts $|0\rangle$ to $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and $|1\rangle$ to $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. The Hadamard gate can be thought as rotations of the qubit by $\pi/2$ about both the x and z axis. The phase gate can be written as

$$S = \begin{bmatrix} 1 & 0 \\ 0 & \mathrm{i} \end{bmatrix}$$

and $\Pi/8$ by

$$T = \begin{bmatrix} 1 & 0 \\ 0 & \exp(\mathrm{i}\pi/4) \end{bmatrix}.$$

It can be easily shown that $S^2 = Z$ and $T^2 = S$.

Another gate that acts on two qubits is the Controlled-NOT (C-NOT) gate which is a counterpart of the classical XOR gate. The C-NOT gate as shown in figure 2.3 acts on two qubits which are called the control qubit and target qubit. If the control qubit is 1, the target qubit is flipped otherwise it is left unchanged. The gate can be summarised as

$$|a\rangle|b\rangle \xrightarrow{\text{C}-\text{NOT}} |a\rangle|a \oplus b\rangle,$$

where $\oplus$ is addition modulo 2.

One can show [15] that every multi-qubit gate can be decomposed into two-qubit gates. Two-qubit gates can be built by C-NOT and single qubit gates. Moreover any single qubit gate can be generated by the set of Hadamard, phase gate and $\Pi/8$. As a result any quantum gate can be built by the set of the two-qubit C-NOT gate and single qubit gates of Hadamard, phase gate and $\Pi/8$ (see figure 2.4). This set of gates are called *universal quantum gates* [17].

2.2.4 Error Correction

QIP, due to its quantum nature, is far more susceptible to errors than classical information processing. A qubit suffers from decoherence unless it is truly

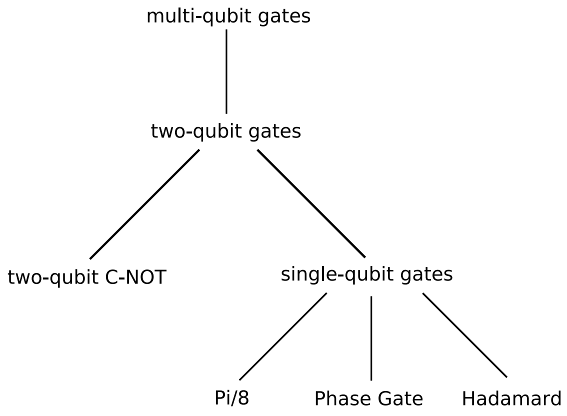

Figure 2.4. The tree of quantum gates. Any gate can be built by the set of Phase, Hadamard, Π/8 and C-NOT gates.

isolated from the environment which is an extremely difficult task experimentally. Decoherence is not the only source of error in the system. Quantum computers are fundamentally analogue-type devices. Single-qubit gates such as Π/8, Phase and Hadamard gates are inherently analogue operations. There will be inevitable inaccuracies in the quantum gate operations due to experimental imperfections. If a quantum computer lacks proper error correction protocols the computer will fail after a limited number of calculations.

Fortunately quantum error correction protocols do exist in principle [18, 19, 20] and a very simple error correcting code has been demonstrated experimentally [21]. Nevertheless this may not be sufficient to ensure our quantum computer works reliably. The error correction protocols require encoding the information and repeatedly recovering it. This itself induces errors in the system which need to be corrected. Therefore the error correction algorithm needs to be robust to the errors that are made in the process of error correction itself.

The errors in a quantum computer propagate throughout the system. For example as we saw in the previous section it is a key requirement that at least some gates operate on two or more qubits. If an error occurs in the first qubit it transfers to the second qubit. Therefore our algorithms must be designed in such a way that minimises the error propagation.

Error correction generally requires some redundancy of the information and therefore the number of elementary qubits increases. Moreover performing quantum gates on encoded information and also performing periodic error recovery steps will slow down the performance of the quantum computer.

Therefore it is not obvious that error correction will improve the performance of the quantum computer.

Fortunately it has been shown that *fault tolerant* quantum error correction is possible and a quantum computer can be in principle *fault tolerant* [22, 23]. By fault tolerant we mean that although the elementary components such as quantum gates and qubits are noisy, the quantum computer itself can in principle work effectively.

2.3 QIP Candidates

Computation is a physical process regardless of whether it is carried out by methods governed by classical or quantum mechanical laws. However the realisation of a quantum computer is challenging because the physical system relies on *coherent* operations. According to the DiVincenzo criteria [3] a physical system suitable for QC should fulfil the following requirements:

1. It should be scalable with well characterised qubits.
2. It should be possible to initialise the state of the qubit.
3. It needs to have long decoherence time compared with the gate operation time.
4. It must provide a universal set of gates.
5. It should have read-out or qubit measurement capability.

Besides these rules there are two other necessary requirements for quantum networking: (6) the ability to inter-convert stationary and flying qubits (photons) and (7) the ability to faithfully transmit flying qubits between specified locations. These seven rules form the necessary and sufficient conditions for a physical realisation of quantum information processing.

There are a number of serious candidates for quantum computing which include:

1. **Trapped ions.** In this scheme the internal levels of trapped ions in a linear configuration form the qubits. The qubits interact with each other using the phonon modes of the ions, collective oscillation acting as a “bus channel”.

2. **Cavity Quantum Electrodynamics (Cavity QED).** The qubit is the photon mode of an optical cavity. The modes are linked by the excitation of, or emission by, an atom passing through the cavity.

3. **Nuclear Magnetic Resonance.** The orientation of spin-1/2 nuclei in a molecule forms the qubit and the spin-spin interactions provide the bus channel for the qubit interactions.

4. **Superconducting Quantum Interference Devices (SQUIDS).** The quantisation of the flux in a superconducting circuit provides a two-level system as the qubit.

5. **Quantum Dots.** By placing artificial inhomogeneity in a crystal electrons can be confined and the orientation of the spin then can be used as the data register.

All these candidates have their own strengths and weaknesses. There is no ideal candidate as yet which can be picked with confidence for use in a large-scale quantum computer. In the ion trapping group of Imperial College we are trying to realise QIP with trapped ions in the Penning trap. In the next chapter I describe why trapped ions are a promising candidate for QIP purposes and also the shortcomings in current ion trap QIP technology.

CHAPTER 3

Realisation of QIP in Ion Traps

3.1 Introduction

Ion trapping has been one of the most successful options for doing quantum computation so far and cavity quantum electrodynamics (CQED) has been the most obvious choice for quantum networking. It seems that ion trapping and CQED can offer complementary elements for building a quantum computer.

In 1995 Cirac and Zoller [24] proposed a feasible scheme to do quantum computing using trapped ions. Since then ion trapping has become one of the most serious candidates for quantum computing. Effectively they noticed that trapped ion research at the time, which was mainly directed to the realisation of frequency standards, coincidently fulfilled most of the DiVincenzo criteria [3] even though they had not been written down by DiVincenzo yet!

3.2 State Detection Using Quantum Jumps

The basic idea of *electron shelving* was proposed first by Dehmelt [25]. The ion is considered to be a three level system consisting of a ground state $|g\rangle$, a metastable excited state $|d\rangle$ and a short lived excited state $|e\rangle$. A laser pulse resonant with the transition $|g\rangle \rightarrow |d\rangle$ drives the ion to a superposition $\alpha|g\rangle + \beta|d\rangle$. Subsequently the transition $|g\rangle \rightarrow |e\rangle$ is addressed by a second laser. The ion is excited to the short lived state $|e\rangle$ and collapses to $|g\rangle$ with the probability $|\alpha|^2$ and scatters a photon. The ion is *shelved* in state $|d\rangle$ with probability $|\beta|^2$. We are assuming that there is no coupling between $|d\rangle$ and $|e\rangle$. When the ion is in the state $|d\rangle$ no fluorescence is detected on the $|g\rangle \rightarrow |e\rangle$ transition. Since the $|e\rangle \rightarrow |g\rangle$ transition is a fast dipole transition the ion scatters millions of photons per second. In practice the solid angle for detection is $\sim 0.1 \times 4\pi$ and the efficiency of the photomultiplier tube is $\sim$0.1 which gives about $\sim$0.001 efficiency in total considering the losses in the optics (see section 5.6). It is possible to detect thousands of photons per second even given the restricted detection optics and the low efficiency of the photomultiplier. A trace of quantum jumps of a single Ca$^+$ ion is shown in figure 3.1 (Details of the specific energy levels in Ca$^+$ and the experimental procedure are given in the next chapter).

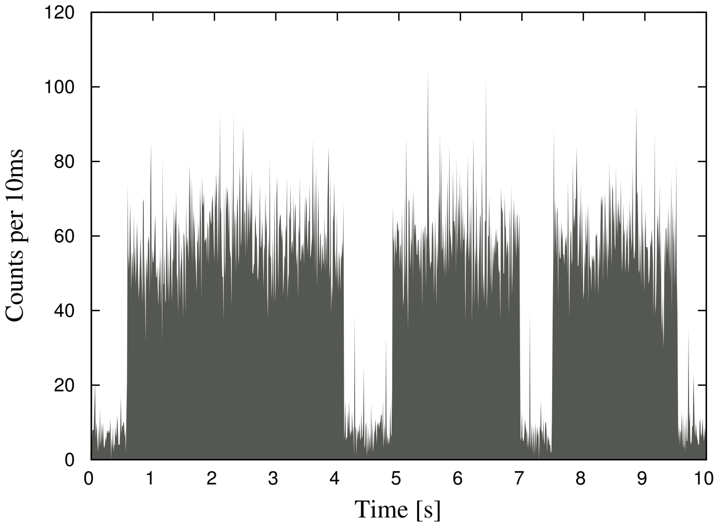

Figure 3.1. Quantum jumps of a single Ca$^+$ ion in a RF trap. When the ion makes a transition to the metastable state the fluorescence drops to the background level. After an average time equal to the lifetime of the metastable state, spontaneous emission returns the ion to the ground state and the fluorescence reappears. The time bin is 10ms.

In the language of quantum computation the states $|g\rangle$ and $|d\rangle$ form the states of the qubit. Supposing the signal to background ratio is about 1 for a time bin of 10ms (i.e. 50% error in detection), 100ms readout will give an accuracy with which the ground state can be determined as high as 99.9%. The higher the signal to background is the higher the efficiency will be at a given readout time and of course the readout time is limited by the lifetime of the metastable state.

3.3 Semi-Classical Theory of Atom-Field Interaction

Ion trap quantum computing is based on entangling the motional state of the ion to its internal atomic state. The motional state of the ion is quantised because the ion is trapped in a harmonic potential. The ion can be trapped using two alternative approaches, the Penning trap or the Paul (radio frequency) trap and its variants [26]. The Penning trap is discussed later on in chapter 4. The manipulation of the qubits is achieved by the interaction of the ion with an electromagnetic field which is provided by a highly stabilised single mode laser or a pair of lasers whose relative frequency separation is highly stabilised.

Consider a single mode classical electromagnetic field propagating along the z axis interacting with a single two-level trapped ion. The ion is assumed to move in the z direction in a harmonic well with frequency ω_z. The interaction Hamiltonian is given by

$$\begin{aligned}
\mathcal{H} &= \mathcal{H}_0 + \mathcal{H}_I, \\
\mathcal{H}_0 &= \hbar\nu(a^\dagger a + \frac{1}{2}) + \sum_i E_i \sigma_{ii}, \\
\mathcal{H}_I &= -\boldsymbol{\mu}_d \cdot \mathbf{E}, \qquad (3.1)
\end{aligned}$$

where in general $\sigma_{ij} = |i\rangle\langle j|$ is the atom transition operator with transition frequency of ω_a, E_i is the atomic energy of level $|i\rangle$, $\boldsymbol{\mu}_d = e\mathbf{r}$ is the electric dipole operator for the internal atomic transition and a ($a^\dagger$) are the creation (annihilation) operators of the harmonic oscillator potential well [27]. We can expand $\boldsymbol{\mu}_d$ in terms of the atomic transition operators to obtain

$$\boldsymbol{\mu}_d = \sum_{i,j} e|i\rangle\langle i|\mathbf{r}|j\rangle\langle j| = \sum_{i,j} \mathcal{D}_{ij}\sigma_{ij},$$

where $\mathcal{D}_{ij} = e\langle i|\mathbf{r}|j\rangle$ is the electric dipole transition matrix element. Suppose the electric field $\mathbf{E}$, which is provided by the laser field, constitutes a wave with frequency ω propagating in the z direction and polarised in the x direction

$$\mathbf{E} = \mathcal{E}\hat{x}[e^{i(kz-\omega t+\phi)} + cc.]/2.$$

Diagonal elements of the electric dipole operator $\mathcal{D}_{ii}$ are zero due to the fact that it has negative parity. Equation 3.1 can be written as

$$\mathcal{H}_I = -\frac{1}{2}\hbar\Omega(\sigma_+ + \sigma_-)(e^{i(kz-\omega t+\phi)} + cc.),$$

where $\Omega = \mu\mathcal{E}/2\hbar$ is the Rabi frequency, $\sigma_+ = |j\rangle\langle i|$ and $\sigma_- = |i\rangle\langle j|$. Now if we go to the interaction picture the unperturbed Hamiltonian can be ignored and the evolution operator in the interaction picture is written as

$$U'(t) = I - \frac{i}{\hbar}\int_0^t dt'\mathcal{H}_I^{'}\ ,$$

where U' is the interaction evolution operator to the first order and $\mathcal{H}_I^{'}$ is the interaction Hamiltonian operator. Next we invoke *the rotating wave approximation.* In this approximation the fast rotating terms $(\exp(\pm i(\omega_z + \omega_a)))$ are ignored. These rapidly rotating terms cancel out in the integration of $U'(t)$. Using the rotating wave approximation we get

$$\mathcal{H}_I^{'} = -\frac{1}{2}\hbar\Omega(e^{i(kz-\delta t+\phi)}\sigma_+ + e^{-i(kz-\delta t+\phi)}\sigma_-).$$

Here $\delta = \omega - \omega_0$ is the detuning frequency and $\omega_0 = \omega_z + \omega_a$. Using the fact that $z = \sqrt{\hbar/2\omega_z m}(a + a^\dagger)$ the Hamiltonian takes the form

$$\mathcal{H}_I^{'} = -\frac{1}{2}\hbar\Omega(e^{i(\eta(a+a^\dagger)-\delta t+\phi)}\sigma_+ + e^{-i(\eta(a+a^\dagger)-\delta t+\phi)}\sigma_-), \tag{3.2}$$

where the η is known as the Lamb-Dicke parameter and is given by $\eta = \sqrt{E_R/E_z}$ in which $E_R = (\hbar k)^2/2m = \hbar\omega_R$ is the classical recoil energy of the atom and $E_z = \hbar\omega_z$ is the vibrational energy quanta in the trap. Writing it in another form gives rise to

$$\eta = \sqrt{\frac{\omega_R}{\omega_z}}\ . \tag{3.3}$$

In order to make the distinction between the motion and the internal atomic states we consider the internal atomic state as a fictitious spin where

$|e\rangle \equiv |\uparrow\rangle$ and $|g\rangle \equiv |\downarrow\rangle$. Expanding the wave function in terms of the spin and the motional states

$$\Psi' = \sum_{s=\uparrow,\downarrow} \sum_{n=0}^{\infty} C_{s,n} |s\rangle |n\rangle.$$

Solving the Schrödinger equation $i\hbar\partial\Psi'/\partial t = \mathcal{H}_I'\Psi'$ results in two coupled differential equations [28]

$$\begin{aligned} \dot{C}_{\uparrow,n'} &= -\, i^{1+|n'-n|}\Omega_{n',n}C_{\downarrow,n}, \\ \dot{C}_{\downarrow,n} &= -\, i^{1-|n'-n|}\Omega_{n',n}C_{\uparrow,n'}, \end{aligned} \tag{3.4}$$

where $\Omega_{n',n} = \Omega|\langle n'|e^{i\eta(a+a^\dagger)}|n\rangle|$ and we have assumed that

$$\Psi' = \begin{bmatrix} C_{\uparrow,n'} \\ C_{\downarrow,n} \end{bmatrix}. \tag{3.5}$$

The evolution of the wave function can be easily investigated in a regime called *the Lamb-Dicke regime.* In this regime the Lamb-Dicke parameter is much less than 1. This occurs when the amplitude of the ion's motion is much less than the inverse of the wave number ($\lambda/2\pi$) or in another word $\langle\Psi|k^2z^2|\Psi\rangle^{1/2} \ll 1$. In this regime $\Omega_{n',n}$ can be calculated by the expansion of the exponential [29] which gives rise to

$$\Omega_{n',n} = \Omega\exp[-\eta^2/2](n_<!/n_>!)^{1/2}\eta^{|n-n'|}\mathcal{L}_{n<}^{|n'-n|}(\eta^2).$$

Here $\mathcal{L}$ is a generalised Laguerre polynomial and $n_<$ ($n_>$) is the lesser (greater) of n and n'. The coupled differential equations 3.4 can be solved easily to obtain

$$\Psi'(t) = \begin{bmatrix} \cos\Omega_{n',n}t & -ie^{i[\phi+\pi/2|n'-n|]}\sin\Omega_{n',n}t \\ -ie^{-i[\phi+\pi/2|n'-n|]}\sin\Omega_{n',n}t & \cos\Omega_{n',n}t \end{bmatrix}\Psi'(0). \tag{3.6}$$

Ideally we assume that the electromagnetic wave is in resonance with the harmonic oscillation of the ion ie. $\delta = \omega_z|n' - n|$. What we are usually interested in are three types of transitions: the carrier ($n' = n$), the first red sideband ($n' = n - 1$) and the first blue sideband ($n' = n + 1$), whose Rabi frequencies, in the Lamb-Dicke regime are given by $\Omega, \eta n^{1/2}\Omega$ and $\eta(n+1)^{1/2}\Omega$ to the first order in η. As an example consider the case when $|\Psi(0)\rangle = |\downarrow, n\rangle$

and radiation at the first blue sideband is applied to the ion. Equation 3.6 implies

$$\Psi'(t) = \cos(\Omega_{n+1,n}t) \mid \downarrow, n\rangle + e^{i\phi} \sin(\Omega_{n+1,n}t) \mid \uparrow, n+1\rangle.$$

This interaction entangles the internal and motional states of the ion.

When red sideband radiation is applied to an ion in the Lamb-Dicke regime, equation 3.2 can be expanded to the first order in η. Using the rotating wave approximation again gives rise to

$$\mathcal{H}_I = \hbar\eta\Omega^{'}(\sigma_+ a + \sigma_- a^\dagger), \tag{3.7}$$

where $\Omega^{'} = -\Omega/2$. This is the same as the Jaynes-Cummings Hamiltonian in Cavity QED with the difference that the quantised single mode of the radiation field is replaced with the harmonic motion of the atom in the trap.

The atom-field coupling requires an inhomogeneous field and the stronger this field the better the coupling becomes and hence the more strongly the motional and internal states couple together. In fact the interaction Hamiltonian 3.1 can be rewritten as

$$\begin{aligned}\mathcal{H}_I &= -\mu_x E_x(z,t)\\ &= -\mu_x(E(z=0,t) + \frac{\partial E(z,t)}{\partial z}|_{z=0} + ...).\end{aligned}$$

For plane waves $\partial E(z,t)/\partial z$ is proportional to k. As a result a smaller wavelength field results in a stronger interaction. This however can be problematic in the microwave regime where the field gradient is not high enough to excite the motional mode. This issue can be resolved by a technique which is called *stimulated Raman transitions.*

3.4 Stimulated Raman Transitions

As mentioned in the last section we need to have a large field gradient to be able to couple the internal atomic states to the motional state of the atom. If the internal state transition frequency is not high enough we cannot use just single photon transitions to gain a strong interaction. A useful technique that can be employed in such cases is stimulated Raman transitions through a third auxiliary level whose transition frequency to the atomic levels we want to couple is sufficiently large. Instead of one laser there are two laser beams

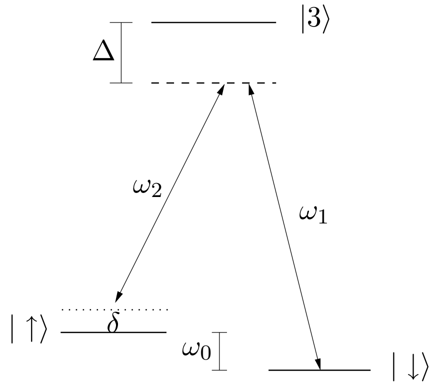

Figure 3.2. Stimulated Raman transition is shown. States $|\uparrow\rangle$ and $|\downarrow\rangle$ could be for example the ground state hyperfine levels.

applied to the atom (see figure 3.2). Laser 1 is tuned to the $|3\rangle \leftrightarrow |\uparrow\rangle$ transition and laser 2 is tuned to the $|3\rangle \leftrightarrow |\downarrow\rangle$. The two lasers are detuned from an auxiliary level $|3\rangle$ by the amount of Δ and from the level $|\uparrow\rangle$ by δ. Δ is much bigger than δ and the motional frequency of the atom. It can be shown that if these conditions are met, state $|3\rangle$ can be eliminated and the interaction Hamiltonian of the system will be

$$\mathcal{H}_I = -\frac{\hbar\Omega}{2}[\sigma_+ e^{-i[(\overrightarrow{\mathbf{k}_1}-\overrightarrow{\mathbf{k}_2})\cdot\overrightarrow{x}-(\omega_1-\omega_2)t+(\phi_1-\phi_2)]} + \text{H.c.}].$$

Here $\overrightarrow{\mathbf{k}_1}$ and $\overrightarrow{\mathbf{k}_2}$ are the two lasers' wave vectors. The equation above is the same as that for the two-level system if $\phi_1 - \phi_2$ is replaced by ϕ, $\overrightarrow{\mathbf{k}_1} - \overrightarrow{\mathbf{k}_2}$ is replaced by k and $\omega_1 - \omega_2$ is replaced with ω.

Another advantage of using stimulated Raman transitions is that it is resilient to the frequency fluctuations of the addressing laser. In an experiment, a single laser can be passed through a Accousto Optic Modulator (AOM) and the frequency of one beam is precisely shifted with respect to the other. Any phase or frequency fluctuations will be passed on to both beams and will be cancelled out because the dynamics is sensitive to the differences of the phase and frequency of the two beams. This is of course true so long as the frequency fluctuations are tiny compared to the detuning from the auxiliary state ($\delta/\Delta \ll 1$).

3.5 Laser Cooling to the Ground State of Motion

There are several reasons why cooling the ions in the trap is an essential part of experimental quantum information processing. The first is to make the storage time of the ion longer. Atoms, once generated from the oven, collide with the electrons to be ionised (or are photo-ionised). The ionisation process does not necessarily happen in the center of the trap and as the result the ions are initially hot. Collisions with the background gas molecules tend to bring the ions into thermal equilibrium with the vacuum enclosure which is usually at room temperature. Initially this can be seen as a cooling mechanism however the temperature limit is the temperature of the vacuum enclosure which is still very hot. The hotter the ion is, the higher is the probability of overcoming the potential barrier and losing it. Energy can be removed from the ion to a great extent by *Doppler laser cooling* [30]. The Doppler limit, as discussed in detail in section 4.4, is proportional to $\hbar\gamma/k_B$ where γ is the natural linewidth of the cooling transition and k_B is the Boltzmann constant.

Secondly the motional state of ion is treated as a "bus" qubit as a means for qubits to "talk" to each other for reasons that will become clear in section 3.6. The energy of the ion in the potential well of the trap is quantised. For sufficiently low temperature this quantisation is revealed. The obvious choice for this "bus" qubit is the ground state and the first excited state of motion. The motion of the Doppler cooled ion is still very classical ($\langle n \rangle \simeq 10$). The ion needs to be cooled further for QIP purposes. An ion can be cooled to the ground state of motion by several techniques. *Resolved side-band cooling* is one such method and is explained in section 3.5.2.

3.5.1 Cooling to the Doppler Limit

Suppose we have an unbound atom with an electric dipole transition frequency ω_0. The atom is subjected to a laser beam with frequency ω which is slightly detuned below the transition resonance frequency of the atom. If the atom is moving towards the beam it will see a Doppler shift of $\omega(1 + v/c)$ where v is the speed of the atom and c is the speed of light. Atoms in this velocity class will absorb the incoming photon with higher probability than atoms which copropagate with the laser beam. In each scattering event the atom loses $\hbar\overrightarrow{\mathbf{k}}$ of momentum where $\overrightarrow{\mathbf{k}}$ is the wave vector of the laser beam. If the atom, which is now in the excited state, emits a photon spontaneously, it will be kicked by the same amount of momentum but in a random direction. However, this absorption and emission happens many times so the average amount of momentum transferred to the atom by the emission process is nearly zero whereas the effects of the absorption events are cumulative and the atom slows down.

The atom performs a random walk in momentum space with steps equal to $\hbar k$. This imposes a limit on the amount by which the atom can be cooled. A more fundamental limit is that the kinetic energy must be at least equal to the recoil energy of the atom since if the cooling is switched off the atom is left with at least the momentum it gained from its final spontaneous emission event. The random walk in momentum space imposes a much higher limit to the temperature and it turns out that it is equal to $\hbar\gamma/2k_B$ where γ is the natural linewidth of the atomic transition and k_B is Boltzmann's constant [29]. The reason is that if the atom walks n random steps in momentum space by scattering n photons the resulting mean momentum is $\sqrt{n}\hbar k$. This constitutes a heating effect which counteracts the cooling process. As we will see in section 4.4, the maximum cooling rate is achieved when $\omega = \omega_0 - \gamma/2$.

It can be shown that for Doppler laser cooling of trapped ions when binding is weak (i.e. the trap frequency is much smaller than the natural linewidth of the cooling transition), the free atom approximation is still valid [29]. For typical trapping parameters this imposes the condition of $\langle n\rangle \gg 1$ where n is the oscillation quantum number. In fact this condition is dependent on the strength of the trap as well as the cooling transition linewidth and can be improved using a stronger trap. As an example a $^{40}\mathrm{Ca}^+$ ion with a natural linewidth of $\gamma = 2\pi \times 24\mathrm{MHz}$ has a Doppler cooling limit of $\sim 0.5\mathrm{mK}$.

3.5.2 Resolved Side-band Cooling

A trapped ion can be treated to a good approximation as a quantum mechanical harmonic oscillator. If the spontaneous decay rate is much smaller than the vibrational frequency of the ion or in other words the transition is weakly allowed, then the energy levels of the system can be resolved as internal levels each corresponding to a ladder of vibrational states (see figure 3.3). Efficient laser cooling of the harmonic motion occurs when the frequency of the laser beam is tuned to the red sideband i.e.

$$\nu = \nu_0 - \nu_z.$$

Here ν_0 is the internal atomic transition frequency and ν_z is the harmonic oscillation frequency of the atom. In this case the atom undergoes the transition

$$|g, n\rangle \xrightarrow{1} |e, n-1\rangle.$$

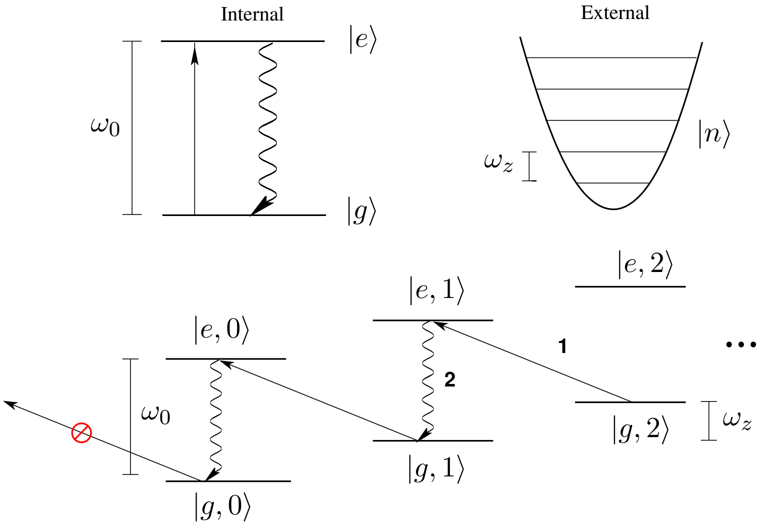

Figure 3.3. Resolved side-band cooling

A photon of energy $h(\nu_0 - \nu_z)$ is absorbed. Subsequent spontaneous emission occurs predominantly at the carrier frequency if the recoil energy of the atom is negligible compared with the vibrational quantum energy. This process can be viewed as

$$|e, n-1\rangle \xrightarrow{2} |g, n-1\rangle$$

in which a photon of energy $h\nu_0 - R$ will be spontaneously emitted where $R = (\hbar k)^2/2m = h\nu_R$ is the recoil energy of the ion . If $\nu_z \gg \nu_R$ this will result in a net cooling of $h\nu_z$. This condition is already satisfied in ion traps if we are in the Lamb-Dicke regime where $\eta^2 = \nu_R/\nu_z \ll 1$. The average effect of this mechanism is cooling the ion by one vibrational energy level. It can be shown [29] that the cooling proceeds until the average quantum number of harmonic well is

$$\langle n\rangle = 5(\gamma/\nu_z)^2/16 \ ,$$

where γ is the natural linewidth of the cooling transition.

In the experiment ν_z is fixed by the trap parameter and is usually of the order of 1MHz. Also an extra laser is implemented to repump state $|1\rangle$ to speed up the process. If cooling to the ground state of motion is needed one needs to choose a very narrow linewidth transition ($\gamma \ll \nu_z \simeq$ 1MHz). When these steps are repeated a sufficient number of times ground state occupation is reached with a high probability. To confirm this, the amplitudes of the sidebands are measured in the following way. The laser is detuned from the

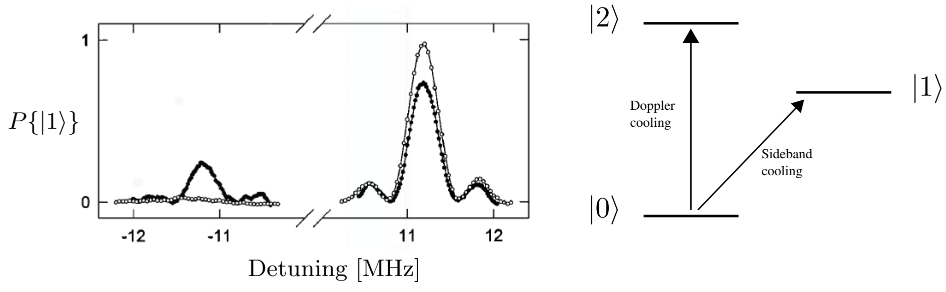

a) b)

Figure 3.4. (a) Absorption spectrum of a single $^9\text{Be}^+$ ion after Doppler precooling (solid points) and after five cycles of additional resolved sideband cooling (hollow points). Image taken from [31]. (b) The schematic level diagram of sideband cooling. The sideband cooling is usually a very slow, narrow linewidth, quadrupole transition whereas the Doppler cooling transition is a fast dipole transition.

sideband cooling transition by δ and the transition is addressed for a certain time. Then in the next step the Doppler cooling transition will be turned on and the state of the ion will be detected by electron shelving. If the ion is in the metastable state there will be no fluorescence whereas if it is in the ground state it will scatter lots of photons. This experiment is repeated many times while the detuning δ is scanned. The probability of finding the ion in the metastable state versus the detuning is obtained as shown in figure 3.4. When the ion is in the ground state of motion the red sideband disappears. The average occupancy of the ion in the ground state of motion can then be derived by $I_{red}/I_{blue} = \langle n\rangle/(1+\langle n\rangle)$, where I is the intensity ($P\{|1\rangle\}$) [31]. To make the cooling faster a laser tuned to the $|1\rangle \rightarrow |2\rangle$ transition is implemented to speed up return to the ground state $|0\rangle$.

3.6 C-NOT Gate

The C-NOT gate that is described here was proposed by Cirac and Zoller [24] and has been realised by Blatt's group in Innsbruck [4]. This scheme is based on a string of trapped ions where the atomic states of the ions are the qubits. We first introduce the C-Phase gate as the route to the C-NOT-gate. For simplicity suppose that there are only two ions in a trap and they are cooled to their collective motional ground state. We assume they are in the Lamb-Dicke regime so that the normal modes of motion are well separated in frequency space. The Coulomb force is considerable at the separation distance at which the ions find themselves and as a result a common motional mode

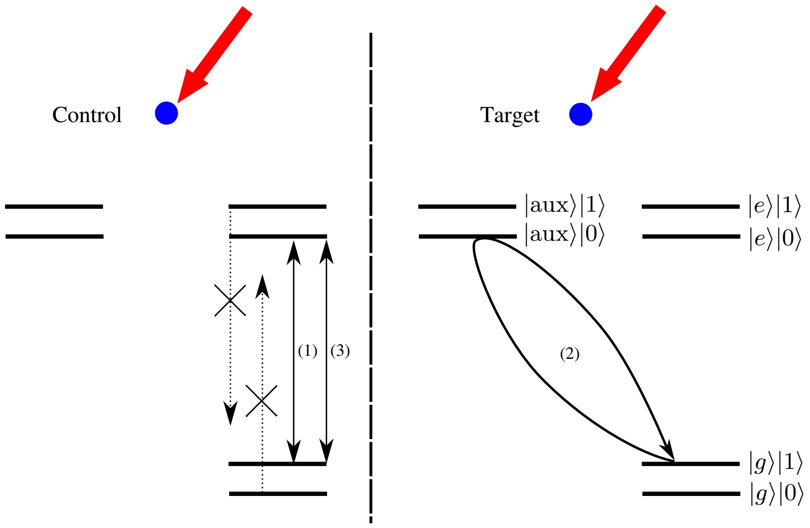

Figure 3.5. Three steps to perform the controlled-phase gate are shown. The dotted line is a non-resonant red-sideband laser for the case when the ion is in the ground state of motion and its internal state.

can be used as the "bus" to make a connection between the internal states of the ions. The obvious choice is the center of mass (CM) mode. Similar to side-band cooling one can address this mode by detuning the laser frequency to the red sideband of the CM mode (ω_z). The Hamiltonian describing the situation is essentially the same as equation 3.7 with the difference that the effective mass is $2m$ and therefore

$$\mathcal{H} = \hbar\frac{\eta}{\sqrt{2}}\frac{\Omega}{2}[|e\rangle\langle g|ae^{-\mathrm{i}\phi} + H.c.] \; .$$

We shine a π-pulse laser which is essentially a pulse with the time duration equal to a quarter cycle of a Rabi oscillation. The evolution operator for the pulse is

$$U = \exp[-\mathrm{i}\frac{\pi}{2}(|e\rangle\langle g|ae^{-\mathrm{i}\phi} + H.c.)] \; .$$

It can be shown easily that for a π-pulse this operator leaves $|g\rangle|0\rangle$ unaltered (because the laser is red-detuned) while

$$\begin{aligned}
|g\rangle|1\rangle &\rightarrow \cos(\pi/2)|g\rangle|1\rangle - \mathrm{i}e^{\mathrm{i}\phi}\sin(\pi/2)|e\rangle|0\rangle = -\mathrm{i}e^{\mathrm{i}\phi}|e\rangle|0\rangle, \\
|e\rangle|0\rangle &\rightarrow \cos(\pi/2)|e\rangle|0\rangle - \mathrm{i}e^{-\mathrm{i}\phi}\sin(\pi/2)|g\rangle|1\rangle = -\mathrm{i}e^{-\mathrm{i}\phi}|g\rangle|1\rangle,
\end{aligned}$$

where $|0\rangle$ ($|1\rangle$) represents the CM mode state. Suppose that there is another state $|\mathrm{aux}\rangle$ which can be addressed by another red detuned laser to the $|g\rangle \leftrightarrow |\mathrm{aux}\rangle$ transition which we call the auxiliary transition. The total state of the system is shown by $|a\rangle|b\rangle|c\rangle$ where $|a\rangle$ and $|b\rangle$ are the states of the first and the second ion respectively and $|c\rangle$ is the CM state. Supposing that the CM mode is in the ground state and the equilibrium position of the ion is at the node of the laser ($\phi = 0$), the C-Phase gate can be achieved in three steps as follows:

1. a π laser pulse is applied driving the qubit transition in the first ion (control qubit). This step cannot affect the system when both the CM mode and the ion are in their ground states. However when either the ion or the CM mode are excited it flips the states of the ion and CM mode and introduces a $-i$ phase factor. In other words it maps the state of the control qubit onto the CM vibrational mode and uses it as a quantum memory.
2. a 2π laser pulse is applied on the second ion (target qubit) driving the auxiliary transition. This does not have any effect when the ion is excited or when both the ion and CM mode are in the ground state since there is no state below the motional ground state to which the ion can be pumped. However it changes the sign when the ion is in the ground state and the CM mode is excited. In other words, conditional on the vibrational mode state, it introduces a minus phase factor. The above discussion makes it clear why cooling to the motional ground state is a prerequisite for this work. The key process in the logic gate is conditional on whether the ion is in the motional ground state or not.
3. same as the step 1, a π laser pulse is applied driving the qubit transition in the control qubit and the vibrational state is mapped back onto the control qubit.

The whole process can be summarised as

$$\begin{array}{lcrclcl}
|g\rangle|g\rangle|0\rangle & \xrightarrow{1} & |g\rangle|g\rangle|0\rangle & \xrightarrow{2} & |g\rangle|g\rangle|0\rangle & \xrightarrow{3} & |g\rangle|g\rangle|0\rangle \\
|g\rangle|e\rangle|0\rangle & \xrightarrow{1} & |g\rangle|e\rangle|0\rangle & \xrightarrow{2} & |g\rangle|e\rangle|0\rangle & \xrightarrow{3} & |g\rangle|e\rangle|0\rangle \\
|e\rangle|g\rangle|0\rangle & \xrightarrow{1} & -i|g\rangle|g\rangle|1\rangle & \xrightarrow{2} & i|g\rangle|g\rangle|1\rangle & \xrightarrow{3} & |e\rangle|g\rangle|0\rangle \\
|e\rangle|e\rangle|0\rangle & \xrightarrow{1} & -i|g\rangle|e\rangle|1\rangle & \xrightarrow{2} & -i|g\rangle|e\rangle|1\rangle & \xrightarrow{3} & -|e\rangle|e\rangle|0\rangle
\end{array}$$

For the four different input states. The C-Phase gate is equivalent to the C-NOT gate if we change the basis to $|\pm\rangle = (|g\rangle \pm |e\rangle)/\sqrt{2}$ for the second ion by a $\pi/2$ pulse at the carrier frequency. In this basis the phase gate acts as

$$\begin{aligned} |g\rangle|\pm\rangle &\rightarrow |g\rangle|\pm\rangle \\ |e\rangle|\pm\rangle &\rightarrow |e\rangle|\mp\rangle \end{aligned}$$

The realisation of the proposed Cirac-Zoller gate demands that the ions are cooled close to zero temperature. This had already been realised many years earlier [32]. However even if the ion is sitting in the ground state of motion for 99% of the time the fidelity* of the gate operation is still not high enough to allow fault tolerant quantum computing. Moreover scaling will become an issue since by scaling up the number of ions the number of modes of motion will scale up too and they need to be cooled as well otherwise it becomes harder to address a single sideband. There are other schemes that have been proposed [34] and implemented [5] that do not require the zero temperature condition.

3.7 Ion Traps as a Candidate for QIP

Why ion traps

As mentioned briefly at the beginning of this chapter DiVincenzo in his paper in 2001 proposed 5 criteria that any potential QC candidate system needs to satisfy. These five rules form the necessary and sufficient conditions for a physical realisation of quantum computing. We now examine how ion trap realisations fare in comparison with the criteria.

1. **It should have a scalable and well characterised qubit system.** This can be achieved using the hyperfine or Zeeman ground state sub-levels in for example Be^+ or a ground and metastable excited state of a weakly allowed optical transition such as Ca^+. The qubits can directly interact with each other by the weak Coulomb repulsion in a common motional state quantum “bus”. Scalability can be achieved using arrays of ion traps which are interconnected by interchanging photons in a cavity QED arrangement or by shuttling the ions from one trap to another by time varying electric potentials.

---

*Fidelity is a measure of ‘closeness’ of two quantum states [33] and fidelity of a quantum gate is a measure of reliability of the performance of the gate.

2. **It should be possible to initialise the state of the qubits.** Spin qubits (internal states) can be initialised with high probability using standard optical pumping techniques. Using the sideband cooling technique the motional state can be prepared to the ground state of the harmonic oscillation of the ion inside the trap. There are alternative gate schemes where cooling to the ground state is not required [5, 34].

3. **The qubit needs to have a long coherence time compared with the gate operation time.** The spin qubit decay time can be as long as 10min compared to typical 10μs gate operation times if one chooses a hyperfine transition which is magnetic field independent to the first order. The weakly allowed optical transitions have a typical lifetime of 1s (for example Ca^+). The motional state coherence time is however much smaller ($\sim 100\mu$s) which is due to thermal radiation noise and fluctuating electric fields resulting from "patch" potentials. Both could be reduced with better surface qualities of the electrodes and going to cryogenic temperatures.

4. **It must provide a universal set of gates.** Single-qubit rotation can be easily accomplished by addressing the qubit transition with a single laser or a pair of Raman lasers for a certain time of illumination. Fidelity can be improved by suppressing the frequency, phase and intensity fluctuations of the laser. Two-qubit gates can be accomplished using a wide range of schemes available notably: 1) Cirac and Zoller 2-qubit C-NOT gate 2) Cirac and Zoller 2-qubit "push" gate 3) Mølmer and Sørensen 2-qubit geometric phase gate.

5. **It should have read-out or qubit measurement capability.** State sensitive detection can be accomplished using the standard quantum jump technique with nearly 100% efficiency.

Ion traps are an ideal candidate for QC since they fulfil all of the DeVincenzo criteria. However there are many complications which I shall address in the next sections that make the realisation of a true quantum computer a very challenging task.

Current status in ion trap QIP

A few months after the Cirac-Zoller proposal in 1995 the NIST group in Colorado demonstrated for the first time a quantum logic gate on a single ion

[35]. They realised the C-NOT gate on a single $^{9}\mathrm{Be}^{+}$ ion in a RF trap. The two qubits were the two internal hyperfine states and the two external motional states of the ion. This was different from what originally had been proposed by Cirac and Zoller because it was not scalable to more ions but it was nonetheless a key step to implementing the two ion C-NOT gate. Three years later they demonstrated deterministic (high certainty and on demand) entanglement of two ions in the Bell-like singlet and triplet states [36]. It is noteworthy that this was done without single ion addressing. They illuminated each ion equally but changed the Rabi frequency of individual ions by putting micromotions on them as a means of single ion addressing. Two years later they experimentally demonstrated four ion entanglement without directly exciting the motional states of the ions and still without single ion addressing [37].

Major progress came in 2003 when both the NIST group and Blatt's group in Innsbruck demonstrated two-ion-qubit gates based on two different schemes [4, 5]. The Controlled-Phase gate in NIST was based on the proposal by Milburn [34]. The qubits were the two internal states of two $^{9}\mathrm{Be}^{+}$ ions trapped in a cylindrical RF trap. The state dependent displacement force was due to a spatially varying a.c. Stark shift from two independent lasers which were detuned with respect to each other by the frequency of the stretch mode (as apposed to the CM frequency in Cirac-Zoller scheme) plus some small additional detuning δ. The frequencies and polarisations of the two lasers were adjusted such that when both ions were in the same state the Stark shifts were equal but were different when they were in opposite states. The effect is that when the ions are in different states there is a time dependant differential force (wobble) between them that can excite the stretch mode. The result is that the wave function of the system accumulates a state dependent phase. It is noteworthy that in this way the force does not depend on the motional state so long as the normal mode excursions are in the Lamb-Dicke regime. So this gate is more resilient to the decoherence due to the heating of the ions and in fact its fidelity (F) is six times better ($F \sim$97%) than in a similar experiment done by the Innsbruck group which was based on the Cirac-Zoller scheme ($F \sim 71\%$).

At the same time the Innsbruck group reported the realisation of the Cirac-Zoller C-NOT gate on two $^{40}\mathrm{Ca}^{+}$ ions in a linear RF trap with a fidelity of ∼80%. The low fidelity was mainly due to the laser frequency noise and addressing error which were improved later. It is worth mentioning that 2% of the infidelity was due to the residual thermal excitation even at $\langle n \rangle_{bus} < 0.02$ which indicates the difficulty in reaching fault tolerant QC if one scales up the number of ions.

In 2004 both the NIST and Innsbruck groups demonstrated three-ion entanglement using techniques similar to those they used for the two-qubit gates. In NIST they produced maximally entangled GHZ states ($\{|0\rangle|0\rangle|0\rangle + i|1\rangle|1\rangle|1\rangle\}/\sqrt{2}$) [38] with a fidelity of 89%. At the same time Blatt's group demonstrated the GHZ and W states entanglement ($\{|1\rangle|0\rangle|0\rangle + |0\rangle|1\rangle|0\rangle + |0\rangle|0\rangle|1\rangle\}/\sqrt{3}$) with fidelities of 72% and 83% respectively [39]. In more recent experiments in 2005 the NIST group demonstrated "Schrödinger cat" states for up to 6 ions ($\{|000\ldots\rangle + |111\ldots\rangle\}/\sqrt{2}$) with fidelities of $F_4 = 0.76, F_5 = 0.60$ and $F_6 = 0.51$ for four, five and six ions respectively [40]. At the same time the Innsbruck group reported W states of up to 8 ions with the fidelities of 0.85, 0.79 and 0.72 for 4, 6 and 8 ions respectively [41]. The lower fidelities of the NIST experiment are thought to be due to lack of single ion addressing.

Future prospects

Thanks to these impressive experiments ion trapping research has attracted a great deal of attention. However it should be mentioned that to perform meaningful fault tolerant quantum computing one needs much higher fidelities (better than 99.9%) and many more qubits (more than 1000) [42]. This shows how far ion trapping QIP is from the reality of a working QC.

Despite being successful at showing the basic operations of quantum computing, strings of ions in a linear RF trap are very limited in terms of scalability. This is because one needs to cool all the modes of the chain, which grow linearly with the number of ions and this gets more difficult as more qubits are introduced into the system. Additionally, the longer the string is, the weaker the trap potential needs to be and, as a result, the effectiveness of sideband cooling goes down. This problem can be addressed using microarray RF traps. These traps could be interconnected in a number of ways for example: 1) photons 2) a moving "head" such as a trap or a link that transfers the information to different traps 3) shuttling the ions into trap nodes in the array.

Interconnecting the traps by photons can be realised using existing Cavity QED techniques. The Innsbruck group and Langé's group in Sussex are trying to couple the internal state of a trapped ion to the cavity mode [43, 44].

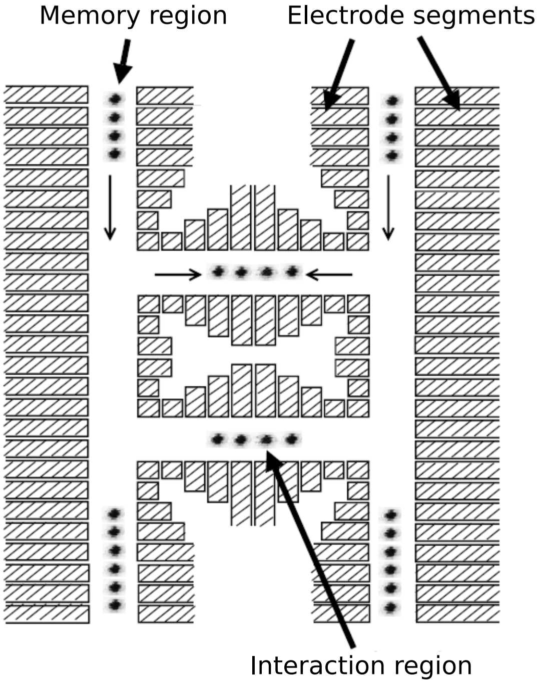

Figure 3.6. Diagram of the Quantum CCD is shown. The ions are stored in the memory regions and are moved to the interaction region for gate operations. Taken from [46].

Another possible scheme is a two-dimensional array of independent ion traps with a moving ion trap head that can approach a particular ion and apply a phase gate ("push" gate) [45]. This is a also a very promising scheme since most of the ions do not need to be moved around and hence are not heated. It is also resilient to the temperature fluctuations of the ion therefore the ions do not need to be cooled to zero temperature. However in practice it could be slow since the trap needs to be moved mechanically.

A "Quantum CCD" is the scheme suggested by the NIST group. A "Quantum CCD" is an array of small ion traps acting as registers which are interconnected via junctions [46]. The proposed device, as shown in figure 3.6, consists of two regions: a memory region in which the ions are stored and an interaction region in which logic operations are performed. By changing the operating voltages of the segmented electrodes, the ions can be trapped in a single trap or shuttled from one region to another by time varying potentials. Since the gate operation time is limited by the trap strength, shuttling times should be just a fraction of the gate operation time. Shuttling ions in a linear array of segmented electrodes has already been realised [47] for relatively large

RF traps in a time approximately equal to the gate operation time and with the heating of about 1 motional quantum. The cause of this heating is not clear yet but it could be due to contaminations on the surface of the electrodes, “patch” potentials and the thermal Johnson noise on the electrodes. Since heating will occur sympathetic cooling is needed in the gate trap.

To allow large numbers of ions the traps need to be microfabricated in, for instance, semiconductor chips. The electrodes can be planar [48] with ions above the electrodes or made conventionally in 3D [49]. The 3D design is more challenging but has a deeper trapping potential.

A two dimensional large scale “Quantum CCD” needs to enable shuttling ions through T and X junctions. Recently Chris Monroe’s group in Michigan [50] built a microfabricated T junction ion trap with an ion transport rate of 98%. It appears that it is impossible to have a pseudo potential that extends along all three direction of the junction. There are three small “humps” (potential barriers) and a bigger time-varying potential is required to overcome these humps. This leads to significant motional heating of the ion and the ions need to be Doppler cooled again.

There are some practical difficulties as well, such as heat dissipation due to the RF voltages applied to the thin layer of electrodes. The time-varying voltages also need to be precisely tailored to minimise the heating of the ions. The voltages in the junctions need to be further improved to have less heating during the shuttling of ions.

Another issue is the qubit decoherence time. Weakly allowed optical transitions such as the $S_{1/2} \rightarrow D_{5/2}$ quadrupole transition in Ca^+ have a decay time of ∼1s. This can be improved using the hyperfine levels that have much longer decay times instead and use first order magnetic field independent transitions for further coherence. This can be achieved by applying a low ambient magnetic field such that the energy difference between the qubit states passes through an extremum [51].

Moreover the single qubit rotation errors must be reduced significantly to enable fault tolerant quantum computing and that demands more laser frequency, power and position stabilities.

Although ion trapping is one of the most successful candidates for QIP it is still quite far off from a realistic quantum computer.

CHAPTER 4

QIP in the Penning Trap

4.1 Introduction

In this section we introduce the Penning trap and the potential advantages of using it for QIP purposes. The trap structure and the equations of motion are explained. The laser cooling in the Penning trap, which is slightly different from that in the RF trap, is also explained. 'Pad traps', our proposal for the scalable Penning trap, are introduced. The results of simulations which I carried out to find the optimum electrode voltages for 'hopping' the ions from one trap to another while they are axially confined are presented at the end of this chapter.

4.2 Background

In the previous chapter the current issues in ion trap QIP particularly with regard to RF traps were discussed. There are some complications which are specific to RF traps. The RF trap uses a radio frequency oscillating electric potential to confine the ion in 3 dimensions. The RF field drives micromotion which increases with the distance of the ion from the centre of the trap. RF heating, which refers to an increase in the ion's secular motion, arises from

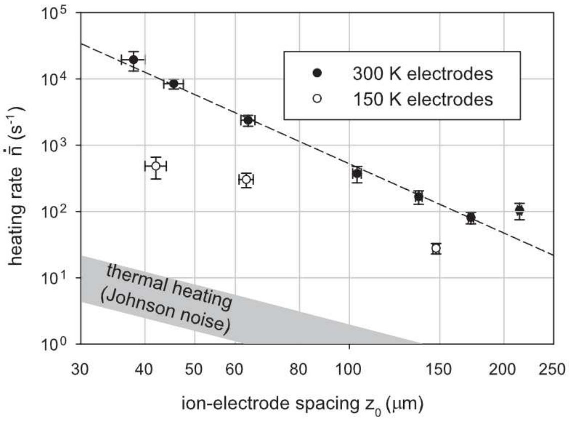

Figure 4.1. Image taken from [53] shows the heating rate of the ion in a needle RF trap against the trap size. The heating is one order of magnitude less when the electrodes are cooled to 150K, but the heating rate is still a few orders of magnitude bigger than the thermal noise expected (the shaded region). The fitted line is $\propto (\frac{1}{z_0})^{-3.5}$.

the presence of "patch potentials". The exact mechanism of RF heating is not entirely understood but it is certainly a well documented effect [52]. The cause of these patch potentials is also not very well understood but contamination on the surface of the electrodes and also the finite temperature could be amongst them. The heating due to these patch potentials increases as the dimensions of the trap become smaller (scaling as a high power of $1/z_0$) [53] (see figure 4.1). In a recent experiment by Chuang's group in MIT a significantly lower heating rate has been demonstrated for a cryogenic RF trap at 6K [54]. However trapping and laser cooling at cryogenic temperatures is a very experimentally demanding task and may not be feasible for a large scale QC. In order to trap larger numbers of ions, the axial potential needs to be weak to allow a chain of ions to be kept along the axis. The heating becomes a bigger issue when shuttling ions through complex electrode structures where the RF traps are even weaker. This issue can be partly solved using independent ion traps with only one ion in each trap.

Another issue is dephasing due to ambient magnetic fields. Ambient magnetic fields could induce severe decoherence into the system (even with active compensation of the magnetic field) if one is not using magnetic field-independent optical transitions [55]. Blatt's group in Innsbruck found that even when the magnetic field is actively compensated there are sources of magnetic field fluctuations localised in the vicinity of the trap which are not picked up by the sensors.

At the Ion Trapping Group in Imperial College London we have opted to use a completely different trap, the Penning trap, to avoid many of these problems. In the Penning trap ions are confined using only static electric and magnetic fields and thus, in principle, do not suffer RF heating effects. Because the trap frequencies are dependent on the magnetic field strength, the Penning trap can in principle be bigger than a RF trap whilst still having high frequencies of the ion motion.

The magnetic field can be provided using superconducting magnets which give field stability better than 1 part in 10^7 and can easily go up to 10T. In fact superconducting magnets provide us with an excellent ambient field shield and their stability could be improved further by active stabilisation. Moreover we have the possibility of running the trap with a magnetic field that would make the optical transition field-independent to the first order.

With these advantages we expect a better coherence time for the Penning trap. The Penning trap is not widely used for QIP because of some perceived disadvantages. Much of the work of the IC group in recent years has been aimed at overcoming these difficulties. The first complication associated with the Penning trap is the Zeeman splittings due to the magnetic field. Similar to the case in using hyperfine ground states in RF traps, more lasers are required to address the Zeeman sublevels. Since many of the RF trap groups are turning to odd isotopes with hyperfine structure this should no longer be seen as a major disadvantage of the Penning trap. The second is that one degree of motion (the magnetron motion) is unstable and the cooling rate is slow (see section 4.3). Fortunately using “axialisation”, a technique which couples this motion to the cyclotron motion, it is possible to cool the magnetron motion as well. Finally the axial frequency of the ion is indirectly limited by the magnetic field strength as well as the volume of the trap. As a result there is a technical limit on using the geometric phase gate or the pushing gate whose speeds depend on the axial frequency of the ion in the trap. However the expected longer coherence time in the Penning trap might compensate for this.

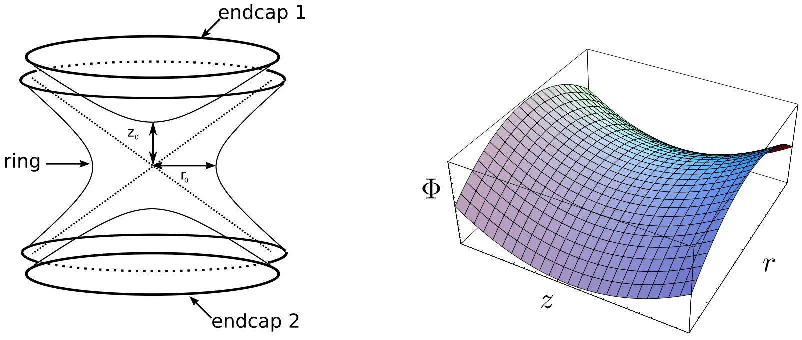

(a) (b)

Figure 4.2. a) Ideal Penning trap structure b) Saddle potential of the Penning trap. The axial motion is stable but the radial motion is unstable. The ion is confined in the radial plane by applying a magnetic field perpendicular to the radial plane.

4.3 Trap Structure and Equations of Motion

An ideal Penning trap employs a combination of a homogeneous magnetic field $\mathbf{B} = B\hat{e}_z$ and an electric field $\mathbf{E} = -\nabla\Phi$ where the potential Φ is given by

$$\Phi(x, y, z) = \frac{U}{2z_0^2 + r_0^2}(2z^2 - x^2 - y^2). \tag{4.1}$$

This static electromagnetic potential can be generated using a set of three electrodes shown in figure 4.2(a). The ring and endcaps 1 and 2 are hyperboloids of revolution. z_0 is half the distance between the endcaps and r_0 is the radius of the ring electrode. A potential difference of U is applied between the ring and endcap. This potential produces a saddle point at the centre of the trap which traps the ion only along the z axis (figure 4.2(b)). In order to confine the ion in the radial plane a homogeneous magnetic field $\mathbf{B}$ is applied along the z direction to force the ion into circular orbits and keep them in the trap. A positively charged particle q sees the Lorentz force

$$\mathbf{F} = -q\nabla\Phi + q(\dot{\mathbf{x}} \times \mathbf{B}),$$

where **x** is the position vector. The classical equations of motion in Cartesian coordinates follow as [56, 57]

$$\begin{aligned}\ddot{x} - \omega_c \dot{y} - \frac{1}{2}\omega_z^2 x = 0,\\ \ddot{y} + \omega_c \dot{x} - \frac{1}{2}\omega_z^2 y = 0,\\ \ddot{z} + \omega_z^2 z = 0,\end{aligned}$$

where $\omega_c = qB/m$ is the (true) cyclotron frequency and

$$\omega_z = \left(4qU/m(2z_0^2 + r_0^2)\right)^{1/2}$$

is the axial frequency. The cyclotron frequency does not depend on the geometry of the trap. The first two equations are coupled and can be written in a compact form as one complex differential equation

$$\ddot{u} + i\omega_c \dot{u} - \frac{1}{2}\omega_z^2 u = 0,$$

where $u = x + iy$. The general solution of this equation is

$$u(t) = R_+ e^{-i\omega_+ t} + R_- e^{-i\omega_- t}, \tag{4.2}$$

$$\omega_\pm = \omega_c/2 \pm \omega_p,$$

where $\omega_p = \sqrt{\omega_c^2 - 2\omega_z^2}/2$. This gives us the trapping condition

$$\omega_c^2 - 2\omega_z^2 > 0. \tag{4.3}$$

Here ω_+ is called the modified cyclotron frequency, ω_- is the magnetron frequency and $R_\pm$ are arbitrary constants which are determined by the initial conditions (initial position and velocity of the ion where it is created in the trap). Although ω_z does not depend on the magnetic field strength B, equation 4.3 puts an upper limit on how high the axial frequency can be in a specific trap employing a given magnetic field. The trajectory of the particle can be written as

$$\begin{aligned}x &= R_+ \cos(\omega_+ t) + R_- \cos(\omega_- t)\\ y &= R_+ \sin(\omega_+ t) + R_- \sin(\omega_- t).\end{aligned} \tag{4.4}$$

The resulting trajectory in the radial plane is depicted in figure 4.3. Under

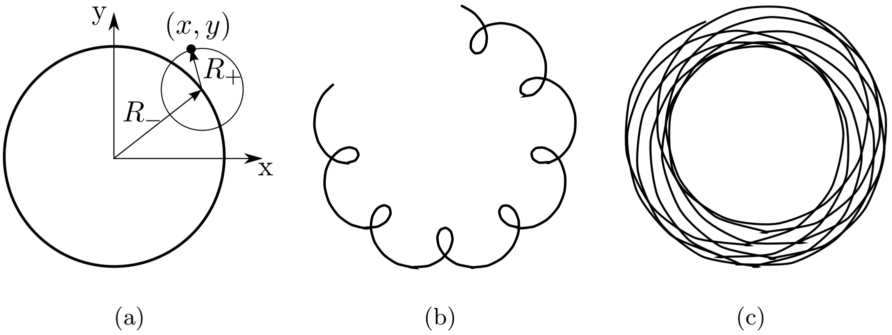

Figure 4.3. a) Visualisation of equation 4.2 b) The trajectory of an ion when $R_+ < R_-$. c) The trajectory for the case when $R_+ > R_-$.

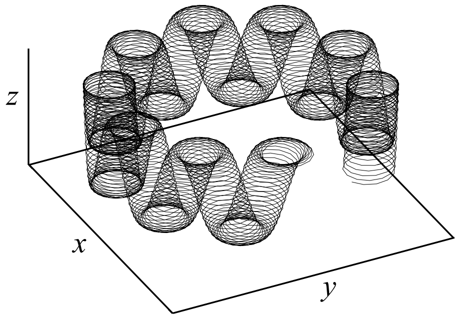

Figure 4.4. Trajectories of a single ion in three dimensions with $\frac{R_-}{R_+} = 10$, $\frac{\omega_+}{\omega_-} = 500$ and $\frac{\omega_-}{\omega_z} = 10$.

normal conditions for the Penning trap we usually have $\omega_c \simeq \omega_+ \gg \omega_z \gg \omega_-$. A trajectory in 3D is shown in figure 4.4.

The Hamiltonian describing the system is

$$\begin{aligned} H &= \frac{1}{2m}(\mathbf{p}-\mathbf{A})^2 + q\Phi \\ &= \frac{1}{2m}\mathbf{p}^2 + \frac{1}{8}m\omega_p^2(x^2+y^2) + \frac{1}{2}m\omega_z^2 z^2 - \frac{1}{2}\omega_c L_z, \end{aligned} \tag{4.5}$$

where we have used the symmetric gauge $\mathbf{A} = \frac{1}{2}\mathbf{B}\times\mathbf{r}$ for the magnetic vector potential and $L_z = xp_y - yp_x$ is the z-component of the angular momentum. Replacing 4.4 into 4.5 gives the total energy of the particle

$$\begin{aligned} E_{tot} &= E_+ + E_- + E_z, \\ E_+ &= \frac{m}{2}\omega_p\omega_+ R_+^2, \\ E_- &= -\frac{m}{2}\omega_p\omega_- R_-^2, \\ E_z &= \frac{m}{2}\omega_z^2 z_{max}^2. \end{aligned} \tag{4.6}$$

The magnetron motion has a negative contribution to the Hamiltonian and hence its “cooling” process must be the opposite of other degrees of freedom. Therefore in order to “cool” the magnetron motion, the laser should be tuned to the blue instead of the red*. However this is incompatible with simultaneous cooling of the “normal” axial and cyclotron modes. As we will find in the next section simultaneous cooling of all three motions can be achieved by offsetting the laser beam position in the radial plane to the side of the trap where the ion moves along the direction of the beam.

4.4 Laser Cooling in the Penning Trap

There are a few proposals for investigating laser cooling in the Penning trap [58, 59]. This section is based on the paper by Itano and Wineland [58]. We are looking for the damping rates due to the laser light which is close to resonance with the internal transition frequency ω_0 of a single ion. For a given laser frequency this force depends on a) the position of the ion and b) the velocity of the ion. Consider the case where a photon of wave vector $\mathbf{k}$ hits an ion

---

[*] We adopt a pragmatic definition of “cooling” as any process that tends to minimise the extent of the motion of an ion in the trap. For the magnetron motion this actually means adding energy to this degree of motion.

with velocity v and a photon of wave vector $\mathbf{k}_s$ is spontaneously emitted and the velocity of the ion is changed to $\mathbf{v}'$. By the conservation of energy and momentum we have

$$\begin{aligned} m\Delta\mathbf{v} &= \hbar(\mathbf{k}-\mathbf{k}_s), \\ \Delta E &= \frac{1}{2}m(\mathbf{v}+\Delta\mathbf{v})^2 - \frac{1}{2}m\mathbf{v}^2 \\ &= \hbar^2(\mathbf{k}-\mathbf{k}_s)^2/2m + \hbar(\mathbf{k}-\mathbf{k}_s)\cdot\mathbf{v}. \end{aligned}$$

ΔE is the total energy change per scattering event. Written in Cartesian form

$$\Delta E_{i=x,y,z} = R(\hat{k}_i^2 - 2\hat{k}_i\hat{k}_{si} + \hat{k}_{si}^2) + \hbar(k_i - k_{si})v_i,$$

where $R = (\hbar k)^2/2m$ is the recoil energy of the ion and $\hat{k}_i = k_i/k$ is the unit vector of the incident wave. Assuming the spontaneously emitted photon distribution is isotropic $\hat{k}_{si}$ is on average zero and $\langle \hat{k}_{si}^2 \rangle = 1/3$. Therefore the average total energy change per scattering event is

$$\langle \Delta E \rangle = 2R + \hbar\mathbf{k}\cdot v.$$

It can be shown [29] that the absorption scattering cross section of the photons of the particular internal atomic transition ω_0 of the ion is

$$\sigma(\omega,\mathbf{v}) = \frac{\sigma_0(\gamma/2)^2}{(\omega_0 + \mathbf{k}\cdot\mathbf{v} + R/\hbar - \omega)^2 + (\gamma/2)^2},$$

where $\sigma_0 = 6\pi c^2/\omega_0^2|\hat{\epsilon}\cdot\hat{d}_{eg}|^2$ and γ is the transition natural linewidth. The recoil energy R is usually much less than the detuning energy $\hbar(\omega_0 - \omega)$. Therefore we can drop the recoil energy and rewrite the cross section as

$$\sigma(\omega,\mathbf{v}) \simeq \frac{\sigma_0(\gamma/2)^2}{(\omega_0 + \mathbf{k}\cdot\mathbf{v} - \omega)^2 + (\gamma/2)^2}.$$

The scattering rate γ_s is equal to the number of photons per unit area per unit time $I/\hbar\omega$ times the scattering cross section σ. The rate equations of the energies in Cartesian coordinates hence become

$$\begin{aligned}\frac{dE_i}{dt} &= \langle\gamma_s \Delta E_i\rangle \\ &= \frac{I}{\hbar\omega}\langle\sigma(\omega,\mathbf{v})\Delta E_i\rangle \\ &= \frac{I}{\hbar\omega}\langle\sigma(\omega,\mathbf{v})[\hbar k_i v_i + R(\hat{k}_i^2 + \hat{k}_{si}^2)]\rangle.\end{aligned}$$

Now consider the case where there is a uniform laser beam propagating along the z axis and covering the whole ion orbit. The rate equations become

$$\begin{aligned}\frac{dE_z}{dt} &= \frac{I}{\hbar\omega}\langle\sigma(\omega,\mathbf{v})[\hbar k v_z + R(1 + \hat{k}_{sz}^2)\rangle, \\ \frac{dE_y}{dt} &= \frac{I}{\hbar\omega}\langle\sigma(\omega,\mathbf{v})R\hat{k}_{sy}^2\rangle, \\ \frac{dE_x}{dt} &= \frac{I}{\hbar\omega}\langle\sigma(\omega,\mathbf{v})R\hat{k}_{sx}^2\rangle.\end{aligned}$$

Assuming we have cooled the ion considerably so that the Doppler broadening is less than the natural linewidth or the detuning, we can approximate the cross section as

$$\sigma(\omega,\mathbf{v}) \simeq \sigma_0(\frac{\gamma}{2})^2 \frac{1 - 2(\omega_0 - \omega)\mathbf{k}\cdot\mathbf{v}/[(\gamma/2)^2 + (\omega_0 - \omega)^2]}{(\gamma/2)^2 + (\omega_0 - \omega)^2}. \tag{4.7}$$

Let γ_s be the average scattering rate

$$\gamma_s = \frac{I}{\hbar\omega}\langle\sigma(\omega,\mathbf{v})\rangle.$$

The rate equations now can be written as

$$\begin{aligned}\frac{dE_z}{dt} &= \gamma_s\{\frac{-2(\omega_0 - \omega)\hbar k^2\langle v_z^2\rangle}{(\gamma/2)^2 + (\omega_0 - \omega)^2} + 4R/3\}, \\ \frac{dE_y}{dt} &= \gamma_s R/3, \\ \frac{dE_x}{dt} &= \gamma_s R/3, \end{aligned} \tag{4.8}$$

considering the fact that $\langle\mathbf{k}\cdot\mathbf{v}\rangle = 0$ and that the spontaneous emission is isotropic. Equation 4.8 shows that the axial motion can be cooled if the laser frequency is set to the red (negative detuning). The motions in x and y plane then will be heated by the recoil heating due to the spontaneous emission.

The steady state occurs when the right hand side of 4.8 is zero ie.

$$E_{kz} = \frac{1}{2}m\langle v_z^2\rangle = \frac{\hbar[(\gamma/2)^2 + (\omega_0 - \omega)^2]}{6(\omega_0 - \omega)}.$$

When the laser is detuned to the red by half the natural linewidth of the transition i.e. $\omega - \omega_0 = -\gamma/2$, the minimum value of energy

$$\min(E_{kz}) = \frac{1}{6}\hbar\gamma$$

can be achieved. It will be more useful to obtain the rate of reduction of the magnetron and cyclotron radii instead of the E_x and E_y rates. In order to do this consider the case in which an ion with amplitudes r_z, r_m and r_c scatters a photon at time t_0 such that its motion amplitudes change to r'_z, r'_m and r'_c. The trap frequencies would not be affected by this scattering since they only depend on the trap geometry and the trapping potentials. We also assume that this interaction is instantaneous because the cooling transition is usually an electric dipole transition whose decay rate is much bigger than the trapping frequencies ω_z, ω_c and ω_m. Using equations 4.4 and their time derivatives one can show the amplitude change per scattering event is

$$\Delta r_z^2 = (r'_z)^2 - r_z^2 = (\Delta v_z/\omega_z)^2 - (2r_z\Delta v_z/\omega_z)\sin(\omega_z t_0).$$

Similarly for the magnetron and the cyclotron motions we can write

$$\begin{aligned}\Delta r_m^2 &= \frac{(\Delta v_x)^2 + (\Delta v_y)^2}{4\Omega^2} + \frac{r_m}{\Omega}[\sin(\omega_m t_0)\Delta v_x + \cos(\omega_m t_0)\Delta v_y],\\ \Delta r_c^2 &= \frac{(\Delta v_x)^2 + (\Delta v_y)^2}{4\Omega^2} - \frac{r_c}{\Omega}[\sin(\omega'_c t_0)\Delta v_x + \cos(\omega'_c t_0)\Delta v_y], \qquad (4.9)\end{aligned}$$

where $\Omega = (\omega'_c - \omega_m)/2$ and ω'_c is the modified cyclotron frequency. The conservation of momentum requires $\mathbf{\Delta v} = \hbar(\mathbf{k} - \mathbf{k}_s)/m$. The rate of change of the magnetron and cyclotron radii are

$$\frac{dr_i^2}{dt} = \frac{I}{\hbar\omega}\langle\sigma(\omega, \mathbf{v})\Delta r_i^2\rangle,$$

where Δr_i^2 is given by the 4.9. The second terms in 4.9 cancel out in the averaging because the scattering is isotropic. Using the fact that $\mathbf{\Delta v} = \hbar(\mathbf{k} - \mathbf{k}_s)/m$ and the isotropic scattering condition, the rate equations can be written as

$$\frac{d\langle r_m^2\rangle}{dt} = \frac{d\langle r_c^2\rangle}{dt} = \gamma_s R/3m\Omega^2.$$

Now consider a uniform beam along the x axis. Substituting equation 4.4 into 4.7 yields

$$\sigma(\omega, \mathbf{v}) \approx \frac{\sigma_0(\gamma/2)^2}{(\gamma/2)^2 + (\omega_0 - \omega)^2}\{1 + \frac{2(\omega_0 - \omega)k[r_m\omega_m \sin(\omega_m t) + r_c\omega_c' \sin(\omega_c' t)]}{(\gamma/2)^2 + (\omega_0 - \omega)^2}\}.$$

The rate equations of the radii can be written now as

$$\begin{aligned}
\frac{d\langle r_z^2\rangle}{dt} &= \frac{2\gamma_s R}{3m\omega_z^2}, \\
\frac{d\langle r_m^2\rangle}{dt} &= \gamma_s\{\frac{(\omega_0 - \omega)\hbar k^2\omega_m\langle r_m^2\rangle}{m\Omega[(\gamma/2)^2 + (\omega_0 - \omega)^2]} + \frac{5R}{6m\Omega^2}\}, \\
\frac{d\langle r_c^2\rangle}{dt} &= \gamma_s\{-\frac{(\omega_0 - \omega)\hbar k^2\omega_c'\langle r_c^2\rangle}{m\Omega[(\gamma/2)^2 + (\omega_0 - \omega)^2]} + \frac{5R}{6m\Omega^2}\}.
\end{aligned}$$

Negative detuning results in cooling the cyclotron mode while heating the magnetron and vice versa positive detuning results in heating the cyclotron mode while cooling the magnetron. It is impossible to cool both motions with a uniform laser beam in the xy plane. This problem can be solved if we use a non-uniform beam instead. Suppose that the beam is Gaussian instead and that the ion orbit is small compared to the waist of the beam. We can make a linear approximation for the laser intensity at the position of the ion (see figure 4.5)

$$I(y) = I_0(1 + y/y_0).$$

By substituting 4.4 for y into the linear approximation one can show that the new rate equations become

$$\begin{aligned}
\frac{d\langle r_z^2\rangle}{dt} &= \frac{2\gamma_s R}{3m\omega_z^2}, \\
\frac{d\langle r_m^2\rangle}{dt} &= -\gamma_s\frac{\hbar k\langle r_m^2\rangle}{2m\Omega y_0} + \gamma_s\{\frac{(\omega_0 - \omega)\hbar k^2\omega_m\langle r_m^2\rangle}{m\Omega[(\gamma/2)^2 + (\omega_0 - \omega)^2]} + \frac{5R}{6m\Omega^2}\}, \\
\frac{d\langle r_c^2\rangle}{dt} &= +\gamma_s\frac{\hbar k\langle r_c^2\rangle}{2m\Omega y_0} - \gamma_s\{\frac{(\omega_0 - \omega)\hbar k^2\omega_c'\langle r_c^2\rangle}{m\Omega[(\gamma/2)^2 + (\omega_0 - \omega)^2]} + \frac{5R}{6m\Omega^2}\}.
\end{aligned}$$

The first term in the magnetron radius rate equation shows cooling due to the intensity gradient of the laser whereas by contrast the gradient results in heating the cyclotron motion. The second term shows this heating because of the negative detuning of the laser in order to cool the cyclotron motion. The

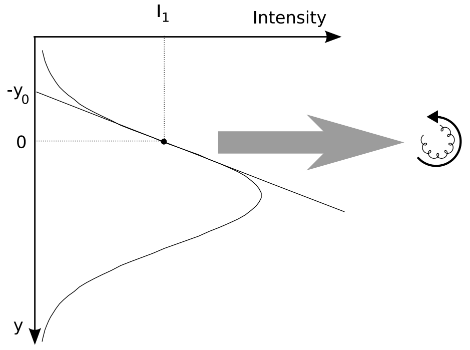

Figure 4.5. If the ion orbit is much smaller than the waist of the beam one can make the linear approximation for the laser intensity. The magnetron motion cooling occurs when the laser beam is offset to the side where the ion is moving along the beam direction.

last term is the recoil heating from the spontaneous emission. If we ignore the recoil heating, cooling both modes for positive y_0 is possible if

$$\omega_m < \frac{(\gamma/2)^2 + (\omega_0 - \omega)^2}{2k(\omega_0 - \omega)y_0} < \omega_c'.$$

Positive y_0 means that the laser beam should be offset to the point where the ion moves in its magnetron orbit along the direction of the beam 4.5. In order to cool all three modes of motions we need to have either another laser along the z direction to cool the axial motion or equivalently use an oblique beam such that there is a component along the z direction as well.

4.5 Axialisation

As mentioned previously, the advantage of the Penning trap over the RF trap is that it uses only static electromagnetic fields and therefore there is no micromotion which can result in the heating of the ion. However the cooling process is much more complex than in the Penning trap. This is due to the characteristics of the magnetron and the cyclotron motions in the radial plane. Because the magnetron motion total energy is negative, making this degree of freedom unstable, energy must be supplied in order to reduce the magnetron radius. Thus the ion should interact with the laser only when the magnetron motion is moving it in the same direction as that in which the laser propagates. As we saw in the previous section the most efficient laser cooling for all degrees of freedom is when the laser beam is directed diagonally to the radial plane to cool the axial motion and offset from the centre of motion to cool the magnetron and cyclotron motions. Even so the cyclotron motion tends to cool much more efficiently than the magnetron motion due to the small gradient of laser intensity across the centre of the trap (see section 4.4). Moreover this puts a limit on the size of the magnetron radius which is related to the size of the laser beam. Axialisation makes the cyclotron and magnetron motions couple so that cooling the cyclotron motion results in cooling the magnetron motion. This makes the localisation of the ion in the Penning trap approach that of an ion in the RF trap but without the micromotion [60].

4.6 Semi-Classical Theory of Axialisation

This section gives a simple description of the axialisation theory which can also be found in [61]. A more detailed treatment of axialisation, which considers the laser detuning as well, can be found in reference [62].

In order to apply this technique the ring electrode is split into four segments. Pairs of diagonally opposing segments are driven with an oscillating voltage at the true cyclotron frequency ω_c*. The RF voltage applied is in the order of a few tens of mV for single ions and therefore the ion, in principle, does not suffer from RF heating. There is a phase shift of π between adjacent segments' voltage. By doing this the cyclotron and the magnetron motions are coupled. Assuming that the voltage applied is $V_d = \sin(\omega t + \phi)$, to a first approximation the electric field can be written as

$$\mathbf{E} = (Cx\sin(\omega t + \phi), -Cy\sin(\omega t + \phi)),$$

where C is a constant that depends on the geometrical shape of the trap and is proportional to the voltage applied to the electrodes. The instantaneous power absorption is equal to

$$\frac{dE_{tot}}{dt} = \mathbf{F} \cdot \mathbf{v} = q\mathbf{E} \cdot \mathbf{v}.$$

Using the equations of motion for an ion in the Penning trap 4.4, the instan-

---

[*] The reader should note that the true cyclotron frequency ω_c is equal to the modified cyclotron frequency ω_+ plus the magnetron frequency ω_- (see section 4.3).

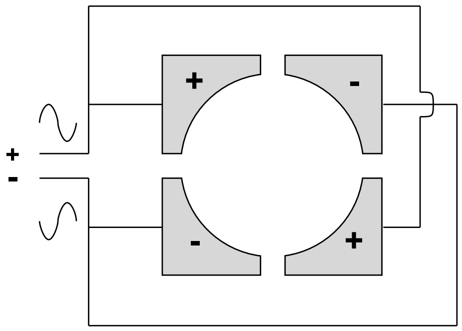

Figure 4.6. For axialisation the ring cap is split into four segments.

taneous power absorption rate can be expressed as

$$\begin{aligned}\dot{E}_{tot} = qC\sin(\omega t + \phi)[R_+^2\omega_+ \sin(2\omega_+ t) + R_-^2\omega_- \sin(2\omega_- t) \\ + R_+R_-\omega_c \sin(\omega_c t)],\end{aligned}$$

in which we have used the relation $\omega_+ + \omega_- = \omega_c$. Of interest to us here is driving the ion at the cyclotron frequency ω_c. In this case, the average power absorption is equal to

$$\dot{E}_{tot} = \frac{qCR_+R_-\omega_c}{2}. \tag{4.10}$$

Axialisation can also be treated quantum mechanically by quantising the harmonic oscillations of the magnetron and the cyclotron motions. Consider that the ion is at the energy level n_+ and n_- of the cyclotron and magnetron motions respectively. The total energy is the addition of the magnetron and the cyclotron energies i.e. $E_{tot} = E_+ + E_-$. Because the magnetron motion has a negative contribution to the total energy, absorption of a photon at the frequency $\omega_c = \omega_+ + \omega_-$ lowers the energy level of the magnetron motion and raises the cyclotron motion by one level

$$(n_+ + \tfrac{1}{2})\hbar\omega_+ - (n_- + \tfrac{1}{2})\hbar\omega_- \rightarrow (n_+ + \tfrac{3}{2})\hbar\omega_+ - (n_- - \tfrac{1}{2})\hbar\omega_-.$$

Similarly the stimulated emission of a photon lowers the cyclotron and raises the magnetron motion respectively by one level. In either case we obtain

$$\dot{E}_+ = \frac{\omega_+}{\omega_-}\dot{E}_-.$$

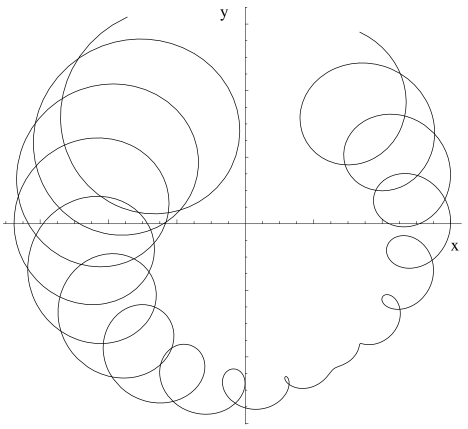

Figure 4.7. The coupling of the magnetron and cyclotron motion by the axialisation. The ion trajectory oscillates between totally magnetron and totally cyclotron.

Using the relation that $\omega_c = \omega_+ + \omega_-$ yields

$$\begin{aligned} \dot{E}_{tot} &= \frac{\omega_c}{\omega_-}\dot{E}_- \\ &\text{and} \\ \dot{E}_{tot} &= \frac{\omega_c}{\omega_+}\dot{E}_+. \end{aligned}$$

This set of equations combined with 4.10 and 4.6 give the following coupled differential equations between the magnetron and cyclotron motion radii

$$\begin{aligned} \dot{R}_+ &= \delta R_- \\ \dot{R}_- &= -\delta R_+, \end{aligned}$$

where $\delta = qC/2m\omega_p$ is called the *beat* frequency. These equations can be easily solved as

$$\begin{aligned} R_+(t) &= R_-(0)\sin(\delta t) + R_+(0)\cos(\delta t) \\ R_-(t) &= R_-(0)\cos(\delta t) - R_+(0)\sin(\delta t). \end{aligned}$$

This resulting trajectory is depicted in figure 4.7. Near resonance the force of the laser cooling in the Doppler limit can be simply taken into account by including a frictional force term $\mathbf{F} = -\gamma\mathbf{v}$. The equations of the cyclotron and

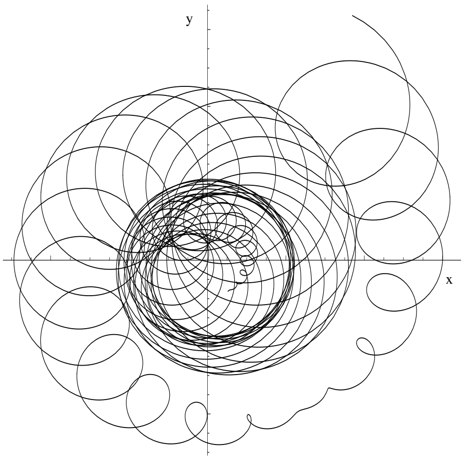

Figure 4.8. Laser cooling the axialised ion leads to smaller magnetron orbit.

magnetron motion can be rewritten as

$$\begin{aligned}\dot{R}_+ &= \delta R_- - \gamma_+ R_+ \\ \dot{R}_- &= -\delta R_+ - \gamma_- R_-.\end{aligned}$$

Here R_+ and R_- are the radii of cyclotron and magnetron motions respectively, γ_+ and γ_- are the corresponding damping coefficients and δ is the beat frequency or coupling constant of the axialisation drive. Using the ansatz that the solutions are in the form of $e^{-\Gamma t}$, the expression for the damping coefficient Γ can be calculated as:

$$\Gamma = \frac{(\gamma_+ + \gamma_-) \pm \sqrt{(\gamma_+ + \gamma_-)^2 - 4(\delta^2 + \gamma_+\gamma_-)}}{2}.$$

Axialisation occurs when both solutions of Γ are real and positive and hence the cyclotron and magnetron motion radii tend to zero as time tends to infinity. Being negative and real is met when either γ_+ or γ_- is positive ($\gamma_+ > 0 > \gamma_-$ or $\gamma_- > 0 > \gamma_+$) such that $\delta^2 > -\gamma_+\gamma_-$. The ratio and sign of the two damping rates can be varied by changing the laser frequency and position.

When the axialisation drive frequency is scanned from below to above the cyclotron frequency a resonance peak is obtained in the ion fluorescence intensity because the maximum laser cooling occurs at this frequency [63] (see figure 4.9). This measurement can be used in determining the cyclotron frequency experimentally.

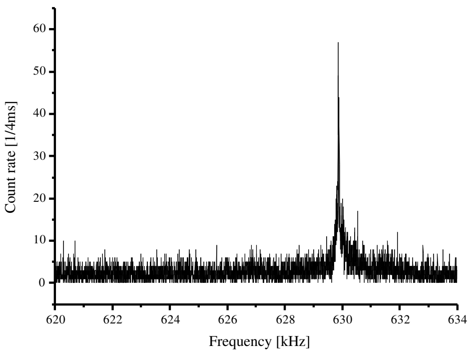

Figure 4.9. The cyclotron frequency resonance when the axialisation driving frequency is scanned from below to above the cyclotron frequency. Taken from [63].

4.7 Scalable Penning Traps

Several proposals for scalable Penning traps for QIP purposes have been put forward in the last few years. By scaling we mean putting a large number of ion traps close together in an array. We also define a trap geometry to be scalable if it can be easily expanded to large number of qubits in such a way so as to allow quantum mechanical interactions between the qubits. A scheme has been proposed to trap electrons within a linear array of Penning traps [64]. The trap consists of several ring electrodes with alternate positive and negative voltages that make trapping regions in between (see figure 4.10a). The gates are performed using microwave techniques rather than optical addressing. Two qubit gates could be implemented by the common axial mode between the two adjacent electrons. However this scheme is scalable only in one dimension and only the two adjacent qubits can directly interact with each other. To address this issue Stahl et al [65] proposed a new scheme which consists of an array of planar Penning traps that could be linked together via superconducting wires. The trap consists of a circular disk as the endcap which is concentric with a ring electrode on a substrate. Contrary to the conventional Penning trap the open geometry allows good optical access (see figure 4.10b). A trap based on this geometry has been built and trapped electrons have been successfully detected [66].

Another design which has been recently proposed by J.R. Castrejon-Pita and R.C. Thompson [67] uses two perpendicular sets of three parallel rods providing good optical access (see figure 4.10c). By applying electric potentials on the outer and central wires axial confinement can be achieved. The magnetic field perpendicular to the wires provides the confinement in the other

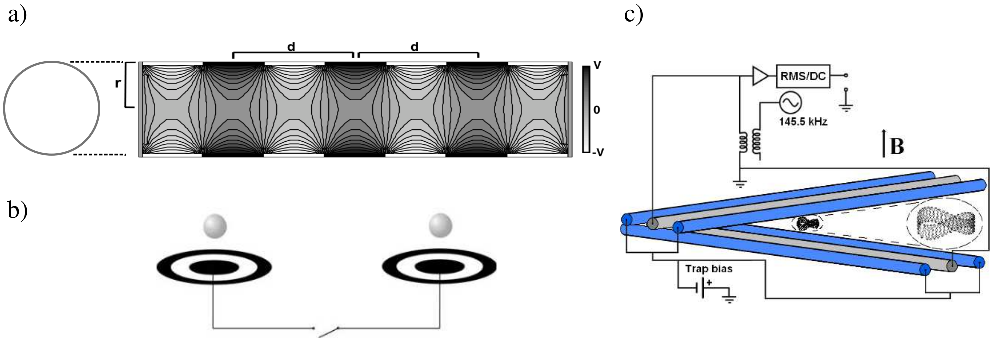

Figure 4.10. Recent proposals for the scalable Penning trap is shown. Images taken from [64, 65, 67].

two dimensions. The ions can adiabatically "drift" in a desired direction along the electric field lines by a kind of modified cyclotron motion. This trap has been built and trapped ions detected electronically.

In this section I briefly introduce "Pad traps", another scheme for QIP in the Penning trap that we have recently proposed [6]. In this scheme ions are "hopped" from one trap another. In order to see how this can happen first consider the motion of an ion in crossed homogeneous electric and magnetic fields. The magnetic field points in the positive z direction and there is a homogeneous electric field in the positive y direction. The total force on the ion could be described by the Lorentz force $\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$. Supposing the ion is initially at rest at $t = 0$ it follows a trajectory which can be described by the parametric equations

$$x = -\frac{V_D}{\omega_c} \sin \omega_c t + V_D t,$$

$$y = \frac{V_D}{\omega_c}(1 - \cos \omega_c t),$$

where $V_D = E/B$ is the "drift velocity" and $\omega_c = qB/m$ is the cyclotron frequency. The trajectory as shown in figure 4.11 is a series of cycloid loops with a constant drift in the x direction. Now consider a pair of traps whose axes are along the z direction but the centres are displaced in the positive x direction. If the trapping electric potential can be replaced with a homogeneous electric field, the ion can be "hopped" from one trap to the other in a single cycloid loop. After that the trapping voltages will be applied again. A single loop takes a time $t_{hop} = 2\pi/\omega_c$ to complete. For $^{40}\text{Ca}^+$ in a magnetic field of 1 tesla we have $t_{hop} = 2.6\mu\text{s}$.

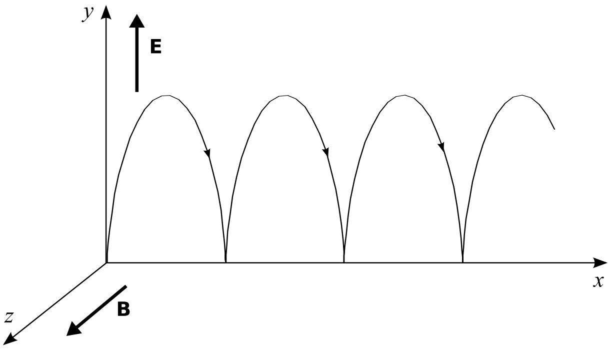

Figure 4.11. The cyclotron motion of a positive particle in a linear electric field in the presence of a magnetic field is shown. The particle comes to rest at $y = 0$.

The separation of the two traps should be equal to $x_l = (E/B)t_{hop}$. In order to have 5mm displacement in a magnetic field of 1 tesla we need to apply the electric field of $E = 1900\text{V/m}$ which is on the order of 20V for a trap with the characteristic dimension of 1cm. We need to have the capability to switch a few tens of volts faster than $0.1\mu\text{s}$. The ion comes to rest again after one cycloid loop, therefore in principle the ion is not heated after the "hopping". An individual trap is composed of a hexagonal array of 6 pads arranged in a ring form with one pad in the centre acting as the endcap, with an identical set of pads over it (see figure 4.12). The geometry is complicated in order to allow to 'hop' the ion from one trap to another in a cycloid loop. By applying positive voltages on the endcaps relative to the ring, a quadrupole potential is generated in the centre. Because of the hexagonal geometry of the trap it could be scaled to two dimensions easily.

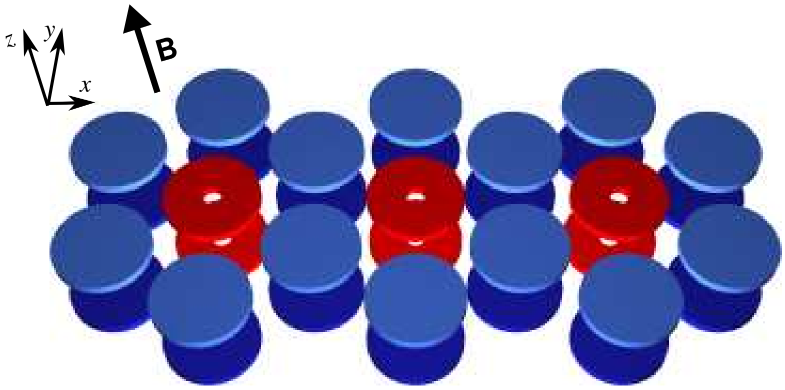

Figure 4.12. Schematic layout of three pad traps.

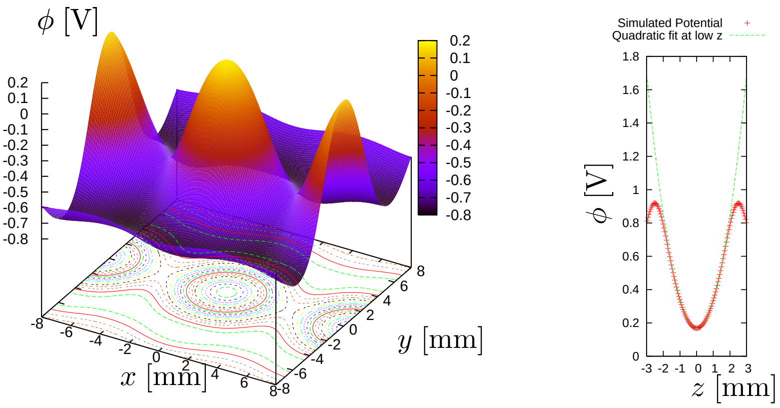

(a) (b)

Figure 4.13. a) Plot of the potential in the radial plane at the midpoint between the substrates. The potential is circular close to the centre of the trap. b) The potential along the axis for $\gamma = 0.9$.

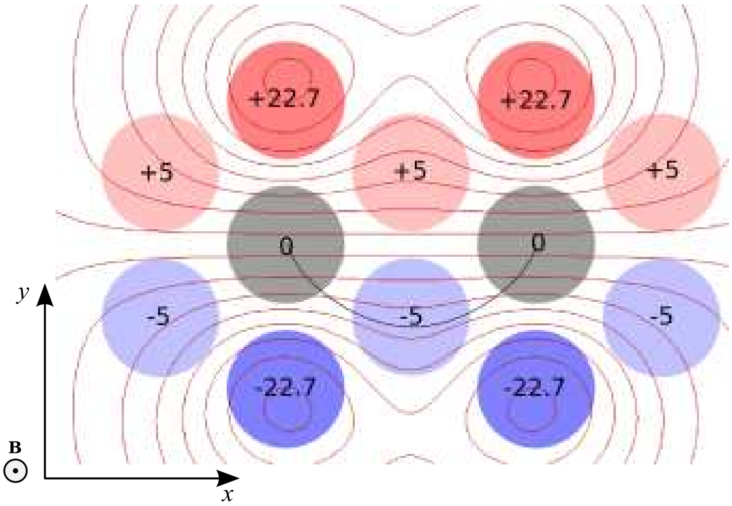

Figure 4.14. A nearly linear electric field can be generated by applying the voltages shown. The ion in the left-hand trap is initially at rest. After switching the voltages it hops to the right-hand trap and comes to rest again.

The pads are chosen to have a radius of 4mm and separated by 5mm from centre to centre. The ratio of the separation of the pads to the distance between the two layers γ is optimised to achieve the closest potential to a quadrupole. The potential turns over close to the endcaps (due to the holes in the endcap electrodes) but matches very well to a harmonic potential in the middle. This is shown in figure 4.13(b). Figure 4.13(a) shows the potential in the radial plane at the midpoint between the substrates. In an ideal Penning trap the contours projected on the x, y plane are circular. The deviation from circular symmetry in the pad trap is small over the trapping region despite the non-ideal trap geometry and the effect of the adjacent traps.

In order to 'hop' an ion from one trap to another a nearly linear electric field is applied in the region between the plates. This is achieved by applying a gradient of voltages on the electrodes. The voltages are set such that a single ion 'hops' in a single cycloid loop from one trap to another in a magnetic field of 1 tesla. It is found that despite the electric field not being perfectly linear it is still possible to drive the ion in a cycloid-like loop provided that the electric field strength is such that the ion does not have any y-component of velocity at the midpoint of the two traps. The voltages are symmetric about the x axis with a gradient that determines the amount of displacement. The fact that the field need not be strictly linear allows an extra degree of freedom in setting the voltages. Voltages may be chosen that impose a degree of axial confinement during the hop. The required voltages to achieve this are shown in figure 4.14. In order to demonstrate the axial confinement, plots of the axial position of the ions z are shown in figure 4.15 for a range of velocities in different directions.

During the 'hopping' the ion should be confined in the axial direction. If the ion has an initial velocity along the axis it will be displaced after the hopping. The axial confinement is assured by optimising the voltages such that there is a minimum of potential at $z = 0$ all the way along the trajectory of the ion in the radial plane. This can be shown if we define s as the total displacement along this trajectory and plot the potential along the z axis as a function of s as depicted in figure 4.16. The depth of this potential varies along the ion's trajectory so the z motion is not harmonic during the 'hop'. Nevertheless the minimum ensures that the ion re-focusses into the left-hand trap or in another word the ion ends up in the center of the left-hand trap. The effect of the anharmonicity is visible in the top chart in figure 4.15 with high energy ions spreading along the axis after the hopping. The voltages are

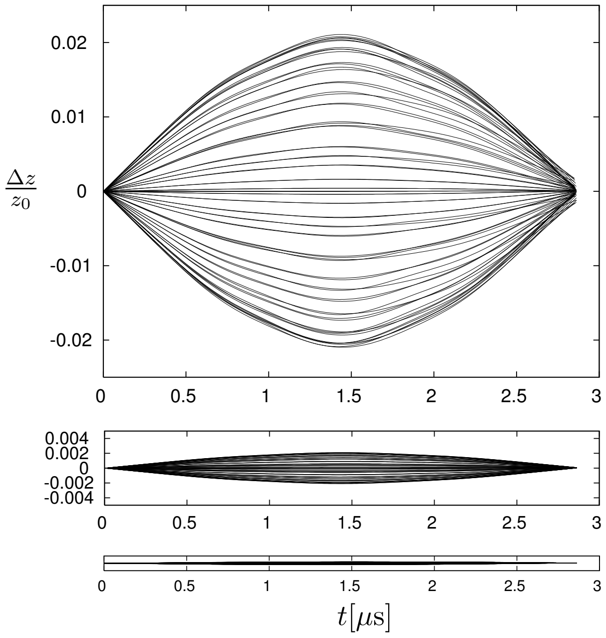

Figure 4.15. Simulation (performed using SIMION [68]) of the axial position of the ion as a function of time for a single cycloid for many ions. The ions start at the centre of one trap and have velocities in different directions. The top panel is for ions with a kinetic energy of 10 meV(≅ 100K). The dispersion of the ions after the hop is noticeable. The middle panel is for 0.1 meV(≅ 1K) and the lower panel is for 0.001 meV (≅ 10mK). The ions almost do not heat up (disperse) at these temperatures after hopping. The initial direction of velocity is varied over the full range of azimuth and declination in steps of 10 degrees.

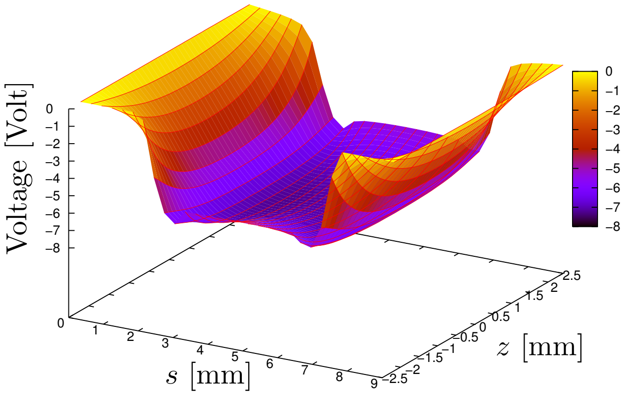

Figure 4.16. The potential in the z direction evaluated at equidistant points along the ion hopping trajectory. The ion sees a minimum along the axis during its hopping.

further adjusted such that the spreading is minimised.

In order to perform quantum gates a conventional cylindrical Penning trap

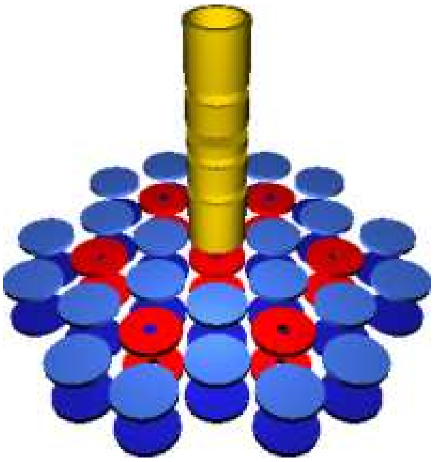

Figure 4.17. An array of 7 padtraps with a conventional cylindrical trap which sits on the top of a padtrap is shown. In this arrangement it is possible to move around ions and perform two-qubit gates by injecting them into the cylindrical trap.

will be positioned on the top of a specific pad trap. The ion can be extracted from the pad trap and injected into the cylindrical Penning trap to perform two-qubit gates (see figure 4.17). The whole array of the ions could be shifted in two dimensions to allow for a different ion to be extracted. The trap can be miniaturised to make the 'hopping' time less. Faster 'hopping' induces less heating as well as faster shuttling of the ions. Small traps have also the advantage of having more homogeneous magnetic field. The magnetic field is never completely homogeneous along the whole pad traps array, nevertheless it does not change with time. As a result it is possible to characterise each trap with its own set of trapping parameters and laser frequencies. The limit to the size of the trap is the radial confinement of the ion which is dependent on the strength of the magnetic field. This can be improved further using the axialisation technique (see section 4.5).

Work on making a padtrap out of circuit board is currently being pursued by another member of the group (D. Crick). The initial design contains three traps with hexagonal shape rather than the circular geometry that we presented here. This is to minimise the non-coated surfaces of the board in order to avoid possible charging of the surfaces by the electron emitted by the filament [*]. The simple planar geometry makes the manufacture very simple. The electrodes have been made of commercially available vacuum compatible

---

[*] The filament is needed for ionisation. More details are available in section 5.3 and 5.10.

circuit boards (FR-4 [*]). We have constructed a prototype trap using this technique. Both sides of the board are copper plated. The printed circuit board (PCB) is then milled with a PCB milling machine. The electric connections are made using Kapton wires at the back of the board. This can be achieved by drilling the board with a tiny hole ($\sim 200\mu$m) and then copper plating the wall of the hole to make a connection between the two sides of the board. This process which is called *through-hole plating* minimises the distortion of the electric potential. There is a 1mm hole at the centre of the endcaps for the ionisation of atoms with electron beams. The filament is placed behind the endcap. The hole could also provide the means to load and extract a single ion along the axis into another trap in which a two qubit gate can be performed, though the cylindrical trap required for this operation has not been included in the first prototype.

---

[*] Flame retardant 4 (FR-4), is a type of material used for making printed circuit boards which is ultra high vacuum (UHV) compatible.

CHAPTER 5

Initial Status of the Experiment

5.1 Introduction

Before I started my PhD, the ion trapping group in Imperial College had been working with ^{40}Ca$^+$ for 3 years. A split-ring Penning trap had been built and had successfully trapped Ca$^+$ ion clouds. Laser cooling had been achieved using two home built extended cavity diode lasers (ECDL) in the infrared along with two commercial ECDLs in the UV. Four cavities had been built and temperature stabilised to which the lasers were locked. To ease the problem of finding the wavelengths of the cooling lasers a wavemeter had been developed. The group had also implemented the axialisation technique to improve cooling of the magnetron motion. Furthermore, with the magnetic field switched off and the trap run as a radio-frequency trap the group were able to work with single Ca$^+$ ions. However, for reasons that were unclear at the time, the experiment was unable to trap a single ion in the Penning trap.

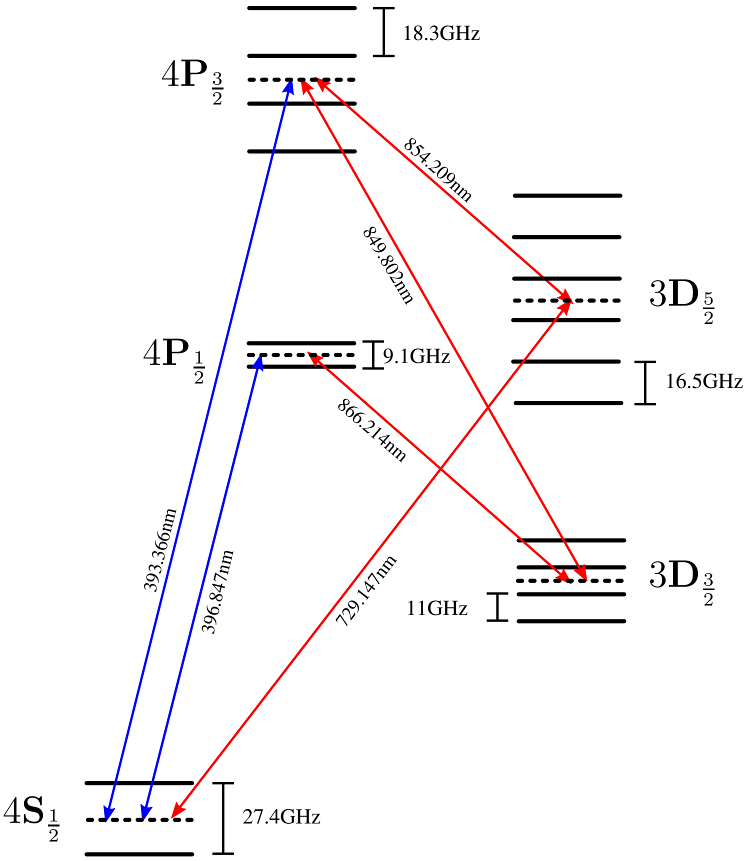

Figure 5.1. Ca^{+} energy levels at the magnetic field of 0.98T is shown.

This chapter describes briefly the status of the experiment at the time that I joined the group.

5.2 ^{40}Ca^{+} Energy Levels

The energy levels for ^{40}Ca^{+} in the presence of a magnetic field are shown in figure 5.1. Dotted lines show the degenerate energy levels in the absence of magnetic field. Due to the magnetic moment of the electron, when a magnetic field is introduced the levels Zeeman split. The Zeeman shift for alkali-like ions with zero nuclear spin is proportional to the magnetic field for weak field*. The

---

*By weak field we mean a magnetic field much smaller than the internal magnetic field which is responsible for the fine structure. The magnetic field which gives rise to the energy difference of $P_{1/2}$ and $P_{3/2}$ in Ca^{+} ion is ∼200T. A typical magnetic field in the Penning trap is in the order of a few tesla.

shift in energy is equal to

$$\Delta E = m_j \mu_B g_J B \ ,$$

for a magnetic field of B. Here $\mu_B = e\hbar/2m_e$ is the Bohr magneton for an electron in SI units and g_J is the Landé g-factor. For the orbital angular momentum of l and total angular momentum of $\mathbf{J} = \mathbf{L} + \mathbf{S}$ the Landé g-factor is given by

$$g_I = 1 + \frac{j(j+1) - l(l+1) + s(s+1)}{2j(j+1)} \ .$$

The $\mathrm{S}_{1/2} \rightarrow \mathrm{P}_{1/2}$ transition at 397nm is an allowed dipole transition (A $\simeq$ $1.4{\times}10^{8}\mathrm{s}^{-1}$) which is fast enough for efficient laser cooling. The ion in the $\mathrm{P}_{1/2}$ state can decay to either of the two $\mathrm{S}_{1/2}$ Zeeman sublevels and therefore two lasers need to be used for the cooling transitions to avoid optical repumping into a 'dark' Zeeman sublevel of the ground state. As shown in figure 5.2(a) we could either use the π or σ transitions. Due to mostly historical reasons we use the π transitions however π transitions do have the advantage of being less sensitive to magnetic field fluctuations.

Once in the $\mathrm{P}_{1/2}$ state the ion can decay into the metastable $\mathrm{D}_{3/2}$ state with a probability of 1/13 [69]. The $\mathrm{D}_{3/2}$ state has a lifetime of ~1s and therefore in order to avoid the ion spending a large fraction of its time outside of the cooling cycle the ion needs to be pumped back to the $\mathrm{P}_{1/2}$ state. There are 6 transitions from $\mathrm{D}_{3/2}$ to $\mathrm{P}_{1/2}$ and at least 4 of these transitions must be addressed as shown in figure 5.2(b) to re-pump all four metastable sublevels of $\mathrm{D}_{3/2}$. We use the four σ transitions. In order to address the sigma transitions either we need to use four lasers at 866nm or alternatively use two lasers with sidebands whose carrier frequencies are at the middle of each of the sigma pairs (see dotted line in figure 5.2(b)). Sidebands are imposed by modulating the current of a diode laser. Because the Zeeman splitting of the levels is large ($\sim$ 900MHz), the intensity ratio of the sidebands to the carrier cannot be higher than $<$ 10% without laser instabilities. Using four separate lasers has the advantage that one gets more power but on the other hand it requires two more stable lasers which is more experimentally demanding. At the time that I joined the group it had been decided to use the sidebands approach for obvious reasons. The two pairs of sigma transitions are separated in frequency by 1.83GHz. Two diode lasers at 866nm whose currents were modulated at 915MHz were used to address each pair of the sigma transitions.

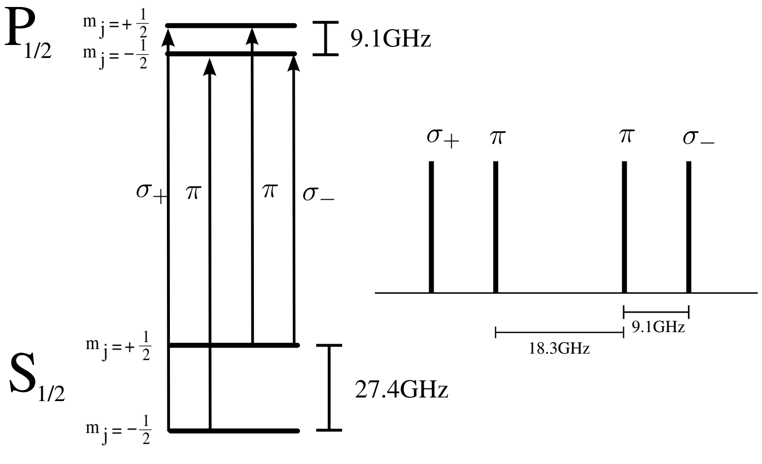

(a)

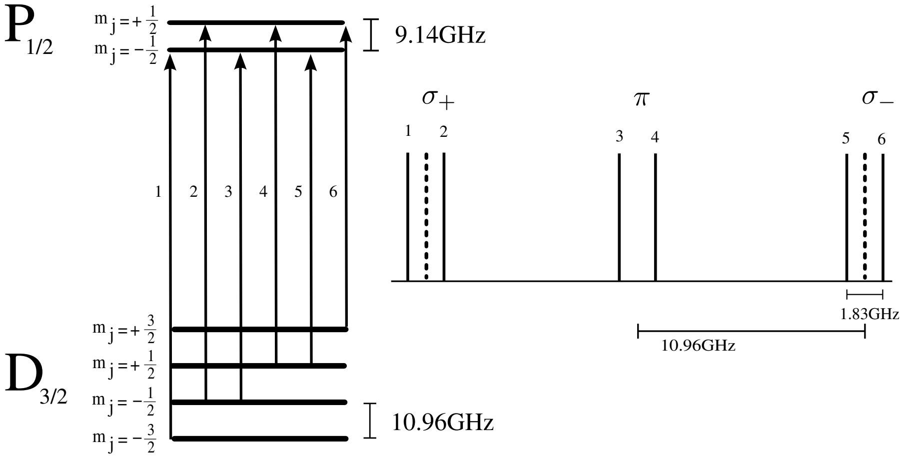

(b)

Figure 5.2. a) The $S_{1/2} \rightarrow P_{1/2}$ transition. The π transition is the least affected by the magnetic field fluctuations. b) The $D_{3/2} \rightarrow P_{1/2}$ re-pumping transitions. The dotted lines are the carrier frequency of the IR lasers at 866nm.

Amplified spontaneous emission (ASE) is light produced by spontaneous emission which is optically amplified in a single pass through the gain medium of the laser. ASE contributes to the noise of the laser because it is not coherent with amplified stimulated emission of the laser and it is not controllable. The ASE wavelength range is as wide as the gain profile of the laser itself. The ASE in the UV lasers could resonantly drive the $S_{1/2} \rightarrow P_{3/2}$ transition at 393nm. The ASE in the IR lasers could also drive the $D_{3/2} \rightarrow P_{3/2}$ transition at 850nm. The ion then decays to the metastable $D_{5/2}$ state and therefore an additional laser at 854nm is needed to re-pump this state. The amount of the ASE light in our experiment is very low and we found that only one laser at 854nm suffices to re-pump all the 6 metastable levels and it is not very sensitive to the frequency of the laser compared to the 866nm repumping lasers (see section 8.3). We believe that the ASE light comes mostly from the UV lasers although we have not shown this conclusively.

Finally the $S_{1/2} \rightarrow D_{5/2}$ quadrupole transition at 729nm will be used for sideband cooling as well as addressing the qubit transition. ASE light is 'useful' in early stages of the experiment since it allows quantum jumps to be seen and by looking at the quantum jumps one can infer the number of ions in the trap. However, ASE needs to be removed from the laser beams to preserve the coherence of the qubit[*]. The $D_{3/2} \rightarrow P_{3/2}$ transition at 850nm can be used to induce more quantum jumps to the $D_{5/2}$ metastable state and measure the number of ions in the trap.

5.3 Trap Structure

The trap was originally designed by M. van Eijkelenborg for laser cooling Magnesium ions. Technical drawings can be found in reference [70]. Later the end-cap holders were re-machined to give more freedom for lateral and angular adjustments [71]. A schematic diagram of the trap is shown in figure 5.3. The electrodes are triangular in cross section rather than hyperbolic to allow easier machining. Although different from an ideal Penning trap the potential in the middle of the trap matches closely to a quadratic potential [72]. The trap is made of beryllium-copper with the characteristic dimensions of r_0 =5mm and $2z_0$ =7mm (see figure 4.2(a)). The ring is segmented into four sections with a 1mm gap in between, to allow axialisation and also radial excitation of

---

[*] This can be done by passing the light through a prism and coupling it into a single mode optical fibre at some distance from the prism or alternatively use a grating.

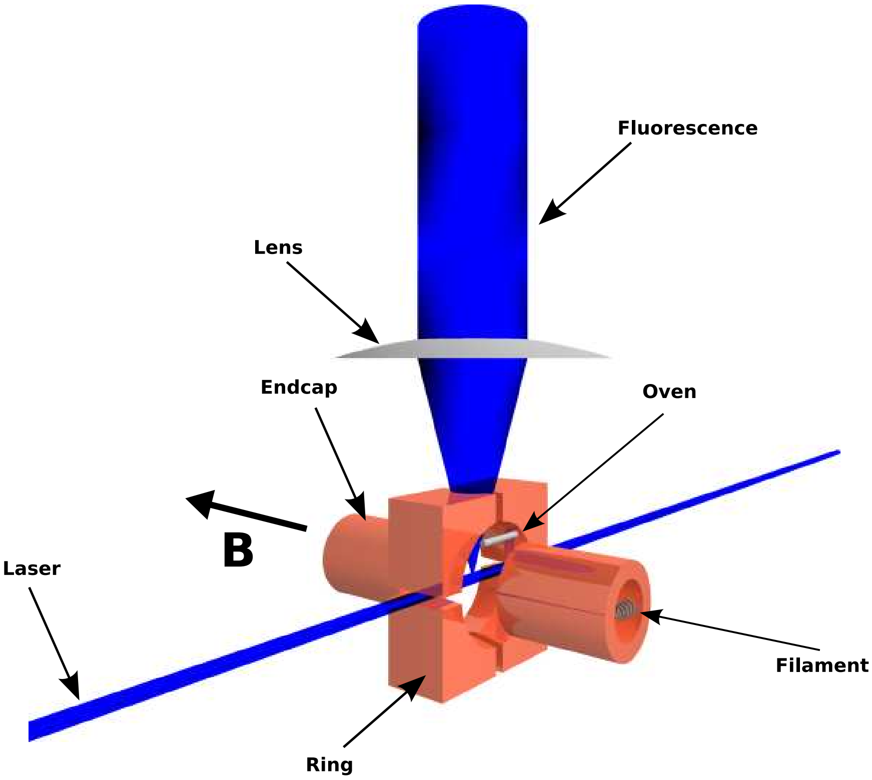

Figure 5.3. Schematic diagram of the split-ring trap (conventional magnet trap).

characteristic trap motions. The gaps between the ring segments are too small for the input laser. The laser also needs to exit the trap cleanly otherwise it scatters light into the detection system. Therefore in order to reduce the laser scattering, two holes are made between the lateral segmented rings. There is also a hole in the upper segments of the ring to increase the solid angle of imaging. A similar hole is made between the lower segments to preserve the symmetry and also to reduce any other scattering that could be imaged onto the detection system. The ionisation is achieved by collisions of Ca atoms with electrons generated from a hot filament. After the electrons are generated they become focused along the magnetic field in a cyclotron motion. Therefore the filament needs to be placed along the axis of the trap. The endcaps are bored to allow the placement of the filament inside. The electrical connections are made via 5A-rated Kapton coated wires and whenever there was a need to make an electrical insulation Macor and Alumina screws were used. Macor, which is a machinable glass ceramic, although very fragile, is ultra high vacuum (UHV) compatible.

The filament is a 200μm diameter thoriated tungsten wire which is coiled in the middle. Thoriated tungsten has a work function of 2.6eV as opposed to that of tungsten which is 4.5eV. As a result it has a more efficient thermionic emission at a given temperature. The ovens are made of 1mm diameter tantalum tubes which are spot welded to a tantalum wire in several places. The oven is filled with lumps of high purity Ca (99%) and the tube is crimped closed at each end. By passing current through the wire the oven heats up and a spray of Ca atoms comes out of a small hole made in the middle of the oven.

5.4 Vacuum System

The pressure inside the trap should be kept at ultra high vacuum (UHV), otherwise the ions are heated or ejected from the trap by collisions or charge transfer with the background gas. The full procedure to get to the required UHV is described in [71, 73]. Copper surfaces oxidise in air and this makes the surface have poor electrical conductivity. Therefore all copper pieces need to be cleaned using “Copper Brite” solution*. To get to the required UHV, before assembling the trap all the parts of the trap including screws, ceramics, Macor, Alumina screws and Kapton wires, are washed first with acetone and then with methanol in a warm ultrasonic bath.

A schematic of the vacuum apparatus is shown in figure 5.4. To get to the desired UHV pressure the trap is first pumped with a turbo pump which is backed by a rotary pump. This gets us down to $\sim$ few $\times\, 10^{-7}$ mbar pressure after 2-3 days. After that the whole vacuum system is placed in an oven and is baked to $\sim$ 200°C. The baking speeds up the ‘outgassing’ process of the contaminations on the surfaces. After baking for about a week the trap is cooled down to room temperature and the pressure goes down to $\sim 10^{-9}$ mbar. At this stage the ion pump is turned on and after 2-3 days the pressure will be as low as $\sim 2 \times 10^{-10}$ mbar.

---

[*] Copper Brite is composed pf 75% concentrated nitric acid, 23% concentrated sulphuric acid, 2% concentrated hydrochloric acid. The solution is very strong and should be diluted to 1:10 with distilled water.

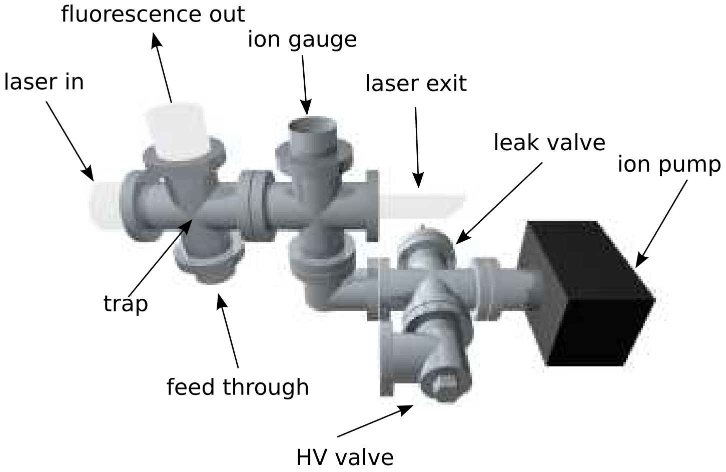

Figure 5.4. Schematic diagram of the vacuum system.

5.5 Magnetic Field

The magnetic field is provided by a conventional electro-magnet (Oxford Instruments N100 A38 TIP) which is driven by a current-stabilised power supply (KSM SCT-220-20). The closer the poles of the magnet are the higher the magnetic field becomes. Therefore the trap has been designed to be as small as possible to allow for a stronger magnetic field. The poles are 4cm apart and each coil has a cold resistance of 4.5 ohms. The pole coils are in series. We run the trap at 0.98 tesla. To get to this magnetic field 15 amps continuous current is passed through each coil. The resistance of each coil goes up to 5.5 ohms at 15 amps as they get to the final temperature. The heat dissipation is ~ 2.5kW and as a result cooling water is necessary.

Because the resistance of the coils could vary in time due to temperature fluctuations, the current must be stabilised. The power supply can stabilise the current to 1 part in 10^4. There is a 1mΩ shunt in series with the coils and by measuring the voltage across it we can measure the current to ± 0.1 amps precision. The power supply is floated at ~ 160V above ground.

5.6 Detection System

The fluorescence of laser cooled ions is collected and transmitted to the photomultiplier tube (PMT) using three quartz lenses in a symmetrical arrangement as shown in figure 5.5. The first lens is inside the vacuum chamber. It is fol-

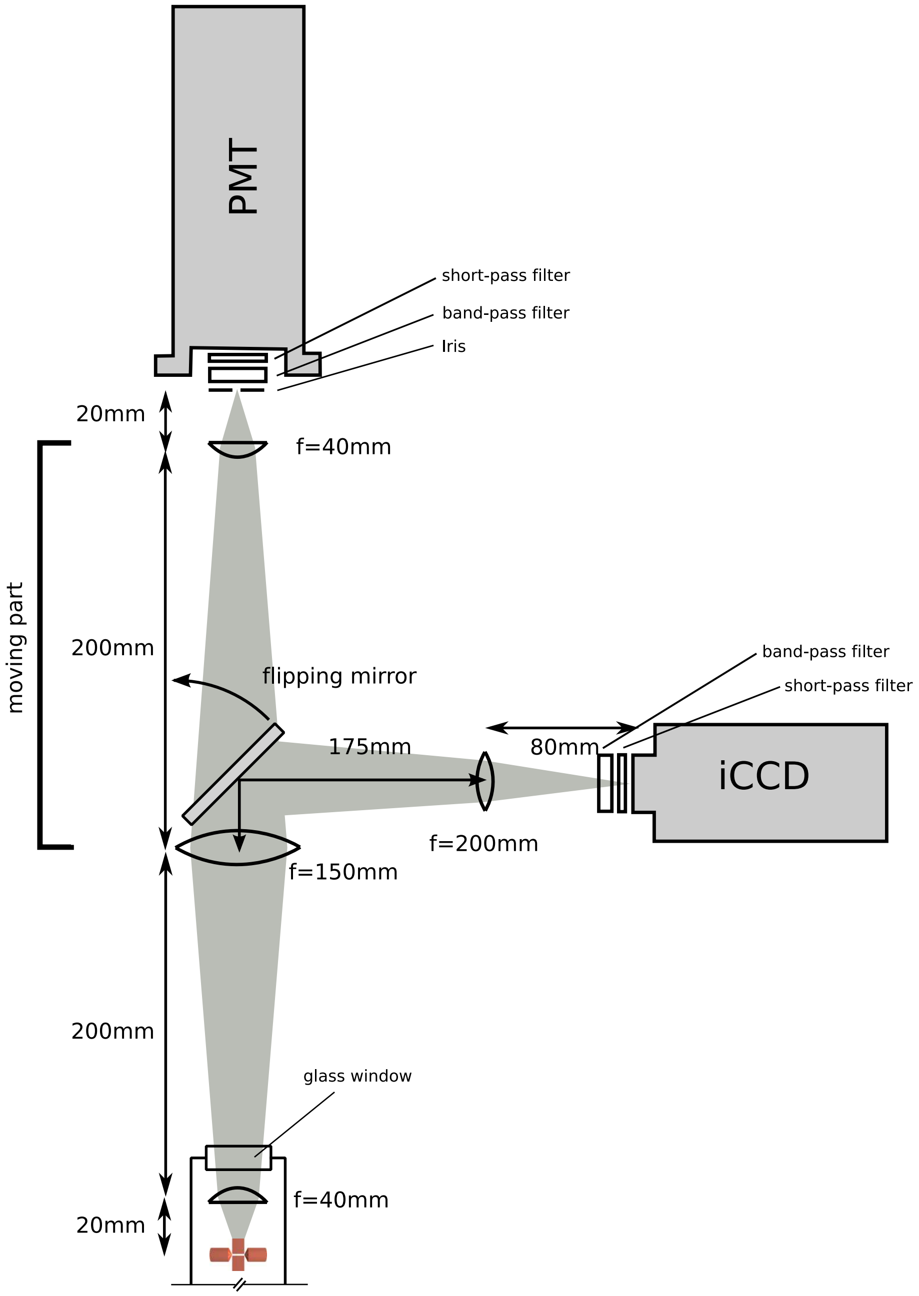

Figure 5.5. The imaging system to the PMT and the iCCD is shown. The magnification to the iCCD camera is calculated to be ∼1.7. The flipping mirror switches the fluorescence to the iCCD camera or PMT.

lowed by two other lenses outside the vacuum chamber which are attached to an aluminium block. The aluminum block can be moved vertically by a translation stage. This symmetrical arrangement was first designed to be used with 280nm light from Mg^+ ions. This means that the focal length of the first lens is greater than 40mm at 397nm and as a result the light is not focused as strongly as it should be. However it was found that this effect was not noticeable in the count rate detected by the PMT if the second lens was placed closer to the trap [73]. Fluorescence is focused into an iris in order to stop scattered laser light, background light from the room and also the filament light from entering the PMT. The focus at the iris is adjusted by moving the block. The PMT is mounted on two translation stages to allow movement in the horizontal plane. To get lower background count rate a short-pass dichroic filter (Comar 435IK25) and a band-pass glass colour filter (Comar 395GB25) are placed after the iris to block the residual background light from the room and the filament.

The detection efficiency is limited by a range of factors. The first is the hole in the frame that carries the ring electrode. The hole has a diameter of 5.5mm and is 13mm from the centre of the trap. This gives a solid angle of $\Omega = 0.14$ steradians ($\sim 11\%$ of 4π). The transmission of the quartz lenses and the trap window are each 92% at 397nm. The filters have a transmission of 0.8 each which gives transmission T=0.46 in total. The PMT has a quantum efficiency of $\epsilon = 0.23$ at 397nm. This gives the total efficiency of

$$\eta = \epsilon \times T \times \frac{\Omega}{4\pi} = 1.2 \times 10^{-3}.$$

The scattering rate of the $S_{1/2} \rightarrow P_{1/2}$ cooling transition is $1.4 \times 10^8 s^{-1}$ [69]. The maximum laser cooling occurs when the laser is red-detuned by half the natural linewidth of the cooling transition which reduces the scattering rate by a factor of 2. Cooling the magnetron motion requires the laser beam to be offset from the centre of the trap which reduces the scattering rate by another factor of 2. If all the $S_{1/2}$, $P_{1/2}$ and $D_{3/2}$ levels are saturated the ion spends an equal amount of time in each of the eight states. In another word the ion spends a maximum of 1/4 of its time in the two $P_{1/2}$ states. Therefore the maximum scattering rate possible would be $8.75 \times 10^6 s^{-1}$. Considering the calculated efficiency we expect a maximum of $\sim$ 10k count/s for a single ion.

The PMT (Thorn-EMI 9893QB) is connected to an amplifier discriminator (Electron Tubes AD2). The amplifier-discriminator gives an output pulse of fixed width and amplitude whenever the photomultiplier input pulse exceeds a set threshold level (0.5mV-2.5mV). The signal is then passed through a emitter-coupled logic (ECL) to transistortransistor logic (TTL) pulse converter (Electron Tubes ET1). The TTL signal is connected to a EG&G Ortec Multi Channel Scaler (MCS) that measures the TTL count rate. The MCS is controllable by the computer and the time bin width is adjustable. The dark count rate of the PMT when running at 2.3kV is $< 10\text{s}^{-1}$. The background count rate rises due to the scattering of the blue lasers inside the vacuum chamber. The background count rate can be as low as 10s^{-1} when the laser intensity at 397nm is optimised to get the best signal to background for a single ion (see chapter 8).

There is a mirror on a flippable mount that can direct the fluorescence to an image intensifying CCD camera (Andor DH534-185-03). The CCD head has 1024×1024 pixels with $(13\mu\text{m})^2$ pixel size. The quoted average spatial resolution (full width at half maximum) is 22μm. The maximum gain of the intensifier is $\sim$ 4000 counts per photo electron with a maximum quantum efficiency of 19.8%. The dark current decreases exponentially with the temperature of the CCD. The minimum dark current achievable is quoted as 0.057 electrons/pixel×sec at $-25°$C. The CCD head is cooled with a Peltier heat-pump supported by water cooling. The resolution of the iCCD camera before shipping was quoted to be 22μm. The image magnification to the iCCD head is calculated to be 1.73 at 397nm.

5.7 Lasers

5.7.1 Diode Lasers

As we saw in section 5.2, in order to laser cool Ca^+ ions in the Penning trap we need to have at least 4 lasers. Diode lasers are very convenient for this purpose because they are reliable and relatively inexpensive. Free-running diode lasers have poor tunability and large linewidth ($\sim$100MHz). However these problems can be dramatically improved using optical feedback to the laser diode. This can be achieved either by coupling the light to a high Q cavity (linewidth <10kHz) or using a diffraction grating (linewidth $\simeq$ 1MHz). The Ca^+ dipole transition natural linewidth is 24MHz. Therefore optical feedback by diffraction grating suits us well. Grating feedback based on Littrow's design is shown in figure 5.6. The first order diffraction of the light reflects back from

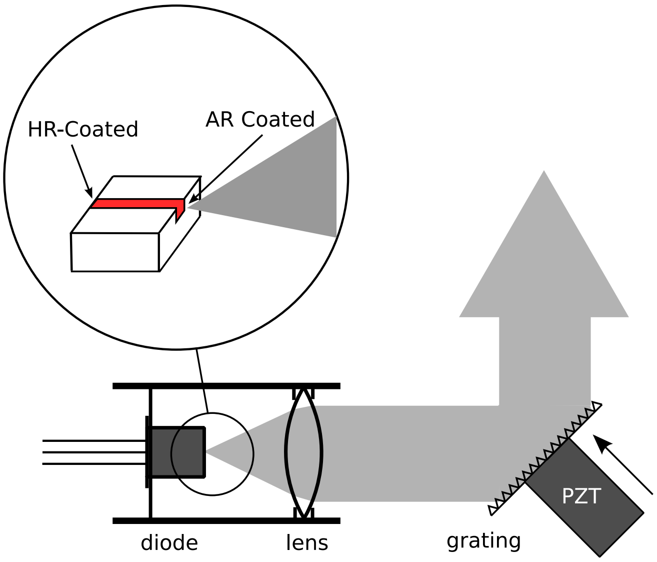

Figure 5.6. The optical feedback based on Littrow's design. The angle of the grating can be adjusted coarsely by two screws whereas the fine adjustment of the cavity length is achieved by a Piezoelectric transducer (PZT) at the back of the grating. The diode chip front facet is anti reflection (AR) coated to reduce optical feedback whereas the rear side has a high reflection coating. The anti reflection coating is important because it improves the stability of the laser and increases the mode-hop free tuning range. The lens is a high numerical aperture AR coated aspherical lens.

the grating and couples to the diode. The rear facet of the diode and the grating form an external cavity. The light will only resonate at the longitudinal mode frequencies of this cavity.

The frequency of the laser can be scanned by changing the length of the cavity. Fine adjustment of the cavity length can be achieved using a Piezoelectric transducer (PZT) mounted at the back of the grating. The manufacturing process of diode chips makes it impossible to set the wavelength to the exact atomic wavelength that is required. However the precise wavelength of the laser can be set by the angle and longitudinal position of the grating. Temperature also changes the free-running wavelength of the laser by a factor of 0.3nm/K [74]. It also changes the cavity length and as a result the whole setup needs to be temperature stabilised. The diode chip itself acts as an etalon inside the external cavity which restricts the mode-hop-free tuning range. The anti reflection coating of the front facet of the chip reduces this unwanted effect. The forward current to the diode changes the temperature of the diode as well as the intensity of the laser.

We use two extended cavity diode laser (ECDL) designs based on Lit-

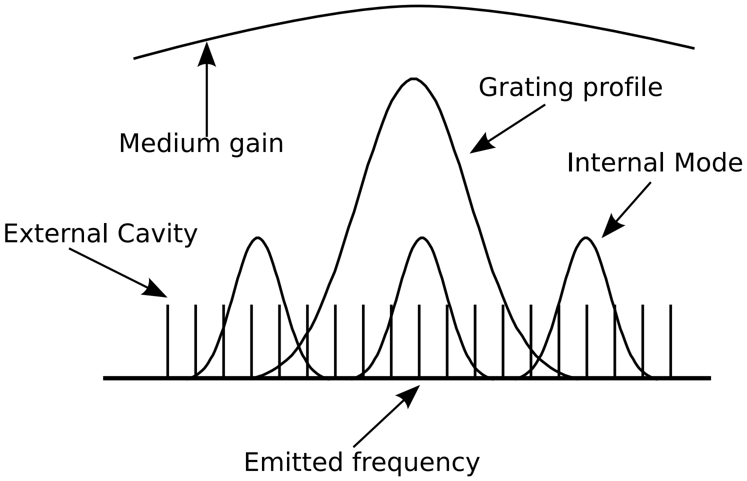

Figure 5.7. The ECDL gain profile is shown. It is usually possible to run the laser at a specific wavelength by changing the combination of temperature, grating, current and Piezo.

trow's configuration. The first one is the Hänsch design [75] on which our two commercial blue lasers (Toptica DL 100) are based. The second design is Boshier's design [76] on which the infrared home-built diode lasers are based. The Hänsch design has better stability since the grating is mounted on a monolithic piece of metal whereas the Boshier design uses a commercial mirror mount with springs and is easier to build. The screws of the mount can relax over time. This is particularly noticeable when the laser is operated at elevated temperature.

At the time that I joined the group we were using only two lasers for the repumping infrared lasers at 866nm. The necessary addressing of 4 transitions was achieved by modulating the current of the diodes at 900MHz to impose sidebands on the laser output. The four sidebands deplete the metastable $D_{3/2}$ states on the four σ transitions seen in section 5.2. The full details of the setup can be found in reference [77]. However due to problems that I describe later on, we are now using four separate lasers instead. An extra laser at 854nm is also used to repump the $D_{5/2}$ state.

The laser diodes for the repumping lasers in the infrared are from a range of manufacturers. We are using SDL5411 at 852nm, Laser 2000 LD1377 at 860nm and Eagleyard EYP-RWL-0850 at 850nm. In order to shift the laser output to the desired wavelength the diode is heated to 30-50°C by a resistor attached to the diode mount. The pulling range of the diodes is approximately

±10nm. The temperature tuning coefficient is ~ 0.3nm/K. A thermistor is attached to the laser diode mount and the temperature is stabilised by a Peltier heat pump beneath the laser diode mount. The temperature controllers are all commercial. The PZT controllers are home-built voltage references (28V) with potential dividers for fine and coarse control. The maximum power of the bare diode at the free-running wavelength is 100mW for SDL5411 and 120mW for LD1377 and EYP-RWL-0850. The free-running wavelength is the wavelength at which the diode operates at room temperature without the external cavity. The output power of each diode after the extended cavity varies but it is on average ~12mW. The relatively low power after the extended cavity is because we use a lot of feedback to have a wider pulling range.

The blue diodes at 397nm are Nichia NLHV3000E and are rated for maximum 30mW power at their free-running wavelength. In the extended cavity we measure about 9mW and 7mW power at a current of 55mA. The measured linewidth from the beat frequency is 1.5MHz for lower power diodes [73]. The natural width of the Doppler cooling transition is 24MHz. As a result the linewidth is well below the natural linewidth as required for efficient Doppler cooling. The laser frequencies slowly drift in time leading to a drop in fluorescence rate. Tracking the fluorescence by manually compensating the drift gives the frequency drift of these lasers to be 300-400 MHz per hour [73]. That means that every few minutes the frequency of the lasers needs to be adjusted when not locked to a reference cavity.

5.8 Laser Preparation

The diode lasers need to be running in a single mode to ensure that the laser linewidth is the narrowest and the maximum power is delivered at a well defined frequency. To monitor the single mode of the IR lasers we pick off portions of these beams using beam splitters and send them to a spectrum analyser (see Appendix A). The spectrum analyser is a scanning Fabry-Perot interferometer with a free spectral range of 7.5GHz (Tec-Optics SA7.5). Because the lasers run at different frequencies they resonate at different cavity lengths and as a result it is possible to monitor all IR lasers on a scanning spectrum analyser although the peaks do not all occur in the same order of the cavity. When the laser is running in a single mode, sharp peaks appear at the resonance frequency of the cavity. Similarly, for the UV lasers we have a Fabry-Perot etalon with a large surface area. Sharp ring patterns appear on

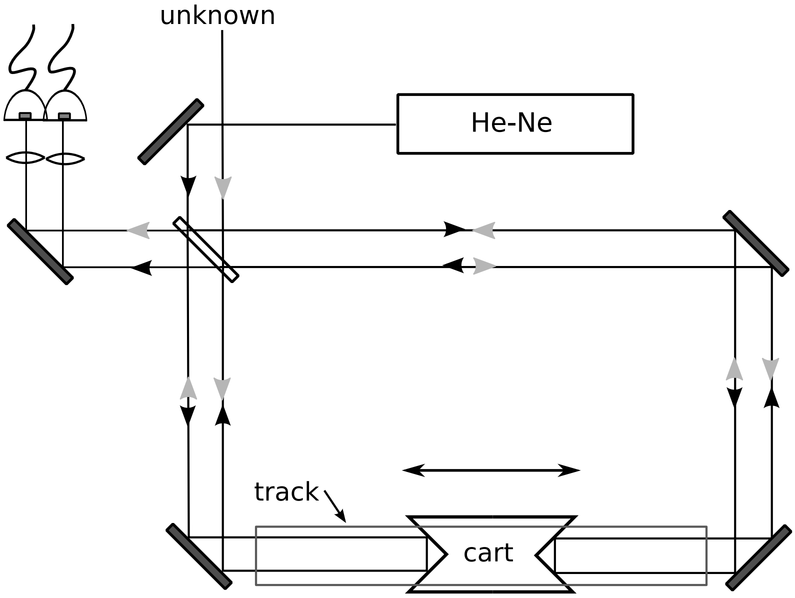

Figure 5.8. The optical layout of the wavemeter. The reference laser and the unknown laser overlap but counter propagate. The number of counted fringes is twice as if the cart had only one reflecting cube corner. The cart moves on an air-track with a solenoid and a spring at two ends. The solenoid gives a push to the cart inversely proportional to the speed of the cart in order to compensate for friction.

the far side of the etalon when the laser is running in single mode. Portions of UV lasers are sent to the etalon and are simultaneously monitored with CCD cameras. It is usually possible to force single mode operation by adjusting the current. If the laser is unstable and it mode hops, the grating needs to be optimised for feedback. This is usually a time consuming process and is caused by mechanical instabilities in the laser box such as in the nubs of the laser mount. In order to minimise this the temperature controllers are kept on at all times.

For day to day operation the lasers' wavelengths need to be approximately set to their corresponding atomic transitions wavelength with a precision of $<$ 0.001nm. The cooling transition natural linewidth is 24MHz. 24MHz in frequency corresponds to $\sim$0.00001nm at wavelength of the UV laser. However the transition linewidth is Doppler broadened for 'hot ions' leading to a linewidth in the GHz range. Therefore this precision is enough to see a small amount of fluorescence the first time we run the trap at the beginning of the day. After that the laser frequencies can be individually adjusted to maximise the fluorescence. We use a home-built Michelson interferometer wavemeter

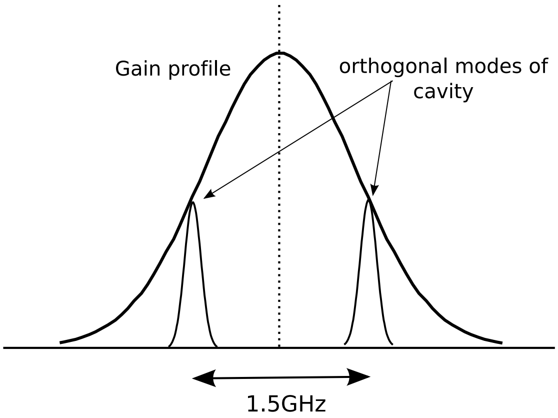

Figure 5.9. The two adjacent modes of the He-Ne laser are separated by the Free Spectral Range (FSR) of the external cavity. The FWHM of the gain profile of the He-Ne laser is ∼1.5GHz. The external cavity length is chosen such that only two modes can lase inside the cavity.

to find the wavelengths, based closely on the design by the Innsbruck group. The full details can be found in references [71, 78]. A schematic diagram is shown in figure 5.8. The reference laser beam is split into two and travels two separate paths. They join again at a beam splitter and are combined on a fast photodiode. The two light beams interfere and produce interference fringes as the difference in the two paths changes. The same process happens for another unknown laser whose wavelength is to be measured. If the path difference is the same for both lasers the unknown wavelength can be inferred from

$$\frac{\lambda_x}{\lambda_R} = \frac{N_x}{N_R},$$

where N_x is the unknown laser fringe count, N_R is the reference laser fringe count and λ_x and λ_R are the unknown and reference laser wavelengths respectively.

The path difference is provided by a moving cart on an air track. There are two corner reflectors on each side of the cart which offset the beam horizontally. The wavemeter is designed in such a way that the paths of the reference and unknown laser overlap. This ensures that the path difference of the two lasers remains the same and also it has the benefit of easy alignment of the unknown laser to the wavemeter.

The reference laser in our system is a home-built stabilised He-Ne, made by another PhD student of the group (D. Crick). The laser tube is inside

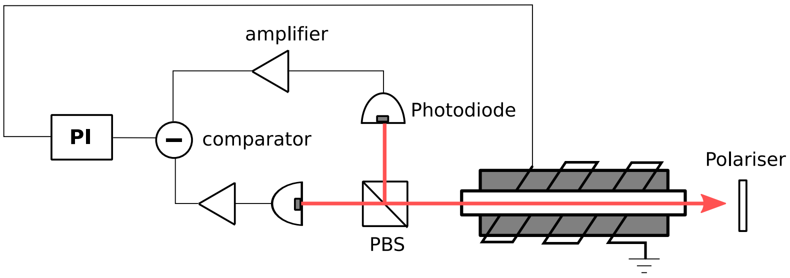

Figure 5.10. The schematic circuit diagram of the electronic feedback circuit for stabilising the He-Ne reference laser is shown. The intensities of the two orthogonal modes of the laser are compared and an error signal is achieved. The error signal is then fed into a typical proportional-integral (PI) gain circuit and fed back into the heater around the tube.

an external cavity with a length of 10cm. The light after the tube is randomly polarised, however the adjacent longitudinal modes have alternating orthogonal polarisations (see figure 5.9). The mode spacing is equal to the free spectral range of the external cavity. The cavity length is such that only two of these modes can exist within the gain profile. The frequency stabilisation circuit tries to keep these two modes symmetrical about the centre of the gain profile by changing the cavity length. The cavity length can be changed by adjusting the temperature of the laser tube with a heater wrapped around it. A schematic diagram of the circuit is shown in figure 5.10. The frequency stability is estimated to be better than 7MHz.

The electronics were provided by the ion trapping group in Innsbruck to read out the fringes with a precision of 1 part in 10^7. The electronic design is described in reference [78]. To increase the precision of the wavemeter either one needs to increase the length of the cart track to count more fringes or equivalently take into account the phase of the fringes in counting. The circuit starts counting when the phases of the two waves match within a 1/100th of a fringe and continues counting until the next time phase matching occurs. An integer number of fringes is counted for both lasers and the ratio N_x/N_R is determined to within a 1/100th of a fringe.

5.9 Laser Locking

5.9.1 Cavities

As mentioned in section 5.7.1 the lasers drift at the rate of 300-400MHz per hour. This is obviously not desirable and so the laser frequency should be

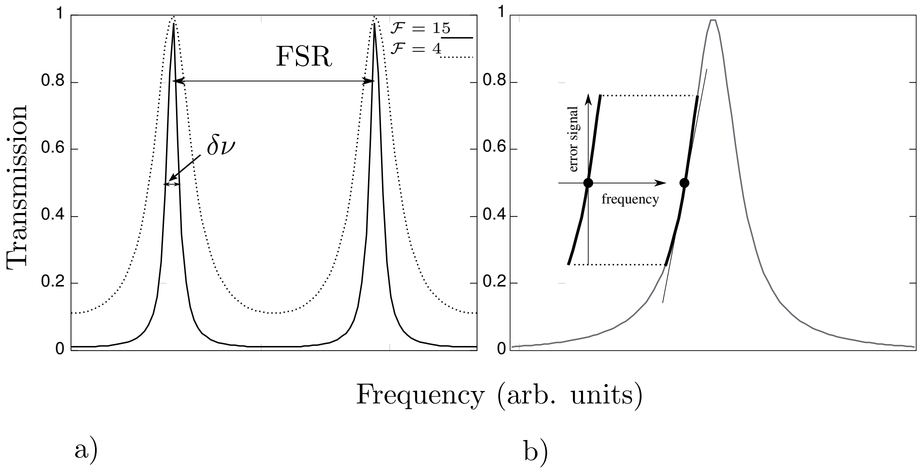

Figure 5.11. a) Fringes of a cavity. The peaks become narrower as the finesse becomes greater. b) Side-fringe lock is shown. The setpoint is halfway between the background and the top of the peak.

locked to a stable reference. The reference could be an atomic transition such as in saturated absorption locking techniques. This relies on counter propagating pump and probe beams to measure a single longitudinal velocity class of atoms within a Doppler-broadened profile. This technique is very robust in terms of drift but unfortunately there is no convenient vapour with transitions close to the 397nm or 866nm Ca^+ transitions. A hollow cathode lamp (HCL) in a magnetic field is another option. However that requires that a stabilised magnetic field be set very close to the magnetic field of the trap and also it may not be very precise. Another viable option is locking to a stable reference cavity and it is this option that we employ. The cavity consists of two reflecting mirrors which are separated by a distance L. The light incident on the cavity only resonates at frequencies for which the length of the cavity is equal to an integer multiple of half the wavelength or

$$f = \frac{nc}{2L}.$$

The separation between adjacent modes is the free spectral range ($\Delta\nu_{fsr}$) and is equal to $c/2L$. The full width at half maximum $\delta\nu$ of the mode depends on the reflectivity of the mirrors (R) and is given by

$$\delta\nu = \frac{\Delta\nu_{fsr}}{\mathcal{F}},$$

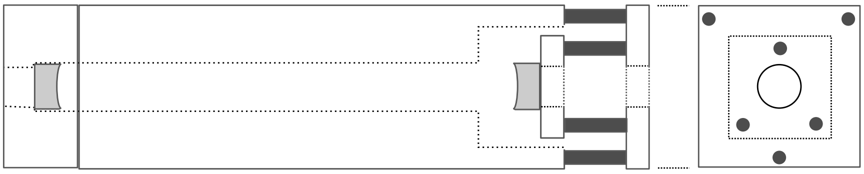

Figure 5.12. The cavities for locking the diode lasers. The finesse of these cavities is about 60 with the FSR of 1.5GHz. They are actively temperature stabilised to <10mK.

where $\mathcal{F} = \pi\sqrt{R}/(1-R)$ is the Finesse of the cavity. The more reflective the mirrors are, the sharper the cavity modes become (see figure 5.11(a)).

For each longitudinal mode there are also transverse electromagnetic modes (TEM) resonating in the cavity. These modes are the particular intensity patterns of the laser perpendicular to the direction of propagation. They are solutions of the paraxial wave equation and resonate at different frequencies. The fundamental TEM_{00} mode has a Gaussian intensity profile. This mode is usually chosen for locking.

The frequency stability of the cavity modes depends on how stable the cavity length is. Supposing the cavity length is stable, one mode is picked for locking. The most straight forward method is the side-of-fringe lock. In this technique the locking point is on one side of the fringe and halfway between the peak and the background level (see figure 5.11(b)). The intensity of the laser after the cavity is monitored by a fast photodiode. The signal on this photodiode is used as an error signal which can then be amplified using ordinary proportional-integration-derivative (PID) techniques and fed back to the laser PZT controller to compensate for the drift.

One downside of side-of-fringe locking technique is that any intensity fluctuations result in frequency fluctuations. The reason is that for a given frequency the locking point moves in frequency space when the intensity of the laser changes. A workaround for this problem is to normalise the error signal. This can be done by picking a beam off the laser by a beam splitter before going to the cavity and sending it to a second photodiode. The error signal is then divided by the second photodiode signal to normalise it. This does not completely solve the problem since the photodiodes' sensitivity and dark current can drift over time.

The reference cavities were built by E. Phillips [71]. A schematic is shown in figure 5.12. Since the aim was only to compensate for the drift and not to narrow the linewidth, a low finesse cavity was designed so as to have a large capture range. The capture range is the range in which the laser frequency can be offset from the cavity mode and still be locked. The mirrors were custom made (CVI Technical Optics) to give a finesse of about 60 for wavelengths of 866nm and 397nm. The radius of curvature is 30cm and the two mirrors are approximately 10cm apart. This gives a FSR of ∼1.5 GHz and a cavity linewidth of ∼ 25MHz. The spacers are made of Zerodur, a glass with a very low thermal expansion coefficient (0.02×10^{-6}/K at 0-50 °C).

Since the frequency of the atomic transition does not necessarily match that of the mode of the cavity, the cavity needs to be tunable. To tune the cavity length three Piezoelectric transducers (PZT) are glued by a thin layer of Torr Seal to the back of the mirror holder. Unfortunately the PZT crystal has a large thermal expansion coefficient compared to Zerodur. In order to reduce this three more non-tuning PZTs are attached to the cavity spacer which then hold the mirror mount. With this re-entrant design these latter three PZTs counteract the thermal expansion of the tunable PZT at the back of mirror mount. The PZT (P-830.10) is of the low voltage variety (0-100V) with a travel range of ∼ 15μm.

Even though the thermal expansion coefficient of the Zerodur is very small and the PZTs are in a counteracting arrangement, the mirror substrate and the epoxy between the PZT and mirror mount could expand or contract with temperature. Therefore it is very important to stabilise the temperature of the cavity. The cavity is covered with soft Neoprene and placed in an aluminium diecast box. A wire is wrapped around the box with a thermistor attached to the box. The whole box is then covered by a layer of hard Neoprene.

The current controller for the temperature is explained in reference [79]. The principle is very simple. The thermistor is in a Wheatstone bridge and the voltage difference across it is used as a means for monitoring the temperature changes. Temperature changes the resistance of the thermistor and that changes the voltage difference across the Wheatstone bridge. This voltage difference is amplified using two comparing operational amplifiers. Since the resistance of the Wheatstone bridge itself could change with temperature, for more stability the sensor electronics are placed inside the cavity box. The signal is then compared to an offset voltage to get the error signal. The error signal is amplified by a proportional-integral circuit and then goes to a power transistor that controls the current through the heating wire. Because the servo control can only heat the cavity the offset temperature needs to

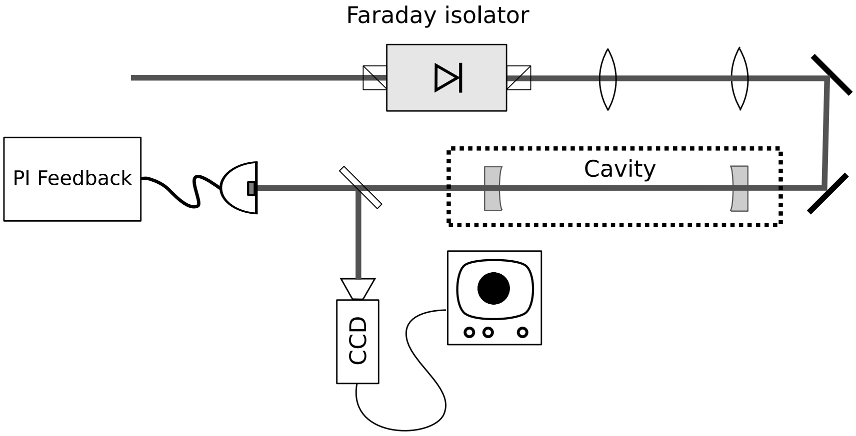

Figure 5.13. The optical layout of the reference cavity for locking the diode lasers. The Faraday isolator stops the back-reflections from the cavity mirrors from going into to the diode laser. Two lenses match the waist and the divergence of the laser beam to that of the cavity.

be higher than the room temperature. The temperature stability achieved is better than 10mK.

The optical layout of our cavity locking is shown in figure 5.13. Some light is picked off from the laser beam with a beam splitter and enters a Faraday isolator. This is to stop any back-reflection from the mirrors of the cavity from going into the laser diode. The waist of the beam is then adjusted using two lenses to match to the waist of the light resonating inside the cavity. Two mirrors allow adjustments to the direction and angle of the beam entering the cavity. Some light leaks out of the cavity and this is sent via a beam splitter to a CCD camera and a photodiode. The CCD camera helps in matching the light to the TEM_{00} mode.

The photodiode signal is fed into the servo control. The feedback circuits are based on reference [80] and were built by R. Hendricks. The full details can be found in reference [77]. The signal is compared with a setpoint to get the error signal. This signal is then amplified by a proportional-integral-derivative (PID) operational amplifier. This has three stages in parallel: one for P, one for I and one for D. This signal is then fed back to the PZT of the diode laser. The bandwidth of the circuit is very low (∼2kHz). This is because the PZT resonance frequency when loaded is very low (∼7kHz). Moreover, we only intend to compensate for the thermal drift of the laser not to narrow the linewidth. These home-built circuits are only used for the IR lasers. We use the commercial PID module of Toptica PL 100 to lock the UV lasers to the

UV reference cavities.

5.10 Trap Loading

The ions are generated by electron collision with neutral Ca atoms. The first ionisation energy of Ca is ~6.11eV. The electron beam needs to have energy higher than this. To achieve this the filament is negatively biased relative to ground. The higher the bias is, the greater the ionisation cross section becomes. The second ionisation energy of Ca is 11.87eV. The bias should be kept below this voltage to avoid double ionisation of the atoms because any ion species other than Ca^+ being trapped, slows the cooling process. We usually bias the filament to ~11V. The electron current density J (amp/cm^2) is described by the Richardson-Dushman equation

$$J = AT^2 e^{-\phi/k_B T},$$

where A is a constant depending on the metal, T is the temperature, k_B is the Boltzmann constant and ϕ is the work function of the metal. Temperature is proportional to the current squared. Therefore by changing the current through the filament one can control the electron current and therefore the number of the trapped ions. When we run the trap for the first time in the day a large cloud of ions ($\sim$ 10k counts/10ms) is loaded with a filament current of 4.7A.

Photoionisation is in principle a better alternative and has a number of advantages. With photoionisation it is easier to load single ions because it is a much more efficient process than electron bombardment and can be very precise in loading only a specific species (isotope) as it is a resonance-enhanced effect [81]. It also does not suffer from the charging up problems associated with electron bombardment. Insulating parts of the trap structure can charge up with electrons and result in ambient electric fields that need to be carefully compensated [82]. With photoionisation the background pressure inside the trap does not increase during the loading because it does not involve firing a filament and therefore heating the system. Finally because photoionisation is essentially a more efficient process, a very weak beam of atoms is enough to generate ions. Therefore the electrodes do not suffer from contamination produced by the oven.

Photoionisation is a two-step process requiring radiation at 423nm and

390nm. The 423nm step is resonant and requires a laser. Unfortunately laser diodes at this wavelength are no longer available. Therefore to produce this wavelength one needs to frequency double a Ti:Sapphire laser which would be expensive. Therefore, it has not been implemented.

The oven is very close to the centre of the trap so it is usually grounded in order to avoid the possibility of charging up by the action of the filament. The oven is heated by passing current through it and a beam of Ca atoms comes out of a small hole in it. The oven current is usually fixed at 1.5-1.6A. Care must be taken not to run the current at more than 2A otherwise it coats the trap electrodes and decreases the lifetime of the oven.

To load the trap the oven is first turned on for 1-2 minutes. The filament is turned on at the end of this period for 20-25 seconds (while the oven is still on). For large clouds, fluorescence is usually seen a few seconds before turning off the filament. For small clouds ($<$ 1000 counts/10ms) the ions take some time to cool and therefore the fluorescence takes some time to rise and may not be visible until many seconds after the filament and oven have been switched off.

5.11 General Procedure

The general procedure for laser cooling is as follows. The UV lasers are turned on and left on for an hour to settle down. The IR lasers are usually left on over night after setting them close to the right mode by adjusting the current and PZT controller. The He-Ne laser is turned on and left for half an hour to stabilise. After the He-Ne reference laser is locked the lasers are set to the correct wavelengths[*] according to table 5.1.

One of the UV lasers is set to scan the frequency range of ∼2GHz below the transition frequency with a scanning rate of ∼1Hz. This is done by putting a sawtooth voltage on the PZT controller of the UV diode laser. It is important to check that the laser operates on a single mode throughout the scan range. If this is not the case adjustment of the laser current is needed.

We run the magnet with a current of 15A which corresponds to a magnetic field of 0.98 tesla at the centre of the trap. The magnetic field can be calculated precisely by measuring the motional frequencies of ions. The resistance of the

---

[*] These wavelengths are what we have found to be optimal for laser cooling relatively small cloud of ions. They are not absolutely calibrated —they are simply the numbers that the wavemeter shows.

Laser	Wavelength [nm]
Red #1	866.2379
Red #2	866.1837
Red #3	866.2425
Red #4	866.1885
Red #5	854.209
Blue #1	396.8417
Blue #2	396.8514

Table 5.1. Optimised wavelengths for a small cloud

coils goes up as the temperature of the magnet coils rises. Therefore the magnet needs to be left on for at least half an hour in order to be stabilised.

As seen in chapter 4 the position of the laser inside the trap is a crucial factor in the cooling process. The radial position of the laser should be offset to the side where the ions move along the direction of the beam. As the ions get colder and their magnetron orbit gets smaller the beam can be pulled closer to the centre of the trap. As the cloud gets smaller the position of the laser becomes more critical. There is a motorised mirror before the trap for this purpose which can be controlled by a LabVIEW program.

The whole trap and the PMT are covered with thick black cloth to shield it against stray background light. The MCS program (see section 5.6) is set to a bin time of 10-20ms and the PMT is turned on. Initially the fluorescence rate per ion may be quite small due to the relatively low precision with which we are able to set the wavelengths of the lasers. Therefore the trap needs to be loaded with a relatively large cloud to compensate for this. The trap is loaded with 1.5-1.6A oven current and 4.5-4.7A filament current.

At this stage a small fluorescence signal (∼100 counts/10ms) should be visible which is usually stable if the beam position is not on the heating side. There are then many parameters to adjust in order to improve the fluorescence level. The first is the frequency of the lasers. This is done in an iterative procedure. When the ions are ‘hot’ the signal level is more sensitive to changes in the tuning of the UV lasers. The spatial overlap of the lasers at the centre of the trap is another important factor. The UV and IR beam alignment can be further adjusted to acquire better signal. The focusing into the trap should be checked as well as the PMT position. Finally the polarisation of the IR and UV lasers need to be optimised.

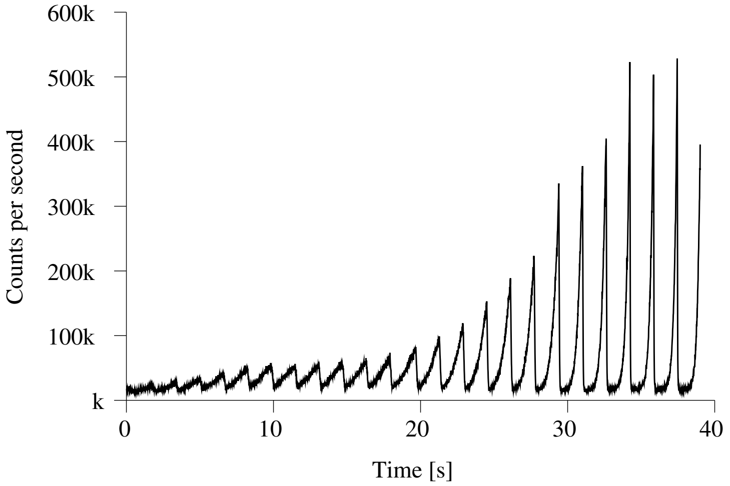

Figure 5.14. The fluorescence signal when all parameters of laser cooling are optimised. The peaks become sharper as the ions get colder.

This process then is repeated by loading a smaller cloud. To get a smaller cloud usually the filament current is lowered while keeping other loading parameters constant. If everything is optimised, loading the trap again should give rise to a series of peaks as shown in figure 5.14 which get sharper as the ions get colder. Although the laser parameters are optimised it takes several frequency scans of the laser to laser cool the cloud and achieve a steady signal. This is mainly due to laser beam offset and can be eased by temporarily increasing the offset during the loading and moving the beam towards the optimum position whilst the cloud is cooling to the centre of the trap. The lasers and the beam positions are further adjusted for the smaller cloud to get narrower peaks. The peaks are half the Doppler broadened profile whose width depends on the temperature of the ions. At this stage the scanning of the UV laser is switched off and the UV lasers can be locked to the reference cavities. Their frequencies are then adjusted using a home-made fine tuning PZT controller for the cavities.

We could also run the split-ring trap as an RF trap by applying a RF potential to the ring electrode and a DC bias to the endcaps. The RF potential is generated by passing the output of a signal generator (HP 3325B) through an amplifier (ENI 325LA) and then to a home-made transformer. We usually apply a RF potential with a frequency of $\sim$ 2MHz and an amplitude of $\sim$ 200V p-p to the ring electrode and a DC bias of 4V to the endcaps.

Before discussing the experimental improvements implemented since I joined the group and the main results, which are presented in chapter 8, a number of other aspects of the project to which I have contributed are discussed in the following two chapters: the frequency stabilisation of the Ti:Sapphire laser and the coherent population trapping modelling.

CHAPTER 6

Ti:Sapphire Laser Stabilisation

6.1 Introduction

This chapter explains the work which I carried out to frequency stabilise a Ti:Sapphire laser by locking it to a high finesse Fabry-Perot reference cavity. The technique used was originally developed by Pound for the stabilisation of microwaves [83]. It was later adapted for the optical regime by Drever and Hall [84]. It is usually referred to as the Pound-Drever-Hall (PDH) technique.

A previous attempt by a previous PhD student in our group (J. Sudbery) in locking the laser was unsuccessful due to an unknown noise in the laser frequency [79]. Several major improvements to the setup, specifically, building a helical resonator (see section 6.3.3) and a loop filter (see section 6.3.6) as a feedback control circuit led to a successful lock.

The laser was developed for use in two experiments planned at the Ion Trapping group in Imperial College; firstly the laser will be used to coherently drive the $S_{1/2} - D_{5/2}$ transition and observe Rabi oscillations. In the next stage it will be used in resolved side-band cooling of a single $^{40}Ca^+$ ion in a Penning trap. The internal atomic transition used in cooling is the $S_{1/2} - D_{5/2}$ transition at 729nm.

The PDH technique is a widely used and powerful modern laser frequency stabilisation method. As in the side-of-fringe technique, reference laser light is mode matched to a stable reference cavity (see section 5.9.1). However unlike many other techniques, the *reflected light* from the cavity is used to derive an error signal. The error signal derived is expressed as the derivative of a peak of the Airy function. The error signal is then fed back to the laser via a servo control system in order to suppress the frequency fluctuations.

One advantage of this technique over other techniques (such as side-of-fringe locking) is that the laser intensity fluctuations are decoupled from frequency fluctuations. In side-of-fringe locking, when the intensity changes the locking point shifts and therefore intensity fluctuations are transformed into frequency fluctuations (see figure 6.1). This problem can be reduced by normalising the intensity of the laser. Normalising is achieved by splitting the beam into two and using one beam to provide an intensity reference while the other is used to lock to the cavity from which an error signal is derived. The intensity of the two beams then can be divided to normalise the signal. However it requires that the photodiodes have similar responses with frequency changes and moreover they do not drift over time.

In contrast locking can be achieved in the PDH technique using only one photodiode. Another advantage of the PDH technique is that it is inherently faster because it is not limited by the response time of the cavity. This is because in the PDH technique the reflected light from the cavity is used for acquiring the error signal rather than the light after the cavity as it is in the side-of-fringe technique. The response time becomes a major limit to the side-of-fringe technique particularly for high finesse cavities. The cavity decay time for a high finesse cavity of length L and finesse of $\mathcal{F}$ is given by

$$\tau_c = \mathcal{F}\frac{L}{2\pi c}.$$

The cavity averages any jitter with frequency components higher than $\sim 1/\tau_c$. Therefore the frequency bandwidth of the lock cannot be substantially greater than the cavity linewidth. On the other hand the discriminator slope is lower if the cavity linewidth becomes bigger. The PDH technique is a good alternative to overcome this dilemma. Another benefit is a relatively large capture range

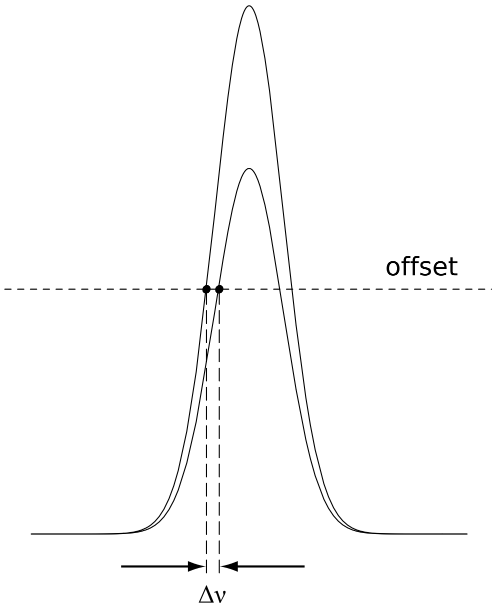

Figure 6.1. Intensity fluctuations shift the locking point if the Airy function itself is used.

as we will see.

6.2 Theory

The principle is simple and is shown in figure 6.2. Instead of locking the laser to a side of the Airy function, it is locked to the middle point of the derivative. The locking point corresponds to the maximum of the intensity peak. Intensity fluctuations only cause changes in the *slope* of the error signal and are not therefore transferred to the frequency domain.

It is instructive to first give an intuitive understanding for the PDH locking scheme. We assume the reference cavity is a symmetric Fabry-Perot cavity. The phase of the incident light is modulated using an electro-optic modulator (EOM). This gives two sidebands spaced from the carrier by the modulation frequency. Some fraction of the incident light (sidebands) is not resonant with the cavity and is directly reflected from the first mirror. A fraction of the light resonating in the cavity (carrier) leaks back out of the input mirror and beats with the sidebands. Signal is the sum of the beat of the upper and lower sidebands with the carrier. The electric fields of the two sidebands are 180° out of phase with respect to each other (see figure 6.3). The result is that when the laser is in resonance with the cavity these two beats are 180° out of

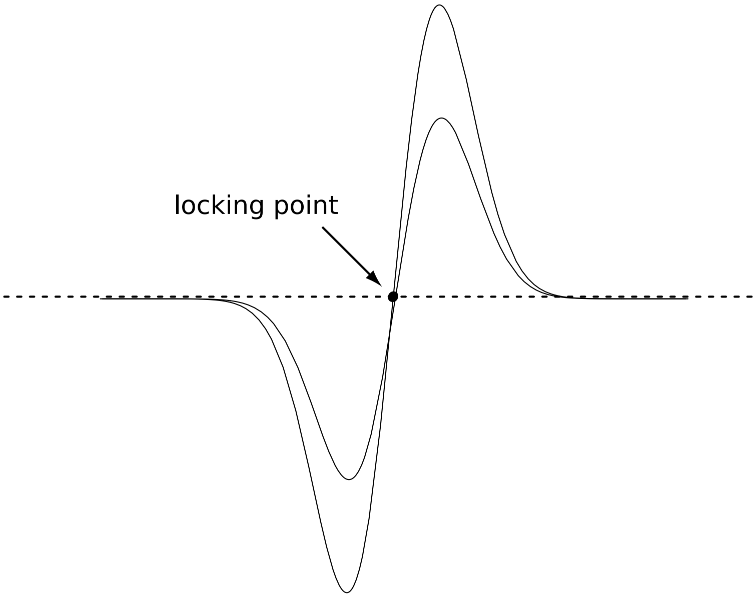

Figure 6.2. The derivative of the Airy function. The intensity fluctuations do not shift the locking point in frequency provided that the derivative of the Airy function is used.

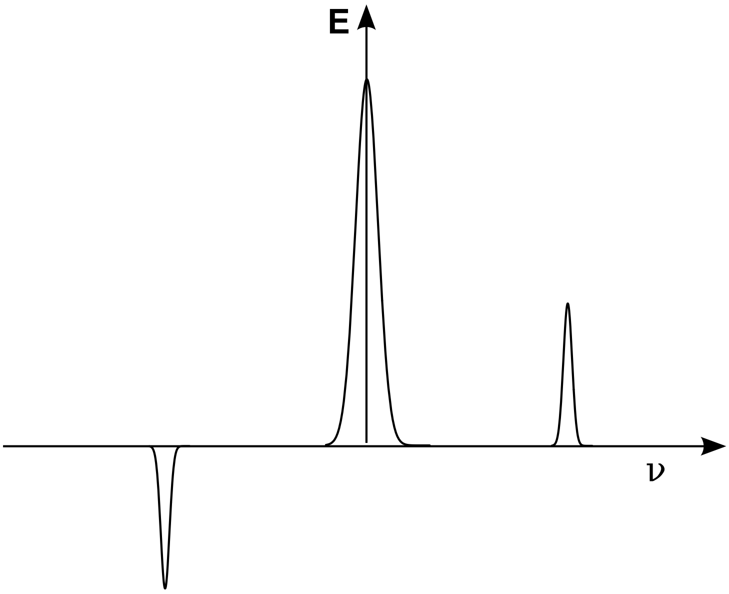

Figure 6.3. The reflected electric field from the reference cavity with imposed sidebands. The sidebands are anti-phase.

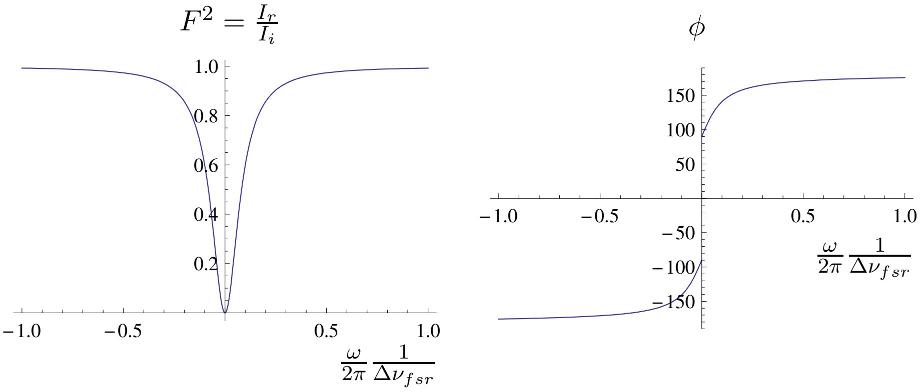

Figure 6.4. *Left:* The reflection coefficient vs. frequency for a cavity with a finesse of ~77. *Right:* The phase of the reflected beam vs. frequency.

phase and destructively interfere. But when the laser is off resonant, the two beats do not completely cancel out. Measuring the beat amplitude gives an error signal depending on which side of the fringe the laser frequency is.

Suppose that the electric field of the incident beam at a point before a Fabry-Perot cavity is $E_i = E_0 e^{\mathrm{i}\omega t}$ and the electric field of the reflected beam at the same point is $E_r = E_1 e^{\mathrm{i}\omega t}$. One can show that for a symmetric cavity with no losses, the reflection coefficient $F(\omega)$ is equal to

$$F(\omega) = \frac{E_r}{E_i} = r\frac{\exp(\mathrm{i}\omega/\Delta\nu_{fsr}) - 1}{1 - r^2\exp(\mathrm{i}\omega/\Delta\nu_{fsr})},$$

where r is the reflectivity of the mirrors and $\Delta\nu_{fsr} = c/2L$ is the free spectral range of a cavity with length L [85]. The intensity and phase of the reflection coefficient is shown in figure 6.4. By measuring the phase it is possible to find out whether the frequency is below or above the resonance frequency. The problem is that the phase cannot be directly measured. Instead the phase is indirectly measured by dithering the laser frequency. Modulating the frequency makes sidebands which have different phases from the carrier. The sidebands get reflected from the cavity and interfere with the carrier which results in a beat pattern at the modulation frequency. The phase of the beat, which now *can* be measured, gives information about the phase of the reflected beam and therefore enables locking to the maximum of the peak.

Analysis in this section can also be found in [86]. The simplified setup is depicted in figure 6.5. The laser beam first passes through an electro-optic modulator (EOM). The EOM, which is essentially a Pockels Cell, rapidly modulates the phase of the electric field at the frequency of a local oscillator.

To be more precise if the incident electric field is $E_i = E_0 e^{i\omega t}$, the EOM shifts the phase giving

$$E_i' = E_0 e^{i(\omega t + \beta \sin \Omega t)}.$$

Here Ω is the phase modulation frequency and β is the modulation depth which is usually smaller than 1. Under these conditions the equation above can be approximated by the first order Bessel functions as follows:

$$E_i' = E_0[J_0(\beta)e^{i\omega t} + J_1(\beta)e^{i(\omega+\Omega)t)} - J_1(\beta)e^{i(\omega-\Omega)t}]. \qquad (6.1)$$

Here we used the Jacobi-Ager expansion equation

$$e^{\mathrm{i}z\cos\theta} = \sum_{n=-\infty}^{n=\infty} \mathrm{i}^n J_n z e^{in\theta}.$$

This is equal to a carrier at the frequency ω and two sidebands at frequencies $\omega - \Omega$ and $\omega + \Omega$. In other words modulating the phase of the incident electric field gives rise to two sidebands which are separated in frequency from the carrier by the modulation frequency of the EOM. The sideband amplitudes are proportional to the modulation depth when $\beta < 1$. Ignoring the interference terms the power of the carrier beam is

$$P_c = J_0(\beta)^2 P_0, \qquad (6.2)$$

where $P_0 = |E_0|^2$ is the power of the incident beam. Similarly for the sidebands we get

$$P_s = J_1(\beta)^2 P_0. \qquad (6.3)$$

When the modulation depth β is small most of the power is distributed between the carrier and the first sidebands ie.

$$P_0 \simeq P_c + 2P_s.$$

For incident light whose electric field is given by

$$E_i = E_0 e^{\mathrm{i}\omega t},$$

the reflected field is given by

$$E_r = F(\omega) E_0 e^{\mathrm{i}\omega t}.$$

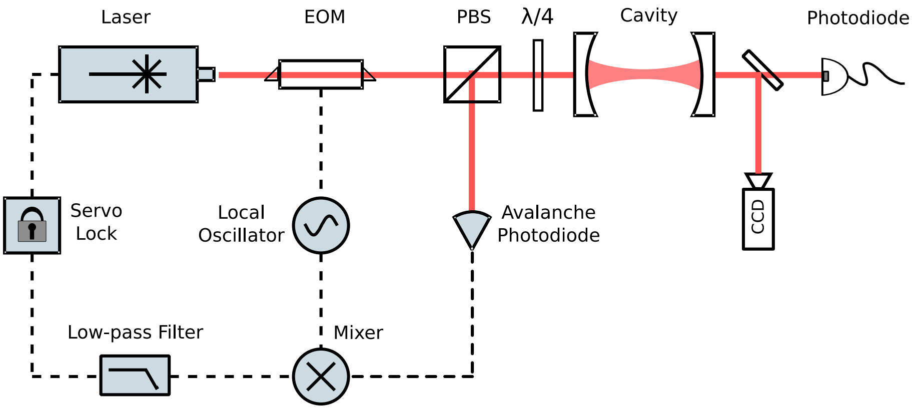

Figure 6.5. The Basic setup for the PDH laser locking. Solid lines are optical paths and dashed lines are signal paths.

Expressing all the components in equation 6.1 in this way, the electric field of the reflected beam is

$$\begin{aligned} E_r' =& E_0\{F(\omega)J_0(\beta)e^{\mathrm{i}\omega t} + F(\omega+\Omega)J_1(\beta)e^{i(\omega+\Omega)t} \\ &- F(\omega-\Omega)J_1(\beta)e^{\mathrm{i}(\omega-\Omega)t)}\}. \end{aligned} \tag{6.4}$$

What is measured by the photodiode is the reflected beam power which is proportional to $|E^2| = EE^*$. Algebraic manipulation of equation 6.4 gives the reflected beam power

$$\begin{aligned} P_r =& 2\sqrt{P_c P_s}\{\mathrm{Re}\{F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\}\cos\Omega t \\ &+ \mathrm{Im}\{F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\}\sin\Omega t\} \\ &+ (\text{constant terms}) + (2\Omega\text{ terms}). \end{aligned} \tag{6.5}$$

The Ω terms come from the interference of the sidebands with the carrier and the 2Ω terms come from the interference of the sidebands with each other. The phase information is carried only in the Ω terms and the 2Ω terms should be discarded. A low-pass filter with a cutoff frequency below 2Ω is well suited for this purpose. In order to determine the reflected power we need to determine the value of the term

$$\mathcal{E}(\omega) = F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega).$$

There are two regimes in which this term can be easily expressed: slow modulation and fast modulation. We will consider both cases and compare them with each other.

6.2.1 Slow Modulation

When the frequency modulation Ω is slow enough for the internal field of the cavity to respond ($\Omega \ll \Delta\nu_{fsr}/\mathcal{F}$), the reflection coefficient for the sidebands can be approximated by

$$F(\omega+\Omega) = \frac{dF(\omega)}{d\omega}\Omega.$$

Therefore $\mathcal{E}(\omega)$ can be approximated by

$$\begin{aligned}\mathcal{E}(\omega) =& F(\omega)F^*(\omega+\Omega) - F^*(\omega)F(\omega-\Omega)\\ =& F(\omega)\frac{dF^*(\omega)}{d\omega}\Omega + \text{cc.}\\ =& 2\text{Re}\{F(\omega)\frac{dF^*(\omega)}{d\omega}\Omega\}\\ =& \frac{d|F(\omega)|^2}{d\omega}\Omega\end{aligned}$$

which is real. As a result only the cosine term in equation 6.5 survives and the intensity of the reflected light can be rewritten as

$$P_r = 2\sqrt{P_sP_c}\frac{d|F(\omega)|^2}{d\omega}\Omega\cos(\Omega t) + (\text{constant terms}) + (2\Omega\text{ terms}).$$

Constant terms which give a DC signal are easily discarded by AC coupling the photodiode. High frequency terms are discarded using a low-pass filter with a cutoff frequency at 2Ω. Now if this signal is mixed with a $\cos(\Omega t)$ signal using a mixer we get

$$\begin{aligned}\mathcal{E} =& 2\sqrt{P_sP_c}\frac{d|F(\omega)|^2}{d\omega}\Omega\cos^2(\Omega t)\\ =& \sqrt{P_sP_c}\frac{d|F(\omega)|^2}{d\omega}\Omega + (2\Omega\text{ term}).\end{aligned}$$

$\cos(\Omega t)$ is simply the modulation signal with a phase difference of 90°. The error signal is shown in figure 6.6a.

6.2.2 Fast Modulation

When the modulation frequency Ω is high ($\Omega \gg \Delta\nu_{fsr}/\mathcal{F}$) and the laser frequency is near the resonance of the cavity we can assume that both sidebands

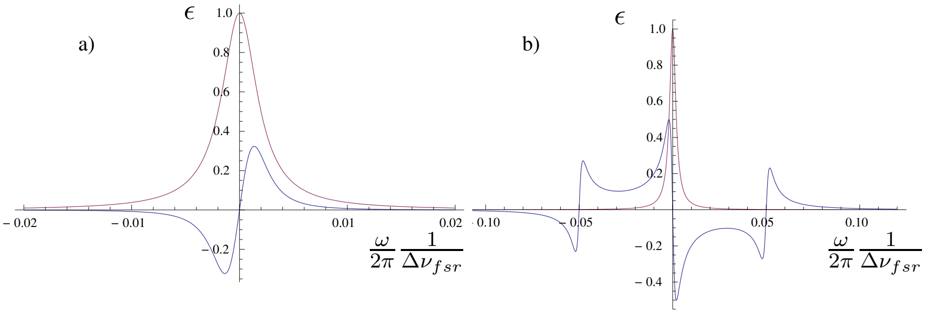

Figure 6.6. **a)** Error signal with slow modulation. The modulation frequency is 0.1% of the free spectral range (about half the cavity linewidth). **b)** Error signal with fast modulation. The modulation frequency is 5% of the free spectral range. The cavity has a finesse of 500. Red lines show the Airy fringe of the cavity. Note that when the modulation is fast the linear slope of the error signal corresponds to the linewidth of the cavity fringe. Also the signal has opposite signs for the two cases of fast and slow modulations.

are totally reflected from the cavity i.e.

$$\begin{aligned} F(\omega \pm \Omega) &= 1, \\ \mathcal{E}(\omega) &= F(\omega) - F^*(\omega). \end{aligned}$$

In this approximation the cosine part in equation 6.5 is negligible and only the sine part survives

$$P_r = -4\sqrt{P_c P_s}\,\mathrm{Im}\{F(\omega)\} \sin \Omega t + (2\Omega \text{ terms}).$$

This signal is mixed with the local oscillator signal whose phase is matched to the error signal to get

$$\begin{aligned} \epsilon &= -4\sqrt{P_c P_s}\,\mathrm{Im}\{F(\omega)\} \sin^2 \Omega t + (2\Omega \text{ terms}) \\ &= -2\sqrt{P_c P_s}\,\mathrm{Im}\{F(\omega)\} + (2\Omega \text{ terms}). \end{aligned}$$

The error signal is shown in figure 6.6. Phase matching is a key part in getting a good error signal. Unmatched phases lead to a diminished and peculiarly shaped signal (see figure 6.7). In practice one can simply change the lengths of the signal cables to shift the phase.

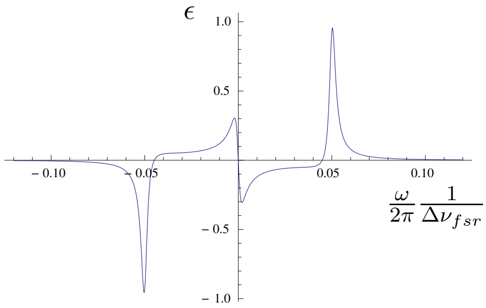

Figure 6.7. Error signal for 80° phase shift. Note how the carrier signal is diminished compared with figure 6.6b

When the laser frequency is near resonance we have

$$\frac{\omega}{\Delta\nu_{fsr}} = 2\pi N + \frac{\delta\omega}{\Delta\nu_{fsr}},$$

where N is a big number and $\delta\omega$ is small compared to the free spectral range. Using this the error signal can be approximated to the first order of $\delta\omega$ as

$$\epsilon = -2\sqrt{P_c P_s}\frac{r}{1-r^2}\frac{\delta\omega}{\Delta\nu_{fsr}}.$$

For a high finesse cavity $r \simeq 1$ and $\mathcal{F} \simeq \pi/(1-r)$. Therefore the error signal can be rewritten as

$$\epsilon \simeq \frac{\sqrt{P_c P_s}}{\pi}\frac{\delta\omega}{\gamma},$$

where $\gamma = \Delta\nu_{fsr}/\mathcal{F}$ is the linewidth of the cavity. The error signal is linearly proportional to the frequency offset when the laser is near resonance. The linear frequency response of the error signal means that the frequency noise can be suppressed using normal feedback control techniques. However as we will see in section 6.3.6, in order to fully take advantage of the PDH technique one needs to use a “loop filter” for the feedback.

The error signal is proportional to the square root of the carrier power times the sideband power. When the modulation depth is small the power is distributed between the first order sidebands and the carrier, that is $P_r = P_c + 2P_s$. Maximum slope for the error signal is achieved when the sideband power is approximately half that of the carrier. To be more precise if the term $\sqrt{P_c P_s}$ is expressed in Bessel functions the maximum occurs at modulation depth $\beta = 1.08$. For this modulation depth the ratio of the sidebands to the carrier is 0.41. This is shown in figure 6.8. It can be seen that the graph is

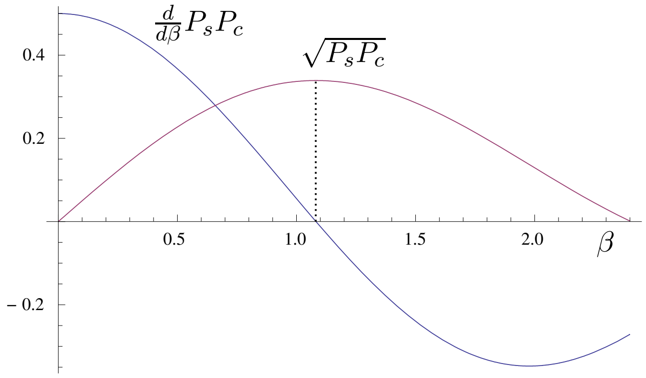

Figure 6.8. $\sqrt{P_cP_s}$ as a function of modulation depth β. Maximum occurs at $\beta = 1.08$. At this modulation depth the ratio of the sidebands to the carrier is ~41%.

not a sharp function of the modulation depth at $\beta \simeq 1$. Therefore the relative power between the sidebands and the carrier is not a very critical parameter.

6.3 Experimental Setup

The outline of the optical setup is shown in figure 6.9. The output of the Ti:Sapphire laser firstly passes through a pair of Faraday isolators. By doing this any back reflection from optical elements is stopped from going back into the laser. The light is then focused onto the first acousto-optic modulator (AOM) used in a double-pass configuration. This AOM ultimately compensates for frequency fluctuations of the laser using the error signal that is generated from the reference cavity. The light then passes through an electro-optical modulator (EOM). The EOM modulates the phase of the light as explained above. Modulating the phase applies mainly two sidebands which are separated from the carrier by the EOM modulation frequency. The light passes through a polarising beam splitter and then a quarter waveplate. Two lenses are then used to match the waist of the beam to the mode of the cavity. The reflected light from the cavity passes again through a quarter waveplate which makes the light now reflect from the polarising beam splitter. Finally the reflected light is focused to an avalanche photodiode (APD).

6.3.1 Ti:Sapphire Laser

The laser is a Coherent MBR-E 110 Ti:Sapphire with maximum power at 864nm. The Ti:Sapphire crystal is pumped by a 5W Coherent Verdi at 532nm.

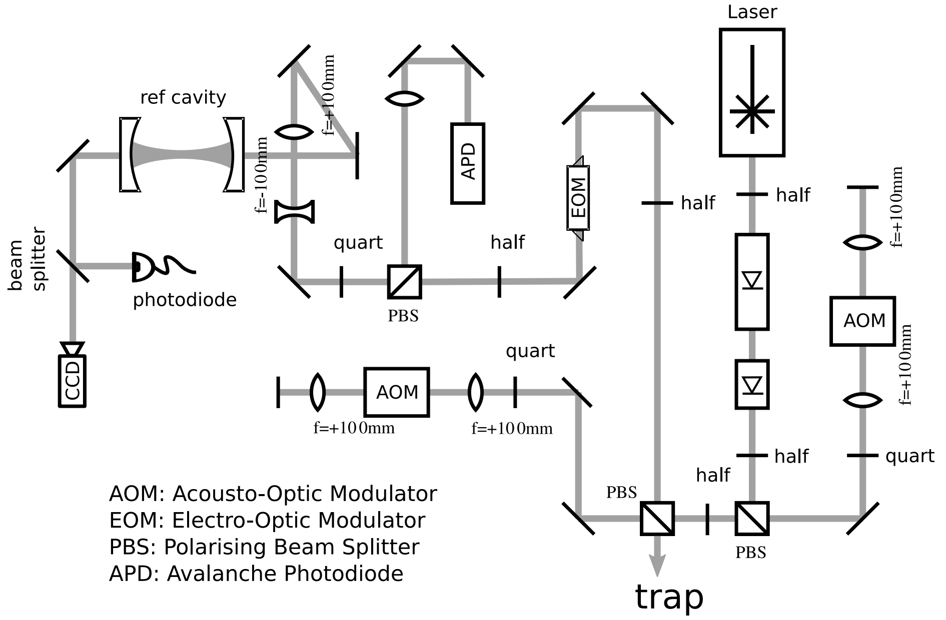

Figure 6.9. Optical setup of the PDH locking system for the Ti:Sapphire laser.

The optical schematic of the laser is shown in figure 6.10. The laser has a monolithic ring cavity resonator with a compact optical diode to keep the laser unidirectional. The gain profile of the Ti:Sapphire crystal is very broad (700nm-1100nm). The laser mirrors narrow this gain profile but not enough to keep the laser running in single mode. The pump beam focus to the crystal ensures that the laser runs in a single transverse mode (TEM_{00}). However it can still run in many different longitudinal modes. Therefore a number of other elements need to be used to run the laser in single mode.

A three-plate birefringent filter is used to tune the wavelength of the laser. Light passing through the filter gets split into ordinary and extraordinary rays which travel with different velocities. The polarisation of the light rotates by

$$\Delta\phi = 2\pi\frac{l}{\lambda}\Delta n,$$

where l is the thickness of the plate and Δn is the difference of the ordinary and extraordinary indices of the medium. The plates are placed at Brewster's angle so that the loss is high for any light whose wavelength is not an exact integer multiple of the effective optical path difference. Three of these plates with different thicknesses are placed to make the gain profile narrower and the losses even higher. Losses occur on Brewster angled components elsewhere in

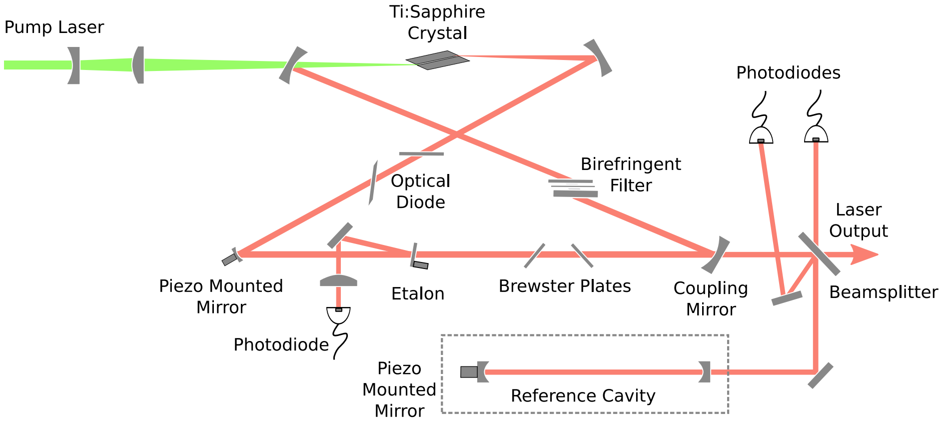

Figure 6.10. The schematic diagram of the Ti:Sapphire laser.

the cavity for linear polarisation components orthogonal to the plane of the Brewster surfaces. By changing the angle of the filter the operating wavelength of the laser can be appropriately adjusted.

A thin Piezo-mounted etalon along with a servo lock eliminates mode hopping and keeps the laser running in single mode. The etalon is a tilted thin plate with its own transmission peaks which depend on its angle. A schematic diagram of the transmission profile for different elements in the laser is shown in figure 6.11. Multiplying these transmission curves gives the total transmission curve versus frequency as shown in figure 6.12. The etalon mode spacing is designed in such a way that only one of the modes can be above the threshold and acquire enough gain to lase. In order to lock the laser to a mode of this etalon, the intensity of the reflected light from the etalon is monitored by a photodiode. The etalon is dithered at a frequency of 88kHz and an error signal is acquired*. The feedback signal is then fed back into the etalon. As part of the setup procedure for the etalon lock, the etalon is slowly tilted by manually adjusting the voltage to the Piezo. The error signal varies during this process as is shown in figure 6.13. The error signal changes smoothly for a period and then jumps to a new value whenever a mode hop occurs. The voltage offset for the lock is set to the middle of this sawtooth error signal and the etalon is then locked to the middle point.

In laser resonators with linear cavities the counter-propagating light can interfere with the propagating light and form a standing wave interference

---

*This can be thought of as the low-frequency limit of the PDH scheme discussed before. This is often called a 'dither' lock.

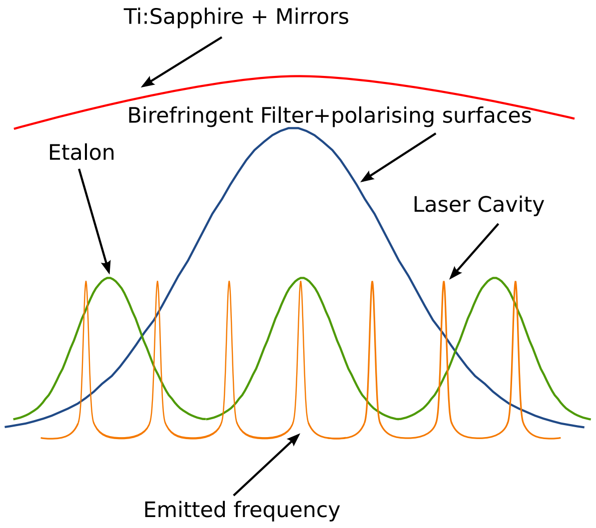

Figure 6.11. Schematic of frequency profile of the various elements in the Ti:Sapphire laser is shown.

pattern whose period is half the wavelength of the laser light. The result is that the gain is saturated at the antinodes and the lasing mode experiences stronger gain saturation than the other competing non-lasing modes (for example a mode which is $\lambda/4$ longer). This effect, which is called “spatial hole burning” [85, 87], causes mode hopping. To make the laser unidirectional the resonator is made in a ring form with an element put in one arm which makes losses higher for one direction than the other. This element could be an optical diode such as a Faraday rotator and a waveplate which is the case in our system. The basic principle is shown in figure 6.14. The diode changes the polarisation of the light only for the backward going light which then has high losses due to the various Brewster’s angle elements inside the laser’s ring cavity.

Inside the laser enclosure there is a temperature stabilised Fabry-Perot interferometer with a finesse of >25 which acts as a reference cavity. The cavity is sealed to eliminate air pressure changes. To compensate for intensity fluctuations, two photodiodes are used. Part of the beam is picked off before going to the cavity and is used to normalise the fringes of the cavity. The laser is locked to the middle point of a side of a fringe. One mirror of the ring cavity (the “tweeter”) is Piezo-mounted and is connected to the servo controller of the commercial electronics box provided for feedback.

In order to scan the laser frequency by up to 40GHz, two tilting Brewster

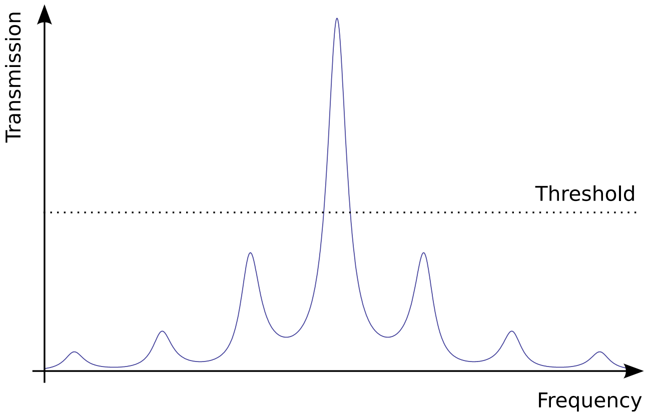

Figure 6.12. The etalon is designed such that only one mode is above the threshold thus keeping the laser running in single mode.

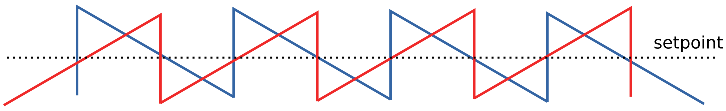

Figure 6.13. The etalon error signal when it is turned clockwise or counter-clockwise is shown. The signal suddenly drops when a mode hop occurs. The etalon is locked to the middle of the sawtooth waveform.

plates are used to change the effective length of the cavity. A single plate would suffice in principle but using two plates at opposing angles cancels out beam position shifts. The two Brewster plates are synchronised with the etalon and allow slow scanning the laser whilst the laser itself is locked to the reference cavity. The Brewster plates are also responsible for correcting the slow drift of the laser. When the laser is running, the servo and the etalon both need to be locked.

The optical table is enclosed and a laminar flow unit blows clean air down onto the laser in order to protect the laser optics from dust. The laser mirrors need to be cleaned and realigned occasionally to get the maximum power. This gives an output laser power up to ~300mW at 729nm when pumped by a 5W Verdi.

The short-term linewidth of a laser can be estimated by measuring the beat pattern of two similar lasers. The temporal interference pattern of two similar lasers gives information about the frequency difference of the two lasers. The FFT spectrum of the beat signal gives the short-term linewidth of the laser. The company that manufactures the laser, Coherent UK, quotes a linewidth

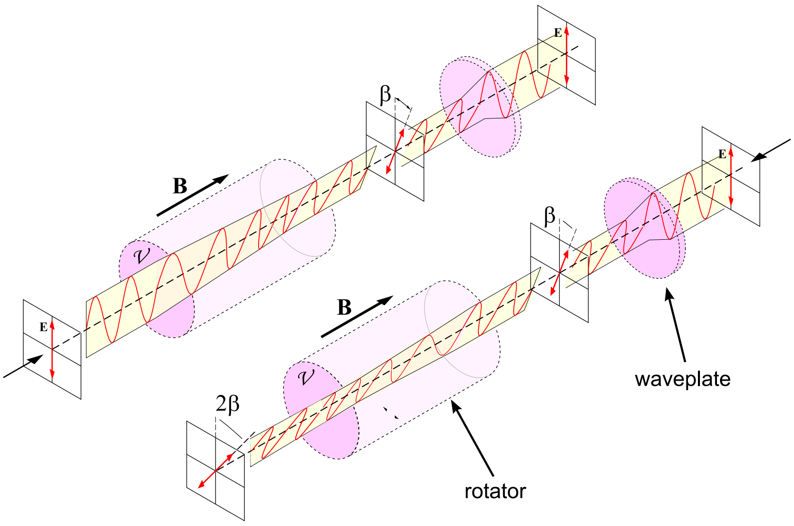

Figure 6.14. The Faraday rotator and the waveplate leave the polarisation of the forward light unrotated and rotate the polarisation for the backward light. The magnetic field is produced by permanent magnets around the rotator.

of 100kHz when the etalon and the servo (reference cavity) are locked. They estimate this linewidth from the error signal of the laser's reference cavity not from a beat signal with another laser. In reality we found that the linewidth is an order of magnitude worse than that quoted by Coherent which we estimated by simply measuring the short term fluctuations of a fringe of our own high finesse cavity.

6.3.2 AOM and EOM

The drivers for the AOMs are Isle Optics (SD-100T and MD250). The first AOM is a Brimrose (TEF-80) modulated at 80MHz with a tuning range of 20MHz. The second AOM is an Isle Optics (LM250) with a tuning range of 350MHz. Since the high finesse cavity is not tunable the second AOM is intended to shift the frequency of the light to the $S_{1/2} \rightarrow D_{5/2}$ transition frequency. The two AOMs are in a double-pass configuration so that the frequency shift does not affect the beam position. The AOM efficiency[*] in a double pass configuration is $\sim 35\%$ for both AOMs.

The free spectral range of the high finesse cavity is ∼1.5GHz (see section 6.3.4). The second AOM can only cover a frequency range of 1.4GHz between

---

[*] By efficiency we mean the ratio of the intensity of the first order light to that of the zero order.

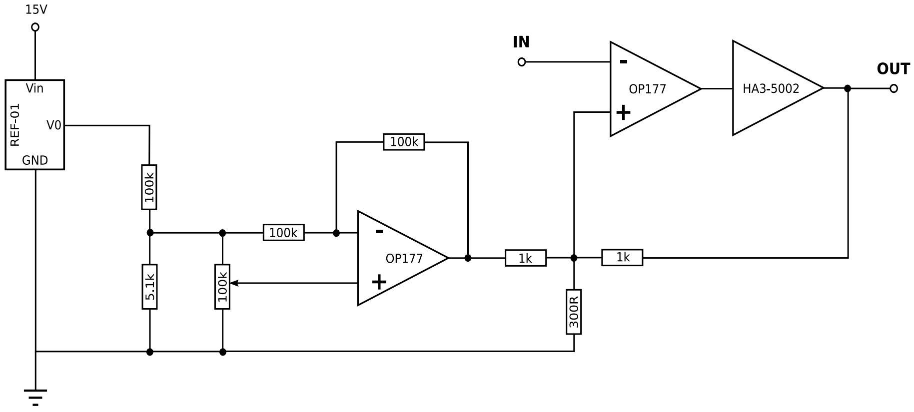

Figure 6.15. Buffer circuit for the feedback circuit. There is an offset remover before the buffer.

the two adjacent cavity modes. There is a 100MHz gap in the free spectral range of the cavity which cannot be addressed using the second AOM. If in the experiment the transition frequency happens to be in this frequency gap, the magnetic field can instead be changed to shift the transition frequency to the range that can be addressed by the AOM.

The EOM (Linos PM25) is modulated at 6.3MHz by a signal generator whose amplitude is amplified using a helical resonator, backed up with a commercial power amplifier (ENI 325LA), to 1200V. The helical resonator, which acts as an inductor, along with the EOM, which acts as a capacitor, form a LC circuit. The modulation signal is amplified when the modulation frequency matches to the resonance frequency of the LC circuit. With the assistance of the amplifier the ratio of the sidebands to the carrier can be as high as 100%. However the ratio is kept to 50% to get an optimum error signal (see figure 6.19(a)).

A perfect PDH error signal does not have any voltage offset when the laser is resonant with the cavity. In practice however a small offset can be observed which shifts the locking point. Possible sources of the offset are as follows:

1. Polarisation of the light incident on the EOM does not match to the axis of the EOM crystal. In this case the polarisation of the light is modulated by the EOM. Therefore the intensity of the beam after the polarising beam splitter will be modulated and that induces a voltage offset in the error signal. The polarisation of the light should be carefully matched to the crystal axis of the EOM.

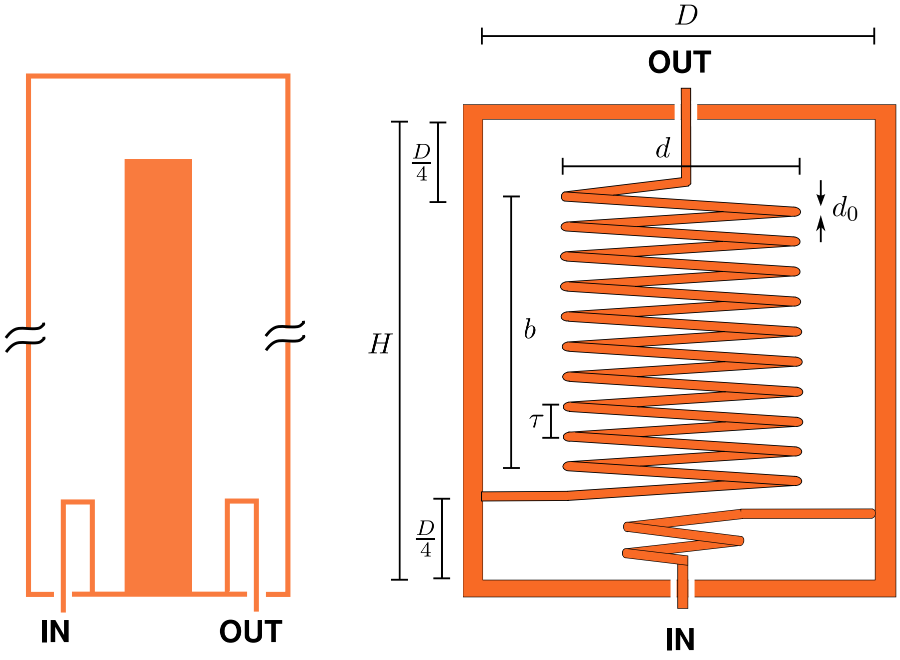

Figure 6.16. *Left:* Coaxial line resonator *Right:* Helical resonator with a secondary coil acting as antenna.

2. Residual amplitude modulations (RAM) in the EOM [88]. RAM arises from multipass-effects inside the modulator. The Brewster angle of the two crystals inside the EOM minimises the back reflections and thus the RAM.

3. Pick-up of the RF field modulating the EOM on the APD. This can be minimised by shielding the photodiode and the EOM.

In the experiment we found that the offset is usually on the order of a few mV. In order to remove this residual offset, an offset voltage is added to the error signal before feeding to the feedback circuit (see figure 6.15). Moreover the circuit acts as a buffer for the feedback circuit as well. The circuit diagram is shown in 6.15.

6.3.3 Helical Resonator

The modulation depth of the EOM depends on the driving voltage. The higher the voltage is, the bigger the modulation depth becomes. A transformer is needed to reach a high voltage of a few hundreds of volts. Building a transformer which has a high quality factor and is at the same time efficient for high frequencies is challenging. Coaxial line resonators are suitable for ultra high frequencies (see figure 6.16). They consist of a central antenna which is excited by a coupling loop at the resonance frequency and energy is removed by another loop. The problem is the length of the antenna that makes it impractical for intermediate and high frequencies.

A helical resonator is very similar to a coaxial-line quarter-wave resonator with the difference that the inner conductor is wound in a helix [89, 90]. They both share the feature of high quality factors at high frequencies. However a helical resonator is much smaller in size than a coaxial-line resonator. For example a 10MHz helical resonator with an unloaded Q of 1000 is about 20cm in length whereas the same coaxial line resonator is nearly 8m long.

The helical resonator is made of a solenoid which is enclosed in a shield. The shield can have a circular or square cross section. One end of the helical winding is directly connected to the shield and the other end is open circuit. Coupling into the helical coil can be done by two alternative methods: loop, and tap coupling. For loop coupling a loop of approximately one turn acting as an antenna is placed slightly below the helical coil. Alternatively the input can be directly tapped to the helical coil [91].

The equivalent inductance of the coil in μH is given by

$$L = 0.025 n^2 d^2 [1 - (d/D)^2] \text{ (μH per axial inch)}, \tag{6.6}$$

with all the relevant dimensions shown in figure 6.16. Empirically the capacitance of the coil for an air dielectric is

$$C = \frac{0.75}{\log(D/d)} \text{ (pF per axial inch)}, \tag{6.7}$$

when $b/d = 1.5$. All equations are valid when

$$\begin{aligned} &\frac{b}{d} = 1.5, \\ &0.45 < \frac{d}{D} < 0.6, \\ &0.4 < \frac{d_0}{\tau} < 0.6. \end{aligned} \tag{6.8}$$

The axial length b of the coil is approximately equivalent to a quarter wavelength inside the resonator*. But this length is much shorter than the free-

---

*To be more precise in practice it is 6% less than the quarter wavelength.

space length which is given by

$$\frac{\lambda}{4} = \frac{c}{4f_0}10^{-6}.$$

Here f_0 is the operating frequency in MHz and c is the speed of light. The helical resonator is similar to an LC circuit. The wave speed inside the coil is expressed by

$$v = f_0\lambda = \frac{1000}{\sqrt{LC}}.$$

Now if the length of the coil is equal to a quarter wavelength we have

$$b = \frac{\lambda}{4} = \frac{250}{f_0\sqrt{LC}}, \tag{6.9}$$

when C is in pF. Three equations 6.6, 6.7 and 6.9 along with the constraints in equations 6.8 determine all the dimensions of the helical resonator.

When we used to use a conventional transformer for driving the EOM, the sidebands were at maximum $\sim$ 5% of the carrier. It was decided to build a helical resonator instead which made a major improvement to the ratio of the sidebands to the carrier and then to the error signal.

A helical resonator was built with the calculated dimensions of D=4in, d=2.2in, b=3.3in, H=5.3in, N=16, $\tau = 0.2$in and $d_0 = 0.1$in. The secondary coil has 1.5 turns and a thickness of 2mm. The unloaded quality factor of the resonator was measured to be $\sim$ 85. The amplification factor is $\sim$ 30. The resonator is followed by a 50cm coaxial cable and the EOM. The EOM has capacitance of 30pF. The cable is a RG58 coaxial cable which has a capacitance of 80pF/m. The loaded resonance frequency of the resonator is 6.3MHz and this is highly dependent on the cable length. The resonator is 5 times more efficient than a similar transformer working around the same resonance frequency as seen by the ratio of the sidebands to the carrier.

6.3.4 Reference Cavity

The cavity has a length of 10cm while the radius of curvature of the mirrors is 35cm. This corresponds to a free spectral range of $\sim$ 1.5GHz. The mirrors are custom built by REO and the cavity has a quoted finesse of $\sim 250,000$ at 729nm*. The mirrors are made from ultra low expansion (ULE) glass and

---

*We have not directly measured this finesse and we are relying on the measurement performed by REO.

are optically contacted onto a 10cm ULE spacer. The thermal expansion coefficient of ULE is three orders of magnitude lower than that of fused silica if controlled at a temperature less than a degree away from its zero-crossing temperature at $\sim 30°$C [92]. The cavity rests on an aluminium V-groove with Viton O-ring in between.

The cavity is in a vacuum chamber pumped by a mini ion pump which keeps the pressure at $\sim 10^{-6}$mbar. The vacuum chamber helps in two ways. It firstly insulates the cavity from the temperature fluctuations of the room. At this pressure the mean free path of particles is $\sim$10cm which is on the order of the spacing between the vacuum chamber and the cavity, hence reducing the thermal conductivity dramatically. Secondly it keeps the pressure inside constant. Pressure fluctuations give rise to changes of the refractive index and therefore instabilities in the resonance frequency of the cavity.

The vacuum chamber is thermally insulated with hard Neoprene. The temperature is stabilised at 30°C using a PID controller similar to that used for stabilising the cavities for 866nm and 397nm lasers (see 5.9.1).

In order to achieve a reasonable error signal, the laser needs to be "mode matched" to the TEM_{00} mode of the cavity as well as possible. By "mode matching" we mean that the waist and the spatial position of the incident beam correspond to that of the cavity mode. The reason is that the error signal is proportional to the square root of the reflected carrier power times the reflected sideband power. The better the mode matching is, the stronger the reflection coefficient for the sidebands becomes. If a high-order mode of the cavity happens to be at a frequency where a sideband is, the sideband does not totally reflect back from the cavity. But most importantly the carrier light that leaks back from the cavity becomes stronger when the mode matching is better.

Mode matching

Two lenses with focal lengths of $f = -100$mm and $f = +100$mm along with two mirrors are used to match the waist of the beam to that of the cavity. The waist of the beam before the cavity was measured to be $\sim 510\mu$m using the knife edge method. Mode matching is achieved when the waist of the beam is adjusted to the waist of a standing wave mode inside the cavity. The distance between the two lenses and also the distance between the second lens and the cavity can be adjusted for this purpose. For a cavity with a length of L and

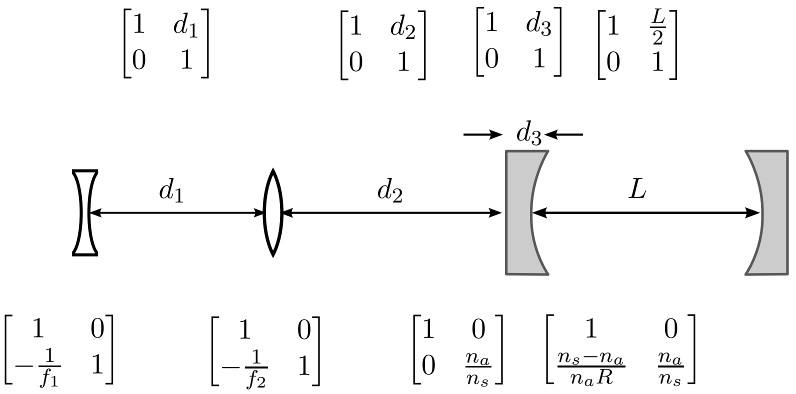

Figure 6.17. The corresponding ABCD matrices used for calculating the beam waist at the centre of the cavity.

radius of curvature of mirrors, R, the waist of the mode is equal to [93]

$$w_0 = \sqrt{\frac{\lambda}{\pi}}\left[\frac{L}{2}\left(R-\frac{L}{2}\right)\right]^{1/4}.$$

For a symmetric cavity with a length of $L = 0.1$m and mirror radius of curvature of $R = 0.35$m the equation above gives a waist of 168μm at the centre of the cavity.

In much the same way that the angle and position of a ray passing through optical elements can be determined using ABCD matrices [93], the waist of a Gaussian beam can also be derived using ray transfer matrix analysis by invoking the ABCD law for Gaussian beams. By applying the ABCD matrices and matching the waist of the beam in the cavity to the natural waist of the TEM_{00} mode in the cavity, the distance between the lenses and between the second lens and the cavity can be calculated. The corresponding matrices used in the calculation are shown in figure 6.17. A Mathematica script was written to calculate the optimal values for d_1 and d_2 (see appendix C). Assuming the input beam is collimated, the calculations give two separate lines which are the mode matching solutions for the respective beam waist and cavity specifications. For a given value of d_1 there are two values of d_2 that give good mode matching and vice versa. The calculated diagram using two lenses with focal lengths of $f_1 = -100$mm, $f_2 = +100$mm and a waist of initial beam of 510μm is shown in figure 6.18. One of the lenses is mounted on a translation stage to allow adjustment of the distance between the lenses.

The lateral position of the beam can be adjusted by two mirrors right before the cavity (see figure 6.9). The mirrors are after the two lenses so that

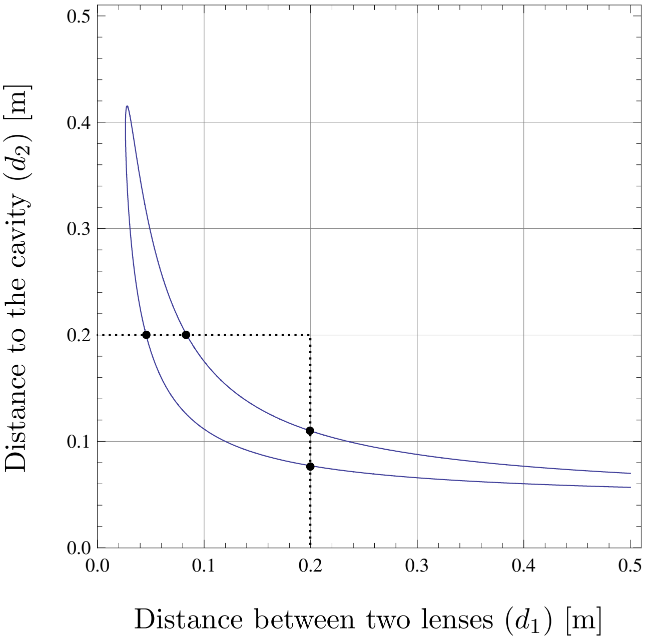

Figure 6.18. For any given separation between the two lenses there are two values for the distance to the cavity at which the best coupling can be achieved and vice versa.

adjustment of the position of the beam does not affect the waist of the beam.

To facilitate the mode matching procedure a photodiode and a CCD camera are placed after the cavity to detect the light which passes through the cavity. The laser is scanned over 1GHz in frequency using an external signal generator running at a frequency of ∼2Hz. The scanning frequency should be kept low otherwise the internal locking systems of the laser go out of lock. Images of different transverse mode shapes are seen transiently on the CCD camera by manually shifting the centre of the scan range. The TEM_{00} mode is identified and the scan range is reduced so that it encompasses only this mode. The TEM_{00} mode which is now the brightest mode in the scan range is also the biggest peak of the photodiode signal after the cavity.

The TEM_{00} intensity can then be improved using the two mirrors before the cavity by making the beam as central as possible to the cavity. The distance between the lenses can be further adjusted until the best mode matching is achieved. In the experiment the intensity ratio of the next biggest mode to the TEM_{00} mode is less than 0.01.

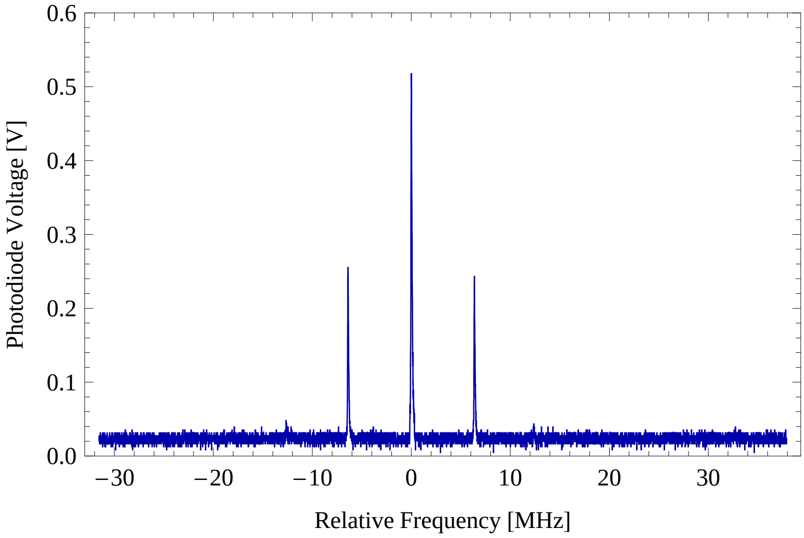

(a)

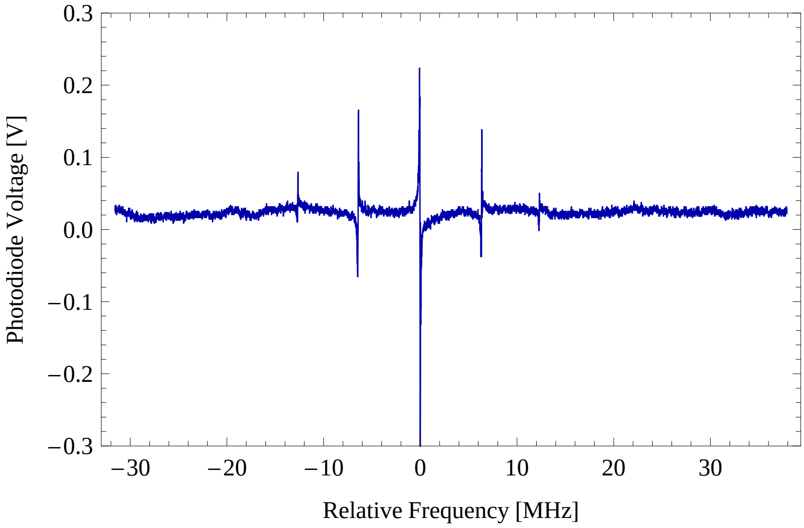

(b)

Figure 6.19. a) Signal of the photodiode after the cavity. Sidebands are half the carrier to maximise the error signal. b) The error signal

6.3.5 Error Signal

The APD signal is amplified using an amplifier (Mini-Circuits MAN1-LN). The output is then mixed with the signal from the signal generator that drives the EOM by a Mini-Circuit mixer (ZRPD-1). The mixed signal is filtered with a low-pass filter (Mini-Circuits BLP-5) with a cutoff frequency at 6MHz to get the error signal. The error signal, is monitored on an oscilloscope and by changing the length of the cable between the mixer and APD or between the signal generator and the RF amplifier, the phase of the error signal can be adjusted. The error signal achieved with this configuration is shown in figure 6.19(b).

6.3.6 Loop Filter

Normally feedback controllers employ the PID technique. PID controllers mainly comprise of three stages which work in parallel: 1) P: where the control signal is proportional to the difference between the input signal and the set-point 2) I: the integral of the difference between the input signal and the set-point and 3) D: the derivative of the difference between the input signal and the set-point. Each stage has its own adjustable gain. The signal is then inverted and fed back to the compensator.

Some implementations of PDH (for e.g. Toptica) use the PID technique. In the beginning a PID controller was designed and built, however it was not successful in locking the laser to the cavity. The original Drever-Hall paper suggests a different approach “loop filter” [84]. Its function is less intuitive than a PID controller but is very well matched to the peculiarities of the PDH method i.e. it tailors the gain in frequency space to match the characteristics of the optical system. Being unsuccessful using a PID controller, we opted for the Loop filter approach originally suggested for the feedback circuit.

As mentioned before, the error signal is made of two components. The first is the field (which includes the sidebands) that is directly reflected from the first mirror of the cavity and is in phase with the laser. The second component is the leakage of the internal field from inside the cavity. When the frequency jitter of the laser is within the cavity linewidth, the laser has enough time to stabilise inside the cavity and therefore the frequency of the light inside the cavity follows the frequency of the laser itself. In this range, the error signal acts as a frequency discriminant. However, when the jitter is faster than the half-width-half-maximum (HWHM) of the cavity resonance ν_c, the

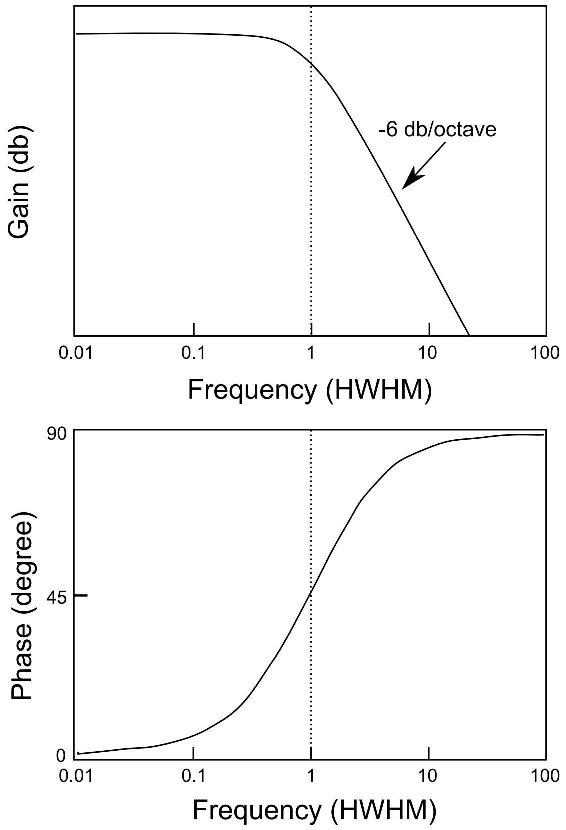

Figure 6.20.

cavity averages the input field and it acts as a phase storage element. In this frequency range the system acts as a phase discriminator due to the storage time of the reference cavity. In other words, for Fourier frequency jitter ν faster than the cavity linewidth ν_c, the phase of the leakage field ϕ is equal to

$$\phi = \int \nu dt. \tag{6.10}$$

As a result there is a 90° phase lag between the leakage field and the laser. Also when the laser jitter is faster than the cavity linewidth the cavity response drops smoothly by 6 db/octave [94]. The Bode plot of the reference cavity response is shown schematically in figure 6.20.

In order to take advantage of both of these features i.e. frequency and phase discriminator of the PDH lock a servo amplifier tailored to this application is needed. This is usually called a “loop filter”. The special role of the loop filter is to introduce a 90° phase shift and also to compensate for the roll-off of the cavity response for Fourier frequencies higher than the cavity linewidth.

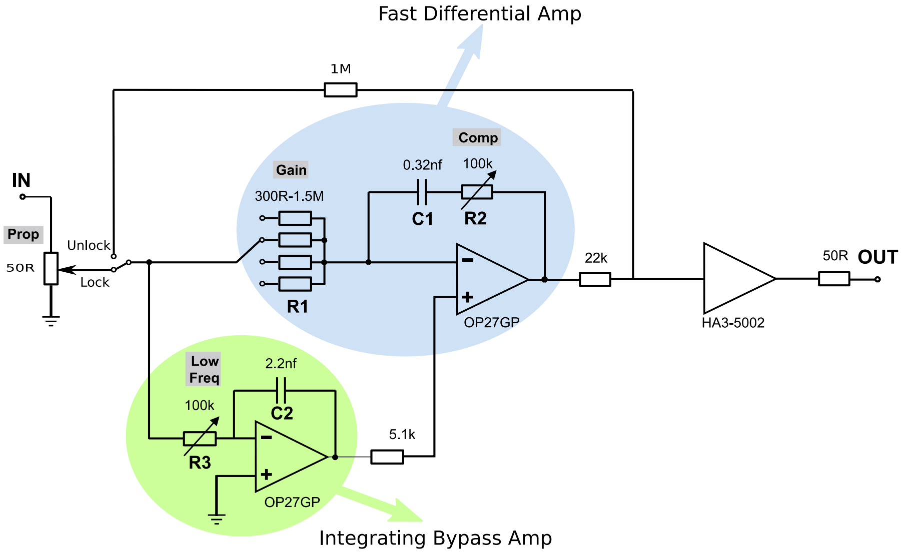

Figure 6.21. Loop filter with bypass topology

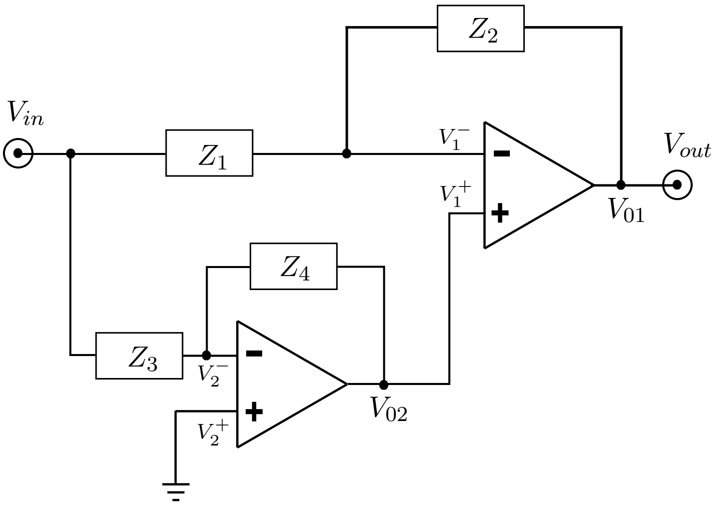

Figure 6.22. Loop filter simplified circuit diagram

A loop filter with bypass topology is the circuit used for the feedback which utilises the two key features of the PDH technique for locking. It has two stages which together act over a frequency range which is divided into two sub-ranges. The breakpoint between the two sub-ranges is set at the cavity linewidth. For frequency components lower than the cavity linewidth the signal passes through an integrating op-amp and is fed into the fast differentiator amplifier's input. For frequency components higher than the cavity linewidth the integrating op-amp is bypassed and the signal passes only through a fast differentiator op-amp. The loop filter circuit diagram used in the experiment is shown in figure 6.21.

To analyse the response of the circuit with respect to the input frequency we can simplify the circuit as in figure 6.22. The two golden rules of operational amplifiers are:

1. The voltage difference between the inputs of an ideal op-amp is zero.
2. The input of an ideal op-amp does not draw any current.

Using these two rules one can derive the set of equations

$$\begin{aligned}
&\frac{V_{in}-V_1^-}{Z_1}+\frac{V_{01}-V_1^-}{Z_2}=0,\\
&\frac{V_{in}-V_2^-}{Z_3}+\frac{V_{02}-V_2^-}{Z_4}=0,\\
&V_1^-=V_1^+=V_{02},\\
&V_2^-=V_2^+=0,
\end{aligned}$$

where the various voltages are depicted in figure 6.22. The impedances of the simplified circuit are as follows (assuming $C_1 = C_2 = C$):

$$\begin{aligned}
Z_1 &= R_1,\\
Z_2 &= \frac{1}{C\omega}+R_2,\\
Z_3 &= R_3,\\
Z_4 &= \frac{1}{C\omega},
\end{aligned}\tag{6.11}$$

where R_1, R_2, R_3, C_1 and C_2 are shown in figure 6.21 and ω is the Fourier frequency of the circuit. By solving this set of equations one can show that

$$\frac{V_{01}}{V_{in}}=-\left(\frac{Z_2}{Z_1}+\frac{Z_2}{Z_1}\frac{Z_4}{Z_3}+\frac{Z_4}{Z_3}\right).$$

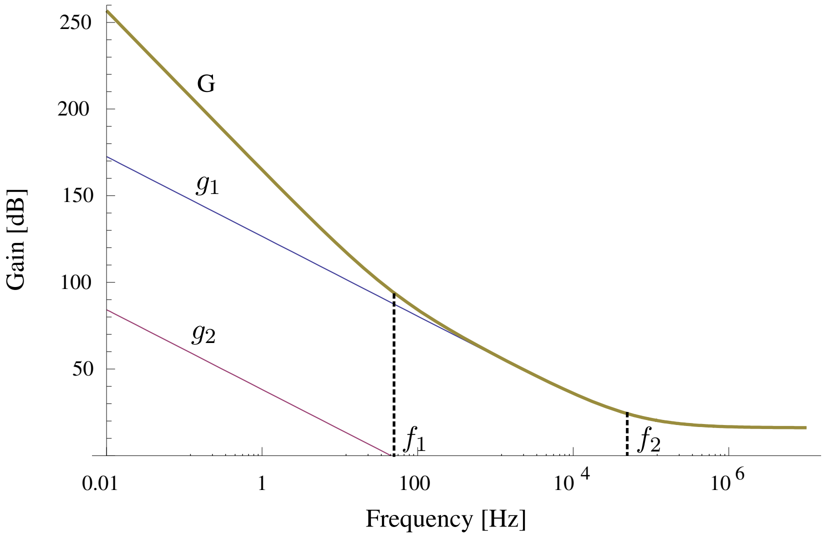

Figure 6.23. The Bode plot for ideal op-amps is shown for the case when $R_1 = 10\text{k}\Omega, R_2 = 50\text{k}\Omega, R_3 = 1\text{k}\Omega, C_2 = 0.32\text{nF}$ and $C_4 = 22\mu\text{F}$. g_2 is a straight line but g_1 has asymptotic behaviour to a constant gain at high frequencies.

In general the total gain G can be written as

$$G = g_1 + g_1 g_2 + g_2, \tag{6.12}$$

where $g_1 = Z_2/Z_1$ is the gain of the first op-amp and $g_2 = Z_4/Z_3$ is the gain of the second op-amp. The first op-amp is a differential amplifier for fast frequency components and the second op-amp is an integrating bypass amplifier. Therefore g_1 is high for high frequencies while g_2 is only high for slow frequencies. The cross term gives gain in intermediate frequencies.

Using the PDH lock combined with the loopfilter allows having the unity gain at a higher frequency than the cavity roll-off and provides corrections for frequency fluctuations over a broad bandwidth. The Bode plot (gain vs. frequency) of the circuit is shown in figure 6.23. The first breakpoint happens at the frequency f_1 which should correspond to the cavity linewidth δ_ν. With no breakpoint the circuit oscillates due to the high integral gain. The second breakpoint happens at f_2. It compensates for the roll-off of another element in the loop, which is the AOM in our case. Therefore it enables one to increase the whole gain of the system and have the unity gain at a frequency higher than f_2.

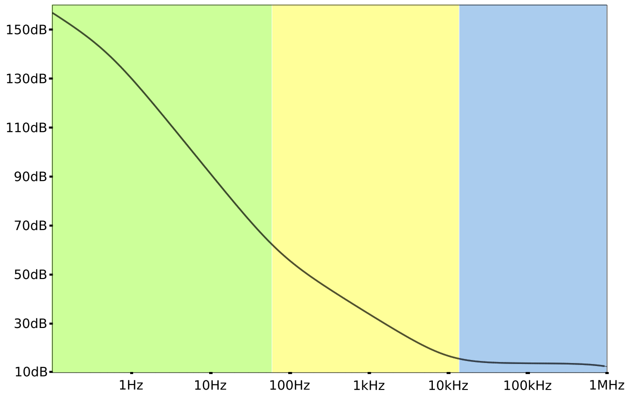

Figure 6.24. Simulated Bode plot for gain=10kΩ, comp=50kΩ and LowFreq=1kΩ. Plot is divided into three sections. Low frequency coloured as green whose gain is high, intermediate frequency as yellow with middle gain and high frequency as blue with low gain.

Substituting corresponding impedances into the equation above gives

$$G = \frac{1}{R_1 R_3}\frac{1}{(C\omega)^2} + \frac{R_1 + R_2 + R_3}{R_1 R_3}\frac{1}{C\omega} + \frac{R_2}{R1}.$$

The first term is large when the frequency is low while the second term is significant for intermediate frequencies and the last term for high frequencies. Decreasing R_1 increases the total gain and decreasing R_3 gives higher gain for low frequency. Increasing R_2 increases the high frequency gain. With the combination of these three parameters the gain profile can be adjusted appropriately.

In equation 6.12 it was assumed that the op-amps are ideal. To look at the circuit more precisely we simulated the circuit using an open-source circuit simulation program called Spice [95]. Plotting the total gain versus frequency gives the Bode plot shown in figure 6.24. The Bode plot consists of three regions as the calculations show. First region working at low frequencies has a very high gain. The gain is -12dB/octave. The middle region, effective in intermediate frequencies, has lower gain of -6dB/octave and finally unit gain for high frequencies. In practice the circuit is built with a number of potentiometers acting as controls. The names given to the various controls are not entirely transparent. Increasing the impedance of "compensation" increases the very high frequency gain. Lowering the impedance of the "gain" increases the total gain of the circuit. Finally lowering the "low frequency" impedance increases the low frequency gain.

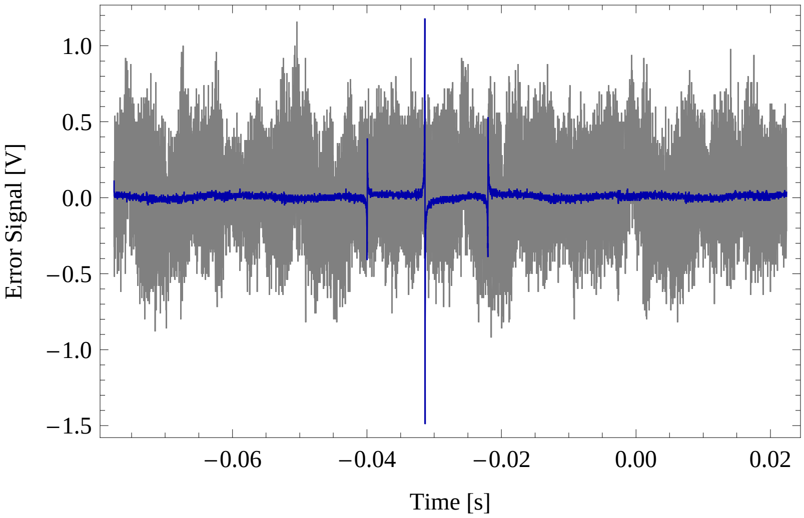

(a) The error signal after the mixer before locking (blue) and after locking (gray) is shown. The scanning frequency is 70MHz per second.

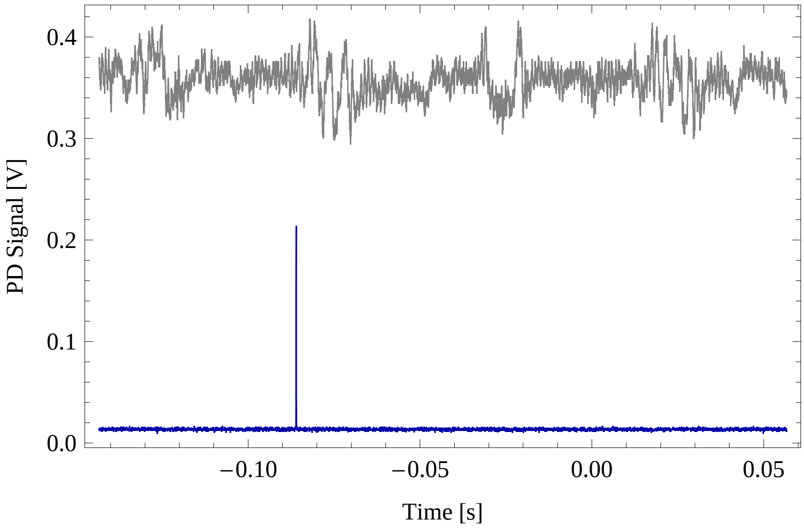

(b) The photodiode signal after the cavity when the laser is scanned (blue) and after it is locked (gray). The smaller peak for the case of scanning compared to the locked level is due to the slow response of the photodiode as the laser was scanned over the narrow linewidth cavity resonance.

Figure 6.25.

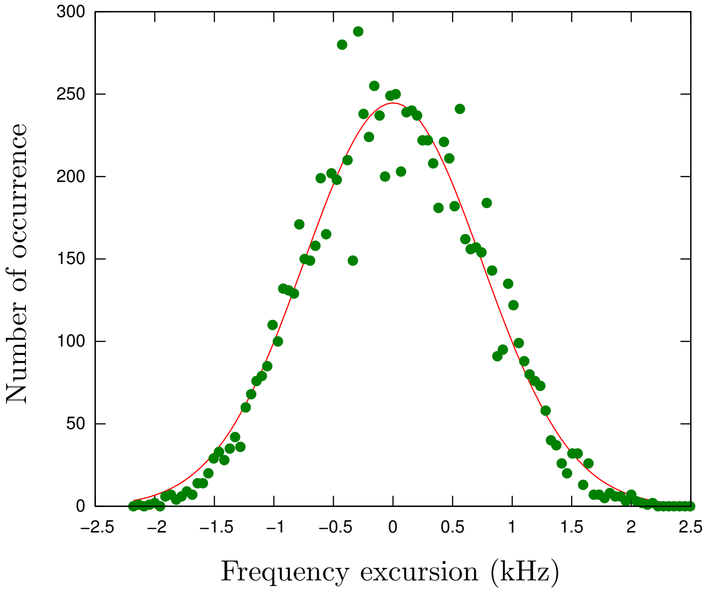

Figure 6.26. Histogram of the frequency excursions of the laser when locked. The data is fitted to a Gaussian function with a full width at half maximum (FWHM) of 1.7kHz.

6.3.7 Results and Future Directions

Using the loop filter we achieved a successful lock to the cavity (see figure 6.25(a)). The linewidth that was estimated was ∼1.7kHz* with a maximum excursion of ∼3kHz. This was inferred from the error signal when the laser was locked and comparing that to the linewidth of the cavity which was assumed to be 6kHz. Frequency excursions of the laser when the laser was locked is shown in figure 6.26. Although working, the lock was clearly not as tight as that achieved by other groups [97].

Taking a fast Fourier transform of the error signal showed clear and significant oscillations at 88kHz (see figure 6.27(a)). This frequency is in fact the dithering frequency of the thin etalon inside the ring cavity of the laser. As discussed in section 6.3.1 the thin etalon is responsible for locking the laser to a mode of its ring cavity. This dithering results in an unnecessary ripple of 88kHz on intensity and frequency. In principle laser electronics should make these intensity fluctuations disappear when the etalon is locked. Discussion with other users of these lasers suggest that these fluctuations are indeed a significant problem. The frequency noise can be minimised by empirical adjustments of the lateral position of the etalon in the cavity. Although the PDH

---

[*] This is only an estimate for the linewidth. A more accurate method to measure the linewidth can be found in [96].

technique is resilient to intensity fluctuations this is a second order effect and can be transfered into frequency fluctuations if the intensity fluctuations are big enough.

The etalon noise can be reduced by lowering the dithering amplitude of the etalon. The etalon dithering signal is a sawtooth and has a height of 1V peak to peak. This amplitude can be reduced to a tenth of that while still keeping the laser locked, although by lowering the dithering amplitude the laser will not be able to scan the full range without mode hopping. Also the beam position on the etalon can be adjusted to minimise the frequency noise. Furthermore the first AOM can be used to compensate for the intensity fluctuations (using it as a “noise eater”). This could be achieved by monitoring the intensity of the light on a fast photodiode to generate an error signal and feeding it to the AOM driver to compensate for the AOM modulation power.

To see how big the effect of the 88kHz noise is we have removed the 88kHz peak in the FFT and then reconstructed the error signal by taking an inverse Fourier transform. Figure 6.27(b) shows the result of this. It is worth mentioning that a large part of the locked error signal is at 88kHz, which is beyond the cavity bandwidth. Therefore the frequency excursion at 88kHz are even larger due to the roll-off in the cavity response. The laser lineshape has a Gaussian carrier with two sidebands at 88kHz approximately half the carrier height. If the etalon dithering problem was fixed first the locking would be tighter. Without this noise it might be also possible to increase the gain of the circuit and as a result a tighter lock would be achieved. Although the laser linewidth achieved is not the limit, we believe a 1.7kHz linewidth should be enough for observing the 729nm transition and also sideband cooling. Given the very narrow linewidth, an effective test requires a Penning trap with a very stable magnetic field. This means having to operate a trap in our superconducting magnet rather than a conventional electromagnet.

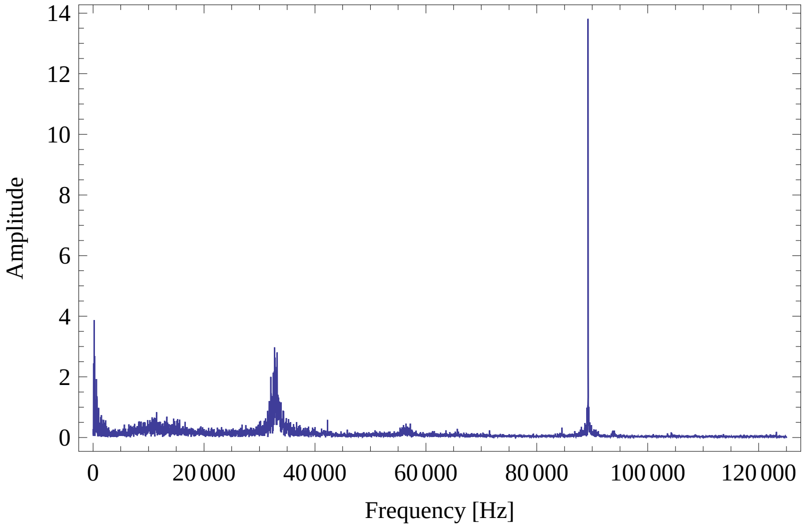

(a)

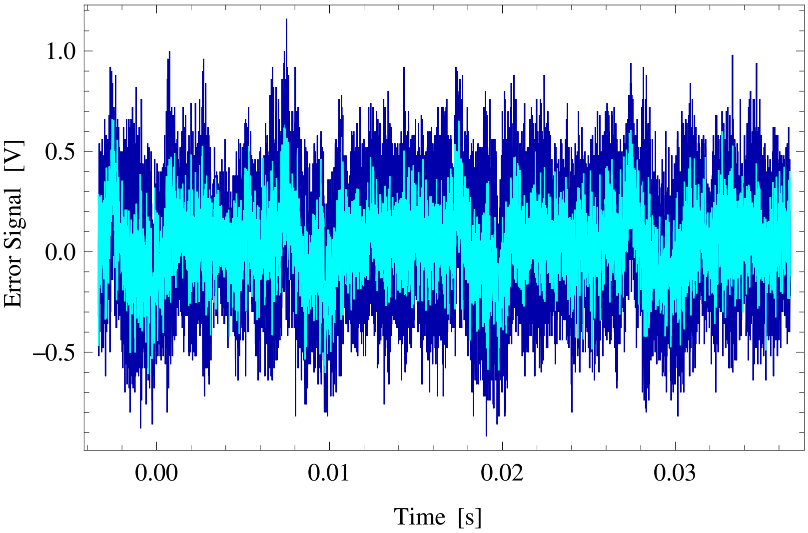

(b)

Figure 6.27. a) The FFT of the error signal when locked to the cavity. b) The reconstructed error signal without the 88kHz noise is shown here (with cyan colour). The blue signal is the original error signal observed with the presence of the 88kHz noise in the laser.

CHAPTER 7

Coherent Population Trapping

7.1 Introduction

We were able to trap and laser cool single Ca^+ ions when running the split-ring trap in the RF mode (see figure 3.1). However, as also mentioned before, the signal count rate per ion was significantly lower ($\sim 1/10$) in the Penning mode than that of the RF mode. This prompted us to investigate the possibility that "coherent population trapping" might decrease the laser cooling efficiency in the Penning trap. Coherent population trapping (CPT) is a process that can hide population in the so called "dark states". Given the low fluorescence count rate that we experienced during the first attempts to trap single Ca^+ ions it was important to consider this process. This chapter explains the work which I carried out to investigate CPT in our laser cooling scheme.

In 1976 in an experiment done by G. Alzetta and coworkers [98] something unexpected was observed. The experiment was laser optical pumping of sodium atoms in a vapour cell. The level scheme involved comprised the two hyperfine sublevels of the $3S_{1/2}$ ground state and one hyperfine sublevel of the $3P_{1/2}$ excited state. A non-homogeneous magnetic field was applied along the vapour cell. There were two lasers which were in resonance with the two S→P transitions. They were attempting to see more fluorescence by repumping an optically pumped level. However the opposite occurred. If only one laser was turned on bright fluorescence was emitted in the vapour cell. But when both

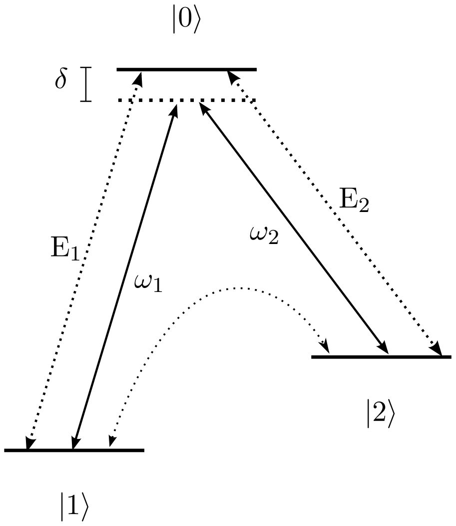

Figure 7.1. Coherent population trapping (CPT) for the simple case of a three-level system in the Λ configuration. When the two lasers have the same detuning the two photon resonance condition is satisfied ie. $E_1 - E_2 = \hbar(\omega_1 - \omega_2)$. In this case level $|0\rangle$ is not populated and the state of the system goes to a superposition of the $|1\rangle$ and $|2\rangle$ states.

lasers were simultaneously shone into the vapour cell the fluorescence disappeared in a certain region in the cell. The non-fluorescing lines appeared as dark lines inside a bright fluorescence path hence the term “dark state” was coined.

We define a “dark state” to be a state which, when populated, results in a drop in fluorescence. In the laser cooling of Ca^+ ions the ion, once in the fluorescing $P_{1/2}$ excited state can decay to the metastable state $D_{3/2}$ and stop fluorescing[*]. Therefore a repumping laser is used to deplete this state. In terms of the above definition the $D_{3/2}$ state is considered as a “dark state”. However the “dark state” of Alzetta et al experiment was subtly different. It turned out to be a superposition of states that gives rise to fluorescence reduction although each individual state was repumped by a laser. This is referred as coherent population trapping (CPT).

CPT is quantum interference between the atomic sublevels. Seeing quantum coherence in such a macroscopic system was in fact quite surprising. This phenomenon has many applications in modern optics. Lasing without inversion [99], laser cooling below the recoil limit [100], electromagnetically induced transparency [101] and laser isotope separation [102] are among them.

The simplest case for CPT occurs in a three-level system. The energy

---

[*] This process is generally called “optical pumping”.

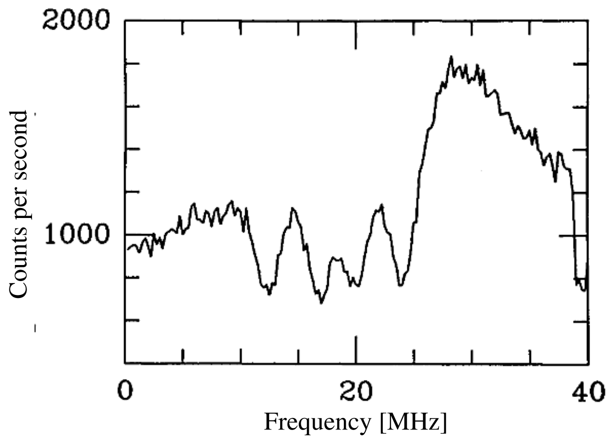

Figure 7.2. The fluorescence signal of a trapped Ba$^+$ ion in an RF field with 0.86G magnetic field. The blue laser detuning is 10MHz. 4 dips in the signal can be seen. The signal drops when the repumping and cooling laser make a trapped state. Taken from [103].

levels are shown in figure 7.1. In this scheme which is the so called Λ system, two lasers with frequencies ω_1 and ω_2 couple the ground states $|1\rangle$ and $|2\rangle$ to an excited state $|0\rangle$. When the two lasers are detuned from the excited state by the same detuning δ, the two photon resonance condition

$$E_1 - E_2 = \hbar(\omega_1 - \omega_2)$$

is satisfied and the population is transferred to a steady state which is a coherent superposition of the ground states $|1\rangle$ and $|2\rangle$. Since the population in the steady state is trapped in the ground state, the excited state is not populated at all and fluorescence disappears completely [*]. This situation is similar to the Raman transition discussed in section 3.4.

CPT has also been observed in laser cooling of ions in the RF trap. This was first observed experimentally with Ba$^+$ ions two decades ago [103] in the Dehmelt group. The level scheme of Ba$^+$ ion is very similar to Ca$^+$. The laser cooling transition is the $S_{1/2} \rightarrow P_{1/2}$ transition at 493nm. The ion decays to the metastable level $D_{3/2}$ therefore a repumping laser driving the $D_{3/2} \rightarrow P_{1/2}$ transition at 650nm is needed to deplete the population of this level. The ion is trapped in a RF trap. A small magnetic field is introduced in order to avoid the dark states on the S-P transition. There are 4 possible ways

---

[*] Initially state $|0\rangle$ is in fact populated but as the system evolutes, after a few steps the excited state depletes and a steady state is achieved with no population in state $|0\rangle$.

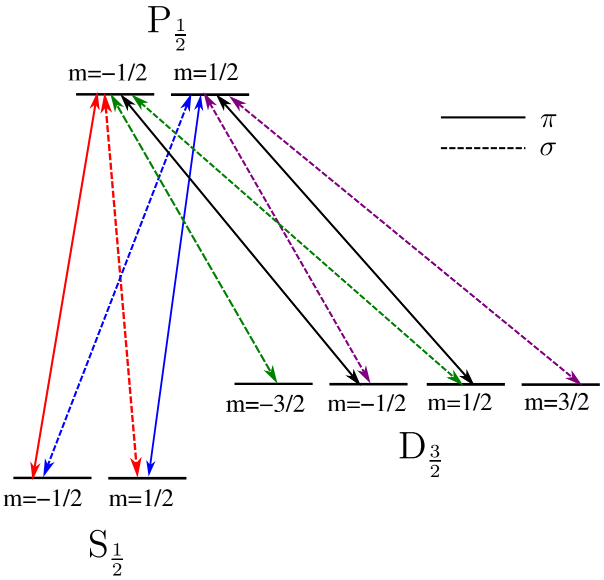

Figure 7.3. Ca^+ ion energy levels in zero magnetic field. Solid lines show π transitions and dotted lines show the σ transitions. There are two possible Λ systems in the D → P manifold with σ transitions which are shown with the green and purple arrows. There are also two CPT states for the UV laser in the S → P manifold with a mixture of σ and π transitions shown with the blue and red arrows. These could be avoided if only σ or π transitions are used.

that the two lasers form a Λ system in S-P-D transitions (see a similar level diagram for Ca^+ shown in figure 7.3). The experimental result is shown in figure 7.2 as 4 dips in the fluorescence signal as the repumping laser frequency is scanned. There are four frequencies at which the repumping laser comes into two-photon resonance with the laser cooling laser and the population gets trapped in a superposition of the S and D states.

The energy level diagram of Ca^+ in the absence of any magnetic field is shown in figure 7.3. This situation is relevant to our system if we switch off the magnetic field and trap as an RF trap instead. In the experiment the cooling laser in the UV, addressing the $S_{1/2} \rightarrow P_{1/2}$ transition, is detuned to the red and the repumping IR laser is resonant with the $D_{3/2} \rightarrow P_{1/2}$ transition to repump the $D_{3/2}$ metastable state. The detuning of the blue laser should be equal to half the transition linewidth for optimum laser cooling. Because the frequency detunings of the two lasers are not the same there is no Λ configuration in the S-P-D levels when the lasers are optimised for laser cooling. Also for the S-P transition, provided that either σ or π transitions are used, no CPT can occur. Possible CPT systems that can happen with a mixture of σ and π transitions are shown in figure 7.3 by blue and red arrows.

The situation is different for the D-P transitions. To repump all 4 metastable $D_{3/2}$ states at least two σ transitions must be used in order to repump the $m = -3/2$ and $m = 3/2$ states. The repumping laser can make two possible Λ systems in the D-P transition. This is depicted in figure 7.3 by purple and green arrows. In order to avoid this, a very small magnetic field (a few gauss) must be introduced to split the levels. The Zeeman splitting should be big enough to avoid CPT but not bigger than the transition linewidth itself to reduce the scattering rate. For frequency standards experiments, applying a magnetic field is not desirable because it shifts the clock transition. Instead, the polarisation of the IR laser is rapidly modulated by acousto-optic modulators (AOM) [104].

As mentioned in chapter 5, we were not able to laser cool single Ca^+ ions in the Penning trap due to the low signal per ion. The number of levels involved in laser cooling Ca^+ ions is significantly bigger in the Penning trap than that in the RF trap (see figure 5.1). Trapped states in the laser cooling scheme might reduce the cooling efficiency. In big clouds the ions move in and out of the beam, breaking the coherence. Thus for big clouds the laser cooling process might be compromised by trapped states but not stopped completely. However this mechanism could be significant for small numbers of ions (which stay within the laser beam) and could stop the cooling process. Naively one would not expect CPT to cause a problem since the blue laser is deliberately detuned for cooling while the repumper is put on resonance, therefore CPT should not happen. However the system was modelled in order to check this thoroughly.

We use π transitions for the laser cooling transition and σ transitions for repumping. The energy levels of Ca^+ in our laser cooling scheme can be divided into two four-level systems which are interconnected by a spontaneous emission decay (see figure 7.4). Due to the magnetic field required for trapping ions in the Penning trap the Zeeman splitting is many orders of magnitude bigger than that required to avoid CPT in the D-P transition. We routinely observe CPT, similar to that seen in the RF trap on the S-P-D transition (see figure 7.5). However just as in the RF trap, it is always possible to avoid these dark states by adjusting the laser frequencies. So this type of dark state should not concern us. However as mentioned in chapter 5 we were using only two lasers with imposed sidebands for the $D_{3/2} \rightarrow P_{1/2}$ transition at 866nm. We would not expect to have CPT by the mere fact that sidebands were used rather than four separate lasers. The sidebands were always detuned from the

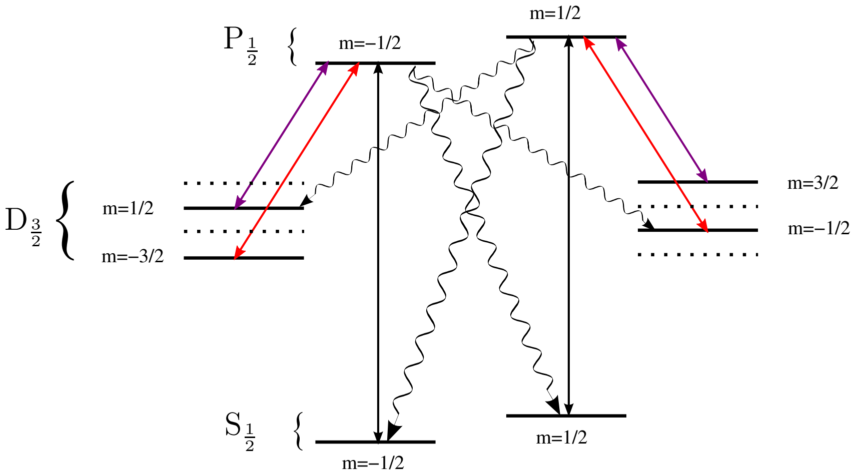

Figure 7.4. The 8-level cooling scheme of Ca^+ ion in the Penning trap. Cooling can be split into two 4-level systems which are connected with spontaneous emission decay. The sidebands of the diode laser address different systems. The sidebands of the same laser are shown with same colours.

D-P transition components by the same amounts but they were addressing two levels of different 4-level systems. The spontaneous decay breaks any coherence in the system.

However CPT can happen in more complex systems than just a simple Λ system. To investigate this, numerical calculations based on the density matrix approach were done to calculate the population of the various levels in the presence of a magnetic field.

7.2 Theory

The CPT theory of a 3-level system can be found in detail in the review paper by E. Arimondo [105]. This section merely extends this to 4-level systems. In order to simplify the calculations, rather than simulating an 8-level system, we split the system into two similar systems which are connected by a spontaneous decay. There is no coherence between the two systems due to the spontaneous decay. Therefore we can reduce the problem to a 4-level system as shown in figure 7.6. The excited state ($P_{1/2}$) is labeled as $|0\rangle$ with energy E_0. The population in this level determines the amount of fluorescence. The relevant lower states are labeled as $|1\rangle$, $|2\rangle$ and $|3\rangle$ with corresponding energies of E_i ($i = 1, 2, 3$). It is assumed that only three transitions of $|i\rangle_{i=1,2,3} \rightarrow |0\rangle$ are possible. The energy difference between the two levels $(0, i)$ is denoted

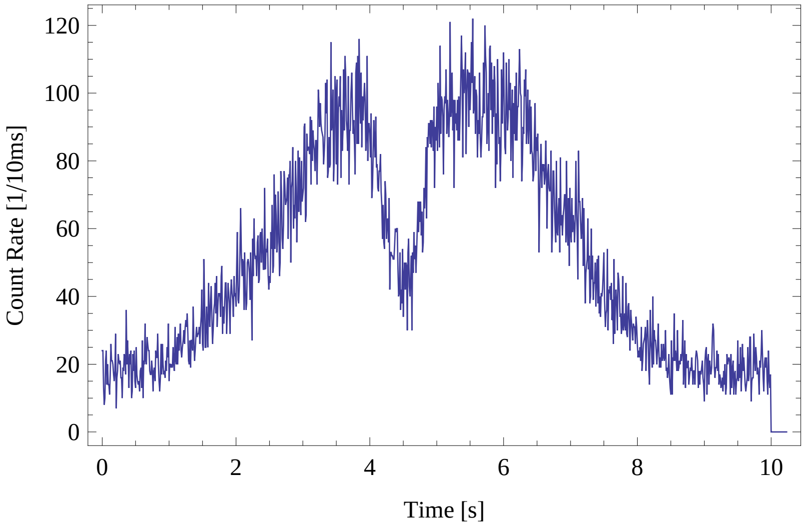

Figure 7.5. Dark state on the S-P-D transition of Ca$^+$ in the Penning trap. This trace is taken by scanning one blue laser while all other lasers are locked to their reference cavities. As the blue laser comes into a two-photon resonance with an IR laser CPT occurs and the fluorescence rate drops. This figure is courtesy of D. Crick and I. Bhatti.

as $E_0 - E_i = \hbar\omega_i$ $(i = 1, 2, 3)$. The frequencies of the lasers are denoted as ω_{Li} $(i = 1, 2, 3)$. The detunings of the lasers from the excited state $|0\rangle$ are defined as

$$\delta_{Li\ (i=1,2,3)} = \omega_{Li} - \omega_i.$$

We also define the Raman detuning δ_{Rij} from the two-photon transition of $|i\rangle \rightarrow |j\rangle$ as

$$\begin{aligned}\delta_{Rij} =& (\omega_{Li} - \omega_{Lj}) - (\omega_i - \omega_j)\\ =& \delta_{Li} - \delta_{Lj}.\end{aligned}$$

The Rabi frequencies characterising the laser interaction are

$$\Omega_{Ri} = -\frac{\mu_i \mathcal{E}_{Li}}{\hbar},$$

where μ_i is the electric dipole moment for the $|i\rangle_{i=1,2,3} \rightarrow |0\rangle$ transition and $\mathcal{E}_{Li}$ is the electric field for the corresponding laser. Following the argument in section 3.3 the atom-laser interaction can be written with the Hamiltonian

$$\mathcal{H} = \mathcal{H}_0 + \mathcal{V}_I,$$

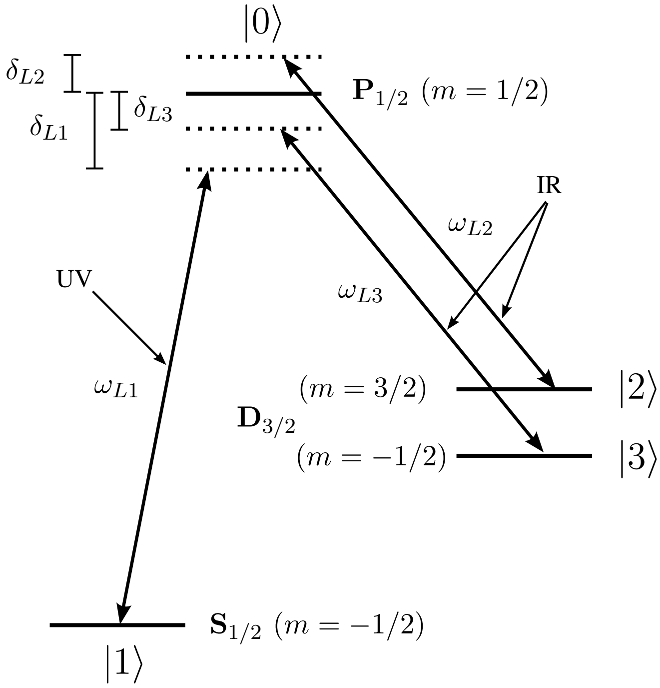

Figure 7.6. The four-level system which is modelled is shown.

where $\mathcal{H}_0$ is the unperturbed Hamiltonian of the atom and is equal to

$$\begin{aligned}\mathcal{H}_0 =&\sum_i E_i|i\rangle\langle i|\\ =&E_0|0\rangle\langle 0| + E_1|1\rangle\langle 1| + E_2|2\rangle\langle 2| + E_3|3\rangle\langle 3|.\end{aligned}$$

$\mathcal{V}_I$ is the interaction Hamiltonian of light with the atom and is given by

$$\begin{aligned}\mathcal{V}_I =&\frac{\hbar\Omega_{R1}}{2}\exp[-i(\omega_{L1}t + \phi_1)]|0\rangle\langle 1| +\\ &\frac{\hbar\Omega_{R2}}{2}\exp[-i(\omega_{L2}t + \phi_2)]|0\rangle\langle 2| +\\ &\frac{\hbar\Omega_{R3}}{2}\exp[-i(\omega_{L3}t + \phi_3)]|0\rangle\langle 3| + \text{h.c.},\end{aligned}$$

where we have used the standard rotating wave approximation to eliminate the non-resonant terms (for reference see [106]). The state of atom is described by the wave function

$$|\psi\rangle = \sum_i c_i|i\rangle.$$

Rather than solving the Schrödinger equation directly, a density operator is defined as

$$\rho = |\psi\rangle\langle\psi|.$$

and an evolution equation is solved accordingly. To clarify what the density operator is let us consider a two-state system with a normalised wave function such as

$$|\phi\rangle = c_0|0\rangle + c_1|1\rangle.$$

Since the wave function is normalised, $c_0^2 + c_1^2 = 1$. Expanding the density operator for this system yields

$$\rho = c_0^2|0\rangle\langle 0| + c_0 c_1^*|0\rangle\langle 1| + c_0^* c_1|1\rangle\langle 0| + c_1^2|0\rangle\langle 0|.$$

The diagonal terms show the population of the states and the off-diagonal terms give the *coherences* of the states. These terms distinguish a mixed state from a coherent state. It can be seen that

$$\begin{aligned} \rho_{ij} &= \rho_{ji}^*, \\ \text{Trace}[\rho] &= 1. \end{aligned} \tag{7.1}$$

Taking the time derivative of the density operator gives rise to

$$\dot{\rho} = |\dot{\psi}\rangle\langle\psi| + |\psi\rangle\langle\dot{\psi}|. \tag{7.2}$$

Substituting the Schrödinger equation for the general wave function $|\psi\rangle$

$$\mathrm{i}\hbar|\dot{\psi}\rangle = \mathcal{H}|\psi\rangle$$

into equation 7.2 gives rise to the Liouville equation

$$\begin{aligned} \mathrm{i}\dot{\rho} =& \frac{1}{\hbar}(\mathcal{H}\rho - \rho\mathcal{H}) \\ =& \frac{1}{\hbar}[\mathcal{H}, \rho]. \end{aligned}$$

In our four-level system the density matrix equation can be rewritten as

$$\mathrm{i}\dot{\rho}(t) = \frac{1}{\hbar}[\mathcal{H}_0 + \mathcal{V}_I, \rho(t)]\ .$$

This equation does not give the full description for the 4-level laser cooling scheme that we are interested in simulating because it does not take into account the spontaneous emission. In order to include the spontaneous emissions, collisions and any other damping effects, a relaxation operator $\mathcal{R}(t)$ is added to the equation. Since we are interested in simulating just a single ion, collisions are neglected. To include spontaneous emission the density matrix equation is modified to

$$\begin{aligned}
\mathrm{i}\dot{\rho}(t) =&\frac{1}{\hbar}[\mathcal{H}_0 + \mathcal{V}_I, \rho(t)] + \mathcal{R}\rho(t) \\
=&\frac{1}{\hbar}[\mathcal{H}_0\rho(t) + \mathcal{V}_I\rho(t) - \rho(t)\mathcal{H}_0 - \rho(t)\mathcal{V}_I] + \mathcal{R}\rho(t) \\
=&I(t) + \mathcal{R}\rho(t). \qquad (7.3)
\end{aligned}$$

The equation above gives the full description of the system. For a four-level system the density matrix has 16 elements but the constraints in equations 7.1 reduce the number of independent differential equations to 10. As an example the differential equation for the component ρ_{ij} of the density matrix is

$$\mathrm{i}\dot{\rho}_{ij}(t) = \langle i|I(t)|j\rangle + \langle i|\mathcal{R}\rho(t)|j\rangle.$$

It is easier to express the density matrix equation in the interaction picture to eliminate the dependencies on the time evolution operator

$$\begin{aligned}
\tilde{\rho}_{0i} =&\rho_{0i} \exp[\mathrm{i}(\omega_{Li}t + \phi_i)] \quad (i = 1,2,3), \\
\tilde{\rho}_{ij} =&\rho_{ij} \exp\{-\mathrm{i}[(\omega_{Li} - \omega_{Lj})t + (\phi_i - \phi_j)]\} \quad (i,j = 1,2,3)\ .
\end{aligned}$$

It is instructive to derive the differential equation for the term ρ_{00}:

$$\begin{aligned}
\langle 0|\mathcal{H}_0\rho|0\rangle =&\langle 0|\rho\mathcal{H}_0|0\rangle = E_0\rho_{00}, \\
\langle 0|\mathcal{V}_I\rho|0\rangle =&\frac{\hbar}{2}[\Omega_{R1}\tilde{\rho}_{10} + \Omega_{R2}\tilde{\rho}_{20} + \Omega_{R3}\tilde{\rho}_{30}], \\
\langle 0|\rho\mathcal{V}_I|0\rangle =&0.
\end{aligned}$$

Therefore the complete form of the differential equation is

$$\mathrm{i}\dot{\rho}_{00} = \frac{1}{2}(\Omega^*_{R1}\tilde{\rho}_{01} + \Omega^*_{R2}\tilde{\rho}_{02} + \Omega^*_{R3}\tilde{\rho}_{03}) + \text{c.c.} + \Re_{00},$$

where $\Re \equiv \mathcal{R}\rho$. The full set of equations can be derived similarly and can be found in appendix B.

The relaxation term in the equations above in its most general form is given by

$$\begin{aligned}\Re_{ij} &= \sum_{i'j'} \mathcal{R}_{ij}^{i'j'} \rho_{i'j'},\\ \mathcal{R}_{ij}^{i'j'} &= -\mathrm{i}(1-\delta_{ij})\delta_{ii'}\delta_{jj'}\Gamma_{ij} + \mathrm{i}\delta_{ij}\delta_{i'j'}\Gamma_i^{i'}.\end{aligned}$$

$\Gamma_i^{i'}$ when $i' \neq i$ is the transition rate constant from state $|i'\rangle$ to state $|i\rangle$ due to dissipative processes and we have

$$\Gamma_i^{i'} = -\Gamma_{i'}^{i}.$$

Γ_i^i is the total decay rate from state $|i\rangle$. Therefore for a closed system we have

$$\Gamma_i^i = \sum_{i' \neq i} \Gamma_{i'}^{i}.$$

Γ_{ij} is the decoherence decay rate of state $|i\rangle$ to state $|j\rangle$ and is equal to

$$\Gamma_{ij} = \frac{1}{2} \sum_{k \neq i \text{ or } j} (\Gamma_k^i + \Gamma_k^j). \tag{7.4}$$

For the simple system considered here most of these relaxation terms are equal to zero. The $^{40}\mathrm{Ca}^+$ ion whilst in the $\mathrm{P}_{1/2}$ state can decay to the ground state $\mathrm{S}_{1/2}$ or the metastable states $\mathrm{D}_{3/2}$ with the branching ratio of 12 to 1. This gives the decay rates of $\gamma_{sp}/26$ for each metastable state and $12\gamma_{sp}/13$ for the ground state where γ_{sp} is the total spontaneous decay rate of level $|0\rangle$. The transition decay rates are therefore equal to

$$\begin{aligned}\Gamma_1^0 = -\Gamma_0^1 &= \frac{12}{13}\gamma_{sp}\\ \Gamma_2^0 = -\Gamma_0^2 &= \frac{1}{26}\gamma_{sp}\\ \Gamma_3^0 = -\Gamma_0^3 &= \frac{1}{26}\gamma_{sp}.\end{aligned}$$

The decays from the metastable states are ignored since the lifetimes of these states are much longer than the time to reach to a steady state. Because there are no transitions other than those from level $|0\rangle$ we have

$$\Gamma_2^1 = \Gamma_3^2 = \Gamma_3^1 = 0.$$

Calculating the decoherence rate from equation 7.4 gives

$$\Gamma_{10} = \Gamma_{20} = \Gamma_{30} = \frac{1}{2}\gamma_{sp},$$
$$\Gamma_{12} = \Gamma_{13} = \Gamma_{23} = 0.$$

7.3 Results

A Mathematica script was written to numerically calculate the evolution of the density matrix. The script can be found in appendix B. It was assumed that the ion is in the ground state $S_{1/2}$ when the laser cooling starts i.e. $\rho_{11}(t=0) = 1$. The steady state solution is achieved after many steps (~ 500) of time evolution. The required time for the evolution to reach the steady state is checked by looking at the change of populations in time.

The results show two-photon resonances and Stark-shifts. With sufficient laser power and appropriate laser detuning, a maximum population in the excited state $\rho_{00} = 0.25$ can be achieved. We will examine these results in detail but we note at the outset that the results show that, similar to the case of the RF trap with appropriate laser detuning, CPT *can* be avoided. This verifies our original assumption that CPT is not responsible for the low signal per ion experienced in the Penning trap Ca^+ ion laser cooling.

A typical result is shown in figure 7.7. In all results shown we assume that the spontaneous emission rate $\gamma_{sp} = 1$. The Rabi frequencies characterising the laser are $\Omega_{R1} = 2$, $\Omega_{R2} = \Omega_{R3} = 0.5$ in units of γ_{sp}.

The most important graph in figure 7.7 is the first one (upper left) which shows the population ρ_{00}. In all the graphs presented the detuning $\delta_{L3} = 0$. The population in state $|0\rangle$ is proportional to the scattering rate and therefore the amount of fluorescence detected. There are two stripes in this figure where the population is zero. The vertical stripe is due to the two-photon resonance between the ground state $|1\rangle$ and metastable state $|3\rangle$ when $\delta_{L1} = \delta_{L3}$. A similar vertical stripe, this time of high population can be seen in the graphs of ρ_{11} and ρ_{33}. The diagonal stripe in ρ_{00} graph is when $\delta_{L1} = \delta_{L2}$ and therefore a two-photon resonance occurs between ground state $|1\rangle$ and metastable state $|2\rangle$. Complementary diagonal stripes of high population can be seen in graphs of ρ_{11} and ρ_{22}.

The Rabi frequencies Ω_{L2} and Ω_{L3} characterise the intensity of the repumping lasers L2 and L3. With higher intensities the dip in the signal broadens. In other words higher intensity broadens the two-photon transition. It also

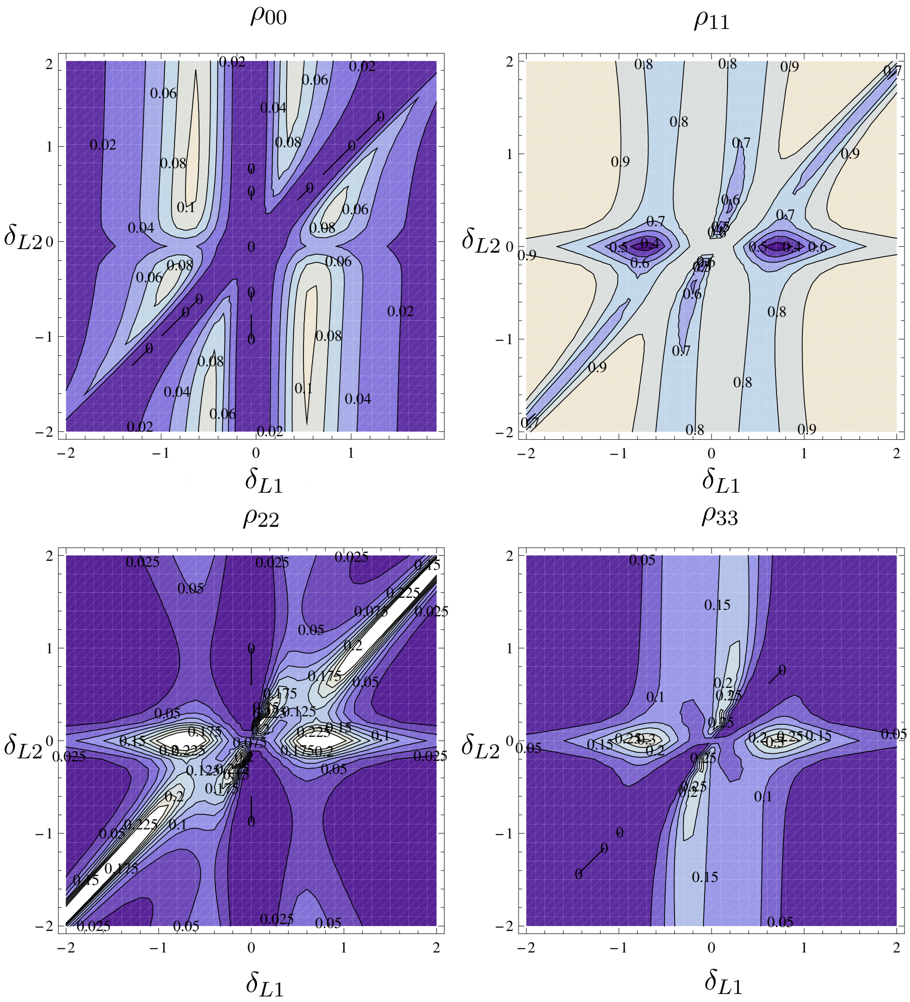

Figure 7.7. Typical contour plots for the population of the levels are shown. Lighter areas have higher population. The Rabi frequencies are $\Omega_{R1} = 2, \Omega_{R2} = \Omega_{R3} = 0.5$ and $\delta_{L3} = 0$.

shifts the resonance peaks of the transitions due to the a.c. Stark shift caused by the external electric field of the IR lasers. This can be clearly seen in figure 7.8. Over-saturating the IR transitions does not broaden the cooling transition significantly; it only shifts the one-photon resonance frequencies.

Increasing the intensity of the L1 laser (the UV laser) broadens the cooling transition as can be seen in figure 7.9. The final temperature of the laser cooling is proportional to the linewidth of the cooling transition, so broadening the transition decreases the cooling efficiency. However it can be seen that the population increases with more UV power. Therefore there is a trade-off between the transition narrowing with low laser intensities and population

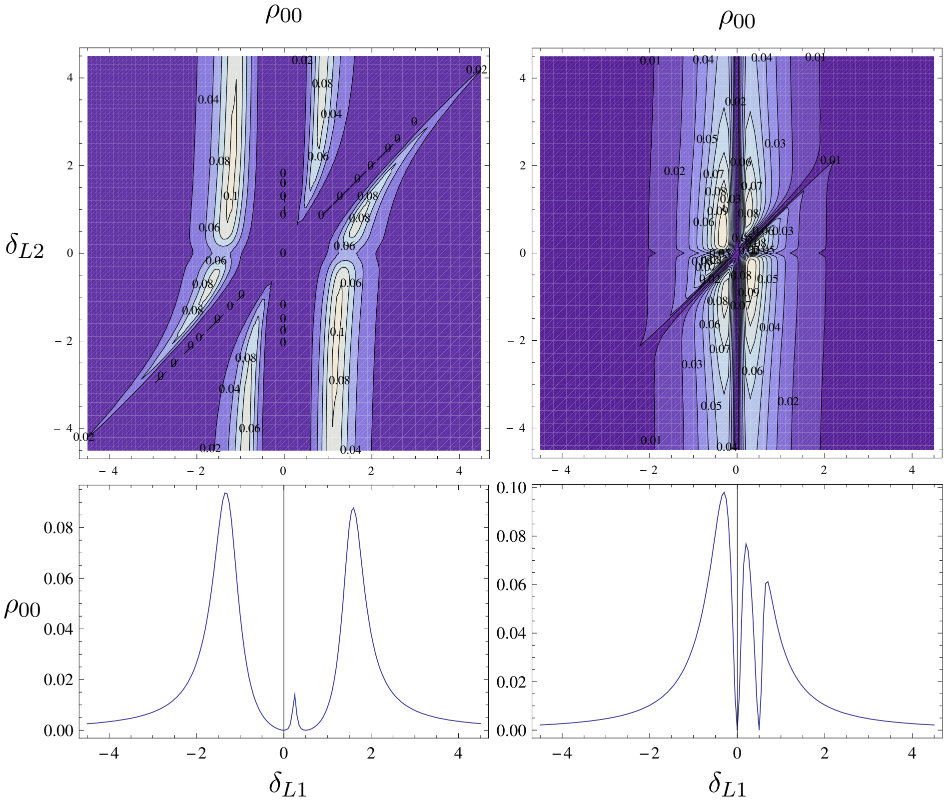

Figure 7.8. IR power broadening. Contour plot for population in the excited state $|0\rangle$ is shown for two conditions: *Left:* $\Omega_{L2} = \Omega_{L3} = 1$ and *Right:* $\Omega_{L2} = \Omega_{L3} = 0.2$. In both cases $\delta_{L3} = 0$ and $\Omega_{L1} = 0.2$. Lighter areas have higher population. The two plots below show a cut through the contour plots at $\delta_{L2} = 0.5$.

increase with high laser intensities.

For laser cooling, the UV lasers are frequency detuned to the red. Naively one would think that the repumping lasers should be in resonance. However that makes a Λ system between the repumping lasers and stops the repumping process. Therefore at least one of the IR lasers needs to be detuned from the transition. With sufficient power for both IR and UV lasers and appropriate detuning, 0.25 population in the excited state $P_{1/2}$ can be achieved. This confirms that there is nothing fundamentally wrong with the laser cooling scheme.

In practice the blue laser is set off-resonance for cooling and the detuning of the IR lasers is adjusted for maximum signal. This process should home in on the best signal, not some local maximum.

The simulations show that maximum population in state $|0\rangle$ occurs when

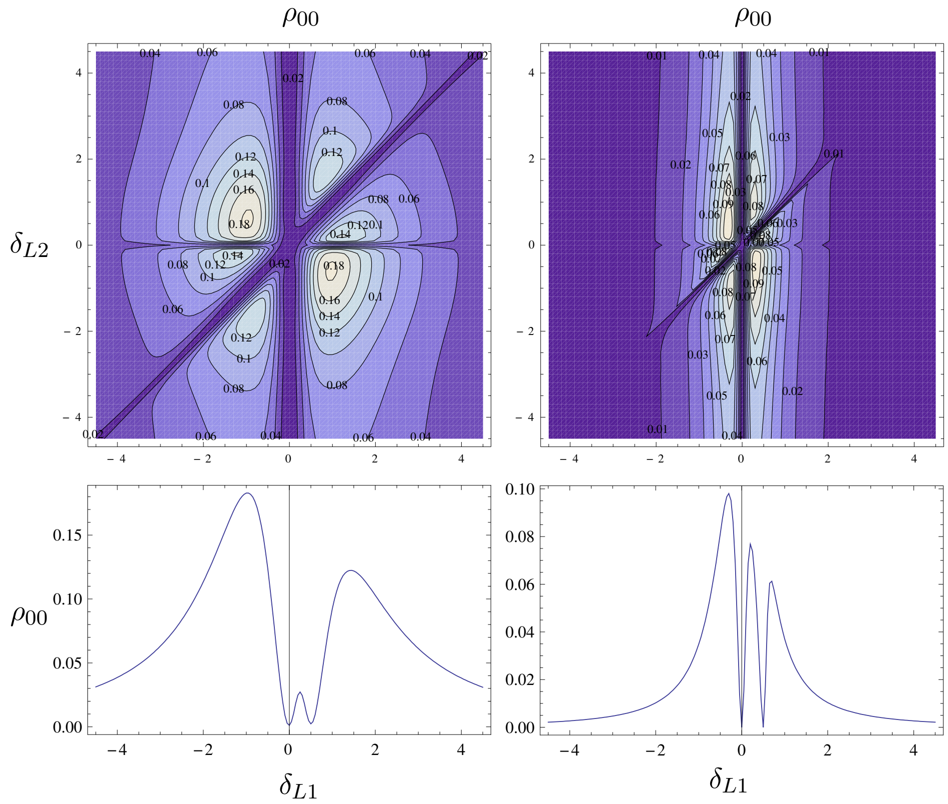

Figure 7.9. UV power broadening. Contour plots for population in the excited state $|0\rangle$ is shown for two conditions: *Left:* $\Omega_{L1} = 0.8$ and *Right:* $\Omega_{L1} = 0.2$. Lighter areas have higher populations. In both cases $\delta_{L3} = 0$ and $\Omega_{L2} = \Omega_{L3} = 0.2$. The two plots below show a cut through the contour plots at $\delta_{L2} = 0.5$.

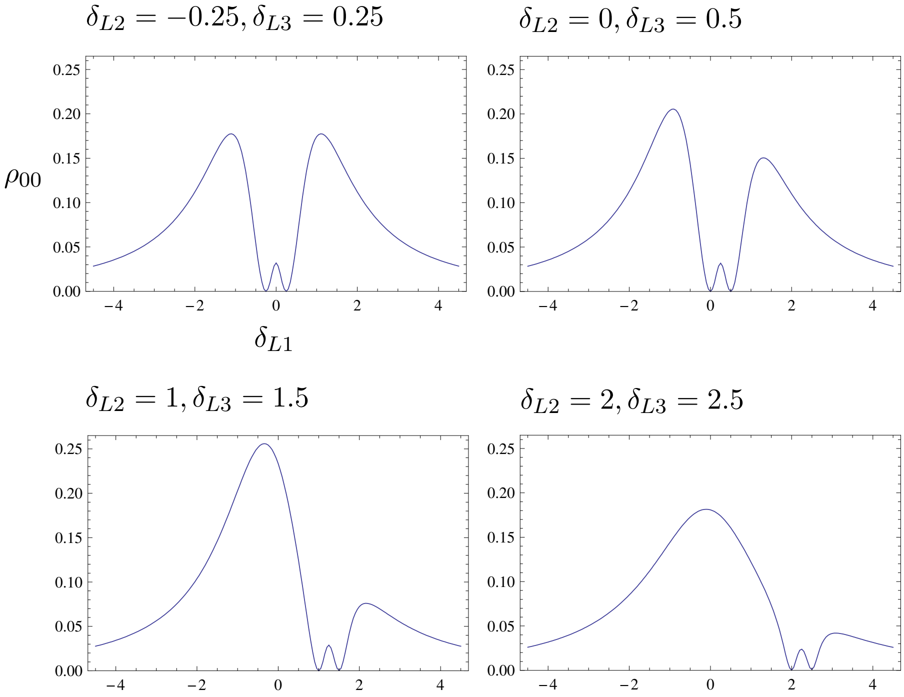

Figure 7.10. Population in the excited state $|0\rangle$ is shown for four different conditions. In all cases $\Omega_{L1} = 0.75$ and $\Omega_{L2} = \Omega_{L3} = 0.35$. The maximum population in the excited state occurs when $\delta_{L2} = 1, \delta_{L3} = 1.5$ ie. both IR lasers are detuned to the blue. Although the maximum population can be achieved using these parameters, the transition is significantly broadened .

$\delta_{L2} = 1, \delta_{L3} = 1.5$ (see figure 7.10). The corresponding Rabi frequencies are $\Omega_{R1} = 0.75$ for the UV laser and $\Omega_{L2} = \Omega_{L3} = 0.35$ for the IR lasers. The rather large Rabi frequency for the UV laser broadens the transition as can be seen by comparing figure 7.10 with previous ones (similar to lower left graph in 7.9).

7.4 Estimate for Laser Power

It is instructive to calculate the power required for the transitions in the specific laser cooling scheme that we use. The irradiance I is related to the electric field amplitude $\mathcal{E}$ by

$$I = \frac{1}{2} c \epsilon_0 \mathcal{E}^2. \tag{7.5}$$

As shown before, the Rabi frequency is related to the electric field by

$$\omega_R = \frac{\hbar \mathcal{E}}{\mu}, \tag{7.6}$$

where μ is the electric dipole moment of the specific transition which is addressed by the laser. From a quantum electrodynamic treatment of spontaneous emission one can show [107] that electric dipole moment is related to Einstein's A coefficient by

$$A_{ij} = \frac{2e^2\omega_{21}^3}{3\epsilon_0 hc^3}\mu^2, \tag{7.7}$$

where ω_{21} is the angular frequency of the transition and ϵ_0 is the permittivity of free space. Substituting equations 7.7 and 7.6 into equation 7.5 gives rise to

$$I = \frac{2\pi ch}{3A\lambda^3}\omega_R^2,$$

where λ is the transition wavelength.

For the transition $S_{1/2} \rightarrow P_{1/2}$ at 396.8nm in Ca^+ ion the Einstein A coefficient is $\sim 1.4 \times 10^8 s^{-1}$. For the special case where the Rabi frequency $\omega_R = \gamma_{sp}$ i.e. when the transition is saturated, the irradiance is equal to

$$I = 0.9 \text{mW/mm}^2.$$

Supposing the beam waist is 50μm at the centre of the trap the required laser power is $\sim 7\mu$W. However since the beam position needs to be offset to a side as we saw in chapter 4 and also that the laser frequency is detuned by half the natural linewidth of the cooling transition, four times as much power is needed.

Similar calculations show that much smaller power ($\sim 1/8$) is needed for the repumping lasers in order to saturate the transition. In the experiment it is found that we are saturating the UV transition but using only the sidebands is not enough to saturate the IR transitions (see chapter 8).

7.5 Conclusion

We have successfully modelled the coherent population trapping in a 4-level system of Ca^+ ion. We also demonstrated that the ideal 0.25 population in the excited state can be achieved despite the large number of levels and lasers.

From the results shown here we found no evidence that coherent population trapping could be a fundamental problem in laser cooling a single Ca^+ ion in high magnetic fields. However due to the number of the transitions involved, the number of situations in which CPT occurs is higher than for the RF trap. This makes the laser cooling of Ca^+ ions in the Penning trap a more challenging but still tractable problem. One should remember that there are two four-level systems which are connected incoherently by a spontaneous emission; if one system is not efficiently addressed either by having incorrect laser frequency detunings or insufficient laser power, the population is optically pumped into it and fluorescence reduces even though the other system is efficiently addressed. Therefore all six lasers need to be properly tuned to achieve optimum laser cooling. This is indeed more rigorous than a two-laser cooling scheme in RF traps.

CHAPTER 8

Observation of Individual Ions in the Penning Trap

8.1 Introduction

As mentioned in chapter 5, in order to address the four hyperfine sublevels of the $D_{3/2}$ metastable state two lasers with sidebands imposed by direct modulation of the current at ∼900MHz were used [77]. The modulation frequency corresponded to the half of the splitting between the adjacent σ components of the transitions. The four sidebands were addressing the $S_{1/2} \rightarrow D_{3/2}$ transitions at 866nm. The two lasers were combined with the two blue lasers using beam splitters and were sent to the trap. Although fluorescence was successfully observed, the laser cooling was not efficient enough to detect a single ion.

With the split-ring trap it was possible to run it as both a Penning trap and a RF trap. With the trap in RF mode we were able to trap a single ion and see clear quantum jumps. The signal count rate per ion was ∼ 7kHz. The lowest number of ions we were able to successfully laser cool in Penning mode was ∼5 ions with a maximum signal rate per ion of ∼ 500 counts per second [77]. This number was estimated by looking at the statistics of the fluorescence signal rather than directly observing discrete quantum jumps [73]. Comparing this with the RF trap results showed that the signal per ion was less than a tenth of that observed when the same trap was run as an RF trap. With the signal per ion of ∼500 counts per second, the quantum jumps are lost in the Poisson noise. Due to this small signal per ion we were not able to resolve quantum jumps in the fluorescence signal.

8.2 Possible Reasons for the Low Signal Rate

There were a number of issues that we identified as possible causes of this low count rate. The first issue was the presence of contaminant ions in the trap such as CaO^+, NO^+, $^{44}Ca^+$ etc., which could reduce the cooling efficiency of the Ca^+ ions. Since the ionisation is by thermionic emission, ions are not selectively loaded. In the laser cooling process other ionised species are also sympathetically cooled by electrostatic interactions with laser cooled Ca^+ ions. Previous experiments in our group show that this significantly increases the final temperature of the cloud [108]. We could imagine a situation in which the amount of contaminants was not reducing proportionately as we were loading smaller clouds. It should be noted that this is less of a problem in the RF trap because the trapping frequency and bias for stable trapping depends on the mass of the ion. Therefore the trap can in principle be made unstable for any ions other than Ca^+. Also when attempting to load a single ion in the RF trap, the RF heating may help to remove contaminant ions that are not laser cooled. However this is not necessarily the case in the Penning trap because there is no RF heating and the Penning trap is stable for a broad range of mass at any given set of trapping parameters (see section 4.3). To investigate this, a mass spectrometry technique was used by the previous PhD student in our group (R. Hendricks). However, no evidence of high proportions of contaminants was found in the acquired spectrum for relatively small clouds [77].

The second possible problem was coherent population trapping (CPT) that might reduce the laser cooling efficiency when trapping lower numbers of ions. We investigated this issue in detail in chapter 7 and found that although there were regions of the parameter space where CPT could certainly happen, according to the calculations one should be able to avoid it by appropriate laser detunings. Nevertheless, the most obvious method for ensuring that there was no CPT in the more complex 8-level system using sidebands was to actually use individual lasers for each transition at 866nm.

The other possible issue was insufficient laser power for the IR repumping lasers. The current modulation frequency of the diode lasers was already very high to impose sidebands bigger than ∼10% of the carrier without causing mode hopping and laser instabilities. This seems to be more of a problem for higher modulation frequencies. Much larger sidebands can be imposed with modulation frequencies in the 10s of MHz range. In fact $\sim$ 900MHz was the maximum frequency that we could go to and stable operation was only possible with two of the four laser diodes that we had available. The laser power in the center of the trap was 0.5mW. This means that at most there was only 50μW power for each sideband-driven transition, which was 1/10 of what we had when running the trap as RF trap. According to the calculations this amount of power should be enough for single ion laser cooling (see chapter 7) but we decided to try using individual lasers for addressing each transition in order to be sure this was not a problem.

Another issue was the magnetic field stability. The magnetic field with the conventional electromagnet is quoted to be stable to 1 part in 10^4 per hour. However in practice the magnet tends to drift slowly and sometimes the electronics of the magnet power supply, which stabilises the current, unlocks and the current needs to be readjusted. The group also has a superconducting magnet which is stable to 1 part in 10^7 per hour. It should be noted that the geometry of this magnet is completely different from that of the conventional magnets. It is a long solenoid with only axial access and as a result traps used in one magnet cannot be used in the other. The lasers were transferred to the trap by optical fibres and were combined using optics mounted on the small breadboard onto which the trap was mounted. Before I joined the group a cylindrical trap had already been built to be used in the superconducting magnet. The trap had already been successfully used with Mg^+ ions. The ring was split into four segments to allow axialisation. Getting this trap to work turned out to be difficult. Due to its cylindrical geometry the trap was not suitable for running as a RF trap. Therefore, one had to make sure that the frequencies of all 4 lasers were set up correctly in the beginning. Furthermore the filament needed to be located along the axis and since the PMT was also along the axis the background light whilst loading was a bigger issue. The background level was increasing so high ($\sim$ 1000 counts/10ms) that one had to wait until the filament was turned off to detect fluorescence. Therefore the lasers' frequencies also needed to be correctly detuned otherwise the fluorescence would disappear very quickly after loading.

To make sure that the frequencies of the lasers were properly detuned the split-ring trap (conventional magnet trap) was simultaneously run with the cylindrical trap and the lasers were shared between the two traps using flipper mirrors. Since the experiment was unsuccessful at seeing fluorescence, the trap was disassembled and rebuilt with the split ring replaced with one monolithic ring instead to make sure the ring segments' alignment was not an issue. Unfortunately the trap did not work. The magnification of the original system was very large and this made alignment difficult since an image of the trap electrodes would not be formed at the top of the bore of the superconducting magnet. The optics collimating the fluorescence, which was specially designed to minimise aberrations using 4 lenses, was replaced with a single lens instead in order to have a lower magnification. However after several unsuccessful attempts over a period of 4 months the trap was put to one side and we concentrated instead on trying to see single ion in the conventional magnet trap acting as a Penning trap.

8.3 Improvements

Work on the split-ring trap in the conventional magnet was resumed but this time with 4 lasers addressing the 866nm repumper transitions. Occasionally the ion is excited to the $P_{3/2}$ state by ASE from the IR diode lasers at 850nm driving the $D_{3/2} \rightarrow P_{3/2}$ transition as well as ASE from the UV diode lasers by driving the $S_{1/2} \rightarrow P_{3/2}$ transition at 393nm. The ion then can decay to the $D_{5/2}$ metastable state. This effect results in quantum jumps in the fluorescence of a single ion. A new laser at 854nm was bought to address the $D_{5/2} \rightarrow P_{3/2}$ transition to repump the hyperfine sublevels of the $D_{5/2}$ metastable state. Addressing this transition increases the average signal count rate for a cloud of ions by 20%. The laser diode used was a 130mW Eagleyard EYP-RWL-0850 at 850nm. As explained below we found that only one laser is enough to repump all 6 sublevels of the $D_{5/2}$ state and it is not as sensitive in frequency as the 866nm lasers are. In fact the rate of the quantum jumps is on the order of 1Hz. Considering the linewidth of $\sim$ 24MHz for the $D_{5/2} \rightarrow P_{3/2}$ transition and a transition rate of $\sim 9.9 \times 10^6 \mathrm{s}^{-1}$ and laser detuning of 17GHz the transition rate is $\sim 75\mathrm{s}^{-1}$ at saturation power. The ASE in the IR lasers also helps to repump the $D_{5/2}$ states. As we will see in section 8.4, the average time for quantum jumps in our experiment is much shorter ($\sim$ 70ms) than the lifetime of the $D_{5/2}$ state ($\sim$ 1s). Therefore the off-resonance addressing is enough for repumping the $D_{5/2}$ metastable state considering that the transition is over saturated too.

Abandoning the sidebands gave the freedom to set the frequency of each laser independently during the laser cooling. This was indeed very helpful especially in loading small numbers of ions as the frequency of the IR lasers became very sensitive. With the sidebands it was not possible to easily change the modulation frequency without inducing laser instabilities and mode hopping.

Since there were 7 lasers altogether involved in the laser cooling, the relative alignment of the lasers in the centre of the trap became an important issue. For that reason the lasers were coupled to optical fibres. Two single-mode fibres were used: one for the IR to be used for all 5 IR lasers and one for the UV to which the two UV lasers were coupled. This turned out to be a major improvement in practice since one needed to adjust the overlap of only two beams rather than 7 whilst the laser cooling was performed. The fibres were custom-made by Oz Optics. Both were angled with FC/PC connectors in order to minimise the back reflection of the input light off the front facet of the fibre. The UV fibre had a 3μm core with an operating wavelength of 320-400nm. The IR fibre had a core of 5μm operating at a wavelength range of 810-980nm. Both fibres had focus-adjustable input couplers.

The UV fibre initially worked well but after a few months the coupling efficiency degraded over a two-day period and eventually decreased to less than 1%. In order to re-polish the fibre it was sent back to Oz Optics. However the coupling efficiency did not improve after re-polishing the fibre tips and it was never found what the real cause was. It seems to be a degradation of the core which increases rapidly once it begins. This process halted the experiment for a few weeks. The Oz Optics fibre was abandoned and we decided to build the fibre ourselves and save time in case it broke again. A fibre polishing kit was bought from Thorlabs and a polarisation maintaining fibre (Nufern PM-S350-HP) was manually polished.

The coupling efficiency of the lasers to the fibres varies with different diodes and is about 30%-40% at maximum. This is mainly due to the low quality of the laser beam intensity profile. We get $\sim 250\mu$W power for each blue laser and $\sim 400\mu$W for each IR laser addressing the $D_{3/2} \rightarrow P_{1/2}$ transition at the last mirror before the trap. Neutral density filters are used to equalise the laser powers. The repumping laser for the $D_{5/2}$ metastable state has a power of $\sim 100\mu$W after the fibre. The relatively low laser powers are because half the power is used for monitoring and locking of the lasers. A high proportion of the rest is lost due to the use of several beam splitters in combining the lasers, mode matching to the fibres and reflections from other surfaces such as mirrors and lenses.

The fibres were originally implemented to take lasers to the superconducting magnet trap. However using fibres has another important advantage. The single-mode fibres only allow one spatial mode to propagate which is the TEM_{00} mode. This is very useful since diode lasers have a very poor beam quality. The fibre acts as a spatial filter for the UV lasers and reduces the laser scattering inside the trap. As a result the background noise from the scattered 397nm light is reduced significantly by using the fibres.

The fibres can also be used in removing amplified spontaneous emission (ASE) light from the laser. This is very important especially for the future QIP experiments that we intend to do, because ASE causes quantum jumps and destroys the coherence of the qubit. At the moment we do not mind because quantum jumps are in fact useful for determining the number of ions in the trap but we will have to avoid them later. For this purpose a prism can be used before coupling the light to the fibre. The core of the fibre is very small ($\sim 3\mu$m) and a short distance is enough for the ASE light to be displaced from the core of the fibre. For instance the 393nm ASE has a wavelength difference of $\Delta\lambda = 4$nm with the 397nm input light. For a 45° prism of BK7 glass and an incident angle of 45° this gives an offset angle of $\sim 4 \times 10^{-4}$ radian. A distance of 25cm between the fibre tip and the prism is enough to displace the light at 393nm by 100μm and stop it from coupling to the fibre.

In order to stabilise the magnet power supply, I designed a circuit to actively stabilise the magnetic field close to the trap. A Hall effect integrated circuit (Honeywell SS94A2D), which gives an output voltage proportional to the magnetic field intensity, was placed in close proximity to the trap. The probe sensitivity is 1mV/gauss and it is temperature compensated. The probe output is fed to a feedback circuit shown in appendix D. The circuit is very simple and has only one stage: an operational amplifier acting as an integrator. The magnet power supply can be controlled externally by a voltage between 0-1V. The output of the circuit is fed to the external input of the magnet power supply. The magnet then can be locked to the Hall probe. The electromagnet current is stabilised to 15 ± 0.1A when it is locked.

The trap is run with a filament current of 4.5-4.7A and an oven current of 1.5-1.6A for the first time in the day. The oven is left on for two minutes and then the filament is turned on for a time long enough to see some fluorescence appear (∼20-25s). After that all parameters, as described in section 5.11, such as the laser frequency, position, alignment, focus, etc. are adjusted. This process is then repeated with a smaller cloud. To get a smaller cloud we usually lower the filament current in steps of 0.05-0.1A. It is important to go in steps because the laser parameters such as the frequency and position of the beam depend on the size of the cloud. As we get to smaller clouds, improved laser cooling means that the level of fluorescence becomes very sensitive to these parameters. If we attempt to run the trap with small clouds straight away, either the laser parameters are so far off their optimum values that we barely see any signal, or due to the heating of the ions, the fluorescence fades away so quickly that we do not have enough time to adjust them.

When trapping large numbers of ions, due to the small trapping volume, the interactions between the ions are strong and frequent. Provided we cool a small number of them the whole cloud can be cooled to the centre of the trap very quickly. However when trapping small numbers of ions they are more independent and it becomes possible for some ions to spend a relatively long time in large orbits without actually interacting with the other ions or with the laser beam. For this reason, when trapping a single ion, it usually takes a few minutes for the fluorescence to appear. This is what we actually see when trapping single ions; we have to wait 1-5min after stopping the filament to see a signal appear. This is another reason why stepwise optimisation is necessary to make sure that all parameters are set correctly.

8.4 Results

The various improvements to the setup described above have enabled us to trap single Ca^+ ions in the Penning trap. At a relatively low oven current of 1.5A and a very low filament current of ~ 3.55A we can consistently trap single ions. The fluorescence signal of a single ion in the split-ring Penning trap is shown in figure 8.1. The signal rate is at maximum $\sim$ 3500 counts per second for a single ion. When there is only one ion in the trap the ion is sometimes shelved in the $\mathrm{D}_{5/2}$ state and stops fluorescing. This is why we see steps in the fluorescence. 'One-step' quantum jumps indicate that the there is only one ion in the trap. If there are two ions in the trap, at some times both ions are fluorescing or only one ion is fluorescing and the other is shelved (in which case the signal is half) or both are shelved (in which case the signal drops to the background level). This is why we see two steps in the fluorescence of two ions (see figure 8.3). This trace is acquired by blocking the 854nm laser, which repumps the $\mathrm{D}_{5/2}$ metastable state and therefore allows

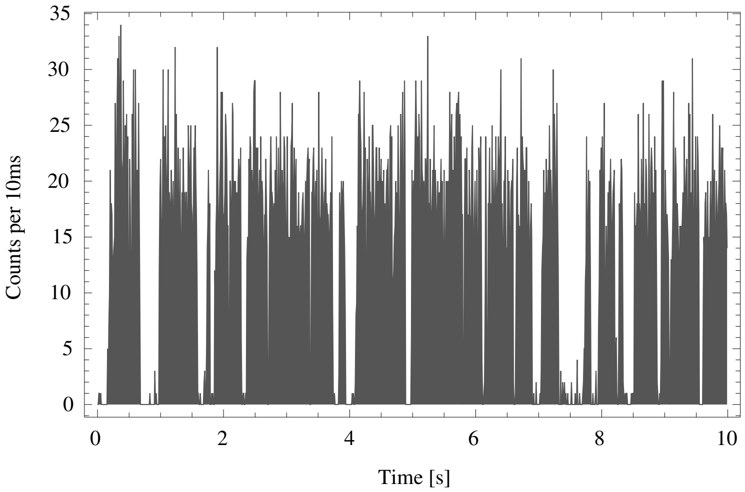

Figure 8.1. A typical fluorescence trace of a single Ca^+ ion in the split-ring Penning trap. The IR lasers have 400μW power each and the UV laser power is 65μW for each laser. The signal to background ratio is ∼ 32.

quantum jumps to happen. A histogram of the signal count rate for a time bin of 10ms is shown in figure 8.2(a). In this histogram the count rate bin is 10ms. As it can be seen there are two Poissonian peaks in the histogram: The first is due to the background and the second is due to the signal. The FWHM of the peaks corresponds to the noise [109].

The probability of the ion staying in the $P_{5/2}$ excited state is proportional to $\exp(-t/\tau_r)$, where τ_r is the radiative decay time (lifetime) of the state and is equal to

$$\tau_r = \frac{1}{A}.$$

Here A is the Einstein's A coefficient. The lifetime of the $D_{5/2}$ is 1168 ± 7ms [110]. In our experiment however the lifetime is much shorter due to the ASE in the IR lasers which drives the $D_{5/2} \rightarrow P_{3/2}$ transition at 854nm. A histogram of the quantum jumps time duration is shown in figure 8.2(b). The number of quantum jumps is ∼ 450 taken in one run. The measured lifetime is 70 ± 5ms.

To improve the background one needs to use an aperture (pinhole) before the PMT. The amount of background is proportional to the surface area of the aperture, therefore the smaller the aperture is the lower the background becomes. For an imaging system with a magnification of 1 the aperture before the PMT can be in principle as small as the orbit of a laser cooled ion.

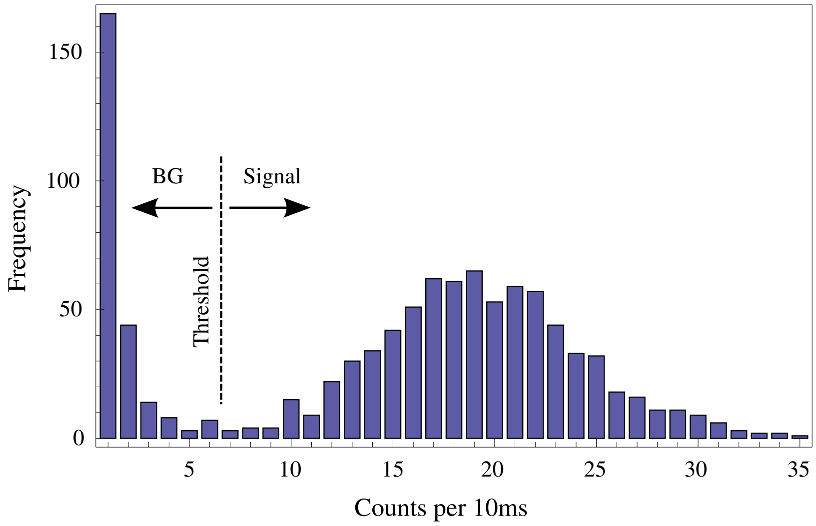

(a)

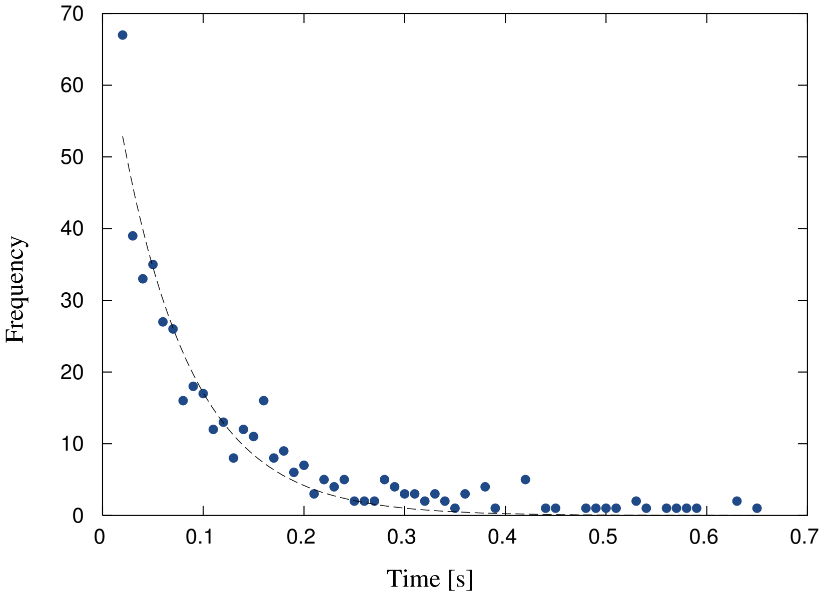

(b)

Figure 8.2. a) The histogram of the signal in figure 8.1 is shown. b) Histogram of time length of ~ 450 quantum jumps for a single ion in the Penning trap. The time bin is 10ms.

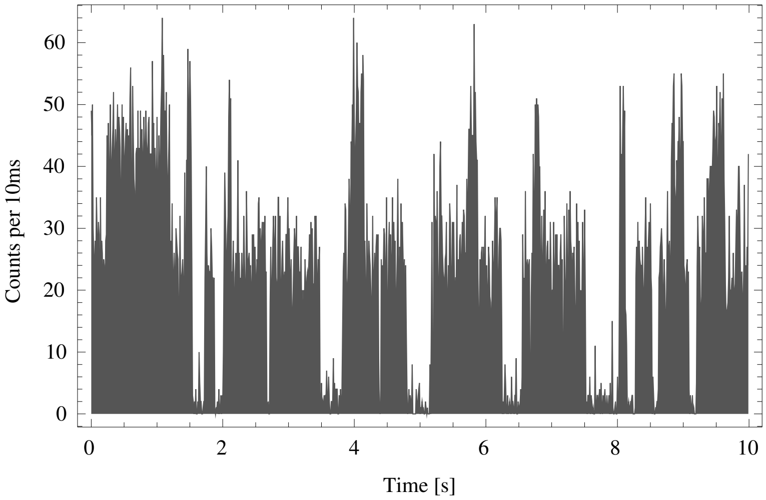

Figure 8.3. Quantum jumps for two Ca^+ ions in the Penning trap.

Considering our imaging magnification which is ~ 1.7 the minimum size of the aperture is $\sim 50\mu$m. In practice however it has to be bigger than this due to the aberrations of the imaging system. In the experiment the pinhole before the PMT is 200μm. It is a thin steel sheet with a central drilled hole. This pinhole improved the signal to background significantly compared with the 500μm commercial iris that we were using previously.

We can hold single ions for a few minutes without laser adjustments. But the drift of a laser even when it is locked to a cavity is noticeable by a slow drop in signal. We believe that the drift of the cavities is due to the pressure changes inside the cavity. The refractive index of the light changes when the pressure changes. This effect, although it seems to be negligible for big clouds, becomes a big issue when trapping single ions.

In a conventional RF trap, when a pair of ions is trapped, because of the micromotion, after a while one ion would be heated and could eventually be thrown out of the trap. However this is not the case in the Penning trap because there is no micromotion. Also there seems to be no good way of deliberately removing just one ion. Therefore one has to trap only a single ion and this is a matter of luck if one is using thermionic emission from a hot filament as a source for ionisation. This could be an easier task using photoionisation. A trace of fluorescence for two trapped ions is shown in figure 8.3. The signal level is almost twice of that of a single ion and two-step quantum jumps show that there are two ions in the trap.

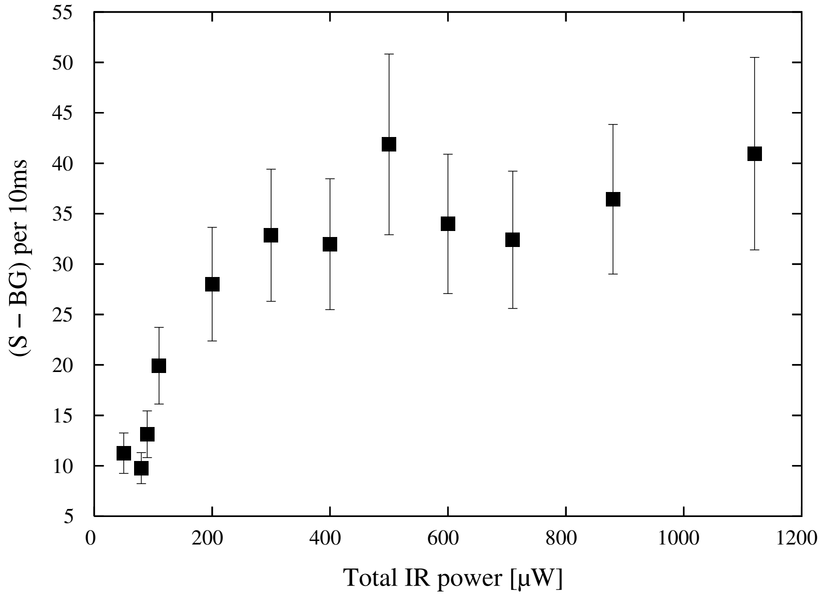

(a)

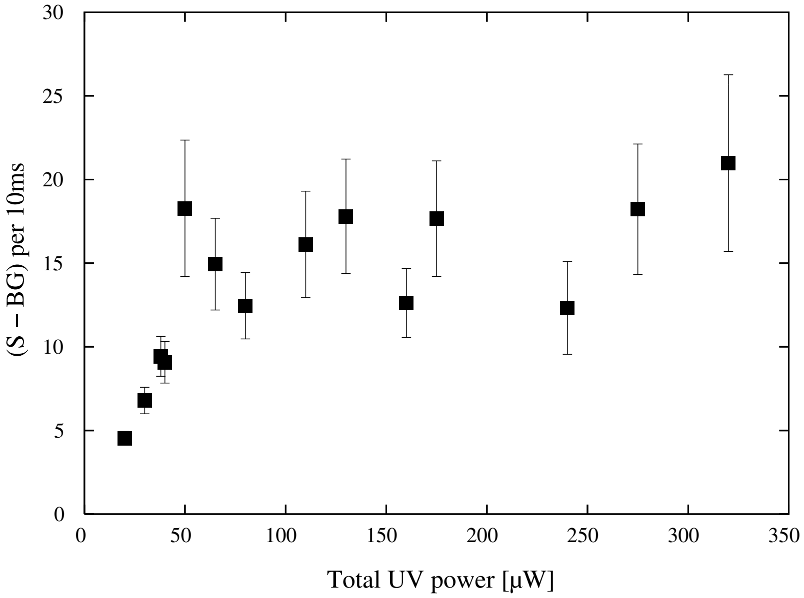

(b)

Figure 8.4. a) The saturation curve of the IR transition. The UV power is ~ 100μW for each laser ('S-BG' stands for signal minus background level). b) The saturation curve for the UV transition. The IR power is ~ 400μW for each laser. Both curves are for a single ion in the Penning trap. The lower signal count rate in the UV curve is because a 200μm pinhole was used whereas for the IR curve a 500μm was used.

We thought we were saturating the IR transitions with sidebands but clearly we were not. In order to measure the power that is needed to saturate the IR transitions, a single ion was trapped and the signal minus background in the quantum jumps was measured for various IR laser powers. A linear variable neutral density filter mounted on a translation stage before the input coupler of the fibre was used to adjust the laser power without changing the laser direction. The resulting graph is shown in figure 8.4(a). The results show that altogether 300μW i.e. 75μW power for each laser is enough to saturate the 866nm transition. This is larger than the $\sim 50\mu$W power that the sidebands were providing. Nevertheless it should not have been stopping us from seeing single ions and quantum jumps.

A similar experiment was performed to estimate the UV power required to saturate the 397nm transition in the Penning trap. The resulting graph for a single ion is shown in figure 8.4(b). The required laser power for saturation was $\sim 25\mu$W for each laser. This result is close to an estimated laser power of 28μW derived in section 7.4. The maximum count rate per ion achieved in the Penning trap was $\sim 4,200$ counts per second. In the same trap when it was run as an RF trap, a maximum count rate per ion of $\sim 12,000$ counts per second was achieved. The fact that we obtained a count rate of a third of that measured in the RF trap could be justified considering that in the Penning trap the laser beam needs to be offset from the centre of the trap in order to cool the magnetron motion.

To increase the signal-to-background (S/BG) in the quantum jumps, the iris before the PMT was replaced with a 200μm pinhole, as discussed above. A graph of S/BG in the Penning trap is shown for a single ion in figure 8.5(a). A maximum S/BG of 62 ± 14 was achieved at the saturation power which was 25μW power for each UV laser. This result can be compared with previous results in the RF experiment where a 0.5mm commercial iris was used (see figure 8.5(b)). In the RF experiment the maximum S/BG achieved was 7.8 ± 1.5 at a saturation power of 75μW. The higher power for the UV laser in the RF experiment was because we were not using the fibres for spatial filtering. Because the beam quality of the diode lasers were rather poor, a part of the light was not being properly focused through the final lens before the trap and therefore was wasted.

As mentioned in chapter 5 the ion can be imaged by a flipper mirror to an iCCD camera. Figure 8.6 shows a single ion imaged by the iCCD camera. The image is an accumulation of 10 shots with durations of 0.1s. A horizontal

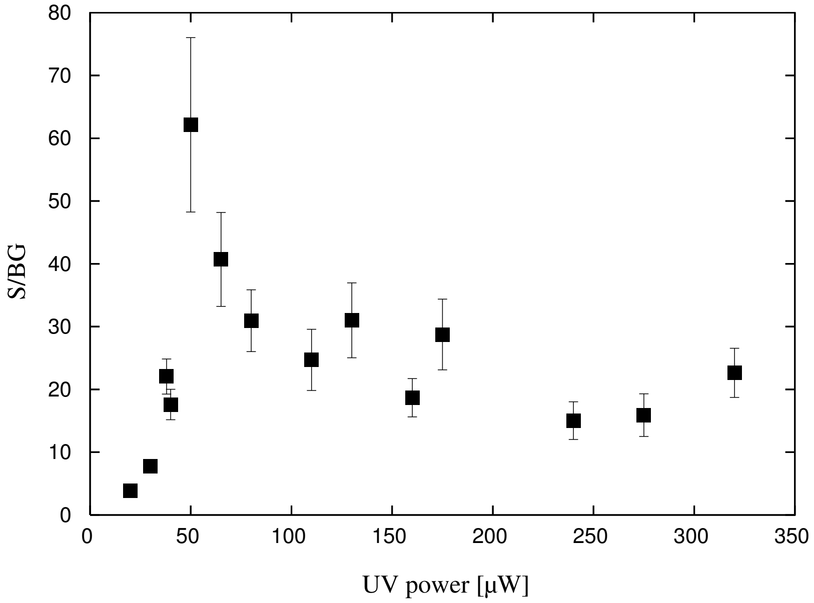

(a)

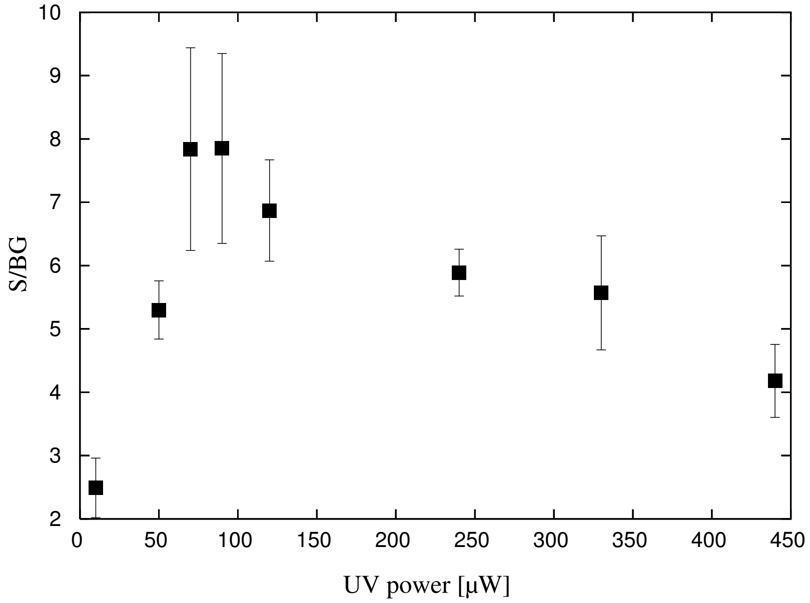

(b)

Figure 8.5. a) Signal to background vs. the UV power for a single ion in the Penning trap. The UV power here is the total power of both UV lasers. b) Signal to background vs. the UV power for a single ion in the RF trap.

intensity profile of the image is shown in figure 8.7. The width of the peak is ~ 2 pixels. The resolution of the iCCD camera measured before shipping was quoted to be 22μm. Considering the fact that the iCCD camera has a pixel size of $(13\mu\text{m})^2$, the single ion image width is very close to the resolution of the camera. In other words the upper bound of the orbit size of the ion is the resolution of the camera. Currently work is in progress to increase the magnification of the optical imaging in order to obtain better image resolution.

Although the imaging system is not specially designed for large magnification (~ 1.7), we managed to resolve two ions on the iCCD camera with a trap bias of 3.4V. This voltage bias is smaller than the usual 4V bias that we normally use. When the voltage bias is high it squeezes a cloud of ions into a pancake shape and when it is low the cloud goes into a cigar shape. To get a crystal of two ions along the axis the bias voltage needs to be low otherwise they collapse into the radial plane. Because the trap is not perfectly symmetrical, changing the trap bias voltage shifts the minimum of the trap. Therefore the beam position needs to be shifted too. This is a very difficult task especially in the Penning trap because if the ion moves to the heating side of the beam the ions will be instantly lost. Figure 8.9 shows two ions and the corresponding surface plot. The ions are along the axis of the trap.

We can estimate the distance between the two ions simply by rewriting equation 4.1 for an ion along the z axis giving the potential

$$\phi = \frac{2Uz^2}{2z_0^2 + r_0^2}.$$

Supposing the two ions are separated by a distance l the force F_ϕ due to the electric potential of the trap on an ion is expressed by

$$\begin{aligned} F_\phi &= Ee \\ &= q\left.\frac{\partial\phi}{\partial z}\right|_{z=l/2} \\ &= \frac{2Ule}{2z_0^2 + r_0^2}. \end{aligned}$$

When the ion is in equilibrium, the Coulomb force between the two ions

$$F_q = \frac{e^2}{4\pi\epsilon_0 l^2}$$

is equal to the force F_ϕ. The distance l between the two ions can then be derived as

$$l = \left(\frac{e(2z_0^2 + r_0^2)}{8\pi\epsilon_0 U} \right)^{\frac{1}{3}} .$$

Substituting the trap parameters of $z_0 = 3.5$mm and $r_0 = 5$mm and trap bias of $U = 3.4$V gives an ions separation of $\sim 22\mu$m. Given our imaging system magnification ($\sim$ 1.7) and camera's resolution ($\sim$ 2 pixels) this corresponds to $\sim$ 3 pixels on the iCCD camera which is very close to the ion separation acquired from iCCD images (see figure 8.8 and 8.9).

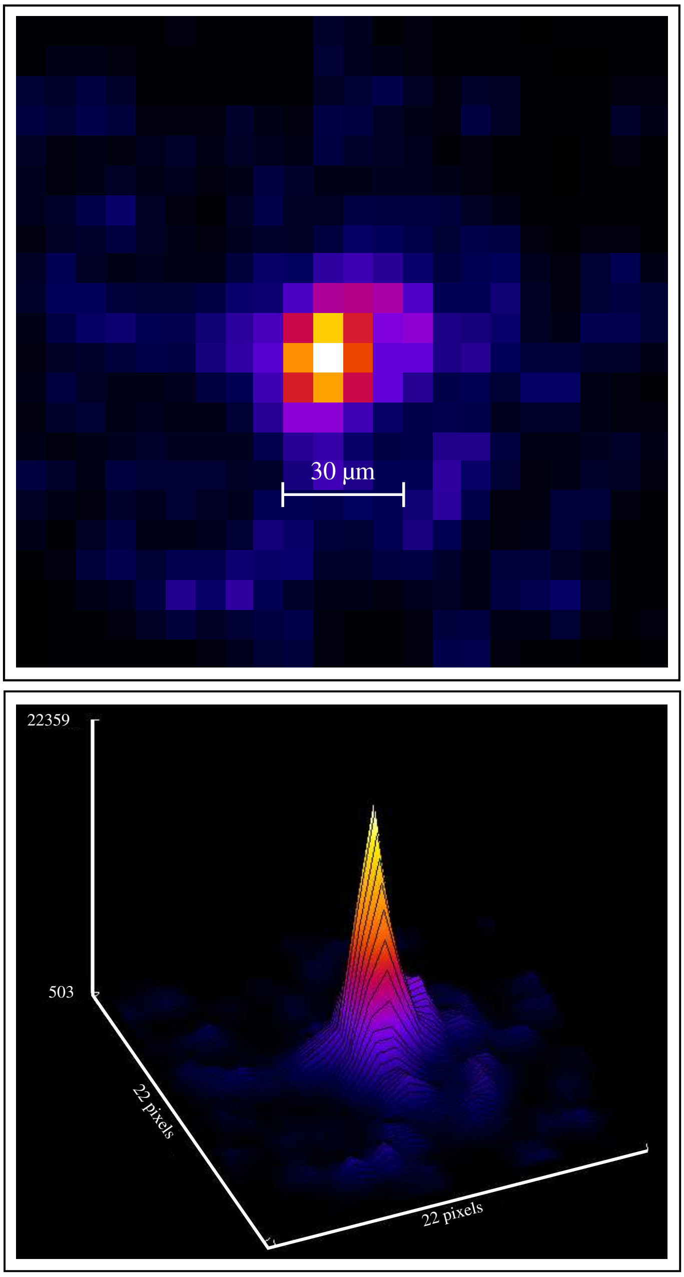

Figure 8.6. *Top:* The iCCD image of a single ion in the Penning trap. This picture is taken by the accumulation of 10 shots of 0.1s duration. *Bottom:* A surface plot of the image.

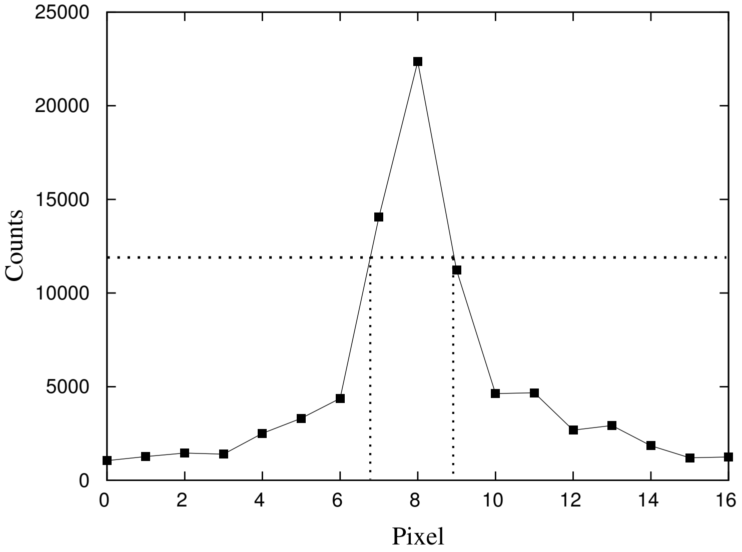

Figure 8.7. A horizontal intensity profile of the single ion shown in figure 8.6. The FWHM of the profile is ~ 2 pixels.

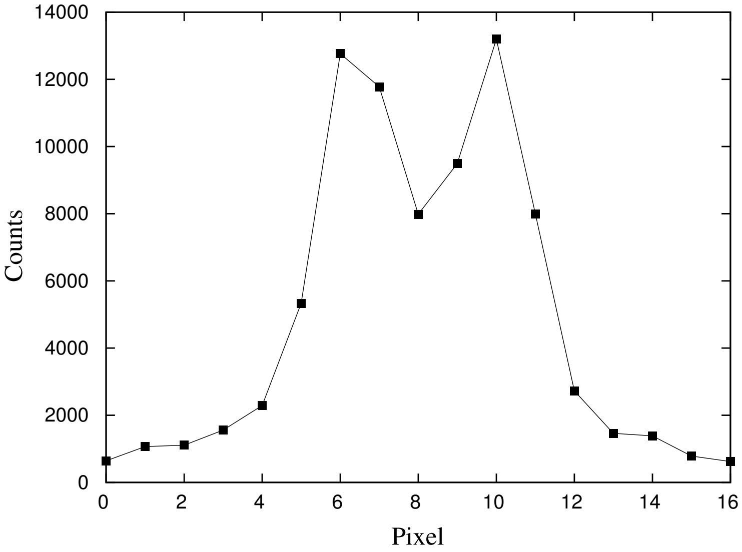

Figure 8.8. A horizontal intensity profile of a pair of ions shown in figure 8.9.

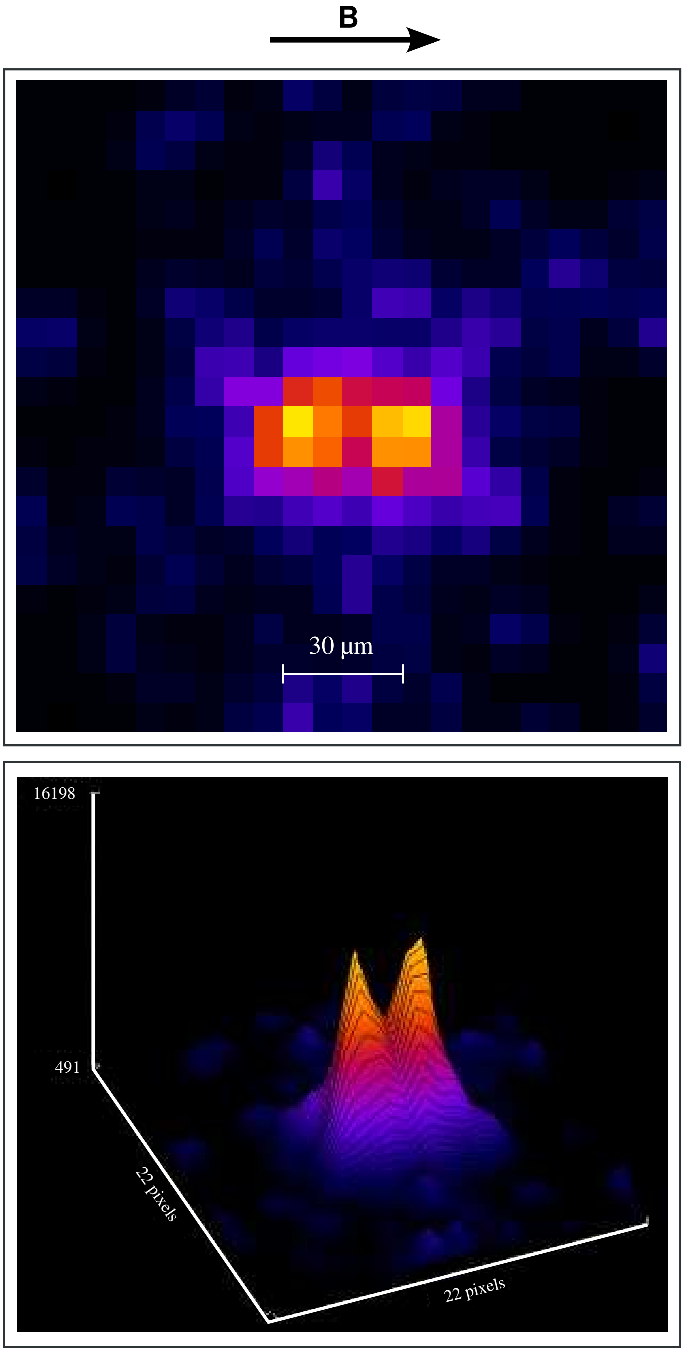

Figure 8.9. *Top:* The iCCD image of two ions in the Penning trap. The two ions are actually separated by $\sim 30\mu$m along the axis.

CHAPTER 9

Conclusions and Suggestions for Future Work

9.1 Conclusions

The project was aimed at laser cooling individual $^{40}\mathrm{Ca}^+$ ions in the Penning trap for QIP applications. Contrary to the RF trap, the Penning trap uses only static electric and magnetic fields. Therefore decoherence in the Penning trap should be in principle less than that in the RF trap.

In chapter 4 we proposed 'Pad traps' a novel Penning trap design for QIP applications. Pad traps are a hexagonal array of pad electrodes which are scalable in 2D. The ions can be in principle hopped in a cyclotron motion from one trap to another by imposing a linear gradient of electric field between adjacent traps. A prototype of the trap has been recently built and will be tested soon.

The first step in our project was to actually trap a single Ca^+ ion in a Penning trap and laser cool it. Although trapping and laser cooling single Ca^+ ions in RF traps is routinely achieved in a number of groups including our own, in the Penning trap this had never been demonstrated prior to the work reported here. This is due to the non-trivial laser cooling process in the presence of a magnetic field. The number of transitions that need to be addressed is greater in the Penning trap due to the Zeeman splitting of the levels. The Ca^+ ion laser cooling scheme in the Penning trap involved 7 diode lasers altogether. We used two UV lasers at 397nm to address the $\mathrm{S}_{1/2} \rightarrow \mathrm{P}_{1/2}$ transition, four IR lasers at 866nm to repump the $\mathrm{D}_{3/2}$ metastable sublevels and one IR laser at 854nm to repump the $D_{5/2}$ metastable sublevels.

Initial attempts at trapping a single ion with only two repumper lasers with imposed sidebands was unsuccessful because of the low signal count rate per ion. We investigated whether coherent population trapping could decrease the laser cooling efficiency and scattering rate per ion as smaller numbers of ions are trapped. We showed that due to the large number of lasers involved in the laser cooling of Ca^+ ion in a Penning trap, CPT can occur in various situations. However we were able to show that with appropriate laser detuning, maximum laser cooling could in principle be achieved just as in the case of the RF trap.

Abandoning sidebands and using individual lasers for each of the 866nm repumping transitions enabled us to successfully trap and laser cool single Ca^+ ions in the Penning trap. Performing a saturation curve experiment for the IR transitions showed that $\sim 70\mu$W was enough to saturate IR transitions at 866nm. This power was indeed higher than the power that sidebands were providing. However we believe that using individual lasers, giving the freedom to adjust each laser frequency independently was also important. The use of optical fibres to ensure good beam overlap at the ion position was equally important. Once we had achieved single ion loading, the signal to background of the quantum jump trace was significantly improved using a smaller pinhole ($\sim 200\mu$m) before the PMT.

We were also able to resolve crystals of pairs of ions along the axis of the trap using an iCCD camera. The separation of the two ions was in agreement with the calculated distance given by the trap bias and trap geometry.

In order to perform the next stage of the experiment, which will be to coherently drive a single ion on the $S_{1/2} \rightarrow D_{5/2}$ transition (the qubit transition), a Ti:Sapphire laser was stabilised in frequency using the Pound-Drever-Hall locking technique. The laser frequency was locked within $\sim$ 3kHz to a temperature stabilised high finesse ($\mathcal{F} \sim 250,000$) cavity.

9.2 Future Work

As mentioned earlier the lasers tend to drift even when locked to the cavities. We believe that the main reason is air pressure changes that give rise to changes of the refractive index of the air and therefore the resonance frequency of the reference cavity.

To overcome these pressure fluctuations the reference cavities have now been placed inside a vacuum chamber. At the moment there is work in progress to lock one IR laser to a cavity and then lock the other 4 lasers to it by a computer-controlled feedback using the scanning Fabry-Perot interferometer signal. This is obviously not a very fast lock but it will compensate for long term drifts and it also saves us from having 4 individual temperature stabilised cavities and feedback control circuits. Preliminary results show that this system is effective in compensating for the long term drifts previously observed.

Conventional electromagnets cannot provide the field stability that is required for the QIP experiments that we intend to perform in the Penning trap. Therefore getting the cylindrical trap to work in the superconducting magnet is the next crucial step. Due to the special cylindrical geometry of the old design it was difficult to run the trap as a RF trap. Therefore one had to be confident about the magnetic field strength at the centre of the trap to tune the lasers' frequencies accordingly. To decrease the number of uncertainties in the parameter space, it has now been decided to redesign the cylindrical trap so that it can initially be run as a RF trap. Working on the cylindrical trap is a priority for the experiment in the near future.

The Ti:Sapphire laser linewidth needs to be improved further by removing the 88kHz noise resulting from the etalon dithering. This can be achieved firstly by reducing the dithering amplitude of the etalon and then finding the optimum lateral position of the etalon where the noise is at a minimum. To further reduce the intensity noise an intensity feedback circuit needs to be built and fed to the first AOM.

The group is also working on the pad trap proposal. A prototype array of three pad traps has been designed and built using a circuit board and is ready to be tested in the conventional electromagnet. Realisation of pad traps is essential to demonstrate the scalability of the Penning trap for QIP purposes. If the prototype works in the conventional electromagnets, a miniaturised version of it suitable for the superconducting magnet will be built and tested.

APPENDIX A

Optical Layout

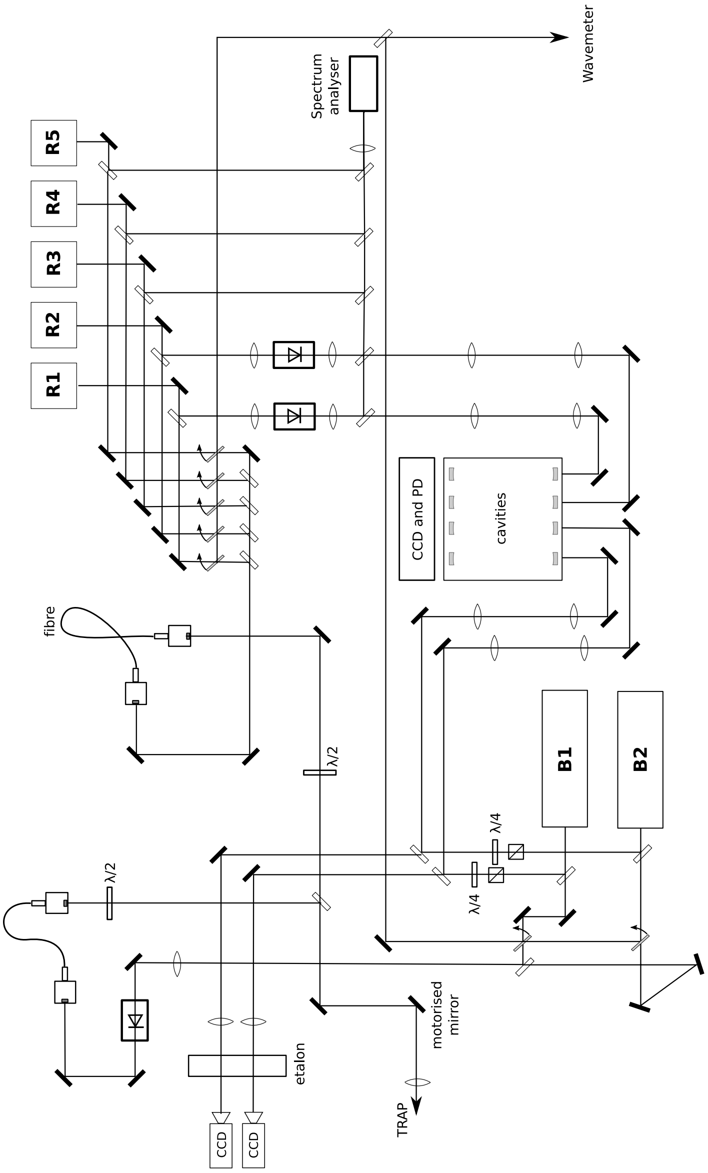

APPENDIX B

CPT Simulation

The complete set of equations describing CPT for a 4-level system (see section 7.2) is as follows:

$$\begin{aligned}
\mathrm{i}\dot{\rho}_{00} &= +\frac{1}{2}(\Omega^*_{R1}\tilde{\rho}_{01} + \Omega^*_{R2}\tilde{\rho}_{02} + \Omega^*_{R3}\tilde{\rho}_{03}) + \text{c.c.} + \Re_{00} \\
\mathrm{i}\dot{\rho}_{11} &= -\frac{1}{2}\Omega^*_{R1}\tilde{\rho}_{01} + \text{c.c} + \Re_{11} \\
\mathrm{i}\dot{\rho}_{22} &= -\frac{1}{2}\Omega^*_{R2}\tilde{\rho}_{02} + \text{c.c} + \Re_{22} \\
\mathrm{i}\dot{\rho}_{33} &= -\frac{1}{2}\Omega^*_{R3}\tilde{\rho}_{03} + \text{c.c} + \Re_{33} \\
\mathrm{i}\dot{\tilde{\rho}}_{01} &= -\delta_{L1}\tilde{\rho}_{01} + \frac{\Omega_{R1}}{2}(\rho_{00} - \rho_{11}) - \frac{1}{2}\Omega_{R2}\tilde{\rho}_{21} - \frac{1}{2}\Omega_{R3}\tilde{\rho}_{31} + \Re_{01} \\
\mathrm{i}\dot{\tilde{\rho}}_{02} &= -\delta_{L2}\tilde{\rho}_{02} + \frac{\Omega_{R2}}{2}(\rho_{00} - \rho_{22}) - \frac{1}{2}\Omega_{R1}\tilde{\rho}_{12} - \frac{1}{2}\Omega_{R3}\tilde{\rho}_{32} + \Re_{02} \\
\mathrm{i}\dot{\tilde{\rho}}_{03} &= -\delta_{L3}\tilde{\rho}_{03} + \frac{\Omega_{R3}}{2}(\rho_{00} - \rho_{33}) - \frac{1}{2}\Omega_{R2}\tilde{\rho}_{23} - \frac{1}{2}\Omega_{R1}\tilde{\rho}_{13} + \Re_{03} \\
\mathrm{i}\dot{\tilde{\rho}}_{12} &= +\frac{1}{2}\Omega_{R2}\tilde{\rho}^*_{01} - \frac{1}{2}\Omega^*_{R1}\tilde{\rho}_{02} - \delta_{R12}\tilde{\rho}_{12} + \Re_{12} \\
\mathrm{i}\dot{\tilde{\rho}}_{13} &= +\frac{1}{2}\Omega_{R3}\tilde{\rho}^*_{03} - \frac{1}{2}\Omega^*_{R1}\tilde{\rho}_{03} - \delta_{R13}\tilde{\rho}_{13} + \Re_{13} \\
\mathrm{i}\dot{\tilde{\rho}}_{23} &= +\frac{1}{2}\Omega_{R3}\tilde{\rho}^*_{02} - \frac{1}{2}\Omega^*_{R2}\tilde{\rho}_{03} - \delta_{R23}\tilde{\rho}_{23} + \Re_{23},
\end{aligned}$$

where ρ_{ij} is the ij term of the density operator, Ω_{Ri} is the Rabi frequency of laser i, and $\mathcal{R}_{ij}$ is the relaxation term.

```
----------------------------------
    Script for Mathematica v6.0
    Last modified 05-06-2007
----------------------------------

Γ01 = 12/13;
Γ02 = 1/26;
Γ03 = 1/26;
Γ21 = Γ31 = Γ32 = 0;
δ12 = δ13 = δ23 = 0;

Ω1 = 0.2;
Ω2 = Ω3 = 0.5;
tf = 200;
δ3 = 0;

γ10 = γ20 = γ30 = 0.5;
list1 = {};
biglist1 = {};
biglist2 = {};
biglist3 = {};
biglist4 = {};
list2 = {};
list3 = {};
list4 = {};
list5 = {};
list6 = {};

EAD = {
{
iρ00′[t] == −i(Γ01 + Γ02 + Γ03)ρ00[t] + Conjugate[Ω1]ρ10[t]
−Ω1Conjugate[ρ10[t]] + Conjugate[Ω2]ρ20[t] − Ω2Conjugate[ρ20[t]]
+Conjugate[Ω3]ρ30[t] − Ω3Conjugate[ρ30[t]]
},
{
iρ11′[t] == iΓ01ρ00[t] + iΓ21ρ22[t] + iΓ31ρ33[t]
+Ω1Conjugate[ρ10[t]] − Conjugate[Ω1]ρ10[t]
},
{
iρ22′[t] == iΓ02ρ00[t] − iΓ21ρ22[t] + iΓ32ρ33[t]
+Ω2Conjugate[ρ20[t]] − Conjugate[Ω2]ρ20[t]
},
{
iρ33′[t] == iΓ03ρ00[t] − iΓ31ρ33[t] − iΓ32ρ33[t]
+Ω3Conjugate[ρ30[t]] − Conjugate[Ω3]ρ30[t]
},
{
iρ10′[t] == −(δ1 + iγ10)ρ10[t] + Ω1ρ00[t] − Ω1ρ11[t]
−Ω2ρ12[t] − Ω3ρ13[t]
},
{
iρ20′[t] == −(δ2 + iγ20)ρ20[t] + Ω2ρ00[t] − Ω1Conjugate[ρ12[t]]
−Ω2ρ22[t] − Ω3ρ23[t]
},
{
iρ30′[t] == −(δ3 + iγ30)ρ30[t] + Ω3ρ00[t] − Ω1Conjugate[ρ13[t]]
−Ω2Conjugate[ρ23[t]] − Ω3ρ33[t]
},
{
iρ12′[t] == (δ2 − δ1 − iδ12)ρ12[t] + Ω1Conjugate[ρ20[t]]
−Conjugate[Ω2]ρ10[t]
},
{
iρ13′[t] == (δ3 − δ1 − iδ13)ρ13[t] + Ω1Conjugate[ρ30[t]]
−Conjugate[Ω3]ρ10[t]
},
{
iρ23′[t] == (δ3 − δ2 − iδ23)ρ23[t] + Ω2Conjugate[ρ30[t]]
−Conjugate[Ω3]ρ20[t]
```

```
}
};
counter = 0;
Do[
Do[
sol = NDSolve[
{
EAD,
ρ00[0] == 1, ρ11[0] == 0, ρ22[0] == 0, ρ33[0] == 0, ρ10[0] == 0,
ρ20[0] == 0, ρ30[0] == 0, ρ12[0] == 0, ρ13[0] == 0, ρ23[0] == 0
},
{
ρ00, ρ11, ρ22, ρ33, ρ10,
ρ20, ρ30, ρ12, ρ13, ρ23
},
{t, 0, tf}
];

list1 = Append[list1, Evaluate[Re[ρ00[tf]]/.First[sol]]];
 list2 = Append[list2, Evaluate[Re[ρ11[tf]]/.First[sol]]];
 list3 = Append[list3, Evaluate[Re[ρ22[tf]]/.First[sol]]];
 list4 = Append[list4, Evaluate[Re[ρ33[tf]]/.First[sol]]];
, {δ1, −2, 2, 0.1}
];

biglist1 = Append[biglist1, list1];
biglist2 = Append[biglist2, list2];
biglist3 = Append[biglist3, list3];
biglist4 = Append[biglist4, list4];

list1 = {};
list2 = {};
list3 = {};
list4 = {};
counter = counter + 1;
, {δ2, −2, 2, 0.1}
]
a = ListPlot3D[
biglist1,
PlotRange → All,
Mesh → True
]
b = ListPlot3D[
biglist2,
PlotRange → All,
Mesh → True
]
c = ListPlot3D[
biglist3,
PlotRange → All,
Mesh → True
]
d = ListPlot3D[
biglist4,
PlotRange → All,
Mesh → True
]
Block[{$DisplayFunction = Identity},
a = ListContourPlot[
biglist1,
DataRange → {{−2, 2}, {−2, 2}},
ContourLabels → Automatic
];
b = ListContourPlot[
biglist2,
DataRange → {{−2, 2}, {−2, 2}},
ContourLabels → Automatic
];
c = ListContourPlot[
biglist3,
DataRange → {{−2, 2}, {−2, 2}},
ContourLabels → Automatic
];
d = ListContourPlot[
biglist4,
```

```
DataRange → {{−2, 2}, {−2, 2}},
ContourLabels → Automatic
]
];

Show[GraphicsGrid[{{a, b}, {c, d}}]]
Export["trapped-contour-all4.eps", %, "EPS", ImageSize → 400];
```

APPENDIX C

Cavity Coupling

```
----------------------------------
    Script for Mathematica v6.0
    Last modified 05-02-2007
----------------------------------
```

```
(*Coupling*)
(* Mathematica 6.0 script *)

L = 0.1; (* Length of cavity *)
f1 = −0.10; (* Focal length of positive lens *)
f2 = 0.10; (* Focal length of negative lens *)
d1min = 0.01; (* Distance between two lenses *)
d2min = 0.01; (* Distance between second lens and cavity *)
d1max = 0.5; (* Max separation of lenses in metres *)
d2max = 0.5; (* Max separation of lenses and cavity in metres *)
nz = 1.5; (* Refractive index of mirror *)
na = 1.0; (* Refractive index of air *)
d3 = 0.01; (* Thickness of mirror *)
R = 0.35; (* Radius of curvature of mirror *)
lambda = 729 ∗ 10^(−9); (* Wavelength of laser *)
wbeam = 0.00051; (* waist of beam *)
wbeammin = 0.0001;
wbeammax = 0.005;

(* First lens *)
M1 = {{1, 0}, {−1/f1, 1}};

(*Lens − LensFreeDistance*)
M2 = {{1, d1}, {0, 1}};

(* Second Lens *)
M3 = {{1, 0}, {−1/f2, 1}};

(*Lens − CavityMirror*)
M4 = {{1, d2}, {0, 1}};

(* Air-Mirror *)
M5 = {{1, 0}, {0, na/nz}};

(* Mirror Distance *)
M6 = {{1, d3}, {0, 1}};

(* Mirror -Cavity *)
M7 = {{1, 0}, {(nz − nz)/(na ∗ R), nz/na}};

(* Half of Cavity Free Distance *)
M8 = {{1, L/2}, {0, 1}};

(*Waistofbeamincavity(min .spotsize), seeMilonni&Eberly, p.498*)
wcav = Sqrt[lambda/(2 ∗ Pi)] ∗ (2 ∗ R ∗ L − L^2)^0.25;
qcav = −I ∗ Pi ∗ wcav^2/lambda;
```

```
(* Waist of laser beam *)
qbeam = −I * Pi * wbeam^2/lambda;

M = M8.M7.M6.M5.M4.M3.M2.M1;
qvector = M.{qbeam, 1};

qNewInverse = qvector[[2]]/qvector[[1]];
wBeamNew =  Sqrt[lambda/(Im[qNewInverse] * Pi)];

myplot = Plot3D[wBeamNew, {d1, d1min, d1max}, {d2, d2min, d2max}, PlotPoints → 100, Mesh → False, FaceGrids → All,
AxesLabel → {"distance between two lenses [m]", "distance to the cavity [m]", "Beam Waist"},
PlotLabel → "f1 = 0.25, f2=-0.1, Waist of Beam = 2.5 mm"];

Show[myplot, Plot3D[wcav, {d1, d1min, d1max}, {d2, d2min, d2max}], PlotRange → {0, 10wcav}, ViewPoint → {0.001, 0, 1}]
Export["out.ps", %, "EPS", ImageSize → 300];

ContourPlot[wBeamNew == wcav, {d1, d1min, d1max}, {d2, d2min, d2max},
FrameLabel → {"distance between the lenses [m]", "distance to the cavity [m]"},
PlotLabel → "f1 = -10cm, f2 = 10cm, Waist of Beam = 510 micro meter", PlotPoints → 300, GridLines → Automatic]
Export["out2.ps", %, "EPS", ImageSize → 300];
```

APPENDIX D

Electromagnet Stabilisation Circuit Diagram

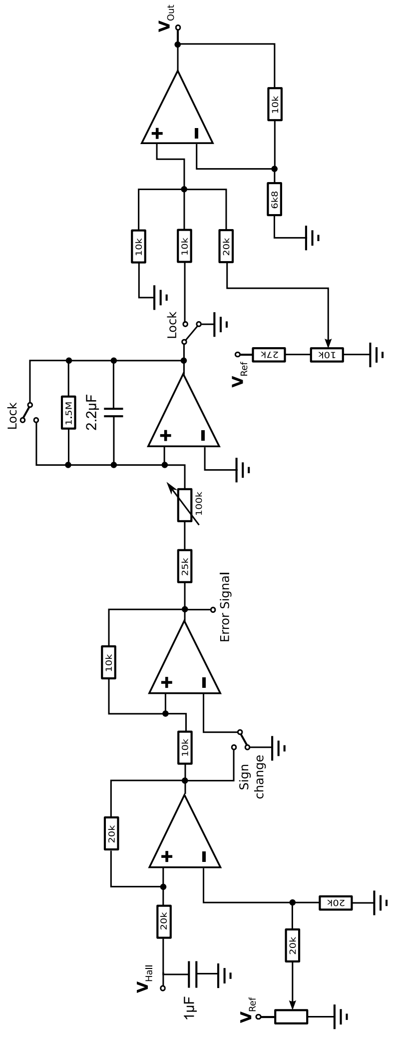

Figure D.1. Feedback circuit for electromagnet stabilisation

Bibliography

[1] G.E. Moore, *Electronics*, **83**(8) (1965). 1

[2] R. Feynman *et al.*, *Int. J. Theor. Phys.*, **21**:467 (1982). 1

[3] D.P. DiVincenzo, *Quant. Inf. Comp.*, **1 (special)**:1 (2001). 1, 2.3, 3.1

[4] F. Schmidt-Kaler *et al.*, Realization of the Cirac-Zoller Controlled-NOT quantum gate, *Nature*, **422**:408 (2003). 1, 3.6, 3.7

[5] D. Leibfried *et al.*, Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate, *Nature*, **422**:412 (2003). 1, 3.6, 2, 3.7

[6] J.R. Castrejon-Pita, H. Ohadi, D.R. Crick, D.F.A. Winters, D.M. Segal, and R.C. Thompson, Novel designs for Penning ion traps, *J. Mod. Opt.*, **54**(11):1581 (2007). 1, 4.7

[7] D.R. Crick, H. Ohadi, I. Bhatti, R.C. Thompson, and D.M. Segal, Two-ion Coulomb crystals of Ca^+ in a Penning trap., *submitted to Optics Express* (2008). 1

[8] D. Deutsch, *Proc. R. Soc. London A*, **400**:97 (1985). 2.1

[9] P.W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings of the Symposium on the Foundations of Computer Science, 1994, Los Alamitos, California*, page 124, IEEE Computer Society Press (1994). 2.1

[10] L.K. Grover, Quantum mechanics helps in searching for a needle in a haystack, *Phys. Rev. Lett.*, **79**:325 (1997). 2.1

[11] C.H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin, Experimental quantum cryptography, *Journal of Cryptology*, **5**(1):3 (1992). 2.1

[12] R. Landauer, Is Quantum Mechanics useful?, *Phil. Trans. R. Soc. London, Ser. A*, **353**:367 (1995). 2.1

[13] D. Deutsch and R. Jozsa, Rapid solution of problems by quantum computation, *Proc. R. Soc. London A*, **439**:553 (1992). 2.2.1

[14] A. Einstein, B. Podolsky, and N. Rosen, Can quantum-mechanical description of physical reality be considered complete?, *Phys. Rev. Lett.*, **47**:777 (1935). 2.2.1

[15] M.A. Nielson and I.L. Chuang, *Quantum Computation and Quantum Information*, Cambridge (2000). 2.2.1, 2.2.3

[16] W.K. Wooters and W.H. Zurek, A single quantum cannot be cloned, *Nature*, **299**:802 (1982). 2.2.3

[17] D.P. DiVincenzo, Two-bit gates are universal for quantum computation, *Phys. Rev. Lett.*, **51**:51 (1995). 2.2.3

[18] P.W. Shor, Scheme for reducing decoherence in quantum computer memory., *Phys. Rev. A.*, **52**:2493 (1995). 2.2.4

[19] A.M. Steane, Error correcting codes in quantum theory., *Phys. Rev. Lett.*, **77**:793 (1996). 2.2.4

[20] A.M. Steane, Multiple particle interference and quantum error correction, *Proc. Roy. Soc. Lond. A*, **452**:2551 (1996). 2.2.4

[21] J. Chiaverini *et al.*, Realization of quantum error correction., *Nature*, **432**(7017):602 (2004). 2.2.4

[22] P.W. Shor, Fault-tolerant quantum computation, *Proceedings of the Symposium on the Foundations of Computer Science* (1996), `arXiv:quant-ph/9605011`. 2.2.4

[23] J. Preskill, Fault-tolerant quantum computation (1997), `arXiv:quant-ph/9712048v1`. 2.2.4

[24] J.I. Cirac and P. Zoller, Quantum computations with cold trapped ions, *Phys. Rev. Lett.*, **74**:4091–4094 (1995). 3.1, 3.6

[25] H.G. Dehmelt, Proposed 10^{14} $\Delta\nu > \nu$ laser fluorescence spectroscopy on Tl^{+} mono-ion oscillator II (spontaneous quantum jumps), *Bull. Am. Phys. Soc.*, **20**:60 (1975). 3.2

[26] P.K. Ghosh, *Ion traps*, Oxford (1995). 3.3

[27] J.J. Sakurai, *Modern Quantum Mechanics*, Addison-Wesley (1985). 3.3

[28] D.J. Wineland *et al.*, Experimental issues in coherent quantum-state manipulation of trapped atomic ions, *J. Res. Natl. Inst. Stand. Technol.*, **103**:259 (1998). 3.3

[29] D.J. Wineland and W.M. Itano, Laser cooling of atoms, *Phys. Rev. A: At., Mol., Opt. Phys.*, **20**:1521 (1979). 3.3, 3.5.1, 3.5.2, 4.4

[30] S. Stenholm, The semiclassical theory of laser cooling, *Rev. Mod. Phys.*, **58**:699 (1986). 3.5

[31] C. Monroe, D.M. Meekhof, B.E. King, S.R. Jefferts, W.M. Itano, and D.J. Wineland, Resolved-sideband Raman cooling of a bound atom to the 3D zero-point energy, *Phys. Rev. Lett.*, **75**(22):4011 (1994). 3.4, 3.5.2

[32] F. Diedrich, J.C. Bergquist, W.M. Itano, and D.J. Wineland, Laser cooling to the zero-point energy of motion, *Phys. Rev. Lett.*, **62**:403 (1988). 3.6

[33] R. Jozsa, Fidelity for mixed quantum states, *J. Mod. Opt.*, **41**:2315 (1994). ∗

[34] G.J. Milburn, S. Schneider, and D.F.V. James, Ion trap quantum computing with warm ions, *Fortschritte der Physik*, **48**:801 (2000). 3.6, 2, 3.7

[35] C. Monroe, D.M. Meekhof, B.E. King, W.M. Itano, and D.J. Wineland, Demonstration of a fundamental quantum logic gate, *Phys. Rev. Lett.*, **75**(25):4714 (1995). 3.7

[36] Q.A. Turchette *et al.*, Deterministic entanglement of two trapped ions, *Phys. Rev. Lett.*, **81**:3631 (1998). 3.7

[37] C.A. Sackett *et al.*, Experimental entanglement of four particles, *Nature*, **404**:256 (2000). 3.7

[38] D. Leibfried *et al.*, Toward Heisenberg-limited spectroscopy with multiparticle entangled states, *Science*, **304**:1476 (2004). 3.7

[39] C.F. Roos *et al.*, Control and measurement of three-qubit entangled states, *Science*, **304**:1478 (2004). 3.7

[40] D. Leibfried *et al.*, Creation of a six-atom ‘Schrödinger cat’ state, *Nature*, **438**:639 (2005). 3.7

[41] H. Häffner *et al.*, Scalable multiparticle entanglement of trapped ions., *Nature*, **438**:643 (2005). 3.7

[42] A.M. Steane, Space, time, parallelism and noise requirements for reliable quantum computing, *Fortschritte der Physik*, **46**:443 (1998). 3.7

[43] J. Eschner *et al.*, Light interference from single atoms and their mirror images, *Nature*, **413**:495 (2001). 3.7

[44] M. Keller *et al.*, Single photon source on the basis of a trapped ion, *Nature*, **431**:1075 (2004). 3.7

[45] J.I. Cirac and P. Zoller, A scalable quantum computer with ions in an array of microtraps, *Nature*, **404**:579 (2000). 3.7

[46] D. Kielpinski, C. Monroe, and D.J. Wineland, Architechture for a large-scale ion-trap quantum computer, *Nature*, **417**:709 (2002). 3.6, 3.7

[47] M.D. Barrett *et al.*, Deterministic quantum teleportation of atomic qubits, *Nature*, **429**:737 (2004). 3.7

[48] S. Seidelin *et al.*, Microfabricated surface-electrode ion trap for scalable quantum information processing., *Phys. Rev. Lett.*, **96**(25):253003 (2006). 3.7

[49] D. Stick, W.K. Hensinger, S. Olmschenk, M.J. Madsen, K. Schwab, and C. Monroe, Ion trap in a semiconductor chip, *Nature*, **2**:36 (2006). 3.7

[50] W.K. Hensinger *et al.*, T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation, *Appl. Phys. Lett.*, **88**:034101 (2006). 3.7

[51] M.D. Barrett *et al.*, Long-lived qubit memory using atomic ions, *Phys. Rev. Lett.*, **95**:060502 (2005). 3.7

[52] Q.A. Turchette *et al.*, Heating of trapped ions from the quantum ground state, *Phys. Rev. A*, **61**:063418 (2000). 4.2

[53] L. Deslauriers *et al.*, Scaling and suppression of anomalous heating in ion traps., *Phys. Rev. Lett.*, **97**:103007 (2006). 4.1, 4.2

[54] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibrandt, K.R. Brown, and I.L. Chuang, Suppression of heating rates in cryogenic surface-electrode ion traps (2007), `arXiv:0706.3763v1[quant-ph]`. 4.2

[55] F. Schmidt-Kaler *et al.*, The coherence of qubits based on single Ca^+ ions, *J. Phys. B: At., Mol. Opt. Phys.*, **36**:623–636 (2003). 4.2

[56] P.K. Gosh, *Ion Traps (Clarendon, Oxford, UK, 1995).*, Clarendon Press, Oxford (1995). 4.3

[57] F.G. Major, V.N. Gheorghe, and G. Werth, *Charged particle traps*, Springer (2005). 4.3

[58] W.M. Itano and D.J. Wineland, Laser cooling of ions stored in harmonic and Penning traps, *Phys. Rev. A: At., Mol., Opt. Phys.*, **25**:35 (1982). 4.4

[59] R.C. Thompson and J. Papadimitriou, Simple model for the laser cooling of an ion in a Penning trap, *J. Phys. B: At., Mol. Opt. Phys.*, **33**(17):3393–3405 (2000). 4.4

[60] E.S. Phillips, R.J. Hendricks, A.M. Abdulla, H. Ohadi, D. Crick, K. Koo, D.M. Segal, and R.C. Thompson, Dynamics of axialized laser-cooled ions in a Penning trap, *in preparation* (2007), `arXiv:0710.4401v1[quant-ph]`. 4.5

[61] H.F. Powell, S.R. de Echaniz, E.S. Phillips, D.M. Segal, and R.C. Thompson, Improvement of laser cooling of ions in a Penning trap by use of the axialization technique, *J. Phys B: At. Mol. Opt. Phys.*, **36**:961–970 (2003). 4.6

[62] R.J. Hendricks, E.S. Phillips, D.M. Segal, and R.C. Thompson, Laser cooling in the Penning trap: an analytical model for cooling rates in the presence of an axializing field (2007), `arXiv:0709.3817v1[quant-ph]`. 4.6

[63] H.F. Powell, D.M. Segal, and R.C. Thompson, Axialization of laser cooled Magnesium ions in a Penning trap, *Phys. Rev. Lett.*, **89**:093003 (2002). 4.6, 4.9

[64] G. Ciaramicoli, I. Marzoli, and P. Tombesi, Scalable quantum processor with trapped electrons, *Phys. Rev. Lett.*, **91**(1):017901 (2003). 4.7, 4.10

[65] S. Stahl *et al.*, A planar Penning trap, *Euro. Phys. J. D*, **32**:139 (2005). 4.7, 4.10

[66] M. Vogel and F. Galve, private communication, GSI and University of Mainz. 4.7

[67] J.R. Castrejon-Pita and R.C. Thompson, Proposal for a planar Penning ion trap, *Phys. Rev. A: At., Mol., Opt. Phys.*, **72**:013405 (2005). 4.7, 4.10

[68] SIMION, URL `http://www.simion.com/`. 4.15

[69] Online NIST atomic spectra database, Version 3, URL `http://physics.nist.gov/PhysRefData/ASD/`. 5.2, 5.6

[70] M.E.M. Storkey, *Studies of laser cooled trapped ions*, Ph.D. thesis, Imperial College London (2001). 5.3

[71] E.S. Phillips, *Controlled dynamics of laser-cooled ions in a Penning trap*, Ph.D. thesis, Imperial College London (2004). 5.3, 5.4, 5.8, 5.9.1

[72] M.A. van Eijkelenborg, M.E.M. Storkey, D.M. Segal, and R.C. Thompson, Ion dynamics in a novel linear combined trap, *Int. J. Mass Spectrom.*, **188**:155 (1999). 5.3

[73] K. Koo, *Laser cooling and trapping of Ca^{+} ions in a Penning trap*, Ph.D. thesis, Imperial College (2003). 5.4, 5.6, 5.7.1, 8.1

[74] C.E. Wieman and L. Hollberg, Using diode lasers for atomic physics, *Rev. Sci. Instrum.*, **62**:1 (1991). 5.7.1

[75] L. Ricci *et al.*, A compact grating-stabilized diode laser system for atomic physics, *Opt. Commun.*, **117**:541 (1995). 5.7.1

[76] A.S. Arnold, J.S. Wilson, and M.G. Boshier, A simple extended-cavity diode laser, *Rev. Sci. Instrum.*, **69**:1236 (1998). 5.7.1

[77] R.J. Hendricks, *Spectroscopy and Dynamics of Laser-Cooled Ca^{+} Ions in a Penning Trap*, Ph.D. thesis, Imperial College London (2006). 5.7.1, 5.9.1, 8.1, 8.2

[78] G. Blasbichler, *Ein Lambdameter mit* 10^7 *Meunsicherheit*, Master's thesis, University of Innsbruck (2000). 5.8, 5.8

[79] J. Sudbery, *Studies of laser cooled Calcium ions in the Penning and combined traps*, Ph.D. thesis, Imperial College London (2003). 5.9.1, 6.1

[80] K.B. MacAdam, A. Steinbach, and C. Wieman, A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb, *Am. J. Phys.*, **60**(12):1098 (1992). 5.9.1

[81] D.M. Lucas *et al.*, Isotope-selective photoionization for calcium ion trapping, *Phys. Rev. A: At., Mol., Opt. Phys.*, **69**(1):012711 (2004). 5.10

[82] M. Brownnutt, V. Letchumanan, G. Wilpers, R.C. Thompson, P. Gill, and A.G. Sinclair, Controlled photoionization loading of $^{88}\mathrm{Sr}^+$ for precision ion-trap experiments, *Appl. Phys. B: Lasers Opt.*, **87**:411 (2007). 5.10

[83] R.V. Pound, Electronic frequency stabilization of microwave oscillators, *Rev. Sci. Instrum.*, **17**:490 (1946). 6.1

[84] R.W.P. Drever, J.L. Hall, F.V. Kowalski, J. Hough, G.M. Ford, A.J. Munley, and H. Ward, Laser phase and frequency stabilization using an optical resonator, *Appl. Phys. B: Lasers Opt.*, **31**:97 (1983). 6.1, 6.3.6

[85] A.E. Siegman, *Lasers*, University Science Books (1986). 6.2, 6.3.1

[86] E.D. Black, An introductaion to Pound-Drever-Hall laser frequency stabilization, *Am. J. Phys.*, **69**:79–87 (2001). 6.2

[87] W. Demtroder, *Laser Spectroscopy: Basic Concepts and Instrumentation*, Springer-Verlag (2002). 6.3.1

[88] E.A. Whittaker, M. Gehrtz, and G.C. Bjorklund, Residual amplitude modulation in laser electro-optic phase modulation, *J. Opt. Soc. Am. B*, **2**:1320 (1985). 2

[89] W.W. Macalpine and R.O. Schildknecht, Coaxial resonators with helical inner conductor, *Proceedings of the IRE*, **47**:2099 (1959). 6.3.3

[90] A. Zverev and H. Blinchikoff, Realization of a filter with helical components, *IRE Transactions on Component Parts*, **8**:99 (1961). 6.3.3

[91] A.I. Zverev, *Handbook of Filter Synthesis*, Wiley-Interscience (1967). 6.3.3

[92] J.W. Berthold and S.F. Jacobs, Ultraprecise thermal expansion measurements of seven low expansion materials, *Applied Optics*, **15**:2344 (1976). 6.3.4

[93] P.W. Milonni and J.H. Eberly, *Lasers*, Wiley-Interscience (1988). 6.3.4, 6.3.4

[94] M. Zhu and J.L. Hall, Stabilization of optical phase/frequency of a laser system: application to a commercial dye laser with an external stabilizer, *J. Opt. Soc. Am. B.*, **10**:802 (1993). 6.3.6

[95] Spice, URL `http://bwrc.eecs.berkeley.edu/Classes/IcBook/SPICE/`. 6.3.6

[96] D.S. Elliott, R. Roy, and S.J. Smith, Extracavity laser band-shape and bandwidth modification, *Phys. Rev. A*, **26**:12 (1982). ∗

[97] G. Kirchmair, *Frequency stabilization of a Titanium-Sapphire laser for precision spectroscopy on Calcium ions*, Ph.D. thesis, University of Innsbruck (2006). 6.3.7

[98] G. Alzetta, A. Gozzini, L. Moi, and G. Orriols, An experimental method for the observation of r.f. transitions and laser beat resonances in oriented Na vapor, *Nuovo Cimento*, **36**:5 (1976). 7.1

[99] M.O. Scully, S.Y. Zhu, and A. Gavrielides, Degenerate quantum-beat laser: Lasing without inversion and inversion without lasing., *Phys. Rev. Lett.*, **62**:2813 (1989). 7.1

[100] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. Cohen-Tannoudji, Laser cooling below the one-photon recoil energy by velocity-selective coherent population trapping, *Phys. Rev. Lett*, **61**:826 (1988). 7.1

[101] S.E. Harris, J.E. Field, and A. Imamoglu, Nonlinear optical processes using electromagnetically induced transparency., *Phys. Rev. Lett.*, **64**:1107 (1990). 7.1

[102] A. Kasapi, Enhanced isotope discrimination using electromagnetically induced transparency., *Phys. Rev. Lett.*, **77**:1035 (1996). 7.1

[103] G. Janik, W. Nagourney, and H. Dehmelt, Doppler-free optical spectroscopy on the Ba^+ mono-ion oscillator, *J. Opt. Soc. Am. B*, **2**:1251 (1985). 7.2, 7.1

[104] M.G. Boshier, G.P. Barwood, G. Huang, and H.A. Klein, Polarisation-dependent optical pumping for interrogation of a magnetic-field-independent “clock” transition in laser-cooled trapped $^{87}\mathrm{Sr}^+$, *Appl. Phys. B: Lasers Opt.*, **71**:51 (2000). 7.1

[105] E. Arimondo, Coherent population trapping in laser spectroscopy, *Progress in Optics*, **35**:259 (1996). 7.2

[106] M.O. Scully and M.S. Zubairy, *Quantum Optics*, Cambridge University Press (1997). 7.2

[107] R. Loudon, *The Quantum Theory of Light*, Oxford University Press (2000). 7.4

[108] M.A. van Eijkelenborg, M.E.M. Storkey, D.M. Segal, and R.C.. Thompson, Sympathetic cooling and detection of molecular ions in a Penning trap, *Phys. Rev. A*, **60**:3903 (1999). 8.2

[109] H.F. Powell, M.A. van Eijkelenborg, W. Irvine, D.M. Segal, and R.C. Thompson, Quantum jumps in singly ionized magnesium, *J. Phys B: At. Mol. Opt. Phys.*, **35**:205 (2002). 8.4

[110] P.A. Barton, C.J.S. Donald, D.M. Lucas, D.A. Stevens, A.M. Steane, and D.N. Stacey, Measurement of the lifetime of the $3d^2D_{5/2}$ state in $^{40}\mathrm{Ca}^+$, *Phys. Rev. A*, **62**:032503 (2000). 8.4