

Spectroscopy and Dynamics of
Laser-Cooled Ca^+ Ions
in a Penning Trap

Richard James Hendricks

October 2006

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

Abstract

I report on work geared towards improving the efficiency of laser cooling in the Penning trap by manipulating the dynamics of the trapped ions. The wider aim of this project is to measure heating and decoherence rates in order to determine the feasibility of quantum information processing using ions trapped in a Penning trap. In this thesis I introduce the basic concepts of quantum information processing and its implementation in radiofrequency trapped-ion systems. I report on the difficulties currently being faced and how many of these obstacles can be circumvented by using Penning traps for ion confinement. Laser cooling in the Penning trap is usually rather inefficient because of the instability of the magnetron radial mode of motion. By extending an analytical model for laser cooling in the Penning trap I show that a technique called ‘axialisation’ can overcome this inefficiency by coupling the magnetron mode to another, stable radial mode. I present the results of an experiment to implement the axialisation technique and measure the improvement in laser cooling rates of calcium ions, and I assess the extent to which these results agree with those derived from our model. Finally I document our efforts to trap and cool single calcium ions in the Penning trap and discuss modifications to the experimental procedure that may make this possible in the near future.

Contents

1	Introduction	11
2	Quantum information processing	15
2.1	Properties of quantum systems	17
2.1.1	Quantum bits	17
2.1.2	Multi-qubit systems	19
2.1.3	Decoherence	20
2.2	Quantum algorithms	23
2.2.1	Universal computers	23
2.2.2	Quantum gates	24
2.2.3	Shor's quantum factoring algorithm	25
2.2.4	Other quantum algorithms	27
2.3	Quantum error correction	28
2.3.1	Classical error correction	29
2.3.2	Principles of QEC	31
2.3.3	Application of QEC	32
2.3.4	Requirements of QEC	34
2.4	Criteria for the implementation of QIP	35
2.4.1	System characterisation and scalability	36
2.4.2	System initialisation	37
2.4.3	Long decoherence times	38
2.4.4	A universal set of quantum logic gates	38
2.4.5	Qubit measurement	39
2.5	Conclusions	40
3	Quantum information processing in ion traps	42
3.1	Quantum computation with trapped ions	43

CONTENTS

3.1.1	Ion cooling in the linear RF trap	43
3.1.2	Electronic states as qubits	45
3.1.3	Quantum gates for trapped ions	48
3.2	Quantum computing potential of trapped ions	51
3.3	Realisation of ion trap quantum information processing	53
3.3.1	Previous developments in ion trap QIP	53
3.3.2	Current issues in ion trap QIP	61
3.3.3	Future directions	64
3.4	Penning trap quantum information processing	70
3.5	Conclusions	72
4	Trapping and laser cooling in the Penning trap	74
4.1	Electrode configuration	75
4.2	Ion motion in the Penning trap	77
4.3	Laser cooling in the Penning trap	81
4.3.1	The laboratory frame	81
4.3.2	The rotating frame	89
4.4	Axialisation	91
4.4.1	A quantum picture of axialisation	92
4.4.2	A classical analogy for axialisation	93
4.4.3	A model for laser cooling in the presence of axialisation	97
4.4.4	Amplitude and phase response to a driving field	109
4.5	Ion Clouds and the Effects of Space Charge	117
5	Experimental setup	120
5.1	Ca ⁺ electronic structure	120
5.2	Experimental overview	123
5.3	The split-ring trap	125
5.3.1	Trap structure	125
5.3.2	Ion generation	127
5.3.3	Vacuum system	127
5.3.4	Magnetic field generation	129
5.3.5	Operation as Penning trap	129
5.3.6	Operation as RF trap	130
5.4	Doppler cooling lasers	132
5.5	Repumper lasers	134

CONTENTS

5.6	Laser diagnostics	137
5.7	Laser locking	142
5.7.1	Optical reference cavities	142
5.7.2	Locking setup	144
5.7.3	Locking procedure	146
5.7.4	Cavity stability	147
5.8	Fluorescence detection	149
5.9	General procedure	151
6	Axialisation	155
6.1	Experimental methods	156
6.2	Results	165
6.2.1	iCCD images	165
6.2.2	Cooling rate measurements	170
6.2.3	Classical avoided crossing	177
6.3	Conclusions	183
7	Towards single-ion work in the Penning trap	185
7.1	Loading and detection of small numbers of ions	186
7.2	Single ions in the RF trap	189
7.3	Small clouds of ions in the Penning trap	194
7.4	Contaminant ions	203
7.5	Coherent population trapping	212
7.6	Laser and magnetic field instability	229
8	Discussion	231
8.1	Discussion of results	231
8.2	Future directions	232
8.2.1	Improving magnetic field stability	232
8.2.2	Improving laser cavity stability	234
8.2.3	Sideband cooling and beyond	235
	Bibliography	237

List of Figures

2.1	The Bloch sphere	18
3.1	Linear RF trap structure	44
3.2	Sideband cooling	46
3.3	Partial energy level diagram of Ca^+	47
3.4	Cirac and Zoller CNOT gate	50
3.5	Magnetic field insensitive qubit states	63
3.6	Microfabricated planar RF trap	66
3.7	Microfabricated array of RF traps	66
3.8	Heating rate measurements in RF traps	68
3.9	Variation of heating rate with endcap separation in a RF trap	69
4.1	Electrode structure of an ‘ideal’ Penning trap	76
4.2	Ion motion in the Penning trap	80
4.3	Schematic of the model for laser cooling in the Penning trap	82
4.4	Linear approximations used in the laser cooling model	85
4.5	Damping rates derived from the laser cooling model	90
4.6	Energy level diagram showing Penning trap motional states	93
4.7	Ion motion in the rotating frame	95
4.8	Motion of an axialised ion in the laboratory frame	96
4.9	Motional frequencies and damping rates when $(\epsilon/\omega_1)^2 \ll M^2$	105
4.10	Motional frequencies and damping rates when $(\epsilon/\omega_1)^2 = M^2$	106
4.11	Motional frequencies and damping rates when $(\epsilon/\omega_1)^2 \gg M^2$	108
4.12	Avoided crossing measured with a FT-ICR technique	110
4.13	Response of a weakly axialised ion to a dipole drive	114
4.14	Response of a strongly axialised ion to a dipole drive	115
4.15	Argand plots showing ion response to a dipole drive	116

LIST OF FIGURES

5.1	Partial energy level diagram of calcium at 0.98 tesla	121
5.2	Zeeman splittings of relevant transitions at 0.98 tesla	122
5.3	Experimental setup	124
5.4	Scale drawings of the split-ring trap	126
5.5	Ultra-high vacuum apparatus	128
5.6	Application of the quadrupolar and dipolar drives	130
5.7	Setup for operating as a RF trap	131
5.8	The Littrow configuration extended cavity diode laser	133
5.9	Spectrum analysis of the 866nm laser sidebands	136
5.10	Optical layout of the wavemeter	139
5.11	Spread in successive wavemeter measurements	141
5.12	Basic design of the optical reference cavities	144
5.13	Optical setup for laser locking	145
5.14	Block diagram of the laser locking electronics	146
5.15	Plots of optical cavity frequency drift	148
5.16	Diagram of the fluorescence detection systems	150
5.17	Fluorescence trace from a large cloud of ions	153
5.18	Fluorescence from ions during and after loading	154
6.1	Harmonic oscillator phase relative to a driving force	159
6.2	Experimental setup for RF–photon correlation	161
6.3	Example RF–photon correlation traces	163
6.4	Example motional phase measurements of an ion cloud	164
6.5	iCCD images of ion clouds with and without axialisation	167
6.6	Fitting of Gaussian functions to iCCD images	168
6.7	Variation of ion cloud aspect ratio with axialisation amplitude	169
6.8	Fluorescence rates as a function of axialisation amplitude	171
6.9	Motional phase measurements showing two resonances	172
6.10	Variation of magnetron cooling rate with axialisation amplitude	174
6.11	Variation of motional cooling rates with axialisation amplitude	175
6.12	Variation of magnetron cooling rate with axialisation frequency	176
6.13	Measurements of an avoided crossing	178
6.14	Motional frequency shifts as a function of axialisation amplitude	180
6.15	Motional frequency shifts with an off-resonant axialisation drive	181
6.16	Measured and predicted two-resonance motional phase plots	182

LIST OF FIGURES

7.1	Quantum jumps of ~ 5 ions with a low signal rate	189
7.2	Variation of RF trap signal rate with filament current	191
7.3	Quantum jumps of 1, 2 and 3 ions in a RF trap	192
7.4	iCCD images of a single ion in a RF trap	193
7.5	Fluorescence traces from ion clouds in a Penning trap	197
7.6	Variation of Penning trap signal rate with filament current	198
7.7	Ion temperatures as a function of filament current	199
7.8	Power broadening of the Doppler cooling transition	201
7.9	Variation of axialised ion signal rate with filament current	202
7.10	Quantum jumps of ~ 4 ions in a Penning trap	203
7.11	Experimental setup for mass spectrometry	206
7.12	Mass spectra with a filament bias of -30 volt	209
7.13	Driving of an ion species from the trap by resonant excitation	210
7.14	Mass spectra with a filament bias of -6.5 volt	211
7.15	Energy level diagram of a Λ system	213
7.16	Partial energy level diagram of Ca^+ at very low \mathbf{B} fields	214
7.17	Dark states in the Ca^+ system at very low \mathbf{B} fields	215
7.18	Observations of dark states in the barium system	217
7.19	Observation of dark states in calcium in the Penning trap	218
7.20	Energy level diagram of the modelled states	219
7.21	Steady-state populations as a function of laser detuning	224
7.22	Effect of repumper laser detuning on the cooling transition	225
7.23	Effect of cooling laser intensity on the cooling transition	227
7.24	Effect of repumper laser intensity on the cooling transition	228

List of Tables

5.1 Experimentally determined wavelengths for laser cooling	152
---	-----

Acknowledgments

Firstly I would like to express my very great thanks to my supervisors, Danny Segal and Richard Thompson. Besides leading this research, they have provided years of advice, encouragement and friendship. I would also like to thank all my colleagues of the last four years: Kingston, Jake, Eoin, Rafael, Hamid, Dan, Sebastián, Tim, Arno, Danyal and Bakry. Not only have they played a crucial role in this research, but they have made working in the group a great pleasure indeed.

For technical assistance I would like to thank Brian and Bandu, from the mechanical workshop and QOLS electronics section respectively. I gratefully acknowledge financial support from the Engineering and Physical Sciences Research Council.

On a personal level I am deeply grateful to my parents and to Andrew and Carli for their emotional and financial support over the years. Thanks to Lisa for being a great friend and housemate throughout the period of this research, and to Chloe and Mario for helping to keep me sane during the writing of this thesis. Finally I would like to thank all those other friends who have made the last four years so enjoyable. Though there is not the room here to give them all the credit they deserve, I am truly dependent on the friendship and support they continue to provide.

Chapter 1

Introduction

Over the last forty years or so, the development of popular computing has changed the way the world works. Dramatic and rapid though this revolution has been, it is based upon the physical implementation of ideas that were first developed back in the 1940s. These emerged from the study, in a mathematical framework, of the nature of algorithms and the efficiency with which they can be carried out — just one aspect of a field known as information science.

At about the same time that computers were beginning to find a place in the world at large, people had begun to realise that, since information is encoded in physical systems and manipulated by physical processes, it is fundamentally governed by physical laws. The development of quantum mechanics and relativity theory that had in the twentieth century so altered the face of physics therefore had a direct bearing on how we could store and manipulate information. It is this realisation that has led to the development of a new branch of physics, that of quantum information science.

At first glance it might seem that quantum information science looks at the limits imposed by the laws of physics on how we can process information. In fact, however, the reverse is also true — an understanding of the properties of information is helping to shed light on many problems in physics. As an example, the solution to the long-standing puzzle of why Maxwell's demon appears to violate the laws of thermodynamics requires an understanding of the laws governing the manipulation of information. Information, it seems, can be viewed as a fundamental resource and many of the laws of physics can be written in terms of the laws of information processing, rather than vice versa.

From classical information theory it has long been known that many computation tasks can *never* be carried out efficiently by a computer running a classical algorithm. Encryption techniques take advantage of this fact by making sure that it requires an extremely long time to perform the computation required for decoding, unless the ‘key’ to the code is known. Another key group of problems that falls into this category is the simulation of quantum systems. As the complexity of the system being modeled increases, the timescale of the computation rapidly becomes overwhelming. In the early 1980s, however, it was realised that a quantum system could be efficiently simulated using another, more isolated and easily manipulated, quantum system. It is not difficult to see that such a ‘quantum computer’ might also be able to solve some of the other problems that are ‘impossible’ to solve using classical computers.

That quantum information science may enable the simulation of quantum systems has captured the imagination of physicists around the world. That it might allow the breaking of codes (and, in fact, the development of truly secure encryption) has aroused the interest of government and commercial sponsors of research. It is ultimately this common interest from different communities that has created and allowed the widespread growth in quantum information science. As an experimental effort, research into quantum information processing is carried out in hundreds of laboratories around the world. The range of disciplines involved is also staggering — since much of modern experimental physics probes systems that are directly governed by the laws of quantum mechanics, it follows that any of these systems can be used as the basis of a quantum computer. The race has begun to find the system in which we can most easily perform quantum information processing.

In order to establish the suitability of a particular technology, it is necessary to know what the requirements for quantum information processing are. To this end I shall in chapter 2 provide an introduction to the basic principles of quantum computing.

Of the various candidates, it is trapped ion research that is currently leading the field of experimental quantum information processing. Here the internal states of electromagnetically trapped ions are used to store information, and laser manipulation enables us to carry out processing of this information. In chapter 3 I shall look at the techniques that have so far been developed to do this. I shall also provide a little background into the progress towards quantum information processing that has already been made using trapped ions.

Almost all the work so far carried out has been performed in radiofrequency traps, where oscillating electric fields are used to confine ions. Although the progress made using these traps has been extraordinary, the large trapping fields can lead to heating of the motion of the ions and to a degrading of the stored information called ‘decoherence’. By looking at the directions currently being taken by the radiofrequency ion trap community I shall highlight the fact that these problems are likely to become even more severe in the near future.

There is another type of ion trap called the Penning trap, which requires only static electric and magnetic fields for ion confinement. We believe that many of the problems associated with the radiofrequency trap can be eliminated or eased through the use of these traps. It is for this reason that the Ion Trapping Group at Imperial College has initiated a project to measure the heating and decoherence rates of ions stored in Penning traps. Since this information already exists for ions stored in radiofrequency traps, measurements from the Penning trap would enable us to evaluate the benefits of performing quantum information processing using ions in Penning traps.

There are, however, different problems associated with the use of Penning traps. For the most part these relate to difficulties in efficiently laser cooling ions stored in them. It is the development of techniques to overcome this difficulty that forms the main focus of this thesis. In chapter 4 I shall discuss in some detail how the Penning trap works and why it leads to less efficient laser cooling. I will introduce a technique called ‘axialisation’ that can overcome this problem. I shall describe an analytical model for laser cooling in the Penning trap that highlights the cause of the inefficiency, and extend it to show the improvement that arises from axialisation.

The experimental setup that we have developed for laser cooling and axialisation in the Penning trap will be described in chapter 5. This will lead to a presentation in chapter 6 of the results of experiments to measure the improvement in laser cooling that occurs as a result of axialisation. Other interesting effects that axialisation has on the dynamics of the trapped ions will also be discussed, and comparisons between our model and experimental results will be made.

In order to achieve our aim of measuring heating and decoherence rates of ions, it will be necessary to isolate individual ions in the Penning trap. In chapter 7 I will look at current experimental progress towards this aim, and

CHAPTER 1: INTRODUCTION

the problems we have encountered. I shall describe the attempts we have made to understand these difficulties, both through experiment and modeling of the system.

Finally in chapter 8 I shall discuss the current state of the experiment in terms of the overall aim of measuring heating and decoherence rates of ions in a Penning trap. Much of this discussion will focus on modifications to the experimental setup that are currently being implemented in an attempt to enable us to resolve some of the difficulties we have faced thus far.

Chapter 2

Quantum information processing

It has been known for some twenty years now that quantum mechanical systems have the potential to form the basis of a computational system. The initial ideas were closely related to the development within classical information science of reversible computation. This provided a model for computation where the input (and indeed the state of the system at any point during the computation) could be derived exactly from the output of the system. The original impetus behind these developments was the idea that energy dissipation is not necessarily required by reversible logic gates, ie. gates where there is a one-to-one correspondence between input and output. However these models also turn out to describe the form of computation that is possible with a quantum mechanical system — within quantum theory transformations of a system are described by operators which are unitary and therefore reversible.

In 1980 Toffoli introduced the Controlled-NOT (or CNOT) gate — a gate which has proved very important in the development of quantum computation. He also showed that there exist reversible three-bit gates that are universal for computation, ie. finite numbers of the same gate can be combined to produce networks equivalent to any other (reversible) gate. It was at this point that Benioff [1] and later Feynman suggested that bits represented by states of a quantum mechanical system could evolve under the action of quantum operators to give reversible computation. Such a system was shown to be capable of computation efficiency at least equivalent to that of a classical computer. In 1985 Deutsch [2] published what is generally regarded to be the

first complete prescription for a quantum computer (QC) and showed that it can perform certain tasks faster than any classical machine.

At the same time that these theoretical developments of quantum computing were taking place, classical computers were just beginning to take up their role as the ubiquitous and versatile devices they are today. The rate of increase in their speed is staggering — computing power has more or less doubled every eighteen months and shows no imminent signs of slowing. One might argue that this success indicates that our efforts would be better focused on continuing the development of classical computers than pursuing an entirely distinct computational paradigm. There are, however, two reasons to believe that non-classical computational systems will be of use in the future.

Firstly, increase in classical processing power has been achieved largely through miniaturisation. As devices become smaller and smaller we will eventually reach a scale where quantum effects become significant. This appears to place a fundamental limit on the speed of an entirely classical computer. As we approach this limit it would seem wise to investigate the possibility of using these quantum effects to our advantage.

Secondly we have the fact that large classes of problems are inherently insoluble on a classical computer. More accurately, we find that the number of computational steps required to solve these problems increases exponentially with the size of the input. As an example, let us look at the problem of factorising large numbers. Using the Menezes sieve method we have to carry out $\sim e^{2L^{1/3}(\ln L)^{2/3}}$ computational steps, where L is the natural logarithm of the number being factorised. Thus a 75 digit number would take about 10^{14} steps to solve — about a days work for a modern PC. If we double the number of digits, however, the number of steps required is closer to 10^{20} and the same PC would be kept busy for more than 600 years. Algorithms such as these are termed ‘exponential-time’ and rapidly become effectively impossible to solve as the input size is increased.

The relevance of this for quantum computing is that there are some problems for which only exponential-time algorithms exist in classical computing, but better scaling ‘polynomial-time’ algorithms exist for a computer based on a quantum mechanical system. So far there are few such problems and their usefulness is generally limited, but the fact that there are any at all suggests there may be more to be discovered.

In this chapter I shall begin by introducing the qubit and looking at the

features of quantum mechanical systems that enable computers based on them to carry out tasks more efficiently than classical machines. I shall then describe how such a system can be used as the basis of a universal quantum computer. At this stage I shall look at some of the problems that can be solved more efficiently, using Shor's factorisation algorithm as an example of a quantum algorithm. Next the basic ideas of quantum error correction (QEC) and how they can enable us to carry out quantum computing under imperfect conditions will be discussed. Finally I shall look at the properties that a real quantum system must possess if it is to form the basis of a quantum computer.

2.1 Properties of quantum systems

2.1.1 Quantum bits

All real quantum mechanical systems are complicated, with infinite numbers of potentially available states that will differ from system to system. In order to develop a general quantum information theory it is necessary to *define* some fundamental sub-system that can be built up to form a larger system suitable for use as a quantum computer. Besides being useful from a theoretical point of view, such a sub-system must be physically realisable and must not restrict the computational power that can be achieved. Perhaps unsurprisingly the definition chosen is the 'qubit' — a two-level system, which therefore has a general state

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle. \quad (2.1)$$

The above state represents a 'coherent' superposition in that there is some measurement basis in which it has a single, well-defined value. For example in the state where $\alpha, \beta = 1/\sqrt{2}$, a measurement in the basis at 45° to the computational basis will always yield the same value.

The state of any qubit can be written in the form

$$|\psi\rangle = e^{i\gamma} \left(\cos \frac{\theta}{2} |0\rangle + e^{i\varphi} \sin \frac{\theta}{2} |1\rangle \right). \quad (2.2)$$

The overall phase factor $e^{i\gamma}$ is immeasurable and has no effect on the evolution of the qubit. Thus we can effectively specify the state of a qubit with just the values of θ and φ and these describe a point on the surface of the 'Bloch sphere' (figure 2.1). The poles of this sphere are by convention taken to be the




Figure 2.1: The state of a qubit can be represented as a point on the surface of the Bloch sphere (taken from ref. [3]).

computational basis states $|0\rangle$ and $|1\rangle$, but clearly we could choose any other axis about which to take θ and φ . This represents a change of basis, and a measurement in a given basis projects the state vector onto one or other of the points where this axis intersects the sphere. The idea of a coherent state as being one which has a definite value in some particular basis is emphasised by the Bloch representation — whatever the state of a qubit we can always choose an axis parallel to its state vector. We will return to the Bloch sphere representation when we discuss quantum gates in the next chapter.

The ability of qubits to exist as a coherent superposition of states is the key difference between quantum and classical bits. Whereas a classical bit can exist in either of two states, there are an infinite number of possible values for α and β in equation 2.1 that satisfy the normalisation condition. Though it cannot be directly accessed there is in some sense an infinite capacity to store information within a single qubit. Upon measurement this ‘hidden information’ disappears and we are left only with one of two results. As we shall see, however, during evolution of a register of many qubits under some transformation this information is ‘used’ to determine the final state of the system. Thus we can make use of this storage capacity during a computation,

but can only ever find a single state during measurement. The challenge of quantum computing is to engineer a computation such that the particular state observed at the end tells us something about the solution to the problem we are interested in.

That the choice of the qubit does not reduce the power of a quantum computer compared to one comprising of, for example, a series of multiple-level systems ('qudits') can be shown by analogy with the use of bits in classical computing when we could quite conceivably build a computer that used decimal digits instead. We would require less of these 'decits' to specify an input number than we would bits and so we could reduce the number of computational steps required to solve a problem. However the number of digits is reduced by a constant factor (in this case $1/\log_{10} 2$) so the complexity class of a given problem remains the same. This analogy breaks down if we consider systems that have an infinite number of available states and there have been proposals for computational schemes based on these. It is not yet fully understood whether such schemes can be used to solve problems that would otherwise be intractable. It is worth noting that some real systems will actually be more suited to multi-digit quantum computation and such systems are being investigated experimentally. It would be fair to say, however, that most of the current research in this field is based on two-level systems and it is perhaps easier to convey the basic theoretical principles if we restrict ourselves to a discussion of such systems.

2.1.2 Multi-qubit systems

We can now combine an arbitrary number of qubits to build up a quantum system of the form $|\psi_1\rangle \otimes |\psi_2\rangle \otimes \dots \otimes |\psi_n\rangle$ (henceforth we shall use the abbreviated notation $|\psi_1\psi_2\dots\psi_n\rangle$ for such systems). The general form of a combined system is also a superposition of states (with appropriate normalisation):

$$|\psi\rangle = \sum_{x \in \{0,1\}^n} c_{x_1x_2\dots x_n} |x_1x_2\dots x_n\rangle. \quad (2.3)$$

Such a quantum state is specified by 2^n complex coefficients, a huge number for all but very small n . This is compared to a classical n -bit register, for which we obviously require only n binary numbers.

These multi-qubit superpositions have some very useful properties. If we

carry out some operation on a subset of qubits, *all* the possible states are affected simultaneously. This effect is known as ‘quantum parallelism’ and is one of the resources that enables quantum computers to be more efficient than their classical counterparts. For example if we have a system of n qubits in a superposition of all different states, we essentially have a variable with 2^n possible values. If we then carry out an operation on those qubits equivalent to applying some function, the final state of the system will be a superposition of the values of the function for each initial state. We have carried out 2^n calculations simultaneously!

Unfortunately quantum parallelism is not quite as miraculous as it first appears. The problem is that as soon as we try and measure the state of the system the superposition collapses and we are left with the value of the function for some random input value — a useless piece of information. Fortunately we can make use of another important feature of quantum systems — quantum interference. In the above example there may be several different initial states which lead to a given final state, or value of the function. These will all interfere and the probability of measuring a particular output value will depend on the sum of the coefficients for all the states yielding that value. By carefully designing the sequence of operations we apply to our qubits, we can create a situation where only final states corresponding to the desired solution to a problem interfere constructively. Then through repeated computations we can find the correct solution with arbitrarily large probability. We will look at examples of such algorithms in section 2.2.

It is possible to view the concepts of quantum parallelism and interference described above in terms of ‘entanglement’, which can in some sense be regarded as the truly fundamental resource used in quantum information processing. In essence, a system of qubits possesses entanglement if the wavefunction describing it (equation 2.3) cannot be expressed as a product of all the individual qubits. Entanglement is by no means well understood, and there is currently a great deal of theoretical work geared towards determining why it proves to be so essential for quantum information processing.

2.1.3 Decoherence

Though by no means a desired feature of quantum systems, a discussion of decoherence is relevant here since it is present in all real systems and is one

of the primary difficulties in implementing quantum computation. Indeed decoherence is not so much an attribute of quantum systems as a process by which the features utilised in quantum computation are lost. It is considered to be in some sense the mechanism by which classical physics emerges from quantum mechanics when we look at systems on a broader scale.

The term decoherence refers specifically to processes that transform a system from a coherent superposition to an incoherent mixed state. In fact we often use the term to describe any uncontrolled transformation of the system so that upon completion of some sequence of gate operations the system is not necessarily left in the intended state. Decoherence can therefore stem from non-ideal gate operations and unintended coupling between qubits, but usually when we talk about decoherence we refer to the coupling of a quantum system to its environment. To demonstrate this let us look at the interaction between a single qubit and the environment in some initial state $|E\rangle$:

$$(\alpha |0\rangle + \beta |1\rangle) |E\rangle \longrightarrow \alpha |0\rangle |E_0\rangle + \beta |1\rangle |E_1\rangle. \quad (2.4)$$

Essentially we now have entanglement between the qubit and the environment, which has therefore acquired information on the state of the qubit. The density matrix of the qubit has become

$$\rho = \begin{bmatrix} |\alpha|^2 & \alpha\beta^* \langle E_1 | E_0 \rangle \\ \beta\alpha^* \langle E_0 | E_1 \rangle & |\beta|^2 \end{bmatrix}.$$

More accurately $|E_0\rangle$ and $|E_1\rangle$ are functions of time, becoming more orthogonal as the interaction continues and therefore possessing more information about the state of the system. We can write this as

$$\langle E_0(t) | E_1(t) \rangle = e^{-\Gamma(t)}, \quad (2.5)$$

where the form of $\Gamma(t)$ depends on the coupling between the qubit and the environment. Clearly as the decoherence process continues the off-diagonal elements tend to zero and our coherent superposition has become an incoherent mixture. The characteristic timescale for this to occur varies hugely between systems and is one of the key factors when considering the potential of different systems for quantum computing.

The extent to which the decoherence time scales with the size of the system

is important for determining the scalability of a quantum computer based on that system. Various models have been put forward to calculate this. One general approach [4] models the environment as a collection of harmonic oscillators and uses the following interaction Hamiltonian:

$$H = \frac{1}{2} \sum_i \sigma_{z,i} \omega_0 + \sum_{\mathbf{k}} b_{\mathbf{k}}^\dagger b_{\mathbf{k}} \omega_k + \sum_{i,\mathbf{k}} \sigma_z (g_{i,\mathbf{k}} b_{\mathbf{k}}^\dagger + g_{i,\mathbf{k}}^* b_{\mathbf{k}}).$$

Here ω_k , $b_{\mathbf{k}}^\dagger$ and $b_{\mathbf{k}}$ are the frequency, creation and annihilation operators of a mode \mathbf{k} of the collection of harmonic operators. The first term represents the free evolution of the system of qubits, the second that of the environment and the third is the interaction between the two. The results of this kind of model depend largely on the extent to which fluctuations of the harmonic oscillators are correlated. Where the lengthscale over which the fluctuations are correlated is much smaller than the size of the qubit system, each of the n qubits will interact separately with the environment. The most affected elements of the density matrix of the system therefore decay as the product of n separate qubit decays

$$\langle E_0 | E_1 \rangle^n = e^{-n\Gamma(t)}.$$

In the opposite case where the harmonic oscillator fluctuations are correlated over the entire length of the qubit register, all the qubits interact collectively with the environment. The interaction of each qubit then depends on the total number of qubits in the register. Overall the decay of coherence thus scales as $e^{-n^2\Gamma(t)}$. In both cases we find that the decoherence time is rapidly reduced as we increase the size of a system. This represents one of the greatest difficulties that must be overcome if we are to successfully carry out quantum computation. In section 2.3 we shall be looking at quantum error correction, a technique that may enable us to reduce the sensitivity of a quantum computer to decoherence.

2.2 Quantum algorithms

2.2.1 Universal computers

Any algorithm, be it classical or quantum, is carried out by applying some sequence of operations on a system that results in the system ending up in some desired state. In classical computing this is done using transistor networks, or ‘gates’, which represent various (classical) logical operations, but these operations can be thought of as theoretical devices independent of any real manifestation. A logic gate takes some number of bits and, depending on the values of these, sets a number of (possibly the same) bits in a particular state. A classical algorithm is essentially a list of logical operations and the bits to which they must be applied.

The above description implies that, since there are an infinite number of possible logic gates and a physical computer cannot possess them all, any particular computer can only carry out some subset of the possible algorithms — those for which it possesses the appropriate gates. This problem was addressed in the early days of information theory, leading to the proposal of the theoretical universal computer, or ‘Turing machine’. The reasoning is that if we can build a computer capable of efficiently simulating (ie. in polynomial-time) the processes of any other possible computer, we can effectively carry out any possible algorithm. This argument was embodied by the Church–Turing hypothesis:

Every function which would naturally be regarded as computable can be computed by the universal Turing machine.

Digital computers constitute a physical representation of the theoretical Turing machine. Their universality stems from their ability to simulate any possible logical operation and therefore to carry out any possible algorithm. It is here that we encounter the idea of ‘universal gates’ — a gate, or set of gates, that can be strung together to simulate any other gate. In classical computing it can be shown that any n -bit logic gate can be built up from a number of NAND gates. Such gates return a single output value ‘0’ iff both of two inputs have value ‘1’.

2.2.2 Quantum gates

The idea of gates carries over into quantum computing. Just like their classical equivalents, we think of n -bit quantum gates as being operations that take the state of a set of n qubits and transform it into some other state. Within the rules of quantum mechanics the only restriction on possible operations, and therefore quantum gates, is that they must be unitary. ie. $U^\dagger U = I$, where I is the identity operator. This is equivalent to saying that the gate must be reversible such that no information is lost during its application. The unitarity restriction implies that large classes of classical gates, including the universal NAND gate, do not have quantum equivalents. For example no classical gate with fewer outputs than inputs can be used in quantum computing.

The fundamental difference between classical and quantum bits is that whereas classical bits either have value ‘0’ or ‘1’, qubits can exist in any superposition of these. This is a significant complication since even for a single qubit there are an infinite number of possible inputs and outputs and therefore possible quantum gates. Returning to the Bloch representation, a single-qubit gate alters the direction of a qubit’s state vector in some way dependent on its original direction. As an example, consider the Pauli X operator

$$X = |1\rangle\langle 0| + |0\rangle\langle 1|. \quad (2.6)$$

Clearly this corresponds to the quantum equivalent of the NOT gate and represents a reflection about the equator of the Bloch sphere in the computational basis. More generally it is a 180° rotation about the x axis (see fig. 2.1). Another important single-qubit gate is the Pauli Z operator, which rotates the state vector by 180° about the z axis and thereby flips the phase of the qubit state:

$$Z = |0\rangle\langle 0| - |1\rangle\langle 1|. \quad (2.7)$$

An important result is that by exponentiating the X , Z and $Y(\equiv iXZ)$ we can form general ‘rotation operators’, that carry out rotations through arbitrary angles in the same sense as X , Y and Z . These can be used together to form any single-qubit operation.

The concept of universality applies equally to quantum computers. In order to build a computer capable of carrying out any possible quantum algorithm, we must be able to efficiently simulate any possible quantum gate with arbi-

trary precision. It might seem unlikely that this should be possible, since the set of possible operations is infinite. In fact a universal three-qubit gate was proposed by David Deutsch in 1989 [5] and since then a number of two-qubit gates have been found to form universal sets when combined with single-qubit rotations. An even more surprising result emerged in 1995 when it was shown that the only quantum gates that *aren't* universal are the single-qubit gates and the classical gates [6, 7].

Probably the most important universal set, from the point of view of ion trap QIP at least, is that comprising of single-qubit rotations and the CNOT gate. This gate carries out the quantum NOT operation on a target qubit iff a control qubit is in the $|1\rangle$ state. The significance of this stems from the fact that it was the first such two-qubit gate for which a physical means of implementation was suggested [8]. We shall look in more detail at this scheme when we look at the implementation of quantum computation in chapter 3.

2.2.3 Shor's quantum factoring algorithm

In 1994 Peter Shor published a complete algorithm for factorising integers in polynomial-time [9]. The difficulty of this task is historically very important, since it provides the security for a large class of very widely used cryptographic codes. The method uses quantum computing to achieve an exponential speedup compared to the best classical algorithm, relying on classical computation to carry out only trivial, polynomial-time tasks.

Shor's algorithm is significant not only because it provides a means for solving an important problem so efficiently, but because a number of subsequent algorithms for different tasks take a very similar approach. The algorithm works by reducing the problem to one of finding the periodicity of a function — a task at which quantum computers seem to be very efficient. It does this by selecting a random integer, a , smaller than the number to be factorised, N . Either a is a factor of N (in which case we have a solution already) or, by Euler's theorem, there is a power of a that has remainder 1 when divided by N :

$$a^r \equiv 1 \pmod{N}. \tag{2.8}$$

If we find r to be odd the algorithm has failed. If r is even we can rewrite

equation 2.8 as a difference of squares:

$$(a^{r/2} - 1)(a^{r/2} + 1) \equiv 0 \pmod{N}.$$

Providing neither of these terms is a multiple of N (in which case the algorithm again fails) the greatest common divisor of N with each will be a factor of N . Shor showed that for random a there is a better than 50% chance of success using the above technique.

The part of the algorithm that requires quantum techniques is the determination of r in equation 2.8. This can be reduced to a period-finding problem, since

$$f(x) = a^x \pmod{N} = a^x a^r \pmod{N} = a^{x+r} \pmod{N}.$$

ie. $f(x)$ is periodic with period r . We divide our qubit register into two parts, each of size $\log N$. The first we set into an equal superposition of states by applying the Hadamard operator ($H = 1/\sqrt{2}[|0\rangle\langle 0| + |0\rangle\langle 1| + |1\rangle\langle 0| - |1\rangle\langle 1|]$) to each qubit in turn. We then apply a sequence of gates to project the second register onto the value of $f(x)$ for our a and N , taking advantage of the quantum parallelism effect. Thus we have the final state

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle |f(x)\rangle.$$

A measurement of the second register yields a random value y_0 and through entanglement collapses the first register into a superposition of all the states $|x\rangle$ that satisfy $f(x) = y_0$:

$$|\psi\rangle = \frac{1}{\sqrt{N/r}} \sum_{k=0}^{(N/r)-1} |x_0 + kr\rangle |y_0\rangle.$$

In order to extract the periodicity from the first register we make use of the quantum analogue of the discrete Fourier transform, which takes a state $|a\rangle$ into a state

$$\frac{1}{\sqrt{N}} \sum_{b=0}^{N-1} |b\rangle e^{2\pi i \frac{ab}{N}}.$$

This transform is achieved through a series of applications of the Hadamard

operator and a second, two-qubit operator

$$S_{j,k} = |00\rangle\langle 00| + |01\rangle\langle 01| + |10\rangle\langle 10| + e^{\frac{i\pi}{2^{k-j}}} |11\rangle\langle 11|,$$

where j and k identify the qubits being operated on. The overall transformed state of the first register is then

$$\mathcal{F}|\psi\rangle = \frac{1}{r} \sum_{j=0}^{r-1} e^{2\pi i \frac{x_0 j}{r}} |j \frac{N}{r}\rangle.$$

It is important to note that this transform can be carried out in a number of steps of order $(\log N)^2$ [10], so can be carried out in polynomial-time. A measurement will now yield a random multiple of N/r , so dividing by N we have

$$\frac{c}{N} = \frac{\lambda}{r}.$$

Now assuming that λ and r have no common factors we can determine r by cancelling down c/N . It has been shown that the chance of λ and r being coprime is greater than $1/\log N$, so we can find r with arbitrarily high probability by repeating the process of order $\log N$ times.

Thus we have a method for factorising large integers for which every step is manageable in polynomial-time. It has never been proven that a classical algorithm cannot achieve this, so we cannot say that quantum computing necessarily provides an exponential speedup for this problem. On the other hand, the use of cryptographic methods that rely on this problem being intractable implies that the existence of a classical polynomial-time algorithm is deemed unlikely.

2.2.4 Other quantum algorithms

Besides Shor's method of factorisation, there are a handful of other important quantum algorithms. Perhaps the most famous is Grover's search algorithm, which provides a quadratic speedup over classical methods for searching an unstructured list. This provides a useful tool for the NP class of problems, where testing solutions can be carried out in polynomial-time but finding them cannot.

Perhaps the most important use of quantum computing from the physicist's point of view is in the simulation of quantum systems. To simulate a state

vector in a 2^n Hilbert space on a classical machine requires manipulation of 2^n -component vectors — a huge quantity of data for all but the smallest quantum systems. A quantum computer is much more efficient, requiring n qubits to specify the same system. Simulation of the evolution of quantum systems is generally still a hard computation, since we will normally require an exponentially large number of gates. There are, however, some systems that we would expect to be able to model efficiently.

There are still relatively few quantum algorithms that are more efficient than the best classical techniques. It remains to be seen whether this class of algorithms is fundamentally limited, or whether there are still significantly more to be found. Either way, the path towards implementation of these algorithms is bound to provide a better understanding of the physical processes at work in quantum computation.

2.3 Quantum error correction

Thus far we have only considered quantum computing in an ideal system, free from decoherence and quantum gate imperfections. In any real system we will always have to contend with these and other difficulties caused by, for example, the presence of unwanted additional states for the qubits. We must therefore address the issue of how sensitive our quantum computer will be to these problems.

One way to approach imperfections in the system is to think of them as some extra transformation applied to the system following ‘ideal’ evolution. We can represent the state of a system of n qubits by a vector in a 2^n dimensional Hilbert space. The purpose of a quantum algorithm is ultimately to project the system into some particular desired state or states, the measurement of which tells us something about the solution to a given problem. As the state vector is rotated under the influence of the transformations that make up the algorithm, decoherence and errors in quantum gates will cause slight additional rotations and projections such that when the algorithm is complete the system may be in a different superposition of states. Thus when we read out the result of our computation there will be a non-zero probability that the state we measure will not be one corresponding to a solution to our problem.

The concern from the point of view of effective quantum computing is that these effects will in general be cumulative, such that the overall likelihood of an error in the final result will scale with the size of the computation. Using a quantum algorithm to solve some large problem might require the application of many billions of quantum gates and would take a timescale long in comparison with the decoherence time of the qubits. Thus we might expect that even with very precise, rapid transformations quantum computing could never accurately tackle problems complex enough to be intractable with classical computing.

It has been pointed out [11] that these difficulties are not dissimilar from those faced and overcome in classical computing. Bits stored in classical memory gradually relax to random values over time in a process similar in some ways to that of decoherence and any computer is constantly reading and rewriting its entire memory. Since direct cloning of quantum states is impossible it was generally believed that such methods of preventing errors were forbidden by the principles of quantum mechanics. In 1995–6, however, a number of papers were independently published describing potential methods to apply state corrections to quantum systems [12, 13]. In the years since, a number of further advances have been made and it is generally accepted that QEC is not only permitted by the rules of quantum mechanics, but will ultimately be required if we are to build a complete operational quantum computer.

Since many of the methods devised for quantum error correction are analogous to techniques routinely used in classical computing, we shall begin by reviewing some of the key features of classical error correction [11, 14].

2.3.1 Classical error correction

We begin by representing our k -bit register (a binary ‘word’) as a k -component vector in a space where operations are carried out modulo 2. We can then write the effect of noise on a word as

$$u \longrightarrow u' = u + e, \tag{2.9}$$

where e is the error vector indicating which bits were flipped by the noise. We then seek a set of codewords C such that

$$u + e \neq v + f \quad \forall u, v \in C (u \neq v) \quad \forall e, f \in E,$$

where E is the set of correctable errors (which should include the case of no error at all).

C is chosen to be a set of n -bit codewords (where $n > k$) in which every word differs from every other in at least d places. d represents the ‘minimum distance’ of the code and for random noise a code can correct any errors affecting less than $d/2$ bits. k/n is the ‘rate’ of the code, since it indicates how many encoded bits must be used for each original bit of information. When specifying the parameters of a code we often use the notation $[n, k, d]$. It can be shown that the maximum number of members of the correctable error set E is 2^{n-k} , so the more additional bits used to encode the information the more errors can potentially be corrected. This ultimately leads us to Shannon’s theorem for data transmission, which states that however noisy the channel we can find a binary code allowing transmission with arbitrarily small error probability, providing k/n is sufficiently small and n sufficiently large.

The goal when generating codes is to maximise both the distance d and the rate k/n . Since these two parameters are incompatible we must find an acceptable compromise. Often the noise we are dealing with is not entirely random and certain errors are more likely than others. In such cases we can tailor the code so that the set of correctable errors E is the set of most likely errors. This enables us to achieve a larger rate whilst maintaining the same insensitivity to errors.

The encoding method we have looked at so far is not directly applicable to a quantum information processing system since error correction involves measuring the actual states of the bits, comparing them with the set of codewords and so determining what they were most likely to have been before the action of the noise. Quantum mechanically we cannot do this without disrupting the computation, since the measurement process itself will project the system into a particular state.

Fortunately there is a modification of the above technique that is of use to us — a process known as parity checking. We define the ‘weight’ of a binary number to be the total number of ‘1’s it contains and if its weight is even we say it has even parity. Thus we can determine parity directly by adding all the components of a vector together modulo 2. Now we look for a set of k n -bit codewords for which there are $n - k$ linearly independent subsets all of which have even parity for every codeword. If we measure the parity of a subset of bits from a received word to be odd then we know there have been

an odd number of errors within that subset. The parities of all $n - k$ subsets taken together completely specify which errors have occurred (providing we have received a member of the set of correctable errors).

If we represent the bits included in each subset by the ‘1’s in a ‘parity check vector’, h , then we can calculate the parity of that subset for a particular codeword by taking the scalar product of the two vectors modulo 2. For example the parity of the subset comprising the last four bits of the seven-bit word (1101001) is even, since

$$(0001111).(1101001) = 2 \bmod 2 = 0.$$

So for the even parity code described above we have the property that

$$Hu^T = 0,$$

for all codewords u , where H is the ‘parity check matrix’ made up of all the check vectors h . Now since we can write any noise-affected word as a sum of the original word and an error vector e we have the property for our code that

$$H(u + e)^T = Hu^T + He^T = He^T.$$

The vector He^T is the ‘error syndrome’ and there is a one-to-one correspondence between the set of possible vectors and the set of correctable errors.

The parity checking technique has one extraordinary property that is of great use to us for quantum error correction; it enables us to determine which of the correctable errors has occurred without giving us any information about the codeword we have received. In theory we can therefore ascertain which corrective transformation we need to apply to our quantum system without affecting the computation itself.

2.3.2 Principles of QEC

We have seen that the sources of errors in quantum information processing are decoherence due to coupling of the system to the environment and the imprecision of quantum gate operations. These can be thought of as ideal evolution and transformations followed by additional processes due to the noise on some or all of the qubits. The interaction of a quantum system with

the environment can be written as

$$|E\rangle |\phi\rangle \longrightarrow \sum_s |E_s\rangle M_s |\phi\rangle, \quad (2.10)$$

where $|E\rangle$ is the initial state of the environment and each M_s is some unknown unitary operator. The purpose of QEC is to restore the system to the state $|\phi\rangle$ without gaining any information about the system.

We begin by coupling the main system q to an ancillary system a of qubits prepared in a state $|0\rangle_a$. The interaction A between q and a is arranged to proceed as

$$A(|0\rangle_a M_s |\phi\rangle) = |s\rangle_a M_s |\phi\rangle.$$

Here all the states $|s\rangle_a$ corresponding to different error operators M_s are orthonormal and are not dependent on the state $|\phi\rangle$, so there is a one-to-one correspondence between $|s\rangle_a$ and M_s . Clearly an interaction with these properties will only be possible for certain error operators. This set of error operators S is the quantum analogue of the set of correctable errors in classical error correction.

After application of the ‘syndrome extraction operator’ A the overall state will be

$$A \sum_s |E_s\rangle |0\rangle_a M_s |\phi\rangle = \sum_s |E_s\rangle |s\rangle_a M_s |\phi\rangle.$$

By now measuring the ancilla in the $\{|s\rangle_a\}$ basis we find the value of s and project the entire system into one particular state $|E_s\rangle |s\rangle_a M_s |\phi\rangle$. We can now deduce M_s from s and apply its inverse to obtain the final state

$$|E_s\rangle |s\rangle_a |\phi\rangle. \quad (2.11)$$

We can therefore recover our original state providing the error operator is indeed a member of the correctable set S .

2.3.3 Application of QEC

The challenge of QEC is to find sets of recoverable states $|\phi\rangle$ and the syndrome extraction operator A corresponding to the dominant noise terms in a given quantum system. Such a set of recoverable states forms a subspace of the overall Hilbert space of a system. QEC can be thought of as a projection of

the state of the system back onto this subspace. From analogy to classical error correction the recoverable states are called quantum codewords. Just like their classical counterparts quantum codewords must obey the property that two words must remain distinguishable after each has been acted on by correctable noise. ie.

$$\langle u | M_i^\dagger M_j | v \rangle = 0.$$

In order to describe real quantum noise in the form of equation 2.10 we consider a general interaction between a single qubit and the environment:

$$\begin{aligned} |E\rangle |\phi\rangle &= |E\rangle (a|0\rangle + b|1\rangle) \\ &\longrightarrow a(c_{00}|E_{00}\rangle|0\rangle + c_{01}|E_{01}\rangle|1\rangle) \\ &\quad + b(c_{10}|E_{10}\rangle|0\rangle + c_{11}|E_{11}\rangle|1\rangle). \end{aligned}$$

This can be written in terms of the operators I , X , Y and Z where X corresponds to a bit-flip error, Z to a phase error and Y to both errors together. We obtain

$$|E\rangle |\phi\rangle \longrightarrow (|E_I\rangle I + |E_X\rangle X + |E_Y\rangle Y + |E_Z\rangle Z) |\phi\rangle. \quad (2.12)$$

Thus we need only be able to correct bit-flip errors, phase errors and the two combined to correct any general quantum noise. The overall error operators for a quantum system, M_s are therefore simply the product of single-qubit operators from the set $\{I, X, Y, Z\}$.

Syndrome extraction for purely bit-flip errors is relatively straightforward, since the errors are essentially the same as those faced in the classical case. We need operators corresponding to the relevant parity checks for the code. Each of these can be accomplished by a series of CNOT gates in which the relevant qubits of the system are the controls and an ancillary qubit is the target. For every control qubit in the '1' state the ancillary qubit is flipped, so if the final ancillary qubit state is '0' the system has even parity. All the parity checks taken together identify which subset of qubits has been flipped (providing the flipping of this subset is one of the correctable errors) and we can apply X to these to restore the system to its correct state.

Correction of phase errors takes place in a similar manner. We can use the fact that $Z = HXH$, where H is the Hadamard operator representing a change to the basis $\{|0\rangle + |1\rangle, |0\rangle - |1\rangle\}$. Thus we can use the same technique

as for bit-flip errors, providing we apply a Hadamard rotation before and after the syndrome extraction process.

To eliminate the effects of general quantum noise we need a code that is able to correct both bit-flip and phase errors. This reduces to the problem of finding a classical error-correcting code which transforms under a Hadamard rotation to another classical error-correcting code. Thus a general recoverable state for full quantum error correction would be of the form

$$|\phi\rangle = \sum_{u \in C} |u\rangle = H \sum_{v \in C'} |v\rangle,$$

where C and C' are classical error correcting codes. Note that it has been shown [15] that quantum error correction can be carried out even if the gate operations that form the correction process are themselves imprecise.

2.3.4 Requirements of QEC

The demands of full quantum error correction are truly daunting and it seems likely that a practical demonstration will not be seen for many years. The reason for this is simply that the correction process involves many times more qubits and quantum gate operations than the computation itself, so the scale of the computer needs to be significantly larger. As an example one of the codes so far discovered for QEC is the [55,1,11] Calderbank, Shor and Steane code, which uses 55 qubits to store a single logical qubit. The syndrome extraction operation for this code requires 660 quantum gates to correct this single logical qubit.

Despite its demanding requirements, QEC will necessarily play a crucial role if a quantum computer is ever to solve problems intractable classically. Taking the factorisation of large integers as an example, networks of classical computers can currently manage numbers with about 130 digits in a reasonable time. A quantum computer running Shor's algorithm might require 1000 qubits and 10^{10} quantum gates to solve this same problem. In the absence of QEC we would therefore need a noise level of the order of 10^{-13} per qubit per gate — an almost certainly impossible requirement. With QEC we might need (using the code described above) 55000 qubits and $\sim 10^{13}$ gates, but the acceptable noise level can be estimated [16] to be of the order of 10^{-5} per qubit per gate.

The critical point is that in the absence of quantum error correction, the acceptable noise level per qubit per gate falls as the scale of the problem increases. With QEC this is not the case. We find that there is a fault-tolerance threshold requirement and if we can satisfy this an increase in the number of gate operations no longer poses a problem. Many quantum error-correcting codes have now been proposed, each requiring different resources and providing protection against different amounts and types of error. Estimates of the fault-tolerance threshold vary, but are most commonly in the range 10^{-3} – 10^{-5} .

2.4 Criteria for the implementation of QIP

So far we have seen that there are features of quantum mechanical systems that can in theory be exploited for computation and we have looked at various algorithms that can make use of these. These ideas were based on the assumption that we could build up a quantum computer from simple two-level systems and could apply certain transformations to them. In reality most systems are extremely complicated, with many available states and interactions with the environment. Clearly in order to physically implement quantum computing schemes we will need a system that possesses certain crucial properties. There has been much discussion of these requirements within the quantum computing community, most notably reviewed by David DiVincenzo [17], and it is his conclusions that we consider here.

In general, the criteria for quantum computing are not dissimilar from those of classical computing:

1. We require a scalable system with well defined qubits on which to carry out the computation. This has obvious parallels with the classical bits of a conventional computer.
2. We must be able to prepare the system in some initial state with a high degree of certainty.
3. The system must be sufficiently decoupled from the environment to allow the qubits to evolve coherently throughout the computation. This is comparable to the classical computing requirement that the timescale for noise to cause unwanted flipping of bits must be much greater than that taken to carry out the computation.

4. We must be able to apply a universal set of gates such that we can simulate any unitary transformation of the system. In classical computing this requirement is satisfied by the ability to apply the NAND gate to any pair of bits.
5. At the end of the computation we must be able to read out the result by measuring the final state of the system.

We shall now proceed to look in more detail at each of these from the quantum computing perspective.

2.4.1 System characterisation and scalability

In order to practically carry out quantum computing, we require a system consisting of a number of well-defined two level quantum systems for use as qubits. The key to realising the level of control necessary is to have a precise knowledge of the physical properties of the qubits. To understand how their states will evolve we need a good knowledge of the internal Hamiltonian of our chosen qubits and the interactions that exist between them. Precise manipulation of qubit states requires an accurate knowledge of their coupling both to the external fields used to control them and to other states of the system. In reality any system is likely to possess a number of unwanted accessible states. To reduce the error rate it is therefore important that any quantum computer is designed such that the probability of transitions into these states is minimised.

There are currently a number of different systems being investigated for use in quantum computation and it is beginning to seem likely that quantum information processing with at least a few qubits will be possible in many of them. However we have already seen that there is a limited set of problems for which quantum computers have an advantage over their classical counterparts and that even for these we generally require of the order of hundreds of qubits before we reach the limits of classical computation. It is for this reason that the issue of scalability is one of the fundamental parameters of any potential quantum computer.

Typically we find that it becomes significantly harder to fulfil the requirements for quantum computing as the effective number of qubits is increased. Possible reasons can be increased decoherence of the system due to increased

coupling to the environment, as is the case for strings of ions in traps [18]. Alternatively we might find that the ability to accurately measure the state of the system is reduced, as with nuclear magnetic resonance techniques where the signal level scales as $1/2^n$ for n qubits. It remains to be seen whether improvements in experimental technique will make it possible to carry out hundred qubit manipulations in any system. Some interesting recent work [19] has addressed the issue of scalability in a different way by manipulating several distinct qubit registers, each of which can be acted on individually by a single processor.

2.4.2 System initialisation

Clearly if we are to obtain any useful results from a quantum computer we need to know what the state of our qubit register was before the processing was started. Typically we would expect to use a simple, well-ordered state such as $|000\dots\rangle$. Besides enabling a standard input, there is another important reason for requiring the ability to set qubits in a given state; all the techniques for quantum error correction currently developed require a continuous supply of qubits in a low-entropy state. The reliability of qubit state measurement (described in section 2.4.5) can also be improved if we have access to extra qubits in the $|0\rangle$ state.

The initialisation criterion poses two difficulties. Firstly we would expect the presence of qubits in the wrong initial state to lead to errors in the result of the computation and hence we must be able to initialise them with a high probability of success. Secondly if the timescale for this process is significantly longer than the time taken to apply quantum gate operations it may prove to be a limiting factor in terms of the speed of computation.

There are two commonly used methods of setting qubits to a standard state. When the state we are interested in is the ground state we can simply cool the system down to a level such that the likelihood of the system being found in any other state is arbitrarily small. Since many techniques will involve cooling times much greater than the decoherence time of the system, this approach may not be able to supply qubits quickly enough for QEC to work efficiently. Alternatively we can make a measurement which projects the system into the required state, or some other known state from which we can generate our desired state. Since by making a measurement we necessarily

provide a strong coupling between our system and the measurement apparatus, we would generally expect the timescale for this kind of projection to be relatively short.

2.4.3 Long decoherence times

As we have already seen, decoherence is caused by the interaction of a quantum system with its environment and is in some sense involved in the emergence of classical behaviour in a system. Hence we need long decoherence times if we are to experience the benefits of QC over classical computing. An important point to note is that any quantum system will possess a range of decoherence times relating to the various degrees of freedom it possesses. Since for the purposes of quantum computing we are interested in the evolution of a particular pair of qubit states, only decoherence processes relating to these states are relevant.

The critical parameter for the purposes of QC is the ratio of the relevant decoherence time to the gate time of the quantum computer. This ratio gives us an indication of the likelihood that we will find the system in the correct state after we have applied some sequence of transformations. As we saw in the previous chapter, we can for a given error correcting code estimate a threshold value for this ratio to enable fault-tolerant computing. Thus long decoherence times are an essential property of any system that is to be used as the basis for a quantum computer.

2.4.4 A universal set of quantum logic gates

In order to carry out an algorithm to process the quantum information in our system, we must be able to apply a specific sequence of unitary transformations to subsets of our qubits. However if we wish to build a universal quantum computer (ie. one which can carry out any possible algorithm) we must be able to apply any possible unitary transformation. Since these form an infinite set we cannot apply all of them directly. Instead we need some finite subset of them which can be strung together to simulate any other transformation with arbitrary precision. We have already seen that most two-qubit quantum gates form such universal sets either by themselves or in combination with more straightforward single-qubit gates.

Despite the wealth of universal sets available, finding methods for implementing them for a given quantum system is not trivial. Often it is difficult or

impossible to apply two-qubit interactions directly and an extra ‘bus qubit’ is required to interact with each in turn. The presence of this extra coupling between qubits, however, can in turn result in an increased rate of decoherence. Also it is beneficial for both QEC and computing speed to be able to carry out operations on different subsets of qubits at the same time, a problem for schemes reliant on a single bus qubit.

The next consideration is essentially the same as that of decoherence. Since we may need to string together large numbers of quantum gates to simulate a single transformation, we would expect to apply perhaps many billions of gates over the course of a computation. As a result it is important that we are able to implement each gate as precisely as possible. We can think of errors in gate application as having similar effects to decoherence in that after the transformation we can no longer be assured of finding the system in the correct state. The similarity is such that the same error correcting codes can be used to eliminate errors due to imperfect quantum gates as those due to decoherence effects. The criterion for reliability of gate operation (in terms of average number of operations between errors) will therefore be of the same order as the acceptable ratio of decoherence time to gate time.

A short gate operation time is clearly beneficial in that it reduces the likelihood of errors developing as a result of decoherence. The gate time itself is very much dependent on the implementation of the processor being used. These might simply be experimental considerations such as limits on the amount of laser power that can be focused on a single ion in a trap, or the speed of the control apparatus. Alternatively there can be fundamental reasons why an adiabatic transformation is required, such as to prevent qubits from entering some other accessible state. For many systems there are likely to be a number of schemes for implementing a universal set of quantum gates. In this case it will be possible to choose the implementation which gives us the best balance of accuracy and rate of gate operation.

2.4.5 Qubit measurement

Once we have applied the required sequence of transformations to our qubit register we must be able to read out the results of the computation by measuring the state of specific qubits. There are two key requirements here. Firstly the act of measuring any given qubit must not cause an unwanted disturbance

of the rest of the system. Algorithms such as Shor’s period-finding method rely on the fact that the measurement of one register will cause the superposition of states of a second register to collapse down onto some specific subset of states. Clearly if the measurement process results in additional changes to the second register the algorithm will fail.

The second issue is that of quantum efficiency; the certainty with which we can say which state a qubit is in. Clearly a system with a low rate of decoherence on which we can apply very precise transformations is of little use as a quantum computer if we cannot reliably read out the results of our computations. Fortunately we can always improve the overall reliability of our measurements by repeating the computation several times. Alternatively, providing it does not disturb the rest of the quantum register, we can simply make copies of the output qubit using a CNOT gate and some additional qubits in the $|0\rangle$ state and measure all of these. The copying proceeds as

$$CNOT(|0\rangle |0\rangle) = |0\rangle |0\rangle$$

$$CNOT(|1\rangle |0\rangle) = |1\rangle |1\rangle$$

Note that this does not constitute a violation of the no-cloning principle since we are only copying the value of the output qubit in the $|0\rangle, |1\rangle$ basis, not the qubit itself. Using this technique we can obtain reliable readout, even with a very unreliable measurement process.

2.5 Conclusions

According to the principles of quantum mechanics, quantum systems possess extra degrees of freedom relative to classical systems — namely the ability to exist as superpositions of entangled states, and quantum interference effects. These can be used as a resource for computation. In the twenty years or so since this idea was first explored there have been significant advances in the understanding of the computing potential of quantum systems. The discovery of Shor’s factoring algorithm marked an important point in the development of QC because it provided for the first time a complete method of solving an important problem with an exponential speedup over classical computing. Indeed the factoring problem is of such importance in fields such as cryptography that Shor’s algorithm has fueled an explosion in interest in (and funding of)

practical implementations of quantum computing. The difficulties involved in realising quantum computation are immense. The development of error-correcting schemes has provided a target to aim at, however, in that if we can satisfy its demands we can carry out computations lasting well in excess of the decoherence times of the system.

Chapter 3

Quantum information processing in ion traps

The physical implementation of quantum computing is extremely difficult. As we saw in chapter 2, fault-tolerant QC is possible, but extremely well isolated qubits and precise gate operations are required for it to be achieved. A number of different implementations have been proposed, perhaps the most advanced of these in terms of processing quantum algorithms being nuclear magnetic resonance (NMR) techniques. However there are significant difficulties involved in scaling up these techniques and it seems unlikely that NMR quantum computing will be possible with more than ten qubits, at least in the near future.

The use of the electronic and nuclear spin states of trapped ions as qubits has become one of the leading technologies for scalable quantum computing. This is largely due to the ability to precisely manipulate the internal states of individual ions and to rapidly prepare and measure qubit states. Trapped ions have long been used for spectroscopy, frequency standards and as a testing ground for various quantum optical effects, so the physics of such systems is relatively well understood. The use of trapped ions as a basis for scalable quantum computing really began with the discovery in 1995 of the first method for applying a CNOT gate to pairs of ions [8]. Since then significant progress has been made in the characterisation of ions as qubits and the manipulation of qubit states. 2003 saw the first two demonstrations of the application of two-ion quantum gates on trapped ions [20, 21], completing a universal set as required for quantum computing.

In this chapter I shall review the proposed schemes for quantum computing with trapped ions and the extent to which they satisfy the requirements for QC discussed in the previous chapter. Next I shall report on the progress that has been made towards realising trapped-ion QC and the obstacles that remain to be overcome. Finally I shall look at how some of these obstacles might be circumvented by using different techniques to trap ions.

3.1 Quantum computation with trapped ions

3.1.1 Ion cooling in the linear RF trap

By far the most commonly used trap for experiments investigating quantum computing is the linear radio-frequency (RF) trap [22, 3], shown in figure 3.1. This consists of four long cylindrical electrodes of which one pair of opposing electrodes is grounded and the other is connected to a rapidly oscillating voltage $V_0 \cos \Omega_T t + U_r$. This creates a RF potential

$$\Phi_{rf} = \frac{1}{2}(V_0 \cos \Omega_T t + U_r)\left(1 + \frac{x^2 - y^2}{R^2}\right), \quad (3.1)$$

where R is a constant dependent on the geometry of the trap. The time average of Φ_{rf} provides a harmonic trapping potential in the x and y directions. Axial confinement is achieved by applying an additional bias voltage U_0 to the ends of each electrode. This results in a harmonic potential along the z axis

$$\Phi_{dc} = \frac{1}{2}rU_0(z^2 - (x^2 + y^2)), \quad (3.2)$$

where r is again a geometrical constant.

To reduce interaction with the environment as much as possible, ion traps are run under high vacuum, typically $10^{-11} - 10^{-10}$ mbar. Overall, the dynamics of the ions within the trap are governed by the Hamiltonian

$$H = \sum_{i=1}^N \frac{m}{2}(\omega_x^2 x_i^2 + \omega_y^2 y_i^2 + \omega_z^2 z_i^2 + \frac{|\vec{p}_i|^2}{m^2}) + \sum_{i=1}^N \sum_{j>i}^N \frac{e^2}{4\pi\epsilon_0 |\vec{r}_i - \vec{r}_j|}, \quad (3.3)$$

for N ions of mass m . The parameters are chosen such that $\omega_z \ll \omega_x, \omega_y$ and the ions lie along the z axis. Assuming good isolation from the environment the motion of ions under this Hamiltonian is of the form of a quantised harmonic




Figure 3.1: Diagram of the basic structure of a linear RF trap. The additional bias voltage applied to each end electrode is not indicated.

oscillator, so the available energy levels increase in units of $\hbar\omega_z$ and we can define corresponding raising and lowering operators a_z^\dagger and a_z . If the ions are sufficiently cool the only vibrational mode available is the centre-of-mass mode, where all the ions are moving as a single body. That this is the only available mode can be useful, since if all the ions are moving in the same way they will all have the same laser–ion coupling. Furthermore it is a requirement of certain quantum computing schemes that the ions are in the motional ground state, ie. $k_B T \ll \hbar\omega_z$. Note that in many cases the trap parameters are also chosen such that the inter-ion spacing is large enough that ions can be individually addressed, but not so great as to complicate laser cooling of all the ions at the same time.

Ion cooling is achieved with a combination of ‘Doppler’ and ‘sideband’ laser cooling. With Doppler cooling the ions are illuminated by a laser red-detuned slightly in frequency from some resonant transition. Due to the Doppler effect, ions moving with a particular velocity component towards the beam will experience light at close to the transition frequency and will thus absorb photons preferentially over other ions. Since these photons are re-emitted in random directions there will over time be a net momentum transfer tending to slow the ions down. Providing the beam direction has a component along each axis of motion, ions can be cooled to the Doppler limit, $k_B T \approx \hbar\Gamma/2$, where Γ is the natural width of the excited state of the cooling transition.

Sideband laser cooling can be used to further cool the ions to the ground state of the centre-of-mass mode. With a sufficiently narrow (and therefore

slow compared to the trap frequency) transition addressed by a laser with sufficiently narrow linewidth, it is possible to resolve the sidebands of a transition due to the vibrational motion of the ions. Because the field an ion “sees” is dependent on its position, there is a coupling between the vibrational state of the ion string and its spin state. Thus, at temperatures close to the Doppler limit, the usual term in the Jaynes–Cummings Hamiltonian for the laser–ion interaction is modified to include an additional term describing this coupling:

$$H_{las} \approx \frac{\hbar\Omega^a}{2}(a_\gamma\sigma^+ + a_\gamma^\dagger\sigma^-) + \frac{\hbar\Omega^b\eta_z}{2\sqrt{N}}(a_\gamma a_z\sigma^+ + a_\gamma^\dagger a_z^\dagger\sigma^- + a_\gamma a_z^\dagger\sigma^+ + a_\gamma^\dagger a_z\sigma^-). \quad (3.4)$$

Here σ^+ and σ^- are the raising and lowering operators associated with the electronic transition, η is the Lamb-Dicke parameter ($\equiv kz_0$) relating the width of oscillation to the wavelength of the laser light and Ω is the Rabi frequency, given by $\Omega \equiv dE/2\hbar$. The second term in this equation enables interactions at a detuning of $\pm\omega_z$ where the motion of the ions is also modified. Let us consider an ion in the state $|a, b\rangle$ where $a \in \{g, e\}$ is the state of some transition and b is the motional state of the ion. A schematic of the energy levels is shown in figure 3.2. By tuning the laser to the frequency of the red sideband a π pulse will stimulate the process

$$|g, b\rangle \longrightarrow |e, b-1\rangle \longrightarrow \begin{cases} |g, b\rangle \\ |g, b-1\rangle \\ |g, b-2\rangle. \end{cases} \quad (3.5)$$

The three decay processes listed have various probabilities, but over time the ion will tend to fall into progressively lower motional states. Using this technique it is possible to cool ions to the motional ground state. At this point, the next lowest accessible energy states correspond to the centre-of-mass mode. Thus, starting from the motional ground state, the absorption of a photon can be made to result in the recoil of the entire string with each ion undergoing the same motion.

3.1.2 Electronic states as qubits

The physical manifestation of qubits for trapped ions used for quantum computation is provided either by a pair of hyperfine ground states, or a metastable electronic state. It is common to use the primary ions of the alkaline earth




Figure 3.2: The various centre-of-mass vibrational states appear as a series of steps (exaggerated for clarity) for each electronic energy level. In sideband cooling, light detuned from an electronic transition by ω_z is used to reduce the vibrational state. Eventually the ion will be left in the vibrational ground state.

metals, since the lone electron in their outer shells tends to result in a relatively simple electronic structure. The choice of ion is based on finding a suitably long-lived qubit state, with appropriate transitions for the frequencies of laser light available.

As an example we shall look at $^{40}\text{Ca}^+$, an ion much used for experiments in QC [23, 24]. The relevant energy levels are shown in figure 3.3. The $S_{1/2}$ and $D_{5/2}$ levels form the computational basis qubit states. At 729nm, the transition between them can be conveniently reached with a frequency-stabilised Ti:Sapphire — providing a linewidth narrow enough to also enable sideband cooling on this transition. The $D_{5/2}$ state has a natural lifetime of about a second. This is sufficiently long that, compared to decoherence processes, spontaneous emission is unlikely to be a limiting factor in the stability of the qubit states. In fact, the long lifetime of this state means that 854nm light may be necessary when sideband cooling, to pump the ions into the short-lived $P_{3/2}$ state and hence back into the cooling cycle. Doppler cooling can be carried out on the $S_{1/2}$ to $P_{1/2}$ transition at 397nm, with 866nm light used to repump any ions that decay into the $D_{3/2}$ state.

The presence of the short-lived $P_{1/2}$ state is also of use for the measurement




Figure 3.3: The relevant electronic energy levels for $^{40}\text{Ca}^+$. The 729nm qubit transition is also used for sideband cooling. Doppler cooling is carried out on the 397nm transition. 854nm and 866nm light is used to pump out the D states to provide a rapid return to the laser cooling cycles.

of qubit states using a technique known as ‘electron shelving’. If a beam at 397nm is shone onto an ion in the $S_{1/2}$ state it will rapidly absorb and re-emit the photons — scattering them in random directions. The short lifetime means that a single ion can scatter many photons within the space of, say, 1 millisecond. Thus even with a low quantum efficiency photomultiplier tube it is possible to detect the scattered photons at a reasonable rate. An ion in the $D_{5/2}$ state is ‘shelved’ and will not scatter photons. A 397nm beam can therefore be used to determine the state of a qubit and cause the collapse of any superpositions as appropriate.

The fact that the $|0\rangle$ qubit state is the ground state of our system enables convenient state preparation. We can use the 854nm beam to stimulate transitions into the short-lived $P_{3/2}$ state and simply wait long enough for the ion to have decayed spontaneously into the ground state with arbitrarily high probability. This ease of measurement and state preparation is another advantage of using trapped ions as qubits.

3.1.3 Quantum gates for trapped ions

Arbitrary manipulations of the system can be achieved through the use of lasers, which transform the ion states via the interaction Hamiltonian in equation 3.4. Applying a laser pulse at the transition frequency and with appropriate phase we can carry out the general rotation operators $R_x(\theta) = e^{-i\theta X/2}$ and $R_y(\theta) = e^{-i\theta Y/2}$ (as defined in section 2.2.2). The angle of rotation is determined by the duration of the pulse multiplied by the Rabi frequency. These rotation operators can then be combined to form any single-qubit transformation.

Since there is no direct coupling between internal states of the ion, two-qubit interactions are necessarily mediated by an additional bus qubit. In the Cirac and Zoller scheme [8] the $|0\rangle$ and $|1\rangle$ states of the centre-of-mass vibrational mode are used as the two states of the bus qubit, so the system must begin in the vibrational ground state. A CNOT gate can be constructed by making use of the same transition sidebands as the cooling technique described in the previous section. The relevant pulses for the scheme are indicated in figure 3.4. The following three pulses are used, where c is the control qubit ion and t is the target:

1. A π pulse of frequency $\omega_0 - \omega_z$ on the control qubit induces a transition

iff it is in the excited state. This has the effect of transferring the qubit state of c to the bus qubit (with a phase term if a transition is made), leaving c in the ground state;

$$\begin{aligned} |g\rangle_c |0\rangle &\longrightarrow |g\rangle_c |0\rangle \\ |e\rangle_c |0\rangle &\longrightarrow -i |g\rangle_c |1\rangle. \end{aligned} \quad (3.6)$$

2. A 2π pulse on the target ion of frequency matching the transition between $|g\rangle_t |1\rangle$ and $|aux\rangle_t |0\rangle$, where $|aux\rangle$ is some auxiliary electronic energy level of the system. The auxiliary state is never occupied, but a factor of -1 arises if the rotation through it occurs:

$$|g\rangle_t |1\rangle \longrightarrow -|g\rangle_t |1\rangle. \quad (3.7)$$

3. The first pulse is repeated. This maps the bus qubit back onto the control qubit, leaving the system back in the vibrational ground state:

$$\begin{aligned} |g\rangle_c |0\rangle &\longrightarrow |g\rangle_c |0\rangle \\ |g\rangle_c |1\rangle &\longrightarrow -i |e\rangle_c |0\rangle. \end{aligned} \quad (3.8)$$

The effect of these three pulses taken together is to introduce a phase factor of -1 only if both qubits are in the $|e\rangle$ state, ie. it is a controlled-phase gate:

$$|e\rangle_c |e\rangle_t |0\rangle \longrightarrow -|e\rangle_c |e\rangle_t |0\rangle. \quad (3.9)$$

This is equivalent to a CNOT gate if we recall that $Z = HXH$; phase gates are equivalent to CNOT gates if we use the Hadamard operator to switch basis states. Thus by applying a Hadamard gate (built up from single-qubit rotations) before and after the pulses described above, we have a complete CNOT gate. We can show this explicitly by considering the target qubit in the state $|\pm\rangle_t = (|g\rangle_t \pm |e\rangle_t)/\sqrt{2}$. The above pulse sequence results in the transformation:

$$\begin{aligned} |g\rangle_c |\pm\rangle_t &\longrightarrow |g\rangle_c |\pm\rangle_t \\ |e\rangle_c |\pm\rangle_t &\longrightarrow |e\rangle_c |\mp\rangle_t. \end{aligned} \quad (3.10)$$




Figure 3.4: Energy level diagram showing the transitions induced by the three pulses in the Cirac and Zoller scheme. Note that the absence of a vibrational state below $|0\rangle$ gives each of the transitions shown a unique frequency.

The Cirac and Zoller scheme uses the fact that there is no vibrational state below $|0\rangle$ to bring about transitions dependent on the state of the bus qubit. As a result it is critical that the system is initially in the vibrational ground state. Though sideband cooling has been used in recent experiments to achieve an average ground state occupation of about 99% [20], this is still far short of the gate precision that is required for fault-tolerant quantum computing to be viable.

There have been a number of proposed two-qubit gate schemes that attempt to remove the stringent requirement for ground state occupation, but all have some other disadvantageous feature. One of the most obvious is to carry out the Cirac and Zoller scheme with the ground state of some higher energy collective vibrational mode than the centre-of-mass mode. These have been found to exhibit heating times as much as 50 times longer than the centre-of-mass mode [25]. However the fact that each ion is moving in a different way means that the laser-ion couplings vary and different pulse durations are required for each ion. The difficulty of eliminating ion heating is replaced by the problem of a more complex control system.

Another conditional quantum logic gate, proposed by Mølmer and Sørensen [26], has the advantage of being insensitive to phonon excitation even

during the gate operation itself. The scheme involves shining two beams, detuned by the same amount from the two sidebands of an electronic transition, on both the control and target ion at the same time. Interaction between the ions occurs via the exchange of a virtual phonon and no change in the vibrational state actually occurs. The drawback of gates based on this method is that the necessarily large detuning (to prevent actual phonon exchange) results in a slow interaction. This places severe restrictions on the acceptable decoherence rate for the quantum system being used. A gate in many ways related to that of Mølmer and Sørensen has been developed and demonstrated by the NIST group [21]. This gate is described in section 3.3.1 and so we shall not discuss it here.

Other gate schemes include that of Poyatos, Cirac and Zoller [27], where a laser pulse provides a momentum transfer to a control ion in a direction dependent on its internal state. Thus the superposition describing the internal state of the ion is transferred to a superposition in position space for the entire ion string. A precisely aligned laser pulse can be made to interact only with ions in one or other of these position states, providing the ability to carry out controlled versions of any single-qubit gate. This approach does not require the ions to be in a vibrational ground state, but a very high degree of precision is required to resolve the separate position states of the ions and any unintended spatial variations in electric potential will result in phase differences developing between the two states.

3.2 Quantum computing potential of trapped ions

Having seen that many of the fundamental techniques and processes required for quantum computation with trapped ions can in principle be performed, let us now discuss the potential of trapped ions to form the basis of a scalable, universal quantum computer. We can do this by considering the extent to which ion trap systems satisfy the criteria for QC laid down in section 2.4 [28]:

1. **Scalability:** The number of ions that can be used for QIP in an ion trap is principally limited by two effects. The first is the increasing complexity of the sideband spectrum of transitions, due to the addition of new modes of oscillation of an ion string for each ion added. Each of these modes has an infinite number of resonances associated with it,

and these combine to form what Chris Monroe is said to refer to as “the spectrum of death”. The complexity of this spectrum does not provide a fundamental limit on scalability, but does make it harder to prevent unwanted transitions during qubit manipulations.

Changes in the structure of ion strings in traps as the ion number is increased may also cause difficulties in carrying out quantum computing. Once a certain number of ions is reached there is a phase transition to other configurations, such as zig-zags. Again it remains to be seen if this poses an insurmountable barrier or whether techniques can be developed to cope with complex string configurations. What is certain is that as the ion number is increased the inter-ion spacing is reduced. Increasingly precise, tightly-focused beams are therefore required if we are to resolve individual ions. To overcome the problem of scalability in ion traps, proposals have been made for quantum processors consisting of arrays of traps or trapping regions [19, 29]. Qubits can then be extracted from particular registers and operated on before being returned, so each register need only contain a manageable number of ions.

2. **Initialisation:** Internal qubits can be zeroed very reliably by optically pumping the ions into short-lived states that decay only into the zero state. The main difficulty in terms of initialisation is preparing and maintaining the system in the vibrational ground state with sufficiently high probability. Sideband cooling is an effective method for obtaining the ground state, but heating of the ions during computation means that additional cooling will be required. Development of the proposed gate schemes for “hot” ions may help alleviate this problem.
3. **Decoherence:** The key sources of decoherence for qubits in ion traps are the heating of ions due to imperfections in the RF field used to trap them and stray magnetic fields that are difficult to eliminate entirely from the apparatus. The fundamental limit on the decoherence time of ions in a trap is set by the natural lifetime of the qubit states, and depending on the ion and states chosen this can be extremely large. Thus ion traps do at least have the potential to be used for long quantum computations.
4. **Universality:** We have already seen that arbitrary single qubit manipulations can in theory be carried out via ion–laser interactions. This

generally requires the ability to address individual ions and control pulse duration, intensity, frequency and phase with very high precision. Having also shown that two-qubit gates are in principle possible in ion traps, all the requirements for universal gate operations are available. The issue is one of gate reliability and whether we can meet the stringent requirements for quantum error correction to be implemented.

5. **Readout:** The ease with which we can carry out measurement of individual qubit states is one of the primary benefits of quantum computing with ion traps. Whereas most other QC schemes suffer from poor readout reliability, the fluorescence measurement technique described in section 3.1.2 has been used in many experiments to achieve measurement reliabilities of close to 100% [30].

3.3 Realisation of ion trap quantum information processing

3.3.1 Previous developments in ion trap QIP

Universal quantum computing with any system is still a long way from becoming a reality. Recent experiments have been geared towards demonstrating the individual components of a quantum information processor, and investigating the extent to which various systems fulfill the criteria for QC. Experiments with trapped ions have proved to be some of the most successful in these tasks, having already demonstrated many aspects of the schemes described in the previous chapter [31].

The first step on the long road to realising universal quantum computing was taken just a few months after the original proposal by Cirac and Zoller. In 1995 Dave Wineland and his associates at NIST in Colorado applied a CNOT gate to two qubits represented by the hyperfine ground states of a beryllium ion and the $n = 0, 1$ states of one of its vibrational modes [32]. The operation was essentially equivalent to the second step in the Cirac and Zoller scheme, flipping the phase of the atomic qubit iff the vibrational qubit is in the $|1\rangle$ state by applying a 2π rotation through an auxiliary state. A Hadamard gate was applied to the atomic qubit before and after the 2π rotation such that the overall effect was to flip its state in the computational basis.

The operation was carried out in approximately 50 microseconds with a reliability of better than 80%. Decoherence times of a few hundred microseconds were measured, short compared to the times that had previously been observed for qubits based on hyperfine ground states without the application of the CNOT quantum gate. Important though this experiment was for demonstrating the first two-qubit quantum gate, there was no straightforward way in which it could be scaled up to allow computations with more than two qubits. Since only one qubit can be represented by the centre-of-mass mode, scalability requires CNOT gates that operate on arbitrary pairs of atomic qubits. Besides this the limited gate reliability and low ratio between characteristic decoherence and gate operation times would prevent computations of more than a few gates.

In 1998 the NIST team succeeded in generating on demand entangled states of pairs of trapped beryllium ions [33] using techniques similar to those involved in the Mølmer and Sørensen CNOT gate scheme [34, 26]. In particular they prepared ion pairs in two of the Bell states, $(|\downarrow\uparrow\rangle \pm |\uparrow\downarrow\rangle)/\sqrt{2}$. States such as these have historically been very important in demonstrating a variety of strange quantum effects such as the ability of a measurement on one particle to immediately affect the state of another a large distance away. Bell states had only previously been produced by selecting out appropriately entangled states from large numbers of particle pairs produced by random processes. The trapped ion method therefore has the advantages that states can be produced at a desired time and that the technique can be scaled up to larger numbers of particles without a significant increase in the timescale for production. Indeed the same group went on to demonstrate the entanglement of four trapped ions using a similar technique [35]. Entangled states are important for quantum computing since they enable us to make use of the quantum parallelism effect. Referring back to Shor's algorithm as an example, it is a requirement that the states of the two quantum registers used for period finding are entangled. This causes the measurement of the $|f(x)\rangle$ register to result in the collapse of the superposition of states of the $|x\rangle$ register onto a certain subset of states.

After several years of relative quiet (in terms of publications), 2003 brought significant developments in ion trap QIP, with the independent demonstration by two groups of two ion-qubit gates. The group at NIST carried out a geometric phase gate, using the hyperfine ground states of beryllium ions [21]. Here the ions experience a phase shift $e^{i\phi}$ iff they are both in the $|1\rangle$ state.

The phase is introduced by addressing the ions with counter-propagating laser beams far detuned from resonance with a fast electric dipole transition. The resulting Stark shifting of the qubit levels leads to a state dependent force on each ion, which varies depending on its position in the standing wave. The ion positions are set such that each ion lies in an identical section of the standing wave, so that if the two ions have the same internal state the centre-of-mass motion can in principle be excited and if the two ions are in different spin states the stretch mode can be excited. By introducing a frequency difference between the two laser beams close to the stretch mode frequency we can prevent centre-of-mass mode excitation. If the lasers' frequency difference is in fact very slightly different from the stretch mode frequency then over time the excitation force will oscillate in and out of phase with the stretch motion and so with appropriate pulse duration the motion can be returned to its initial state. The internal state, however, retains a phase shift $e^{i\phi}$ as a result of this temporary motional displacement.

The angle ϕ can be changed by altering laser intensities, so for $\phi = \pi$ we have a controlled-phase gate (strictly speaking we must also apply two single qubit gates). As we saw earlier, this is equivalent to a CNOT gate if we apply a Hadamard gate before and after the two-qubit operation. The reliability of the controlled-phase gate was measured to be better than 95%, and the gate time about 40 microseconds. Although this is a remarkable achievement for one of the first demonstrations of two ion-qubit gates, it is significantly short of the 10^{-3} to 10^{-5} per qubit per gate error rate estimated as being required for fault-tolerant QC.

At the same time as the above result was published, Rainer Blatt's group in Innsbruck announced the successful operation of a CNOT gate following closely the Cirac and Zoller proposal [20]. They used electronic states of calcium-40 ions in a linear RF trap as atomic qubits, and measured a ground state occupation of the vibrational centre-of-mass mode of about 99%. Their gate reliability was found to be about 75%, with the main error being the frequency stability of the lasers used. This fidelity was significantly worse than the NIST result, but it is nevertheless impressive given that the method used is essentially the first to have been developed. In subsequent experiments the measured fidelity of this gate has been improved slightly. A more detailed description of the techniques used can be found in reference [36]. At about the same time the Innsbruck group were able to perform a simple quantum

algorithm, the Deutsch–Josza algorithm, using one ion and the vibrational mode as qubits [37].

During this period the NIST group began working with ions stored in segmented linear RF traps — traps with several extra electrodes that generate a number of different trapping zones. With segmented traps, strings of ions can be separated and ions can be shuttled around between trapping zones. With sufficiently small electrodes placed close enough together (the NIST trap from this era had electrodes as small as $100\ \mu\text{m}$) it was found to be possible to separate ions located in the same trapping region without any measurable failure rate. After separation both sets of ions usually remained in the ground state of the stretch mode of motion, but would typically have a mean motional quantum number for the centre-of-mass mode of about 1. It is generally found to be true that the stretch motional mode is less sensitive to heating than the centre-of-mass mode. The timescale for ion separation was typically in the region of 2 milliseconds.

The demonstration of two ion qubit gates had paved the way for performing a number of basic quantum information processing experiments on small numbers of ions in linear strings and in segmented traps. The trend of the Innsbruck and NIST groups' simultaneous publishing of results was continued in June 2004 with the appearance in *Nature* of back to back papers describing deterministic teleportation of qubit states between trapped ions [38, 39]. Besides being interesting in its own right as the first demonstration of teleportation that is non-probabilistic, this technique is useful because it might be used in quantum computing to rapidly move quantum information around without the need to shuttle ions over large distances during a computation. The first step is to entangle the internal states of a pair of ions by using a two qubit gate to create a Bell pair. These can then be distributed to separate locations. By entangling one part of the Bell pair with the ion to be teleported and making appropriate measurements of their internal states it is possible to determine what single qubit operation must be applied to the remaining part of the Bell pair in order to put it in an identical state to the teleported qubit.

By using a segmented trap the NIST group were able to move individual ions around so that only those required for each operation were present in a given trapping region. They found the average teleportation fidelity to be $78\pm 2\%$, with the key sources of infidelity being the initial state preparation, motional heating from the separation process and dephasing of the final state

due to fluctuating magnetic fields. With smaller electrodes in the trap the time taken to carry out the teleportation would have been reduced significantly from what was then 4 ms, and this would have helped improve the fidelity of the process.

Since their two-qubit gate involves individual ion addressing, the group in Innsbruck were able to perform teleportation using three ions in a single linear RF trap. They observed teleportation fidelities ranging between 73% and 76%, depending on the state being teleported.

In the same month as these teleportation results were published the NIST and Innsbruck groups also published back to back papers in *Science*. The Innsbruck work related to the deterministic generation of Greenberger–Horne–Zeilinger (GHZ) and W maximally-entangled states of three ions [40]. The analysis of these states showed that GHZ states (of the form $|0\rangle|0\rangle|0\rangle + |1\rangle|1\rangle|1\rangle$) are extremely sensitive to decoherence, since they are superpositions of states with maximally different energies. Measurements showed coherence times of about 1 millisecond — much shorter than the 200 milliseconds measured with W states (of the form $|0\rangle|1\rangle|1\rangle + |1\rangle|0\rangle|1\rangle + |1\rangle|1\rangle|0\rangle$). Also interesting is the fact that measurement of a single qubit necessarily destroys all coherences in a GHZ state, but for a W state will sometimes reduce the system to a Bell state.

The NIST publication detailed the use of spin squeezing to improve the signal to noise ratio of a Ramsey type spectroscopy experiment [41]. They used a three ion GHZ state to obtain a measurement precision 1.45 times greater than that possible in a system without entanglement. Though they acknowledge that dephasing of GHZ states due to magnetic field fluctuations is rapid, their technique can in principle be scaled up to more ions and enable spectroscopic resolution limited only by the Heisenberg uncertainty.

Using exactly the same experimental techniques as for their demonstration of spin squeezing and teleportation, the NIST team also performed a demonstration of a simple quantum error correcting code [42]. They encoded one logical qubit onto three physical qubits before applying additional pulses designed to simulate the effect of spin-flip errors on the qubits. As expected they observed an improvement in the fidelity of the logical qubit when the QEC code was applied, providing the induced error was not so large that unprotected multiple errors were likely. Since there were significant errors in the gate operations required to perform the code, for small errors it would often

have been better not to perform QEC at all. They also found that certain input states were less well corrected by the code, though this was an expected feature of the code they were using.

A few months later the NIST group performed a variant of the quantum Fourier transform [43] — a key step in Shor’s factoring algorithm. The variant used requires only single qubit rotations and measurements, where the rotations made on each successive qubit depend on previously made measurements. Once again this was done with three ions in a segmented trap, so that they could represent the numbers 0 to 7 (in a binary representation). They initially place the ions in a superposition of states with some periodicity, such as $|0\rangle + |2\rangle + |4\rangle + |6\rangle$. For this example the fourier transform should produce a superposition $|0\rangle + |4\rangle$. They found transform accuracies ranging between 87% and 99%, depending on the input state used. They attribute the errors predominantly to initial state preparation and the dephasing that occurs during the ~ 4 milliseconds needed to carry out the Fourier transform operations.

The most recent achievements by these two main groups were published in *Nature* in December 2005. Both groups demonstrated the production of maximally entangled states of larger numbers of ions — the Innsbruck group created W states of up to 8 ions [44], whilst the NIST group had produced GHZ states of 6 ions [45]. Both groups used standard linear RF traps. In Innsbruck the fidelities measured were between 0.85 (for a 4-ion state) and 0.72 (for an 8-ion state). At NIST, fidelities of better than 0.76, 0.60 and 0.51 were observed for GHZ states of 4, 5 and 6 ions respectively. The somewhat poorer fidelities in the latter experiment indicate the fact that the NIST gate, which incorporates no individual ion addressing, is much less suited to operations involving many ions in one trap.

These results are technically very impressive. Reconstructing the density matrices for these entangled states requires generating them a huge number of times and manipulating them so that different coherences can be measured. In the case of the Innsbruck 8-ion experiment the state was created 656 100 times, requiring 10 hours of experimental time, in order to fully characterise it. Although this work is very interesting as a study of entanglement itself, it represents the natural limit to what can be performed using current techniques. The number of density matrix components increases exponentially with the number of ions, so even repeating these experiments with just a few more ions would become technically next to impossible. It seems that new direc-

tions must be taken if we are to move further down the road towards scalable quantum information processing with trapped ions. Groups will need to concentrate on improving operations on one and two qubits, and use relatively small qubit registers as building blocks for larger devices.

Thanks largely to the results from the Innsbruck and NIST groups over the last ten years, ion trapping leads the field in experimental QIP. This has led to a dramatic increase in interest in ion trapping research. A number of other groups have since entered the fray and begun investigating various aspects of ion trap QIP. These groups are all taking a slightly different approach to the problem — they use a whole menagerie of different ion species and plan to use different quantum gates. In order to provide an impression of the breadth of research that is ongoing, I shall briefly describe the key features and progress of the work being done by a few of the groups around the world.

Chris Monroe's group, now at the University of Michigan, have chosen the hyperfine ground states of $^{111}\text{Cd}^+$ as their qubit levels and make use of the Mølmer and Sørensen gate described earlier [26, 46]. Their choice of ion is based on a number of factors, including the presence of a magnetic field-insensitive qubit transition (which we shall discuss later) and a large fine structure splitting that has implications for the Raman processes used to manipulate their ions. Their measured heating rates for cadmium ions also seem to be one or two orders of magnitude better than those reported by groups using other species in similar traps [47], due in part to the rather large mass. The main difficulty with cadmium is that the strong dipole transition used for laser cooling and detection lies at 214.5 nm and requires the technically demanding frequency-quadrupling of a Ti:sapphire laser. Key recent work has included sympathetic cooling of a qubit ion by addressing a different cadmium isotope [48] and the probabilistic entanglement of ion qubit states with photons [49]. They are also very much involved in the drive towards smaller, more scalable traps [50] and techniques to move ions around in them [51].

The group of Derek Stacey and Andrew Steane in Oxford have also recently entered the field of trapped-ion QIP [52]. They ultimately intend to use the odd isotope calcium-43, since it possesses hyperfine states that they can use as long-lived qubits. Thus far they have produced entangled states of two calcium-40 ions using their own variant on the NIST spin-dependent phase gate. They are currently working on implementing the same gate with $^{43}\text{Ca}^+$. The Oxford group are also about to begin work with a segmented trap that

will enable them to shuttle ions around and address groups of ions individually. From a theoretical point of view the group is investigating the possibility of using composite laser pulses to more rapidly trace out a closed loop in phase space and hence dramatically increase gate speed. Devising and generating these pulse sequences will be very demanding and it will most likely be some time before they are investigated experimentally.

Recently moved to Siegen in Germany, the group of Christof Wunderlich are taking a rather novel approach to qubit manipulation. They are using ytterbium-171 ions, which like $^{43}\text{Ca}^+$ and $^9\text{Be}^+$ possess long-lived hyperfine ground states suitable for use as qubits [53] and can be laser cooled on a strong dipole transition. Instead of using direct optical or Raman manipulation of their qubit states, however, the Siegen group use microwave techniques. Although the phase and frequency control that can be achieved with microwaves is much better than that which can easily be achieved with optical radiation, microwaves cannot be focused to the small spot sizes required for individual ion addressing. Furthermore, the long wavelength results in a very small value for the Lamb–Dicke parameter, which expresses the ability of radiation to couple to the motional degrees of freedom of an ion. Thus it is not generally considered possible to carry out multi-qubit gates, where mutual coupling to the motional state is used to achieve entanglement. In order to combine the useful features of optical and microwave sources, Raman processes are often used. Here acousto-optical modulators are used to imprint the properties of the microwaves onto optical laser sources.

The alternative approach used by the Siegen group is to apply a magnetic field gradient across a string of ions. The Zeeman shifting of the energy levels then means that the microwave frequency required to manipulate the spin qubit is different for each ion — so individual addressing of ions becomes possible. Furthermore, the effective (spin-dependent) Zeeman forces acting on the ions results in a coupling between the spins of different ions analogous to the spin–spin coupling used in nuclear magnetic resonance QIP [54, 55, 56]. The selective microwave addressing of individual ions in a string has now been demonstrated, along with the ability to prepare qubits in a known initial state [53]. The magnetic field gradient can also be used to perform an additional trick — by appropriately tailoring the magnetic field it is possible to obtain a situation where the difference in Zeeman shifts between successive ions corresponds exactly to the frequency spacing between different motional

modes. Thus all the motional modes can in principle be sideband cooled using a single laser frequency [57]. It now remains to be seen whether the spin coupling technique can be used to generate multi-particle entanglement.

To provide an exhaustive list of all the research that is currently ongoing in trapped-ion QIP would require much more space than is available here. It is however worth noting that even at this stage the field is still attracting new research groups. The groups of Isaac Chuang at MIT, Michael Drewsen at the University of Aarhus in Denmark and Patrick Gill at the NPL in London are all in the early stages of setting up and demonstrating 2-qubit quantum logic gates with various ions.

Other groups are investigating other sides to trapped-ion QIP, such as the interfacing of ions with photons in order to enable the distribution of entanglement over large distances. This can be done by placing a trapped ion at an antinode of some mode of a high-finesse optical cavity, so that photon emission is predominantly into that cavity mode and the ion state is coupled to the photon number of the cavity. Photons can be allowed to leak out of the cavity and can travel large distances through optical fibres before interacting with another ion via another high-finesse cavity. These ideas are currently being investigated at, for example, the University of Innsbruck and Sussex University [58, 59, 60].

Finally, the group of John Bollinger at NIST is attempting to generate entanglement of very large numbers of ions by performing a phase gate on a trapped ion crystal in a Penning trap. These crystals can be comprised of millions of ions and although it is not clear how such mass entanglement might easily be used for quantum computing it may find interesting application in precision spectroscopy [61].

3.3.2 Current issues in ion trap QIP

Having reviewed the key results that have thus far been obtained in ion trapping experiments aimed at QIP, it would be easy to make the mistake of thinking that researchers were well on the way towards building a fault-tolerant universal quantum computer. It is true that the development of ground state cooling, precise single-qubit manipulation and multi-qubit entanglement provides compelling evidence for the potential for ion trap quantum computing. These techniques have enabled the demonstration of a basic quantum algo-

rithm and teleportation, making them interesting in their own right. But all of these have been very much proof-of-principle experiments — to achieve anything like universal quantum computing will require qubit lifetimes and gate fidelities that are orders of magnitude better than those currently measured. Most important, however, is the requirement for scalability.

Fault-tolerant quantum computing will probably require thousands of gates to be carried out on thousands of physical qubits (many more than there are logical qubits if we are implementing quantum error correction). As we have seen we cannot hope to do this with a single ion string, but must look towards some kind of multiplexed architecture with multiple trapping zones. In such an arrangement ions could be shuttled around from one trap to another, so that the required ions could be brought together when a gate operation is required, and stored whilst other operations are being carried out [19]. These processes themselves place significant extra demands on ion control — in order to perform entanglement of a pair of logical qubits each encoded on seven physical qubits using a Calderbank, Shaw and Steane code it has been estimated that hundreds of ion movements and two-ion gates need to be performed. Clearly this demonstrates how far ion trap experiments have to go before they can really be considered a viable system for QIP.

More reliable, longer-lived qubits will be necessary if qubit relaxation and dephasing are to be prevented from destroying the quantum information in the system. More specifically, the *internal* states of the ions must be unaffected by long waiting periods, RF heating and magnetic field fluctuations (possibly caused by shuttling of the ions near electrodes). To this end most of the groups that haven't already done so are now moving towards using hyperfine ground states of ions, though this generally makes single and two-qubit gate operations more complicated. The effectively infinite lifetimes of these states makes them invulnerable to long waiting periods and hence very long computations can be carried out. As an example, the Innsbruck team have measured coherence times for Bell states based on the $S_{1/2}$ and $D_{5/2}$ levels of calcium-40 of about one second [62]. This time is limited by the spontaneous decay lifetime, and if they transfer the Bell state into a superposition of the ground state Zeeman sublevels they observe coherence times of about 20 seconds [63].

Besides having inherently longer lifetimes, hyperfine state qubits have another significant advantage. The more complicated Zeeman structure means that at certain magnetic fields the rate of change of the energy shifts for two




Figure 3.5: Partial energy level diagram for ${}^9\text{Be}^+$ showing the Zeeman splitting of the hyperfine ground states, taken from [65]. At various magnetic field strengths the gradients for different levels is equal, so the corresponding transitions (indicated by arrows) have reduced sensitivity to magnetic field fluctuations.

hyperfine levels can be the same (see figure 3.5). This means that the transition frequency between these two levels is to first order independent of magnetic field fluctuations. These so called ‘clock transitions’ have been known about and used by the frequency standards community for some time. Indeed, as long ago as 1991 the NIST group measured qubit coherence times of over 10 minutes using such a transition in ${}^9\text{Be}^+$ [64]. The transition they used was at a rather high magnetic field of about 0.8 tesla, so a Penning trap was required. They are currently investigating the use of clock transitions at field strengths more compatible with radiofrequency traps [65].

In order to combat RF heating effects, groups will be looking to recool ions before operations are carried out on them. In order not to alter the internal states of the qubits, it will most probably be necessary to cool them sympathetically by addressing some other ion species or isotope present in the same trap. This technique has already been implemented by a number of groups using atomic [48, 66, 67] and molecular [68] ions.

As well as improvements in coherence times of qubits, the fidelity of gate operations needs to be increased by several orders of magnitude. It is not quite clear how such a gain will be made. The Innsbruck group believe that by working with calcium-43 and making use of recent technical improvements to their experiment they will be able to achieve a two-qubit gate fidelity of about 98.5%. It is possible that in order to move beyond this kind of level it will be necessary to implement some other gate that is less sensitive to ion heating and laser frequency noise.

Improved ion shuttling techniques will also be required if scalable ion trap structures are to be used. Besides the NIST team, who have been working with segmented traps for some time, other groups such as those at Michigan and MIT have all performed preliminary experiments in various traps [51, 69]. The main problems that have been encountered relate to moving ions around T-junctions and crosses in order to swap them around. It is impossible to generate a pseudopotential minimum that extends along all three directions in a T-junction. Instead there are always three small local maxima near the point where the arms join together [51]. Although this doesn't prevent movement of ions around corners, it does make the process more demanding and can induce ion heating. The optimum way to perform this kind of ion movement will need to be determined before the full potential of multiplexed ion traps can be explored.

3.3.3 Future directions

Ion traps have been chosen by several funding agencies as one of the most promising routes to QIP and hence a large amount of funding has been earmarked for ion trap research. Understandably these funding agencies want to invest in areas that provide the most overall benefit. Hence funding is targeted at aspects that are deemed necessary for ion trap QIP, regardless of an individual research group's choice of ion or gate. Since the choice of ion dictates the laser wavelengths that are required, funding has not yet been channelled into laser technology — even though this currently represents one of the main limitations to gate fidelities.

Instead the principal focus of current investment is in the production of a new generation of scalable miniature traps that are produced using modern microfabrication techniques. In the United States the DTO have initiated a

‘Foundry’ programme, which establishes links between selected ion trap research groups and organisations with established microfabrication facilities, such as Sandia National Laboratories and Bell Laboratories. The scope of the design briefs given to these facilities extends beyond fabrication of the trap structures, to development of the complex control systems that are required to operate them. The first generation of these microtrap ‘chips’ is currently in production, and consists of multi-layered structures with multiple trapping zones on the scale of a few 10s of μm . The European Commission has recently responded with the approval of its own ‘Microtraps’ project. This somewhat smaller-scale operation seeks to involve European microfabricators in the development of miniature trap architectures. Since some research groups are involved in both schemes it seems probable that they will take different approaches to trap production and investigate different aspects of the problems encountered.

The Foundry programme is currently at the stage of producing a number of prototype traps. Although these don’t necessarily have a large number of trapping zones, or the means to shuttle ions between traps, they can be used for the testing of, for example, heat dissipation and material properties. Images of one of these prototype traps can be seen in figure 3.6. Although less directly applicable to QIP, the array of 10^6 cylindrical traps in figure 3.7 shows the extent to which microfabrication techniques can take ion trapping into a new regime. The basic structure of these traps is similar to the traps we use in our experiments, but with a volume 10^{-11} times smaller. It is believed that the trap array shown may find application in mass spectrometry.

Although the resources and techniques are available to fabricate trap structures on very small scales, doing so introduces a whole new set of problems which will need to be overcome. Some of these are simply practical difficulties associated with having such a large number of electrodes on a single chip. How, for example, does one make the required large number of connections between ion-trap chip electrodes inside a vacuum system and the control electronics outside?

A similar practical concern is how to generate the time-varying potentials that are required to shuttle ions between trapping zones. These potentials must be tailored with great precision if we are to avoid excessive heating of the ions, particularly at junctions and points where we wish to move ions around corners. The problem of scheduling these movements will also need to




Figure 3.6: Two-level planar trap being produced by Bell Laboratories. (a) The tweezer-like structure across the diagonal carries the RF, whilst the other electrodes are used for trap biasing and ion shuttling. (b) Close-up showing the RF electrode structure (taken part way through the fabrication process).




Figure 3.7: Images taken [70] of cylindrical traps produced by Sandia National Laboratories for mass spectrometry. (a) Part of the array of 10^6 traps. (b) Cross-section through an individual trap structure.

be addressed, since we know that a great deal of ion movement will become necessary as systems are scaled up beyond a few qubits. Classical algorithms will need to be developed to calculate efficient schedules. It seems likely that much of the control will eventually be performed by conventional electronic systems on board the ion chip, and indeed there are plans to include a CMOS layer on later generations of Foundry chips. Achieving sufficient isolation of the CMOS layer from the large RF voltages present just a few microns away will undoubtedly be a difficult task.

Dissipation of the heat generated by the RF voltages is also a major problem. Excessive heat can lead to failure of the extremely thin insulating layers between electrodes. With such small structures transferring this heat elsewhere can be difficult and hence there may be limits imposed to the RF voltages that can be used.

As well as these easily quantified difficulties there remain more sinister problems to be overcome. The most significant of these as far as current research goes is that of the anomalously large heating rates of ions that have been observed. With ideal electrode materials, the only source of ion heating in an RF trap should be that due to residual Johnson noise in the electrodes. In principle this effect should scale as the inverse square of the distance between an ion in a trap and the nearest electrode. The rate of motional heating due to Johnson noise should be less than a quantum per second when the trap dimension is about $100\mu\text{m}$. Measured heating rates in all experiments so far have been two or more orders of magnitude greater than this, however, and there have been significant variations when different ions are used [47] (see figure 3.8).

One possible culprit is contamination of the electrodes by other metals, leading to the presence of small patch potentials. It is expected that if these patch potentials are small compared to the trap dimension they would cause a heating effect that scaled with the trap dimension, r , as r^{-4} . An experiment has been performed by the Michigan group using an RF ‘endcap’ trap in which the distance between the electrodes and trap centre could be varied continuously over the range 25 to $250\mu\text{m}$ [71]. The measured heating rate was indeed observed to vary almost as r^{-4} (see figure 3.9). It was also found to vary strongly with the temperature of the electrodes, leading to the intriguing possibility that the effect could be quenched through cryogenic cooling. The actual cause of this heating remains something of a mystery, but will need to

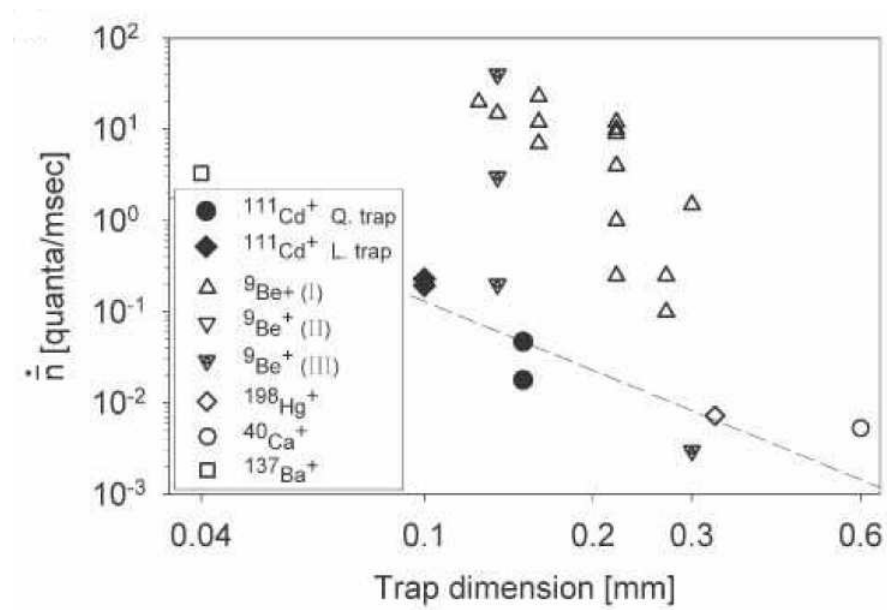


Figure 3.8: Plot taken from [47] showing ion heating rates for a variety of species in traps of different sizes. For comparison, the effect of Johnson noise should be just 10^{-3} quanta per millisecond when the trap dimension is $100\ \mu\text{m}$.




Figure 3.9: Heating rate plot taken from [71], showing the variation in heating rate with trap dimension and temperature.

be understood if microtraps are to operate effectively.

Finally we need to consider the fact that many of the techniques that have been demonstrated so successfully in macroscopic ion traps cannot easily be implemented in a multiplexed architecture. A trivial example is that single ion loading cannot realistically be performed in every trapping zone, so a separate loading area will need to be used. This will need to rapidly generate large numbers of single ions that can be shuttled to other regions of the ion chip as required.

More worrying is the fact that optically addressing individual ions will become increasingly difficult as the number of ions is increased. It has been suggested that MEMS fabricated mirrors could be placed on board the chip to direct beams as required. Such solutions have already been implemented in other applications (for example the optical router developed by Bell Laboratories [72]), but the requirements for ion addressing would be much more stringent. The development of precise single- and multi-qubit gates that require no optical addressing of individual ions and can be scaled up to large numbers of ions would be greatly appreciated by the ion trap community.

Another major difficulty to be faced is that of ion detection and qubit readout. It is interesting that this should represent a problem, since the ability to accurately perform readout has historically been one of the most touted benefits of ion trap QIP. The difficulty lies in the fact that one cannot easily simultaneously detect light from ions in many different parts of an ion chip at once. Although conventional and intensified CCD technology is capable of detecting the fluorescence levels from a single ion, imaging of this light onto the CCD is not trivial. Obtaining the magnification required to resolve different ions necessarily involves having a small field of view and hence we cannot image over a large area without severe aberrations. Arrays of microlenses have been suggested as a possible solution, but once again the fabrication and positioning of such an array would be extremely demanding. Another proposed solution is the use of an array of solid state detectors that are positioned sufficiently close to the ions that a large fraction of the fluorescence can be collected without the need for optical imaging.

It is intriguing to note that the proposed solutions to many of the problems detailed above involve the integration of more and more components onto the ion chip itself. Taking this trend further one might envisage an ion chip on which semiconductor laser sources, imaging optics, detectors and control electronics are all fabricated directly on board. Such a device is an enticing prospect indeed, and represents a stark contrast to the amounts of space and resources that are currently used in large laboratories to manipulate strings of just a few ions.

3.4 Penning trap quantum information processing

In the last section we discussed many of the problems that are currently being faced by the ion trap community in its effort to achieve fault-tolerant quantum computing. Many of these result from deficiencies in the experimental processes, such as the linewidths of lasers used for qubit manipulation. Though these may be difficult to overcome, it is not anticipated that they pose any fundamental limit to our ability to carry out QIP. There are also environmental effects, most notably fluctuations in the ambient magnetic field. The qubit dephasing due to this is a serious issue in current experiments — even with active magnetic field stabilisation the Innsbruck group typically see line-frequency fluctuations of just below a milligauss. Finally there are the

problems associated with the trapping technique itself. RF heating is one of the key issues of the day, since it is observed to increase dramatically as the trap size is reduced. At present there is little understanding of why the observed heating rates are so high.

Instead of tackling these problems directly, let us consider the possibility of avoiding some of them altogether — or at least exchanging them for different problems that may be more easily overcome. The Ion Trapping Group at Imperial College is investigating the possibility of using Penning traps instead of RF traps for QIP. The Penning trap uses only static electric and magnetic fields to confine ions and so suffers no RF heating effects. Though the ion motion in this trap is somewhat complicated, the axial motion is simple harmonic — so all the multi-qubit gates described above can be performed in exactly the same way. Since we require a stable magnetic field, we would certainly envisage using a superconducting magnet for experiments involving coherent qubit manipulation. The level of non-predictable field fluctuations in such magnets can be extremely low — certainly much better than the 1 part in 10^7 required to match the residual fluctuations in the ambient field after active stabilisation. In fact there is a potential gain to be made here, since superconducting coils with certain configurations can provide excellent exclusion of external magnetic flux. An additional benefit is that since we are flexible in the magnetic field that we use for confinement, we can easily choose a field which matches the field-insensitive point of a clock state.

It is anticipated that for the reasons above it ought to be possible to measure significantly reduced heating rates and lower decoherence rates for ion qubits in a Penning trap, compared to those in RF traps. There are, however, some serious disadvantages to working with ions in a Penning trap. The large Zeeman splittings resulting from the high magnetic field mean that it is no longer possible to address more than one Zeeman sublevel with a single (narrow-linewidth) laser. Thus more lasers are required to perform laser cooling and qubit state detection. The rather different ion motion is also a problem — the radial motion is unstable and hence laser cooling is much less efficient. The first of these problems is simply a technical issue in that more lasers must be correctly tuned and stabilised. The second is more fundamental, but as we shall see in the next chapter there are techniques that can be used to overcome it. Another significant difference lies in the way that the motional frequencies scale. Unlike the RF trap, the frequencies in the Penning trap

do not scale with trap size, but instead with magnetic field strength. This means that there is less drive towards trap miniaturisation than there is in the RF trap. Though this means that the frequencies achievable will ultimately be lower, there have been several phase gate proposals whose speed is not dependent on ion motional frequencies. Furthermore, it is possible that the benefits of lower heating and decoherence rates might outweigh the effect of increased computation times.

We are currently working on isolating single $^{40}\text{Ca}^+$ ions and cooling them to the Doppler limit. This involves overcoming the two key difficulties outlined above. Once this is achieved we plan to implement sideband cooling using techniques now standard in other laboratories. We will then be in a position to make precise measurements of the heating rates of single ions in Penning traps for the first time. It will also be possible to measure the lifetimes of the internal states of the ions, and hence investigate the nature and scale of the decoherence mechanisms that are present for ions in Penning trap. The fact that our ion system and plans for sideband cooling are the same as a number of other groups working in RF traps means that we will be in a good position to compare the results obtained in the two types of trap.

3.5 Conclusions

It is by no means certain that a large-scale quantum computer capable of out-performing a classical machine will ever be built. There are a number of promising schemes being investigated, each of which has both benefits and disadvantages compared to the others. That the field of practical quantum computing is so open is testament to the extreme precision of control required. It also seems that the criteria for QC are to some extent mutually incompatible — qubits that interact strongly with each other tend to interact with everything else and therefore decohere rapidly.

Of the alternative implementations proposed, the use of trapped ions remains one of the most promising. All of the requirements for quantum computing can potentially be satisfied and many of the necessary processes have been demonstrated, at least individually. There are, however, many experimental difficulties still to be overcome. The successful operation of two ion-qubit gates has shown that ion trap systems can in principle be used to carry out arbitrary operations on large numbers of qubits. It remains to be seen how

far the limits on gate reliability can be pushed before we meet fundamental barriers.

In order to achieve fault-tolerant quantum computing the key issue will be one of scalability, and much research is already being focused on this. It is apparent that it will be many years before the myriad anticipated difficulties with scaling up ion trap QIP experiments have been resolved. Furthermore it is probable that a whole suite of unforeseen problems will present themselves as research continues. The still increasing scale of the ion trap community, and the interest and funding that it succeeds in attracting, indicates that many researchers believe these problems can indeed be solved. The involvement of professional microfabricators will add a whole new dimension to ion trap research in that it will harness the years of experience that have been gained from microchip manufacturing. It also seems likely that this collaboration with external institutions will be extended so that, for example, professional electronic and software engineers are involved in the development of efficient control systems.

Though it is of course possible that at some stage some other technology will overtake the ion trap as the medium of choice for QIP, the fact that ion traps are the first to enter this advanced phase of development suggests that for the foreseeable future at least they will continue to lead the field.

Chapter 4

Trapping and laser cooling in the Penning trap

We have seen in the previous chapter that we can use time-varying electric fields to confine charged particles. We have also seen, however, that the large amplitudes of the trapping fields in these radio-frequency (RF) traps can result in significant ion heating rates — particularly as the size of the trap is reduced. In this chapter we will shall consider the Penning trap as an alternative to the RF trap that eliminates this problem by using only static fields. Since it is impossible to generate an electric field that has at any point a potential minimum in all three directions, the Penning trap incorporates an additional magnetic field to provide three-dimensional confinement of ions.

This alternative approach to ion trapping results in an ion motion that is very different from that in an RF trap. I shall begin by describing the ‘ideal’ Penning trap and will briefly discuss the effects of non-ideal trap structures. I shall then calculate the motion of an ion in a Penning trap. Next I shall introduce a classical model for Doppler cooling in the Penning trap and show that the cooling rates that can be achieved are greatly restricted by the complicated ion motion. I will then look in detail at a technique called ‘axialisation’, which can be used to improve the efficiency of laser cooling in the Penning trap. Finally I will consider the effect of having more than one ion in the trap in terms of cooling and axialisation.

4.1 Electrode configuration

The electrode structure of an ideal Penning trap is in fact the same as that of the ideal radio-frequency trap. In both cases the ultimate aim is to create a 3-dimensional electric potential ϕ that provides a force on a charged particle that scales linearly with the distance of the particle from the centre of the trap. The force on a charged particle in an electric potential field is given by $-e\nabla\phi$. In order for this force to vary linearly with the distance from trap centre, we therefore require that ϕ varies in each direction as the square of the distance from the centre of the trap. We can write this requirement as

$$\phi = \frac{U}{R_0^2} (ax^2 + by^2 + cz^2).$$

Where U is an applied potential, R_0 is a geometric factor dependent on the electrode structure and a, b and c are constants. Laplace's equation requires that $\nabla^2\phi = 0$ and hence $a + b + c = 0$. If we choose a trap design that is symmetrical in the x and y directions we can set $a = 1, b = 1$ and $c = -2$ so that

$$\phi = U \frac{x^2 + y^2 - 2z^2}{x_0^2 + y_0^2 + 2z_0^2}.$$

Here x_0, y_0 and z_0 are the distances from the centre of the trap to electrodes in the x, y and z directions. In order to generate such a potential the electrodes must follow the equipotential surfaces. Furthermore, if we want the potential at the centre of the trap to be midway between the electrode potentials we must have $r_0^2 = x_0^2 + y_0^2 = 2z_0^2$. An arbitrary equipotential, and therefore the shape of the electrode structure, is given by

$$\phi = \frac{x^2 + y^2 - 2z^2}{r_0^2} = \frac{r^2}{r_0^2} - \frac{z^2}{z_0^2} = \pm 1.$$

These equations represent hyperbolae that have been revolved about the z axis, as shown in figure 4.1. The trap therefore consists of two endcap electrodes held at a potential U relative to a central ring electrode, where U is positive if we are trapping positively charged ions. Typical trap dimensions would be $r_0 = \sqrt{2}z_0 = 5\text{mm}$.

The described electrode configuration provides a trapping potential in the axial direction, but necessarily forms an anti-trap in the radial directions. In




Figure 4.1: Electrode structure of the ‘ideal’ Penning trap. The electrodes are hyperbolae that have been revolved about the z axis, with the radius of the ring electrode r_0 equal to $\sqrt{2}z_0$ where z_0 is the distance from the centre of the trap to the endcaps.

order to confine ions in all three directions we apply a static magnetic field parallel to the z direction. The principle is to force ions into sufficiently small cyclotron loops in the radial plane that they never reach the edges of the trap. The magnetic field strength required to achieve this depends on the trap bias U and the trap dimensions, but typical fields used in experiments would be in the region of 1–6 tesla.

Since the electrodes in the ideal structure derived above are infinite surfaces, all real Penning traps are necessarily non-ideal. Indeed since electrode manufacture and alignment is never perfect and we usually require holes and breaks in them to allow atoms, electrons and lasers to enter the trap, electrodes often deviate significantly from the ideal. As well as the quadratic term in the potential function, we therefore have a whole series including higher-order terms. The two effects of this are that the amplitude of the quadratic term is less than it would be for the same applied potential in an ideal trap and that the trap is not perfectly harmonic. The former effect does not present a problem since we can always compensate with the applied bias. An equivalent R_0^2 value (ideally equal to $r_0^2 + 2z_0^2$) can usually be determined through mea-

measurements of motional frequencies and the quadratic term is then equivalent to that of an ideal trap with the measured R_0^2 .

Anharmonicities can be a more serious problem since they can dramatically reduce the quality factor of the motional frequencies. Many experiments on ions in Penning traps rely on precise measurements of mass [73, 74, 75] and this is done by measuring motional frequencies in the trap. Other experiments [76] couple ion motion to external electronic circuits in order to cool them and this again requires high quality factors for the coupled motions. By adding extra ‘compensation’ electrodes more degrees of freedom are obtained and some of the higher-order terms can be eliminated [77, 78].

For the purposes of quantum information processing, we would require only a few ions to be trapped and they would be cooled to temperatures $\ll 1\text{K}$. Under these conditions ions will typically only occupy a very small region of space close to the centre of the trap, in which the electric potential is quadratic to good approximation in all but the most non-ideal traps. With compensation there is even greater flexibility in electrode structure and other design requirements such as scalability and optical access can take priority. This has led to the development and proposal of a large number of novel Penning trap designs which are far from the ideal shape, with probably the most extreme being planar and 3-dimensional arrays of traps [79, 80].

4.2 Ion motion in the Penning trap

We begin by looking at the motion of an uncooled ion in an ideal Penning trap. For a constant magnetic field \mathbf{B} and electric potential $\phi(\mathbf{r})$ we have for a positively charged particle the Lorentz force

$$\mathbf{F} = e\mathbf{v} \times \mathbf{B} - e\nabla\phi.$$

If \mathbf{B} points along $-\mathbf{z}$ and the electric field is quadrupolar so that $\phi = U_0(2z^2 - x^2 - y^2) / (2z_0^2 + r_0^2)$ the acceleration in the z direction is given by

$$\begin{aligned} \ddot{z} &= -\frac{e}{m} \frac{\partial\phi}{\partial z} \\ &= -\frac{4eU_0}{m(2z_0^2 + r_0^2)} z. \end{aligned} \tag{4.1}$$

This is just simple harmonic motion with a frequency ω_z where

$$\omega_z = \sqrt{\frac{4eU_0}{m(2z_0^2 + r_0^2)}}. \quad (4.2)$$

In the x direction

$$\begin{aligned} \ddot{x} &= -\frac{e}{m}B\dot{y} - \frac{e}{m}\frac{\partial\phi}{\partial x} \\ &= -\frac{e}{m}B\dot{y} + \frac{2eU_0}{m(2z_0^2 + r_0^2)}x \\ &= -\omega_c\dot{y} + \frac{\omega_z^2}{2}x, \end{aligned}$$

where we have defined the true cyclotron frequency $\omega_c = eB/m$. Similarly for the y direction

$$\ddot{y} = \omega_c\dot{x} + \frac{\omega_z^2}{2}y.$$

We can solve these coupled equations for the radial plane by defining a new variable $u = x + iy$ such that

$$\begin{aligned} \ddot{u} &= \ddot{x} + i\ddot{y} \\ &= i\omega_c(\dot{x} + i\dot{y}) + \frac{\omega_z^2}{2}(x + iy) \\ &= i\omega_c\dot{u} + \frac{\omega_z^2}{2}u. \end{aligned} \quad (4.3)$$

Trying a general solution corresponding to circular motion

$$u = u_0 e^{i\omega t}$$

we obtain

$$u_0 e^{i\omega t} \left(\omega^2 - \omega_c \omega + \frac{\omega_z^2}{2} \right) = 0.$$

This has the roots

$$\begin{aligned} \omega &= \frac{\omega_c}{2} \pm \sqrt{\frac{\omega_c^2}{4} - \frac{\omega_z^2}{2}} \\ &= \frac{\omega_c}{2} \pm \omega_1. \end{aligned} \quad (4.4)$$

Where ω_1 has been defined to be equal to the square root term. Taking the positive sign in the sum we obtain a motion with a frequency $\omega'_c = (\omega_c/2) + \omega_1$, called the modified cyclotron. The negative sign gives us the usually much slower magnetron motion at frequency $\omega_m = (\omega_c/2) - \omega_1$. Plots showing axial and radial motion in the Penning trap for typical experimental parameters can be found in figure 4.2. Note that the real part of a solution for ω gives us a motional frequency and the imaginary part tells us its corresponding damping rate. For stable trapping we require no negative damping for either motion and thus ω_1 must be real. We therefore have the condition

$$\frac{\omega_c^2}{4} - \frac{\omega_z^2}{2} > 0$$

so that for a magnetic field B there is an upper limit on the trap bias U_0 given by

$$U_{max} = \frac{eB^2 (2z_0^2 + r_0^2)}{8m}. \quad (4.5)$$

It is interesting to note that we can transform into a rotating reference frame in which the magnetic field is replaced by an effective static quadratic potential. Transforming to a frame v rotating at ω_r relative to the u frame gives us the relations

$$\begin{aligned} u &= v e^{i\omega_r t}, \\ \dot{u} &= (\dot{v} + i\omega_r v) e^{i\omega_r t}, \\ \ddot{u} &= (\ddot{v} + 2i\omega_r \dot{v} - \omega_r^2 v) e^{i\omega_r t}. \end{aligned}$$

Substituting these into our equation of motion (4.3) and simplifying we obtain

$$\ddot{v} = (i\omega_c - 2i\omega_r) \dot{v} + \left(\frac{\omega_z^2}{2} + \omega_r^2 - \omega_r \omega_c \right) v.$$

If we set the rotation frequency to be half the cyclotron frequency then $\omega_r = \omega_c/2$ and

$$\begin{aligned} \ddot{v} &= - \left(\frac{\omega_c^2}{4} - \frac{\omega_z^2}{2} \right) v \\ &= -\omega_1^2 v. \end{aligned} \quad (4.6)$$

The velocity dependent term has disappeared and we are left with independent




Figure 4.2: Plots of ion motion in the Penning trap in (a) the axial direction and (b,c) radial plane. In plot (b) the amplitude of the magnetron motion is larger than that of the modified cyclotron motion. In plot (c) the opposite is true.

simple harmonic motion in each direction. This leads to a general rotational motion in a static harmonic potential, with frequency $\pm\omega_1$. In order to obtain the frequencies in the laboratory frame we need only add $\omega_c/2$ to transform back into the static frame. This gives us the solutions $(\omega_c/2) \pm \omega_1$ as before.

4.3 Laser cooling in the Penning trap

4.3.1 The laboratory frame

We now wish to investigate the effects of laser cooling on the motion of an ion in a Penning trap. There are a great many numerical and analytical models for laser cooling in ion traps that are valid in the limit of various approximations [81, 82, 83]. Here we will consider a classical analytical model developed by Thompson & Papadimitrou [84], which deals with the case of a single ion in a Penning trap that has already been laser cooled to some low temperature.

The situation considered in the model is shown diagrammatically in figure 4.3. A narrow-linewidth laser with frequency ν_l is directed parallel to the x axis, but offset in the y direction by a distance y_0 . For a cold ion the radius of its orbit will be small compared to the waist w of the laser beam, and the Doppler shift due to its motion will be small compared to the linewidth δ_ν of the cooling transition. The force on the ion due to interaction with the laser beam is given by

$$F = \hbar k R,$$

where $\hbar k$ simply gives the momentum of each scattered photon and R is the photon scattering rate. The assumption here is that the re-emission of photons occurs in a random direction and hence the net force due to this is zero. This assumption is only valid when we average over many scattering events. Since the limit on the temperature that Doppler cooling can achieve is governed by photon re-emission processes this model cannot be used to calculate ultimate temperatures. What it can tell us are the average damping rates of the various modes of ion motion.

The photon scattering rate depends on the intensity of the laser, I , and




Figure 4.3: Diagram showing the experimental situation assumed in the cooling model. The Gaussian on the left shows the intensity of the beam as a function of y and indicates the changing intensity experienced by the ion as it oscillates in the trap.

its effective detuning from resonance (including Doppler shifts), Δ , as

$$R = \frac{I}{I_{sat}} \frac{A}{4} \left(\frac{1}{1 + (I/I_{sat}) + (2\Delta/\delta\nu)^2} \right).$$

A in the above is the Einstein A coefficient for the transition, so that for the particular case of calcium (with its accessible metastable $D_{3/2}$ levels) the maximum scattering rate is given by $A/4$. I_{sat} is the saturation intensity and the second term in the final denominator accounts for the effects of power broadening.

For a Gaussian laser beam the intensity seen by an ion as a function of its position in the trap is given by

$$I = I_0 e^{-2((y_0 - y)/w)^2},$$

where w is the waist of the beam.

If ν_0 and λ are the centre frequency and wavelength of the cooling transition respectively, then taking into account Doppler shifts the effective detuning of

the beam from resonance is

$$\Delta = \nu_l - \nu_0 - \frac{\dot{x}}{\lambda}.$$

We can express the scattering rate equation more simply as

$$R = \frac{SA}{4} \left(\frac{1}{1 + S + D^2} \right), \quad (4.7)$$

where we define

$$\begin{aligned} S &= I/I_{sat} \\ &= S_0 e^{-Y^2} \\ Y &= \sqrt{2} (y_0 - y) / w \\ D &= 2(\nu_l - \nu_0 - (\dot{x}/\lambda_0)) / \delta\nu. \end{aligned}$$

The force on the ion due to the laser is always in the x direction and can be separated into three components such that

$$F = F_0 + F_\alpha + F_\beta.$$

Here F_0 is the force due to the photon scattering rate that would occur for a particle at rest in the centre of the trap. F_α is the change in this force if the ion is at some other position in the y direction since the intensity is different. F_β is the change in the force as a result of the Doppler shift due to a change in the ion's velocity. We can think of the effect of the constant force F_0 as being simply to shift slightly the centre of the trapping potential in the direction of the beam. This has no effect on the ion other than the slight change in the centre of its motion and hence we can neglect it. We thus need to calculate only the F_α and F_β terms.

We now make our key assumption that the ion is cold and therefore occupies small radii orbits close to the centre of the trap. This means that the ion only ever occupies a small region of space and hence only samples a small slice of the intensity profile of the laser beam. If this is the case then the force on the ion resulting from its position in the y direction varies approximately linearly with y . The magnitude of this force is scaled by the gradient of the scattering rate in the y direction near the centre of the trap (see figure 4.4).

This gradient can be calculated as

$$\begin{aligned}\frac{\partial R}{\partial y} &= \frac{\partial R}{\partial S} \frac{\partial S}{\partial Y} \frac{\partial Y}{\partial y} \\ &= \frac{AS}{\sqrt{2}w} \frac{Y(1+D^2)}{(1+S+D^2)^2}.\end{aligned}$$

The component of the force on the ion, F_α , is given by

$$\begin{aligned}F_\alpha &= \hbar ky \left(\frac{\partial R}{\partial y} \right)_{y=0, \dot{x}=0} \\ &= \frac{hA}{\sqrt{2}\lambda w} S_0 e^{-Y_0^2} \left(\frac{Y_0(1+D_0^2)}{\left(1 + \left(S_0 e^{-Y_0^2}\right) + D_0^2\right)^2} \right) y.\end{aligned}\quad (4.8)$$

Where we have defined Y_0 and D_0 to be the values of Y and D at $y, \dot{x} = 0$. We also define α such that the constant term in the above equation is equal to $-2\alpha m$ where m is the mass of the ion. Inclusion of the factor 2 proves convenient in later calculations. This gives us simply

$$F_\alpha = -2\alpha m y. \quad (4.9)$$

The other effect of the ion being cold is that its velocity is small and hence the Doppler shift due to its motion can be assumed to be small compared to the linewidth of the cooling transition (figure 4.4). We can therefore use another linear approximation for the change in force F_β as a function of \dot{x} . Here the gradient is

$$\begin{aligned}\frac{\partial R}{\partial \dot{x}} &= \frac{\partial R}{\partial D} \frac{\partial D}{\partial \dot{x}} \\ &= \frac{SAD}{\lambda \delta_\nu (1+S+D^2)^2}.\end{aligned}$$

The component of the force on the ion, F_β , is given by

$$\begin{aligned}F_\beta &= \hbar k \dot{x} \left(\frac{\partial R}{\partial \dot{x}} \right)_{y=0, \dot{x}=0} \\ &= \frac{hA}{\lambda^2 \delta_\nu} S_0 e^{-Y_0^2} \left(\frac{D_0}{\left(1 + \left(S_0 e^{-Y_0^2}\right) + D_0^2\right)^2} \right) \dot{x}.\end{aligned}\quad (4.10)$$




Figure 4.4: *Left:* if the amplitude of an ion's motion is small compared to the waist of the beam, the variation in intensity (and therefore scattering rate) as a function of y is approximately linear. *Right:* similarly, if the ion velocity is small then the Doppler shift as the ion oscillates is small compared to the linewidth of the transition and so the change in scattering rate as a function of \dot{x} is approximately linear.

As before we define β such that the constant terms are equal to $-2\beta m$ and

$$F_\beta = -2\beta m\dot{x}. \quad (4.11)$$

Note that we can change the values of α and β by changing the laser beam offset y_0 and frequency ν_l . However the two are not independent — increasing the beam offset may increase α but will always reduce the overall scattering rate and hence reduce β . Likewise, changing the laser frequency will change both α and β .

In order to assess the validity of our approximations for typical experimental parameters let us consider the case of a single calcium ion being laser cooled on a transition at 397nm. A typical magnetron frequency in one of our Penning traps is 30kHz. If we assume an ion temperature of 50mK then the radius of the magnetron orbit is about $17\mu\text{m}$ and the Doppler shift on the laser frequency is about 8MHz. These figures are to be compared with a beam waist of not less than $30\mu\text{m}$ and a natural linewidth on the 397nm transition of 24MHz. Thus it is reasonable to say that the ion samples a fairly small section of the intensity profile of the laser beam and the lineshape of the transition.

We have shown that the effect of interaction between an ion and a laser beam can be expressed in terms of two forces on the ion which vary linearly with y and \dot{x} . Returning to the model of the Penning trap we presented in section 4.2 we can introduce the effects of laser cooling by adding these two extra forces into the equation of motion in the x direction so that

$$\ddot{x} + \omega_c \dot{y} + 2\beta \dot{x} - \frac{\omega_z^2}{2} x + 2\alpha y = 0, \quad (4.12)$$

$$\ddot{y} - \omega_c \dot{x} - \frac{\omega_z^2}{2} y = 0. \quad (4.13)$$

In order to make these equations more symmetric we must rotate our coordinate system by 45° , into what we shall call the (X, Y) frame (note that the definition of Y is different from that used previously in this section). Thus,

$$X = \frac{x + y}{\sqrt{2}}$$

$$Y = \frac{y - x}{\sqrt{2}}.$$

Writing equations 4.12 and 4.13 in terms of X and Y we obtain

$$\begin{aligned} \frac{1}{\sqrt{2}} \left[\ddot{X} - \ddot{Y} + (\omega_c + 2\beta) \dot{x} + (\omega_c - 2\beta) \dot{y} \right. \\ \left. + \left(2\alpha - \frac{\omega_z^2}{2} \right) x + \left(2\alpha + \frac{\omega_z^2}{2} \right) y \right] = 0 \end{aligned}$$

$$\frac{1}{\sqrt{2}} \left[\ddot{X} + \ddot{Y} - \omega_c \dot{x} + \omega_c \dot{y} - \frac{\omega_z^2}{2} x - \frac{\omega_z^2}{2} y \right] = 0.$$

To separate the terms in \ddot{X} and \ddot{Y} we take the sum and difference of this pair of equations, yielding

$$\ddot{X} + (\omega_c - \beta) \dot{Y} + \beta \dot{X} - \left(\frac{\omega_z^2}{2} - \alpha \right) X + \alpha Y = 0$$

$$\ddot{Y} - (\omega_c + \beta) \dot{X} + \beta \dot{Y} - \left(\frac{\omega_z^2}{2} + \alpha \right) Y - \alpha X = 0.$$

Assuming $\beta \ll \omega_c$ and $\alpha \ll \omega_z^2/2$, which corresponds physically to the idea that the laser interaction is sufficiently weak that it does not affect the oscillation

frequencies, we can approximate these as

$$\ddot{X} + \omega_c \dot{Y} + \beta \dot{Y} - \frac{\omega_z^2}{2} X + \alpha Y = 0$$

$$\ddot{Y} - \omega_c \dot{X} + \beta \dot{X} - \frac{\omega_z^2}{2} Y - \alpha X = 0.$$

We can now solve these equations in the same way as in section 4.2, by changing to a variable $u = X + iY$ such that

$$\ddot{u} + (\beta - i\omega_c) \dot{u} - \left(\frac{\omega_z^2}{2} + i\alpha \right) u = 0. \quad (4.14)$$

Again we look for solutions corresponding to circular motion of the form

$$u = u_0 e^{i\omega t},$$

giving us

$$u_0 e^{i\omega t} \left[\omega^2 - (\omega_c + i\beta) \omega + \frac{\omega_z^2}{2} + i\alpha \right] = 0.$$

This has the roots

$$\begin{aligned} \omega &= \frac{\omega_c + i\beta}{2} \pm \sqrt{\frac{(\omega_c + i\beta)^2}{4} - \frac{\omega_z^2}{2} - i\alpha} \\ &= \frac{\omega_c + i\beta}{2} \pm \sqrt{\omega_1^2 - \frac{\beta^2}{4} + \frac{i\omega_c\beta}{2} - i\alpha}. \end{aligned}$$

Where we have made the substitution $\omega_1^2 = \omega_c^2/4 - \omega_z^2/2$ based on the solutions we obtained for the Penning trap without laser cooling. A Taylor expansion of the square root term with the assumption that $\alpha \ll \omega_1^2, \beta \ll \omega_1^2$ allows us to express the solutions as

$$\omega = \frac{\omega_c + i\beta}{2} \pm \omega_1 \left(1 + \frac{i(\beta\omega_c - 2\alpha)}{4\omega_1^2} \right).$$

The real parts of these solutions give us the motional frequencies and the imaginary parts the corresponding damping rates due to the laser cooling.

Thus taking the positive sign we obtain

$$\begin{aligned}\omega &= \frac{\omega_c}{2} + \omega_1 + \frac{i}{2\omega_1} \left(\beta \left(\frac{\omega_c}{2} + \omega_1 \right) - \alpha \right) \\ &= \omega'_c + i \frac{\beta\omega'_c - \alpha}{2\omega_1},\end{aligned}\tag{4.15}$$

and with the negative sign

$$\begin{aligned}\omega &= \frac{\omega_c}{2} - \omega_1 + \frac{i}{2\omega_1} \left(\alpha - \beta \left(\frac{\omega_c}{2} - \omega_1 \right) \right) \\ &= \omega_m + i \frac{\alpha - \beta\omega_m}{2\omega_1}.\end{aligned}\tag{4.16}$$

The frequencies of motion are therefore the same as we obtained in the absence of laser cooling.

Looking at the damping rates we can immediately see the difficulty in achieving efficient laser cooling in the Penning trap; for positive damping of the modified cyclotron motion we require that $\beta\omega'_c > \alpha$ whereas for the magnetron motion we require $\alpha > \beta\omega_m$. It should be remembered that α and β are related by the overall scattering rate — offsetting the beams further may increase the magnitude of α , but will always reduce the overall scattering rate and therefore β . This means that for efficient cooling of the modified cyclotron we need red-detuned lasers (+ve β) with perhaps a slight offset of the beams in the y direction (-ve α). For good cooling of the magnetron motion we require blue-detuned lasers (-ve β) with a small offset in the $-y$ direction. These requirements are incompatible and therefore it is impossible to achieve *efficient* laser cooling in a conventional Penning trap. We can, however, obtain limited damping of both motions if we satisfy the criterion $\omega_m < \alpha/\beta < \omega'_c$ where both α and β must be positive. This corresponds to red-detuned lasers and an offset of the beam such that it propagates in the same directions as the ions travel around the centre of the trap. This is confirmed if we calculate numerically the cooling rates for the modified cyclotron and magnetron motions as a function of the beam offset y_0 and detuning $\nu_l - \nu_0$. Figure 4.5 shows the results of such a calculation for a calcium ion with trapping parameters that are typical for our experiment. A beam waist of $50\mu\text{m}$ is assumed. Different plots are shown for the case where the beam intensity $I_0/I_{sat} = 0.1$ and 10.0 , corresponding to well below and above the saturation intensity. In both cases the only way to achieve cooling of both motions is to red-detune and offset the laser beam,

though the degree of offset and detuning required differs with beam intensity.

4.3.2 The rotating frame

In the next section we shall be looking at a technique that in principle enables us to achieve efficient laser cooling of both radial motions in the Penning trap. In order to extend our model of laser cooling to include this technique, however, it is necessary to consider once again the motion of the ions in a rotating reference frame. Though the motivation for making this transformation will not yet be clear, we will in this section confirm that the results of our laser cooling model do not depend on our choice of reference frame.

We begin from equation 4.14, the equation of motion in the u (laboratory) frame:

$$\ddot{u} + (\beta - i\omega_c)\dot{u} - \left(\frac{\omega_z^2}{2} + i\alpha\right)u = 0.$$

As in the previous section we transform into a frame v rotating at ω_r relative to the u frame, giving us the relations

$$\begin{aligned} u &= v e^{i\omega_r t}, \\ \dot{u} &= (\dot{v} + i\omega_r v) e^{i\omega_r t}, \\ \ddot{u} &= (\ddot{v} + 2i\omega_r \dot{v} - \omega_r^2 v) e^{i\omega_r t}. \end{aligned}$$

Substituting these into our equation of motion we obtain

$$e^{i\omega_r t} \left[\ddot{v} + (\beta - i\omega_c + 2i\omega_r)\dot{v} - \left(\frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta)\right)v \right] = 0. \quad (4.17)$$

Using the trial solution $v = v_0 e^{i\omega t}$ we find that

$$v_0 e^{i\omega t} \left(\omega^2 - (i\beta + \omega_c - 2\omega_r)\omega + \frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta) \right) = 0.$$

Finding the roots of this we obtain

$$\omega = \frac{\omega_c + i\beta}{2} - \omega_r \pm \sqrt{\frac{(\omega_c + i\beta)^2}{4} - \frac{\omega_z^2}{2} - i\alpha}.$$

This is identical to our result for the lab frame with the exception of the ω_r




Figure 4.5: Cooling rates calculated using the described cooling model. Plots (a) and (b) are the modified cyclotron and magnetron motion cooling rates with a beam intensity of $0.1 \times I_{sat}$. Plots (c) and (d) are the same, but with a beam intensity of $10 \times I_{sat}$. The scale bar is in arbitrary units. Lighter shades indicate larger cooling rates and in each plot the zero cooling rate is given by the (not necessarily straight) contour that evenly bisects the plot.

term. We can therefore write the solutions

$$\omega = \omega'_c - \omega_r + i \frac{\beta \omega'_c - \alpha}{2\omega_1} \quad (4.18)$$

and

$$\omega = \omega_m - \omega_r + i \frac{\alpha - \beta \omega_m}{2\omega_1}, \quad (4.19)$$

where ω'_c and ω_m are the motional frequencies in the lab frame. As expected the damping rates are independent of the choice of frame and we need only add ω_r to the motional frequencies to transform back into the lab frame.

4.4 Axialisation

So far we have shown that it is not possible to efficiently laser cool both the magnetron and modified cyclotron motions simultaneously. This is in fact equally true of various other cooling techniques that have been implemented in Penning traps. One very simple method of cooling hot ions to ambient temperatures that has been used for decades in radio-frequency ion traps is buffer-gas cooling. A low pressure of some inert buffer gas such as nitrogen is admitted into the trap. Collisions between the gas molecules and the ions redistribute the kinetic energy of the particles until they approach a common temperature. The difficulty with buffer gas cooling in the Penning trap is that, because it is unstable, taking energy out of the magnetron motion causes an increase in its amplitude. Thus the magnetron motion is effectively heated and ions rapidly leave the trap.

Wineland and Dehmelt realised that interactions between different trapped ions weakly couple the different motions of the ions. They proposed that an oscillating quadrupole drive at the sum of the magnetron and axial frequencies could provide stronger coupling and hence extend cooling techniques for the axial motion to the magnetron motion [85, 86]. The technique was first demonstrated by the Dehmelt group [87] and has since been considered in some detail (see for example [88] and [89]). It has been used in conjunction with buffer gas cooling [90] to cool the magnetron motion to ambient temperatures.

The technique of cooling the magnetron motion by coupling it to another motion was first demonstrated in conjunction with laser cooling by the Impe-

rial College Ion Trapping Group in order to improve the cooling of magnesium ions [91, 92]. It was found that only by making use of this technique, which we shall henceforth call axialisation, could single ions in a Penning trap be cooled close to the Doppler limit. The electronic structure of magnesium is such that there are no suitable energy levels for a qubit that can be manipulated by a direct optical transition. Thus the decision was made to implement laser cooling and axialisation of calcium ions in order to investigate the feasibility of quantum information processing in Penning traps.

There are a number of different ways of describing the axialisation technique, some of which we shall briefly discuss here. Fundamentally the process must be governed by quantum mechanics, and we can construct a picture of axialisation in this framework. It is also possible, however, to obtain equivalent results using an entirely classical model. Indeed a classical view of axialisation has some benefits in that it enables us to draw analogies with other, more familiar, classical systems.

4.4.1 A quantum picture of axialisation

In quantum mechanical terms the effect of the harmonic potential used to confine ions in a trap is to superpose on each electronic energy level a ladder of motional energy states. Since the magnetron motion is unstable, its corresponding ladder is inverted (see figure 4.6). Applying the axialising quadrupolar drive at the true cyclotron frequency corresponds to a transition from some initial state to one with a lower magnetron motional quantum number but a higher modified cyclotron motional quantum number and back again. If there is some probability that the increased quantum number of the modified cyclotron motion can be reduced by a cooling process rather than just by return to the initial state then there will over time be a net reduction in the quantum number of the magnetron motion.

In the dressed state picture of atom–photon interactions, we can think of the effect of an electromagnetic field being to create additional ‘dressed’ atom states with energies given by $\{E\} \pm nh\nu$. Here $\{E\}$ is the set of energy levels of the non-dressed atom, n is some integer and $h\nu$ is the photon energy of the field. Transitions are described in terms of the coupling that occurs when dressed atom states are near degenerate with some non-dressed state. In this picture, the effect of our quadrupolar drive near the cyclotron frequency is to




Figure 4.6: Energy level diagram showing the magnetron and modified cyclotron motional states (not to scale). The arrows indicate typical transitions due to the axialisation drive.

create a set of energy levels corresponding to dressed motional states. Each of these dressed states will be near degenerate with at least one non-dressed state and hence there will be coupling between them.

All of these states can be considered as quantum harmonic oscillators and the coupling we are interested in is analogous to the exchange of energy between two classical coupled oscillators whose natural frequencies are similar. Besides energy exchange, however, there are other phenomena associated with coupled oscillator systems with near degenerate frequencies. Of particular interest is the avoided crossing, whereby if we couple together two oscillators with similar natural frequencies the normal mode frequencies are observed to shift away from each other. We shall look at these effects later in this chapter.

4.4.2 A classical analogy for axialisation

Earlier in this chapter we used a classical model to calculate the two radial motional frequencies of an ion in an unperturbed Penning trap. We saw that these two frequencies are symmetrically placed about $\omega_c/2$ and that the model is unaffected by a change in reference frame. Given these results we now consider the special case of a reference frame rotating with a frequency $\omega_c/2$

in the laboratory frame. In this frame the two radial motions have the same frequency ω_1 , but have different senses of rotation. The overall radial motion is of course simply the sum of the modified cyclotron and magnetron motions. The sum of two counter-rotating circular motions is in general elliptical, with the half length of the ellipse's two axes corresponding to the sum and difference of the amplitudes of the two motions. Figure 4.7 shows ion motion in the rotating frame for various relative motional amplitudes. Special cases occur when we have either pure modified cyclotron or pure magnetron motion, so that we see circular motion in one direction or the other, or if the amplitudes for the two motions are the same, giving us motion along a straight line.

This description of ionic motion in the rotating frame will be instantly familiar if we recall that it is identical to the general classical description of the polarisation state of an electromagnetic wave. This is again the sum of two counter-rotating circular polarisation states, and the resultant polarisation depends on the relative amplitudes of the two components. It is interesting to take this analogy further by considering the effect of retardation on such a polarisation state. The different refractive indices along different axes in a retarder result in the accumulation of a phase difference between the components of the wave along each axis. Assuming the axes of the retarder are aligned at some angle to the major axis of the polarisation state, we will see a variation in the polarisation state as the wave passes through the retarder. With an axis alignment of 45° this change corresponds to pure left-handed circular polarisation changing to linear polarisation followed by pure right-handed polarisation and back again.

As we will shortly show in detail, an oscillating quadrupolar drive can be deemed static in a frame rotating at $\omega_c/2$ providing its oscillation frequency is ω_c . When added to the radially symmetrical trapping potential, a static quadrupolar field has the effect of increasing the trap frequency along one axis, whilst reducing it along the other. This leads to a 'fast' and 'slow' axis analogous to that in the case of the electromagnetic wave passing through a retarder. We can therefore assume an equivalent result whereby the motion of the ion oscillates from pure magnetron motion to pure modified cyclotron motion and back again. Figure 4.8 shows a plot of this motion in the laboratory frame. If we introduce efficient cooling of the modified cyclotron motion we find that axialisation extends this cooling to the magnetron motion.




Figure 4.7: Plots showing ion motion in the rotating frame for the cases of (a) pure modified cyclotron motion, (b) pure magnetron motion, (c) amplitude of modified cyclotron $>$ magnetron amplitude, (d) amplitudes of the two motions are the same, (e) amplitude of magnetron motion $>$ modified cyclotron amplitude.




Figure 4.8: Simulation of the motion of an ion in the laboratory frame in the presence of axialisation. The ion initially has a large magnetron radius and a small modified cyclotron radius, but over time the modified cyclotron radius becomes smaller and the magnetron radius becomes larger. In the absence of laser cooling the system returns back to its initial state after the same amount of time has elapsed again.

4.4.3 A model for laser cooling in the presence of axialisation

We have extended the laser cooling model discussed in section 4.3 to include the effects of axialisation. This work was led by and carried out in collaboration with R.C. Thompson, and is presented for the first time here.

Consider a radially quadrupolar potential oscillating with frequency ω_a in the laboratory frame,

$$\phi_a = V_0 \frac{x^2 - y^2}{2r_0^2} \cos(\omega_a t).$$

In polar coordinates this becomes

$$\begin{aligned} \phi_a &= \frac{V_0}{2} \frac{r^2}{r_0^2} (\cos^2 \theta - \sin^2 \theta) \cos(\omega_a t) \\ &= \frac{V_0}{2} \frac{r^2}{r_0^2} \cos(2\theta) \cos(\omega_a t) \\ &= \frac{V_0}{4} \frac{r^2}{r_0^2} (\cos(2\theta + \omega_a t) + \cos(2\theta - \omega_a t)), \end{aligned}$$

which is equivalent to two counter-rotating quadrupole fields. We can now transform into the frame rotating at a frequency ω_r by making the substitutions $\theta' = \theta - \omega_r t$ and $r' = r$, so that in this frame

$$\phi'_a = \frac{V_0}{4} \frac{r'^2}{r_0^2} (\cos(2\theta' + (2\omega_r - \omega_a)t) + \cos(2\theta' + (2\omega_r + \omega_a)t)). \quad (4.20)$$

Clearly if $\omega_r = \omega_a/2$ then one of these components is static and the other rotates with frequency $-4\omega_r$. We anticipate that the oscillating quadrupole field will give us a coupling between motions and hence we expect to see effects only when the drive frequency is resonant with a sum or difference of motional frequencies. In a frame rotating at $\omega_c/2$, the sum of the motional frequencies is zero. Thus the static component of ϕ_a will be resonant, but the (rapidly) rotating component will in general not be. We therefore ignore the rotating component and calculate the acceleration in the v frame due to the static

component. If x' and y' are cartesian coordinates in the v frame then

$$\begin{aligned}\phi'_a &= \frac{V_0}{4r_0^2} r'^2 \cos(2\theta') \\ &= \frac{V_0}{4r_0^2} r'^2 (\cos^2 \theta' - \sin^2 \theta') \\ &= \frac{V_0}{4r_0^2} (x'^2 - y'^2).\end{aligned}$$

The forces on a positively charged ion can then be calculated as

$$\begin{aligned}F_{x'} &= -e \frac{\partial \phi'_a}{\partial x'} \\ &= -\frac{eV_0}{2r_0^2} x' \\ F_{y'} &= -e \frac{\partial \phi'_a}{\partial y'} \\ &= +\frac{eV_0}{2r_0^2} y' \\ F_v &= F_x + iF_y \\ &= -\frac{eV_0}{2r_0^2} (x' - iy') \\ &= -\frac{eV_0}{2r_0^2} v^*.\end{aligned}\tag{4.21}$$

We introduce $\epsilon = eV_0/2mr_0^2$ so that the force due to axialisation is simply given by $-\epsilon m v^*$.

Returning to our model for laser cooling in the rotating frame we can now add in the effect of axialisation so that equation 4.17 becomes

$$\ddot{v} + (\beta - i\omega_c + 2i\omega_r) \dot{v} - \left(\frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta) \right) v + \epsilon v^* = 0.\tag{4.22}$$

Our trial solution now corresponds to elliptical motion with uniform decay, because the rotational symmetry is broken by the v^* term. This is a superposition of two circular motions in opposite senses with arbitrary amplitudes. If the amplitudes are the same we have motion along a straight line. If one amplitude is zero we return to the case of circular motion.

$$v = Ae^{i\omega t} + Be^{-i\omega^* t}.$$

Note the use of ω^* in the second term — we want the same damping rate (real part of the exponent) for both terms even though they travel in different senses. Substituting our trial solution into equation 4.22 we obtain

$$\begin{aligned}
 & e^{i\omega t} \left[A\omega^2 - A(i\beta + \omega_c - 2\omega_r)\omega \right. \\
 & \quad \left. + A\left(\frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta)\right) - \epsilon B^* \right] \\
 & + e^{-i\omega^* t} \left[B\omega^{*2} + B(i\beta + \omega_c - 2\omega_r)\omega^* \right. \\
 & \quad \left. + B\left(\frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta)\right) - \epsilon A^* \right] = 0.
 \end{aligned}$$

It is straightforward to show that if $c \cdot f(x) + c^* \cdot g(x) = 0$, where c has both real and imaginary parts, then $f(x)$ and $g(x)$ must both be equal to zero. Therefore

$$\begin{aligned}
 & A \left[\omega^2 - (i\beta + \omega_c - 2\omega_r)\omega + \left(\frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta)\right) \right] - \epsilon B^* = 0 \\
 & B \left[\omega^{*2} + (i\beta + \omega_c - 2\omega_r)\omega^* + \left(\frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta)\right) \right] - \epsilon A^* = 0.
 \end{aligned}$$

Previously we saw that in order to be in a frame in which the axialising quadrupole drive is static we must have $\omega_r = \omega_a/2$. We therefore define Δ to be half the detuning of the axialisation drive from the cyclotron frequency, so that $\omega_a = \omega_c + 2\Delta$. Thus the frame rotation frequency is $\omega_r = \omega_c/2 + \Delta$. We also express the variable δ to be the (complex) shift in the frequencies due to the presence of laser cooling, the axialising drive and the shift in the frame rotation frequency Δ , so that $\omega = \omega_1 + \delta$. Our equations for ω can therefore be rewritten as

$$\begin{aligned}
 & A \left[2\delta\omega_1 - i\beta\omega_1 + 2\Delta\omega_1 + i\frac{2\alpha - \beta\omega_c}{2} + (\Delta + \delta)^2 - i\beta(\Delta + \delta) \right] - \epsilon B^* = 0 \\
 & B \left[2\delta^*\omega_1 + i\beta\omega_1 - 2\Delta\omega_1 + i\frac{2\alpha - \beta\omega_c}{2} + (\Delta - \delta^*)^2 - i\beta(\Delta - \delta^*) \right] - \epsilon A^* = 0.
 \end{aligned}$$

The $(\Delta \pm \delta)^2$ and $\beta(\Delta \pm \delta)$ terms are small relative to those containing ω_1 . If we ignore these and take the complex conjugate of the latter equation the

pair of equations becomes

$$A \left[(2\delta\omega_1 - i\beta\omega_1) + \left(2\Delta\omega_1 + i\frac{2\alpha - \beta\omega_c}{2} \right) \right] - \epsilon B^* = 0 \quad (4.23)$$

$$B^* \left[(2\delta\omega_1 - i\beta\omega_1) - \left(2\Delta\omega_1 + i\frac{2\alpha - \beta\omega_c}{2} \right) \right] - \epsilon^* A = 0. \quad (4.24)$$

Solving equation 4.24 for A and substituting into 4.23 gives

$$B^* \left[(2\delta\omega_1 - i\beta\omega_1)^2 - \left(2\Delta\omega_1 + i\frac{2\alpha - \beta\omega_c}{2} \right)^2 - |\epsilon|^2 \right] = 0.$$

To simplify this we define $M = (2\alpha - \beta\omega_c)/2\omega_1$, a term related to the laser cooling parameters. In order that we can find the real and imaginary parts of this expression we also set $\delta = \delta_0 + i\gamma_0$ so that

$$B^* \omega_1^2 \left[(2\delta_0 + 2i\gamma_0 - i\beta)^2 - (2\Delta + iM)^2 - \frac{|\epsilon|^2}{\omega_1^2} \right] = 0.$$

All of the constants and variables are now real. Evaluating the squared terms and separating out real and imaginary parts we obtain

$$4\delta_0^2 - 4\gamma_0^2 - \beta^2 + 4\beta\gamma_0 - 4\Delta^2 + M^2 - \frac{|\epsilon|^2}{\omega_1^2} = 0 \quad (4.25)$$

and

$$8\delta_0\gamma_0 - 4\delta_0\beta - 4\Delta M = 0. \quad (4.26)$$

From equation 4.26 we can write down the cooling rates

$$\gamma_0 = \frac{\beta}{2} + \frac{\Delta M}{2\delta_0}. \quad (4.27)$$

Substituting γ_0 back into equation 4.25 we find

$$\delta_0^2 - \left(\Delta^2 - \frac{M^2}{4} + \frac{|\epsilon|^2}{4\omega_1^2} \right) - \frac{\Delta^2 M^2}{4\delta_0^2} = 0,$$

which if we define N to be equal to the term in brackets gives us the quartic equation

$$\delta_0^4 - N\delta_0^2 - \frac{\Delta^2 M^2}{4} = 0.$$

This has four roots

$$\delta_0 = \pm \frac{1}{\sqrt{2}} \sqrt{N \pm \sqrt{N^2 + \Delta^2 M^2}}. \quad (4.28)$$

Since we have defined δ_0 to be real we must take the positive sign for the inner square root and so we can eliminate two of these solutions. Recalling that the motional frequencies in the laboratory frame are given by $\omega = \omega_r \pm (\omega_1 + \delta_0)$ and that $\omega_r = \omega_c/2 + \Delta$ we can see that having two solutions $\pm\delta_0$ gives us components at four frequencies, given by

$$\omega = \omega'_c + \Delta \pm \delta_0 \quad (4.29)$$

and

$$\omega = \omega_m + \Delta \mp \delta_0. \quad (4.30)$$

Returning to the quantum mechanical picture we briefly introduced earlier in this section, we can draw a comparison between these extra motional frequencies and the dressed states we discussed. The extra solution near the modified cyclotron motion can be thought of as the magnetron motion dressed by the axialising field and likewise the extra motion near the magnetron frequency can be thought of as the dressed modified cyclotron motion. It should be borne in mind that we have yet to consider the amplitudes of each of these motions. In the limit of no axialisation we would expect the amplitudes of two of these motions to fall to zero so that we return to the result derived in section 4.3. Note also the flipped \pm sign for the magnetron motion. We shall see the significance of this shortly, when we consider the damping rates corresponding to each solution.

We look now at how the frequency shifts δ_0 vary with the axialisation and laser parameters. Let us first restate together the key results we have derived:

In the rotating frame:

$$\delta_0 = \pm \frac{1}{\sqrt{2}} \sqrt{N + \sqrt{N^2 + \Delta^2 M^2}}$$

$$\gamma_0 = \frac{\beta}{2} + \frac{\Delta M}{2\delta_0},$$

where

$$M = \frac{(2\alpha - \beta\omega_c)}{2\omega_1}$$

$$N = \Delta^2 - \frac{M^2}{4} + \frac{|\epsilon|^2}{4\omega_1^2}$$

and the frequency shifts in the laboratory frame are given by $\delta_0 + \Delta$.

We would anticipate that if we set the amplitude of the axialising quadrupole drive to zero we will obtain in the laboratory frame unshifted frequencies, ie. ω'_c and ω_m . We confirm this by setting ϵ to zero. In this case $N = \Delta^2 - M^2/4$ and we can calculate δ_0 as

$$\begin{aligned} \delta_0 &= \pm \frac{1}{\sqrt{2}} \sqrt{\Delta^2 - \frac{M^2}{4} + \sqrt{\left(\Delta^2 - \frac{M^2}{4}\right)^2 + \Delta^2 M^2}} \\ &= \pm \frac{1}{\sqrt{2}} \sqrt{\Delta^2 - \frac{M^2}{4} + \sqrt{\left(\Delta^2 + \frac{M^2}{4}\right)^2}} \\ &= \pm \Delta. \end{aligned}$$

In the laboratory frame the frequency shifts for each motion are therefore given by $+2\Delta$ or zero and if we assume for the present that the amplitudes for the former solution are zero we are left, as expected, with unshifted frequencies. In agreement with the results from section 4.3 we also find that this is true regardless of the laser cooling parameters. Putting $\delta_0 = \pm\Delta$ into our equation for the damping rate we obtain

$$\begin{aligned} \gamma_0 &= \frac{\beta \pm M}{2} \\ &= \frac{\beta\omega_1 \pm \left(\alpha - \frac{\beta\omega_c}{2}\right)}{2\omega_1}, \end{aligned}$$

so either

$$\gamma_0 = \frac{\alpha - \beta\omega_m}{2\omega_1}$$

or

$$\gamma_0 = \frac{\beta\omega'_c - \alpha}{2\omega_1}.$$

For the modified cyclotron motion it is the latter solution that corresponds to the unshifted motional frequency in the laboratory frame. We know from the flipped \pm sign in equation 4.30 that the solutions for δ_0 are reversed for the magnetron motion, and hence it is the former damping rate that applies to the unaltered frequency. Thus the damping rates we obtain for both motions are the same as those from our model without axialisation.

Interestingly, if we now add in some non-zero amplitude for the axialisation quadrupole drive, ϵ , we find that the behaviour of the system in terms of frequency shifts does become dependent on the laser parameters. The critical factor is the balance of the three terms in our expression for N . Note that if the detuning of the axialising drive is made sufficiently large the Δ^2 term in N will always dominate and δ_0 will reduce once again to $\pm\Delta$, as it did for the case of no drive at all. This demonstrates that axialisation is very much a resonant effect and the drive frequency must be close to the true cyclotron frequency if it is to work. The effective width of the resonance is related to the amplitude of the quadrupole drive and the strength of the laser cooling. Close to resonance Δ is small and the frequency behaviour depends on the relative magnitudes of ϵ/ω_1 and $|M|$. We shall now consider in turn three cases relating to $(\epsilon/\omega_1)^2 \ll M^2$, $(\epsilon/\omega_1)^2 \sim M^2$ and $(\epsilon/\omega_1)^2 \gg M^2$.

$(\epsilon/\omega_1)^2 \ll M^2$: In this regime $N \sim \Delta^2 - M^2/4$ and we are to good approximation in the previously considered case of no axialisation. Numerically calculated motional frequency shifts in the rotating and laboratory frames are plotted for the magnetron motion in figure 4.9, along with the corresponding damping rates. The motion which has unshifted frequency in the laboratory frame has the damping characteristics of the magnetron motion in an unperturbed Penning trap. The other motion is the modified cyclotron motion ‘dressed’ by the quadrupole field. It has the characteristics of the modified cyclotron motion, but in this regime we expect its amplitude to be small. There is no visible coupling of the two damping rates. Note the swapping of the curves at $\Delta = 0$. The differently shaded curves correspond to taking either the positive or negative solution for δ_0 and there is a sudden change in which

root represents which motion. This is not as strange as it may first appear — if we think of the extra motion as being the dressed modified cyclotron motion then if the dressing field has too low a frequency the dressed motion will be below the magnetron frequency, if it is too high the dressed motion will have a higher frequency.

$(\epsilon/\omega_1)^2 \sim M^2$: In this regime we begin to see more interesting effects occurring. Although the Δ^2 term always dominates the expression for N , it is not always the dominant term in the equation for δ_0 . If we write $N = \Delta^2$ then

$$\delta_0 = \pm \frac{1}{\sqrt{2}} \sqrt{\Delta^2 + \sqrt{\Delta^2 (\Delta^2 + M^2)}}.$$

In the region where Δ is small compared to M this expression becomes

$$\begin{aligned} \delta_0 &= \pm \frac{1}{\sqrt{2}} \sqrt{\Delta (\Delta + \sqrt{\Delta^2 + M^2})} \\ &\approx \pm \sqrt{\left| \frac{\Delta M}{2} \right|}. \end{aligned}$$

Therefore for small Δ we expect the shift in motional frequencies to vary as $\pm\sqrt{\Delta}$ in the rotating frame. As Δ becomes larger compared with M the curves will again tend towards $\pm\Delta$.

Figure 4.10 is a plot of the frequency shifts and damping rates calculated numerically with $(\epsilon/\omega_1)^2 = M^2$. With the exception of ϵ , all the parameters are the same as those used to plot figure 4.9. The behaviour is exactly as we expected. In the rotating frame we see $\sqrt{\Delta}$ behaviour near resonance tending towards a linear dependence as we become further detuned. Transforming into the laboratory frame generates an unusual swirling in the shape of the curve.

Looking at the damping rate curves we see for the first time the useful property of the axialisation technique. Near resonance the damping rates of the two motions converge and we are left with an intermediate damping rate for both motions. The benefit of coupling together the magnetron and modified cyclotron motions like this is not just that we increase the magnetron damping at the expense of the modified cyclotron, but that we no longer need to maintain a large beam offset α and small frequency gradient β in order to prevent heating of the magnetron motion. The average cooling rate of the two




Figure 4.9: Motional frequency shifts (solid lines) and damping rates (dashed lines) relative to the magnetron motion in the rotating frame (a) and laboratory frame (b) for the case where $(\epsilon/\omega_1)^2 \ll M^2$. The equivalent situation for the modified cyclotron motion differs only in that the damping rate curves are swapped around. Plotted parameters are for Ca^+ at 0.98 tesla with $(\epsilon/\omega_1)^2 = 0.01$, $|M| = 1$ and $\alpha/\beta = 100$.

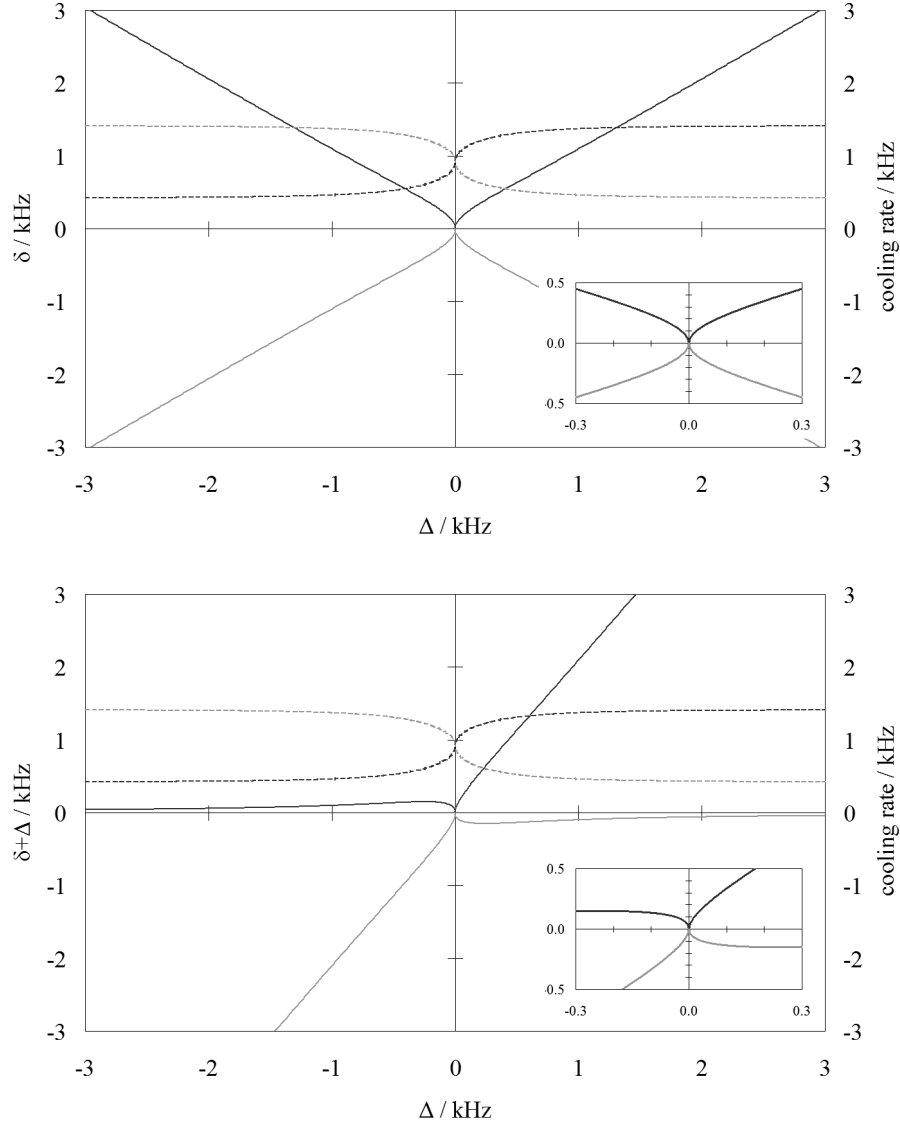


Figure 4.10: Motional frequency shifts (solid lines) and damping rates (dashed lines) relative to the magnetron motion in the rotating frame (a) and laboratory frame (b) for the case where $(\epsilon/\omega_1)^2 = M^2$. The inset plots are enlargements of the central regions in which the behaviour for small Δ can be more clearly seen. As before the equivalent situation for the modified cyclotron motion would differ only in that the damping rate curves are swapped around. Plotted parameters are for Ca^+ at 0.98 tesla with $(\epsilon/\omega_1)^2 = 1$, $|M| = 1$ and $\alpha/\beta = 100$.

motions is

$$\begin{aligned}\gamma_{ave} &= \frac{1}{2} \left(\frac{\beta}{2} \pm \frac{\Delta M}{2\delta_0} + \frac{\beta}{2} \mp \frac{\Delta M}{2\delta_0} \right) \\ &= \frac{\beta}{2}.\end{aligned}\tag{4.31}$$

The α dependence has gone and the optimum case is represented by maximising β . Furthermore reducing the beam offset increases the beam intensity seen by the ions and so increases the overall scattering rate and hence β , giving us yet higher mean cooling rates. We must be careful, however, since this also increases the sensitivity to the frequency of the axialisation drive. If we set up our laser cooling to give large average damping rates but we are not properly tuned to resonance we will rapidly heat the magnetron motion.

The shape of the cooling rate curves near resonance is given by substituting $\delta_0 = \pm\sqrt{|\Delta M/2|}$ into our equation for γ_0 . This yields

$$\gamma_0 \approx \frac{\beta}{2} \pm \sqrt{\left| \frac{\Delta M}{2} \right|}.$$

$(\epsilon/\omega_1)^2 \gg M^2$: In this regime we see the same damping rate coupling behaviour as in the previous case, although the range of detunings in which we see significant coupling is larger. What is different, however is that as soon as $\epsilon/\omega_1 > |M|$ the frequency shifts no longer cross at zero. Instead we see an avoided crossing — the frequency shifts approach each other as before for large Δ but near resonance they split apart. The damping rates still swap over at $\Delta = 0$ as required, though at no point do the motions share the same frequency. Figure 4.11 shows plots of the frequency shifts and damping rates calculated numerically using the same parameters as in the previous cases, but with $(\epsilon/\omega_1)^2 = 10$ ($|M| = 1$ as before).

If $(\epsilon/\omega_1)^2 \gg M^2$ we can make the approximation that $N = \Delta^2 + \epsilon^2/(4\omega_1^2)$. We can then calculate the frequency shifts in the rotating frame as

$$\begin{aligned}\delta_0 &\approx \pm \frac{1}{\sqrt{2}} \sqrt{N + \sqrt{N^2}} \\ &= \pm \sqrt{N} \\ &= \pm \sqrt{\Delta^2 + \frac{\epsilon^2}{4\omega_1^2}}.\end{aligned}$$




Figure 4.11: Motional frequency shifts (solid lines) and damping rates (dashed lines) relative to the magnetron motion in the rotating frame (a) and laboratory frame (b) for the case where $(\epsilon/\omega_1)^2 \gg M^2$. Once again the situation for the modified cyclotron motion differs only in that the damping rate curves are swapped around. Plotted parameters are for Ca^+ at 0.98 tesla with $(\epsilon/\omega_1)^2 = 10$, $|M| = 1$ and $\alpha/\beta = 100$.

The frequency shifts at $\Delta = 0$ are simply equal to $\epsilon / (2\omega_1)$.

Avoided crossings resulting from the coupling together of motional modes of ions have been reported before by Cornell *et al.* [93]. They were working on Fourier Transform Ion Cyclotron Resonance spectroscopy and were coupling the magnetron motion to the axial motion in an attempt to extend resistive cooling of ions to the radial motions. Their measurements of the motional frequencies were made by taking the Fourier transform of the electronic signal induced on the endcaps by ion motion after a driving pulse close to some motional frequency. In some of their traces they noticed two distinct frequencies close together and plotting them against the strength of the coupling field they observed an avoided crossing similar to that we have calculated. Their plot is reproduced in figure 4.12. By a different method they derived the same equation we have formulated for the frequency shifts in the regime where the coupling rate is stronger than the cooling rate. The result here extends theirs into the region of weaker coupling. Furthermore the work by Cornell *et al.* did not consider the phase behaviour of the ions' motion. Since motional phase behaviour forms an integral part of the experimental techniques we use to measure ion cooling rates, and so investigate the efficacy of axialisation, we shall investigate it in our model. We shall also be looking at the relative amplitudes of the different motions by considering their ability to couple to an external driving field.

4.4.4 Amplitude and phase response to a driving field

Thus far we have calculated the motional frequencies and damping rates that we expect to observe for an ion being laser cooled in the Penning trap in the presence of axialisation. We noted, however, that the extra motional frequencies arising from the axialising field exist even if the amplitude of the drive, ϵ , is zero and suggested that these motions would always have zero amplitude. The solutions we obtained were of the form

$$v = Ae^{i\omega t} + Be^{-i\omega^* t}.$$

Using our model we have determined both real and imaginary parts of ω , which correspond to the frequencies and damping rates, but we can say little about the relative amplitudes of the motions. The motions begin with some initial amplitude and tend continuously towards zero, as expected for a damped




Figure 4.12: Result reproduced from [93] showing an avoided crossing resulting from coupling of the axial and magnetron motions in a Penning trap as measured with a Fourier Transform Ion Cyclotron Resonance technique.

harmonic oscillator. Essentially this is another way of saying that our model provides us with cooling rates but tells us nothing about the ultimate limits of the cooling.

The extent to which different motions can be driven with an external force tells us directly about the ability of an ion to exhibit these motions. Hence we can investigate the motional amplitudes by extending the model to include the effects of a driving force. Experimentally this corresponds to an additional dipole field oscillating at a frequency close to that of some motion, so that with a sufficiently large driving force an ion will oscillate at the frequency of the drive. The amplitude of the motion depends on how closely the driving field is tuned to one of the ion's natural motional frequencies. When the drive frequency is exactly resonant with a natural frequency of the ion's motion the amplitude of the response can be extremely large. The phase of the ion's motion relative to the driving field also depends on the tuning of the driving field — the ion will change smoothly from being in phase with the drive to being in antiphase with it as we tune the drive from below to above resonance. The width of this resonance in the response of an ion to a driving field is directly related to the damping present in the system. As we will see in chapter 6 this is useful since it enables us to determine the laser cooling rates in our experiment.

A dipole field oscillating at frequency ω_d in the laboratory frame can be thought of as a pair of static dipole fields rotating at the oscillation frequency, since

$$2F \cos \omega_d t = F (e^{i\omega_d t} + e^{-i\omega_d t}).$$

If the driving frequency is close to the magnetron frequency, so that $\omega_d = \omega_m + \delta'$, then in a frame rotating at a frequency $\omega_r = (\omega_c/2) + \Delta$ the driving force is given by

$$F \left(e^{i(-\omega_1 + \delta' - \Delta)t} + e^{i(-\omega_c + \omega_1 - \delta' - \Delta)t} \right).$$

The first term here is nearly resonant with the magnetron frequency (the detuning in the rotating frame being $\delta' - \Delta$, which we shall define to be δ). The second term will in general be far from resonance and can be ignored. We define ω to be the frequency of the driving field in the rotating frame, so that $\omega = -\omega_1 + \delta$.

If we now add this driving force into the equation of motion in the rotating

frame for an ion being laser cooled with axialisation (equation 4.22) we obtain

$$\ddot{v} + (\beta - i\omega_c + 2i\omega_r)\dot{v} - \left(\frac{\omega_z^2}{2} + i\alpha + \omega_r^2 - \omega_r(\omega_c + i\beta)\right)v + \epsilon v^* = Fe^{i\omega t}. \quad (4.32)$$

We can now look for steady state solutions corresponding to the situation once the system has reached equilibrium. These will have the form

$$v = Ae^{i\omega t} + Be^{-i\omega t},$$

where we now have the situation that ω is real and is some known applied frequency rather than a variable. We proceed in an identical manner to before until we reach the equivalent of equations 4.23 and 4.24:

$$A \left[-2\delta\omega_1 + i\beta\omega_1 - 2\Delta\omega_1 - i\alpha + i\beta\frac{\omega_c}{2} \right] + \epsilon B^* = F, \quad (4.33)$$

$$B^* \left[-2\delta\omega_1 + i\beta\omega_1 + 2\Delta\omega_1 + i\alpha - i\beta\frac{\omega_c}{2} \right] + \epsilon^* A = 0. \quad (4.34)$$

The first of these equations is the component rotating in the same sense as the driving force (in this case the direction of the magnetron motion), with the second equation describing the component of rotation in the opposite sense. Bearing in mind that δ is now a defined constant detuning we can group together the constant terms in square brackets as C_m and C_c respectively. This gives us simply

$$AC_m + \epsilon B^* = F$$

$$B^* C_c + \epsilon^* A = 0.$$

We can easily solve these equations to find A and B , giving us

$$A = \frac{FC_c}{C_m C_c - \epsilon\epsilon^*} \quad (4.35)$$

$$B = \frac{F\epsilon}{\epsilon\epsilon^* - C_m^* C_c^*}. \quad (4.36)$$

Both A and B are therefore functions of the amplitudes and detunings of the excitation and axialisation drives (F , δ , ϵ and Δ), the laser parameters (α and β) and the trap frequencies (ω_c and ω_1).

A and B are complex quantities and we can obtain the amplitude and

phase response of the ion's motion relative to the excitation drive by taking their moduli and arguments respectively. Figure 4.13 shows typical plots of the amplitude and phase response (in the laboratory frame) when the axialisation amplitude is small. In each case the solid line indicates the response of the motion being driven, and hence is the response we can observe experimentally. The dashed lines correspond to the response of the component of the motion we are not directly driving (ie. the counter-rotating motion) and would be difficult to measure experimentally. Looking at the amplitude of A we can see that the response is large when the excitation drive is resonant with the motional frequency that is least shifted, but is small when the drive is resonant with the extra frequency that is generated as a result of axialisation. This validates our previous assertion that the extra frequencies we observe in our model have very low amplitudes when the drive is small or the detuning is very large. The phase response of A is similar to that which we would expect if there were no axialisation, but with a small dip at the extra frequency. Again if the axialisation drive has zero amplitude this dip disappears altogether and we return to the unaxialised case.

If the axialisation amplitude is larger, the amplitude responses to the drive at each of the motional frequencies become more similar. At the same time the extra dip in the phase response increases in size until it covers the full π radians shift we expect as we pass through resonance. Figure 4.14 shows amplitude and phase plots from our model when the axialisation amplitude is large.

In order to gain a better understanding of the relationships between the dipole excitation drive and the various motions, it can be informative to plot the amplitudes A and B on an Argand diagram (see figure 4.15). On such a plot the positive real axis (having zero phase) corresponds to the direction of the applied dipole drive. The motional amplitudes are seen to trace out near-circular paths on the Argand diagram. A begins in phase with the drive and traces out two loops, always lagging behind the excitation drive. B on the hand is almost π radians ahead of the drive for large negative detunings. It traces out a path that takes it into phase, where there is an abrupt change in direction so that during the second loop it lags further behind and ends up once again in antiphase with the drive.

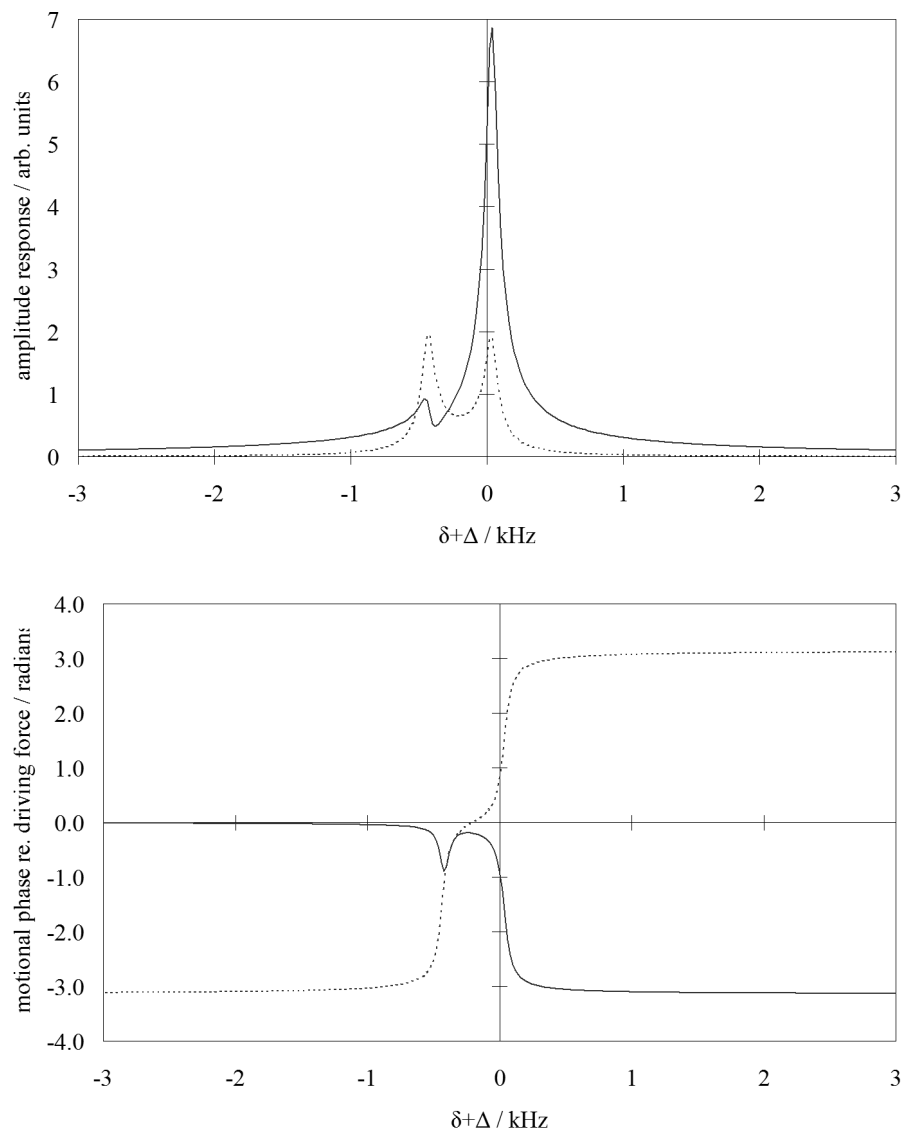


Figure 4.13: Plots showing the amplitude and phase response of a calcium ion to an external dipole drive near the magnetron frequency when the axialisation drive is weak. Parameters used are $B = 0.98\text{T}$, $\alpha = 15.00\text{kHz}^2$, $\beta = 0.08\text{kHz}$, $\Delta = -0.2\text{kHz}$ and $\epsilon = 40\text{kHz}^2$.

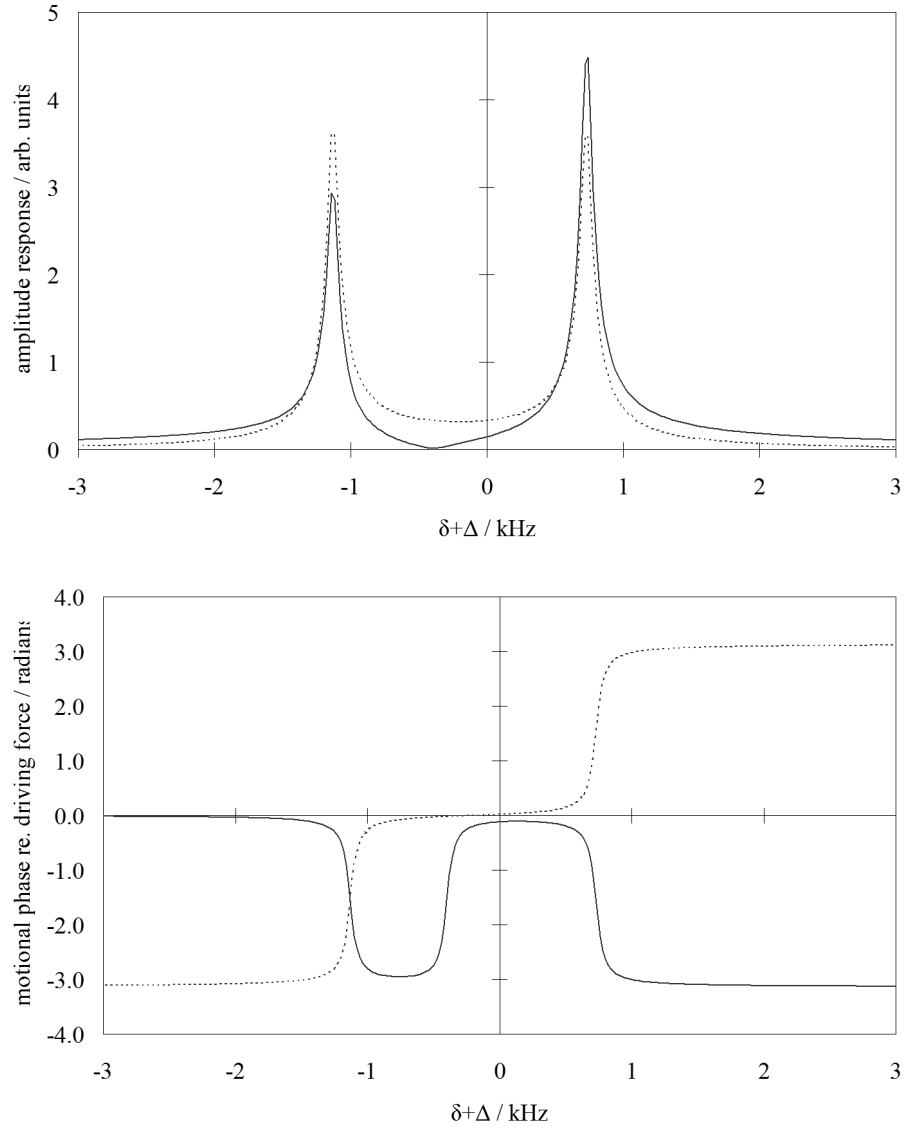


Figure 4.14: Plots showing the amplitude and phase response of a calcium ion to an external dipole drive near the magnetron frequency when the axialisation drive is strong. Parameters used are $B = 0.98\text{T}$, $\alpha = 15.00\text{kHz}^2$, $\beta = 0.08\text{kHz}$, $\Delta = -0.2\text{kHz}$ and $\epsilon = 300\text{kHz}^2$.




Figure 4.15: Argand plot of amplitudes A and B . Parameters are the same as those used to generate figure 4.14.

4.5 Ion Clouds and the Effects of Space Charge

So far in our discussion of laser cooling and axialisation we have considered only the case of a single ion in a Penning trap. In practice there will usually be a large number of ions in the trap at any one time and it is therefore important to consider the effect of this on the results of our model.

The obvious effect of having many ions in the trap is that the harmonic potential imposed by the trap electrodes is distorted by the other charges in the trap. When the mean inter-ion spacing is large, the forces due to the trapping potential will dominate and particle motion will be largely unaffected. If, however, we have a very large number of ions in the trap, or we apply cooling to bring them all closer to the centre of the trap, then the inter-ion spacing will tend to be reduced. Eventually we will reach a stage where the mutual Coulomb repulsion balances the force due to the trapping potentials and we will not be able to reduce the ion separation any further.

With Doppler cooling we are able in principle to cool a single calcium ion to the level of about 0.5mK. Since there is a well-defined relationship between temperature and the radius of an ion's orbit we can calculate the resulting magnetron radius at 1 tesla to be about $2\mu\text{m}$. If the Coulomb interaction with other ions in the trap limits us to a separation greater than this then we will be unable to cool our ions to the Doppler limit. In the axial direction the amplitude of a single ion's motion at 0.5mK is typically about $0.3\mu\text{m}$. The effect of a second ion in the trap is not so serious in terms of temperature, since the two ions can both have low oscillation amplitudes (and therefore temperatures) regardless of their separation.

In the axial direction we can easily calculate the minimum ion separation between two ions lying along the axis of the trap if we assume that the forces due to laser cooling are weak compared to the trapping force. We calculated the force due to the trapping potential at an axial position z in equation 4.1. We now set this to be equal in magnitude to the Coulomb force between two single charged positive ions separated by $2z$ so that

$$\begin{aligned} m\omega_z^2 z &= \frac{e^2}{4\pi\epsilon_0 (2z)^2} \\ z &= \sqrt[3]{\frac{e^2}{16\pi\epsilon_0 m\omega_z^2}}. \end{aligned}$$

For typical experimental parameters $z \sim 10\mu\text{m}$, so the ion separation is about $20\mu\text{m}$.

Calculating the minimum radius if there are two ions in the radial plane is somewhat less straightforward, since we have a potential hill in that plane. We can obtain a lower bound by transforming into a frame rotating at $\omega_c/2$, in which the effect of the magnetic field is to form a potential well confining the ions. From equation 4.6 we have the force due to this potential in such a frame. Equating it as before to the Coulomb force between ions that are $2r$ apart we find that

$$m\omega_1^2 r = \frac{e^2}{4\pi\epsilon_0 (2r)^2}$$

$$r = \sqrt[3]{\frac{e^2}{16\pi\epsilon_0 m\omega_1^2}}.$$

For the same typical experimental parameters this gives us a minimum value of r of about $10\mu\text{m}$. Even with optimal cooling the obtainable radial temperature is therefore limited by the Coulomb interaction to about 10mK. This ideal case corresponds to both ions sitting still in the potential well in the rotating frame. In practice the ions will not be cooled to this extent and it can be shown that the actual separation will always be larger than we have calculated. We estimate that the difference would be at least two or three times in our experiment. We cannot therefore expect to obtain radial temperatures of much less than 100mK with more than one ion in the radial plane of the trap.

Whether several ions will tend to line up along the axis of the trap or instead spread out in the radial plane depends on the relative gradient of the potential well in each direction. In general when the trapping voltage U_0 is large (for a given magnetic field) the ions will be forced into sharing the radial plane. When U_0 is small the ions will form an axial line. Since the ions can in principle possess low temperatures for the axial motion even when they are spread out along the axis of the trap it might seem that this arrangement is better for laser cooling, and hence we should use as low a trap voltage as possible. This is not necessarily the case, since as we have seen the separation of just two ions can easily be $20\mu\text{m}$ or more and we want ideally to address them both with our cooling lasers. Thus a balance, and hence a more intermediate trap voltage, is required.

In terms of quantum information processing with two (or more) ions in a single Penning trap we would envisage working with them lined up along the axis of the trap so that we can cool all of the motions to at least the Doppler limit. Reducing the ion separation in such an arrangement means increasing the trap voltage and doing so without forcing the ions into the radial plane means that the magnetic field must also be increased. A larger trap voltage also leads necessarily to an increased axial frequency, which is important for rapid processing. Given the limited achievable magnetic field strength we will never be able to obtain axial ion separations small enough that we can use large strings of cooled ions for information processing. It is instead imagined that a Penning trap quantum computer might consist of some arrangement of several traps, with only two ions occupying the processor trap at any time. Such a configuration is comparable to the proposals made by the radio-frequency ion trapping community for implementing scalable quantum information processing [19, 29].

Chapter 5

Experimental setup

The Ion Trapping Group at Imperial College has been working on the trapping and laser cooling of calcium for about five years. In this chapter I will describe the experiment that has been developed over that time. Since much of the experimental setup is dictated by the properties of the ion being studied I will begin by looking in more detail at the electronic structure of Ca^+ — this time in the presence of a magnetic field. I shall then provide an overview of the basic setup of the experiment before looking in detail at each of the key components. Finally I shall outline the general procedure used to trap and cool a cloud of ions.

5.1 Ca^+ electronic structure

In the presence of a large (~ 1 tesla) magnetic field the Zeeman sublevels of the Ca^+ system are shifted by an amount much greater than their natural widths (figure 5.1). The consequence for laser cooling is that a much larger number of levels must be addressed in order to form a closed cooling cycle and prevent optical pumping into some non-addressed state. The need to keep an increased number of lasers frequency stabilised is one of the key experimental difficulties with laser cooling in the Penning trap. Indeed it is for precisely this reason that ion species possessing accessible metastable D states, and therefore requiring repumper lasers, have never before been laser cooled in the Penning trap.

The $S_{1/2} \rightarrow P_{1/2}$ transition at 397nm is used for Doppler cooling and qubit state detection. Since the $P_{1/2}$ state can decay with equal probability




Figure 5.1: Relevant energy levels of calcium in the presence of a 0.98 tesla magnetic field, as calculated using data from the NIST Atomic Spectra Database [94].




Figure 5.2: *Top:* Zeeman splitting of the 397nm cooling transition in the presence of a 0.98 tesla magnetic field. *Bottom:* an equivalent diagram for the 866nm transition. The carrier (dashed lines) and sideband (solid black lines) frequencies used to address the σ components are indicated.

to either of the ground state Zeeman sublevels, two lasers must be used. We use the π ($\Delta m_j = 0$) components for largely historical reasons, though there is some small benefit to this in that they are shifted the least from the zero field frequency and are therefore less affected by magnetic field instability (see figure 5.2).

The $P_{1/2}$ states will also decay with a probability of 1/13 to the $D_{3/2}$ manifold [94]. These states have long lifetimes and so repumper lasers are required for efficient Doppler cooling. To pump out all of these sublevels we address the four σ ($\Delta m_j = \pm 1$) components of the $D_{3/2} \rightarrow P_{1/2}$ transition at 866nm. We can take advantage of the fact that the σ components are grouped in pairs close together in frequency by using just one laser with sidebands to address each pair. With appropriate selection of the laser's carrier frequency and sideband spacing we can position the sidebands such that each covers

one of the desired transitions (see figure 5.2). We thus reduce the number of 866nm repumper lasers required from four to two.

Amplified spontaneous emission (ASE) from the 397nm lasers can extend far enough in frequency to drive the 393nm $S_{1/2} \rightarrow P_{3/2}$ transition. Here the upper state has an extra decay channel to the $D_{5/2}$ manifold. Typically the rate of pumping into these levels is small, but if required, ions can be pumped back into the cooling cycle with a laser operating at 854nm to drive the $D_{5/2} \rightarrow P_{3/2}$ transition. We have found that typically a single laser is enough to pump out all of the Zeeman sublevels of the $D_{5/2}$ manifold and that the repumping efficiency is not very sensitive to the exact frequency of this laser.

Our sideband cooling and qubit transition will ultimately be the $S_{1/2} \rightarrow D_{5/2}$ quadrupole transition at 729nm. The ability to drive transitions with $\Delta m_j = \pm 2$ means that there are ten available components. It is worth noting that by driving the 850nm $D_{3/2} \rightarrow P_{3/2}$ transition we can intentionally put some ion population into the $D_{5/2}$ (qubit) state, albeit in an incoherent manner. This is useful for testing our ability to perform qubit state measurement and for observing quantum jumps, which can enable us to determine how many ions are present in the trap.

5.2 Experimental overview

The basic structure of the experiment is shown schematically in figure 5.3. Light from each of the 397nm and 866nm extended-cavity diode lasers is split into three parts; one to be used for laser locking, one for diagnostic purposes and the other to be sent to the trap. The first of these is used to lock each of the lasers to its own low finesse, temperature-stabilised optical cavity to improve frequency stability. For the 397nm lasers the diagnostics beams are sent to a Fabry-Perot etalon which can be used to ensure the lasers are running with just a single mode of output. The equivalent diagnostic for the 866nm lasers is a confocal cavity spectrum analyser, which enables us to check for single-mode operation and observe the amplitude of the sidebands relative to the carrier frequency. Faraday isolators are used to prevent light from the cavities and diagnostics returning to the lasers and causing optical feedback.

Each of the beams being sent to the trap can be individually redirected by flipping mirror mounts to a wavemeter, which enables the initial setting

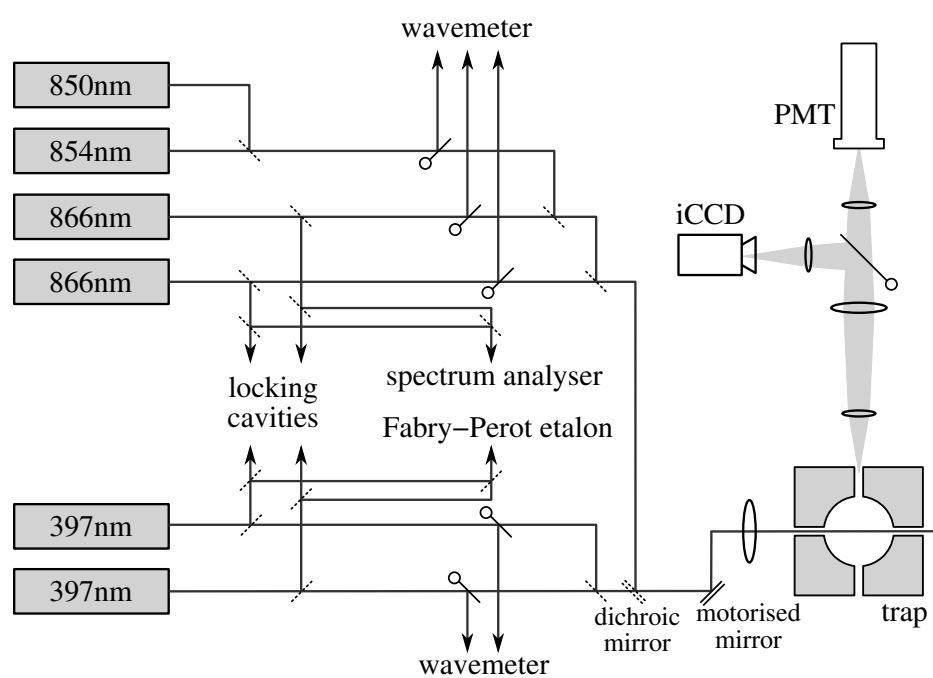


Figure 5.3: Schematic of the key components of the experiment.

of laser wavelengths. After this point the 397nm beams are brought together on a 50:50 beamsplitter before being passed through a half-wave plate and a polarising beamsplitter. This serves the dual purpose of acting as a variable filter, for reducing the beam power entering the trap, and enabling us to redirect the beams to other ion traps when required. A similar setup exists for the lasers at 866nm, only with some light from the 854nm and 850nm lasers being mixed in using a glass plate as a beamsplitter. The arrangement of the polarising beamsplitters is such that the plane of polarisation of the 397nm light entering the trap is parallel to the magnetic field vector, whereas the polarisation of each of the repumpers is perpendicular to it. The infra-red and ultra-violet beams are merged at a dichroic mirror before being focused into the trap. The penultimate mirror before the trap is in a computer-controlled motorised mount. This enables us to make precise and repeatable adjustments to the positioning of the beams in the trap.

Fluorescence from the ions is collected and nearly collimated by a lens placed close to the trap in a direction perpendicular to the path of the laser beams. This fluorescence is focused down onto a small circular aperture in front of a photon-counting photo-multiplier tube (PMT). This is connected to a multi-channel scalar which is used to count the number of photons detected in a selected time interval. Alternatively the fluorescence can be redirected and focused onto the photocathode of an intensified CCD camera for spatial imaging of the ions.

5.3 The split-ring trap

5.3.1 Trap structure

The basic design of the trap used in these experiments is shown in figure 5.4. There are two key differences between this and the conventional Penning trap. Firstly the electrode surfaces are not the ‘ideal’ hyperbolic surfaces, but are instead triangular in cross-section. This makes manufacture of the trap significantly easier. Near the centre of the trap the shape of the potential is still a good approximation to a true quadrupole. The second difference is that the ring electrode is broken into four segments. This allows us to apply additional radial quadrupolar fields as required for axialisation. The radius of the ring, r_0 , is 5.0mm and the distance from trap centre to the endcaps, z_0 is 3.5mm.




Figure 5.4: Scale drawings of the split-ring trap. For clarity all fastenings and one of the laser shields have been omitted. *Left*: vertical section through the line joining the endcaps. *Right*: vertical section through the plane of the ring electrodes.

An ideal Penning trap with these dimensions would have an R_0^2 parameter of 50mm^2 . From measurements of motional frequencies we find an actual R_0^2 of $55 \pm 1\text{mm}^2$.

The trap was originally built for use with magnesium ions. It was designed by M. van Eijkelenborg and technical drawings can be found in reference [95]. After some unsuccessful attempts at rebuilding it for use with calcium, the endcap mounts were altered to allow better alignment of the endcaps. Shields were also added to reduce the amount of laser scatter that can be imaged by the fluorescence detection system.

The electrodes and their holders are made from beryllium-copper. The ring electrode segments are held in place by a square frame, using ceramic screws to keep them electrically isolated. The endcaps are held in their mounts by metal grub screws. Electrical connections to all the electrodes are made by screwing Kapton-coated wires onto them. To allow laser access to the centre of the trap, and to enable fluorescence from the ions to be collected, holes are drilled in all four sides of the ring holder. A lens holder is attached to the top of the trap so that the fluorescence can be imaged onto the detection system.

5.3.2 Ion generation

Ions are created by electron bombardment of calcium atoms. Detailed descriptions of the same arrangement have been given in references [96] and [97], so I shall describe the setup only briefly. The atoms are emitted from one of two ovens mounted on the underside of the lens holder. The ovens are made from 1mm diameter tantalum tubes (wall thickness $80\mu\text{m}$) to which $250\mu\text{m}$ diameter tantalum wire is spot-welded at several points along its length. Each oven is stuffed with small lumps of calcium and a hole is made on one side. Driving current through the tantalum wire heats the oven body and causes sublimation from the surface of the calcium. The calcium is emitted from the hole in the oven, which is pointed towards the centre of the trap.

Filaments for the production of electrons are located in small recesses behind the endcaps. Small holes in the endcaps allow the electrons to pass through and into the centre of the trap. This arrangement is necessary because the electrons are forced to spiral along the magnetic field lines, so the filaments must be located along the axis of the trap.

5.3.3 Vacuum system

To prevent heating of the ions and decoherence of their internal states via interactions with background gas molecules it is important to operate the trap at very low pressures so that the rate of collisions is small. The trap is housed in an ultra-high vacuum (UHV) apparatus as shown in figure 5.5. Quartz and conventional glass windows are used to allow laser access and imaging of the fluorescence from the ions. An ion pump is used to maintain the vacuum and an ion gauge enables us to measure the pressure in the system.

In order to achieve a low background pressure, and to ensure there are no contaminants on the electrode surfaces, it is important that the trap and all the associated parts are thoroughly cleaned before assembly. First the copper parts of the trap are cleaned using a solution of 1 part Copperbrite to 4 parts distilled water. Copperbrite is a mixture of 75% concentrated nitric acid, 23% concentrated sulphuric acid and 2% concentrated hydrochloric acid. These parts of the trap are then rinsed thoroughly with distilled water. They are then cleaned, along with the all the non-copper parts, in acetone and methanol. The cleaning process was the same as has been described before in references [96, 97].




Figure 5.5: Diagram of the ultra-high vacuum apparatus that houses the trap.

The initial pumping down of the system was performed using a Pfeiffer–Balzers TPU-062 turbomolecular pump backed by a Pfeiffer–Balzers DUO1.5A rotary pump. To reduce the outgassing rate from contaminants, the system was baked at a temperature of $\sim 215^\circ$ for about one week. After this the system was returned to room temperature and the ion pump was switched on. A few days later the system pressure settled at approximately 10^{-9} mbar. At this level it was the turbomolecular pump that limited the pressure and the vacuum system was valved off from the pumping station. With only the ion pump engaged the pressure continued to gradually fall towards its current measured value of 2×10^{-10} mbar.

It is expected that the error on our pressure measurements may be as much as a factor of two or three, due to the positioning of the ionisation gauge in a rather narrow section of vacuum tubing. This tubing becomes rather hot when the gauge is in use and it is expected that the local pressure will be somewhat increased as a result. Furthermore this heating effect prevents us from operating the gauge on its largest filament current setting — which gives improved sensitivity at low pressures. We find that the ion gauge is also affected by the magnetic field used for operating the Penning trap. This means that when the field is switched on the indicated pressure drops by a factor of about 10. It is believed that this is a constant factor and hence can be

compensated for when measuring the system pressure during trap operation.

5.3.4 Magnetic field generation

The vacuum housing is designed such that the pole pieces of an electromagnet can be brought as close as possible to the trap itself. The magnet used is an Oxford Instruments N100 driven by a KSM SCT-220-20 high-power current supply. This electromagnet can be used to generate a maximum magnetic field at the centre of the trap of just over 1 tesla. In order not to drive the current supply too hard we generally operate it at a current of 15.0 ± 0.1 ampere, which gives us a magnetic field at trap centre of 0.98 ± 0.01 tesla. A more precise measurement of the magnetic field during a given experiment is made by measuring the motional frequencies of ions in the trap.

The amount of power dissipated in the magnet coils can be as large as 2 kilowatt, so water cooling is used to prevent the coils overheating. Because of the large thermal mass of the iron pole pieces, however, it can take several hours for an equilibrium temperature to be reached — leading to a slow variation in the impedance of the magnet coils. In principle the power supply should stabilise the current so that the magnetic field remains constant. In practice, however, an intermittent fault meant that the output current was sometimes unstable and on occasion it was necessary to use a voltage stabilisation setting. In these instances the output current would vary in inverse proportion to the impedance of the magnet coils.

We can gauge the current passed through the magnet coils by measuring the voltage across a $1\text{m}\Omega$ shunt resistor placed in series with the coils. This gives us a precision of ± 0.1 ampere. Although we can only monitor the current stability to about 1 part in 150, repeated measurements of motional frequencies of ions in the trap suggest that with the power supply operating stably in current-limiting mode the magnetic field stability is in the region of 10^{-3} – 10^{-4} .

5.3.5 Operation as Penning trap

In order to operate the trap as a Penning trap, a bias must be applied to the endcaps relative to the ring electrodes. Generally we use a trap bias in the region of 4V. We choose to ground the ring segments and apply a positive potential to the two endcaps, since application of the quadrupolar axialisation

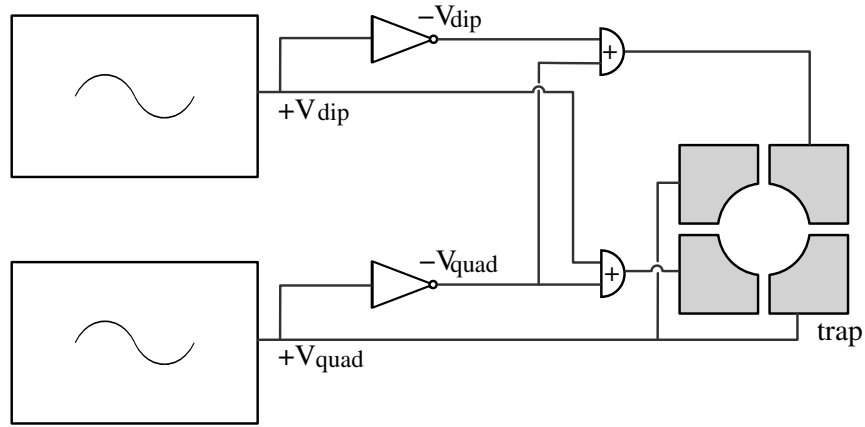


Figure 5.6: Diagram showing how the quadrupolar and dipolar drives are simultaneously applied to the ring segments of the trap.

drive to the ring is more straightforward when it is biased about ground. As we will see in more detail in the next chapter, we often want to apply a dipole excitation in order to drive one of the radial motions. The quadrupolar and dipolar drives are generated using a simple circuit consisting of summing and inverting amplifiers. The signals from two function generators, V_{quad} and V_{dip} , (taken from an HP 3325B, HP 8643A or TTi TG550) are each split into two and in each case one part is inverted. As shown in figure 5.6, one pair of opposing ring segments has $+V_{quad}$ applied to it. The other two segments have $-V_{quad} \pm V_{dip}$ respectively.

5.3.6 Operation as RF trap

In order to test certain systems, and so that we can compare results in the Penning trap with those obtained in the RF trap, it is often helpful to operate our trap as an RF trap. This is easily performed by grounding all four ring segments and applying a biased RF trapping potential to the endcaps. The setup for generating this potential is shown in figure 5.7. The output from a function generator (HP 3325B) is passed through an ENI 325LA RF power amplifier and then through a transformer made in-house, in order to generate the large potentials required. A simple choke circuit, also made in-house, is used to mix this RF with a small dc potential, so that we can apply a bias to the endcaps. A Daiwa CN-101L standing wave ratio meter just after the RF

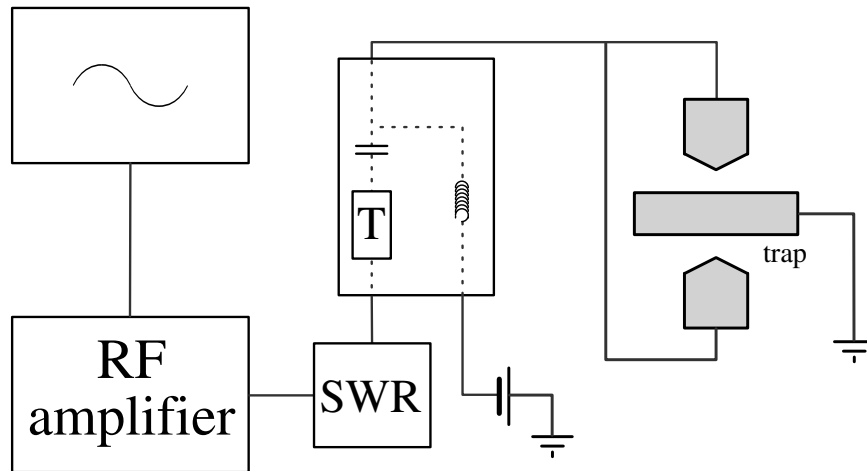


Figure 5.7: Diagram showing the electronic setup for generating the RF trapping potential. The object marked ‘T’ is a transformer.

amplifier enables us to measure the efficiency with which the RF is coupled into the transformer, and hence allows us to find the resonant frequency.

We generally use a drive frequency of about 2.0MHz. The system is highly resonant, so selecting a different frequency involves changing the inductance or capacitance in the system. This can most easily be done by changing cable lengths or altering the number of coils in the transformer. Typically the output from the function generator would be set to $\sim 400\text{mV}$ (peak-to-peak). After the two gain stages the potential applied to the endcaps is then in the region of 200V p-p. We have found that slightly larger RF amplitudes, say 250–300V p-p, are somewhat better for loading of the trap — presumably because of the increased trap depth. Subsequently reducing the drive amplitude, however, reduces RF heating and gives us better signal level. Although still larger amplitudes can be used, compression in the transformer becomes significant and we are limited to about 500V p-p. We generally apply a dc bias to the endcaps of about -4V, though this does not noticeably affect the stability of the trap.

When working with single ions in an RF trap (or strings of ions in a linear RF trap) it is possible to eliminate micromotion of the ions by applying potentials to compensation electrodes. The aim is to counteract any stray dc potentials and position the ions exactly at the node of the RF potential. In

principle we can achieve this by applying small potentials to the individual ring segments and ovens. Although we have worked with single ions to test our system, we have never required micromotion compensation and it has not been performed.

5.4 Doppler cooling lasers

The 397nm Doppler cooling lasers are commercially available TuiOptics DL100 extended-cavity diode lasers (ECDLs). One of these contains a Nichia NLHV500A diode with a nominal 5mW free-running power and the other a Nichia NLHV3000E diode producing 30mW of output power when free-running. Both diodes have a free-running wavelength at room temperature within 1nm of our desired transitions.

The lasers are based on the Littrow extended-cavity design (see figure 5.8), whereby the output from the diode is collimated and passed onto a blazed reflective grating. The grating is angled such that the first-order diffracted light for the desired wavelength passes back into the laser diode, effectively forming a cavity of lower free spectral range and higher finesse than that of the bare diode and so reducing the linewidth of the laser. The zeroth-order diffracted light leaves the cavity and forms the output of the laser. Coarse tuning is possible by varying the angle of the grating such that different longitudinal modes (with different frequencies) are sent back into the diode. Each grating is mounted on a piezo-electric transducer (PZT) so that the length of the diode–grating cavity can be precisely controlled. This allows continuous fine-tuning of the laser’s output.

Since the laser frequency varies with temperature the various components of the laser head are mounted on a metal block, which is in turn attached to a Peltier device. This enables both heating and cooling of the metal block and in conjunction with a temperature sensor can be used to stabilise the temperature. According to the manufacturer’s specifications the temperature stability is at the level of 5mK and the operating range is between -50 and 50°C .

The drive electronics for these lasers consist of integrated temperature, current and PZT modules, along with a proportional–integral–differential (PID) controller used for locking the lasers to their optical cavities (TuiOptics DTC100, DCC100, SC100 and PID100).




Figure 5.8: The Littrow configuration extended cavity diode laser.

With a typical diode current of 50mA the output from the lower and higher power diodes in their extended cavities will generally be about 3mW and 9mW respectively. The more powerful beam is filtered such that the beam powers at the trap are the same, about 0.5mW. The theoretical spot size of the laser in the centre of the trap given the focusing optics is $\sim 75\mu\text{m}$ by $30\mu\text{m}$, but since the beam profile is not a perfect Gaussian we might expect the actual spot size to be a little larger. Measurements of the waists of the two blue beams at focus have been made by K. Koo, yielding about $70\mu\text{m}$ for each. Assuming a beam waist w_0 in each direction of $75\mu\text{m}$ by $30\mu\text{m}$ we can calculate the peak intensity I_0 of the beam to be about $140\text{mW}/\text{mm}^2$. If we calculate the saturation intensity for this transition using the equation

$$I_{sat} = \frac{\hbar\omega^3\gamma}{\pi c^2}, \quad (5.1)$$

we obtain $I_{sat} = 1.0\text{mW}/\text{mm}^2$. Thus we see that at the centre of the beam the available laser power is more than a hundred times that required to saturate the cooling transition. For large clouds this excess laser power effectively increases the area of the beam that possesses enough power to cool the ions and so leads to more fluorescence. We can calculate that the region of the beam that is saturating the transition has major and minor axes of $240\mu\text{m}$ and $90\mu\text{m}$ respectively. However oversaturation also causes power broadening of the transition and therefore increases the minimum temperature which can

be reached through Doppler cooling. For this reason it is necessary to heavily attenuate the 397nm beams when small clouds of ions are being cooled.

Measurements of the linewidths of the 397nm lasers were made by K. Koo by tuning them to very similar wavelengths and superimposing them on a fast photodiode. Taking the Fourier transform of the observed beat signal yields a spike at the difference frequency of the two lasers, the width of which corresponds to the sum of the two linewidths. The measured width of the beat frequency was ~ 3 MHz, which suggests a linewidth per laser of about 1.5MHz. This is well below the 24MHz natural linewidth of the transition and indicates that the laser linewidth does not limit the final temperature that can be achieved through laser cooling.

According to specifications the temperature stability of these laser systems is ± 5 mK. We can estimate the resulting frequency drift by measuring the amount by which each laser needs to be retuned in order to restore maximum signal from a cloud of ions after a given period of time. This is best performed in an RF trap in order to reduce the effect of magnetic field instability. Using this technique the frequency drift of these lasers was measured by K. Koo to be about 300–400 MHz/hr. Since the Doppler cooling transition has a natural width of 24MHz, this would suggest that the lasers would need retuning every few minutes in order to prevent them drifting above or too far below resonance. Given the number of lasers involved in our cooling scheme it is obvious that better frequency stability is required. It is for this reason that the 397nm cooling lasers and 866nm repumpers are locked to external reference cavities. We will look at this technique in section 5.7.

5.5 Repumper lasers

The repumper lasers are all extended-cavity diode lasers built in-house. Although there are some slight differences in the arrangement of the components, the lasers follow the same basic Littrow configuration as the 397nm lasers. A detailed description of the mechanical design can be found in reference [98]. The diodes themselves are SDL 5411 (for the 854nm, 850nm and one of the 866nm lasers) and Laser 2000 LD 1377 (for the remaining 866nm system). The free-running wavelengths of these diodes are quoted as 852 ± 4 nm and 860 ± 4 nm respectively at room temperature. The gain curves for both diode types extend to $\sim \pm 15$ nm from the free-running wavelength. The use of an

extended-cavity arrangement improves our ability to tune the laser output far from the maximum of the gain curve, but we also heat the diodes covering the 866nm transitions to about 40–50°C so that the free-running wavelength is within a few nanometres of the desired transitions. The nominal free-running maximum output powers are 100mW for the SDL 5411 and 120mW for the LD 1377, though we typically run them at currents well below the rated maximum.

The current, temperature and PZT drivers are all separate commercially-available systems from a variety of manufacturers. Typically we operate all of these lasers at currents of 60–80mA. We measure output powers of about 8mW and 4mW for the SDL 5411 and LD 1377 866nm lasers respectively. The 850 and 854nm systems are generally operated at powers well in excess of 10mW.

The repumping transitions have scattering rates and photon momenta much smaller than those of the 397nm transition. As a result all of these lasers can be tuned into resonance with their transitions without causing significant Doppler heating of the ions and it follows that the laser linewidths are not critical. Given the extended-cavity arrangement we would estimate the linewidths to be in the region of 1MHz or less.

The temperature control system for these lasers is essentially the same as those for the commercial 397nm systems. Although a direct measurement has not been carried out, we anticipate that the temperature stability ought to be similar to the ± 5 mK of the 397nm lasers. In order to maintain saturation of the repumping transitions it is important that the 866nm lasers have good frequency stability. To improve upon the inherent stability of the lasers they are locked to optical reference cavities as described in section 5.7. Ultimately it is the thermal stability of these reference cavities that is the key parameter in terms of laser frequency stability.

As discussed in section 5.1 the 866nm lasers are operated with sidebands in order to address all four σ components of the $P_{1/2} \rightarrow D_{3/2}$ transition with just two lasers. The sidebands are generated by direct amplitude modulation of the diode current at a frequency corresponding to half the splitting between adjacent σ components (see figure 5.2). The RF source for both lasers is an Agilent 83650L function generator. An RF splitter (Mini-Circuits ZFRSC-42) and attenuators are used to produce an output level suitable for each laser. Typically the RF power for the SDL 5411 diode is about 7dBm and for the




Figure 5.9: The output of a spectrum analyser showing the carrier and sidebands on one of the 866nm lasers. The slight asymmetry of the peaks occurs because the analyser cavity is not perfectly confocal.

LD 1377 diode is about 10dBm. It would seem likely that this difference is due to the efficiency with which power is transferred to the diodes. In order to achieve good modulation depth of the current it is imperative that there be good impedance matching between RF source and diode. This task is performed by placing a stub tuner in each transmission line. After this a Bias-Tee (Mini-Circuits ZFBT-6G-FT) is used to add the direct current to the RF. The output from this is passed to the laser diode.

In general we observe sidebands with amplitudes about 10% of that of the carrier (see figure 5.9). By increasing the RF power it is possible to increase the sideband amplitude further, but this comes at the cost of stability. Since no particular effort has been made to match the length of the extended cavity of the lasers to the sideband frequency, transferring power from the carrier to the sidebands effectively reduces the amount of feedback to the laser. This makes the laser more likely to jump into a different frequency mode, or to begin operating in a broadband manner.

Due to chromatic aberration in the final focusing lens, the infra-red beams are left slightly defocused at trap centre. This is useful since it reduces the dependence of the cooling efficiency on beam positioning. Just before entering

the trap we measure about 0.5mW of power in each 866nm beam. If we assume 5% of this is in each sideband and that the beam waists in each direction at trap centre are $240\mu\text{m}$ and $100\mu\text{m}$, we find a peak intensity I_0 at each correct wavelength of $0.70\text{mW}/\text{mm}^2$. This is well in excess of the saturation intensity of $0.09\text{mW}/\text{mm}^2$ calculated from equation 5.1. More importantly we can calculate the region of the beam which is saturating the 866nm transitions to have major and minor axes of $480\mu\text{m}$ and $200\mu\text{m}$ respectively.

Since the 854nm laser covers several transitions it is hard to estimate the fraction of its power that actually addresses each one. We have however looked at the effect of introducing a second laser at 854nm which was tuned to each of the transitions in turn. In each case the same fluorescence signal was observed from a cloud of ions with either or both of the lasers present. We can thus infer that we are saturating all of the required 854nm transitions. It is worth noting that we generally observe an increase in the signal from a large ion cloud of about 20% when the 854nm laser is introduced into the trap. Since it is ASE from the cooling lasers that necessitates the 854nm light, we usually find that once we attenuate the cooling beams for small clouds we no longer see any significant effect from the 854nm repumper.

5.6 Laser diagnostics

In order to trap a cloud of ions we must first set all our lasers to their correct wavelengths. Although the initial temperature of ions produced in the trap can be very large, well over 1000K, in the presence of laser cooling they will rapidly reach an equilibrium temperature dependent on the laser parameters. Typically we find that even with a large ion cloud we will only observe enough fluorescence to begin optimising the laser frequencies if our 397nm cooling lasers are initially set to within about 800MHz of their optimum frequencies. This corresponds to a temperature of about 200K and implies that the corresponding Doppler width for the 866nm transition will be about 400MHz. These Doppler widths represent the minimum precision with which we must be able to set our initial laser frequencies.

We set our wavelengths using a Michelson interferometer wavemeter built in-house. The basic operating principle of such a wavemeter is that light from a laser is split into two beams, each of which travels down a separate arm of the wavemeter. The two beams are brought together again at a beamsplitter and

the resulting interference signal is directed onto a fast photodiode. Varying the relative length of the two arms leads to the observation of interference fringes as the two parts of the beam shift in and out of phase with each other. If we vary the path length by a distance d the number of fringes, N , which pass the photodiode is given by

$$N = \frac{d}{\lambda}, \quad (5.2)$$

where λ is the wavelength of the light. By sending two lasers simultaneously into the two arms of the wavemeter we find that the ratio of the number of fringes observed as the path difference is changed is simply given by the ratio of the wavelengths of the two lasers. If one of these is a stable reference laser of known wavelength we can determine the wavelength of the other laser as

$$\lambda_{meas} = \frac{N_{ref}}{N_{meas}} \lambda_{ref}. \quad (5.3)$$

The design of the wavemeter follows closely that of one built by the Ion Trapping Group at the University of Innsbruck [99]. The optical and mechanical elements (see figure 5.10) were built by E. Phillips and a detailed description can be found in reference [96]. The electronics for fringe counting was supplied by the Innsbruck group.

The variation in path length is performed by a cart travelling along an air-track. Cornercube reflectors attached to each end of the cart form part of each arm of the wavemeter such that moving the cart increases the length of one arm whilst reducing the length of the other. This layout doubles the number of fringes that can be counted by moving the cart along a given length of track. It is important that vibrations of the cart are small compared to the wavelengths of the lasers being used, since otherwise they will cause extra fringes and lead to erroneous measurements. In order to ensure that the change in path length of the two arms is the same for both lasers, the wavemeter is designed such that the lasers counter-propagate along the same paths. This has the added advantage that the laser to be measured must simply be overlapped with the reference laser beam leaving the wavemeter.

The reference laser is a Melles-Griot stabilised He:Ne (050-STP-903) with a frequency stability quoted as being better than 1MHz over the course of an hour. This corresponds to an uncertainty due to the reference laser of




Figure 5.10: The optical layout of the wavemeter. The dashed lines represent light from the laser being measured, which follows the path of the reference He:Ne in reverse.

2 parts in 10^8 . The control electronics uses zero-crossing detectors for fringe counting — one is added to the current fringe count every time the AC-coupled photodiode signal passes through zero with a particular gradient. The process is triggered when the cart breaks an IR beam near the start of its travel. The counts are not actually started until there is a near simultaneous zero-crossing in the fringes from the reference laser and the laser being measured. Once a set number of reference fringes has been observed the counts continue until the next near-simultaneous zero-crossing of the two fringe patterns. The ratio is then determined and the unknown wavelength calculated.

Assuming the wavelength of the reference laser is sufficiently well known and that the alignment of the laser beams is good enough that the variation in path difference is the same for each, the resolution of the wavemeter will depend on how close together the zero-crossings of the two channels must be in order to start and stop the counting process. If, for example, the number of fringes counted is 10^6 and the zero crossings must be within 1% of the period of the reference fringes then the resolution would be 0.01 parts in 10^6 .

We have measured the actual resolution of the wavemeter by looking at the statistical deviations of several series of measurements of unknown laser wavelengths. In each case the laser being measured was locked to an external reference cavity and 40 measurements were taken in quick succession so that the laser frequency drift would not be significant. Typical results are shown in figure 5.11. We find standard deviations of about 3×10^{-5} nm at 397 nm (corresponding to an uncertainty of ~ 60 MHz) and 5×10^{-5} nm at 866 nm (corresponding to ~ 20 MHz). At both wavelengths the spread in measurements extends to no further than 1×10^{-4} nm from the mean value, which gives us an idea of the single measurement resolution. Clearly these uncertainties are within the requirements we laid down earlier in this chapter of 800 MHz at 397 nm and 400 MHz at 866 nm.

It is worth noting that since the optical paths taken by the reference laser and the laser being measured are nominally the same any slight changes in optical alignment within the wavemeter do not significantly affect its long-term measurement stability. We expect the main sources of instability to be the reference laser frequency and relative changes in the refractive index of air at the different wavelengths due to temperature. At room temperature the measurement shift due to this effect is just 2×10^{-5} nmK $^{-1}$ at 397 nm and 6×10^{-6} nmK $^{-1}$ at 866 nm.




Figure 5.11: Histograms showing the spread in a series of successive measurements by the wavemeter of (*left*) a 397nm and (*right*) an 866nm laser. The units are 10^{-5} nm, corresponding to about 20MHz at 397nm and 4MHz at 866nm.

Note that the readout from the wavemeter gives the approximate wavelength of a measured laser in air at room temperature. Since the counting electronics were originally designed for a slightly different reference laser we must compensate in order to gain an accurate measurement. Furthermore we are interested in the vacuum wavelengths of the lasers that we measure, since this is what the ions are exposed to. Measurements we have made of the precise wavelengths of the 397nm and 866nm transitions using the (calibrated) wavemeter are consistent with accepted values at the level of about 2×10^{-4} nm. This gives us an indication of the wavemeter's absolute accuracy.

In order to ensure that the 397nm lasers are operating on a single mode of their extended cavities, both beams are focused into a Fabry-Perot etalon (Queensgate Instruments EV85-831). When the lasers are running cleanly a single series of concentric rings is observed. If another frequency is present in the laser's output then intermediate rings become visible. To perform the equivalent function for the 866nm lasers, a spectrum analyser (Tec-Optics SA7.5) is used. This serves the additional function of allowing us to monitor the amplitude of the sidebands on these lasers.

5.7 Laser locking

5.7.1 Optical reference cavities

As we saw earlier, the rate of thermal drift of the cooling and repumper lasers is large compared to the linewidths of the transitions we are addressing. This means that we would need to retune the lasers every few minutes in order to maintain efficient Doppler cooling. In order to alleviate this problem, we lock the 397nm and 866nm lasers to temperature-stabilised optical cavities.

To illustrate the basic principle, consider a cavity comprised of a pair of identical curved mirrors with reflectance R separated by a distance L . If we shine a laser of wavelength λ through one of the mirrors a fraction $\sim (1 - R)$ of the light intensity will enter the cavity and be reflected back and forth by the mirrors. Only if the length of the cavity is close to an integral multiple of $\lambda/2$ will waves on successive passes of the cavity interfere constructively to form a large amplitude resultant wave. Since the same fraction $(1 - R)$ is transmitted through each mirror the output from the cavity will only be large if our constructive interference criterion is satisfied. Thus the cavity forms a frequency selective filter, transmitting a comb of frequencies. The separation of the frequencies that are transmitted (the free spectral range) depends on the cavity length as

$$\text{FSR} = \frac{c}{2L}. \quad (5.4)$$

The greater the reflectance, R , of the mirrors, the longer the storage time in the cavity will be and therefore the closer L must be to $\lambda/2$ in order to maintain constructive interference. Thus the transmission peaks are narrower if R is large. We define the finesse, \mathcal{F} , of the cavity to be the ratio of the free spectral range to the half-width of the transmission peaks. It can be calculated as

$$\mathcal{F} = \frac{\pi\sqrt{R}}{1 - R}. \quad (5.5)$$

Providing the length of such a cavity is kept stable the transmission peaks can be used as reference frequencies for locking a laser to. The relatively simple lock that we perform is a ‘side-of-fringe’ lock. Here we tune the laser such that its wavelength gives us a transmission through the cavity of half the maximum. We then monitor the transmission on a photodiode and electronically alter the laser frequency such that the transmission remains constant. Providing

we have some means of changing the cavity length we can still tune the laser locked to it. The key benefit of using a reference cavity is that since we do not need regular access to the cavity we can more easily isolate it from mechanical and thermal contact with the environment. The length is therefore more stable than that of a the extended cavity of a diode laser.

One complication not considered in the discussion thus far is that an optical cavity can support a whole series of transverse electric and magnetic (TEM) modes. Due to their differing divergences each experiences a slightly different phase shift over each pass of the cavity, which effectively changes the cavity length that each mode ‘sees’. The result is that the frequency combs due to different transverse modes do not in general overlap. In order to obtain a single set of well-defined peaks for use as a frequency reference we need to make sure that we only couple light into one of the transverse modes. This is achieved by tailoring the beam entering the cavity such that its size, position and divergence all match those of one of the modes. In principal the TEM₀₀ mode is optimal for this task since it has the smallest divergence and waist and so we can avoid coupling to other modes more easily.

The reference cavities we use were designed by E. Phillips and a detailed description can be found in reference [96]. The basic design is shown in figure 5.12. Each cavity body is made from several pieces all cut from a single block of Zerodur — a glass-like material designed to have a very low coefficient of thermal expansion. The mirrors (CVI Technical Optics PR1-xxx-95-0537-0.30CC, where xxx is the wavelength 400 or 866) were made to give a finesse of about 60. They are attached to the Zerodur by a thin layer of TorrSeal epoxy resin. In order that we can vary the length of the cavity, one of the mirror pieces is attached to the main body by six PZTs. Since these have a large thermal expansion coefficient they are mounted in sets of three in a re-entrant fashion — they counteract each other such that if they all expand by a given amount there is no net movement of the mirror. The cavity length is changed by altering the length of either one of the sets of PZTs. The distance between the mirrors is approximately 100mm, which leads to a free spectral range of 1.5GHz.

The cavities are situated in a pair of aluminium boxes lined with neoprene for thermal and mechanical insulation. In order that we can stabilise the temperature of the boxes, they are each wrapped in heater wire and thermistors are held in contact with the metal surfaces. The boxes are then wrapped in




Figure 5.13: Schematic of the optical setup for locking each laser.

The 397nm lasers are locked using the PID controller integrated into the laser driver. The photodiode signal is sent into the controller and a setpoint is chosen to lie about half way up the transmission peak of the TEM_{00} mode. Appropriate proportional, integral and differential gains are determined experimentally. The output is passed internally to the PZT driver module of the controller.

The locking electronics for the 866nm lasers were built in-house. A block diagram is shown in figure 5.14. The design follows largely that in reference [101], but has been modified to include proportional, integral and differential stages. A pair of voltage references and a potential divider are used to generate a variable setpoint. The input signal can be inverted if required to allow locking to either a positive or negative slope. It is then added to the setpoint to generate the error signal. A capacitor across the summing amplifier that performs this task acts as a low-pass filter and caps the frequency response of the circuit. The error signal is passed to separate proportional, integral and differential stages. The output from these is added by another summing amplifier and a variable gain is applied. The output from the circuit is used to control the PZT driver for the lasers.

It is important to note that the PZTs have their fundamental resonance frequency at about 7kHz. As with any driven oscillator there will be a phase

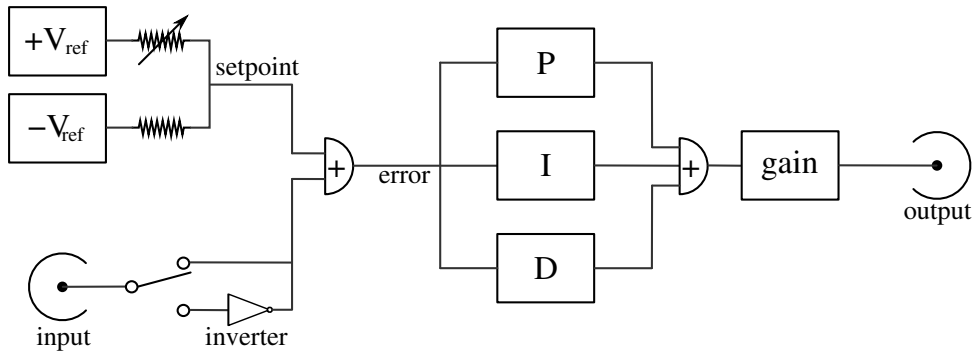


Figure 5.14: Block diagram of the locking electronics for the 866nm lasers.

shift of π in their response if the frequency of the driving force passes through this resonance. In our case this change in phase would mean that we were acting to *increase* high-frequency components of the error signal rather than reduce them. The low-pass filter in the circuit therefore has its cutoff frequency at $\sim 2\text{kHz}$. The cutoff frequencies for the integral and differential stages are initially set to 10% and 90% of this value respectively, following a general rule-of-thumb suggested in electronics texts. Because we are limited to such a low bandwidth, these locks do not perform any linewidth narrowing of the lasers and are intended mainly to cut out thermal fluctuations. As a result the integral term of the locking circuit is dominant.

5.7.3 Locking procedure

When a laser is locked for the first time an appropriate setpoint must be chosen. This is best performed by scanning either the laser or cavity over a frequency range that encompasses the desired transmission peak. If we then monitor the error signal we will see a (possibly inverted) copy of the peak that has been shifted by some amount. By adjusting the setpoint potentiometer we can shift the peak so that its half-maximum is at ground.

Locking of a laser is performed by tuning the cavity until the transmitted signal is close to its maximum. Switching the lock on at this point should result in the transmitted signal rapidly moving to the setpoint. If this is not the case it is often useful to turn the overall gain down to some small value, resulting in a somewhat loose lock. The gain can then be gradually increased until small oscillations are observed. The gain should be reduced to just below

this level.

5.7.4 Cavity stability

The length stability of the cavities directly governs the frequency stability of the lasers locked to them. We have measured the long-term stability by using the wavemeter to make wavelength measurements over the course of several hours. Each measurement consisted of an average of at least ten readings. In order to minimise the effect of ‘creep’ of the PZTs the potential across them was kept constant for at least 24 hours before each series of measurements. Results for each of the four cavities, whose frequency drifts were measured on up to three different days, can be seen in figure 5.15. The error bars represent the standard deviations of each set of readings.

We can see that on any given day the drift of each laser tends to be relatively linear, but that the direction and magnitude of drift varies from day to day. If we take a weighted average of the overall drift rates for each day, ignoring the direction of drift, we find an average drift of $64\pm 13\text{MHz/hr}$ and $97\pm 13\text{MHz/hr}$ for the two 397nm cavities and $29\pm 7\text{MHz/hr}$ and $56\pm 9\text{MHz/hr}$ at 866nm.

We would have expected the cavities for like wavelengths to drift at similar rates, since they are mounted in the same metal boxes. It is also to be expected that the UV cavities should drift the most, since the same change in length leads to a larger change in frequency. A drift in frequency of one free spectral range (in this case 1.5GHz) corresponds to a change in the length of the cavity of half the wavelength of the light being used. We can therefore calculate similar length stabilities of the cavities of $8\pm 2\text{nm/hr}$, $13\pm 2\text{nm/hr}$, $8\pm 2\text{nm/hr}$ and $16\pm 3\text{nm/hr}$ for the two UV and infra-red cavities respectively. We have calibrated and monitored the sensor signal being used for temperature stabilisation. As a result we have found the temperature to be stable at the level of a few mK, although this measurement assumes that the thermistors are accurate to this level and genuinely indicate the temperature of the cavities.

Although the typical drift rates are something like 5–10 times smaller than those of the unlocked lasers, we had hoped to observe significantly lower rates. Given the materials used we had aimed to achieve frequency drifts per hour well below the 24MHz natural linewidths of our transitions. In practice the

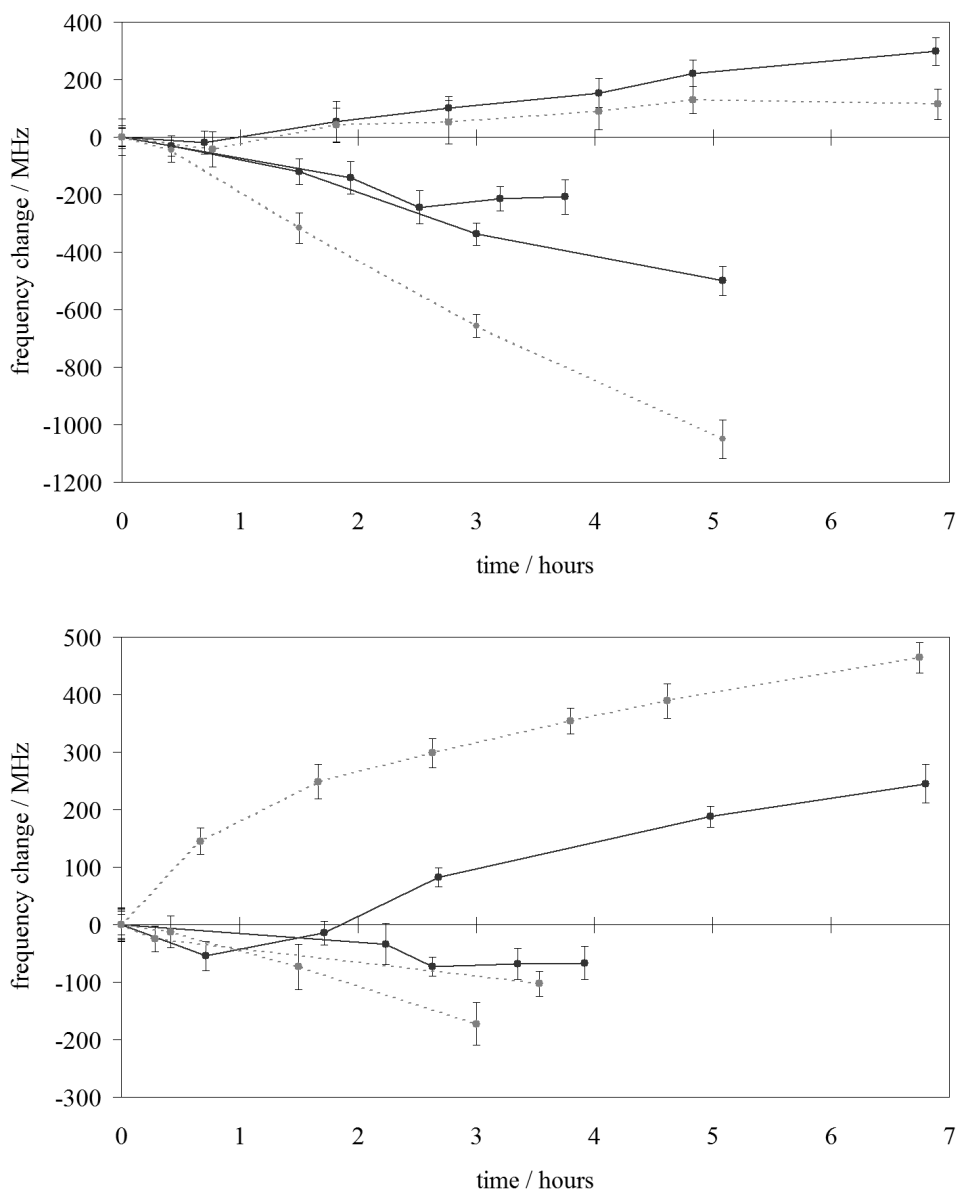


Figure 5.15: Plots showing the cavity frequency drifts on up to three days for (*top*) the 397nm cavities and (*bottom*) the 866nm cavities.

fairly poor long-term stability of the electromagnet used in our experiment means that the transitions we address are only stable at the level of a few 10s of MHz, so the laser drifts are acceptable. At some stage it is planned to begin work with a trap located in a superconducting magnet. Here the field stability will be several orders of magnitude better and the laser stability will become very much the limiting factor. The amount by which we have to correct the locked laser frequencies to maintain maximum signal in the superconducting magnet will of course give us a better test of the cavity stability.

5.8 Fluorescence detection

Ion detection and qubit state measurement are performed by monitoring the fluorescence rate from ions interacting with the 397nm Doppler cooling lasers. Assuming that all cooling and repumping transitions are saturated, each ion will on average spend an equal amount of time in each of the eight available electronic states — namely the two $S_{1/2}$, two $P_{1/2}$ and four $D_{3/2}$ states. The detected fluorescence comes from decay of the two $P_{1/2}$ states, with a rate of $1.4 \times 10^8 \text{s}^{-1}$. This means that the maximum rate of fluorescence for a given ion is about $3.5 \times 10^7 \text{s}^{-1}$. Note that in the Penning trap the cooling lasers are often offset from the centre of the trap and we might not expect the ions to interact with the laser all the time. As a result there will probably be a reduction in fluorescence of ~ 2 . Also it is important not to saturate the cooling transition in order to prevent power broadening of the transition causing a significantly raised Doppler limit. By definition the saturation intensity is that which gives us half the scattering rate, so we can expect to work with somewhat less than this.

Fluorescence that passes through a hole in the top of the ring electrode structure is collected and imaged by a 3 quartz lens system onto a small aperture placed in front of a photo-multiplier tube (PMT) as shown in figure 5.16. The lens system was designed to be a symmetrical system for use with 280nm light from Mg^+ ions. The shift in wavelength for use with calcium means that the system is slightly asymmetric, but the difference is slight and extra aberrations are likely to be correspondingly small. A low-pass filter with a cutoff at 435nm (Comar 435IK25) and band-pass filter (Comar 395GB25) are placed on top of the aperture to reduce background light and light scatter from the 866nm repumper lasers.




Figure 5.16: The imaging system used to detect ion fluorescence (not to scale). The indicated focal lengths of lenses are in mm.

The current from the PMT (Thorn-EMI 9893QB running at 2.3kV) is passed through a discriminating amplifier (Electron Tubes AD2) and an ECL-TTL pulse converter (Electron Tubes ET1). This arrangement produces a TTL pulse every time the photocurrent exceeds a threshold corresponding to a photon detection. This signal is cloned and one copy is sent to a multi-channel scalar (EG&G-Ortec MCSplus), which counts and stores the number of photons detected in a set time interval. The MCS is software controlled, allowing us to view data in real time and change data acquisition settings. The bin time can be set as low as $2\mu\text{s}$, but typically we work with it set to 10ms.

For spatial imaging of the ions, a mirror in a flipping mount can be used to redirect the fluorescence onto the photocathode of a 1 megapixel intensified CCD camera (Andor DH-534:18S-03) as shown in figure 5.16. This combines a high peak quantum efficiency (of about 20%) and gain with a pixel size of $14\mu\text{m}$ and a minimum spatial resolution of $22\mu\text{m}$. The image intensifier can be used as a shutter on a $\ll \mu\text{s}$ timescale, though there is a CCD readout time of at least $1\mu\text{s}$ per pixel that limits the rate at which images can be taken. The magnification of the imaging system for the camera is calculated to be 1.73 ± 0.05 , giving us an image resolution of about $13\mu\text{m}$. Although a larger

magnification would give us improved resolution it is important not to spread the fluorescence signal over too large an area of the CCD since there is a dark count associated with each pixel. An identical pair of filters to those described above is placed in front of the ICCD to reduce unwanted light.

The limiting aperture in terms of light collection is a hole of diameter 5.8mm in the ring holder which lies 13mm from trap centre. This represents a solid angle of $\pi 2.9^2/13^2 = 0.16$ steradians. We can assume that the fluorescence is non-directional and therefore the fraction we collect is $0.16/4\pi = 1.2 \times 10^{-2}$. The transmission of each quartz lens and the vacuum window is approximately 0.92 and the low-pass and band-pass filters have a transmission of 0.70 and 0.80 respectively. Ignoring losses at the aperture we can therefore expect a fraction 0.40 of the collected light to be imaged onto the PMT. Given the specified quantum efficiency of 0.23 we can calculate the overall detection efficiency as being 1.2×10^{-3} . A single ion fluorescing at the maximum possible rate of $3.5 \times 10^7 \text{s}^{-1}$ would give us a signal rate of $4.4 \times 10^4 \text{s}^{-1}$. Since we anticipate a scattering rate less than a quarter of the maximum we might expect a signal rate of a little under 10^4s^{-1} .

5.9 General procedure

In order to demonstrate how the various parts of the experiment described above fit together, I shall now briefly describe the procedure which is carried out when we trap and cool a cloud of ions.

We tend to operate the trap at a magnetic field of 0.98 tesla. This corresponds to a current through the coils of the electromagnet of 15.0 ampere. In order to allow time for the temperature of the magnet coils to stabilise we usually run up the magnetic field early on in the procedure.

Initially we need to set each of the laser wavelengths using the wavemeter. Each laser is in turn sent to the wavemeter and tuned to its transition as detailed in table 5.1. These values are measurements of the wavelengths previously found to be optimal for cooling large ion clouds. They are taken directly from the wavemeter output and have not been corrected to account for its calibration. If used, the 854nm laser is tuned close to the zero-field transition at 854.21nm.

A sawtooth function applied to the PZT of the laser grating is used to repeatedly scan one of the 397nm lasers up to the previously set frequency

396.8422
396.8521
866.1870
866.2424

Table 5.1: Optimal wavelengths in nanometres for cooling large clouds at a magnetic field 0.98 tesla, as taken directly from the wavemeter output.

over a range of $\sim 2.5\text{GHz}$ about once per second. This is important in order to rapidly cool in the hotter ions in the trap. The Fabry-Perot etalon is used to check that this laser remains single-mode throughout the frequency sweep and that the static cooling laser's output is also clean. The spectrum analyser is used to check the spectral purity of the 866nm beams and to ensure that the sideband amplitude is adequate.

It is important that the beams are offset slightly from trap centre in the appropriate direction to prevent magnetron heating of the ions. Furthermore the beam offset helps cool the hotter ions during loading of the trap, since these can have very large orbits and hence will otherwise not spend much time interacting with the beam. Particularly for large clouds we will often load with a large beam offset and gradually sweep the beam back in towards trap centre once a stable signal is observed. This process can be performed in a precise manner using the computer controlled mirror mount.

With the PMT running and the MCS set to a bin time of 10ms, we load the trap. We generally begin with a large cloud since this makes us less sensitive to laser frequencies and beam position. Typically we might run the oven at 1.70A for 20 seconds before also switching on the filament at 4.60A for 15 seconds. Usually we will observe some fluorescence from the ions whilst the filament is running and we can vary the time the oven and filament are left on depending on how much of this we see. After loading, the filament bias can be changed to match the endcap bias in order to minimise distortion of the trapping potential.

We usually find that there is a period of about 10 seconds where the ions cool from their initial temperature distribution to a more steady-state situation. Depending on laser parameters this will often lead to a drop in fluorescence, since the Doppler widths of the transitions become smaller and we are more sensitive to the laser frequencies. Providing we have set the




Figure 5.17: Trace showing the fluorescence signal from a large cloud of ions after optimisation of the laser parameters. The trace is an average of several frequency scans of the laser.

wavelengths precisely enough we can expect some fluorescence to remain and we can use this as our diagnostic to optimise the various parameters. Since the wavemeter can measure infra-red wavelengths more precisely than those in the UV, we tend to find that it is tuning of the 397nm lasers and beam positioning that give the most initial gain in signal. Usually we will have to iterate several times between adjustment of all the laser wavelengths and beam positions before large, stable fluorescence peaks are seen. An example trace of a large cloud of ions after the signal has been optimised can be seen in figure 5.17.

Even if the laser parameters are already optimised for a given cloud size it can take several frequency scans of the laser to cool in the ions after loading. To illustrate this we set the laser parameters to give optimum cooling of a large cloud of ions in the trap. We then expelled these ions and immediately reloaded a new cloud using the same loading parameters and without moving the laser beam. The obtained fluorescence trace is shown in figure 5.18. Although some fluorescence is observed in the period during and immediately after loading, it takes several scans of the laser before a steady state situation is reached. This




Figure 5.18: Trace showing the fluorescence signal during and immediately after loading of a cloud of ions. The laser parameters had previously been optimised for an ion cloud of similar size, but the beam offset was not increased during loading. It takes several frequency scans of the laser for all of the ions to be efficiently cooled and for a steady signal to be achieved.

effect can be eased by temporarily increasing the beam offset during loading.

Once we have optimised the laser parameters for an ion cloud of a given size, we can reload the trap with a smaller cloud of ions and expect to see fairly healthy signal soon after loading. We would expect to need to tune the laser wavelengths and positions slightly as we reduce the cloud size, since we can achieve lower ion temperatures and hence the optimum parameters for laser cooling are different. We tend to use somewhat different oven and filament parameters once we are trying to load small clouds of reproducible size. The motivation for this and details of the approach we use will be discussed in chapter 7.

Once we are at the stage of reducing the cloud size we would usually lock one of the 397nm and both of the 866nm lasers to their reference cavities. It is at this stage that we would typically make use of the axialisation technique for improving the efficiency of Doppler cooling in the Penning trap. I presented a theoretical description of this technique in the previous chapter, and will discuss the experimentally observed effects in the following chapter.

Chapter 6

Axialisation

We showed in some detail in chapter 4 that laser cooling of the radial modes of motion in the Penning trap is inherently inefficient. The magnetron motion is unstable and so requires a blue-detuned laser for efficient Doppler cooling. This is the opposite detuning to that required for cooling of the modified cyclotron motion and hence simultaneous cooling of both motions is not usually possible. By introducing an offset in the position of the cooling lasers it is, however, possible to introduce slight cooling of the magnetron motion at the cost of greatly reduced cooling of the modified cyclotron motion.

Axialisation works by coupling the two radial modes together with a quadrupolar RF drive at the sum of their motional frequencies, and leads to an averaging of the cooling rates for the two motions. Thus we can in principle obtain efficient cooling of both modes without having to offset the beam positions. We also saw in chapter 4 that a side effect of this coupling is the shifting of the motional frequencies involved, so that when the axialisation drive is near resonant an avoided crossing is expected.

Our group at Imperial College has already used axialisation to cool magnesium ions to the Doppler limit [91, 92]. Magnesium possesses no accessible metastable states and so it is easy to form a closed laser cooling cycle. Unfortunately the absence of such states means that magnesium is not well suited to quantum information processing.

In this chapter I shall detail the results of experiments carried out at Imperial College to demonstrate the axialisation technique and hence increase the efficiency of laser cooling of calcium in the Penning trap. These current experiments with calcium represent the first time that axialisation has been

used with an ion that is suitable for quantum information processing. I shall begin by describing the experimental procedure for axialising a cloud of ions. In order to enable us to quantify the improvement in cooling efficiency that results, I shall at this stage introduce an experimental technique for measuring the cooling rates of the various motions. Next I shall present the results we obtained, both in terms of the dimensions of the ion cloud as recorded on an ICCD camera and the measured laser cooling rates. Finally I shall show that we do indeed observe a classical avoided crossing that closely matches that which we obtained from the theoretical treatment presented in chapter 4.

6.1 Experimental methods

In order to couple together the two radial motions we need to apply a quadrupolar RF drive in the radial plane. We do this by splitting the output from a function generator into two parts and inverting one of them. Each of the two resulting signals is applied to a diagonally opposing pair of ring segments. In addition to the quadrupolar drive we will often want to apply a dipolar RF field. To achieve this we add the output from a second function generator to just one ring segment, and its inverse to the diagonally opposing segment. Details of the generation of these signals were given in section 5.3.5 and the setup is shown diagrammatically in figure 5.6. Since the RF is biased about ground, we must also apply a positive trapping potential to the endcaps.

To stimulate the exchange of amplitude between the magnetron and modified cyclotron modes, the photon energy of the axialisation drive must match the difference in phonon energies of the two motions. Since the magnetron mode is unstable its energy relative to the motional ground state is negative. This means that the frequency of the axialising drive must match the *sum* of the two radial motional frequencies. As we showed in section 4.2, this sum is necessarily equal to the true cyclotron frequency, $\omega_c = eB/m$. The correct frequency for axialisation is therefore dependent only on the magnetic field strength at the centre of the trap.

We monitor the magnetic field by measuring the current through the coils of the electromagnet, with a precision a little better than 1%. For large currents the rate of change in magnetic field strength with current drops significantly. This means that in the region of 10–15 ampere a 1% change in current leads to a change in magnetic field strength of less than 0.5%. It is with this

precision that we can initially set the magnetic field. Since we typically operate at about 0.98 tesla, the cyclotron frequency for calcium is close to 380kHz. Thus the initial variation in cyclotron frequency is usually less than 2kHz.

We can fine-tune the axialisation frequency experimentally in either of two ways. The quickest and most straightforward method is to image a cloud of ions on the ICCD camera and look for the changes that occur in the shape of the cloud when the axialisation drive is on resonance. Since the optimum laser parameters are different in the presence of axialisation, the improvement in cooling due to the coupling of the motions may be offset by the fact that the laser positions and frequencies are no longer optimal. Therefore changes in the cloud shape can sometimes be somewhat subtle. A small drive amplitude, perhaps 50mV, gives the best sensitivity to the RF frequency. It is, however, still difficult to discern a difference in cloud shape when the axialisation frequency is changed by less than about 200Hz.

The second technique is to measure the two radial motional frequencies independently by applying a very small (<1 mV) dipolar field. When this excitation drive is resonant with a motional frequency there is a significant change in the amount of fluorescence recorded on the PMT. By using a sufficiently small drive amplitude the motional resonances can be determined with a precision of about 10Hz, though it is worth bearing in mind that magnetic field drifts can result in changes in the motional frequencies much greater than this. The correct axialisation frequency is, of course, determined by adding the two measured motional frequencies together.

Once the correct axialisation frequency has been obtained it is necessary to reposition the laser beams. This will invariably involve moving the lasers closer to trap centre, although it is still possible to heat ions from the trap if the lasers are moved too far to the wrong side of the trap. Either the PMT or the ICCD can be used for detection, but the latter is often the most useful since it provides spatial information in addition to the overall fluorescence rate. The optimal position for the beams will depend on the amplitude and any detuning of the axialisation drive.

As we will see in detail later in this chapter, the amplitude of the quadrupole drive is important for optimising laser cooling. Too small a drive will lead only to weak coupling and a relatively small effect. It is, however, also possible for the drive to be too strong, leading to reduced cooling efficiency. The exact cause of this is not known for sure, but it is possible that with too

strong an axialising field the orbits of the ions are changed so much that they can no longer be addressed effectively by the cooling lasers. It may also be that at large axialisation drive amplitudes there is off-resonant heating of the ions.

In order to quantify the efficacy of axialisation we make use of an RF-photon correlation technique to measure the cooling rate in the system. We showed explicitly in section 4.4.4 the response of an ion's motion to an external driving force close to one of the motional frequencies in the presence of axialisation. The ion was seen to oscillate in phase with the force when the the driving frequency is much lower than the natural frequency of motion. As we pass through resonance there is a change in the relative phase of the motion and the driving force. Once the driving frequency is much larger than the natural frequency the motion is in antiphase with the drive. As we saw earlier, the behaviour near resonance can be rather complicated. When axialisation is not present, however, things are much more straightforward.

Let us consider the motion of a particle in a one-dimensional harmonic potential. In the absence of any damping or driving forces the motion is governed by the equation

$$m\ddot{x} = -kx. \quad (6.1)$$

The solution is, of course, simple harmonic motion with the natural frequency ω_n given by

$$\omega_n = \sqrt{\frac{k}{m}}. \quad (6.2)$$

If we add a damping term γ and driving term $F \sin \omega t$, the equation of motion becomes

$$m\ddot{x} = -kx - \gamma\dot{x} + F \sin \omega t. \quad (6.3)$$

Looking for solutions that oscillate at the driving frequency ω we try

$$x = a \sin \omega t + b \cos \omega t.$$

Differentiating and substituting this into the equation of motion gives

$$(k - m\omega^2)(a \sin \omega t + b \cos \omega t) = -\gamma\omega(a \cos \omega t - b \sin \omega t) + F \sin \omega t.$$

Separating out the sine and cosine components enables us to solve for a and




Figure 6.1: Plot showing the variation in the motional phase of a damped harmonic oscillator relative to that of a sinusoidal driving force as a function of its frequency. The resonant frequency here is 25kHz and the damping rate is 1kHz.

b , obtaining

$$a = \frac{F(k - m\omega^2)}{(\gamma\omega)^2 + (k - m\omega^2)^2} \quad \text{and} \quad b = \frac{-F\gamma\omega}{(\gamma\omega)^2 + (k - m\omega^2)^2}.$$

The phase of the motion relative to the drive, ϕ , is given by

$$\phi = \arctan\left(\frac{b}{a}\right) = \arctan\left(\frac{-\gamma\omega}{m(\omega_n^2 - \omega^2)}\right), \quad (6.4)$$

where we have substituted in for k using equation 6.2.

The arctan function gives us the expected π phase shift as we pass through the resonance at $\omega = \omega_n$, as shown in figure 6.1. The width of the function is given by the γ/m factor, and so gives us a direct measure of the damping rate due to the laser cooling, γ . Since we will always be using the same ion and we are only interested in relative cooling rates we will henceforth drop the factor m for convenience. This means that what we call the cooling rate γ is given simply by the frequency range over which the phase shifts from $+\pi/4$ to $-\pi/4$, relative to the centre of the function.

Providing the motional coupling due to axialisation is strong enough that the two resonances we expect to see for each motion are well separated, the form of each resonance closely matches the arctan function derived above. Thus by fitting this function to the resonances observed when we drive the ions with a dipole field, we can determine how the cooling rates of each motion are affected by the presence of the axialising field.

Measurement of the phase of the ions' motion relative to the driving field is made using a RF-photon correlation technique [102]. Fluorescence from a given ion is modulated at the frequency with which the ion is oscillating in the trap. This modulation occurs for two reasons. Firstly the ion moves about within the intensity profile of the laser, and hence fluoresces more during certain parts of the motional period. Secondly the varying Doppler shift experienced by the ion as it oscillates takes it in and out of resonance with the laser frequency. If we are driving the ions at a given frequency, we will therefore observe correlations between the phase of the driving field and that of the modulation in photon emission.

This correlation is observed by looking at the time delay between some arbitrary phase of the driving field and the time at which the next photon is received. The apparatus for doing this is shown schematically in figure 6.2. With all the laser frequencies static, we trigger the start gate of a time-amplitude converter (EG&G-Ortec TAC 566 or 567) with the TTL output from the function generator producing the excitation drive. The TAC then measures the time that passes until it receives at its stop gate a TTL pulse from the PMT, corresponding to a photon detection. This time interval is converted into the amplitude of a signal pulse sent to a multi-channel analyser (EG&G-Ortec MCA), which measures the amplitude and records the number of times that pulses of a given amplitude have been received.

Photon detection from a given ion is more likely at a particular point in its motional period. This occurs at some arbitrary time after the TAC starts timing, and then again each motional period after that. If we plot a histogram of the number of times that we observe a given time delay we might therefore expect to see a modulation at the drive frequency. When the drive is weak, or not near resonance, there is a significant spread in the phases of different ions' motion. Hence if we have many ions in the trap there is a smearing out of the correlations between emitted photons and the phase of the driving field. If we have a very weak drive we will find that the photons are emitted

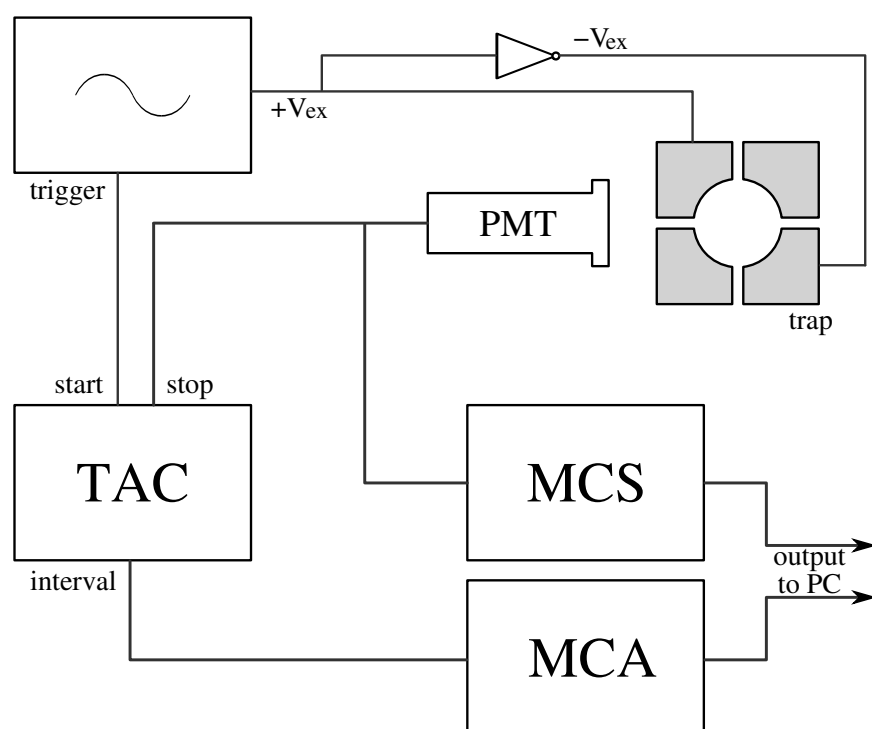


Figure 6.2: Schematic of the apparatus used to measure correlations between photon emission times and the phase of the excitation drive.

at apparently random times and we would expect our histogram to show only an exponential decay corresponding to random emission. When the drive is strong, or near resonance, the spread in the ions' motional phases is reduced and hence we do begin to see modulation at the driving frequency. In general we see a modulating signal superposed on an exponential decay, as shown in the top plot in figure 6.3.

Although the modulation depth is generally smaller when we have more ions, we can increase it by driving the ions more strongly. This tends to make the ions travel round the trap as a single group and leads to very strong correlations between the times at which photons are emitted and the phase of the driving field. The lower plot in figure 6.3 shows a MCA trace where this is the case, and the modulation depth is seen to be close to unity. Driving the ions too hard, however, can heat them more strongly than they are being cooled. Thus when choosing the amplitude of the excitation drive we must balance modulation depth against the risk of losing ions from the trap.

We usually begin by monitoring the fluorescence detected with the PMT as we tune the frequency of a very small ($< 1\text{mV}$) excitation drive. The change in fluorescence that occurs near resonance tells us what the centre frequency is. With the excitation drive set to a chosen frequency near resonance, and still at a very low amplitude, we then start the MCA counting the time delays. After a few seconds it is usually possible to determine whether the modulation depth is sufficient. If it is not we increase the drive amplitude slightly and start counting again. Once the modulation depth is acceptable the MCA is left running until clear oscillations are obtained that will enable us to accurately determine their phase. This will typically take anything between 20 seconds and a minute, depending on the rate of fluorescence from the ions.

In order to extract the phase, ϕ , of the modulation in fluorescence relative to the excitation drive, we fit the obtained signal to the function

$$f(\Delta t) = e^{-a\Delta t} (b + c \sin(\omega\Delta t - \phi)). \quad (6.5)$$

This corresponds to exponential decay with rate a from an initial value b , modulated at the drive frequency with a depth c . As an example the plots in figure 6.3 have had this function fitted to them.

Since we trigger the TAC from an arbitrary point in the cycle of the driving field, the phase we obtain is itself arbitrary. What we are interested




Figure 6.3: Example RF–photon correlation traces with (*top*) a modulation depth typical of those used for cooling rate measurement, and (*bottom*) a modulation depth near unity corresponding to a very strong excitation drive amplitude. In each case the data have been fitted to equation 6.5.

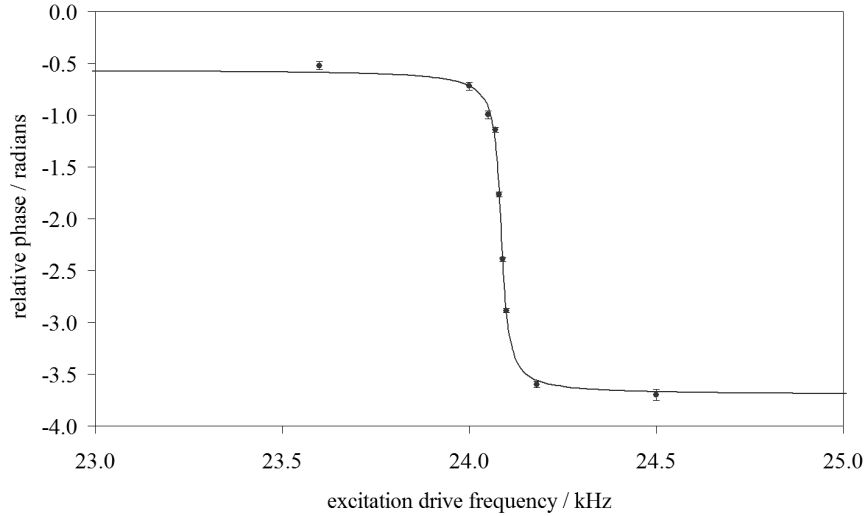


Figure 6.4: Plot showing the experimentally determined phase of ion motion relative to that of an excitation drive near the magnetron frequency. The data have been fitted to equation 6.4 in order to determine the laser cooling rate, γ , to be 27 ± 2 Hz.

in, however, is how that phase changes with the drive frequency. In order to map out the phase function we use a range of excitation frequencies that may extend as far as 10kHz either side of resonance. We sample it with more points where the phase is rapidly changing, so that near resonance we might use drive frequencies separated by as little as 5–10Hz. Further away from resonance we find that we need to use large excitation amplitudes in order to obtain a significant modulation depth. About 10kHz away we might require almost 100mV, whereas very near resonance a fraction of a millivolt can be sufficient. To avoid heating the ions from the trap we begin each measurement using a small amplitude and raise it gradually as required. In general we try to take measurements with different drive frequencies in a random order, so as to smear out the effects of magnetic field drifts.

A typical plot showing the phase difference as a function of drive frequency for the magnetron motion is shown in figure 6.4. We have fitted the data shown to an arctan function in order to determine the width of the resonance, and have used equation 6.4 to deduce the cooling rate from this (ignoring the mass factor).

Finally we note that since the position of the laser beam affects the point in the ions' motional cycle at which the fluorescence is greatest, it is important that the beam is not moved during a given cooling rate measurement. Doing so would introduce an additional change in phase. Between different cooling rate measurements the beams can be moved, as is required to optimise the laser cooling.

6.2 Results

6.2.1 iCCD images

As mentioned in the previous section, it is usually immediately apparent when the correct axialisation drive frequency has been reached because there is a significant change in the shape of the ion cloud as viewed on the iCCD camera. Initially the ion cloud will tend to be more extensive in the radial plane than in the axial direction, since the radial cooling is weaker. This effect is most pronounced when the trap bias is very large, since the cloud is then more compressed in the axial direction and forced to spread out in the radial plane. Once the axialisation is applied this ratio of radial extent to axial extent can be dramatically reduced.

Figure 6.5 shows iCCD images of the same cloud of ions in a Penning trap with various trap biases. In each image the vertical axis corresponds to the axial position in the trap, and the horizontal axis to the radial position. The pixel size is approximately $7.5\mu\text{m}$. Images (a) and (b) are taken before and after the application of a 300mV axialisation drive when the trap bias is 7.8 volt. By integrating across the image we can obtain a near Gaussian intensity profile in each direction, as shown in figure 6.6. The asymmetry that is visible in these plots results from the smearing in one direction that is observed on each iCCD image. This smearing is reproducible and occurs regardless of the size of the ion cloud or the method of trapping. It is attributed to aberrations introduced either by the imaging system or the image intensifier in the iCCD. Despite the asymmetry on the plots, a good fit to a Gaussian function is nevertheless obtained in the region of maximum signal. Fitting only to this central region, and without factoring in the asymmetry, means that a significant fraction of the counts is not included. As we shall see shortly, there is a significant variation in aspect ratio with the laser cooling parameters and

so it is not the precision with which we can determine the aspect ratios that limits the quality of the conclusions we can draw from analysis of the ICCD images.

The Gaussian fits enable us to determine the aspect ratio of the cloud shown to be 4.72 (+0.51, -0.10) before and 1.70 (+0.09, -0.05) after the application of axialisation. That the aspect ratio does not drop further can be attributed to the effects of space charge. Mutual repulsion between the ions means that we cannot keep increasing the ion density indefinitely. Instead the ion cloud is forced to stretch in the axial direction. The effect of the axialisation drive is therefore working against the trap bias and an equilibrium position is reached. When fewer ions are present we would, for these trap parameters, expect this equilibrium aspect ratio to be smaller. With a single ion there is no space charge effect and we ought to be able to use axialisation to reduce the ion's radial motion until we reach the Doppler limit. This has already been performed with magnesium [91, 92].

With a smaller trap bias the change in the aspect ratio of the cloud is less obvious by eye, but from measurements of the aspect ratios the improvement is still found to be significant. Images (c) and (d) in figure 6.5 show a cloud before and after axialisation, with a trap bias of 3.6 volt. The measured aspect ratio drops from 1.73 (+0.09, -0.05) to 1.25 (+0.05, -0.04) with axialisation.

The errors on these aspect ratios are due predominantly to the resolution of the image intensifier on the ICCD. This resolution is specified as $12.7\mu\text{m}$ (taking into account the magnification), about $5\mu\text{m}$ bigger than the pixels in the images. Thus the true extent in each direction can in fact be anything up to $5\mu\text{m}$ *smaller* than that measured. This error is obviously larger on small measurements and so tends to give an artificially small aspect ratio.

It is important to note here that the laser cooling parameters can have a drastic effect on the aspect ratio of an ion cloud. To demonstrate this, we have measured aspect ratios of a cloud before and after optimising the laser parameters to be 1.87 and 1.36 respectively. Although this is only one example, it demonstrates that this measurement procedure is susceptible to errors that are not reflected in the measurement precision. To minimise this effect we reoptimised the laser parameters before each of these images was taken.

We have studied the aspect ratios of clouds of ions as the axialisation amplitude is gradually changed from 0 to 800mV. We find that the most




Figure 6.5: iCCD images of the same ion cloud with and without axialisation at different trap biases. In images (a) and (b) the trap bias is set at 7.8 volt and the aspect ratio drops from 4.72 to 1.70 when axialisation is applied. In images (c) and (d) the trap bias is 3.6 volt and the aspect ratio changes from 1.73 to 1.25. The estimated region over which the beam is saturating the cooling transition is the same in each case, and is indicated by the scale bar.

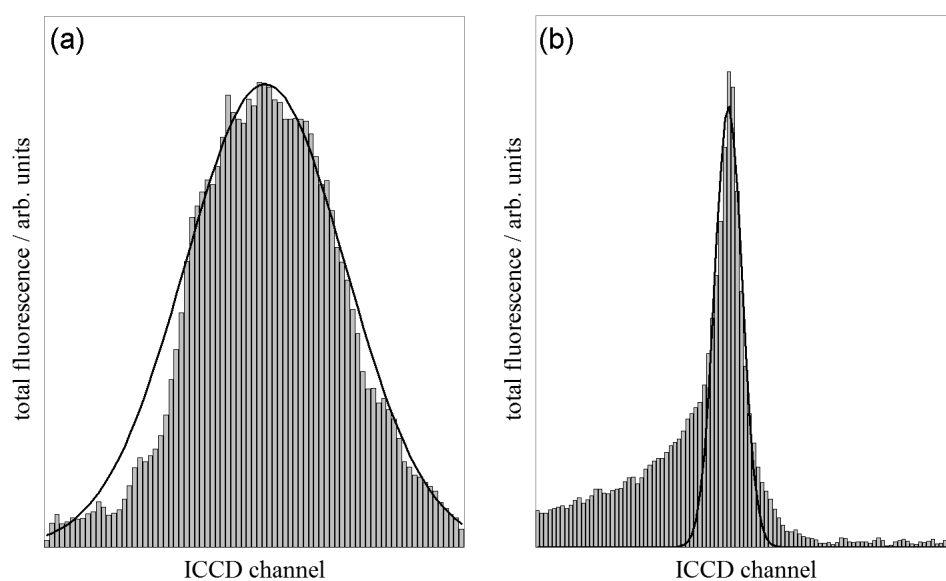


Figure 6.6: Histograms showing the integrated fluorescence from each channel in (a) the radial direction and (b) the axial direction of the image in figure 6.5(a). The full widths at half maximum of the Gaussians fitted to the points near the centre of each histogram are used to calculate the stated aspect ratio of 4.72.




Figure 6.7: Aspect ratios of two ion clouds as a function of axialisation amplitude. The filled circles correspond to the smaller of the two clouds.

rapid change occurs at rather low amplitudes and that after about 150mV the changes are generally small. Ideally we would carry out this experiment with just one or two ions in the trap. Because of difficulties described in chapter 7 this has not been possible and we have been forced to work with clouds of at least a few tens of ions. In order to obtain some indication of the effect of this change in cloud size, we have tended throughout this work to use a number of clouds, each with different sizes.

Figure 6.7 shows the aspect ratios of two different sized clouds of ions as the axialisation amplitude is varied. At each amplitude the laser parameters are reoptimised in an attempt to maximise the laser cooling rate. The typical fluorescence rate from the smaller cloud is about 2×10^4 counts per second when unaxialised, suggesting that this is in the region of 20–40 ions. The fluorescence from the larger cloud is about double this.

With both clouds the aspect ratio is seen to drop rapidly as the axialisation amplitude is first increased. When the drive is 800mV the aspect ratio for the smaller cloud is actually measured to be unity.

In figure 6.8 we plot the total fluorescence detected by the ICCD as a func-

tion of axialisation amplitude for the same series of images of the two clouds. The absolute fluorescence rates cannot be compared between clouds, since we have used different exposure times and the gain on the image intensifier is not necessarily the same. For the larger cloud we find that there is a slight increase in fluorescence as we turn on the axialisation drive. The situation appears to become very slightly worse if we use amplitudes of about 400mV or greater. The fluorescence from the small cloud initially increases by about 25%, but degrades significantly at higher amplitudes. We might expect a slight reduction in fluorescence from a large cloud of ions, since space charge effects mean that reducing the radial extent of the ion cloud necessarily requires increasing the axial extent. If the axial size of the cloud is larger than the beam size then increasing it further will push more of the ions out of the beam and hence lead to less interaction. It is, however, unexpected that a small cloud should exhibit this drop in fluorescence at a lower drive amplitude.

6.2.2 Cooling rate measurements

We have made cooling rate measurements on a number of different clouds with a wide range of different axialisation parameters. Each cooling rate measurement requires a whole series of RF-photon correlation traces to be taken, and so takes something in the region of 30–60 minutes. Over this time even the locked lasers will drift far enough that without reoptimisation a good deal of the signal will be lost. It is therefore necessary for us to retune the lasers throughout the course of these measurements, and since we have only the fluorescence rate as an indicator it is likely that the cooling rate will actually vary by an amount that is hard to quantify. In order to try and gain some handle on the dependence of the cooling rate on the laser parameters, we note that without axialisation the cooling rates measured are typically in the region 5–15Hz, but that on rare occasions rates as high as 30Hz have been recorded. This sixfold variation may be in part due to the fact that different numbers of ions have been used, but even in measurements with the same cloud the cooling rates can vary somewhat. We find that it is much easier to maintain constant fluorescence when the size of the ion cloud is larger and hence we have often been forced to use clouds with more ions than we would ideally like.

Recall that in the presence of axialisation it is expected that we will observe




Figure 6.8: Background-corrected fluorescence rates from two ion clouds as a function of axialisation amplitude. The filled circles correspond to the smaller of the two clouds. The rates for the two clouds are not directly comparable since the exposure times of the two sets of images and the intensifier gains are not the same. Although the errors in measurement are very small compared to the data points on the plot, there is a noticeable variation between the repeated measurements due to uncontrolled changes in the laser cooling parameters.




Figure 6.9: Phase plot near the magnetron frequency with an axialisation amplitude of 400mV. Two distinct resonances are visible, each of which can be fitted to a separate arctan function.

two resonances for each motion, separated by an amount dependent on the degree of coupling. This is indeed found to be the case, as demonstrated by the plot in figure 6.9. Here an axialisation drive of 400mV has been applied and we are looking at the resonances near the magnetron frequency. When the resonances are separated by such a large degree it is possible to fit an arctan function to each resonance and obtain a cooling rate for each motion. In this instance the cooling rates are found to be 53 ± 11 Hz and 90 ± 26 Hz respectively, where the quoted error is obtained from the quality of fit.

Since it is the cooling rate of the magnetron motion that we are most interested in, let us begin by looking at this as a function of axialisation amplitude when the drive is close to resonance. We have taken several such sets of data, two of which are shown in figure 6.10. The first plot shows the results obtained using a cloud that exhibited a maximum fluorescence rate of 1.5×10^4 counts per second. From this we might estimate that there are 30–50 ions present. There is a distinct trend of increasing cooling rate as the axialisation amplitude is increased, and although the gradient is shallow there are several measurements that are higher than the maximum cooling rates

we have obtained without axialisation. That there are some points where the cooling rate dips below this trend is ascribed to changes in the laser parameters so that the cooling is less efficient at these points.

In the lower plot of figure 6.10 the cooling rates are plotted for a cloud of ions with a fluorescence rate of about 7.5×10^4 counts per second, so the number of ions is expected to be significantly greater than in the previous case. For the first few points there is again an increase in cooling rate with drive amplitude, though in this case it is with a steeper gradient. At some point in the region of 200–400mV, however, the signal rate drops back down. Even if we were to ignore the very high cooling rate measurement at 200mV we would see that there is no improvement in cooling rate beyond 200mV. That this degradation does not occur in the data from smaller clouds can be ascribed to the fact that larger clouds can occupy regions of the trap further from the centre, where the oscillations in the electric field are greater. They might therefore be expected to be more sensitive to the axialisation drive. We would therefore need a much larger drive amplitude to see the same effect in a smaller cloud.

It is interesting to look at the cooling rates of the magnetron and modified cyclotron motions of the same ion cloud when the axialisation amplitude is varied. These are plotted in figure 6.11 for the same large cloud as was used in the lower plot in figure 6.10. The cooling rate of the modified cyclotron motion is seen to drop almost exponentially. Once a drive amplitude of 200mV is reached the measured cooling rates for the two motions are rather similar. From our theoretical treatment of axialisation we expect the cooling rates of the two motions to be the same when the coupling rate is strong, though this should occur at the average of the two unaxialised cooling rates. As we increase the drive amplitude further, however, the cooling rates of both motions fall together.

We have also studied the variation in magnetron cooling rate as a function of the axialisation drive frequency, when the amplitude is kept constant. We worked with a cloud similar in size to that used for the previous result, and a drive amplitude of 200mV was used throughout. The results are shown in figure 6.12. The dashed line represents the true cyclotron frequency as determined from measurements of the unperturbed magnetron and modified cyclotron motional frequencies. It does seem that with the frequency closer to resonance the measured cooling rate is slightly larger, though there are some




Figure 6.10: Magnetron motion cooling rates measured as a function of axialisation amplitude for (*top*) small and (*bottom*) large clouds.




Figure 6.11: Cooling rates of both the magnetron (filled circles) and modified cyclotron (open circles) motions, measured as a function of axialisation drive amplitude.

values near resonance where the rate is unexpectedly low. Presumably this is once again due to variation in the laser parameters. The best cooling rate measured occurred when the axialisation drive was detuned by about 700Hz, though this may be due to the rather large errors derived from the fitting of the phase to an arctan function. It is not clear why there should be any requirement to detune the drive, unless its amplitude is too strong. Even if this were the case there is no reason why it should need to be detuned in a particular direction. One thing that is apparent from these results is that we can afford to detune the axialisation drive by 1kHz or more and still see an appreciable improvement in cooling rate over the unaxialised case. Accurate probing of the region within a few hundred hertz of resonance would only really be possible in the presence of a more stable magnetic field than that currently available.




Figure 6.12: Cooling rates of the magnetron motion, measured as a function of axialisation drive frequency. The dashed line corresponds to the true cyclotron frequency as determined from measurements of the unperturbed magnetron and modified cyclotron frequencies. For comparison, measured unaxialised cooling rates are typically in the region 5–15Hz, but have on rare occasions been as high as 30Hz.

6.2.3 Classical avoided crossing

As we have already seen, we do indeed see the additional resonance frequencies we expected from our theoretical analysis of the coupling due to axialisation. In section 4.4.3 we showed that in the strong coupling regime we should observe a classical avoided crossing of the form shown in figure 4.11. As a function of axialisation drive frequency, ω_a , the motional frequencies (in the laboratory frame) should vary as

$$\omega = \omega_0 + \left(\frac{\omega_a - \omega_c}{2}\right) \pm \sqrt{\left(\frac{\omega_a - \omega_c}{2}\right)^2 + \left(\frac{\epsilon}{2\omega_1}\right)^2}. \quad (6.6)$$

Here ω_0 is the unperturbed motional frequency and ϵ/ω_1 is a coupling parameter, the magnitude of which governs the extent to which the cooling rates of the radial motions are coupled when we are near resonance.

The phase plots that we used to measure cooling rates for the radial motions can also be used to provide precise values for the resonance frequencies of the motions with different axialisation frequencies. These have been plotted for the magnetron motion with a 200mV axialisation drive in figure 6.13. The error bars account for errors resulting from the determination of the resonance frequencies, the (rather poor) stability of the function generator used for axialisation and fluctuations in the magnetic field. The data have been fitted to equation 6.6, with ω_c and ϵ/ω_1 as the free parameters. Clearly this equation accurately describes the shifting of the resonance frequencies and confirms that we have indeed achieved strong coupling of the radial motions. The line of best fit is obtained with $\omega_c = 379.54\text{kHz}$ and $\epsilon/\omega_1 = 0.869\text{kHz}$. If we calculate the factor ϵ/ω_1 from our model, using the relevant trap parameters and the amplitude of the axialisation drive, we obtain about 60kHz.

The rather large difference between calculated and measured coupling rates can only partly be explained by the non-ideal trap structure and experimental imperfections. For the most part it is believed to occur because the model assumes that only one ion is present and that it is always in the centre of the trap. When there are many ions the resulting space charge means that they occupy orbits with much larger radii, and so they travel through regions of the trap where the magnitude of the trapping and axialising fields are radically different. Furthermore, the model predicts that the presence of an avoided crossing requires the coupling rate to be large compared to the laser cooling




Figure 6.13: Resonance frequencies of the magnetron motion plotted as a function of axialisation drive frequency. The solid line is a fit to equation 6.6, whilst the dashed line indicates the unperturbed magnetron frequency.

rates. It is possible that in a similar manner the coupling rate must be larger when forces due to ion interactions are present.

With the axialisation drive on resonance, the shifts in the resonance frequencies should be given simply by $\pm\epsilon/2\omega_1$. Since ϵ is proportional to the amplitude of the axialisation drive we should find that near resonance the motional frequencies shift linearly with drive amplitude. We also know that the shifts for the magnetron and modified cyclotron motions should be equal but opposite, so that each pair adds to the true cyclotron frequency. Using the same technique as described above, we have mapped out the shifts in motional frequency for both radial modes as a function of axialisation drive amplitude. The difficulty here is that the frequency of the modified cyclotron motion is much more sensitive to changes in the magnetic field than the magnetron frequency is. This is clearly apparent from the results, shown in figure 6.14. The magnetron frequencies are seen to follow closely the applied linear fit, whereas there is much more scatter in the modified cyclotron data. Although we alternated between measurements of the magnetron and modified cyclotron frequencies, we performed the measurements with lower axialisation drive am-

plitudes first. Thus it appears that the magnetic field was initially rather unstable, but settled down as the experiment continued. This fits with the trend observed in previous experiments. We can estimate the magnetic field at various points by adding the measured frequencies to obtain the corresponding true cyclotron frequency. From this we obtain a magnetic field stability throughout the experiment of about 1 part in 1100, and the error bars on the plots show the resulting uncertainties in the motional frequency shifts, combined with the error from the determination of the resonances themselves.

The straight lines that have been fitted to the plot for the magnetron resonance have gradients 3.1kHz/V and -2.9kHz/V respectively. That the magnitude of these is similar confirms that the axialisation drive is close to resonance. The quality of the fits suggest that there is indeed a linear relationship between the frequency shift and the amplitude of the axialisation drive when this is the case. To confirm that the shifts for the modified cyclotron motion are indeed equal and opposite to those for the magnetron motion, the straight lines on the modified cyclotron plot have been given the same gradients as we derived from the magnetron data, but switched around. Taking into account the increased errors on these points the data do appear to be consistent with our theory.

If the axialisation drive is detuned significantly from resonance we would expect from equation 6.6 that the frequency shift as a function of drive amplitude is no longer linear. Figure 6.15 is a plot of one of the magnetron motional resonances as a function of axialisation drive amplitude when the drive is detuned by about 1kHz . Although the absence of a data point when the drive is turned off is disappointing, it is clearly apparent that equation 6.6 fits the data much better than a simple linear fit.

Finally we shall look in more detail at the phase behaviour of the system when there are two resonances present. This behaviour was discussed from a theoretical point of view in section 4.4.4. So far in this chapter we have looked only at the positions and widths of the two resonances that are observed. Figure 6.16(a) shows the results of a RF-photon correlation experiment where we have taken points at many more excitation frequencies. The dashed line indicates the position of the motional frequency when the axialisation drive is not present, as measured in an experiment immediately prior to this one. We can use the positions of the resonances, and the previous measurement of the unaxialised magnetron frequency, to establish the axialisation drive




Figure 6.14: Resonance frequencies of the magnetron (*top*) and modified cyclotron (*bottom*) motions as a function of axialisation drive amplitude. The solid line is in each case a linear fit based only on the magnetron data.




Figure 6.15: Resonance frequency of the magnetron motion as a function of axialisation drive amplitude when the drive is detuned by about 1kHz. The solid line is a fit to equation 6.6, whereas the dashed line is a linear fit.

parameters. Plot 6.16(b) shows the results obtained when we introduce these parameters into our model. Clearly there is something of a discrepancy in terms of the shapes of the two plots in the region between the resonances. This is perhaps to be expected if we recall that our model assumes that we have only one ion and that it is confined to a region smaller than the waist of the cooling lasers. Unlike the frequencies of the resonances, which are a function only of the coupling parameters, the shape of the phase response is very much related to the cooling behaviour. Thus the fact that the experimental parameters in this instance do not fully satisfy the assumptions of the model would be expected to have a significant effect here.

It is worth noting that the shape that is observed can be obtained using the model if we assume the unperturbed magnetron frequency to be slightly different to that measured experimentally. This extra degree of freedom allows us to generate the plot shown in figure 6.16(c). It is possible that the magnetic field drifted during the experiment, and that the unaxialised magnetron frequency changed in the interval between our measuring it and collecting the data presented here. The field drift required for this to be the case, however,

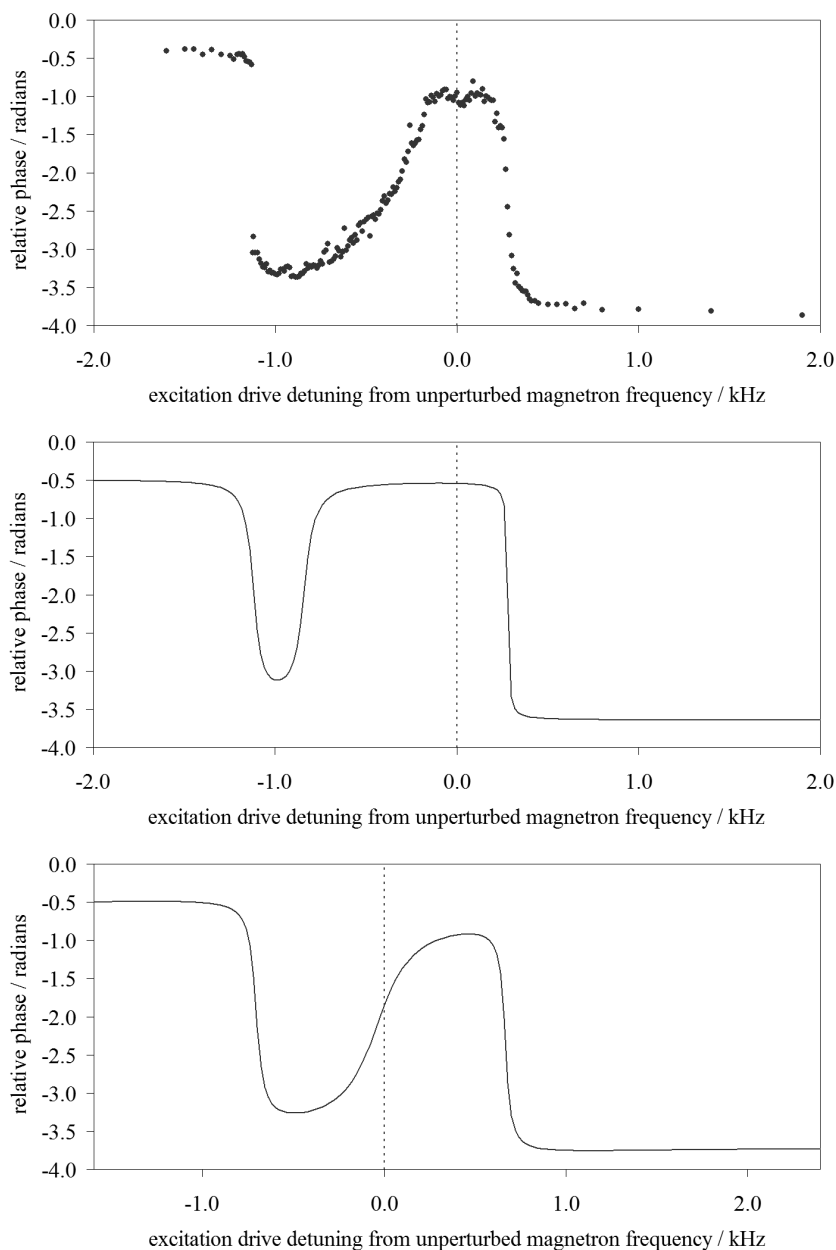


Figure 6.16: (a) Plot showing the measured phase of the ions' motion relative to an excitation drive in the presence of axialisation. The dashed line is the previously measured unperturbed motional frequency. (b) Form of the equations derived in section 4.4.4, using axialisation parameters obtained from the experimental data. (c) Form of the same equations where we have used the unperturbed magnetron frequency as an extra degree of freedom.

would have to be greater than 1%. It seems unlikely that this could have happened so abruptly and that we should then have seen no evidence for magnetic field instability on that scale during the rest of the experiment. Another reason for doubting such a change in magnetic field is the presence in the data of the small dip at the top of the central peak. The position of this dip coincides with the position of the measured unperturbed motional frequency. Though this feature is not generated by our model, it would be strange if its appropriate position was merely coincidence. In general our model allows us to precisely predict the positions of the motional resonances, and can be used to generate the basic form of the phase plots measured experimentally.

6.3 Conclusions

In this chapter we have presented results showing that axialisation improves the measured laser cooling rate of the magnetron motion, and leads to a reduction in the aspect ratio of ion clouds. We have also shown that the resonant frequency shifts and avoided crossing expected from our model in section 4.4.3 very closely match those obtained experimentally.

The measurements we have made of the cooling rates of the radial motions are currently somewhat limited by the fact that there is a good deal of variation in the laser parameters over the course of each measurement. This leads to uncontrolled scatter in the data that makes it hard to observe anything more than general trends in the data. Fluctuations in the magnetic field have also been observed at a level only slightly better than 1 part in 10^3 . These fluctuations affect our ability to keep the axialisation drive within a few hundred hertz of resonance. The fluctuations also have a large effect on measurements of the modified cyclotron motion, the frequency of which is strongly dependent on the magnetic field strength. In order to better parameterise the effects of axialisation, it will be necessary to improve both the laser and magnetic field stability.

With improved laser stability it should also be possible to repeat these cooling rate measurements using much smaller ion clouds. At the moment this is made difficult because smaller clouds are more sensitive to the laser frequency drifts that occur on the timescale of these experiments. As the number of ions is reduced the effects of space charge become smaller and we would expect to be able to cool the ions to temperatures close to the Doppler

limit. This should manifest itself in the reduction of the cloud size as viewed on the ICCD. Ultimately, of course, we would like to axialise a single ion, since it is in this case that we can most easily cool to the Doppler limit itself. We would then be in a position to begin sideband cooling and internal state manipulation of the ion, leading towards the ultimate aim of this project of measuring ion heating and decoherence rates in the Penning trap.

Chapter 7

Towards single-ion work in the Penning trap

In order to achieve our stated long-term goal of investigating the decoherence rates of ion qubits in the Penning trap, it will be necessary to isolate single ions and cool them to close to the ground state of their motion. Trapping of single ions using radio-frequency traps is now routine in many laboratories around the world. Recent progress towards quantum information processing in ion traps has for the most part been made using ions in linear RF traps, where several groups have successfully laser cooled single ions to extremely low temperatures. Most of the experiments working with single ions in Penning traps are geared towards making precise mass measurements. Typically they use electronic circuits connected to the endcaps of the trap to detect and cool the ion by measuring and damping the image currents induced by the ion's motion. This technique is limited in that an ion cannot be cooled to a temperature below that of the electronics [103]. Usually liquid helium is used as a cryogen to cool the electronics, so the lower limit on ion temperature is 4K. Besides the fact that laser cooling is the only way to achieve the much lower temperatures we require, it also has the benefit that we can use the cooling transition to measure the internal electronic state of the ion. It is for these reasons that we are interested in the laser cooling of single ions in the Penning trap.

The Imperial College Ion Trapping Group has already demonstrated the laser cooling of a single magnesium ion to the Doppler limit in a Penning trap [91]. As we have already seen, calcium has an energy level structure that

can more easily be used for quantum information processing. As a result we have for the last few years been attempting to cool single ions of calcium. Initial attempts were made in a conventional Penning trap but single ions were never observed [97]. In recent months, after the development of the split-ring trap, we have been able to retry this experiment with the benefit of axialisation. Unfortunately we have yet to observe single ions. It appears that the fluorescence rate we observe per ion is much smaller than we would expect, so that the signal rate is too poor to observe clouds with less than 5–10 ions.

In this section I shall describe our efforts to trap and cool single ions and our attempts to find the cause of the problems we have encountered. I shall begin by detailing our technique for loading progressively smaller clouds of ions until we observe single ions. In order to demonstrate that this technique does indeed work I shall show that we have successfully used it to trap single ions when we operate our trap as an RF trap. Next I will present data that we have obtained when trying to load single ions in the Penning trap. I shall then describe in detail our investigations of two possible explanations for the apparent drop in signal rate, namely the presence of contaminant ions in the trap and the driving of ions into some non-fluorescing state. Finally I shall look at other possible causes that we would hope to investigate in the future.

7.1 Loading and detection of small numbers of ions

As stated in chapter 5, ions are created in our trap by the electron bombardment of calcium atoms. Since there is no straightforward method of controllably ejecting ions from a Penning trap it follows that we must instead load just one ion. There are five parameters that we can vary to alter the number of ions that are loaded. The first four of these are the currents we pass through the oven and filament to generate the calcium atoms and electrons, and the durations for which we apply these currents. The final parameter is the bias we apply to the filament in order to give the emitted electrons sufficient energy to ionise calcium. When setting these parameters we have two objectives; we want to be able to change the number of ions we load in a reproducible way and we want as high a ratio of calcium to contaminant ions as possible.

Changing the current that we drive through an oven changes its temperature. Since different atomic species vapourise at different temperatures we

want to use a current sufficiently large to produce a significant number of calcium atoms, but not so large that many other contaminants are also produced. The ovens also have large thermal masses (compared to the filaments), so the time taken to reach an equilibrium temperature can be a minute or more. For these reasons we have for the most part used an oven current just large enough that we can easily load Calcium, which is left on for two minutes before the filament is turned on.

The current driven through the filament varies the temperature and hence the number of electrons that pass through the trap. This provides us with a good method of changing the probability of ionisation and it is this parameter that we use to select the size of the cloud of ions that we load. The relationships are non-linear, however, and we want to ensure that we are operating in a region where the rate of change of ionisation probability with current is not too steep. We typically drive current through the filament for twenty seconds, this being a long enough time that we can easily reproduce it from one load to the next.

The final variable is the filament bias, which we define as the difference in potential between the filament and the centre of the trap. This tells us the kinetic energy that an electron will possess near the centre of the trap and hence the energy that is available for ionisation. The ionisation potential of Ca^+ is 6.1 volt and so we must ensure that the filament bias is greater than this. The cross section for ionisation increases with the available energy, so that we can load more calcium if we increase the bias further. However we must be aware that there will always be other atoms and molecules present both as residual gases in the chamber and emitted by the oven and filaments. If these have ionisation energies less than the filament bias there will be a chance of ionising them in addition to calcium. Another important factor is that the potential required to produce Ca^{2+} is 11.9 volt, so we must remain below this value. Generally we use a filament bias of about 7–8 volt when trying to load small clouds of ions.

Detecting the number of ions we have in the trap is difficult when the number is large since we cannot distinguish between many ions with a small fluorescence rate and fewer ions with a large fluorescence rate. Furthermore the difficulty in efficiently cooling larger clouds of ions means that the signal rate per ion will generally fall as the number of ions is increased. When working with very small clouds of ions it becomes much easier to determine

the number of ions present because it is possible to resolve discrete steps in the signal rate corresponding to different numbers of ions fluorescing.

We have already stated that the cooling lasers can be used to determine which internal electronic state an ion is in because only if it is in one of the states addressed by the cooling transition will fluorescence be observed. It follows that if we introduce a weak laser to pump the ion into some metastable state we will see no fluorescence until the ion decays back into the ground state. We therefore see a random series of steps in the signal as the ion makes ‘quantum jumps’ between the ‘on’ and ‘off’ states, with the statistics for the times spent in each state being governed by the probabilities for pumping into the non-fluorescing state and decay back to the ground state. With just a few ions we can easily see different numbers of steps corresponding to different numbers of ions being in the ‘on’ state. The height of these steps tells us the signal rate per ion. If we then remove the laser pumping the ions into the ‘off’ state we will obtain a fluorescence rate corresponding to all the ions being in the ‘on’ state. It is therefore straightforward to determine the number of ions present. In our experiment we use the $D_{5/2}$ state as the ‘off’ state (see section 5.1). Pumping into this state is performed by introducing a laser to drive the 850nm $D_{3/2} \rightarrow P_{3/2}$ transition.

Once we have more than just a few ions the signal rate per ion tends to drop considerably for the reasons described above, and it also becomes more likely that different ions will undergo quantum jumps at similar times. This can make it hard to tell by eye what the signal rate per ion is. As an example, the fluorescence trace in figure 7.1 shows quantum jumps of several ions in a RF trap. The laser parameters are not quite optimal and so the fluorescence rate is slightly lower than usual, making it hard to resolve discrete steps in the fluorescence. Where this is the case we can use statistical analysis to estimate the number of ions present. A detailed description of this analysis process is given in reference [97]. The technique is based on the fact that both random noise and the signal we observe in the limit of a large number of ions are governed by Poissonian statistics, where it can be shown that the mean signal level is equal to the variance. The effect of quantum jumps in a signal is to increase the variance relative to the mean. If we know the variance and the mean of the signal, as well as the maximum and background signal rates, then we can estimate the number of ions present. For the fluorescence trace in figure 7.1 this method yields an ion number estimate of ~ 5 . The uncertainty




Figure 7.1: Fluorescence trace from ~ 5 ions in a RF trap. The poor signal-to-noise ratio makes determining the ion number by eye difficult. The three marked sections are (a) quantum jumps, (b) full signal without quantum jumps and (c) background. The ‘noise’ in region (a) is sub-Poissonian.

in the calculation increases with the number of ions and depends on the signal to noise ratio, but the technique is estimated to be accurate to within one or two ions in the region up to ~ 10 ions.

7.2 Single ions in the RF trap

In order to test the trap loading, laser cooling and fluorescence detection systems we have performed an experiment looking for single ions when operating our trap as an RF trap. There are two main differences in how the loading parameters affect the number of ions loaded in the RF trap, compared to the Penning trap. The first of these is that the very large oscillating trap voltages can be much larger than the applied filament bias. This means that even for an applied filament bias of zero an electron can possess a large potential energy and will be able to ionise a large number of different atomic and molecular species. Secondly this range of available energies and the different effective trap depth in the RF trap means that the probability of a generated

ion falling into the trap can be rather different. Indeed we find that parameters that give us a single ion in the RF trap actually load many ions in the Penning trap.

In order to characterise the loading parameters we used a number of different oven currents between 1.20 and 1.50 ampere and varied the filament currents between 4.05 and 4.40 ampere. For each load the oven was first left to run for 2 minutes. After this the filament was also turned on and 20 seconds later both were turned off. A filament bias of -12 volt was used and a current of 1.0 ampere was passed through the coils of the electromagnet during loading, such that there is a small field of about 0.1 tesla to guide electrons into the trap.

The largest fluorescence rate we recorded with the currents used was just below 4×10^5 counts per second. In previous experiments using larger currents we have observed maximum fluorescence rates of about 6×10^5 counts per second. As we reduce the filament current we generally find that at some point the signal observed for each load begins to fall, with the value at which this begins to occur depending on the oven current. Plots of overall fluorescence rate as a function of filament current for a number of different oven currents are displayed in figure 7.2.

Once the filament currents are sufficiently small we begin to observe the discrete quantum jumps expected and we can estimate the number of ions either statistically or by eye. Examples of fluorescence traces for 1, 2 and 3 ions can be found in figure 7.3. It is possible to reduce the number of ions loaded in a well controlled manner using the filament current, particularly when the oven current is such that the rate of change of ion number with filament current is not too large. Furthermore it is straightforward to continue this process until single ions are obtained.

One interesting observation is that for oven currents of 1.50 ampere or greater we can load small numbers of ions without using the filament. It is assumed that some small number of electrons is emitted by the tantalum wires that heat the oven. Because of the large trapping voltages these electrons can be swept into the centre of the trap and so if we leave the oven on for long enough ions will eventually be produced. A useful trick for trapping single ions has been to keep running the oven at 1.50 ampere until fluorescence is observed, usually slightly under two minutes. It should be remembered that this trick would not work in the Penning trap since there is only a small bias




Figure 7.2: Plots showing the drop in observed fluorescence rate in the RF trap as the filament current is reduced. The different plots correspond to different oven currents as indicated.

of the oven relative to the centre of the trap. An alternative explanation for this effect could be that some small number of ions is emitted directly as a result of heating of the calcium in the oven. Being produced so far from the centre of the trap it would seem unlikely, however, that ions produced in such a way would reach the trapping region and remain confined.

The largest signal rate we observed for a single ion in the RF trap is about 8×10^3 counts per second. This is close to the value we calculated in section 5.8 to be the expected fluorescence from a single ion in a Penning trap. Since there is never any need to offset the lasers from the centre of an RF trap we might have expected slightly more signal. Several of the variables in our calculation were estimated or taken from manufacturer's specifications and we ignored imperfections in the imaging system. This makes our expected fluorescence rate something of an upper limit and it is perhaps not surprising that we observed somewhat smaller signal levels in the experiment. Even if we reduce our estimate of the expected signal in the Penning trap on the basis of the results here we would expect to measure single ion fluorescence rates of at least a few thousand counts per second. This is well above the background

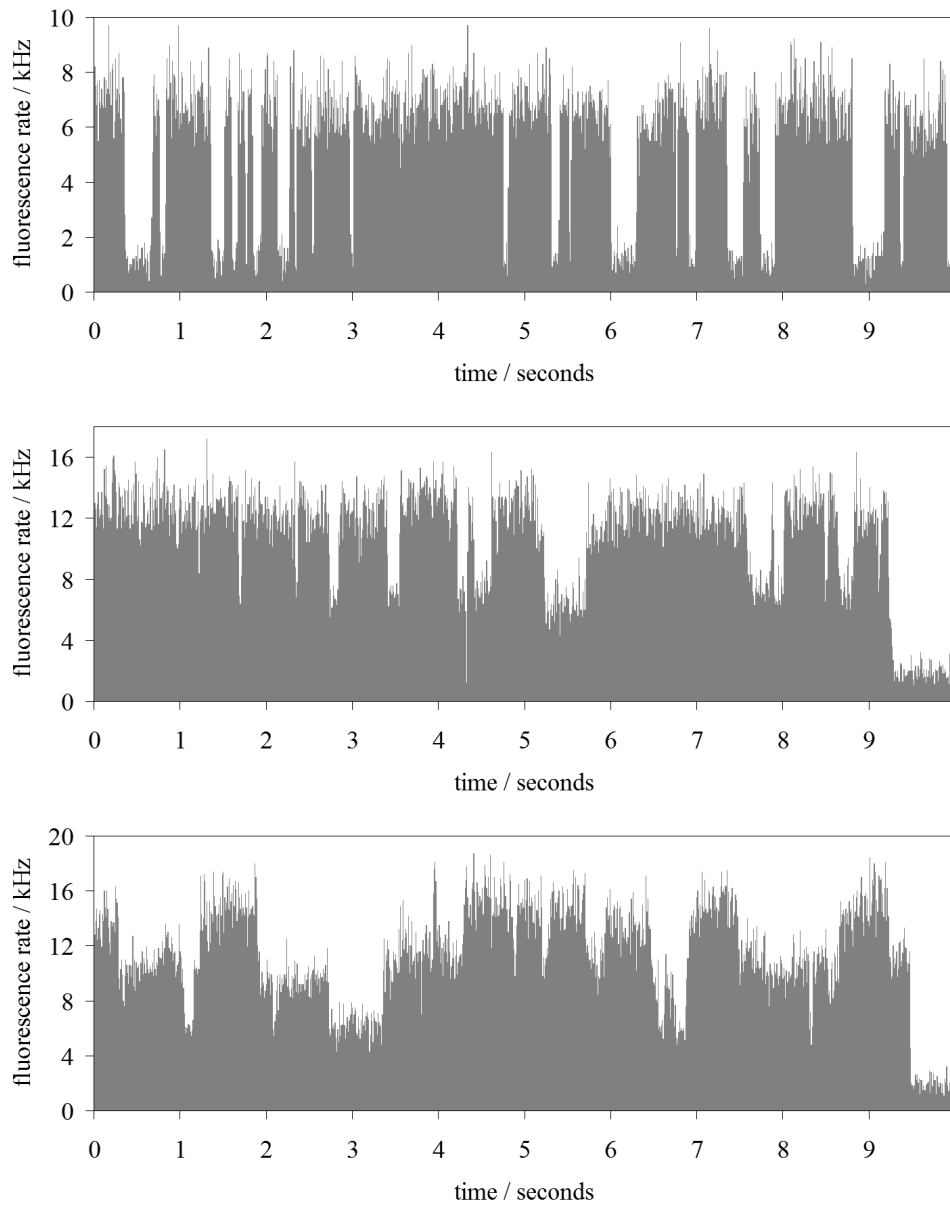


Figure 7.3: Example traces showing fluorescence in the RF trap from 1, 2 and 3 ions respectively. The ion number has been determined from the quantum jumps that are clearly visible in each trace. The end sections of the middle and lower trace show the background noise level.




Figure 7.4: iCCD images of a single ion in a RF trap. The image on the left is a 20 second exposure. The centre and right-hand images are successive frames taken from a sequence of 100 ms exposures of the same ion. Quantum jumps are clearly observable. In each image the pixel size is approximately $7.5\mu\text{m}$.

levels we can achieve of around a few hundred counts per second, as can be seen in figure 7.3.

It is interesting to note that the signal rate per ion in the RF trap remains similar when we increase the number of ions from one to three. The signal rates per ion for the three traces in figure 7.3 are 5.2, 5.4 and 4.3×10^3 counts per second for 1, 2 and 3 ions respectively. Unfortunately it does not necessarily follow that the same will be true in the Penning trap. Since laser cooling in the Penning trap can be strongly dependent on the position of ions in the beam, space charge effects due to having more than one ion present in the trap can greatly reduce the cooling efficiency. It is possible that this situation might be improved somewhat with axialisation.

In order to show that the iCCD imaging system is capable of observing single ions and the changes in their fluorescence due to quantum jumps we have taken several sequences of images of single ions in the RF trap. Examples of such images are shown in figure 7.4. The first of these has an exposure time of 20 seconds. The spatial amplitude of the ion's motion is measured to be about $90\mu\text{m} \times 35\mu\text{m}$. That this is so large is due to micromotion caused by the oscillating trapping potential. Because of additional stray potentials near the trap, the ion will not sit in the exact node of the quadrupole potential and hence there is a component of its motion that oscillates with the RF drive. In principle this micromotion can be eliminated by applying potentials to compensation electrodes in order to shift the ion to the exact centre of the quadrupole potential.

To resolve quantum jumps in an ion's fluorescence we must use an image acquisition time somewhat smaller than the average period for which the ion will remain in the 'off' state. This necessarily means that the signal to noise ratio will be somewhat limited. The second and third images in figure 7.4 are consecutive 100ms exposures of the same ion. Both images were corrected in real time for the background light level and the dark count of the ICCD. The fact that the ion is in the 'off' state for the centre image, and the 'on' state for the right-hand image, is clearly apparent. We can therefore image single ions in near real time and directly observe the effects of quantum jumps in their electronic states.

7.3 Small clouds of ions in the Penning trap

The method of controlling the cloud size by changing the filament current that we used in the RF trap can be used in exactly the same way in the Penning trap. Because of the different effective trap depths we would expect the number of ions trapped to be scaled somewhat relative to the RF trap, so that the loading parameters for a single ion would be different. Also, because the applied trapping potential is now constant, we would expect the ratios of calcium to various contaminant ions to depend more critically on the bias of the filament relative to the centre of the trap.

Typically we would load ions with an oven current of 1.40–1.50 ampere and a filament bias in the region of 6.5–10.0 volt relative to the potential at the centre of the trap. As with loading in the RF trap we apply the oven current for 2 minutes before switching on the filament current. 20 seconds later we switch off both currents. Generally we start with a filament current of about 4.60 ampere, which we would reduce in steps of 0.05–0.10 between successive loads. When we are laser cooling smaller clouds of ions we find that we are usually much more sensitive to the laser beam frequencies and positions. Reducing the cloud size in steps is therefore important because it enables us to gradually home in on the optimum beam parameters.

During loading we would usually leave one of the Doppler cooling lasers scanning repeatedly towards resonance over a range of a few GHz. Once we have loaded a cloud we typically have to wait a short while for the ions to cool in to the centre of the trap (see, for example, figure 5.18 in section 5.9). How long this takes depends on the size of the cloud and the laser cooling

parameters. When there are many ions in the trap, interactions between ions are significant and providing we are laser cooling at least some of them efficiently we can rapidly cool them all to the centre of the trap. For smaller clouds we can consider each ion's motion to be independent and it becomes possible for an ion to remain in a large orbit for periods of a minute or more before it is cooled in to the centre of the trap. The duration of this cooling-in period can be reduced by increasing the offset in the position of the laser beams during loading and then gradually moving them back towards the centre of the trap until signal is maximised. This has the effect of sweeping the ions into the centre of the trap where they can be more efficiently cooled.

Once the trap is loaded and the signal level has become constant we can optimise the laser beam positions and frequencies. Unfortunately these parameters are all related and obtaining maximum signal involves several iterations between them. We generally find that we are more sensitive to the laser parameters when we are working with smaller clouds. It is therefore essential that we reduce the filament current in a stepwise fashion so that the cloud size is reduced in stages. If we try to load too small a cloud of ions we find that there is little or no initial fluorescence and optimising the laser parameters is not possible. It can sometimes be difficult to decide whether ions have simply yet to cool into the trap or whether we have loaded too small a cloud and must reload with a larger filament current.

Figure 7.5 shows typical fluorescence traces from clouds of ions that have been loaded with large and small filament currents. Each spectrum is an average of 20 scans of one of the Doppler cooling lasers and in each case the laser parameters have been optimised to give the largest peak fluorescence. The oven current used was 1.50 ampere. In figure 7.6 the peak fluorescence rates from these and other consecutively loaded clouds has been plotted. The beam power was constant throughout, except for the two points indicated, where it was reduced to about 40% and 25% respectively. There is a clear reduction in signal amplitude as we reduce the filament current below 4.40 ampere, due to the fact that we have fewer ions in the trap to interact with the laser. It is interesting that the peak signal amplitude actually falls slightly if we use filament currents larger than 4.40 ampere. If we have a large number of ions then the physical size of the cloud is large and only a small fraction of the ions interacts with the laser at any one time. Even if we add more ions the number that are actually within the beam only increases very slightly, though

they are likely to be less cold.

When we look at the widths of the fluorescence peaks in these clouds we find that they remain fairly constant until we reach 4.10 ampere, at which point they fall rapidly. This is an indication that small clouds can be cooled more efficiently and a lower Doppler width can be achieved. The observed lineshapes differ significantly from the Voigt profile that one might expect. This difference is due to the fact that the Doppler cooling rate varies with the laser frequency, so the temperature of the ions is not constant. The Doppler width, and the profile we are probing, therefore change throughout the scan. This dynamic effect is also responsible for the small lump in the spectra that sometimes occurs shortly before the main peak. These same features have been generated in numerical models that incorporate a variable ion temperature [100].

Despite this dynamic effect it is still possible to make a crude estimate of the temperature of the ions by fitting Doppler profiles to the fluorescence spectra. This is reasonable providing the Doppler width is much larger than the natural width and we ignore the wings of the spectra. Even so, such a fit gives us only an upper limit to the ion temperature. The estimated temperatures obtained for the spectra discussed above are plotted as a function of filament current in figure 7.7. We find that the observed temperatures are consistently in the region of 35 kelvin when the filament current is greater than 4.10 ampere. Three separate measurements were taken for a cloud loaded at 4.60 ampere, which vary between 28 and 39 kelvin — giving us an idea of the precision with which we can determine the ion temperature. At 4.10 and 4.05 ampere the temperature drops rapidly to 22 and 11 kelvin respectively, indicating that a much smaller cloud is present.

In order to maximise the laser cooling rate the laser frequency must be detuned by half the linewidth of the transition. This detuning also sets the limit of Doppler cooling. When large numbers of ions are in the trap it is observed that ion temperatures are significantly higher than when fewer ions are present. As a result the Doppler width will tend to be large and the laser must be detuned by a correspondingly large amount. Oversaturation of the cooling transition broadens the natural width of the transition, but since the Doppler width is dominant the effect on cooling is small. Instead the extra power means that we can effectively cool more ions and so there is a somewhat improved signal rate.

When the number of ions in the trap is small the natural width forms a

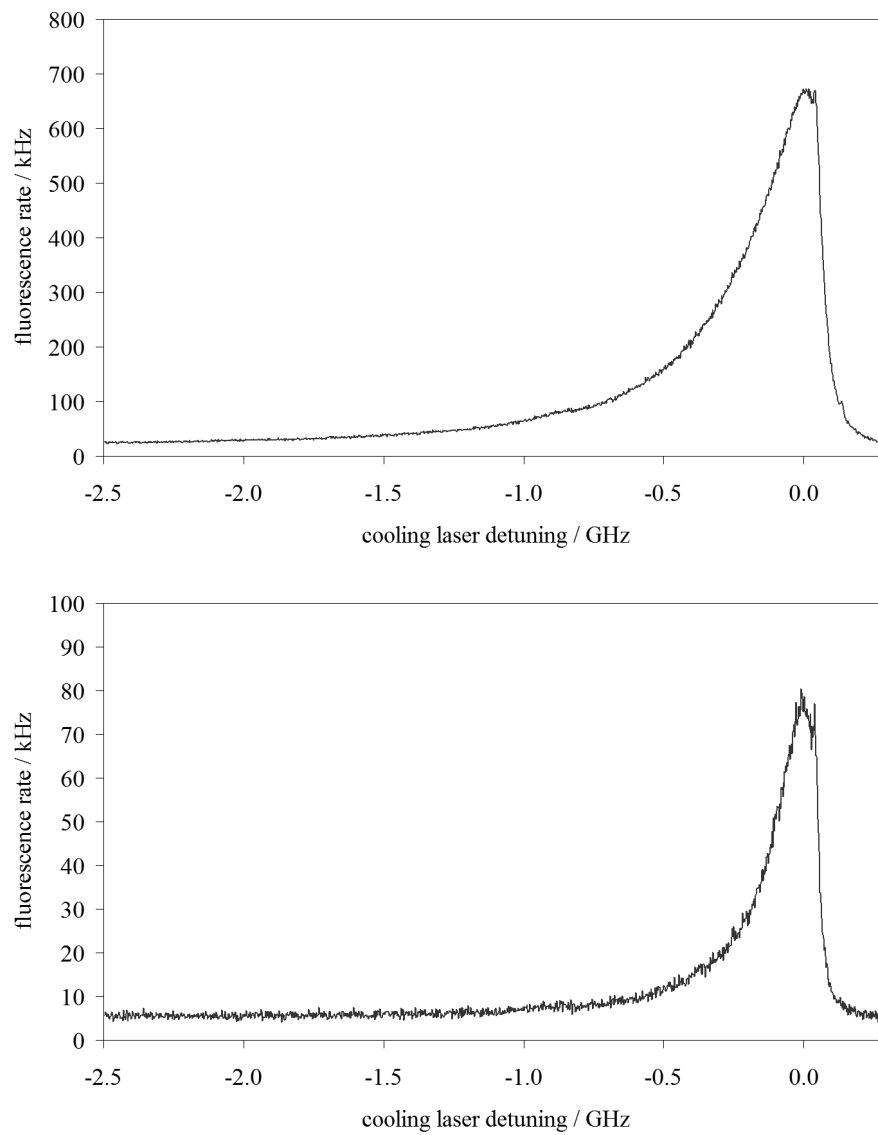


Figure 7.5: Fluorescence traces from ion clouds loaded with filament currents of 4.70 ampere (top) and 4.10 ampere (bottom).




Figure 7.6: Peak fluorescence rates from ion clouds loaded with successively smaller filament currents. Note the maximum in peak fluorescence rate for a current of 4.40 ampere. The percentages next to the two points nearest the origin indicate the relative beam power that has been used. Although the measurement error is very small, there is significant scatter in the repeated measurements at 4.60 ampere due to changes in the laser cooling parameters.




Figure 7.7: Plot of the temperatures of ion clouds as a function of the filament current used during loading. Temperatures have been estimated by fitting Doppler profiles to spectra such as those in figure 7.5. Some indication of the accuracy of this approach is given by the scatter in the repeated measurements at 4.60 ampere. The percentages next to the two points nearest the origin indicate the relative beam power that has been used.

significant part of the overall lineshape. Excess power can then cause significant broadening of the line and hence the temperature we can obtain with the laser cooling is increased. Filtering of the beam can therefore reduce the extent of saturation broadening and also enables us to bring the laser beam nearer to the centre of the trap without causing heating of the magnetron motion. Figure 7.8 shows two spectra of the same (fairly small) ion cloud, where the laser beam power for the second is $\sim 40\%$ that of the first. The peak signal rate is only reduced by about 10%, so the signal to noise ratio is actually improved. The temperature as measured by fitting a Doppler profile is reduced from 22 kelvin to 6 kelvin.

In general we find that as the number of ions loaded is reduced the optimum beam position for maximum signal is seen to move closer to the centre of the trap. When we talk about ion temperature in this regime we are really measuring ion velocities. In the Penning trap there is a well defined relationship between an ion's velocity and its orbital radius. If we assume that the magnetron radius is much larger than the modified cyclotron radius then $v = r\omega_m$. Clearly if the limit to the ion temperature is reduced when there are fewer ions then the beam must be closer to the centre of the trap, since the radius r is smaller.

We have tried to load still smaller clouds of ions whilst making use of the axialisation technique described in chapter 6. Typically we would keep reducing the filament current until the fluorescence rate observed is small (perhaps 1×10^4 counts per second). Based on experience from the RF trap we might expect this level of fluorescence to correspond to just a few ions. We then set the Doppler cooling lasers so that they are both static a little below resonance and look to observe quantum jumps. Although the fluorescence might show evidence of non-Poissonian statistics, we have not seen the obvious discrete steps we have seen with just a few ions in the RF trap. The implication is that the signal rate is much lower than we had expected and we have more ions than we can discern with the quantum jumps technique. If we proceed further with reducing the filament current we can continue loading smaller clouds and the fluorescence rate falls accordingly. Eventually we reach a stage where the fluorescence rate is similar to the background level, even if we filter the 397nm lasers more heavily. With such a poor signal to noise ratio it becomes difficult to optimise the laser parameters and eventually the signal can be lost altogether.

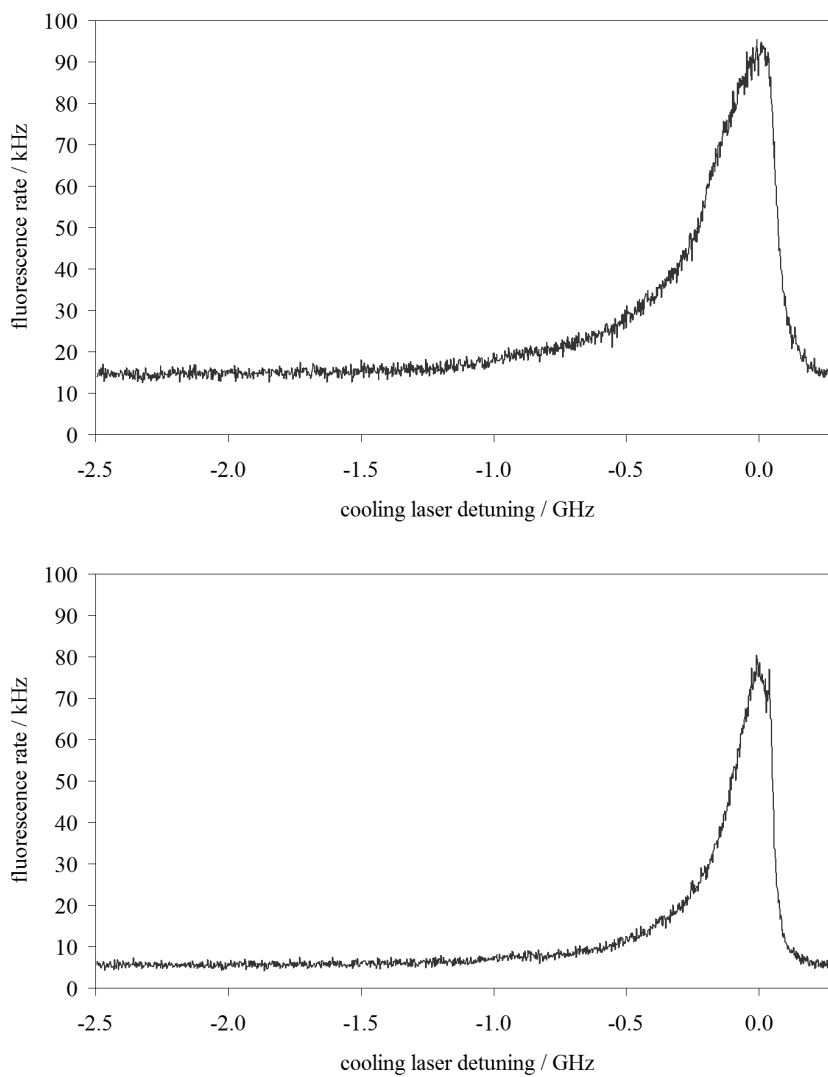


Figure 7.8: Spectra taken of the same cloud of ions before and after filtering of the beam to about 40% of its original power. Filtering the beam is seen to reduce the width of the line, indicating a temperature reduction from 22 kelvin to 6 kelvin.




Figure 7.9: Plot showing the peak fluorescence rate (solid line) from clouds loaded with increasingly small filament currents in a Penning trap with axialisation. The ratio of this signal to the background noise is also plotted (dashed line).

Figure 7.9 shows the observed fluorescence rates from a series of ion clouds loaded with an oven current of 1.40 ampere and progressively smaller filament currents. The clouds are axialised by a 150 mV drive. The peak fluorescence rates are fairly similar to those recorded without axialisation. The ratio of the signal level to the background light level is also plotted. At various points the cooling laser is attenuated, leading to an improvement in the signal to noise ratio. The lowest filament current at which fluorescence was observed was 4.02 ampere, where the peak signal rate was about 2000 counts per second. A quantum jumps trace obtained from this cloud is shown in figure 7.10. The fluorescence rate per ion is sufficiently small compared to the background level of ~ 300 counts per second that it is hard to discern discrete steps by eye. A statistical analysis (as described in section 7.1) suggests that there are 4 ions present in the trap, which leads to a maximum signal rate per ion of about 500 counts per second. We have obtained quantum jumps traces from other ion clouds that indicate the presence of ~ 3 –5 ions, and these indicate similar signal rates per ion.

We have on other occasions also loaded clouds that give even lower overall




Figure 7.10: Quantum jumps trace showing ~ 4 ions in a Penning trap with axialisation. Region (a) is the full signal without quantum jumps, (b) is the signal with quantum jumps and (c) is the background.

fluorescence rates. It is difficult to know whether these are actually smaller clouds or whether the laser parameters are simply not as well optimised. We can only distinguish between these two cases by looking at quantum jumps traces, but if the signal rate is similar to the background it becomes difficult to see clear quantum jumps. Furthermore such small signal is very easily lost, since it is hard to see the effect of changing the cooling parameters.

7.4 Contaminant ions

Results presented in the previous section showed that having more calcium ions in the trap has a detrimental effect on our ability to efficiently cool the ion cloud. This problem would be even worse if we had a large number of ions of some other species which are not being laser cooled. In this situation the Coulomb interactions between ions can lead to the redistribution of ion temperatures until equilibrium is reached. Effectively this means that one species of ion can be sympathetically cooled through interactions with another ion species that is being laser cooled. This process is often useful, since we can use it to cool ions that are not suited to laser cooling, and has been

demonstrated many times (see for example [48]).

Work performed by the Imperial College group on sympathetic cooling of molecules in the Penning trap has shown that when the charge-to-mass ratios of the two ion species are similar there can be a significant increase in the measured temperature of the laser cooled ions [68]. During this work the population of laser-cooled ions was believed to be always greater than that of the sympathetically cooled species. In our experiment here we might imagine a situation where our technique for reducing the number of calcium ions does not lead to an equivalent decrease in the number of contaminant ions. We would then find that the problem of contaminant ions becomes more significant when we try to reduce the number of calcium ions. In order for an ion species to be found in the trap it must have an ionisation potential less than or equal to the bias of the filament relative to the centre of the trap. Fortunately calcium has a fairly low ionisation potential of 6.1 volt, so we can load it without using a very large filament bias. This means that a large number of possible contaminants can immediately be eliminated.

It should be noted that one contaminant ion that will certainly be present in the trap when we load a large cloud of ions is $^{44}\text{Ca}^+$. This isotope has a natural abundance of $\sim 2\%$, about 50 times less than that of ^{40}Ca . Since we have no isotope selectivity in our loading technique, it might seem that the presence of $^{44}\text{Ca}^+$ will always be a problem. In fact the probability of a given calcium ion being $^{44}\text{Ca}^+$ is the same whatever the number of ions in the cloud, so if we load just one or two ions we will most often only have $^{40}\text{Ca}^+$.

That we have no difficulty in trapping and cooling single ions in the RF trap does not contradict this explanation for the difficulties we have in the Penning trap. It should be remembered that the RF voltage and frequency required for stable trapping depend on ion mass. It is therefore quite possible that the trap can be stable for calcium but not for many of the possible contaminant ions. It has on occasion been found that scanning the RF voltage around the stable region for Ca^+ can increase the signal level and this can be attributed to cleaning of the trap.

Selectively removing an ion species from the Penning trap is not as straightforward as it is in the RF trap. The stability condition in the Penning trap (equation 4.5) is that the confining potential U must be between zero and

U_{max} , where

$$U_{max} = \frac{eB^2 (2z_0^2 + r_0^2)}{8m}.$$

Clearly by changing U we can only remove ions heavier than calcium without also removing the calcium. The only way of driving out lighter ions would be to apply an external RF drive tuned to a motional frequency of the selected ion. This can rapidly heat the selected ion species from the trap, but requires that we know which species are present to begin with. Besides calcium-44, some contaminants likely to be present in the trap are those formed from ionisation of (and reactions between) residual gas molecules in the trap. In previous experiments HCO^+ , N_2H^+ and NH^+ have all been observed at air pressures of $\sim 10^{-8}$ mbar [68]. Since these are all lighter than calcium we cannot remove them simply by changing the trap voltage.

In order to determine whether contaminant ions are present in our trap we performed mass spectrometry on clouds of ions loaded using a variety of filament currents and biases. When there are multiple species of ion in one trap some of the resonant motional frequencies can shift slightly due to interaction between the different types of ion [68]. The ‘breathing mode’ frequency, given by the difference between the modified cyclotron and magnetron frequencies, is not seen to shift and is therefore useful for mass spectrometry. In order to excite this motion we need a 3D quadrupole and hence we apply a radio-frequency drive to the endcaps of the trap whilst keeping the ring electrodes at a fixed potential (see figure 7.11). In practice, trap imperfections mean that this arrangement will also excite ion motions other than the breathing mode.

Once a cloud has been loaded and the laser cooling parameters optimised we make both Doppler cooling lasers static at a frequency slightly detuned from the position of maximum fluorescence. We then scan the RF drive over the frequency range we are interested in whilst monitoring the fluorescence rate. When the drive is resonant with some motional frequency of an ion present in the trap, the cloud will be heated and will expand accordingly. This changes the degree of interaction with the laser and a change in fluorescence is observed. Whether this change is an increase or a decrease presumably depends on the position of the laser beams in the trap.

Since we are worried primarily about the idea that the ratio of contaminant ions to $^{40}\text{Ca}^+$ might become larger as we load smaller clouds, we would ideally perform this experiment with the smallest cloud sizes we can achieve.

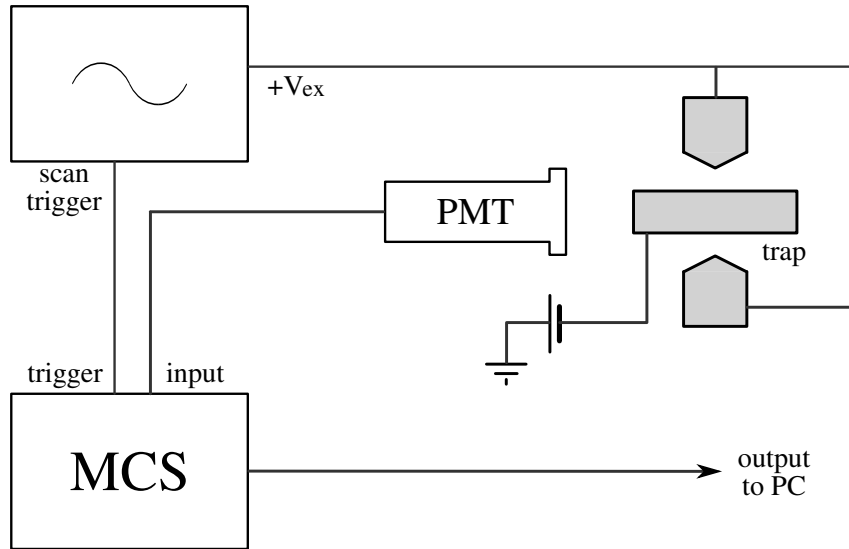


Figure 7.11: Setup for driving the breathing mode in order to perform mass spectrometry of an ion cloud. The dc trap bias is applied to the ring electrodes, with the RF applied about ground to the endcaps.

Unfortunately, as with the observation of quantum jumps, it is difficult to obtain clear results when the fluorescence rate is very low. Furthermore we find that if the size of the cloud is small we tend to lose signal very easily when we excite a calcium-40 motional resonance. If we have a slightly larger cloud we can excite a resonance and see only a temporary dip or peak in the fluorescence rate. With small clouds the signal loss is often total and the fluorescence rate may never recover. It would seem that when several ions are being laser cooled the interactions between them act to rapidly recool an ion that has been excited into some large orbit by the excitation drive. With just a few ions present one or all of them can easily be heated into a large orbit where they do not interact significantly with the laser beams. For these reasons it was for the most part necessary to use clouds large enough to produce fluorescence rates of 1.5×10^5 counts per second or more. Though these clouds might seem very large compared to those used for the measurement of cooling rates in chapter 6, it should be borne in mind that the amplitudes of the excitation drives being used are generally much larger.

In order to check that our arrangement works we first performed mass

spectrometry of a cloud of ions loaded using a large filament current and a filament bias of -30 volt. Using such a large filament bias means that many different atomic and molecular species can be ionised. We would especially expect to observe Ca^{2+} . Selected spectra from those we obtained are shown in figure 7.12. In spectrum (a) the breathing and modified cyclotron mode of $^{40}\text{Ca}^+$ are clearly visible. Also present is what looks to be the modified cyclotron mode of $^{44}\text{Ca}^+$.

Spectrum (b) is a lower frequency scan in which we would expect to see resonances from ions heavier than calcium. The very sharp and strong resonance is the second harmonic of the axial motion of $^{40}\text{Ca}^+$. There is also a resonance which appears to correspond to an ion with mass 56 amu. This is corroborated by scans taken with a slightly different trap bias, in which the resonance is shifted accordingly. One possible candidate for this is calcium oxide (CaO). When exposed to air calcium will oxidise to form CaO — seen as a white layer on the surface of the metal. Although efforts were made to fill the ovens in the trap with clean, unoxidised calcium there will inevitably have been some calcium oxide present. Although we would expect the amount emitted from the oven to be much less than the amount of calcium, CaO has an ionisation energy of 6.7 eV and so ions might easily be produced. It might seem problematic that CaO has an ionisation energy so close to that of calcium itself, since it means we cannot easily prevent loading it. However we must drive this resonance significantly harder than we must drive the calcium resonances, suggesting that the number of ions present is much lower. Since we would expect the rate of loading of CaO to be reduced as we turn down the oven current it is unlikely to cause a problem at the few ion level.

Spectrum (c) is a higher frequency scan, and hence covers higher charge-to-mass ratio ions. The leftmost pair of features, at about 460 and 495 kHz, probably correspond to ions formed from various residual gases. The features are rather broad, making it hard to be certain, but frequencies corresponding to the breathing mode of ions with mass 28 amu (N_2^+) and 30 amu (NO^+) are close to the observed peaks. The ionisation energies of these species are 15.6 eV and 9.3 eV respectively, making them easy to prevent from being loaded. An alternative possibility for what appears to be a mass 28 amu ion is doubly-charged CaO . The highest frequency features, at about 1200 kHz, correspond to either a singly charged ion of mass 12 amu or a doubly charged ion of mass 24 amu. The latter case seems unlikely given that we

see no resonances corresponding to a singly charged ion with a mass of 24 amu. Ionisation of Carbon is possible since it has an ionisation energy of 11.3 eV. Although it is not completely clear why there should be Carbon atoms in the trapping region, one possible source is the Kapton coated wires that are used to make the electrical connections to the trap components. The wires connected to the filaments carry large currents and hence become rather hot during use. It is observed that over time the Kapton on these wires becomes slightly blackened. Interestingly, Kapton coatings have become something of a problem in the aeronautics industry, where they have in the past been very commonly used, because they have been found to form a (conducting) layer of Carbon when they are overheated. It is feasible that atomic Carbon might be emitted from such a layer when our filaments are switched on and the Kapton coated wires become hot.

The middle two resonances in spectrum (c) are the breathing and modified cyclotron modes of the expected $^{40}\text{Ca}^{2+}$ ions.

Figure 7.13 shows a series of scans covering the $^{40}\text{Ca}^{2+}$ resonances, which were taken in quick succession. The features are observed to get smaller in amplitude until they are no longer visible. Even if we then drive the resonances harder the features are still not seen, suggesting that we have heated these ions from the trap. This ability to expel ions of a given species would be useful were we to find that ions other than calcium were still being loaded even with smaller filament biases.

Having demonstrated that the mass spectrometry technique works, we then looked for contaminant ions in a smaller cloud of ions loaded using a more typical filament bias of -6.5 volt relative to trap centre. A trap bias of 4 volt was used. Since different motions of the ions have different sensitivities to the drive it was necessary to perform several frequency scans, using different amplitudes in each region. To average out the effects of laser frequency noise we performed each scan five times at each amplitude. Wherever possible (ie. away from the motional frequencies of calcium) we performed scans with RF amplitudes as high as 800mV peak-to-peak. In case contaminant ions were being driven out we performed scans at lower amplitudes first and gradually moved up in amplitude.

In some cases different motional frequencies of different ions can occur very near each other. For example, with our experimental parameters and a trap bias of 4 volt the breathing mode of $^{40}\text{Ca}^+$ occurs at 329kHz, whilst

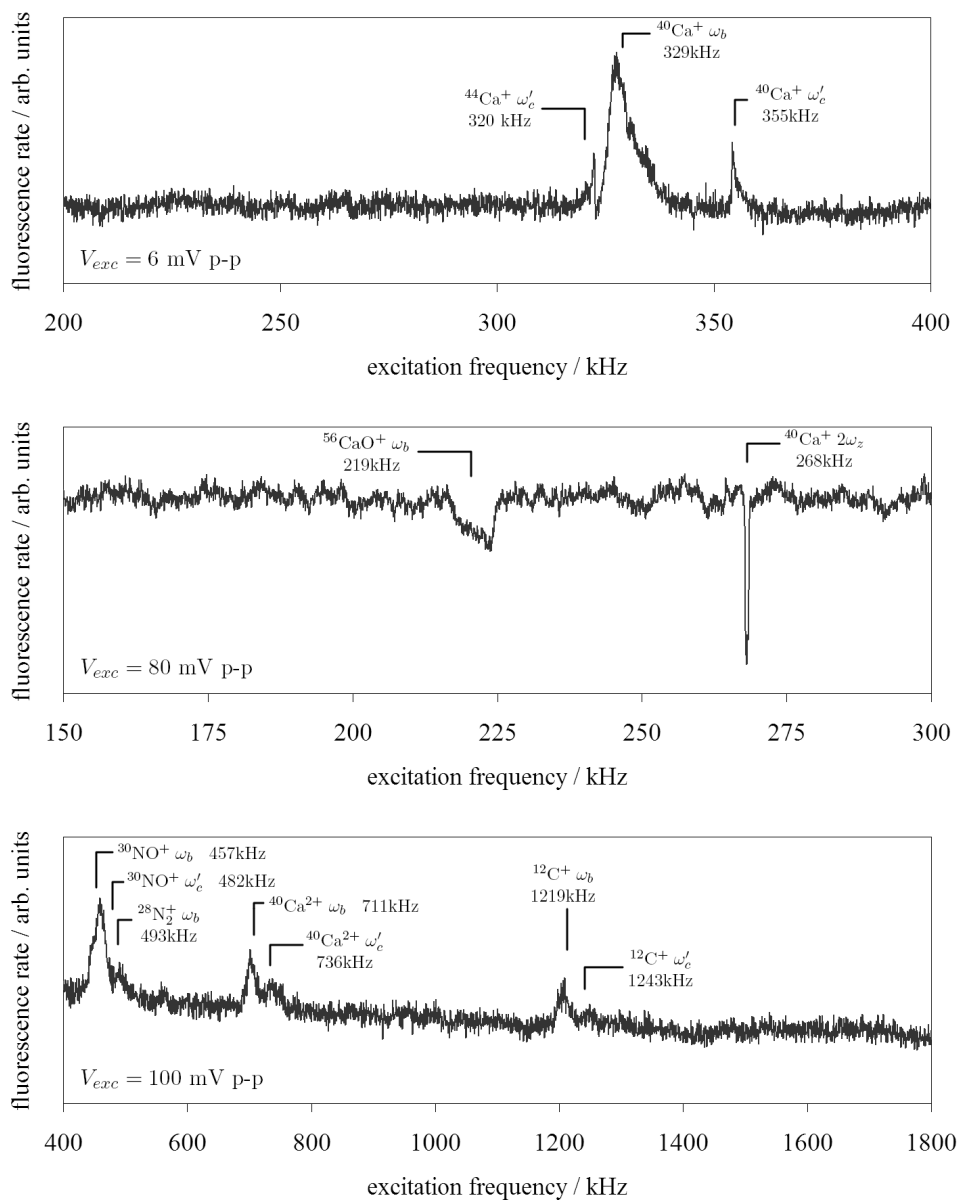


Figure 7.12: Mass spectra for ions loaded with -30 volt filament bias. *Top*: features of Ca^+ . *Middle*: lower charge/mass ions. *Bottom*: higher charge/mass ions.




Figure 7.13: Repeatedly scanning across the breathing and modified cyclotron resonances of $^{40}\text{Ca}^{2+}$ we find that the features gradually disappear. This is attributed to driving of these ions out of the trap.

the modified cyclotron mode of $^{44}\text{Ca}^+$ is at 328kHz. In such instances we can differentiate between the two species by performing another scan with a slightly different trap bias. At 3 volt the $^{40}\text{Ca}^+$ breathing mode lies at 343kHz but the modified cyclotron mode of $^{44}\text{Ca}^+$ has moved only to 334kHz.

Figure 7.14 shows the mass spectrum obtained by piecing together the various scans we performed. The amplitudes of each portion are indicated, and in most cases it is the largest amplitude used that has been shown. The exceptions to this are very near some of the detected resonances, where larger amplitude scans were performed but showed only broader features. All the features observed can be attributed to modes of motion of calcium-40, or multiples and fractions thereof, except the broad feature centred at about 310 kHz. This lies between the breathing and modified cyclotron modes of calcium-44.

As discussed earlier we are forced to use a larger cloud than we would ideally like, and as a result the ratio of contaminants to calcium-40 might be lower than it is for very small clouds. Because of this we cannot completely rule out the idea that contaminants are responsible for our poor signal rate per ion. There is, however, no reason for us to expect the absolute number of

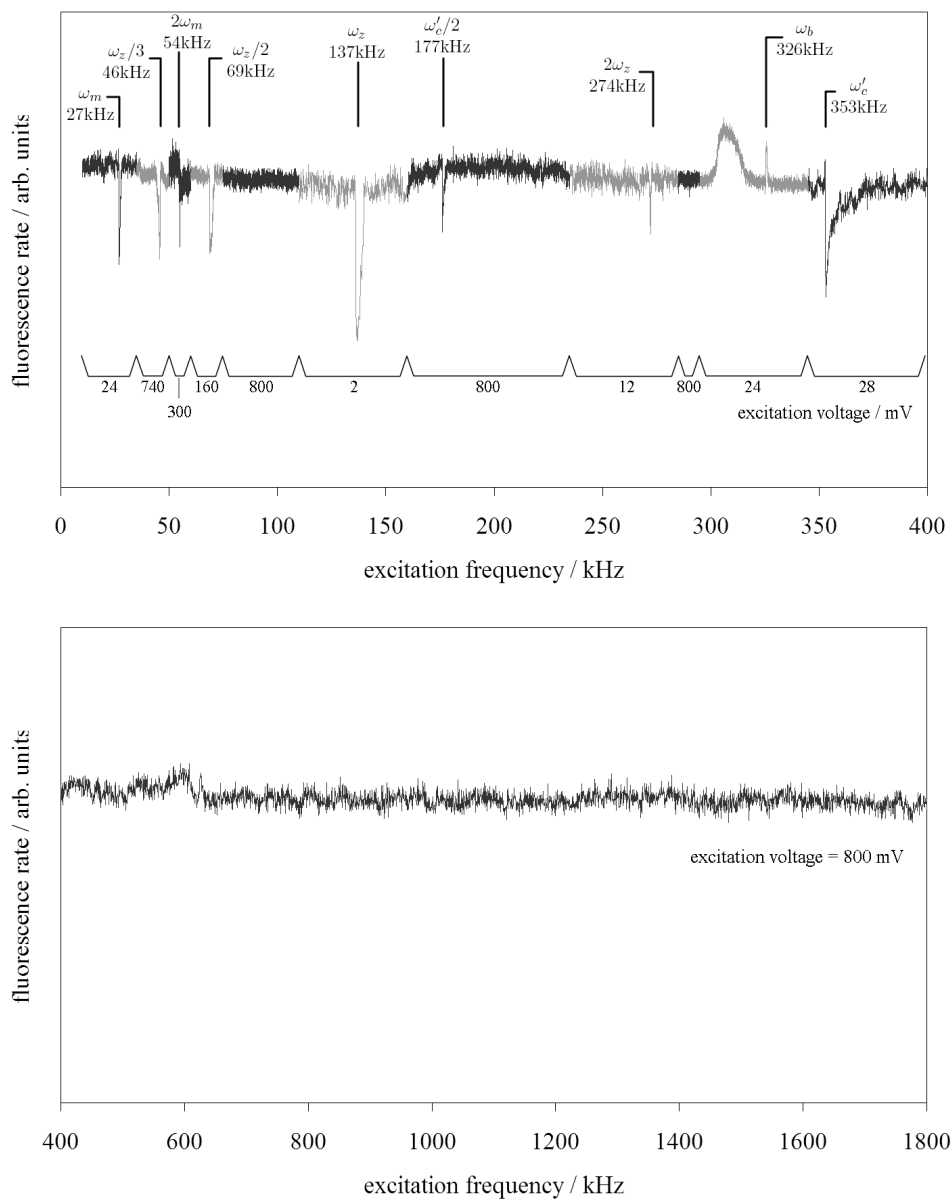


Figure 7.14: Complete mass spectrum for a small cloud of ions loaded with a filament bias of -6.5 volt. Scans with different drive amplitudes have been pieced together so that resonances with different sensitivities can be seen together. The motions believed responsible for each feature are indicated. The broad feature at around 310kHz lies in between the breathing and modified cyclotron modes of $^{44}\text{Ca}^+$.

contaminant ions to increase as we reduce the filament current. The fact that we see no evidence of them at this scale sets an upper limit on the possible extent of the problem at the few ion level, and it would seem unlikely that this could be the sole cause of our poor signal rates.

One technique that could be used to eliminate altogether the possibility of contaminant ions is the use of photoionisation instead of ionisation by electron bombardment as our means of loading the trap. This technique uses a narrow linewidth laser to resonantly drive the required atom into an excited state, from where it can be ionised by a second photon at either the same or a different frequency [104]. The resonant excitation provides the species selectivity, so the second excitation can even be due to a broadband source [105]. Depending on the ratio of the isotope shifts to the transition linewidths, it is even possible to select out specific isotopes with high probabilities [106].

7.5 Coherent population trapping

An interesting alternative explanation for the fact that we observe somewhat lower fluorescence rates than we expect is coherent trapping of the internal state of our ions in some non-fluorescing ‘dark’ state. If this were occurring then even though we might be able to trap single ions we would be unable to cool and detect them efficiently. This would not affect us to the same degree when we are working with large clouds because ions move in and out of the laser beams and so coherent evolution of dark states is inhibited.

For our purposes, a dark state is any internal atomic state that does not couple to our $S_{1/2} \rightarrow P_{1/2}$ laser cooling transition. At the most trivial level, all electronic states other than the $S_{1/2}$ and $P_{1/2}$ levels are dark. If we switch off our 866 nm repumper we drive population into the $D_{3/2}$ and see no fluorescence. We tend to refer to situations such as this, when the ion is simply driven into a non-addressed energy level, as optical pumping. It is, however, also possible to drive the internal state into a superposition of two or more energy levels that does not fluoresce — even if these energy levels would individually couple to the cooling laser. It is in this instance that we refer to coherent population trapping.

The classic example of this effect is in three level systems that are addressed by two lasers (see the so called ‘ Λ ’ system in figure 7.15). When the two lasers are equally detuned from the upper state $|0\rangle$ they are in ‘two-photon resonance’




Figure 7.15: Energy level diagram of a Λ system. Raman transitions between states $|1\rangle$ and $|2\rangle$ are possible without populating state $|0\rangle$ if the two lasers are equally detuned from the upper state.

and a non-fluorescing superposition of states $|1\rangle$ and $|2\rangle$ can be established. Population can also be transferred between these two states without the $|0\rangle$ state ever being occupied. These ‘Raman’ processes are often useful, such as when one wants to drive transitions between Zeeman or hyperfine ground state qubits without having to occupy an intermediate state that is vulnerable to spontaneous decay.

Naïvely one would not expect to see such processes in our system, because ideally we want to have our Doppler cooling lasers detuned from resonance and our repumper lasers near resonance. However, since we actually have a rather large number of energy levels and transitions involved in our cooling cycle it is not inconceivable that some unexpected trapped state might occur. Thus the situation warrants a detailed analysis.

Much theoretical work has been performed on the characterisation of coherent population trapping, though most has been geared towards situations where there is little or no magnetic field present. The much larger Zeeman splitting makes the situation in the Penning trap rather different, and so we cannot necessarily carry across the results from these low-field models directly.




Figure 7.16: Partial energy level diagram for calcium-40 showing the near-degenerate Zeeman sublevels at very low magnetic field strengths. All of the possible 397 nm and 866 nm transitions are indicated. Dashed lines are π transitions, solid lines are σ transitions.

In order to look at the problems that can arise, however, it is interesting to discuss observations of coherent population trapping at low fields.

There is a rather special case of coherent population trapping that occurs when the magnetic field strength present at the position of an ion is very close to zero. When this occurs the Zeeman components of the energy levels are near degenerate and a single laser can be used to carry out what are effectively Raman transitions between these levels. As an example, the relevant levels of calcium-40 are shown in figure 7.16 along with all the possible transitions that can be addressed by single 397 nm and 866 nm lasers. It is not difficult to see that dark states in the $S_{1/2} \rightarrow P_{1/2}$ system are easily avoided. Providing we address only the π or σ components there is no way that a single laser can form the kind of Λ -type system shown in figure 7.15. The only dark states that can possibly be a problem occur if we use a mixture of π and σ polarisations, or through optical pumping into one of the ground states if we make the mistake of using pure σ_+ or σ_- polarisation (see figure 7.17).

The $D_{3/2} \rightarrow P_{1/2}$ system, however, is rather different. Here we must use both σ_+ and σ_- polarised light if we are to pump out all of the Zeeman sublevels of the $D_{3/2}$ state. Even if we use only these circularly polarised




Figure 7.17: Partial energy level diagrams of calcium-40 at very low magnetic field strengths. *Top*: easily avoided dark states that can be formed in the $S_{1/2} \rightarrow P_{1/2}$ system. *Bottom*: Λ -type systems that can be formed with the σ components of the $D_{3/2} \rightarrow P_{1/2}$ transitions.

components, we can see from figure 7.17 that two Λ -type systems are formed and thus we have two dark states. One of these is a superposition of the $|m_j = -1/2\rangle$ and $|m_j = +3/2\rangle$ states whilst the other is a superposition of the $|m_j = -3/2\rangle$ and $|m_j = +1/2\rangle$ states.

In general, of course, we can have π polarised light present as well. This enables us to have more complicated systems that possess more than three levels. A detailed analysis of such systems in ions such as calcium-40 is presented in reference [107]. The authors show that regardless of the polarisation used there will always be two dark states present in the $D_{3/2} \rightarrow P_{1/2}$ system. These dark states can be superpositions of two, three or four Zeeman sublevels, depending on the intensities of the different polarisation components.

Because of the above result, when the magnetic field strength is exactly zero it is usually impossible to observe strong fluorescence on the 397 nm cooling transition. There are two techniques that are routinely used to prevent population trapping in the dark states so that cooling and detection become possible. The first of these is simply to apply a small magnetic field such that the Zeeman sublevels of the $D_{3/2}$ level are separated slightly. Obviously if the Zeeman splitting is too large then it becomes impossible to address all the components with a single laser, so a small magnetic field of just a few gauss is usually used. For some purposes, such as optical frequency standards, unshifted transition frequencies are required and so the application of a magnetic field is undesired. In these instances it is possible to prevent the driving of population into dark states by rapidly modulating the polarisation of the 866 nm laser.

Since in Doppler cooling schemes the laser is often scanned towards resonance over several GHz, we can expect it to be possible for it to be in two-photon resonance with the repumper laser at some point during the scan. If this occurs then notches can be seen in the lineshape of the cooling transition. Such notches have been observed with, for example, barium ions in RF traps in the presence of small magnetic fields [108]. Barium has a similar three-level structure to calcium, and an example barium ion fluorescence trace taken from this reference is shown in figure 7.18. Note that the width of these notches is increased by power broadening of the transitions.

In the Penning trap, with its much larger magnetic field, we would not generally expect to see coherent population trapping in superpositions of the $D_{3/2}$ levels. In principle if we have two or more of the 866 nm laser frequencies




Figure 7.18: Plot taken from [108], showing the fluorescence rate from barium ions as a function of the repumper laser detuning. Dips are clearly visible as the repumper laser passes through two-photon resonance with the Doppler cooling laser.

detuned from resonance by exactly the same amount it is possible to observe such effects. Each of the Zeeman sublevels is effectively addressed by its own laser, however, and these can all be detuned such that they are not in two-photon resonance. Thus simply by optimising the laser frequencies for maximum fluorescence we should be moving away from a situation in which dark states can appear.

What we might expect to see in the Penning trap is coherent population trapping in superpositions of the $S_{1/2}$ and $D_{3/2}$ levels. This is exactly the same situation as that which led to the RF trap barium absorption profile shown in figure 7.18, where notches were seen as the repumper laser passed through two-photon resonance with the Doppler cooling laser at various points in its frequency scan. Indeed we regularly see these small dips as we change the laser parameters when attempting to optimise the signal. An example fluorescence trace taken in the Penning trap and exhibiting two-photon resonance dips is shown in figure 7.19. Bearing in mind that for the lowest ion temperatures we want to have the Doppler cooling lasers detuned by half the transition linewidth and the repumpers (almost) on resonance, we ought not to be affected by these dips once we have found the optimum laser wavelengths.




Figure 7.19: Plot showing the variation in fluorescence rate from ions in a Penning trap as one of the Doppler cooling lasers is scanned towards resonance. Note the loss of fluorescence that occurs as the cooling laser passes through two-photon resonance with one of 866 nm repumper lasers (dashed line).




Figure 7.20: Partial energy level diagram showing the four-level system modelled. Calcium-40 consists of two such systems coupled only by spontaneous decay. The frequencies, detunings and Rabi frequencies for lasers L2 and L3 are defined in the same way as those shown for laser L1.

Since it is possible for coherent population trapping to occur in more complicated state superpositions, we have modelled the calcium-40 system using a numerical density matrix approach in order to determine whether there are any other dark states present. The detailed modeling was performed by H. Ohadi, with my involvement being in the direction of the investigations and the interpretation of the results. The approach taken follows closely that described in references [109, 110]. In order to simplify the problem, we take advantage of the fact that the eight levels in our system can be broken up into two identical four-level systems as shown in figure 7.20. These two systems are coupled only through spontaneous decay processes, where coherences are in any case lost. An extension to eight levels would be relatively straightforward, but significantly increases the computation time required to calculate steady state solutions.

One of the $P_{1/2}$ states of calcium-40 is denoted $|0\rangle$. It is the population of this state that determines the fluorescence rate that we can observe. Since we address the π components of the 397 nm transition we need only include one

of the ground state Zeeman sublevels, $|1\rangle$. The repumper lasers are tuned to the σ components, and so we need to include two of the $D_{3/2}$ sublevels, which we call $|2\rangle$ and $|3\rangle$. States $|0\rangle$ and $|1\rangle$ have a frequency difference ω_{01} and are coupled by a laser with a frequency ω_{L1} so that its detuning is $\delta_{L1} = \omega_{L1} - \omega_{01}$. The Rabi frequency associated with this coupling is Ω_{R1} . Similar definitions are made for the frequencies, detunings and Rabi frequencies of the lasers addressing the transitions between $|0\rangle$ and the two metastable states $|2\rangle$ and $|3\rangle$. We also define the Raman detuning δ_{Rij} to be the difference in detuning between lasers i and j , such that $\delta_{Rij} = \delta_{Li} - \delta_{Lj}$.

The equations describing this system were derived by H. Ohadi, following the approach used in references [109, 110]. The wavefunction of the internal state of our ion is given by

$$|\psi\rangle = \sum_{i=0}^3 c_i |i\rangle. \quad (7.1)$$

We begin by introducing the density operator, $\hat{\rho} = |\psi\rangle\langle\psi|$. Thus $\hat{\rho}$ is defined by a 4×4 density matrix, such that

$$\hat{\rho}_{ij} = \langle i | \hat{\rho} | j \rangle = c_i c_j^*. \quad (7.2)$$

The diagonal elements $|c_i|^2$ are the probabilities of the ion being in the state $|i\rangle$ and the off-diagonals are the coherences $c_i c_j^*$. In order to conserve population, we require that

$$\text{Tr}(\hat{\rho}) = \sum_{i=0}^3 \hat{\rho}_{ii} = 1. \quad (7.3)$$

In order to determine the equations of motion for our system we take the time derivative of the density operator,

$$\dot{\hat{\rho}} = |\dot{\psi}\rangle\langle\psi| + |\psi\rangle\langle\dot{\psi}|. \quad (7.4)$$

We can substitute for $|\dot{\psi}\rangle$ using the time-dependent Schrödinger equation

$$|\dot{\psi}\rangle = -\frac{i}{\hbar} \hat{H} |\psi\rangle, \quad (7.5)$$

to obtain

$$\dot{\hat{\rho}} = -\frac{i}{\hbar} [\hat{H}, \hat{\rho}]. \quad (7.6)$$

CHAPTER 7: TOWARDS SINGLE-ION WORK IN THE PENNING TRAP

We must also add a relaxation operator, \mathfrak{R} , which will account for the effects of spontaneous emission. The matrix elements therefore evolve as

$$\dot{\hat{\rho}}_{ij} = -\frac{i}{\hbar} \sum_{k=0}^3 \left(\hat{H}_{ik} \hat{\rho}_{kj} - \hat{\rho}_{ik} \hat{H}_{kj} \right) + \hat{\mathfrak{R}} \hat{\rho}_{ij}. \quad (7.7)$$

Our Hamiltonian includes an internal energy term, \hat{H}_0 , and an atom-laser interaction term, \hat{V}_{AL} ,

$$\begin{aligned} \hat{H} &= \hat{H}_0 + \hat{V}_{AL} \\ &= \sum_{i=0}^3 E_i |i\rangle \langle i| - \mu_{ij} E, \end{aligned} \quad (7.8)$$

where μ_{ij} is the electric dipole matrix element for a transition $|i\rangle \rightarrow |j\rangle$. Expressing \hat{V}_{AL} in terms of the Rabi frequencies defined earlier we obtain

$$\hat{V}_{AL} = \sum_{i=1}^3 \frac{\hbar \Omega_{Ri}}{2} e^{-i(\omega_{Li}t + \phi_i)} |0\rangle \langle i| + \text{h.c.} \quad (7.9)$$

We can now begin calculating the evolution equations for each matrix element. For example,

$$\begin{aligned} \dot{\hat{\rho}}_{00} &= -\frac{i}{\hbar} \left(\langle 0 | \hat{H}_0 \hat{\rho} | 0 \rangle + \langle 0 | \hat{V}_{AL} \hat{\rho} | 0 \rangle - \langle 0 | \hat{\rho} \hat{H}_0 | 0 \rangle - \langle 0 | \hat{\rho} \hat{V}_{AL} | 0 \rangle \right) + \hat{\mathfrak{R}} \hat{\rho}_{00} \\ &= -\frac{i}{\hbar} \left(E_0 \hat{\rho}_{00} + \frac{\hbar}{2} \sum_{i=1}^3 \left(\Omega_{Ri} \hat{\rho}_{i0} e^{-i(\omega_{Li}t + \phi_i)} + \text{c.c.} \right) - E_0 \hat{\rho}_{00} \right) + \hat{\mathfrak{R}} \hat{\rho}_{00} \\ &= -\frac{i}{2} \left(\Omega_{R1} \tilde{\rho}_{01}^* + \Omega_{R2} \tilde{\rho}_{02}^* + \Omega_{R3} \tilde{\rho}_{03}^* \right) + \text{c.c.} + \hat{\mathfrak{R}} \hat{\rho}_{00} \end{aligned} \quad (7.10)$$

Where we have made the substitution $\tilde{\rho}_{i0} = \hat{\rho}_{i0} e^{-i(\omega_{Li}t + \phi_i)}$, $\forall 1 \leq i \leq 3$ in order to eliminate the rapid components of the time evolution. Proceeding in exactly the same way for the other density matrix elements, and making the equivalent substitution $\tilde{\rho}_{ij} = \hat{\rho}_{ij} e^{-i((\omega_{Li} - \omega_{Lj})t + (\phi_i - \phi_j))}$, $\forall 1 \leq i \leq 3, i \leq j$, we

obtain the full set of equations:

$$\begin{aligned}
\dot{\hat{\rho}}_{00} &= \frac{i}{2} \left(\Omega_{R1}^* \tilde{\hat{\rho}}_{01} + \Omega_{R2}^* \tilde{\hat{\rho}}_{02} + \Omega_{R3}^* \tilde{\hat{\rho}}_{03} \right) + \text{c.c.} + \hat{\mathfrak{A}} \hat{\rho}_{00} \\
\dot{\hat{\rho}}_{11} &= -\frac{i}{2} \Omega_{R1}^* \tilde{\hat{\rho}}_{01} + \text{c.c.} + \hat{\mathfrak{A}} \hat{\rho}_{11} \\
\dot{\hat{\rho}}_{22} &= -\frac{i}{2} \Omega_{R2}^* \tilde{\hat{\rho}}_{02} + \text{c.c.} + \hat{\mathfrak{A}} \hat{\rho}_{22} \\
\dot{\hat{\rho}}_{33} &= -\frac{i}{2} \Omega_{R3}^* \tilde{\hat{\rho}}_{03} + \text{c.c.} + \hat{\mathfrak{A}} \hat{\rho}_{33} \\
\dot{\tilde{\hat{\rho}}}_{01} &= i\delta_{L1} \tilde{\hat{\rho}}_{01} + \frac{i}{2} \left(\Omega_{R1} (\tilde{\hat{\rho}}_{00} - \tilde{\hat{\rho}}_{11}) - \Omega_{R2} \tilde{\hat{\rho}}_{21} - \Omega_{R3} \tilde{\hat{\rho}}_{31} \right) + \hat{\mathfrak{A}} \hat{\rho}_{01} \\
\dot{\tilde{\hat{\rho}}}_{02} &= i\delta_{L2} \tilde{\hat{\rho}}_{02} + \frac{i}{2} \left(\Omega_{R2} (\tilde{\hat{\rho}}_{00} - \tilde{\hat{\rho}}_{22}) - \Omega_{R1} \tilde{\hat{\rho}}_{12} - \Omega_{R3} \tilde{\hat{\rho}}_{32} \right) + \hat{\mathfrak{A}} \hat{\rho}_{02} \\
\dot{\tilde{\hat{\rho}}}_{03} &= i\delta_{L3} \tilde{\hat{\rho}}_{03} + \frac{i}{2} \left(\Omega_{R3} (\tilde{\hat{\rho}}_{00} - \tilde{\hat{\rho}}_{33}) - \Omega_{R1} \tilde{\hat{\rho}}_{13} - \Omega_{R2} \tilde{\hat{\rho}}_{23} \right) + \hat{\mathfrak{A}} \hat{\rho}_{03} \\
\dot{\tilde{\hat{\rho}}}_{12} &= -i\delta_{R12} \tilde{\hat{\rho}}_{12} + \frac{i}{2} \left(\Omega_{R2} \tilde{\hat{\rho}}_{10} - \Omega_{R1}^* \tilde{\hat{\rho}}_{02} \right) + \hat{\mathfrak{A}} \hat{\rho}_{12} \\
\dot{\tilde{\hat{\rho}}}_{13} &= -i\delta_{R13} \tilde{\hat{\rho}}_{13} + \frac{i}{2} \left(\Omega_{R3} \tilde{\hat{\rho}}_{10} - \Omega_{R1}^* \tilde{\hat{\rho}}_{03} \right) + \hat{\mathfrak{A}} \hat{\rho}_{13} \\
\dot{\tilde{\hat{\rho}}}_{23} &= -i\delta_{R23} \tilde{\hat{\rho}}_{23} + \frac{i}{2} \left(\Omega_{R3} \tilde{\hat{\rho}}_{20} - \Omega_{R2}^* \tilde{\hat{\rho}}_{03} \right) + \hat{\mathfrak{A}} \hat{\rho}_{23} \\
\dot{\tilde{\hat{\rho}}}_{ji} &= \dot{\tilde{\hat{\rho}}}_{ij}^*
\end{aligned} \tag{7.11}$$

The relaxation operator is used to account for processes that transfer population in an incoherent way and hence destroy coherences. For the purposes of our model we neglect all such processes except spontaneous emission from the $P_{1/2}$ state, $|0\rangle$. Decay is into the ground state with probability 12/13, and into each of the metastable states with probability 1/26. The metastable levels are assumed not to decay, since the lifetime of these states is much longer than the time taken for an equilibrium solution to be reached. Allowance is made for the fact that spontaneous decay can occur into and out from levels not included in the model.

A Mathematica script was used to perform a numerical calculation of the evolution of the density matrix. The matrix is originally set such that all the terms except $\hat{\rho}_{11}$ (corresponding to the ground state) are zero. The evolution is allowed to continue long enough for a steady-state situation to be reached. This duration is determined by looking at how the result changes if the system is allowed to evolve for a longer period of time.

A variety of relevant Rabi frequencies and detunings was modelled. Al-

though two-photon resonances and ac Stark shifts are clearly visible, at no point are any higher-order effects observed. Significantly, the fraction of the population in the $\hat{\rho}_{00}$ state that can be achieved is unaffected and with sufficient laser power can be close to the optimum value of 0.25. This implies that the scattering rate on the cooling transition is similar to that achieved in the RF trap, providing we set the parameters correctly.

A typical result, corresponding to experimental parameters that are easily realisable, is shown in figure 7.21. Here the Rabi frequencies are $\Omega_{R1} = 0.2$ and $\Omega_{R2,3} = 0.5$, in units of the decay rate of the $P_{1/2}$ state, Γ . The detuning of one of the infra-red lasers (δ_{L3}) is set to zero. The contour plots show the populations of the four levels for different detunings of the cooling and second repumper beams (again in units of Γ).

Perhaps the most important feature is the peak in the upper-left quadrant of the $\hat{\rho}_{00}$ plot, since this tells us the optimum laser detunings. Ideally the cooling laser (L1) should be on the red-detuned side of this peak, since we want a large fluorescence rate and a steep gradient for efficient cooling. The maximum value of the $\hat{\rho}_{00}$ density matrix element also determines how much fluorescence we can expect for ion detection.

The vertical and diagonal stripes on the $\hat{\rho}_{00}$ plot correspond to dark states due to two-photon resonances between the cooling laser and the repumper lasers ($\delta_{R12} = 0$ and $\delta_{R13} = 0$). The less deep trough that runs horizontally across the plot is due to two-photon resonance between the two repumper lasers ($\delta_{R23} = 0$). For each of these three regions where we have two-photon resonance, there are ridges in the plots of $\hat{\rho}_{22}$ and $\hat{\rho}_{33}$. This confirms that the population has been transferred into one of (or a superposition involving) these states. As we have said, the diagonal stripe corresponds to the formation of a superposition of the states $|1\rangle$ and $|2\rangle$. That this appears to be a valley on much of the $\hat{\rho}_{11}$ plot is due simply to the fact that in these far-detuned regions most of the population would never leave the ground state anyway were it not for the presence of the two-photon resonance.

For laser cooling we always want to red-detune the cooling laser, L1, from the lowest frequency peak. It is expected that the optimum detunings for the repumper lasers, L2 and L3, would be to have either one laser on resonance and one blue-detuned (as with the previous results), or both lasers blue-detuned. To compare the lineshape and strength of the cooling transition we have generated similar contour plots to those described above, but with the fixed laser




Figure 7.21: Results of the numerical model, showing the steady-state populations of the four levels as a function of the detunings δ_{L1} and δ_{L2} when $\delta_{L3} = 0$. Lighter shades indicate larger populations. A reduction in the excited state population, $\hat{\rho}_{00}$, is observed when there is two-photon resonance between each pair of lasers.

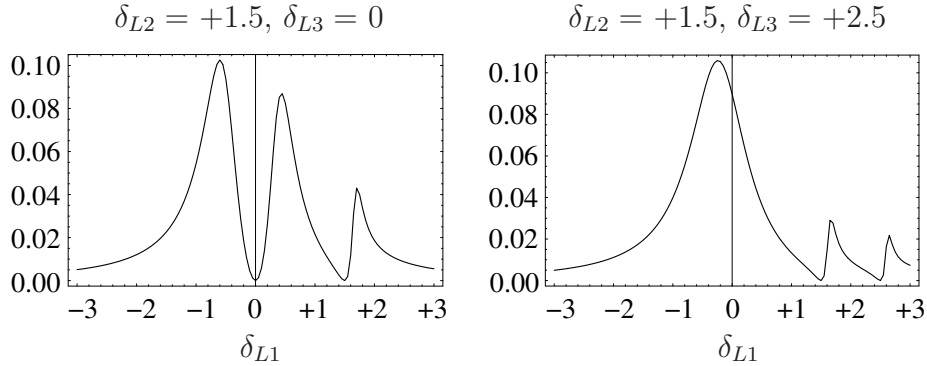


Figure 7.22: Plots of $\hat{\rho}_{00}$ as a function of cooling laser detuning δ_{L1} . *Left:* $\delta_{L2} = +1.5$, $\delta_{L3} = 0$. *Right:* $\delta_{L2} = +1.5$, $\delta_{L3} = +2.5$.

frequency $\delta_{L3} = +2.5$. In each of the two cases we determined the value of δ_{L2} that gives the maximum fluorescence and plotted $\hat{\rho}_{00}$ as a function only of the cooling laser detuning δ_{L1} . This corresponds to a horizontal cut through each $\hat{\rho}_{00}$ contour plot, passing through the peak in the upper left region. The two resulting plots can be seen in figure 7.22.

The positions of the two-photon resonances are of course different, but the maximum population in the $|0\rangle$ state is the same for the two situations. The half-width of the peak in the first case, however, is somewhat smaller than in the second case. Since this width governs the temperature to which we can cool the ions we can conclude that the optimum parameters for laser cooling involve one of the repumper lasers being on resonance, with the other blue-detuned.

For optimal laser cooling we need to make a trade off between maximum scattering rate, which involves large laser powers, and (almost) unbroadened transitions, which require low laser powers. When discussing the fluorescence rates that we could expect from a single ion, we assumed that in order to prevent power broadening we would only be able to work with a fluorescence rate about half of that which is theoretically possible. We have validated this by looking at the effect of increasing the cooling laser's Rabi frequency in our model. We used the same parameters as we did for the plots shown in figure 7.21, but with $\Omega_{R1} = 2.0$ instead of 0.2. The contour plots for $\hat{\rho}_{00}$ and

the cuts through the position of peak fluorescence are shown in figure 7.23 for both cases.

The increased laser power on the cooling transitions leads to almost a doubling of the population of the $|0\rangle$ state, and hence there will be a corresponding increase in the fluorescence rate. There is, however, a dramatic increase in the width of the transition and therefore a raised Doppler limit. That we can in principle achieve populations of the $|0\rangle$ state approaching the maximum possible 25% is comforting since it suggests that nothing unexpected is occurring in the system. More importantly, however, the model supports our assumption that it should be possible to achieve about half of the maximum possible fluorescence rate without significant power broadening of the transition.

We also assumed previously that oversaturation of the repumping lasers did not cause broadening of the cooling transition, and so did not raise the Doppler limit. Once again we have checked this by looking at the effect of increasing the relevant Rabi frequencies in the model. We used the same parameters as we did for the results in figure 7.21, but increased both Ω_{R2} and Ω_{R3} from 0.5 to 2.0. The results are plotted in figure 7.24.

The most obvious feature of the plots is that the regions where the population of the $|0\rangle$ state is very low are much broader. This appears to be in part due to broadening of the two-photon resonances. For the most part, however, the population in the $|2\rangle$ and $|3\rangle$ states remains small in these regions — suggesting that the population is not trapped in these levels. It seems that the increased repumper intensity leads to a significant ac Stark shifting of the cooling transition and hence in the regions where δ_{L1} is small most of the population remains in the ground state. Although the position of peak fluorescence is shifted, the absolute value of the maximum is the same and the transition is not broadened significantly. Thus our ability to efficiently laser cool is not affected by oversaturation of the repumper transitions. It is interesting to note that if the ion moves between regions of high and low repumper intensity the extent of the ac Stark shifting will vary and hence the Doppler cooling transition will, in effect, be broadened. That there is no improvement from the increased repumping efficiency is probably due to the fact that we are using a low cooling laser Rabi frequency to avoid power broadening. The low rate of transfer into states $|2\rangle$ and $|3\rangle$ means that we are not very sensitive to the rate at which we pump out from these states.

From looking at the results of our model we find that there is no evidence

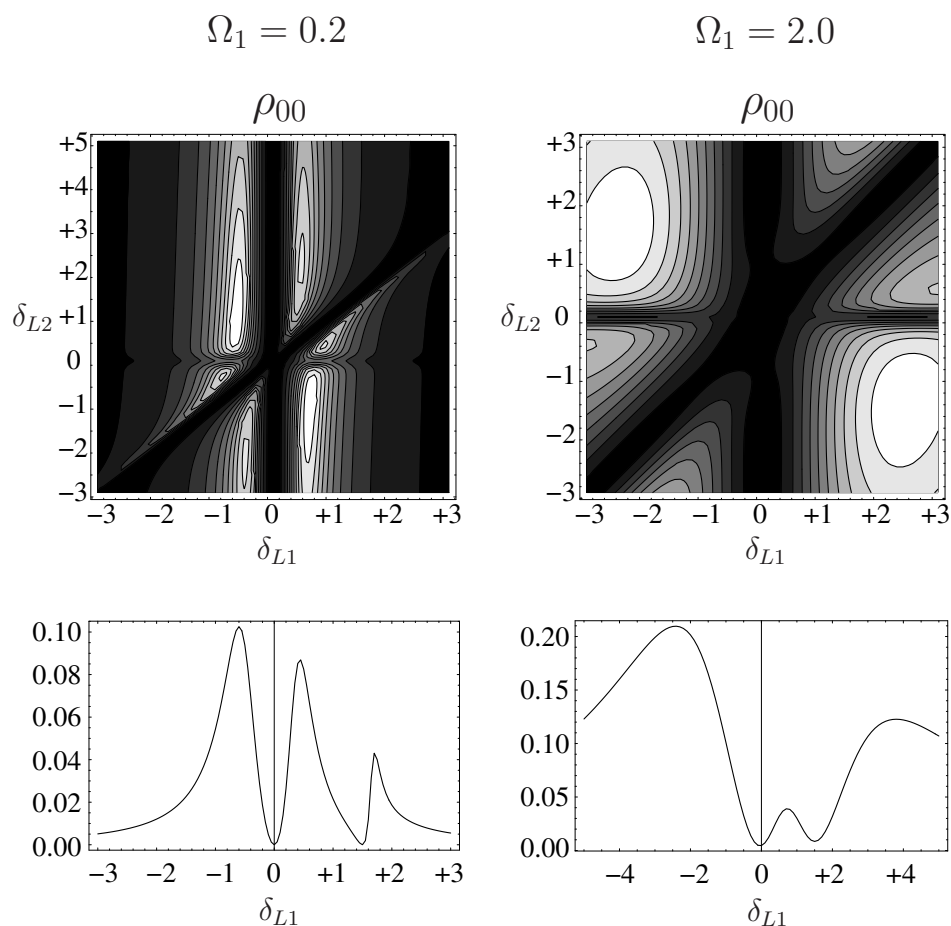


Figure 7.23: Contour plots of $\hat{\rho}_{00}$ and horizontal cuts through the position of maximum fluorescence when (*left*) $\Omega_{R1} = 0.2$ and (*right*) $\Omega_{R1} = 2.0$. All other parameters are the same as in figure 7.21. Note the changes in scale on some of the plots.




Figure 7.24: *Top*: contour plots of the steady-state populations of all four states when $\Omega_{R2,3} = 2.0$. *Bottom*: cuts through the position of maximum fluorescence when (*left*) $\Omega_{R2,3} = 2.0$ and (*right*) $\Omega_{R2,3} = 0.5$. All other parameters are the same as in figure 7.21. Note the change in scale on the bottom-left plot

to suggest that there should be a fundamental problem with laser cooling in the $^{40}\text{Ca}^+$ system at high magnetic fields. Our results do, however, give us an indication of the lack of tolerance that we have in terms of the laser parameters. In order to obtain a large fraction of the population in the $|0\rangle$ state we must keep all three laser frequencies within about one natural linewidth of the transition. Recalling that the true system comprises two such four-level systems, which are incoherently coupled by spontaneous decay, we find that the demands on our laser parameters are stringent indeed. Even if one system is addressed efficiently, population will simply optically pump into the other system. Thus it is imperative that all six laser frequencies be kept constant if we are to see strong fluorescence.

7.6 Laser and magnetic field instability

Having satisfied ourselves that our poor signal rate per ion in the Penning trap is unlikely to be due to the presence of contaminant ions or some non-fluorescing state of the calcium system into which the population is driven, we must look for some other possible cause. Thus far we have not considered how the stabilities of the lasers and magnetic field affect our ability to efficiently cool very small numbers of ions.

We saw in section 7.3 that the temperature of an ion cloud drops very rapidly as we approach the kind of loading parameters that give us what we think are just a few ions. Though the signal level is not sufficient to allow us to make even an estimate of the temperature for these very small clouds, we might expect it to be even lower. Even the lowest temperatures obtained in our earlier measurements correspond to Doppler profiles with a FWHM of 200MHz or more. Once the temperature reaches about 0.2 kelvin the Doppler width is comparable with the 24MHz natural width. A drop to this kind of temperature would make us far more sensitive to the wavelengths of both the cooling and repumper lasers. It is possible that when we have only a few ions in the trap we reach this regime and the laser stability limits the fluorescence rate that we observe.

It is worth noting that although the amount of fluorescence observed is proportional to the scattering rate and therefore related to the cooling rate, the fluorescence does not necessarily relate directly to the temperature of the ions. The temperature of the ions depends also on the heating rate in the

system, and we expect this to be very different for different sized clouds. Even though the scattering rate per ion may be worse when we have only a few ions in the trap, the heating rate can be much less so that the temperature of the cloud is lower. This may explain why we obtain healthy signal from single ions in the RF trap and are not affected by laser instability — the RF heating means that the temperature of an ion may actually be quite large even though the fluorescence rate is high.

Besides the fact that we have measured our laser frequency stabilities per hour to be in the region of 1 to 4 times the natural width of the transitions we are addressing, it may not be the drift in the laser frequencies that makes us unable to keep them addressing their corresponding transitions. The Zeeman splitting of the cooling and repumping transitions at 0.98 tesla is about 10GHz. We have calculated from repeated motional frequency measurements that the magnetic field stability is in the region of 10^{-3} – 10^{-4} over the course of several hours. Though we do not know what the magnetic field stability is like on much shorter timescales, the worst case shift in each transition frequency would be about 10MHz. It is possible that when cooling ions with a very low Doppler width, suddenly moving each laser 10MHz from its transition is enough to dramatically reduce the scattering rate.

We know that single magnesium ions have been trapped and laser cooled to the Doppler limit in a Penning trap using an electromagnet identical to ours [91, 92]. That this experiment was not affected despite having a similar magnetic field stability may be due to the fact that the cooling scheme for magnesium requires only a single transition to be addressed by a narrow linewidth laser, and that this transition has a width double that of ours. Thus they might not have been as sensitive to changes in magnetic field at the 10^{-3} level.

Ultimately the best way to determine if this is the problem is to improve the stability of the lasers and magnetic field. We have already assembled a new split-ring ion trap designed to fit into the bore of a superconducting magnet. This magnet has a nominal stability of 10^{-7} per hour and we expect the short term stability to be even better than this. It is hoped that the use of this magnet will ease the difficulties that we have had with the trapping of single calcium ions. It is expected that laser stability will then become the limiting factor and that we will be better able to gauge the extent of frequency drifts.

Chapter 8

Discussion

8.1 Discussion of results

The stated long-term aim of our project is to measure the decoherence and motional heating rates of calcium ions in a Penning trap. In order to achieve this aim we must first be able to trap single ions and cool them to the Doppler limit. Based on the findings of previous experiments with other ion species, in order to do this it will be necessary to implement axialisation so that ions can be cooled more efficiently.

In this thesis I have presented results showing that we are able to routinely trap, axialise and laser cool clouds of ions in a Penning trap. Analysis of the aspect ratios of the clouds provides us with convincing evidence for the efficacy of the axialisation technique. Moreover, by using a RF-photon correlation technique we can obtain a direct measure of the rate of cooling of the system and we find that in the presence of axialisation this rate is increased. Drawing more detailed conclusions from these results is difficult because the error on the measurements is rather large, due mostly to the instability of the laser systems used and the rather long measurement process. Improvements in laser stability would allow much more consistent cooling and therefore improved measurements of cooling rate. The instability of the magnetic field and the fact that we are forced to work with clouds of ions also hamper our attempts to parameterise the axialisation technique.

We have extended a model for laser cooling in the Penning trap to include the effects of axialisation. This model successfully yields the basic observed features of axialisation, namely a coupling of the radial modes of motion and

hence an improvement in the cooling rate of the magnetron motion. The model assumes that there is only one ion present in the trap, and that it is sufficiently cold that the extent of its motion is small. These assumptions do not match the experimental conditions we can achieve, and hence we do not observe cooling rates that correspond directly to those obtained from the model. The coupling effects, such as the shifting of the motional frequencies when the axialisation drive is near resonant, are accurately predicted by the model. The only discrepancy lies in the absolute magnitude of the coupling rate, and this is attributed to space charge effects in the cloud. If we were able to repeat our measurements using a single ion we would genuinely be satisfying the assumptions made in the model. We would therefore expect to see better agreement between the cooling and coupling rates measured and those predicted by the model.

When attempting to load single ions of calcium in the Penning trap we have found that the signal rate per ion appears to be unexpectedly low. An experimental and theoretical analysis suggests that this is not due to the presence of contaminant ions in the trap, or the transfer of ion population into some non-fluorescing state. The fact that we can easily trap single ions in the RF trap suggests that there is no fundamental flaw in the experimental arrangement. Single ions of a different species have previously been laser cooled in the Penning trap, where the main difference was a greatly simplified laser cooling scheme that required only one laser. It therefore seems reasonable that the cause of our poor signal rate with calcium is due to the rather large drift rates of the many laser systems used. This may be exacerbated by the poor stability of our electromagnet. In conclusion, it seems that improvements in magnetic field and laser stability may hold the key to observing single calcium ions in the Penning trap.

8.2 Future directions

8.2.1 Improving magnetic field stability

We have identified the stability of the magnetic field as being one of the limiting factors in our ability to maintain efficient Doppler cooling for extended periods of time. Thus in the short term an improvement in magnetic field stability might help us to improve fluorescence rates and trap single ions. In the near

future we would also hope to begin addressing our narrow qubit transition. This transition has a natural width of less than 1 hertz, and hence we need a stable laser to address it. A Pound–Drever–Hall stabilised Ti:sapphire laser with a linewidth of about 7kHz has recently been developed by our group. Since the Zeeman splitting of the least shifted 729nm transition is about 5GHz, we require a field stability of about 1 part in 10^6 if it is to be the laser, rather than the magnetic field, that is the limiting factor. Even if we allow the magnetic field to limit us, sideband cooling requires that we be able to resolve the motional sidebands from the carrier transitions. Depending on the motion we are addressing the separation frequencies will vary between about 25kHz (for the magnetron) and 350kHz (for the modified cyclotron), so the limits are only slightly relaxed. This required stability is orders of magnitude beyond what our electromagnet can deliver.

Looking at the long term aims of the project, the decoherence rates we can measure will ultimately be limited by the magnetic field stability. Just as experiments in RF traps are currently affected by qubit dephasing due to fluctuations in the ambient magnetic field, so our qubits will be affected by the trapping field instability. Thus it will in any case eventually be essential that we have a much more stable field.

To this end we have already on three occasions rebuilt a trap that is suitable for use with a superconducting magnet that we have available. The superconducting magnet has a long, narrow, vertical bore to which the magnetic field is parallel. Thus the axis of the trap must be vertical, rather than horizontal, and the vacuum housing must be designed such that the centre of the trap lies at the centre of the magnet coils whilst the other fittings lie outside of the magnet bore. The trap that was rebuilt had previously been used in experiments with different ions and is described in [111, 112]. The first two times we rebuilt this trap we incorporated a split ring electrode comparable to that in the trap used for the experiments described in this thesis. After we failed to observe any fluorescence on these two occasions we decided to rebuild the trap with a solid ring electrode. Although this makes it impossible to apply an axialisation drive, it should still provide valuable information on the improvement that can be gained from having a more stable magnetic field. Ultimately we would hope to reintroduce the split ring electrode.

The limited access to the superconducting magnet trap means that optical fibres are used to deliver the beams to a breadboard located just outside the

bore of the magnet. These fibres support only a single spatial mode and hence the emergent beams have an extremely clean spatial profile. This may prove to be of some benefit when we need to offset the beam from the centre of the trap, since light from the edges of the beam will not lie on the wrong side of the trap and cause magnetron heating. Having a clean beam profile also reduces the amount of light scattered when the beam is sent into the trap and so reduces the background light level. In order to check that the power loss due to coupling into the fibres is not responsible for our failure to observe fluorescence in the superconducting magnet trap, we have recently successfully trapped ions in the conventional electromagnet trap using the fibres.

According to the manufacturer's specifications, the long-term stability of the superconducting magnet is 1 part in 10^7 per hour. This takes the form of a linear drift, and so can be accounted for in our experiments. It is not known for certain what the short-term stability will be, but we hope that it will be one or two orders of magnitude better. In principle, superconducting magnets can exhibit excellent exclusion of external magnetic flux. Thus if the magnet's field is stable enough we may find that we can achieve a better overall stability with our large magnetic field than is possible in a RF trap with a very low, stabilised field.

8.2.2 Improving laser cavity stability

The stability of each laser used for Doppler cooling or repumping is set by the length stability of the cavity to which it is locked. This is in turn set by the cavity's temperature stability and expansion coefficient, and hence it is these two characteristics that must be improved if we are to attain better laser stability.

We have estimated the temperature stabilities of the cavities by monitoring the error signal used to control the temperature of the cavity housings. This error signal is found to be stable at about the millikelvin level. This, of course, assumes that the sensor signal genuinely represents the temperature of the cavities, and some independent means of measuring this would shed some light on the rather poor stability that is observed. Currently we use a single thermistor to generate the temperature signal. The use of several thermistors would enable us to obtain an average temperature for the different regions of the cavity housings. This might improve temperature stability and would give

us some indication of the measurement precision.

If the temperature stability really is as good as our measurements indicate, the implication would be that it is the thermal expansion coefficients of the cavities that are the cause of the problem. The only effective way to improve the temperature stability further would be to evacuate the cavities so that non-radiative heat transfer mechanisms are almost eliminated.

Improving the thermal properties of the cavities necessarily means rebuilding them. The components that are expected to contribute most to the expansivity are the PZT stacks used to tune the cavity and the TorrSeal epoxy used to join the various pieces. Although these components are arranged in a re-entrant configuration, it is possible that each has a different expansion coefficient and so they do not cancel each other out perfectly. When we begin carrying out sideband cooling we will need to perform rapid switching of each laser using acousto-optical modulators. We can also use these to detune the beam going to the trap relative to the beam that is locked to the cavities, and so we will not necessarily need to tune the reference cavities. Thus we may be able to eliminate the problematic components altogether.

An interesting alternative might be to stabilise the cavities themselves to some more stable reference. Thus the four cavity lengths would be locked to two additional lasers that are themselves locked to some atomic reference. Convenient atomic lines exist, for example, in Rubidium at 780nm. For a lock that eliminates only long-term drifts, it would be sufficient to have even a very low finesse at the reference wavelength. One difficulty is the lack of a suitable atomic line near the 397nm transition, but it might prove possible to use calcium ions in a hollow cathode lamp at zero magnetic field as a reference. Note that by changing the cavity mirrors we might be able to obtain a reasonable finesse both at our laser wavelengths and the wavelength of some stable reference laser, such as a stabilised He:Ne laser. The problem with using these kinds of stabilised cavity arrangements would be that at least two more laser locks would need to be maintained and this would add to the experimental complexity.

8.2.3 Sideband cooling and beyond

There are four candidates for the sideband cooling and qubit transition in the $S_{1/2} \rightarrow D_{5/2}$ system. The best in terms of the qubit transition are the

$\Delta m_j = 0$ components, since these are the least Zeeman shifted and hence the least vulnerable to magnetic field fluctuations. For sideband cooling, the optimum transitions are the $\Delta m_j = \pm 2$ components, since the upper states of these transitions always remain in the cooling cycle during quenching. Explicitly, in the first instance our sideband pulse might take us from $S_{1/2}(m_j = -1/2) \rightarrow D_{5/2}(m_j = -1/2)$. We then need to precisely address the quenching transition $D_{5/2}(m_j = -1/2) \rightarrow P_{3/2}(m_j = -3/2)$. Pumping into any other $P_{3/2}$ state will lead to the possibility of decay to the $S_{1/2}(m_j = +1/2)$ level and hence loss of the ion from the cooling cycle. Unfortunately the large frequency separation means that we cannot change rapidly from one of these transitions to the other, and so we must make a choice based on whether we are limited more by sideband cooling efficiency or magnetic field fluctuations.

In either case the sideband cooling sequence will consist of optical pumping into the relevant ground state, followed by a number of cycles of a (near) π pulse on the red sideband of the chosen 729nm transition and a pulse on the required 854nm quenching transition.

Once the ion reaches the motional ground state, it is no longer possible to drive transitions on the red sideband. Thus the relative amplitudes of the motional sidebands can be used to estimate the mean vibrational quantum number of that mode of the ion's motion. More accurate measurements are possible by looking at the (vibrational state-dependent) Rabi frequency of the 729nm transitions. Either of these techniques could be used to measure the heating rates of the ion motion by measuring the ion temperature a known time after sideband cooling is completed.

Measurement of decoherence rates would initially involve measuring the dephasing rate of single ions using Ramsey spectroscopy. Comparison of this rate with those measured using RF traps would prove extremely useful for determining the benefits of using Penning traps for quantum information processing.

All of the techniques presented in this section have not yet been implemented mainly because of the instability of the magnetic field we have available and the fact that we have not yet been able to isolate single ions. The move to a superconducting magnet may occur in a relatively short time, however, and we believe that this may assist us in trapping individual ions. If this is the case then it is not inconceivable that rapid progress could be made in performing preliminary spectroscopy of the 729nm transition. The key part of

the apparatus — namely the narrow linewidth laser — is already available for use. It would remain only to set up a system of acousto-optical modulators, to enable rapid pulsing of each laser, before we could begin sideband cooling. It would seem, then, that the experiment is at an exciting stage where we must overcome a single hurdle in order to open up the way to the ultimate aims of the project.

Bibliography

- [1] P. Benioff. The computer as a physical system: a microscopic quantum mechanical hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics*, 22:563–591, 1980.
- [2] D. Deutsch. Quantum theory, the Church–Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London A*, 400:97–117, 1985.
- [3] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [4] G.M. Palma, K. Suominen, and A.K. Ekert. Quantum computers and dissipation. *Proceedings of the Royal Society of London A*, 452:567–584, 1996.
- [5] D. Deutsch. Quantum computational networks. *Proceedings of the Royal Society of London A*, 425:73–90, 1989.
- [6] D. Deutsch, A. Barenco, and A. Ekert. Universality in quantum computation. *Proceedings of the Royal Society of London A*, 449:669–677, 1995.
- [7] S. Lloyd. Almost any quantum logic gate is universal. *Physical Review Letters*, 75:346–349, 1995.
- [8] J.I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74:4091–4094, 1995.
- [9] P.W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer (revised). *SIAM Journal on Computing*, 26:1484–1509, 1997.

BIBLIOGRAPHY

- [10] A. Ekert and R. Jozsa. Quantum computation and Shor's factoring algorithm. *Reviews of Modern Physics*, 68:733–753, 1996.
- [11] A. Steane. Quantum computing. *Reports on Progress in Physics*, 61:117–173, 1998.
- [12] A.R. Calderbank and P.W. Shor. Good quantum error-correcting codes exist. *Physical Review A*, 54:1098–1105, 1996.
- [13] A. Steane. Multi-particle interference and quantum error correction. *Proceedings of the Royal Society of London A*, 452:2551–2577, 1996.
- [14] A. Steane. Introduction to quantum error correction. *Philosophical Transactions of the Royal Society of London A*, 356:1739–1758, 1998.
- [15] A.Y. Kitaev. Quantum error correction with imperfect gates. In O. Hirota, editor, *Quantum Communication and Computing and Measurement*. Plenum, 1996.
- [16] A. Steane. Space, time, parallelism and noise requirements for reliable quantum computing. *Fortschritte der Physik*, 46:443–458, 1998.
- [17] D.P. DiVincenzo. The physical implementation of quantum computing. In S.L. Braunstein, editor, *Scalable quantum computers : Paving the way to realization*, pages 1–13. Wiley-VCH, 2001.
- [18] D. Wineland et al. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103:259–328, 1998.
- [19] D. Kielpinski, C. Monroe, and D. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, 2002.
- [20] F. Schmidt-Kaler et al. Realization of the Cirac–Zoller controlled-NOT quantum gate. *Nature*, 422:408–411, 2003.
- [21] D. Leibfried et al. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422:412–415, 2003.
- [22] P.K. Ghosh. *Ion Traps*. Oxford University Press, 1995.

BIBLIOGRAPHY

- [23] H.C. Nägerl et al. Investigating a qubit candidate : Spectroscopy on the $S_{1/2}$ to $D_{5/2}$ transition of a trapped calcium ion in a linear Paul trap. *Physical Review A*, 61:023405, 2000.
- [24] D. Bouwester, A. Ekert, and A. Zeilinger, editors. *The Physics of Quantum Information*. Springer, 2000.
- [25] B.E. King et al. Cooling the collective motion of trapped ions to initialize a quantum register. *Physical Review Letters*, 81:1525–1528, 1998.
- [26] A. Sørensen and K. Mølmer. Entanglement and quantum computation with ions in thermal motion. *Physical Review A*, 62:022311, 2000.
- [27] J.F. Poyatos, J.I. Cirac, and P. Zoller. Quantum gates with “hot” trapped ions. *Physical Review Letters*, 81:1322–1325, 1998.
- [28] D.F.V. James. QC with hot and cold ions. In S.L. Braunstein, editor, *Scalable quantum computers : Paving the way to realization*, pages 53–67. Wiley-VCH, 2001.
- [29] D. Leibfried et al. Towards quantum information with trapped ions at NIST. *Journal of Physics B*, 36:599–612, 2003.
- [30] F. Schmidt-Kaler et al. The coherence of qubits based on single Ca^+ ions. *Journal of Physics B*, 36:623–636, 2003.
- [31] C.A. Sackett. Quantum information experiments with trapped ions: status and prospects. *Quantum Information and Computation*, 1:57–80, 2001.
- [32] C. Monroe, D.M. Meekhof, B.E. King, W.M. Itano, and D.J. Wineland. Demonstration of a fundamental quantum logic gate. *Physical Review Letters*, 75:4714–4717, 1995.
- [33] Q.A. Turchette et al. Deterministic entanglement of two trapped ions. *Physical Review Letters*, 81:3631–3634, 1998.
- [34] K. Mølmer and A. Sørensen. Multiparticle entanglement of hot trapped ions. *Physical Review Letters*, 82:1835–1838, 1999.
- [35] C.A. Sackett et al. Experimental entanglement of four particles. *Nature*, 404:256–259, 2000.

BIBLIOGRAPHY

- [36] F. Schmidt-Kaler et al. How to realize a universal quantum gate with trapped ions. *Applied Physics B*, 77:789–796, 2003.
- [37] S. Gulde et al. Implementation of the Deutsch–Josza algorithm on an ion-trap quantum computer. *Nature*, 421:48–50, 2003.
- [38] M.D. Barrett et al. Deterministic quantum teleportation of atomic qubits. *Nature*, 429:737–739, 2004.
- [39] M. Riebe et al. Deterministic quantum teleportation with atoms. *Nature*, 429:734–737, 2004.
- [40] C. Roos et al. Control and measurement of three-qubit entangled states. *Science*, 304:1478–1480, 2004.
- [41] D. Leibfried et al. Toward Heisenberg-limited spectroscopy with multi-particle entangled states. *Science*, 304:1476–1478, 2004.
- [42] J. Chiaverini et al. Realization of quantum error correction. *Nature*, 432:602–605, 2004.
- [43] J. Chiaverini et al. Implementation of the semiclassical quantum Fourier transform in a scalable system. *Science*, 308:997–1000, 2005.
- [44] H. Häffner et al. Scalable multiparticle entanglement of trapped ions. *Nature*, 438:643–646, 2005.
- [45] D. Leibfried et al. Creation of a six-atom ‘Schrödinger cat’ state. *Nature*, 438:639–642, 2005.
- [46] P.C. Haljan, P.J. Lee, K.-A. Brickman, M. Acton, L. Deslauriers, and C. Monroe. Entanglement of trapped-ion clock states. *Physical Review A*, 62:062316, 2005.
- [47] L. Deslauriers, P.C. Haljan, P.J. Lee, K.-A. Brickman, B.B. Blinov, M.J. Madsen, and C. Monroe. Zero-point cooling and low heating of trapped $^{111}\text{Cd}^+$. *Physical Review A*, 70:043408, 2004.
- [48] B.B. Blinov, L. Deslauriers, P. Lee, M.J. Madsen, R. Miller, and C. Monroe. Sympathetic cooling of trapped Cd^+ isotopes. *Physical Review A*, 65:040304, 2002.

BIBLIOGRAPHY

- [49] B.B. Blinov, D.L. Moehring, L.-M. Duan, and C. Monroe. Observation of entanglement between a single trapped atom and a single photon. *Nature*, 428:153–157, 2004.
- [50] D. Stick, W.K. Hensinger, S Olmschenk, M.J. Madsen, K. Schwab, and C. Monroe. Ion trap in a semiconductor chip. *Nature Physics*, 2:36–39, 2006.
- [51] W.K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, and C. Monroe. T-junction ion trap array for two-dimensional ion shuttling, storage and manipulation. *Applied Physics Letters*, 88:034101, 2006.
- [52] D.M. Lucas et al. Oxford ion-trap quantum computing project. *Philosophical Transactions of the Royal Society A*, 361:1401–1408, 2003.
- [53] C. Balzer, A. Braun, T. Hannemann, C. Paape, M. Ettl, W. Neuhauser, and C. Wunderlich. Electrodynamically trapped Yb^+ ions for quantum information processing. *quant-ph/0602044*, 2006.
- [54] F. Mintert and C. Wunderlich. Ion-trap quantum logic using long-wavelength radiation. *Physical Review Letters*, 87:257904, 2001.
- [55] F. Mintert and C. Wunderlich. Erratum: Ion-trap quantum logic using long-wavelength radiation. *Physical Review Letters*, 91:029902, 2003.
- [56] C. Wunderlich. Conditional spin resonance with trapped ions. *quant-ph/0111158*, 2001.
- [57] G. Morigi, D. Reiß, and C. Wunderlich. Simultaneous cooling of axial vibrational modes in a linear ion-trap. *Physical Review A*, 72:023421, 2005.
- [58] A. Kreuter et al. Spontaneous emission lifetime of a single trapped Ca^+ ion in a high finesse cavity. *Physical Review Letters*, 92:203002, 2004.
- [59] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. Continuous generation of single photons with controlled waveform in an ion-trap cavity system. *Nature*, 431:1075–1078, 2004.
- [60] G. Morigi, J. Eschner, S. Mancini, and D. Vitali. Entangled light pulses from single cold atoms. *Physical Review Letters*, 96:023601, 2006.

BIBLIOGRAPHY

- [61] J.J. Bollinger, D.J. Heinzen, W.M. Itano, and D.J. Wineland. Optimal frequency measurements with maximally correlated states. *Physical Review A*, 54:R4649, 1996.
- [62] C. Roos et al. Bell states of atoms with ultralong lifetimes and their tomographic state analysis. *Physical Review Letters*, 92:220402, 2004.
- [63] H. Häffner et al. Robust entanglement. *Applied Physics B*, 81:151–153, 2005.
- [64] J.J. Bollinger, D.J. Heinzen, W.M. Itano, S.L. Gilbert, and D.J. Wineland. A 303-MHz frequency standard based on trapped Be^+ ions. *IEEE Transactions on Instrumentation and Measurement*, 40:126–128, 1991.
- [65] C. Langer et al. Long-lived qubit memory using atomic ions. *Physical Review Letters*, 95:060502, 2005.
- [66] P.O. Schmidt, T. Rosenband, C. Langer, W.M. Itano, J.C. Bergquist, and D.J. Wineland. Spectroscopy using quantum logic. *Science*, 309:749–752, 2005.
- [67] P. Bove, L. Hornekær, C. Brodersen, M. Drewsen, J.S. Hangst, and J.P. Schiffer. Sympathetic crystallization of trapped ions. *Physical Review Letters*, 82:2071–2074, 1999.
- [68] 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. *Physical Review A*, 60:3903–3910, 1999.
- [69] C.E. Pearson, D.R. Leibbrandt, W.S. Bakr, W.J. Mallard, K.R. Brown, and I.K. Chuang. Experimental investigation of planar ion traps. *Physical Review A*, 73:032307, 2006.
- [70] D. Cruz, J.P. Chang, and M.G. Blain. Field emission characteristics of a tungsten microelectromechanical system device. *Applied Physics Letters*, 86:153502, 2005.
- [71] L. Deslauriers, S. Olmschenk, D. Stick, W.K. Hensinger, J. Sterk, and C. Monroe. Scaling and suppression of anomalous quantum decoherence in ion traps. *quant-ph/0602003*, 2006.

BIBLIOGRAPHY

- [72] J. Kim et al. 1100 × 1100 port mems-based optical crossconnect with 4-dB maximum loss. *IEEE Photonics Technology Letters*, 15:1537–1539, 2003.
- [73] R.S. Van Dyck Jr., D.L. Farnham, and P.B. Schwinberg. A compensated Penning trap mass spectrometer and the ^3H – ^3He mass difference. *Journal of Modern Optics*, 39:243–255, 1992.
- [74] T. Beier, H. Häffner, N. Hermanspahn, S.G. Karshenboim, H.-J. Kluge, W. Quint, S. Stahl, J. Verdú, and G. Werth. New determination of the electron’s mass. *Physical Review Letters*, 88:011603, 2002.
- [75] I. Bergström, M. Björkhage, K. Blaum, H. Bluhme, T. Fritioff, S. Nagy, and R. Schuch. High-precision mass measurements of hydrogen-like $^{24}\text{Mg}^{11+}$ and $^{26}\text{Mg}^{11+}$ ions in a Penning trap. *European Physical Journal D*, 22:41–45, 2003.
- [76] M. Vogel, D.F.A. Winters, D.M. Segal, and R.C. Thompson. Proposed precision laser spectrometer for trapped, highly charged ions. *Review of Scientific Instruments*, 76:103102, 2005.
- [77] G. Gabrielse and F.C. MacKintosh. Cylindrical Penning traps with orthogonalized anharmonicity compensation. *International Journal of Mass Spectrometry and Ion Processes*, 57:1–17, 1984.
- [78] G. Gabrielse, L. Haarsma, and S.L. Rolston. Open-endcap Penning traps for high precision experiments. *International Journal of Mass Spectrometry and Ion Processes*, 88:319–332, 1989.
- [79] S. Stahl, F. Galve, J. Alonso, S. Djekic, W. Quint, T. Valenzuela, J. Verdú, M. Vogel, and G. Werth. A planar Penning trap. *European Physical Journal D*, 32:139–146, 2005.
- [80] J.R. Castrejón-Pita and R.C. Thompson. Proposal for a planar Penning ion trap. *Physical Review A*, 72:013405, 2005.
- [81] W.M. Itano and D.J. Wineland. Laser cooling of ions stored in harmonic and Penning traps. *Physical Review A*, 25:35–54, 1982.

BIBLIOGRAPHY

- [82] W.M. Itano, L.R. Brewer, D.J. Larson, and D.J. Wineland. Perpendicular laser cooling of a rotating ion plasma in a Penning trap. *Physical Review A*, 38:5698–5706, 1988.
- [83] G.Z.K. Horvath and R.C. Thompson. Laser cooling of ions stored in a Penning trap: A phase-space picture. *Physical Review A*, 59:4530–4546, 1999.
- [84] R.C. Thompson and J. Papadimitrou. Simple model for the laser cooling of an ion in a Penning trap. *Journal of Physics B*, 33:3393–3405, 2000.
- [85] D. Wineland and H.G. Dehmelt. Line shifts and widths of axial, cyclotron and G-2 resonances in tailored, stored electron (ion) cloud. *International Journal of Mass Spectrometry and Ion Physics*, 16:338–342, 1975.
- [86] D. Wineland and H.G. Dehmelt. Line shifts and widths of axial, cyclotron and G-2 resonances in tailored, stored electron (ion) cloud (erratum). *International Journal of Mass Spectrometry and Ion Physics*, 19:251, 1976.
- [87] R.S. Van Dyck Jr., P.B. Schwinberg, and H.G. Dehmelt. Electron magnetic moment from geonium spectra. In B. Kursunoglu, A. Perlmutter, and L. Scott, editors, *New Frontiers in High Energy Physics*, page 158. Plenum N.Y., 1978.
- [88] S.L. Brown and G. Gabrielse. Geonium theory: Physics of a single electron or ion in a Penning trap. *Review of Modern Physics*, 58:233–311, 1986.
- [89] L. Schweikhard and A.G. Marshall. Excitation modes for Fourier transform–ion cyclotron resonance mass spectrometry. *Journal of the American Society of Mass Spectrometry*, 4:433–452, 1993.
- [90] G. Savard, S. Becker, G. Bollen, H.-J. Kluge, R.B. Moore, T. Otto, L. Schweikhard, H. Stolzenberg, and U. Wiess. A new cooling technique for heavy ions in a Penning trap. *Physics Letters A*, 158:247–252, 1991.
- [91] H.F. Powell, D.M. Segal, and R.C. Thompson. Axialisation of laser cooled magnesium ions in a Penning trap. *Physical Review Letters*, 89:93003, 2002.

BIBLIOGRAPHY

- [92] 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 axialisation technique. *Journal of Physics B*, 36:961–970, 2003.
- [93] E.A. Cornell, R.M. Weisskoff, K.R. Boyce, and D.E. Pritchard. Mode-coupling in a Penning trap: π -pulses and a classical avoided crossing. *Physical Review A*, 41:312–315, 1990.
- [94] NIST Atomic Spectra Database.
<http://physics.nist.gov/physrefdata/asd/index.html>.
- [95] M.E.M. Storkey. *Studies of laser-cooled trapped ions*. PhD thesis, Imperial College, London, 2001.
- [96] E.S.G. Phillips. *Controlled dynamics of laser-cooled ions in a Penning trap*. PhD thesis, Imperial College, London, 2004.
- [97] J.L.K. Koo. *Laser cooling and trapping of Ca^+ ions in a Penning trap*. PhD thesis, Imperial College, London, 2003.
- [98] A.S. Arnold, J.S. Wilson, and M.G. Boshier. A simple extended-cavity diode laser. *Review of Scientific Instruments*, 69:1236–1239, 1998.
- [99] G. Blasbichler. Ein Lambdameter mit 10^{-7} Meßunsicherheit. Master’s thesis, Institut für Experimentalphysik, Universität Innsbruck, Austria, 2000.
- [100] J. Sudbery. *Studies of laser cooled calcium ions in the Penning and combined traps*. PhD thesis, Imperial College, London, 2003.
- [101] 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. *American Journal of Physics*, 60:1098–1111, 1992.
- [102] M.A. van Eijkelenborg, K. Dholakia, M.E.M. Storkey, D.M. Segal, and R.C. Thompson. A driven, trapped, laser cooled ion cloud: a forced damped oscillator. *Optics Communications*, 159:169–176, 1999.
- [103] S. Djekic et al. Temperature measurement of a single ion in a Penning trap. *European Physical Journal D*, 31:451–457, 2004.

BIBLIOGRAPHY

- [104] N. Kjærgaard, L. Hornekær, A.M. Thommesen, Z. Videsen, and M. Drewsen. Isotope selective loading of an ion trap using resonance-enhanced two-photon ionization. *Applied Physics B*, 71:207–210, 2000.
- [105] S. Gulde, D. Rotter, P. Barton, F. Schmidt-Kaler, R. Blatt, and W. Hogervorst. Simple and efficient photo-ionization loading of ions for precision ion-trapping experiments. *Applied Physics B*, 73:861–863, 2001.
- [106] D.M. Lucas et al. Isotope-selective photoionization for calcium ion trapping. *Physical Review A*, 69(012711), 2004.
- [107] D.J. Berkeland and M.G. Boshier. Destabilization of dark states and optical spectroscopy in Zeeman-degenerate atomic systems. *Physical Review A*, 65:033413, 2002.
- [108] Y. Stalgies, I. Siemers, B. Appasamy, T. Altevogt, and P.E. Toschek. The spectrum of single-atom resonance fluorescence. *Europhysics Letters*, 35:259–264, 1996.
- [109] E. Arimondo. Coherent population trapping in laser spectroscopy. In E. Wolf, editor, *Progress in Optics XXXV*, pages 257–354. Elsevier Science B.V., 1996.
- [110] D. McGloin. *Electromagnetically induced transparency*. PhD thesis, University of St. Andrews, 2000.
- [111] J.L. Hernández-Pozos. *Dynamics of Mg^+ and Be^+ trapped ions and design and construction of a cylindrical Penning trap*. PhD thesis, Imperial College, London, 1999.
- [112] H.F. Powell. *Quantum optics with a single trapped ion*. PhD thesis, Imperial College, London, 2002.