



The Imperial College of Science, Technology and Medicine
Department of Physics

Manipulation of Ca^+ Ions in Penning Traps

Daniel R. Crick

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics of Imperial College, London,
01/05/2009

I, Daniel R. Crick, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

The long term aim of this work is to study the suitability of using laser cooled Ca^+ ions in Penning traps as the basic components of a quantum computer. A great deal of progress in the field of quantum computing has been made in recent years with laser cooled ions stored in radio frequency ion traps. Building a useful quantum computer with trapped ions is however extremely challenging. Penning traps offer some possible benefits over radio frequency traps. They also create some additional difficulties. The potential advantages and disadvantages of Penning traps are discussed throughout the thesis.

We show that we are able to overcome the problems associated with laser cooling in Penning traps, and have trapped single ions for extended periods of time. Pairs of Ca^+ ions have been aligned along the axis of a Penning trap, and have been optically resolved.

A novel Penning trap array based on PCB boards has been developed. A prototype was built and tested, along with the electronics required to shuttle ions between different sub-traps. Ions have been shuttled a distance of 10 mm in $2.5\ \mu\text{s}$. A return trip efficiency of up to 75% was seen.

A quantum effect – J -state mixing caused by large magnetic fields – has been observed for the first time in single atomic ions. The magnetic field causes a forbidden $\Delta J = 2$ transition to become weakly allowed. This effect is of general interest in atomic physics, and is also very relevant for quantum computation studies. A quantitative prediction of the magnitude of the J -mixing effect has been derived theoretically. This is compared to experimental data, and is found to be in excellent qualitative and good quantitative agreement.

This thesis is dedicated to my parents.

Acknowledgements

First and foremost I owe enormous gratitude to my supervisor Danny Segal. Thanks for the support, the advice, the much needed optimism, the crate of beer after the first single ion, and a lot more. Thanks of course also to Richard Thompson, for great wisdom, for asking the questions which always seemed to lead to progress, for the summer BBQs, and for making some time even when there wasn't any. If I ever find better people to work for I will be very surprised.

I gratefully acknowledge generous financial support from SCALA, without which this work would not have been possible.

Certainly the best aspect of this PhD has been getting the chance to work with such brilliant colleagues over the past few years. My greatest thanks and respect to Bakry, Rafael, Rich, Danyal, Manuel, Sharath, Shailen, Sandeep and all the undergraduate and 6th form students who passed through the lab. Thanks in particular to Hamid, Ianjit and Sean. Why does the experiment only work well after midnight?

Thanks to Sean T for the brilliant headache-inducing conversations about angular momentum and transition rates. Thanks also to all the rest of the 114 crew, who made these few months of writing a lot more fun than they could have been.

I wish to thank Bandu Ratnasekara, for advice on electronics, advice on life, many cups of coffee, and always making time for me. I also have great respect for Brian Willey in the mechanical workshop.

I would like to thank my parents for continuous support, both financial and otherwise.

Also to Katy, a lot of the above, and everything else in between.

*Why why why do I write so slow?
I sit and think while my beard does grow.
Sentences come and paragraphs go,
As I blankly stare at the monitor's glow.*

*Why oh why do I type so slow?
Errors! Warnings! Badboxes! D'oh!
Why use L^AT_EX? Just for show?
So very slowly, the words do flow.*

*Why why why do I write so slow?
I started this thesis such a long time ago.
What's that Sean? Am I busy? No.
Time for a coffee break? Let's go go go!*

Contents

1	Introduction	14
1.1	Motivation and Description of Thesis	14
1.2	Ion Traps	15
1.2.1	RF Trap	16
1.2.2	Penning Trap	19
1.2.3	Combined Trap	22
1.3	Laser Cooling	23
1.3.1	Doppler Cooling	23
1.3.2	Resolved Sideband Cooling	25
1.3.3	Laser Cooling in a Penning Trap	27
1.4	Quantum Computation	28
1.4.1	The Qubit	29
1.4.2	Quantum Algorithms	31
1.4.3	Quantum Gates	33
1.4.4	Cirac-Zoller Gate 1995	36
1.4.5	Cirac-Zoller Gate 2000	38
1.5	Scalable Quantum Computing	39
1.5.1	The DiVincenzo Criteria	39
1.5.2	Multiple Ion Traps	40
2	Experimental Setup	42
2.1	General Outline	42
2.2	Calcium 40 Ions	43
2.2.1	Ion Generation	45
2.3	Laser Systems	48
2.3.1	Extended-Cavity Diode Lasers	48

2.3.2	Wavemeter	53
2.3.3	Reference Cavities	56
2.3.4	Multiple laser locking system	60
2.4	Split Ring Trap	63
2.4.1	Vacuum System	64
2.4.2	Magnet	64
2.4.3	Fluorescence Collection and Detection	67
3	Single Ions and Small Crystals	70
3.1	Loading Individual Ions	70
3.2	Ion Crystals	73
4	PCB Penning Trap	79
4.1	Design and Manufacturing	79
4.1.1	Boards	79
4.1.2	Superstructure and Vacuum Chamber	82
4.2	Simulations	88
4.3	Electronics	96
4.3.1	Output Stage	97
4.3.2	Voltage Sources	99
4.3.3	Control Pulse Demultiplexer	100
4.4	Results	102
4.5	Future Improvements	104
4.5.1	Shim Coils	104
4.5.2	Electronics	105
4.5.3	Prototype Design	106
4.6	Outlook and Comparison With Other Work	107
5	Quantum Jumps and J-Mixing	109
5.1	Notes on the Experimental Setup	112
5.1.1	Amplified Spontaneous Emission in Diode Lasers	112
5.1.2	Combined Trap	114
5.2	Theory	115
5.2.1	Transition Rates	115
5.2.2	Magnetic Perturbation	117
5.2.3	Wigner-Eckart Theorem	120
5.3	Data Analysis	121
5.4	Results	126

6 Discussion	130
A Miscellaneous Electronic Schematics	134
B Script to Generate SIMION Geometry Files	139
C Data Analysis Script	144

List of Tables

1.1	DiVincenzo criteria fulfilled by different technologies	40
2.1	Ca ⁺ Landé g factors	45

List of Figures

1.1	Schematic view of Paul trap and Penning trap	16
1.2	Paul trap stability diagram	18
1.3	Ion trajectory in a Penning trap	21
1.4	Combined trap stability diagram	22
1.5	Resolved sideband cooling	25
1.6	Laser beam offset	28
1.7	Bloch sphere	30
2.1	General schematic of experiment	43
2.2	Ca^+ energy level structure	44
2.3	Optical table schematic	49
2.4	Laser diode gain curve	50
2.5	Littrow configuration	50
2.6	Laser photographs	52
2.7	HeNe laser modes	54
2.8	HeNe laser schematic	55
2.9	Effect of air pressure	57
2.10	Reference cavity vacuum chamber	58
2.11	Wheatstone bridge	59
2.12	Spectrum of four single mode lasers	61
2.13	Split-ring trap	63
2.14	Split-ring trap vacuum chamber	65
2.15	Magnetic field calibration	66
2.16	Fluorescence detection system	67
3.1	Loading a cloud of ions	71

3.2	Loading a pair of ions	72
3.3	Fluorescence level histograms	73
3.4	Two-ion images	75
3.5	Two-ion orientation sequence	78
4.1	PCB trap electrodes	80
4.2	PCB trap board design	81
4.3	PCB trap diagram	82
4.4	PCB trap photos	83
4.5	Potential along trap axis	85
4.6	Quartic coefficient of axial potential	85
4.7	Oven holder	86
4.8	Vacuum chamber	87
4.9	Electrostatic potential in xy plane	89
4.10	Ion shuttling trajectory	90
4.11	Position error vs voltage error	90
4.12	KE error vs timing error	91
4.13	Optimal voltage and timing parameters	93
4.14	Final conditions of ion	94
4.15	Restoring force along z	95
4.16	Generalised shuttling electronics system	96
4.17	Pulse converter output stage	97
4.18	Voltage pulses for various resistances	98
4.19	Voltage source generator (partial)	100
4.20	Control pulse demultiplexer	101
4.21	Repeatedly shuttled ion cloud	102
4.22	Ion shuttling efficiency	104
4.23	Version 2 pulse converter	105
5.1	Single ion quantum jumps	109
5.2	Observation of coherent population trapping	111
5.3	Single ion quantum jumps with lasers filtered	113
5.4	Dark time histogram	114
5.5	Distribution of single ion signal levels	123
5.6	Shelving rate as a function of magnetic field	126
5.7	Normalised shelving rate	128
5.8	$D_{5/2}$ state lifetime	129

A.1	HeNe stabilisation circuit	135
A.2	Hot cathode filament and oven controller schematic	136
A.3	Voltage source generator	137
A.4	Version 2 pulse converter	138

Chapter 1

Introduction

1.1 Motivation and Description of Thesis

A quantum computer could be a very useful device for performing algorithms, such as factoring large numbers, much faster than classical computers. Some very interesting and very difficult problems such as simulating quantum systems could also be done using a quantum computer. Many different physical systems are being considered as the basis for a quantum computer, including laser cooled ions in radio frequency (RF) ion traps. This work considers the use of laser cooled ions in a different type of ion trap – the Penning trap – as the basic elements of a quantum computer. Penning traps do have disadvantages when compared to RF traps, but they may also offer some advantages.

This chapter first provides a very brief introduction to the general experimental techniques used in the main work of the thesis, namely ion trapping and laser cooling. The basic ideas of quantum computation are then discussed. The introduction to quantum computation is first given for a generalised system, but then emphasis is placed on the specifics of laser cooled trapped ions. Finally, the recent hot topic in quantum computation – scalability – is discussed.

In chapter 2, the apparatus used to trap, cool and detect the ions is explained. The whole setup is described, but some of the the upgrades implemented by the author are covered in more detail.

Chapter 3 presents images of some of the first Ca^+ ions ever to be trapped individually and in pairs in a Penning trap. Chapter 3 also shows the manipulation of the alignment of a pair of ions, parallel and perpendicular to the

trap axis.

The details of a prototype Penning trap containing multiple trapping regions are given in chapter 4. The first demonstration of controlled shuttling of ions between different trapping zones, orthogonal to the trap axis, is also presented.

Chapter 5 discusses the phenomenon of J -state mixing in atomic systems caused by strong magnetic fields. A clear observation of the J -mixing effect in single atomic ions is presented for the first time. A theoretical model is also derived. The model predicts a branching ratio between the forbidden $P_{1/2} \rightarrow D_{5/2}$ transition and the strongly allowed $P_{1/2} \rightarrow S_{1/2}$ transition of $4.2 \times 10^{-7} B^2 \text{ tesla}^{-2}$. This is seen to be in fairly good agreement with experimental results.

Chapter 6 then gives a critique of the benefits and disadvantages of using Penning traps as part of a quantum computer, in light of the new work presented in this thesis.

1.2 Ion Traps

Half of the 1989 Nobel prize in physics was awarded to Hans G. Dehmelt and Wolfgang Paul for “the development of the ion trap technique”*. Ion traps have found use in many areas of physics: precision mass measurements; spectroscopy; g factor measurements; physics of highly charged ions; precision tests of quantum mechanics; study of non-neutral plasmas; study of antimatter; radioactive lifetime measurements; study of quantum phase transitions; ion-electron recombination studies; ion-molecule reactions; atomic clocks; quantum computing, and more. For a more thorough introduction to ion trapping, see [1] or [2].

The effect of the electromagnetic force dominates the motion of an ion. Hence by manipulating electromagnetic potentials it is possible to trap ions indefinitely, as long as they are in a good vacuum. Generally, ultra high vacuums on the order of 10^{-8} mbar or lower are needed, while even harder vacuums are desirable (we typically trap at pressures of a few 10^{-10} mbar).

It is not possible to create a minimum in the electric potential at any point in free space along all three dimensions using only static electric fields [3]. This is simple to deduce by considering Laplace’s equation: $\nabla^2 \phi = 0$. If there is a

*The other half was awarded to N. Ramsey for work on molecular beams.

minimum in two dimensions then the potential along the other dimension must be a maximum, and hence a charged particle at this point will not be trapped. To overcome this, a couple of different methods can be used. An oscillating electric field can be applied, which with appropriate parameters averages out over time to give a minimum in all three directions. Alternatively, a magnetic field can be combined with a static electric field to provide three dimensional confinement. These techniques are the basis of the two important types of ion trap*: the Paul (or radio frequency (RF)) trap and the Penning trap. The two trapping techniques can also be used together to form a combined trap, which is discussed briefly in §1.2.3 and §5.1.2.




Figure 1.1: Schematic view of a Paul trap and a Penning trap. A cross section in the xz plane is shown. The traps have cylindrical symmetry about the z axis (vertical). The typical voltage polarities for trapping positive ions are shown.

1.2.1 RF Trap

A Paul trap consists of two endcap electrodes facing each other, with a ring electrode between the two. An oscillating voltage is applied between the ring and endcaps as shown in figure 1.1 (a). If the ring and endcap electrodes take the shape of infinite hyperboloids of revolution[†], then the potential is purely quadratic along all directions and the equations of motion of an ion inside the

*There are other traps such as the electrostatic ion beam trap, and the charged particle storage ring, which can both store ions but do not trap them in a single fixed position.

[†]Obviously a real trap cannot have infinite size electrodes, however many traps have been built which do closely approximate this form.

trap have a relatively simple form:

$$\ddot{r} = - \left(\frac{e}{mr_0^2} \right) (U + V \cos \Omega t) r \quad (1.1)$$

$$\ddot{z} = \left(\frac{2e}{mr_0^2} \right) (U + V \cos \Omega t) z \quad (1.2)$$

where e and m are the ionic charge and mass respectively, and $U + V \cos \Omega t$ is the potential applied to the ring electrode relative to the endcaps.

The motion of ions in the Paul trap consists of oscillations at many different frequencies. There is generally one oscillation frequency along each direction which is lower than the applied RF frequency. The motion at this fundamental frequency is known as the secular motion (or macromotion) of the ion. The higher frequency oscillations are known as micromotion.

If the parameters a and q are defined as

$$a_z = -2a_r = -\frac{8eU}{mr_0^2\Omega^2} \quad (1.3)$$

$$q_z = -2q_r = \frac{4eV}{mr_0^2\Omega^2} \quad (1.4)$$

then the above equations can be transformed into Mathieu equation form*:

$$\frac{d^2 r}{d\tau^2} + (a_r - 2q_r \cos 2\tau) r = 0 \quad (1.5)$$

$$\frac{d^2 z}{d\tau^2} + (a_z - 2q_z \cos 2\tau) z = 0 \quad (1.6)$$

where $\tau = \Omega t/2$.

It can be seen that the solutions are stable (and hence ions can be trapped) for certain regions in parameter space. For a particular set of values for the trap size, RF frequency and amplitude, and static bias voltage, the solutions are stable for a particular range of ion mass-to-charge ratios. The region of stability is the green shaded area enclosed by the lines in figure 1.2. Outside the shaded region the amplitude of the micromotion grows with time. The black lines enclose the region in which the axial ion motion is stable, and the red lines show where the radial motion is stable.

*See [1] for a full derivation.




Figure 1.2: Paul trap stability diagram. The four lines show boundaries in the aq plane, inside which an ion can be stably trapped.

The frequency of the secular motion is

$$\omega_0 = \frac{\beta\Omega}{2} \quad (1.7)$$

where β can be calculated from an approximate equation, an infinite continued fraction, or numerical simulations. An approximate equation for β is [4]

$$\beta = \sqrt{a - \frac{(a-1)q^2}{(2(a-1)^2 - q^2)} - \frac{(5a+7)q^4}{32(a-1)^3(a-4)}} \quad (1.8)$$

which simplifies to

$$\beta \approx \sqrt{a + \frac{q^2}{2}} \quad (1.9)$$

for $|a|, q \ll 1$.

The frequencies of the micromotion oscillations are given by the secular frequency plus the RF drive frequency (and harmonics),

$$\omega_n = \left(n \pm \frac{\beta}{2}\right) \Omega \quad (1.10)$$

where n is an integer.

The trap with hyperboloid electrodes is but one possible geometry of an RF trap. The geometry most studied for quantum computation is based on the

same principle as the ideal Paul trap, but consists instead of four long, parallel, rod shaped electrodes. Opposite pairs of electrodes are connected together and driven with the RF voltage. This traps ions tightly in two dimensions. Endcap electrodes with a positive DC bias are used to provide trapping in the axial direction. In this geometry, ions have a high secular motional frequency in the direction perpendicular to the electrodes, and a much lower one along the axis of the trap. If a small number of ions are trapped, they tend to become spaced out in a chain-like structure along the axis of the trap. When just two ions are held in a linear RF trap there are two modes of oscillation along the axis of the trap: a mode where the centre of mass oscillates, and a mode where the distance between the two ions oscillates (stretch mode, or breathing mode). A linear string of N ions will have N modes along the trap axis, and another N modes along each of the other two dimensions of the trap.

At extremely low temperatures the motion of the ions is clearly quantised. Since the modes are of the entire ion chain, a group of ions held in a trap like this are correlated and may become entangled with each other via these common modes of motion. Quantum gates can be performed between two or more ions, using this common motion as a means of interaction between the ions.

1.2.2 Penning Trap

The ideal Penning trap has an identical structure to the ideal Paul trap, except that the Penning trap does not use an RF field, but instead uses a uniform magnetic field, B , along the axis of the trap. For positive ions, the endcap electrodes are biased with a positive potential relative to the ring. The ions are then trapped in a potential minimum along the z direction. There is a potential maximum in the radial plane, but as ions move away from the axis they are pulled round by the magnetic field in a cyclotron orbit and are thus confined in the radial direction. The combination of magnetic and electric fields also introduces a second kind of motion in the radial plane, known as magnetron motion. The motion of an ion in a Penning trap is a harmonic oscillation along the z direction, and a precession around the z axis consisting of a superposition of cyclotron and magnetron motion.

In an ideal Penning trap the field Lagrangian can be solved* to give equa-

*See [1] for a derivation.

tions of motion:

$$\ddot{x} = \omega_c \dot{y} + \frac{1}{2} \omega_0^2 x \quad (1.11)$$

$$\ddot{y} = -\omega_c \dot{x} + \frac{1}{2} \omega_0^2 y \quad (1.12)$$

$$\ddot{z} = -\omega_0^2 z. \quad (1.13)$$

Equation 1.13 can be very easily solved to show that the motion along z is a harmonic oscillation at frequency ω_0 . The axial frequency, ω_0 , and the true cyclotron frequency, ω_c , can be shown to be [1]

$$\omega_0 = \sqrt{\frac{4eU}{m(2z_0^2 + r_0^2)}}, \quad \omega_c = \frac{eB}{m} \quad (1.14)$$

where z_0 and r_0 are the axial and radial dimensions of the trap respectively.

A characteristic length R_0 can be defined such that

$$\omega_0 = \sqrt{\frac{4eU}{mR_0^2}}. \quad (1.15)$$

For an ideal trap, $R_0^2 = 2z_0^2 + r_0^2$, but any non-ideal trap will still have the characteristic dimension R_0 defined by the quadratic component of the axial potential.

The x and y motion cannot be solved quite as easily as the axial motion. The position in the xy plane can be made equivalent to the position of a complex number in the Argand plane by defining a new variable $u = x + iy$. Equations 1.11 and 1.12 then combine to give an equation of radial motion:

$$\ddot{u} = -i\omega_c \dot{u} + \frac{1}{2} \omega_0^2 u. \quad (1.16)$$

Letting $u = u_0 \exp(-i\omega t)$ yields a quadratic equation in ω with the roots:

$$\begin{aligned} \omega_{\pm} &= \frac{1}{2} \left(\omega_c \pm \sqrt{\omega_c^2 - 2\omega_0^2} \right) \\ &= \frac{1}{2} (\omega_c \pm \omega_1) \quad \text{where } \omega_1 \equiv \sqrt{\omega_c^2 - 2\omega_0^2}. \end{aligned} \quad (1.17)$$

Thus there are two distinct frequencies of the radial motion:

$$\omega_+ = \frac{1}{2} (\omega_c + \omega_1) \quad (1.18)$$

which is known as the modified cyclotron frequency, and

$$\omega_- = \frac{1}{2}(\omega_c - \omega_1) \quad (1.19)$$

is the magnetron frequency.

An example of the motion of an ion in the radial plane of a Penning trap is shown in figure 1.3. In this example the orbit size of the modified cyclotron motion is smaller than that of the magnetron motion. The axial motion (orthogonal to the radial plane) is simple harmonic motion.




Figure 1.3: Example of the trajectory of an ion in the radial plane of a Penning trap. The magnetic field is normal to the page. The small loops correspond to the modified cyclotron motion, while the slow drift around the centre of the trap is the magnetron motion.

These three different modes of motion with three different frequencies are uncoupled. Usually, $\omega_+ \approx \omega_c \gg \omega_0 \gg \omega_-$. If ω_1 is imaginary, then the ion motion is not stable. Therefore there is an upper limit on the axial trapping frequency, determined by the magnetic field strength:

$$\omega_0 \leq \frac{\omega_c}{\sqrt{2}}. \quad (1.20)$$

Some other useful relations between the various frequencies can also be

derived:

$$\omega_+ + \omega_- = \omega_c \quad (1.21)$$

$$\omega_0^2 + \omega_+^2 + \omega_-^2 = \omega_c^2. \quad (1.22)$$

Thus if ω_+ and ω_- are both measured, ω_c (and hence the magnitude of B) can be determined.

The amplitude of the axial motion and the size of the modified cyclotron orbit both decrease as the energy of each mode decreases. However, the size of the magnetron orbit increases as the energy of the magnetron mode is reduced. This makes the magnetron motion fundamentally unstable whenever a damping mechanism (such as laser cooling) is present.

1.2.3 Combined Trap

A combined ion trap is a trap to which the fields of both the RF and the Penning trap are applied. Adding a uniform magnetic field to a Paul trap increases the stability of ions in the radial plane, and alters the radial frequencies, but does not affect motion in the axial direction [4].

Figure 1.4 shows how the Mathieu stability diagram of an RF trap is altered as a magnetic field is applied. The effect is to increase a_r by an amount $(\frac{\omega_c}{\Omega})^2$, shifting the stability boundary* and increasing the range of stable parameters.




Figure 1.4: Combined trap stability diagram. The a_r boundary is shifted in comparison to the pure RF trap. The $(\frac{\omega_c}{\Omega})^2$ label is scaled to the a_r axis.

*The boundary is moved down (more negative) on the a_r scale, which is up on figure 1.4.

The motional frequencies in the radial plane are altered from the pure RF trap frequencies to become [4]

$$\omega_n^r = \left(n \pm \frac{\beta_r}{2} \pm \frac{\omega_c^2}{2\Omega^2} \right) \Omega \quad (1.23)$$

while the axial motion is unaffected. Note that β_r is also dependent on B , so the magnetic field dependence of ω_n^r is not simply linear or quadratic.

E. Fischer demonstrated the mass selective behaviour of the combined trap 50 years ago [5]. Since then very few experiments have made use of the combined trap. The experiment presented in chapter 5 is one of them.

1.3 Laser Cooling

1.3.1 Doppler Cooling

Laser cooling was first proposed in 1975 by Hänsch and Schawlow [6] and by Wineland and Dehmelt [7]. It relies on the radiation pressure of light, and the Doppler effect. Ions are slowed down (and thus cooled) as they move towards a laser beam, but when they move away from the beam they interact with it less strongly due to the Doppler shift.

If a laser beam is passed into a trap, it can interact with ions which are trapped. The ions move around in the trap, so sometimes they are moving towards the laser source, and sometimes away from it. If the frequency of the laser light is ω_0 in the lab frame, then in the frame of reference of an ion it will be*

$$\omega = \omega_0 - \mathbf{k}_0 \cdot \mathbf{v}, \quad (1.24)$$

where $\mathbf{k}_0$ is the wavevector of the laser light, and $\mathbf{v}$ is the velocity of the ion in the lab frame.

If ω_0 is chosen so that ω will be resonant with a transition of the ion when the ion has a particular velocity, then the ion will absorb light most strongly when it has this velocity. When an ion absorbs a photon, it absorbs momentum and is pushed back. The force which acts on the ion is the rate at which photons are scattered multiplied by the momentum carried by each

*Assuming the ion is non-relativistic.

photon:

$$\mathbf{F} = \gamma \hbar \mathbf{k}_0 \quad (1.25)$$

where the scattering rate, γ , is a function of the laser intensity, the Doppler shift and the rate of spontaneous decay (Γ). For a two level ion,

$$\gamma = \frac{\Gamma}{2} \frac{I/I_{\text{sat}}}{\left(\frac{2(\Delta\omega - \mathbf{k}_0 \cdot \mathbf{v})}{\Gamma} \right)^2 + I/I_{\text{sat}} + 1} \quad (1.26)$$

where $\Delta\omega$ is the laser detuning and I_{sat} is the saturation intensity [8].

If the laser intensity is low, the population of the upper state tends to zero. If the intensity is very high, the population is shared equally between the two states. The saturation intensity is defined as the amount of light incident on a two level ion which causes the average population of the upper level to be half the population of the lower level. It can be related to the spontaneous decay rate and the wavelength of the light:

$$I_{\text{sat}} = \frac{hc\pi\Gamma}{3\lambda^3}. \quad (1.27)$$

For an ion with more than two levels involved in the laser cooling cycle, the situation is more complicated. The scattering rate then depends on several laser intensities and detunings. It can also depend strongly on the relative detuning of two or more laser frequencies.

After the ion has absorbed a photon, it will spontaneously decay back to the ground state* emitting another photon. The emitted photon will go in a random direction, so the momentum change due to this is averaged out over several absorption/emission processes. For a typical transition in a trapped ion, the laser cooling force is $\sim 10^{-20}$ N, compared to a typical trapping force of $\sim 10^{-16}$ N and a gravitational force of $\sim 10^{-24}$ N.

There is a limit on how cold an ion or ion cloud can be made using the Doppler cooling technique. Suppose an ion is brought almost to rest when it absorbs a photon. It will soon spontaneously emit another photon in a random direction, the recoil of which will increase the ion's motion. As this process is repeated it leads to Brownian motion of the ion and hence a temperature increase. When this effect balances out the cooling effect of the laser, no further cooling is possible. Assuming a detuning of $\Delta\omega = \Gamma/2$, this lower

*Or another state of lower energy than the excited state.

limit is [8]

$$k_B T_{\text{limit}} = \frac{\hbar \Gamma}{2}. \quad (1.28)$$

For a typical cooling transition, the Doppler cooling limit temperature is ~ 1 mK (corresponding to an ion velocity of ~ 1 m/s). This temperature is too high for some quantum gates (§1.4.4) to work because the energy of an ion at the Doppler limit is generally much greater than that of a phonon of the ion's vibrational motion. Thus the ion's motion is still classical rather than quantum in nature. The Doppler limit can be overcome using a technique known as resolved sideband cooling.

1.3.2 Resolved Sideband Cooling




Figure 1.5: Quantised energy level diagram of a two-level ion in a harmonic trapping potential.

The energy of an ion in a harmonic trap has two contributions: the internal electronic energy (E_{elec}), and the vibrational energy (E_{vib}). The energy level structure of such an ion is shown in figure 1.5. The state of an ion in this system can be labelled $|a, n\rangle$ where $a = g$ or e for the ground or excited state respectively, and n is the number of phonons. To perform sideband cooling, the different transition frequencies (ω , ω_{red} , ω_{blue} etc.) must be resolved. However, as mentioned above, the linewidth of the Doppler cooling transition is usually broader than the energy of a vibrational phonon, so a sideband cooling transition will usually use different states to those used for Doppler cooling. The transition must be very narrow, so typically a forbidden transition such as an electric quadrupole transition, or a two-photon Raman transition is used.

If the Doppler cooling limit is achieved, then an ion will have $E_{\text{vib}} \approx \hbar\Gamma/2$ and $E_{\text{elec}} = 0$. Since Γ is much larger than the frequency of vibrational motion of a trapped ion (ω_{vib}), the average number of phonons, $\langle n \rangle$, in the ion's motion is much larger than unity.

To perform resolved sideband cooling, a narrow linewidth laser* with angular frequency ω_{red} (first red sideband) is applied. This will induce the transition $|g, n\rangle \rightarrow |e, n-1\rangle$, shifting the phonon number down by one. The excited state has a long lifetime so if nothing further is done the ion will just remain in the excited state. To shorten the lifetime, an additional laser is applied. This laser dresses the states and increases the effective decay rate. The dressing laser should be slightly detuned[†] from a transition between a third auxiliary state and either $|g, n\rangle$ or $|e, n\rangle$. For this to work there must be a dipole allowed transition between the third state and both the ground and excited states [9].

When the ion decays, it could go to $|g, n-1\rangle$, $|g, n\rangle$, $|g, n+1\rangle$ etc. The relative decay rates depend on the Lamb-Dicke parameter, η , which is the ratio between the spatial extent of the ion's motional ground state and the wavelength of the light.

$$\eta = \frac{2\pi z_{\text{rms}}}{\lambda} \quad (1.29)$$

$$\text{where } z_{\text{rms}} = \sqrt{\langle a, 0 | z^2 | a, 0 \rangle}. \quad (1.30)$$

If $\eta \ll 1$ then the spontaneous decay rate on the carrier transition is much higher than the decay rate on any of the sidebands. Therefore the cooling transition cycle goes as

$$\begin{array}{ccccc} \hbar\omega_{\text{red}} \text{ absorption} & & \text{Spontaneous emission} & & \\ |g, n\rangle & \longrightarrow & |e, n-1\rangle & \longrightarrow & |g, n-1\rangle. \end{array}$$

The net effect is to decrease n , shifting the state to the left in figure 1.5. When the ion reaches $|g, 0\rangle$, light on the red sideband will no longer interact with the ion[‡], so (ignoring heating) the ion should end up in the $n = 0$ vibrational and electronic ground state, with energy $\frac{1}{2}\hbar\omega_{\text{vib}}$. In reality, effects such as off-resonant transitions, RF heating, etc., prevent a 100% pure $n = 0$ state from being achieved.

*The linewidth must be small compared to ω_{vib} .

[†]The auxiliary state should not be populated during the cooling process. It will become populated if the dressing laser is on resonance.

[‡] $|e, -1\rangle$ does not exist.

Although this technique was first demonstrated as early as 1988 [10], and has since been used to cool ions such that $\langle n \rangle$, is measured to be zero more than 99% of the time [11], resolved sideband cooling has not yet been achieved in a Penning trap.

If two ions are held in a trap, then vibrational phonons must be considered to belong to one of the modes of the complete two-ion system. The pair of ions share vibrational energy and cannot be treated as two separate systems. This shared motion can be used to carry quantum information between the qubits and so two-qubit gates can be realised. Since demonstrating the alignment of two ions along the axis of a Penning trap [12], we hope in the future to entangle pairs of ions in this way using the axial mode of motion.

1.3.3 Laser Cooling in a Penning Trap

Laser cooling of ions in a Penning trap is complicated by the unstable motion in the radial plane. If energy is removed from the axial and cyclotron modes, the amplitudes of these motions shrink. On the other hand, if energy is removed from the magnetron mode the ion orbit expands as the ions move down the sides of the potential maximum in the radial plane. In order for laser cooling to improve three dimensional confinement, energy must simultaneously be removed from the axial and cyclotron motion and pumped into the magnetron motion.

Very soon after the first demonstrations of laser cooling in other systems such a strategy was devised for ions in a Penning trap [13]. It involves positioning the red detuned laser beam slightly offset from the centre of the trap. If the beam is offset to the correct side of the trap, scattering occurs when the laser opposes the cyclotron motion but is in the same direction as the magnetron motion, as shown schematically in figure 1.6. This results in the amplitude of both motions being reduced. This strategy is effective, but the radial confinement is limited by the size of the laser beam focus, and the resulting motion is never as small as it can be for ions in an RF trap.

A second strategy – axialisation – is also available if the trap is equipped with a ring electrode split into segments. When a weak oscillating radial quadrupole potential is applied across the ring at the true cyclotron frequency $\omega_c = \omega_+ + \omega_-$, the effect is to couple together the two otherwise independent radial modes. In the absence of damping, the result is that energy is periodically exchanged between these two motions. In the split-ring trap (see §2.4) at




Figure 1.6: The cooling laser must be offset from the centre of a trap in order to heat the magnetron motion while cooling the cyclotron motion.

1 tesla, simulations show that the frequency of this energy exchange is approximately 10 Hz per mV of axialisation drive voltage. If a source of damping is present the overall orbit size gradually shrinks [14]. This is essentially the same as a process known as ‘sideband cooling’* or ‘entropy reduction by motional sideband excitation’ first performed by the group of Dehmelt in 1976 [15], with the exception that in this case the damping is provided by laser cooling instead of buffer gas collisions.

Details of theoretical and experimental treatments of the axialisation process can be found in several places [14, 16, 17], so no more details will be given here. The results of chapter 3 demonstrate the implementation of this technique to align a pair of ions along the magnetic field of the trap.

1.4 Quantum Computation

“Quantum computation is one of the greatest challenges facing experimental physics. Going to the moon is nothing compared with it. It is also a very beautiful area of study because it appears to involve practically the whole of physics and it stretches the

*This is not the same as the ‘resolved sideband cooling’ process mentioned in §1.3.2.

theoretical and experimental resources of every branch of physics.
It's cool in that way.” – D. Deutsch [18]

A quantum computer (QC) is a device which relies on coherent quantum mechanical processes to perform computations. A QC is fundamentally different from a classical computer, and will perform some tasks much more efficiently than any possible classical computer. This section gives a very brief introduction to the theory of the QC; reports on the progress and possibilities of using various quantum systems to build a real QC; and discusses what currently appears to be the most viable system for building a QC: laser cooled trapped ions.

1.4.1 The Qubit

The ideas of quantum computation were first explored in the 1970s and 1980s by physicists and computer scientists such as C. Bennett [19], P. Benioff [20], D. Deutsch [21], and R. Feynman [22]. The main idea is that while a classical bit of information can be in just two states – on or off, 1 or 0, true or false, high voltage or low voltage, etc. – a quantum bit (qubit) can be in a state which is a quantum superposition of two orthogonal eigenstates.

$$\begin{aligned} |\psi\rangle &= a|0\rangle + b|1\rangle \\ &= \left(\cos \frac{\theta}{2} |0\rangle + e^{i\phi} \sin \frac{\theta}{2} |1\rangle \right) e^{i\gamma}. \end{aligned} \quad (1.31)$$

A qubit contains more information than a classical bit, since there are more possible states which it can be in. In fact, to describe a pure qubit fully, three continuous real numbers are required*. For a single qubit, $e^{i\gamma}$ is a global phase factor which can never be measured. In this case the qubit can be defined with just two real numbers.

$$|\psi\rangle = \cos \frac{\theta}{2} |0\rangle + e^{i\phi} \sin \frac{\theta}{2} |1\rangle. \quad (1.32)$$

If these two numbers are written as angles then the qubit can be visualised as a point on the surface of a sphere of unit radius (figure 1.7). This is known as the Bloch sphere (or the Poincaré sphere if the qubit is encoded as the

*The amplitudes a and b of equation (1.31) are complex, so one might think that four numbers are required. However, the normalisation condition ($|a|^2 + |b|^2 = 1$) reduces the number of parameters by one.

polarisation of a photon). Note that the state of a qubit is usually written without the $e^{i\gamma}$ factor, but this phase is important when dealing with multi-qubit gates (as described below).




Figure 1.7: Bloch sphere representation of a qubit.

If a superposition collapses due to decoherence, but the resulting state is still unknown, then the state is described by classical probabilities instead of quantum amplitudes. This is known as a mixture. A qubit which is not in a purely coherent quantum superposition state, but partly in a classical mixture, can also be visualised with the Bloch sphere. It is a point inside the volume of the sphere instead of on the surface. To characterise qubits fully (allowing for mixtures as well as pure superpositions) a density matrix is often used. For a single qubit such as the state in equation (1.31), the density matrix would be

$$\rho = |\psi\rangle\langle\psi|. \quad (1.33)$$

This is a 2×2 matrix, with $|a|^2$ and $|b|^2$ as diagonal elements. The off-diagonal elements (a^*b and ab^*) quantify the coherence, and are zero for a purely classical mixture. Arguably, all quantum systems have some part mixture and some part superposition. In fact it has been said that

“There is no such thing as a pure state. They are simply a collective delusion of theorists.” – A. White [23].

Only in an ideal quantum computer is decoherence negligible, and qubits can be approximated as pure states.

It may seem that a single qubit could be used to store an infinite amount of information (an infinite number of bits are required to characterise a continuous number). However, even though a qubit may contain an infinite amount of information, this information cannot be retrieved. Upon a measurement, the wavefunction will collapse to either of the eigenstates of the measurement. If the $|\psi\rangle$ of equation (1.31) was measured and found to be in state $|0\rangle$ then the only knowledge gained is that $|a| \neq 0$. The information contained in the qubit would be accessible if the qubit could be copied and measured many times, but the no-cloning theorem states that this is not possible [24].

To make the most use of qubits, a QC must contain many of them. A multiple qubit system is known as a quantum register. The parts of a quantum register state corresponding to each qubit are formally combined using the Kronecker (tensor) product, $\otimes$. If there were, for example, three qubits in a quantum system, then the state of that system would be defined by $2^3 = 8$ amplitudes.

$$|\psi\rangle = a_0 |0\rangle \otimes |0\rangle \otimes |0\rangle + a_1 |0\rangle \otimes |0\rangle \otimes |1\rangle + \dots + a_7 |1\rangle \otimes |1\rangle \otimes |1\rangle. \quad (1.34)$$

Usually, the Kronecker product is not written, and simply implied.

$$|\psi\rangle = a_0 |000\rangle + a_1 |001\rangle + \dots + a_7 |111\rangle. \quad (1.35)$$

If a computational operation is performed on these qubits, that operation is done on all the terms in the wavefunction at the same time. According to D. Deutsch, this means the computation is performed simultaneously in eight parallel universes. A classical computer would need eight CPUs to perform the operation in a similar way. If the number of qubits is increased to say, $N = 300$, then the number of terms in the wavefunction goes up to $\sim 10^{90}$. Since there are only on the order of 10^{80} atoms in the universe [24], it is clear that a classical computer cannot hope to compete with a quantum computer in this kind of comparison.

1.4.2 Quantum Algorithms

The advantage of the QC is not as simple as the ability to perform 2^N operations simultaneously. At some point the quantum state of the system has to

be measured, and then the wavefunction collapses and only one final solution is left. Even if this was not the case, the problem would remain of searching for something interesting through 2^N solutions of 2^N problems. The trick to performing a potentially useful quantum computation is to make use of the relative phases between the different eigenstates of the quantum register, and achieve some kind of interference between them.

An outline of the general procedure of a standard quantum algorithm is as follows:

1. Set the state of a group of qubits (quantum register) to a superposition state where each eigenstate has the same amplitude.
2. Perform a series of quantum gates (see §1.4.3.) on the qubits. Each gate need only involve one or two qubits.
3. By performing more gates, manipulate the system in such a way that the phase difference between different parts of the wavefunction becomes an amplitude difference.
4. Measure some of the qubits, and use the information gained in the measurement to deduce something useful.

A famous example of a quantum algorithm was proposed by Peter Shor in 1994 [25]. It makes use of the quantum Fourier transform to find the prime factors of large numbers. This is a very interesting idea, because most modern cryptography protocols are based on the fact that it is very computer-intensive to factorise large numbers. The number of quantum computational steps needed to factor a large number is exponentially smaller than the number of classical computational steps which would be required. In 2001, a group at IBM implemented Shor's algorithm in a seven qubit NMR based QC, to factor the number fifteen* [26].

A QC might also be used to simulate another quantum system. A QC based on a large array of qubits can behave very similarly to solid state systems. If each qubit in the system can be measured, and the interaction between qubits can be controlled at will, then a quantum simulator could be very useful even if it cannot perform general quantum algorithms. Large crystals of laser cooled trapped ions have been proposed as a system for performing quantum simulations [27].

*Previous to this, it was only suspected that $15 = 3 \times 5$.

Another idea called cluster-state quantum computation has been proposed. This relies on having a large number of entangled qubits, on which a series of measurements is then made [28]. The sequence of measurements itself determines the path of the computation to be performed, eliminating the need for quantum gates.

Also, algorithms have been developed for solving the Schrödinger equation in a Hilbert space too large for a classical computer to cope with. The development of quantum algorithms is currently an active field with much theoretical work being produced. See reference [24] and references therein for more information.

1.4.3 Quantum Gates

Quantum gates are the quantum equivalent of classical boolean logic gates such as AND and NOT. Many different quantum gates have been studied theoretically. There have been many proposals for the implementation of multi-qubit quantum gates in trapped ion systems [27, 29, 30, 31, 32], and several have been achieved experimentally [11, 33, 34].

A single qubit gate can be seen as a rotation (or series of rotations) around the Bloch sphere. For example, the quantum NOT gate rotates the qubit by π radians around the x axis of the Bloch sphere. For this reason the operator is often written as $\hat{X}$. In reality (assuming the qubit is encoded as a pair of energy levels in a trapped ion) it is performed by applying a ‘ π pulse’ of a laser resonant with the qubit transition of the ion. For a qubit in one of the eigenstates of the z axis projection measurement (i.e. $|0\rangle$ or $|1\rangle$), this gate acts the same as its classical equivalent: converting 0 to 1 and vice versa. Unlike the classical gate, the quantum gate can be used on superposition states. Flipping the $|0\rangle$ and the $|1\rangle$ is equivalent to the $\hat{X}$ rotation (if the global phase of the state is ignored).

$$\hat{X} \left(\cos \frac{\theta}{2} |0\rangle + e^{i\phi} \sin \frac{\theta}{2} |1\rangle \right) = \cos \frac{\theta}{2} |1\rangle + e^{i\phi} \sin \frac{\theta}{2} |0\rangle \quad (1.36)$$

$$= \left(\cos \frac{\theta + \pi}{2} |0\rangle + e^{i(\pi - \phi)} \sin \frac{\theta + \pi}{2} |1\rangle \right) e^{i\phi}. \quad (1.37)$$

If the system contains only one qubit, then this global phase can never be measured and can be ignored. However, if there are several qubits then a

vector on the Bloch sphere for each is not enough to specify the full state. Rotating θ by 2π puts the Bloch vector back to the same position, but causes a phase change of the qubit wavefunction:

$$\hat{X}\hat{X}|\psi\rangle = -|\psi\rangle. \quad (1.38)$$

A rotation of 4π is needed to completely return the state to its original form. Generally, a rotation of φ about x causes a global phase change of $e^{-i\varphi/2}$.

There are other quantum gates which have no classical analogue. For instance, the $\hat{Z}$ gate (which rotates a qubit by π around the z axis of the Bloch sphere), or the Hadamard gate.

A Hadamard gate, usually denoted by $\hat{H}$, will convert a qubit in state $|0\rangle$ to an equal superposition of the two eigenstates: $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. Similarly, a Hadamard gate operating on $|1\rangle$ will yield an equal superposition, but with an important phase difference: $\hat{H}|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. Classically, the Hadamard gate would appear to simply scramble the state. If a measurement of the state is made after the application of a Hadamard then the measurement will randomly give either 0 or 1 no matter what the original state was. However, if two Hadamards are applied in succession, the qubit will return to its original state. Another useful fact is that if two $\hat{H}$ gates are applied to a qubit, with a $\hat{Z}$ (or $\hat{X}$) gate applied in between, the resulting combination is equivalent to a single $\hat{X}$ (or $\hat{Z}$) gate.

$$|\psi\rangle = a|0\rangle + b|1\rangle \quad (1.39)$$

$$\hat{Z}|\psi\rangle = a|0\rangle - b|1\rangle \quad (1.40)$$

$$\begin{aligned} \hat{H}|\psi\rangle &= \frac{1}{\sqrt{2}}a(|0\rangle + |1\rangle) + \frac{1}{\sqrt{2}}b(|0\rangle - |1\rangle) \\ &= \frac{1}{\sqrt{2}}((a+b)|0\rangle + (a-b)|1\rangle) \end{aligned} \quad (1.41)$$

$$\begin{aligned} \hat{H}\hat{Z}\hat{H}|\psi\rangle &= \hat{H}\hat{Z}\frac{1}{\sqrt{2}}((a+b)|0\rangle + (a-b)|1\rangle) \\ &= \hat{H}\frac{1}{\sqrt{2}}((a+b)|0\rangle - (a-b)|1\rangle) \\ &= b|0\rangle + a|1\rangle = \hat{X}|\psi\rangle. \end{aligned} \quad (1.42)$$

The state in equation (1.40), which is only dependent on a phase difference (in the measurement basis) to distinguish it from $|\psi\rangle$, can thus be measured by

making use of the $\hat{H}$ gate. The Hadamard gate can be thought of as rotating the qubit about the y axis of the Bloch sphere by $\frac{\pi}{2}$, and then about the x axis by π . The idea of rotating the state before measuring it is required to make use of the additional information present in a quantum register. If several similar copies of the same qubit are available, then the full density matrix (including off-diagonal elements) can be found by measuring the qubits immediately after applying various rotations. This process is known as quantum state tomography, and can be used to characterise the entanglement of whole quantum registers [35, 36].

If a qubit is encoded using two energy levels of an atomic ion, with energies E_0 and E_1 , then the Bloch vector continuously rotates about the z axis at a frequency of $\omega = (E_1 - E_0)/\hbar$. The qubit transition is addressed using a laser* on resonance with the qubit transition. Therefore the laser phase also rotates at the same frequency. When pictured in a frame which is rotating at frequency ω , the Bloch vector is stationary. The Bloch vector rotates about x at the Rabi frequency, ω_R , when the laser is applied. The Rabi frequency gives the strength of the coupling between the laser light and the transition:

$$\omega_R = \frac{\langle 1 | \mathbf{d} | 0 \rangle \cdot \mathbf{E}}{\hbar} \quad (1.43)$$

where $\langle 1 | \mathbf{d} | 0 \rangle$ is the transition dipole moment and $\mathbf{E}$ is the vector electric field amplitude. The xy plane of the Bloch sphere can be rotated about z by changing the phase of the laser light relative to the phase of the ion. This could be done with an electro-optic modulator† for example. Any fluctuations in the intensity of the laser will cause fluctuations of the Rabi frequency, and hence imperfect x rotations. Fluctuations in the laser frequency and in the ion's energy level splitting (caused by magnetic field fluctuations etc.) will cause unwanted z rotations as well as imperfect x rotations.

The classical NOT gate is reversible‡. However, most classical logic gates are not. The fundamental reversibility of quantum mechanics leads to another difference between classical and quantum gates: quantum gates must be reversible. This means that quantum gates must always have the same number of outputs as inputs, unlike classical gates such as the AND gate which has

*If E_0 and E_1 are close (hyperfine or Zeeman levels) then a pair of Raman lasers or an RF/microwave source is used.

†Applying an electric field across the EOM alters the refractive index of the material, which thus alters the optical path length of the laser beam.

‡The original state of a bit can be deduced by knowing the output.

two inputs and one output, or the FANOUT gate which takes one input and produces multiple outputs. An example of a two qubit gate is the famous Controlled NOT (CNOT) gate. This takes two qubits, a control qubit and a target qubit. It flips the state of the target if the control is in state $|1\rangle$. In other words:

$$|\psi\rangle = a|0\rangle_C|0\rangle_T + b|0\rangle_C|1\rangle_T + c|1\rangle_C|0\rangle_T + d|1\rangle_C|1\rangle_T \quad (1.44)$$

$$\widehat{\text{CNOT}}|\psi\rangle = a|0\rangle_C|0\rangle_T + b|0\rangle_C|1\rangle_T + c|1\rangle_C|1\rangle_T + d|1\rangle_C|0\rangle_T. \quad (1.45)$$

Where $|\phi\rangle_C$ is the control qubit and $|\phi\rangle_T$ is the target qubit. In the many worlds interpretation, the target qubit is flipped in all the universes where the control qubit is in state $|1\rangle$. The pair of qubits becomes entangled.

The Toffoli gate is a three-qubit gate very similar to the CNOT gate. It will flip the state of the target qubit if two control qubits are in state $|1\rangle$. In reality it is often more difficult to perform a three or more qubit gate than a one or two qubit gate. Fortunately it is possible to construct any multi-qubit gate from a small set of one-qubit gates and a single two-qubit gate [37]. So to build a real quantum computer, it is not necessary to perform controlled three-qubit interactions.

For a system of trapped ions, energy levels of individual ions are used as qubit states. Single-qubit gates are achieved by applying precisely controlled laser pulses to the ions, and multi-qubit gates are possible due to the Coulomb interaction between different ions. Most progress so far has been made with small crystals of ions in RF traps. It is hoped that a small chain, or even just two ions trapped along the axis of a Penning trap, can be manipulated in a similar fashion to a chain of ions along a linear RF trap. Also, a scheme has been proposed by Porras and Cirac where ions in a two dimensional crystal, rotating in a Penning trap, are used as qubits [27]. In this scheme the axial motion of the ions is cooled and used as part of the quantum logic, while the cyclotron and magnetron motions are ignored.

1.4.4 Cirac-Zoller Gate 1995

The first two-qubit gate based on trapped ions was proposed by Cirac and Zoller in 1995 [29]. A slightly simpler gate based on the same idea was experimentally achieved later in the same year by Monroe *et al.* [33]. The reliability (fidelity) of the full Cirac-Zoller gate when performed by Schmidt-Kaler *et al.*,

was measured to be approximately 70%, requiring $\sim 500 \mu\text{s}$ to complete [11].

The Cirac-Zoller (C-Z) gate is a controlled $\hat{Z}$ gate. It acts on two ions: one is the control qubit, the other is the target qubit. The state of the system can be written as $|c, t, n\rangle$ where the control qubit, c , can be in the ground or excited state (g or e), as can the target qubit. The number of phonons in the centre of mass mode is n . Suppose the system starts in the motional ground state, but not generally in the electronic ground states:

$$|\psi\rangle = a|g, g, 0\rangle + b|g, e, 0\rangle + c|e, g, 0\rangle + d|e, e, 0\rangle. \quad (1.46)$$

The C-Z gate then works as follows:

1. A π laser pulse is applied to the control qubit on the red sideband. If the ion is in the excited state it can go down to the ground state and the system will gain a phonon; if the ion is in the ground state then it will not interact with the laser pulse since there is no $n = -1$ state for it to go into. If the ion absorbs the π pulse it picks up a phase of $e^{-\frac{i\pi}{2}} = -i$.

$$|\psi\rangle \rightarrow a|g, g, 0\rangle + b|g, e, 0\rangle - ic|g, g, 1\rangle - id|g, e, 1\rangle. \quad (1.47)$$

2. A 2π pulse is then applied to the target qubit, on the red sideband of the transition frequency between the ground state and an auxiliary state. The ion passes through this state but it does not become actually populated. If the ion goes up to this auxiliary state and back it will pick up a $e^{-\frac{i2\pi}{2}} = -1$ phase change. This only happens if the target qubit is in the ground state, *and* there is one phonon in the motional state.

$$|\psi\rangle \rightarrow a|g, g, 0\rangle + b|g, e, 0\rangle + ic|g, g, 1\rangle - id|g, e, 1\rangle. \quad (1.48)$$

3. Another π pulse is then applied to the red sideband of the control qubit transition, providing another phase shift of $-i$ to the c and d terms.

$$|\psi\rangle \rightarrow a|g, g, 0\rangle + b|g, e, 0\rangle + c|e, g, 0\rangle - d|e, e, 0\rangle. \quad (1.49)$$

It is seen that now the target qubit has been changed from $\alpha|g\rangle + \beta|e\rangle$ to $\alpha|g\rangle - \beta|e\rangle$ only in the universes where the control qubit was in the excited state. This is a controlled $\hat{Z}$ gate.

As shown in equation (1.42), the $\hat{Z}$ operator can be transformed into a

NOT gate by simply applying a pair of Hadamard gates. Thus the Cirac-Zoller gate can be used as a CNOT gate by applying a Hadamard to the target qubit before and after the main operation.

1.4.5 Cirac-Zoller Gate 2000

The quantum gate described above relies on the ability to cool the ions down to the motional ground state. This makes it sensitive to heating, and also very difficult to perform. The gate proposed by Cirac and Zoller in 2000 does not rely on cooling to the ground state, and it is also more obviously scalable (it assumes that ions are stored in an array of microtraps) [31].

The two-qubit gate is performed by applying a standing wave of off-resonant laser light on to both of the ions. The standing wave causes an AC Stark shift in each of the ions. The strength of the AC Stark shift is dependent on the ion's internal (qubit) state. If an ion is trapped close to a node of the standing wave, then the dipole force caused by the AC Stark shift is roughly linear. This force causes a small displacement of each ion, so that depending on the internal states, the pair of ions move closer to each other or become more separated. If the standing wave is switched on and off such that the ions move adiabatically, then the motional state is unchanged. However, the energy of the pair of ions depends on the mean distance between the two, and so there is a change in the phase of the wavefunction ($e^{\frac{-iEt}{\hbar}}$) which depends on the qubit states. By tuning the parameters of the standing wave laser (detuning, intensity, pulse duration), the amount of phase can be controlled.

The phase change only depends (to first order) on the mean separation of the pair of ions. Although the spatial wavefunction of each motional state is different, the mean position is always the same. Hence the gate does not depend on the motional quantum number.

The experiment performed by Leibfried *et al.* in 2003 followed a similar scheme but with a small difference [34]. A pair of beams with a frequency difference of 6.126 MHz was used to create a moving standing wave*. This produced an oscillating force at a frequency similar to the secular frequency of the ions (6.1 MHz). When the ions are in the same state, the force on each ion is the same, and so the stretch mode of motion is not excited. If they are in different states, there is a differential force between the ions and the stretch mode is excited.

*AKA a 'walking wave'.

When the beams are applied, the quantum state of the ion pair starts moving round a loop in position-momentum (phase) space. Since the frequency of the oscillating force is different from the frequency of the ion motion, the amplitude of the oscillation oscillates at the difference frequency ($6.126 \text{ MHz} - 6.1 \text{ MHz} = 26 \text{ kHz}$ in this case). After $39 \mu\text{s}$, the motional oscillation and the force oscillation come back into phase, and the loop in position-momentum space is closed. The phase gained by the quantum state of the ions is determined by the area of the loop traced in position-momentum space. Only the area of the loop is important, not the shape, and so the gate is fairly robust against thermal motion and laser intensity fluctuations.

1.5 Scalable Quantum Computing

1.5.1 The DiVincenzo Criteria

There are many physical systems which could be considered as candidates for implementing a quantum computer. To be a serious candidate, a proposed system must have certain properties. The traditional list of essential requirements was first described by D. DiVincenzo. His criteria state that the proposed system must have [38]:

1. A scalable physical system with well characterised qubits.
2. The ability to initialise the state of the qubits to a simple state, such as $|0\rangle$.
3. Long relevant decoherence times, much longer than the gate operation time.
4. A universal set of quantum gates.
5. A qubit-specific measurement capability.
6. The ability to interconvert stationary and flying qubits*.
7. The ability to transmit flying qubits between specified locations.

The last two points are not really needed for quantum computing, but rather in order to build a ‘quantum network’.

*Flying qubits transfer quantum information between different places. A flying qubit would most likely be encoded in the state of a photon.

At the time of writing, there are no physical systems which can achieve all seven (or even just the first five) of the DiVincenzo criteria. Table 1.1 (mostly taken from [39]) shows which potential QC systems are close to fulfilling each of the seven DiVincenzo criteria.

QC Approach	1	2	3	4	5	6	7
NMR	N	O	O	Y	O	N	N
Trapped Ion	O	Y	Y	Y	Y	O	O
Neutral Atom	O	Y	O	O	O	O	O
Cavity QED	O	Y	O	O	Y	O	O
Optical	O	O	Y	O	O	O	Y
Solid State	O	O	O	O	O	N	N
Superconducting	O	Y	O	O	O	N	N

Table 1.1: DiVincenzo criteria fulfilled by different technologies. Y = proof of principle has been achieved; O = potentially viable approaches have been proposed, but there has been no sufficient proof of principle; N = no viable approach is known.

One of the most successful approaches so far is the use of liquid state NMR [40]. In an NMR QC, the nuclear spins of various nuclei in molecules act as the qubits. The spin-spin interaction of adjacent nuclei enable quantum gates to be performed. A liquid sample will contain many of the same type of molecule, and so the sample will act as a large ensemble of many QCs. In order to scale up this system to hundreds of qubits, each individual molecule would need hundreds of spin $\frac{1}{2}$ nuclei, each with a slightly different (and resolvable) transition frequency. Due to this difficulty, and others*, NMR quantum computing is generally considered to be non-scalable.

An approach which does seem to have a potential for scalability is quantum computing with trapped ions.

1.5.2 Multiple Ion Traps

The factorisation of a 200 digit number by a quantum computer requires about 3,500 qubits [41]. The number of qubits required rises by another order of magnitude if error correction is to be implemented[†]. On the other hand, the

*The signal to noise level drops dramatically as the number of qubits in an NMR system is increased. As the molecules are only cooled to a ‘pseudo-pure state’, the number of molecules which are actually used relative to the total number is very small for larger numbers of qubits.

[†]Just as classical computers require additional bits to perform error correction, so too do quantum computers. A classical error correction system typically reinforces a large number of bits using just a few check (parity, checksum etc.) bits. On the other hand, quantum error correction systems require a single logical qubit to be encoded using several ($\gtrsim 5$) physical qubits.

cutting edge experiments at NIST and Innsbruck are producing entangled states of less than ten ions [35, 36]. The motional spectrum of a single chain of ions in a linear RF trap (or along the axis of a Penning trap) becomes increasingly complicated and dense as more ions are added to the chain. A chain of N ions has N normal modes along each direction, with each mode producing sidebands in the spectrum. For this reason, a single linear trap is deemed to be non-scalable in terms of quantum computing.

A useful quantum computer would need many qubits, but the gates need only be performed on single qubits and between pairs of qubits. There is work currently being performed around the world involving RF traps which contain many different trapping regions. Each trapping region would contain only one ion or a few ions. By altering the voltages on the different electrodes, ions can be moved between different trapping zones.

Some theoretical proposals of scalable microtrap arrays have been published [31, 42], and a number of real ion trap arrays have also recently been built. The group of Wineland has demonstrated a working microtrap built from gold on a quartz substrate [43]. Ion trap arrays have also been micro-fabricated from semiconductor materials, both silicon based [44], and gallium-arsenide based [45]. Printed circuit board has also been used to produce RF ion trap arrays [46, 47] (as well as the trap described in chapter 4). Most of these microtraps have dedicated RF electrodes to provide the trapping field, and also ‘DC’ electrodes to push ions between traps.

In general there should be the possibility of performing a gate between any pair of qubits in a quantum computer. For this to be implemented, either qubits could be swapped between pairs of ions (using SWAP quantum logic gates), or the ions could be physically moved around. The group of C. Monroe has entangled pairs of ions where each ion was stored in a completely separate trap [48]. However, physically shuttling ions between traps is currently the most popular approach [42, 49]. Ions in these traps have been rapidly shuttled between different sub-traps [46], and even transported around corners [50]. Little similar work has been done elsewhere with multiple Penning traps for quantum computation. Our experimental efforts to produce a prototype scalable Penning ion trap are detailed in chapter 4.

Chapter 2

Experimental Setup

The work of this thesis was performed with two different ion traps, henceforth referred to as the *split-ring trap* and the *PCB trap*. The split-ring trap is described below in §2.4, and in more detail in [51] and [52]. The PCB trap is the subject of chapter 4. Much of the other equipment (lasers, magnet, etc.) is external to the two traps and is identical or very similar in both cases. This chapter contains a description of this equipment, focusing on the items which have not been detailed elsewhere (such as in [51], [52] and [53]). Possible improvements to the setup are also suggested.

2.1 General Outline

Figure 2.1 shows a highly simplified diagram of the whole experiment. The trap is held inside a vacuum chamber, which is between the pole pieces of a magnet. The trap electrodes are connected to electronics outside the vacuum. Also within the vacuum are an atomic beam oven, and a hot cathode electron beam filament, which together are used for creating calcium ions. Laser beams are sent into the trap. This laser light can cool the ions and cause fluorescence to be emitted. The fluorescence is detected by either a photomultiplier tube (PMT) or an intensified charge-coupled device camera (ICCD) connected to a computer. Some control of both the lasers and the trap electronics can be provided by another computer, which may even receive feedback from the detection system.

To run the experiment, the laser wavelengths are initially measured using a wavemeter (see §2.3.2) and adjusted appropriately. Then the trap is loaded




Figure 2.1: General schematic of experiment. Dashed lines indicate transfer of light. Continuous lines indicate transfer of electrical signals.

with calcium ions by heating the oven to produce a weak atomic beam and also heating the negatively biased filament to produce an ionising beam of electrons. Assuming some fluorescence is observed with the PMT, many parameters can be optimised to increase this level of fluorescence. When the laser wavelengths are correct, they may be locked using the system described in §2.3.3. Progressively smaller clouds of ions are then loaded, down to individual ions, while the parameters are manually optimised even further.

2.2 Calcium 40 Ions

All the work of this thesis has been performed using $^{40}\text{Ca}^+$ ions. Laser cooling work is currently being done elsewhere with the odd isotope, $^{43}\text{Ca}^+$, as well as with the much more abundant $^{40}\text{Ca}^+$ [54, 55]. From this point on, the most abundant isotope is assumed, and Ca^+ is taken to mean $^{40}\text{Ca}^+$.

The energy level diagram of Ca^+ is shown schematically in figure 2.2. The g factors used to calculate the linear Zeeman splittings are given in table 2.1. For an ion in a Penning trap, the energy splitting due to the magnetic field is generally much smaller than the fine structure splitting, but much greater than the linewidth of the transition. Thus to avoid optical pumping into one of the Zeeman sublevels when laser cooling, they must be individually considered

and driven. The transitions addressed by the various lasers in our setup are shown in figure 2.2. The 397 nm transitions chosen are both π ($\Delta m_J = 0$) polarised, while all four of the 866 nm transitions are σ ($\Delta m_J = \pm 1$) polarised.




Figure 2.2: Energy level structure of singly ionised calcium 40 with Zeeman splitting shown at a magnetic field of 1.0 tesla (not to scale). Laser cooling and repumper transitions are shown. The fine structure and Zeeman splitting are both exaggerated for clarity.

The most important transition from the viewpoint of laser cooling is the $S_{1/2} \rightarrow P_{1/2}$ transition at 397 nm (396.84673 nm in vacuum at low B [56]). When an ion absorbs a 397 nm photon it will absorb the momentum of the photon, and then 94% of the time it will emit another 397 nm photon in a random direction. With the appropriate red detuning of applied laser cooling

$S_{1/2}$	$P_{1/2}$	$P_{3/2}$	$D_{3/2}$	$D_{5/2}$
2	2/3	4/3	4/5	6/5

Table 2.1: g factors of the relevant states in Ca^+ .

light, this will cool the ion.

With a branching ratio of 6%, an ion in the $P_{1/2}$ state decays to the $D_{3/2}$ state instead of the ground state. An ion in the $D_{3/2}$ state cannot decay via an electric dipole transition, so it has a long lifetime of around one second. This electron shelving is problematic as it will interrupt the laser cooling cycle. To put the ion back into the $S_{1/2} \leftrightarrow P_{1/2}$ cooling cycle, repumper light at 866 nm (866.2140 nm [56]) is applied.

In theory, the $D_{5/2}$ state will never be populated. In practice however, things are not so simple, and any unwanted decay to $D_{5/2}$ will interrupt the cooling process. Ions can be pumped back into the cooling cycle via the 854 nm transition $D_{5/2} \rightarrow P_{3/2}$. Using just one laser at 854 nm, the large Zeeman splitting makes $D_{5/2} \rightarrow P_{3/2}$ repumping weak and unreliable. However, even modest controlled shortening of the $D_{5/2}$ lifetime can be useful when cooling the ions and optimising the system parameters. Decay to the $D_{5/2}$ state is discussed in more detail in chapter 5.

2.2.1 Ion Generation

Electron bombardment ionisation of a weak atomic beam is used to generate ions in situ in the trap. Both the split-ring trap and the PCB trap use the same technique. A typical loading procedure is discussed in more detail in §3.1.

To construct the atomic beam oven, a tube of tantalum was filled with small lumps of calcium metal, freshly cut with a scalpel. The tube was 1 mm in diameter and approx 10 mm long for the split-ring trap. A 2 mm diameter tube, 5 mm in length was used for the PCB trap. The larger diameter tube was significantly easier to fill. A hole of 0.5-0.8 mm diameter was made in the wall of the tube. The tube was spot welded to a 0.25 mm diameter tantalum wire, then attached to the electrical feedthrough below the trap. For the split-ring trap, the oven wires were simply connected to the feedthrough wires using barrel connectors. The oven in the PCB trap system was held in place using screws and ceramic washers, as shown in figure 4.7 in chapter 4. Care was taken to avoid exposing the calcium to any liquid, or even to air for more than

a few hours as the element is highly reactive. Passing a current through the oven causes it to heat up and emit atoms. It was found that the dimensions of the tube and wire strongly affected the amount of current required to produce a calcium beam. In the split-ring trap, a typical oven current is 1.8 amp. In the PCB trap setup, with a thicker tube and shorter leads, a current of 5.0 amp was normally used. The temperature of the oven when heated was not measured, but observation of a faint red glow indicates a temperature of ~ 500 °C.

The electron beam is generated by passing a current through a filament, which is negatively biased relative to the potential at the centre of the trap. The filament is simply a piece of thoriated tungsten wire, 0.125 mm in diameter. The energy of the electrons as they collide with the atoms, is approximately the same as the potential between the filament and that of the trapping region (times the electron charge). For reliable ionisation of calcium (and to avoid double ionisation) a filament bias of around -10 V was used*, producing an electron beam with an energy of about 10 eV.

Since the electrons are much lighter than the calcium atoms/ions, the collisions do not impart much kinetic energy to the atoms/ions. The typical initial energy of an ion (before cooling) is therefore a combination of the kinetic energy of the atom ($\frac{3}{2}k_B T \approx 100$ meV) and potential energy due to the trap (≈ 500 meV). The latter is minimised if ions are only created close to the trap centre. Unfortunately in our systems the ions are formed over a fairly large volume. This also means that newly created ions in a Penning trap tend to be moving in a large magnetron orbit which takes time to shrink as the ions cool.

As an alternative to electron bombardment, laser photoionisation can be used to load ion traps. This technique has a number of advantages, not least of which is that a photoionisation laser beam can be focused close to the trap centre. Photoionisation of calcium, first achieved by N. Kjaergaard *et al.* [57] in Aarhus also has the following further advantages:

- The efficiency is five orders of magnitude greater than electron bombardment ionisation [58], allowing the use of a lower atomic flux which leads to less contamination of the electrodes.
- The cooler oven, and lack of a hot filament produce less outgassing and

*Since the resistance of the filament is so low, this bias can be applied to either side of the filament.

so a lower background pressure.

- There is no charging up of insulators by an electron beam.
- Fragile filaments tend to break after prolonged usage. A broken filament requires the vacuum system to be opened.
- In a Penning trap the ionising beam is restricted to travelling along the magnetic field direction. Photoionisation removes this constraint at the trap design stage.
- Atoms and molecules other than calcium will not be ionised, and a particular isotope can even be selected.
- Photoionisation allows more control over the loading rate – i.e. a filament stays hot and will produce more electrons if operated twice in rapid succession, while a laser beam can be shuttered and filtered easily.

These advantages are weighed against the complexity (and cost) of a photoionisation system. The simplest of these systems, as demonstrated by D. Lucas *et al.*, is a diode laser operating at 423 nm combined with an LED at 389 nm [54].

Obviously, a photoionisation laser for this experiment is highly desired, but in the short term at least, the final point of the above list can be improved upon. Instead of passing a constant current through the filament, the electron beam is detected as a small (μA) current, and feedback is applied to the filament controller to keep this beam current at a stable value. The feedback circuit, which also allows computer control of the oven and filament, is shown in figure A.2 in the appendix. The control circuitry is optically isolated from the filament, so that any bias can be applied. In the PCB trap, the oven itself is used to detect the electron beam. Since the oven is connected to a power supply, and must be otherwise electrically floating, a large amount of mains noise is present. This problem was solved using the brute force approach of passing the control signal through a multi-pole low pass filter. The circuit was designed by the author, and was mostly built by 6th form students G. Farrell and A. Soltani. This electron beam controller was not used when producing the results presented in this thesis, but will be used in the future*.

*At least until a reliable photoionisation laser is implemented.

2.3 Laser Systems

An advantage of using calcium ions in a laser cooling experiment is that it is possible to produce the required wavelengths (397 nm, 866 nm) using diode lasers. Diode lasers are relatively cheap and easy to maintain compared to other tunable lasers such as dye lasers and doped crystal/glass lasers. They can also be made fairly small. It is not possible to produce very high power beams with diode lasers*, but for laser cooling small ion clouds, only low power is needed.

The optical setup consists of two diode lasers running at 397 nm, four at 866 nm, and one at 854 nm. A schematic of the optical table is shown in figure 2.3.

After all the beams are combined, they are focused into the trap. The beam is focused by a 25 cm focal length plano-convex singlet lens, and is ~ 3 mm in diameter before the lens. This produces a beam waist at the centre of the trap of ~ 100 μm , with a Rayleigh range of ~ 5 cm.

A narrow linewidth titanium:sapphire laser operating at 729 nm is currently being developed for performing spectroscopy and coherent manipulation on the electric quadrupole $S_{1/2} \leftrightarrow D_{5/2}$ transition. This laser, described in some detail by H. Ohadi [52], was not used to obtain any results presented in this thesis, and will not be mentioned further.

2.3.1 Extended-Cavity Diode Lasers

Laser diodes generally have a broad gain curve. That is, it is possible for the laser to run at a range of different wavelengths. This range varies greatly between different diodes, from a few nm up to many tens of nm as shown in figure 2.4. In order to tune the wavelength to the desired value, a diffraction grating is used to feed some light back into the diode, as shown schematically in figure 2.5. The wavelength of the light which goes back into the diode depends on the angle of the grating. Thus the angle can be adjusted to pick the wavelength range to be fed back into the diode. The losses of the light of this wavelength are smaller than that of light which is rejected by the grating, and so the modes of this wavelength range are preferentially amplified.

Adding the grating effectively lengthens the laser cavity from L_0 to L_1 . This decreases the linewidth of a laser mode by a factor of roughly L_1/L_0 .

*No more than a few tens of mW.




Figure 2.3: Schematic view of the main optical table.




Figure 2.4: Emission Spectrum of an Eagleyard Photonics EYP-RWE-0870-06010-0750-SOT01-0000 laser diode without external feedback. Taken from [59].




Figure 2.5: The Littrow configuration for tuning a laser diode using a diffraction grating.

By adjusting θ , one of the modes of the diode-cavity (L_0) is selected. There are several extended-cavity (L_1) modes within the linewidth of the selected diode-cavity mode, but just one of them will preferentially lase. Very fine adjustments to the wavelength can be made by making small adjustments to the length L_1 . This is possible as the grating is mounted on a stack of piezoelectric crystals. The effective length of the diode-cavity (L_0) can also be tuned by adjusting the temperature, and the current passing through the diode. These alter the refractive index of the semiconductor as well as the actual length.

The two lasers used for the 397 nm transitions are commercially available tunable extended cavity diode laser systems (TuiOptics* DL 100). These are seen as the dark blue boxes, B1 and B2, on the left hand side of figure 2.3. The actual diodes used in each are not identical, leading to an output power imbalance between the two. This can easily be solved by filtering the more powerful laser. The final output power of each is approximately 5 mW. To avoid shortening the lifetime of the diodes, they are generally operated at a fairly low current of 57 to 60 mA – only 1 or 2 mA above the lasing threshold.

The diode, collimator and grating holder are mounted onto a base plate which sits on top of a Peltier thermoelectric element. This then sits on a base along with a temperature sensor. The Peltier and the temperature sensor are connected to a temperature controller housed in a supply rack, along with the current controller, piezo controller and proportional-integral-differential (PID) locking controller.

The infrared (IR) repumper lasers are similar to the 397 nm lasers, but built in-house based on the design described by Boshier *et al.* [60]. They are shown as the green boxes on the right hand side of figure 2.3. The lids of the mounts are coated in 3 cm of packing foam to improve temperature stability. Several different IR diodes were tested in the same mounts over a period of several years. The most stable, powerful and reliable diode found so far is the Ridge Waveguide Laser for external cavity setups (EYP-RWE-0870-06010-0750-SOT01-0000) from Eagleyard Photonics. This GaAs laser diode has an anti-reflection coating on the output facet, and is optimised for use in external cavities. The gain curve (figure 2.4) is nominally centred around 870 nm and lases well at 866 nm. Each of the IR lasers has a separate current and temperature controller, each with an individual power supply and display.

*Now Toptica.

These have been acquired from several different sources (mostly Thorlabs and Laser 2000). A suggested improvement would be to replace this with a rack of identical current control and temperature control boards, powered from one power supply and controlled digitally from a single computer.

The laser operating at 854 nm is required to repump a range of transitions and so would ideally be broadband and high power. In order to increase its power and linewidth, the laser (R5 in figure 2.3) has no grating. It is an Eagleyard EYP-RWL-0850 diode (nominally centred at 850 nm), set to produce light centred at 854 nm by adjusting the temperature and current. This offers a small improvement over a grating-stabilised diode for repumping the $D_{5/2}$ state. A suggested alternative would be to rapidly modulate or sweep the frequency of a single laser over the whole Zeeman broadened $D_{5/2} \rightarrow P_{3/2}$ transition. This is easier said than done with a standard diode laser, as the Zeeman splitting at 1 tesla is tens of GHz. Employing six individual lasers at 854 nm is currently considered even more impractical. Photos of the diode laser systems are shown in figure 2.6.




Figure 2.6: Photographs of the diode laser mounts with lids removed. Left: TuiOptics DL 100 with 397 nm diode; Right: 866 nm laser.

The two 397 nm beams are combined on a beam splitter and coupled into the same optical fibre. The fibre (Thorlabs PM-S350-HP) is single mode and polarisation maintaining. Similarly, all of the IR beams are combined and coupled into a single mode fibre (Oz Optics SMJ-A3A,3AF-850-5/125-3-4). The main disadvantage of using fibres is that the laser power is significantly reduced (transmission is between 25 and 35% for both the 397 nm and IR beams). The advantages are:

- The beams can be readily transported off the main optical table and into, for example, a different ion trap.
- The beam shape after the fibre has a high quality Gaussian profile.

- The pair of 397 nm beams are perfectly aligned with each other, as are the four or five IR beams.

The last point in particular is very useful in reducing the number of experimental parameters which require optimisation. Ideally in the future, one further step will be taken by combining the 397 nm beams with the IR beams into just one single mode fibre.

2.3.2 Wavemeter

To tune the lasers initially, the wavelengths are measured using a wavemeter. In this device, shown schematically at the bottom of figure 2.3, the beam is split into two components using a beam splitter. These two components travel along paths of different lengths, are recombined with the same beam splitter and are then overlapped onto a photodiode. The path lengths are constantly changing, because part of the interferometer is mounted on a small car which travels back and forward along a glass rod, floating on compressed air. As the car moves all the way along the rod, interference fringes on the photodiode are counted. Another (reference) beam is also sent parallel to the first beam and interfered on another photodiode. The wavelength of this reference beam is well known, so when the number of fringes of each of the two lasers is compared, the wavelength of the test beam can be found. If n and n_{ref} are the number of fringes counted from the test beam and the reference beam respectively, then the measured wavelength is simply

$$\lambda = \lambda_{\text{ref}} \times \frac{n_{\text{ref}}}{n}. \quad (2.1)$$

By measuring the relative phases of the two sets of interference fringes, an improvement can be made on the precision of the wavelength measurement. This enables wavelengths to be measured to a precision of better than seven significant figures. Usually this is adequate for then observing some fluorescence when loading a cloud of ions in a trap. Improvement can then be made by fine tuning the lasers to optimise the ion fluorescence signal.

The reference laser used is a frequency-stabilised Helium Neon (HeNe) laser at 633 nm. The stabilised HeNe first used in the wavemeter was a 05 STP 903 laser from Melles Griot, costing several thousand pounds. This laser developed faults and became unusable, so one of the first tasks of the author was to build a new stabilised HeNe for a fraction of the cost of the commercial

laser. The home-built laser has a modular design which allows the replacement of individual components.

The gain curve of the HeNe laser is ~ 1 GHz wide. The length of the laser tube is ~ 15 cm, so two longitudinal modes of the light can be above the threshold for lasing. These two modes both lase simultaneously, and have opposite (linear) polarisations. The laser tube (JDS Uniphase model 1007), is designed to operate in this way. Figure 2.7 schematically shows two modes simultaneously above threshold. The intensities of the two modes will be the same order of magnitude, but not identical.




Figure 2.7: Two longitudinal modes can fit under the gain curve of the HeNe laser. (Approximately to scale).

The tube is powered by a 1250 volt supply, with a 100 k Ω ballast resistor close to the anode of tube. As the laser operates, it heats up and expands. This expansion shifts the frequencies of the lasing modes*. The position of the gain curve however – which is determined by the intrinsic properties of atomic neon – will not move. Therefore the modes will slide along under the gain curve, a new one coming above threshold and lasing as an old one passes below threshold.

This frequency drift is removed by observing the intensities of the two modes and providing feedback to the length of the laser tube. Since the two modes have opposite polarisations, this can be done quite easily. Some light is taken from the laser (through the rear mirror) and passed through a polarising

*If the laser tube expands by a fraction, δ , the wavelength of the n th mode will increase by $\delta\lambda$.




Figure 2.8: Schematic diagram of stabilised Helium Neon laser.

beam splitter. The intensity of each of the polarisations is measured by a photodiode. If the HeNe tube has the correct orientation* then the intensities of the two polarisations correspond well to the intensities of the two modes. The difference between the photodiode signals is then used to control a heater which is wrapped around the laser tube. A Minco Thermofoil heater[†] was wrapped around the HeNe tube and held in place with silicon stretch tape. The tube must be warmed up before the stabilisation system is locked, so that the tube expands when the heater is on and contracts when the heater is off. A diagram of the stabilised HeNe system is shown in figure 2.8.

The control electronics are shown in figure A.1 in the appendix. A switch allows the user to change between ‘manual’ mode – where the voltage across the heater is controlled manually by adjusting a potentiometer on the control; and ‘lock’ mode – where the heater voltage is controlled by an error signal obtained from the difference between the two photodiode signals. The manual mode is useful for rapidly warming up[‡] the laser before it can be locked successfully. A voltmeter was built into the control box, which is used to monitor either the error signal or the voltage across the heater.

Another switch allows the signals from the two photodiodes to be easily swapped – effectively flipping the positions of the two modes shown in figure 2.7 about the centre of the gain curve. Rotating the polariser at the output of the

*To optimise the orientation, the tube was rotated until the largest intensity difference reached between the two polarisations is maximised. Once the orientation has been optimised, it should not require changing again because the absolute angle of the polarisations is only determined by the structure of the tube and does not change over time.

[†]Thermofoil heaters are thin, flexible heating elements consisting of an etched foil resistive element laminated between layers of flexible insulation.

[‡]Warm up takes approximately 30 minutes.

laser by 90° allows either one of the two modes to be used by the wavemeter. Thus there are four possible combinations which produce four slightly different output wavelengths. The combination which gave the greatest output power was chosen.

The reference laser works reliably, but it should be noted that any adjustment made to the alignment of the locking optics causes the output wavelength to shift. This shift is always a significant fraction of the (~ 1 GHz) free spectral range, and hence very significant compared to the 22 MHz Ca^+ cooling transition linewidth. Thus the reference laser wavelength must be recalibrated whenever the optics are realigned*

2.3.3 Reference Cavities

All lasers drift in wavelength over time. Despite being temperature stabilised, the frequencies of the diode lasers will often drift by an amount comparable to the Ca^+ transition linewidths within minutes or hours. To improve this situation, the laser wavelengths can be locked to external reference cavities which are more stable than the diode-grating cavities. The lasers are locked by feeding back to L_1 , such that transmission of the laser through a reference cavity is kept at a setpoint on the side of one of the cavity fringes. Separate reference cavities are used for each 397 nm laser, and a third cavity is used to stabilise just one of the 866 nm lasers. The remaining IR lasers are then locked to the stabilised IR laser using computer controlled feedback from the transmission spectrum of a fourth cavity, the length of which is continuously scanning.

The reference cavities are described in detail in [53]. The cavities consist of Zerodur[†] spacers, approximately 120 mm long. A mirror is glued to one end, while the mirror on the other end is mounted on a stack of piezoelectric crystals. The piezo stacks are mounted as a re-entrant design (i.e. two sets with similar length oppose each other to cancel out length changes due to temperature fluctuations). The finesse[‡] of all the cavities is around 60. The

*This can be done first by setting the diode lasers (just one 397 nm and one 866 nm) to approximately the right wavelengths for cooling ions in an RF trap. A small amount of fluorescence should still be seen even if the laser frequencies are ~ 1 GHz off. Then after optimising the fluorescence signal by tuning the lasers, they are measured again with the wavemeter and the corrected calibration is found.

[†]Zerodur[®] is an inorganic, non-porous glass ceramic manufactured by Schott Advanced Optics. It has a very small thermal expansion coefficient ($\sim 10^{-7}$ K⁻¹ [61]).

[‡]The finesse of a cavity is defined as its free spectral range divided by the width (FWHM) of its resonances.

radius of curvature of the cavity mirrors is 30 cm.

The cavities themselves have not changed since 2004. However I have implemented a number of improvements to the container, the temperature stabilisation and the piezo controller. These correspond to the main external fluctuations which can affect the optical length of such a cavity – the temperature; the air pressure (and humidity, and temperature); and the voltage on the piezo stacks.

When the cavities were first used, they were housed in an unsealed aluminium box. A single bead thermistor was attached to the outside of the box, along with a length of heater wire. The box was then insulated with Neoprene. The thermistor was used in a Wheatstone bridge with a potentiometer for setting the desired temperature. The error signal was processed with a PI controller which then passed a current through the heater wire. The piezo voltages were set using a standard high voltage piezo driver. The piezo driver voltage could be set from 0 - 200 V, while the cavity goes through one free spectral range in just 2 V.

Figure 2.9 shows how the refractive index of air changes with pressure. The small change in refractive index over standard atmospheric conditions is seen to cause a frequency shift (about 200 MHz per mbar at 397 nm) which is large compared to the 22 MHz Ca^+ cooling transition linewidth. A frequency shift due to atmospheric humidity fluctuations was also estimated – 80 MHz at 397 nm for 10% humidity change. The shift due to a variation of the air temperature is roughly 700 MHz per degree C.




Figure 2.9: Change in the refractive index of air, and the implied frequency shift at 397 nm, as a function of pressure. Estimated from the Ciddor equation via the NIST Engineering Metrology Toolbox [62].

The obvious method of eliminating these effects is to house the cavities inside a vacuum chamber. A vacuum chamber was constructed, and is shown schematically in figure 2.10. Stainless steel flanges were welded onto a 150 mm diameter steel tube. A pair of PVC rings, bolted onto the end flanges hold glass windows in place. The seal between each window and flange was made using a viton o-ring. BNC sockets were mounted in holes along the top of the main tube. The sockets were coated with TorrSeal epoxy on both sides to provide a high-vacuum seal, and then soldered to the electronic devices inside the chamber. Heater wire is wrapped around the main tube, which is covered with insulating foam (not shown in the diagram). A UHV valve was attached to one side, while the other side was sealed with a blanking flange. The system was pumped out with a turbo pump and then valved off. In the future an ion pump could be attached in place of the blanking flange, but currently the system is left with no continuous pumping.




Figure 2.10: Vacuum chamber for housing the reference cavities. Dimensions in mm.

Inside the chamber, the three cavities are resting on a piece of FR4 circuit board, and glued in place with TorrSeal. Four high-precision thermistors (Minco MS0/15934/1) are glued onto the sides of a cavity, and are electri-

cally connected in series. The thermistors form one leg of a Wheatstone bridge, along with three thin-film resistors of $33\text{ k}\Omega$ (1%) each. The two midpoints of the bridge are directly connected to an instrumentation amplifier IC (AD622ANZ) with a gain of 1000. The schematic is shown in figure 2.11. Outside the vacuum can, the output is passed through a proportional-integral controller, and then returned as a current through the (approx $50\text{ }\Omega$) heater wire using a power transistor (TIP3055). The thermistors each have a resistance of $33\text{ k}\Omega/4$ at a temperature of $29\pm 1\text{ }^\circ\text{C}$. Equilibrium at this temperature is reached when roughly 5 W is dissipated by the heating wire. Since the resistors, the amplifier, and the voltage regulators are all inside the thermally stable chamber, any spurious change in error signal (not due to a change in thermistor resistance) should be very small. The resistance of the thermistors as a function of temperature is quite linear between 20 and $40\text{ }^\circ\text{C}$. The response was experimentally found to be $R \approx (-0.33\text{ }T/\text{K} + 18)\text{ k}\Omega$ for each thermistor. With $\pm 15\text{ V}$ rails and a gain of 1000, this leads to a relationship between the temperature and the error signal of $\Delta T \approx 2\text{ mK}$ per volt. After a warm up time of one day, the error signal voltage is stable to within $\pm 0.05\text{ V}$. Thus the average temperature of the four thermistors appears to be stable to $\pm 0.1\text{ mK}$. This seems remarkably good, but note that there will be thermal gradients and temperature fluctuations away from the thermistors which are much larger than this.




Figure 2.11: Wheatstone bridge plus instrumentation amplifier.

In addition to improving the pressure and temperature stability, the voltage applied to the piezos used to tune the cavity length is now more stable and has finer control. A $10\text{ k}\Omega$ multiturn potentiometer (coarse control) is connected across a 10 V precision voltage reference*. The wiper is loaded

*A REF01 was used.

by a second 10 k Ω pot (fine control) in series with a 100 k Ω resistor. This is repeated three times in parallel for the three cavities. The outputs of the fine control pots are directly connected to the piezo stacks of the cavities. No buffering is required because the piezo is a purely capacitive load, and fast response times are not necessary. The voltage reference drifts by a maximum of 8.5 ppm per $^{\circ}\text{C}$ (3.0 ppm per $^{\circ}\text{C}$ typical) [63]. Assuming a typical output is around 5 volt, and one free spectral range of the cavity is covered by a 2 volt change, the cavity will shift by only ~ 20 kHz (at 397 nm) for a 1 $^{\circ}\text{C}$ change in ambient temperature. For comparison, a typical piezo driver, outputting around 100 V and drifting by ~ 100 ppm per $^{\circ}\text{C}$, gives a cavity drift of ~ 5 MHz per $^{\circ}\text{C}$. This 5 MHz is small compared to the other effects, and even smaller than the transition linewidth, however it was still an unwanted drift which could be (and was) straightforwardly improved upon.

In conclusion, the performance of the cavities is now greatly improved. With the lasers locked, ions have been trapped for many hours without a significant change in fluorescence level. However, whenever the tuning voltage is altered, the piezo lengths then drift/overshoot very noticeably over a timescale of minutes, and continue to drift for many hours afterwards. This is a common effect with piezo tunable cavities. In fact, the low drift etalons produced and sold by the National Physical Laboratory (NPL) are advertised as having a much lower overshoot than conventional cavities [64]. In conclusion, our cavities have somewhat poorer performance than the NPL cavities, but the system is more than an order of magnitude less costly than three NPL etalons. If tuning the cavities can be kept to a minimum, then the drift is low enough to work with single trapped ions for extended periods of time.

2.3.4 Multiple laser locking system

To stabilise the frequencies of the four 866 nm lasers, instead of locking each individual laser to its own stable reference cavity, just one laser is locked. The stability of this master laser is then transferred to several slave lasers using a scanning Fabry-Perot cavity, and computer controlled feedback.

The free spectral range of a confocal cavity of length L is

$$\text{FSR} = \frac{c}{4L}. \quad (2.2)$$

Several laser beams are combined and coupled into the same cavity. When the length of the cavity is scanned over a free spectral range then each of the

different laser wavelengths appears as a peak in the transmission of the cavity. An example of this is shown in figure 2.12.




Figure 2.12: The transmission of a cavity (black) coupled to four lasers as the cavity length is scanned by scanning the voltage (red) of a piezoelectric crystal.

Assume lasers A and B produce peak maxima when the piezo voltage is V_A and V_B respectively, and the voltage varies linearly with time during each sawtooth scan, $V_A \propto t_A$. If laser A has a constant wavelength of λ_A , then λ_B is stabilised by adjusting laser B such that $t_{AB} \equiv t_B - t_A$ is kept constant. Several more lasers can be added, and locked to λ_A in the same way. Such schemes have been implemented elsewhere using a temperature controlled scanning cavity, and a stable helium-neon laser as a reference [65, 66].

The scanning cavity, shown in the upper right of figure 2.3, is a confocal cavity (TecOptic SA-7.5) with a free spectral range of 7.5 GHz. The cavity is scanned by one free spectral range by applying a piezo voltage change of roughly 400 volt. The reflectivity of the mirrors is 99.3% at 866 nm, however, the apparent finesse depends strongly on the input beam shape, the alignment, and the discrepancy between the cavity length and the true confocal length. An appropriately placed aperture before the cavity significantly improves the quality of the laser beam shape, and hence improves the quality of the peaks.

The cavity is scanned with a sawtooth voltage (shown in red in figure 2.12) over less than one free spectral range, roughly 5 times per second. The transmission is detected by a photodiode built into the cavity body. This signal is amplified and then measured by a National Instruments USB-6008 data acquisition card. The card also receives a trigger pulse in phase with the scanning

voltage. A LabView* program then attempts to detect four peaks within one scan. The peaks must be above a certain threshold and also over a certain width.

The spectrum shown in figure 2.12 is detected and displayed by the LabView program. Four large peaks can be seen corresponding to the four lasers. Thinner peaks are also seen on the right hand side which are due to the rapid flyback of the cavity length. When the lock is engaged, the data containing the four peaks is split into four subarrays. The program saves the positions of each peak – one per subarray. The program also knows which peak corresponds to the master laser (λ_A ; R2 on figure 2.3). When the cavity is scanned again, the inter peak distances are compared to values calculated from the saved ones. If one of them has drifted by more than a certain amount, a signal is sent to increment (or decrement) a digital potentiometer which (added to a manual control pot) sets the voltage across the piezo controlling the laser cavity length, L_1 . This feedback stabilises the drift of the lasers, though only on slow timescales (on the order of the scanning period, 0.2 seconds).

The locking program has the additional feature of an audible alarm to call attention whenever there is a problem with the system of 866 nm lasers. Whenever the alarm is on, the feedback is temporarily deactivated. The alarm will sound unless each of the subarrays contains one and only one peak. Thus if a laser mode-hops it will probably activate the alarm. The alarm is also sounded when one of the digital potentiometers moves close to the limit of its tuning range. This can be fixed easily by adjusting the appropriate manual control.

If the scanning cavity itself drifts so far that the peaks move out of their respective subarrays, the program will deactivate feedback and engage the alarm. This drift can easily be compensated for by manually changing the DC offset of the sawtooth scanning voltage. A future improvement will be to make the program adjust the offset voltage automatically to keep the master laser peak centred at its original position.

In our setup, the lasers are of very similar wavelengths (within 100 GHz). A stable HeNe laser is not used as the reference. This has a couple of advantages:

1. As mentioned above, a change in air pressure causes a significant change in refractive index. The change in refractive index is different for different wavelengths. This difference causes a slight shift in the peak to peak

*Laboratory Virtual Instrumentation Engineering Workbench.

separation when the refractive index changes, but the shift is suppressed if the wavelengths are very similar. For 866 nm and 632 nm, the wavelength dependent refractive index change leads to a ~ 10 MHz shift per mbar pressure change.

2. If the laser wavelengths were all very different then the spectrum would appear very differently whenever the p in equation (2.2) changed slightly. It is not uncommon for the cavity length to change by a few free spectral ranges over the course of days/weeks.

On the other hand, at the risk of increasing complexity, it may be better to follow the work of A. Rossi *et al.* [66] more closely by temperature stabilising the scanning cavity and using a HeNe laser as the master laser. This would eliminate the overshoot problem of the reference cavity mentioned above. Also the PID lock of laser R2 is more sensitive to acoustic noise and it is much more likely to become unlocked than a stabilised HeNe laser.

The majority of this multiple laser locking system was built as part of an M.Sc. project by I. Bhatti under my day to day supervision. More detailed information can be found in [67].

2.4 Split Ring Trap




Figure 2.13: Schematic of the split-ring trap. Figure courtesy of H. Ohadi [52].

The results of chapter 3 (and [12]) were obtained using the split-ring trap. This trap has already been described in detail in [51], [52], [68] and elsewhere, so only a brief outline is given here..

The trap was designed by M. van Eijkelenborg [69]. It has a standard Penning trap geometry, but with the ring electrode split into four segments – hence the name. The electrodes are made from solid oxygen-free copper. They are not hyperbolic, but machined with conical shapes to give a good approximation to a hyperbolic trap. The distance between the endcaps, $2z_0$, is 7 mm. The radius of the ring, r_0 is 5 mm. Thus the characteristic trap dimension, $R_0^2 = 2z_0^2 + r_0^2$ would be 50 mm² if the trap was ideal. Instead, by measuring motional frequencies of small ion clouds, a value of 55 ± 1 mm² was found [51].

The ring segments are 1 mm apart. The gaps between the ring segments are enlarged with cylindrical holes to improve optical access. The laser beam passes between two pairs of ring segments. Fluorescence escapes the trap between another pair of ring segments, at right angles to the laser beam and the B field. Fluorescence detection is described in §2.4.3. A pair of calcium filled ovens are located between the ring and an endcap on each side. Heating either of these ovens produces a weak beam of atomic calcium directed into the trap. A filament is located within a recess at the rear of each endcap electrode. A small hole in each of the endcaps (not visible in figure 2.4) allows the beam of electrons generated by each filament to pass through the trap.

2.4.1 Vacuum System

The vacuum system is shown, approximately to scale, in figure 2.14. The tubes and flanges are all standard CF40 size. They are made from 304 and 316 stainless steel, both of which display a low enough permeability to be used in high magnetic fields. The ion pump is a Varian triode with a pumping speed of 30 l/s. After cleaning*, assembling, and pumping with a turbomolecular pump (PfeifferBalzers TPU-062), whilst baking at around 215 °C for a week, the background pressure stabilised at 2×10^{-10} mbar.

2.4.2 Magnet

The vacuum system sits between the pole pieces of an Oxford Instruments N100 water cooled electromagnet. The magnet is shown (approximately to

*Cleaning is described in more detail in [53].




Figure 2.14: Vacuum chamber used to house the split-ring trap.

scale) in the upper left of figure 2.3. Up to 15 amp can be passed through each of the two coils*, providing a field between the poles of about 1 tesla. The iron pole pieces saturate, so that the field is not a linear function of the current. The field measured with a Hall probe between the pole pieces, as a function of the current through the coils (connected in parallel) is shown in figure 2.15. The pole pieces were 40 mm apart.

A much more accurate and precise method of measuring B involves measuring the motional frequencies of a small cloud of ions. Details of this method can be found in [51] and [53].

The coils each have a resistance of $6\ \Omega$ when warm ($4.5\ \Omega$ when cold), so a few kW of power is dissipated. Initially, the current was provided by a KSM SCT-220-20 forced air cooled high-power supply. This supply was used for many years, but it drifted (around one part per 10^3 per hour after warming up, almost a part in 10^2 per hour when cold); produced significant electrical (and acoustic) noise; was sensitive to electrical spikes; and eventually failed catastrophically. It was fixed in house several times[†], but has since been

*A higher current can be applied for a short period, but the temperature of the coils then increases dramatically.

[†]The repair work was performed by myself and H. Ohadi. It was very time consuming and mainly involved testing and replacing a large number of awkwardly mounted power transistors.




Figure 2.15: Magnetic field strength as a function of electromagnet coil current.

replaced by a much superior water-cooled current supply (Danfysik System 8000 Model 853). The Danfysik supply is digitally controlled (18 bit) and can provide up to 100 A at up to 105 V. It has a drift of ± 3 ppm over 30 minutes, ± 10 ppm over 8 hours [70]. Also, each of the (104) power transistors is individually fused, which allows for easier repair, and less explosive failure modes than the old supply.

For a Zeeman splitting of roughly 10 GHz, a drift of the magnetic field strength by a part in 10^3 or 10^5 gives a frequency drift of 10 MHz or 100 kHz respectively. Recalling that the linewidth of the 397 nm transition is 22 MHz, the KSM power supply was only just adequate, while the Danfysik supply is more than adequate, for laser cooling and spectroscopy on the 397 nm transition. However, for future spectroscopy of the narrow linewidth $S_{1/2} \leftrightarrow D_{5/2}$ transition the drift level of 10 ppm may be too large. Work is being done to move towards using a superconducting magnet or temperature stabilised permanent magnets. Permanent magnets offer high stability when temperature stabilised, and both have the additional advantage of lower power consumption. The group has operated a trap in a superconducting magnet already for Be^+ and Mg^+ , but not yet with Ca^+ . The combination of a $B \approx 1$ tesla field with good optical access means that the conventional electromagnet is still in use.

2.4.3 Fluorescence Collection and Detection

Figure 2.16 shows the lens system used to image light from a cloud of ions onto a PMT* or (by means of a flipping mirror) an ICCD camera†. Both the split-ring trap and the PCB trap use the same system, combining one singlet lens inside the vacuum chamber with another two (or three) outside. This seemingly unwieldy system is the result of the fact that the PMT (and camera) must be kept away from the large magnetic field in order to function properly. With the PMT about 50 cm above the magnet, the stray field has a small effect.




Figure 2.16: Schematic of the imaging system for fluorescence collection. All dimensions in mm.

When the mirror is down, an image of $\times 1$ magnification is formed 20 mm

*The PMT is a Thorn-EMI 9893QB. The photocurrent is converted to TTL pulses which are then time binned.

†The camera is an Andor DH-534:18S-03.

beyond the third lens, focused onto an iris or pinhole to reduce the background light. The light then passes through a filter to remove unwanted wavelengths, and hits the PMT photocathode. The filter consists of a low-pass filter with a cutoff at 435 nm (Comar 435IK25), and bandpass filter (Comar 395GB25). The filter transmissions at 397 nm are 70% and 80% respectively*. The iris can be adjusted between 0.8 and 10 mm diameter. When working with single ions, a 200 μm pinhole aperture gives a good compromise between background reduction and sensitivity to misalignment.

Flipping up the mirror sends the light to form another unmagnified image which is then re-imaged by a multiplet lens (Minolta MD 50mm F1.7) onto the ICCD camera. By moving the camera and the final lens the magnification can be varied.

All of the 40 mm focal length lenses are 25 mm diameter plano-convex lenses. The 150 mm focal length lens is bi-convex and has a diameter of 45 mm. The multiplet lens is a standard SLR camera lens. It has a poor 60% transmission at 397 nm due the many surfaces with no anti-reflection coating, but does offer low aberrations for low cost and complexity.

The detection efficiency, η , has not been measured directly, but can be estimated as: (solid angle fraction) $\times$ (lens and window transmission) $\times$ (filter transmission) $\times$ (PMT efficiency). The solid angle is limited by the diameter of the 150 mm focal length lens[†]. This lens, focusing on a virtual image at a distance of 240 mm (20 mm below the centre of the trap), is 45 mm diameter and limits the solid angle fraction to

$$\frac{\pi (22.5 \times 40/240)^2}{4\pi 20^2} = 0.0087. \quad (2.3)$$

The lenses and windows each have transmission of $\sim 90\%$. The transmissions of the filters are 70% and 80% at normal incidence, but the bandpass filter has a reduced transmission at other angles. The finite solid angle passing through the filter reduces the effective transmission to $\approx 50\%$. Assuming a

*This filter has since been replaced with a Semrock FF01-406/15-25 bandpass filter, with a transmission of 90%. The superior Semrock filter was initially purchased to filter out any 393 nm light present in the laser beam (see chapter 5).

[†]In future experiments this lens will be replaced with a larger one such that the solid angle is limited instead by the aperturing of the trap itself, as was mistakenly assumed in [51] and [52].

PMT efficiency of 0.2 (and an aperture transmission of 100%):

$$\eta \approx 0.0087 \times 0.90^4 \times 0.7 \times 0.5 \times 0.2 \quad (2.4)$$

$$\approx 4 \times 10^{-4}. \quad (2.5)$$

An error can be roughly estimated by combining error estimates on the various factors in equation 2.5. The fractional error on the detection efficiency is:

$$\frac{\Delta\eta}{\eta} \approx \sqrt{0.1^2 + 4 \times 0.03^2 + \left(\frac{0.05}{0.7}\right)^2 + \left(\frac{0.1}{0.5}\right)^2 + \left(\frac{0.05}{0.2}\right)^2} \quad (2.6)$$

$$\approx 0.35. \quad (2.7)$$

So $\eta = 4 \pm 1.4 \times 10^{-4}$. Some other factors which may affect the detection efficiency are:

- Optical misalignment and aberrations.
- Non-optimal PMT voltage setting.
- PMT performance reduction due to stray magnetic fields*
- PMT performance reduction due to dirt on the window, ageing, etc.
- Photons not being counted when more than one appears very close together in time[†].
- Photons not being counted due to imperfect transmission of the electrical pulses through the various elements (Discriminator/amplifier, ECL to TTL converter, splitter/buffer, counter).

*Using a different PMT, P. Ranin *et al.* found that the detection efficiency could vary by up to 25% depending on the orientation in the Earth's magnetic field [71]. As our PMT is close to a large magnet this effect could be even more pronounced.

[†]As we have no problem recording count rates of $\sim 10^7 \text{ s}^{-1}$, this effect is surely negligible.

Chapter 3

Single Ions and Small Crystals

In theory, the combination of single-qubit rotations and a universal two-qubit quantum logic gate is sufficient to build any quantum logic network [37]. Thus with the ability to control single ions and two-ion crystals, a Penning trap QIP scheme should be possible. A two-ion crystal along the axis of a Penning trap behaves in the same way as a similar structure in a linear RF trap and so the techniques developed for RF traps could be employed in Penning traps.

As of late 2005, when the work submitted in this thesis was begun, individual Ca^+ ions had never been laser cooled in a Penning trap. Several improvements have since been made, and it is now possible to load, trap, cool and image single ions and also two-ion strings/crystals in the Penning trap*. In this chapter, images of individual Ca^+ ions are presented. It is also shown that the unstable radial motion can be overcome to align a pair of ions along the trap axis. All the results presented in this chapter made use of the split-ring trap described in §2.4.

3.1 Loading Individual Ions

Before loading ions into the trap, the wavelengths of all the various lasers are checked to be at their correct values using the wavemeter described in §2.3.2 and the beams are aligned centrally in the aperture formed by the gap in the ring segments. One of the two 397 nm lasers is scanned from ~ 1 GHz below resonance, up to resonance, with a sawtooth function at about 1 Hz, while

*Loading single ions in the Penning trap is still not trivial. Some further improvements to the apparatus, as described in chapter 2, could be implemented to make single ion work more routine.

all the other lasers are kept at a fixed frequency. Doing this means that the laser interacts periodically with the very hot ions that would otherwise not be strongly cooled. This increases the cooling rate of a large, hot ion cloud. Also, imperfect laser tuning is less likely to cause rapid ion heating if one laser has a large red detuning for most of the time. It was found that scanning both lasers in phase with the same amplitude did not yield a significant improvement in the cooling. A large cloud of ions is loaded by running the atomic beam and electron filament simultaneously until some fluorescence signal is observed*. Initially the fluorescence level may be low but fine adjustment of the six laser frequencies and the two beam positions allows this signal to be optimised. A fluorescence level trace observed immediately after loading a cloud of ions is shown in figure 3.1. The peaks, corresponding to the laser scans, get higher and narrower as the cloud cools into the centre of the trap.




Figure 3.1: Fluorescence trace from loading a cloud of ions while scanning one laser.

The optimum beam positions and detunings depend on the size of the cloud, so smaller clouds of ions are loaded (by lowering the filament current) and the fluorescence level is re-optimised. When the laser parameters have been optimised, the scanning is stopped and the lasers can be locked.

For a very small cloud of ions the character of the fluorescence signal changes visibly due to the presence of quantum jumps. These are dark periods caused by ions becoming shelved in the metastable $D_{5/2}$ state (see chapter 5

*If the filament is operated at a very high current, then some stray light from the filament will be detected. No stray light is seen during a regular loading procedure.

for a discussion of this). The procedure for loading very small ion clouds is as follows: The oven is heated for 30 seconds; then, with the oven still on, the filament is heated for another 20 to 30 seconds, after which both are then switched off. At this stage there is usually no fluorescence detected, however an ion (or ions) may have been loaded into the trap. If the ion is in a large magnetron orbit, then it spends very little of its time in the focused laser cooling beam. It can therefore take a significant time for the ion to cool. As the ion is slowly cooled it moves closer to the centre of the trap and the cooling rate and fluorescence level then increase dramatically. It is not uncommon to wait several minutes for the fluorescence from a single ion to become visible above the background level. However, when this does happen the fluorescence increases to its maximum value very suddenly. Figure 3.2 shows the fluorescence rate during such a loading procedure. The trace begins just after the filament is switched off. After 35 seconds an ion cools to the centre of the trap and the fluorescence level rises to ~ 4000 counts per second. A second ion joins the first at $t \approx 70$ s. Occasional quantum jumps can be seen*.




Figure 3.2: Fluorescence rate after loading two ions. Time bins are 10 ms wide.

An interesting feature to note in figure 3.2 is the temporary loss of signal at $t \approx 70$ s. This is interpreted as being due to the second (hot) ion coming into the centre of the trap and temporarily heating the cold ion that is

*The laser at 854 nm was used to repump from shelved state. There are still many quantum jumps, but they are too short to be seen because of the resolution of the figure.

already there. The two ions then re-cool resulting in the subsequent two-ion fluorescence level.

Figure 3.3 shows histograms of fluorescence rates corresponding to different numbers of ions in the trap. The example in figure 3.3(d) shows that the number of ions (two in this case) can be simply determined by eye in real time.




Figure 3.3: Histograms from 40 seconds of data collection with 30 ms time bins show discrete levels of fluorescence rate for small numbers of trapped ions. (d) shows ten seconds of fluorescence (10 ms time bins) from two ions; two levels of signal (plus the background level) can clearly be seen.

3.2 Ion Crystals

When a small number of ions are trapped, they can undergo a number of distinct types of motion. If they are hot the motion of the ions is effectively uncorrelated and they collide with each other at random times. If the ions are cold they can form a ‘crystal’ such that the separation of the ions remains approximately constant. Work involving small numbers of ions in RF traps showed that the change between these two types of motion is abrupt

and may be thought of as a phase transition [72]. Also of interest, both for quantum computation and for developing a deeper understanding of simple quantum systems, is the transition between different shapes/orientations of small crystals of trapped ions. These have been studied theoretically in both RF traps [73] and Penning traps [74]. Experimental studies of these phase transitions have been performed in RF traps [75, 76, 77], and with relatively large crystals in Penning traps [78], inspiring theoretical interest in the use of ion crystals in Penning traps as quantum simulators [27]. The work presented below and in [12] shows control over very small ion crystals, paving the way for a Penning trap quantum computer similar to those planned in RF traps.

Consider two ions in a Penning trap. The orientation of the crystal that forms at low temperature depends on the external trap parameters. Applying a high axial potential forms the ions into a dumbbell shape in the radial plane (radial crystal). Due to the magnetic field this dumbbell shape rotates about the trap centre at a frequency close to the magnetron frequency, ω_- (the Coulomb repulsion between the pair of ions leads to a shift in ω_-). Strong cooling brings the ions closer together but their orientation does not change. On the other hand for trap voltages below some critical value, and as long as the magnetron motion is cooled effectively, it is energetically favourable for the dumbbell shape to form along the axis of the trap (axial crystal). In this case, especially under the influence of axialisation, each ion is expected to have its radial motion in the trap minimised. Two ions can line up along z if [74]

$$6\omega_0^2 < w_c^2. \quad (3.1)$$

This leads to an upper limit on the trapping voltage of

$$U < \frac{eR_0^2 B^2}{24m}. \quad (3.2)$$

For the split-ring trap at $B = 1$ tesla, $U < 5.5$ volt.

Since fluorescence is collected in the plane of the ring electrode, and an image can only be obtained over timescales much greater than $\frac{2\pi}{\omega_-}$, the images of a radial crystal are blurred into an elongated shape in the radial plane. On the other hand an axial crystal should appear as two well resolved spots in the image plane. Figure 3.4 shows images of two ions under the two different conditions, together with a single ion for comparison.

Figure 3.4(a) is a radial structure taken at a trapping potential of 3.5 V.




Figure 3.4: Images of ions taken at a magnification of ~ 4 . Magnetic field axis is vertical. A constant background level has been subtracted from every pixel.

The trap motional frequencies were $\omega_- = 21.8$ kHz, $\omega_0 = 125$ kHz and $\omega_+ = 357$ kHz. Although the axial trapping potential is below the critical voltage needed to align two ions along the axis, the radial confinement/cooling is not strong enough to form the ions into an axial crystal. Note that although the image is somewhat blurred the fluorescence forms two bright zones. This can simply be explained by projecting the fluorescence of an individual ion in a circular orbit onto the radial plane. The ion spends more time at the extrema of the projected motion and so the fluorescence appears brighter here. If the size of the ion orbit has a constant radius, r , and the fluorescence rate is constant, then the intensity incident on a pixel of the camera is

$$I \propto \sin^{-1} \left(\frac{X + \frac{\Delta X}{2}}{r} \right) - \sin^{-1} \left(\frac{X - \frac{\Delta X}{2}}{r} \right), \quad (3.3)$$

where X is the position of the pixel relative to the centre of the image, and ΔX is the width of a pixel. The pixel size of the camera is $13 \mu\text{m}$, with a minimum spatial resolution quoted as $22 \mu\text{m}$ ($\sim 4 \mu\text{m}$ and $\sim 6 \mu\text{m}$ respectively with our magnification). Figure 3.4(b) shows a theoretical image obtained by convolving the pattern given by equation (3.3) at an ion separation of $20 \mu\text{m}$ with a Gaussian to add some simple optical aberration.

Figure 3.4(c) shows two ions with a trap potential of 2.0 V. With the application of a weak axialisation drive (50 mV peak-peak at 376 kHz) to couple the unstable magnetron motion to the modified cyclotron motion, the ions form an axial crystal and the two spots apparent are genuinely the fluorescence from two different ions. In both cases the fact that two ions were present in the trap was corroborated by observing two-ion quantum jump traces using the PMT (e.g. figure 3.3(d)) before flipping the mirror to send light to the ICCD camera. It is possible to form an axial crystal without an axialisation drive, but the crystal formed is then much more sensitive to the positions of the laser beam foci and wavelengths than in the axialised case.

The expected distance between the pair of ions in an axial crystal can be estimated by balancing the repulsive force between the ions with the axial trapping force. If the ring electrode is at zero voltage and the endcap voltage is U , then the potential along the axis of the trap is

$$\phi \approx \frac{U}{2} \left(1 + \frac{z^2}{z_0^2} \right). \quad (3.4)$$

Thus the trapping force along z is

$$F_{\text{trap}} = -e \frac{\partial \phi}{\partial z} = -\frac{eUz}{z_0^2}. \quad (3.5)$$

The two ions are held symmetrically either side of the trap centre ($z = 0$). The distance between the ions, $2z$, is found by setting the total force on an ion to zero.

$$F_{\text{ion-ion}} + F_{\text{trap}} = 0 \quad (3.6)$$

$$\frac{e^2}{4\pi\epsilon_0 (2z)^2} - \frac{eUz}{z_0^2} = 0 \quad (3.7)$$

$$2z = \left(\frac{ez_0^2}{2\pi\epsilon_0 U} \right)^{\frac{1}{3}}. \quad (3.8)$$

For $U = 2$ volt and $z_0 = 3.5$ mm, the ion-ion separation is $2z = 26$ μm .

In another experimental sequence, two ions are loaded into the trap at a trapping bias of $U = 2.0$ V. This corresponds to an axial frequency of

$$\omega_0 = \sqrt{\frac{4eU}{mR_0^2}} = 2\pi \times 94 \text{ kHz}. \quad (3.9)$$

Images of the ions are obtained at intervals as the level of the axialisation drive amplitude is changed in steps. The results of this process are shown in figure 3.5 and more clearly in the supplementary animation in [12]. At low axialisation amplitude the ions move in the radial plane and their separation, which is related to the rotation frequency, is relatively large. As the axialisation drive strength is increased, the ions move closer together radially as expected. At a certain strength it becomes energetically favourable for the ions to change their orientation while keeping their separation constant [74]. Once the ions are aligned along the axis of the trap, their separation is fixed and does not change as the drive strength is increased further*. The transition is not expected to be abrupt because stable motion at an intermediate alignment is possible. The images seen are all consistent with the expected evolution of the system as the strength of cooling is increased, and the minimum observed separation in the images corresponds roughly to the expected value of 26 μm , although uncertainties in the precise value of the magnification prevent an

*However, when the drive strength is increased above about 1 V peak-peak, heating effects cause the crystal to fall apart.

exact comparison.




Figure 3.5: Images of a two-ion crystal as the axialisation voltage is varied. Magnetic field axis is vertical. Peak to peak axialisation voltage is shown in yellow. A constant background level has been subtracted from every pixel.

As mentioned at the beginning of the chapter, only single-qubit and two-qubit gates are required to build a universal quantum computer, so in theory just one and two ion processes will suffice. However, a realistic quantum computation scheme will likely make use of quantum error correction, decoherence free subspaces etc., which require several ancilla qubits for each primary qubit [24]. Quantum error correction has been demonstrated experimentally with three beryllium ions (one primary plus two ancilla) in a linear RF trap [79]. Thus despite what is theoretically possible using just two ions, it may be advantageous to use longer strings of ions in the processor region of an ion trap QC. Using the 1 tesla magnetic field, we were unable to align three ions along the trapping axis, however this should be possible if a higher field strength is used [74].

Chapter 4

PCB Penning Trap

This chapter describes a novel trap prototype containing three Penning traps. The idea behind this scalable trap has been published in [80]. The electrode structure allows ions to be trapped in multiple trapping zones, and also allows ions to be transported between the different trapping regions. Single ions have been trapped and clouds of ions have been transported from one trap to the next, and back again, with a return-trip efficiency of up to 75%. Although other groups have shuttled ions along the axis of a cylindrical Penning trap [81], this is the first report of ion transport from one Penning trap to another in a direction perpendicular to the magnetic field.

4.1 Design and Manufacturing

4.1.1 Boards

The PCB trap, as the name suggests, is made from printed circuit board. Two boards are mounted facing each other, with the magnetic field direction normal to the board surface. Copper pads on the board surfaces act as the trap electrodes. Close to the centre of the trap, the potential is very similar to an ideal Penning trap.

The electrode design is shown in figure 4.1. The full board design is shown in figure 4.2. Almost all of the board area is covered with copper so that any effects due to charge buildup on the insulating surface are kept to a minimum. The electrodes are arranged in rows of hexagons. The hexagonal pads in each row are electrically connected together with copper vias running through the board on to tracks on the rear side. Three hexagonal pads on each board,

labelled C in figure 4.1, act as the endcap electrodes. Each of these has a 1 mm diameter hole at the centre. Above and below the endcap electrodes are rows of pads labelled B and D . The remaining electrodes are labelled A and E . The trap is made from two parallel boards facing each other with a gap of 5 mm. The design on each is a mirror image of the opposing board. In normal operation the A , B , D , E electrodes on the pair of boards act as the ring electrodes for the three traps.




Figure 4.1: Upper part of the PCB trap board design, with electrodes labelled.

As shown in figure 4.2, there are seven 2 mm diameter holes below the trapping regions. Five of these are used to make connections to the five electrodes. These are connected to the vacuum feedthrough via wires bolted onto the rear of the boards. The outermost two holes are required to mechanically attach the trap to the rest of the superstructure.

The axialisation method mentioned in the previous chapter can be performed by putting the oscillating voltage on A and E , along with a signal of opposite polarity on B and D . Since symmetry of this trap is hexagonal instead of square as in the split-ring trap, a good quadrupole field is not created by simply applying out of phase signals between the electrode pairs. To create a good quadrupole field at the centre of the trap the magnitude of the axialisation voltage applied to B and D must be 0.62 times* the voltage applied to A and E .

The boards were manufactured in the Imperial College Biomedical Engineering PCB Prototyping Laboratory. An LPKF H60 CNC Milling/Drilling machine was used to drill the holes, then the boards were electroplated with copper using an LPKF through-hole plating system. This process connects

*This value was calculated with numerical simulations using SIMION 7.




Figure 4.2: PCB trap board design to scale, with dimensions in mm. The green side faces into the trapping region, the orange side is the rear. The x axis is horizontal while the y axis is vertical. Note that a mirrored version is used for the opposing board.

the front and rear sides of the board through the holes. The CNC machine was then used to mill away gaps between the electrodes. The gaps are approximately $200\text{ }\mu\text{m}$ wide. Finally, the same machine was used to saw around the edges of the boards to produce the design shown in figure 4.2. An A4 size piece of copper coated FR4 board (1.8 mm thick) was used, so four pairs of trap boards could be produced in a single run*.

4.1.2 Superstructure and Vacuum Chamber




Figure 4.3: Diagram of the PCB Trap. Some components (screws etc.) have been excluded for clarity.

*Only one pair was actually used.




Figure 4.4: Photographs of the PCB trap. Background has been removed digitally.

A labelled diagram of the trap is shown in figure 4.3. Photographs of the assembled trap are shown in figure 4.4. A pair of copper legs is bolted to a CF40 flange with an 11 pin electrical feedthrough. Stainless steel plates are attached to the copper legs. The steel plates hold the trap boards, the filament and oven, and the lens holder. All of the screws, metal washers and nuts are size M2 non-magnetic stainless steel.

The boards are held 5.0 mm (± 0.01 mm) apart by stainless steel spacers. The spacers also provide an electrical connection between the pair of boards. The spacer which would connect electrode C on each board was not used so that the endcaps on each side could be connected independently*. The spacer connecting electrode E was also removed after discovering that this reduces the laser scatter. Kapton coated wires are used to connect the electrodes to the feedthrough pins underneath the trap. Each wire is hooked at the end, and fastened tightly between a washer and a large copper area on the rear of one of the boards. The other end of each wire is attached to a feedthrough pin using a UHV crimp connector. Kapton wires are also used to connect to the oven and the filament.

The 5.0 mm spacing between the pair of boards was chosen to make the trapping potential as harmonic as possible. Figure 4.5 shows the potential along the axis of the middle trap, as computed numerically using SIMION 7†. Close to $z = 0$, the potential is quadratic. The separation between the boards was varied (in steps of 0.1 mm), and the potential was calculated for each case. A function

$$\phi = a_0 + a_2 (z - z_0)^2 + a_4 (z - z_0)^4 \quad (4.1)$$

was fit to the simulated potential (with a_0 , a_2 , a_4 and z_0 as free parameters) over the range $|z| \leq 1$ mm. For a purely quadratic potential (as in an ideal Penning trap), the quartic coefficient, a_4 , would be zero. Figure 4.6 shows how the quartic component of the potential varies. The fits were performed using gnuplot‡. It can be seen that the axial potential close to the trap centre is

*Having separate control of the endcap electrodes provides the opportunity for exciting the axial motion of ions. This technique can be used to detect ions purely electronically. Also, a DC bias can be applied to move the centre of the trap, which may be useful for compensating for any charge buildup, or for reducing micromotion if RF trapping is used.

†SIMION® is a software package primarily used to calculate electric fields and the trajectories of charged particles in those fields when given a configuration of electrodes with voltages and particle initial conditions.

‡Gnuplot is a command-driven interactive and scriptable function plotting program.

almost purely quadratic when the inter-board separation is around 5 mm. The characteristic trap dimension in this case can be estimated as $R_0^2 = 2/a_2 = 26.3 \text{ mm}^2$.




Figure 4.5: Potential along the central axis of the PCB trap. The separation of the boards is 5.0 mm. The simulation grid size was 0.1 mm.




Figure 4.6: Quartic coefficient of the potential along the central axis of the PCB trap for $-1 \leq z \leq 1$. The error bars show the uncertainty in the fit parameter.

A schematic of the atomic beam oven holder is shown in figure 4.7. Each steel plate has a pair of holes, 12 mm apart, countersunk on both sides. M2 alumina washers sit in the countersinks, with steel washers outside to spread the force of the nut/screw. The alumina washers prevent each metal screw

from contacting the steel plate (which is grounded to the vacuum chamber). Each side of the oven is held between a pair of metal washers, as are the wires which connect the oven to the feedthrough. The filament is connected in an identical way on the opposite side of the trap. A 1.0 mm diameter hole in the steel plate, midway between the screws, is concentric with a 1 mm diameter hole in each of the trap boards. These holes allow the atomic beam (and electron beam) to pass through the centre of the trap. When the trap is aligned correctly with the magnetic field, the ionising electron beam passes all the way through and hits the oven. Detecting current on the oven can be used to measure the strength of the electron beam, as mentioned at the end of §2.2.1, and also to aid in the process of aligning the trap with the B field.




Figure 4.7: Schematic of the system used to hold the atomic beam oven in place. The electron beam filament is held in an identical way.

The lens is a 40 mm focal length singlet lens. It is positioned 20 mm above the centre of the middle trap, such that the optical system is very similar to that described in §2.4.3. Unfortunately this only allows the middle trap to be imaged. The original design used three lenses (6 mm diameter, 9 mm focal length) – one above each trap – so that ions in any of the traps could be imaged. The old design was tested several times and no ion fluorescence was ever seen above the background level, so the system was made more similar to the split-ring trap system which was known to work. A piece of constantan foil with a 6 mm diameter hole is used as a baffle above the trap to block scattered light coming from areas not directly underneath the lens.

FR4 (flame retardant 4) board is a composite of an epoxy resin reinforced with a woven fibreglass mat. The vacuum properties of FR4 board have been studied, and the outgassing rate found to be low enough for UHV applications [82]. The resin in the board material changes from a glassy to a plastic state (glass transition) at a temperature of between 120 °C and 180 °C depending on the particular resin chemistry [83]. This limits the temperature at which the vacuum system can be baked. In order to improve the final vacuum, the chamber was pumped out for a week at 250 °C with a blanking flange used instead of the trap feedthrough. The chamber was then cooled, vented with dry nitrogen, and within an hour the trap was bolted into place and the air pumped out. The whole system including the trap was then baked at 120 °C for another week. The vacuum chamber containing the PCB trap is very similar to that of the split-ring trap. It is shown schematically in figure 4.8. A 20 l/s ion pump (Meca 2000) was used, in combination with a Getter pump (SAES Getters CapaciTorr D 400). The getter pump has a particularly high pumping speed for hydrogen (~ 100 l/s). No ionisation gauge or leak valve was used. The final pressure, as measured by the ion pump, reached below 10^{-9} mbar (below the bottom of the scale).




Figure 4.8: Vacuum chamber used to house the PCB trap.

4.2 Simulations


To shuttle ions from one trap to another, the voltages on the various electrodes are switched. By applying ascending (or descending) voltages to A , B , C , D and E , an approximately linear electric field can be produced, perpendicular to the magnetic field direction. When the linear electric field is applied, an ion initially at rest will start to move along this direction but be pulled round in a cycloid loop by the B field. By appropriately choosing the magnitude of the electric field, the size of the cycloid loop will be the same as the spacing between adjacent traps. By applying these voltages for the correct amount of time ($\sim 2.5 \mu\text{s}$), an ion will move out of one trap and come to rest at the centre of the next trap. Applying the opposite voltages for a similar duration will cause an ion to hop between traps in the opposite direction.

Figure 4.9 shows the electrostatic potential in the mid-plane between the pair of the trap boards, when the voltages are set to ‘trap mode’ and to ‘shuttle mode’ respectively. The typical trapping potential is formed when the voltages on the electrodes are $(V_A, V_B, V_C, V_D, V_E) = (0, 0, 2.5, 0, 0)$ volt. Note from figure 4.9(a), that even though the electrodes are hexagonal, the shape of the potential close to the centre of the trap is almost perfectly circular. The potential in figure 4.9(b) is produced when $(V_A, V_B, V_C, V_D, V_E) = (20, 12.5, 2.5, -7.5, -15)$ volt*. Figure 4.10 shows an example of the trajectory of an ion as it moves between a pair of trapping regions.


In the simulations the electrode voltages, time durations, initial conditions, etc., can be made unrealistically perfect. To learn roughly how sensitive the ion shuttling system is to experimental imperfections, various parameters were changed and the effect on a simulated ion was observed. The simulations were performed numerically using SIMION 7. The electrode geometry definition file was generated using a Perl script which is listed in appendix B.

The data points in figure 4.11 show how far an ion misses the centre of the target trap if one of the voltages is not correct. The ion begins at rest at the centre of the first trap. Electrodes A , B , C and D were set to 20, 12.5, 2.5 and -7.5 V respectively. Since the ion passes closest to electrode E during most of its trajectory, a small change in V_E has a bigger impact on the path of the ion than a small change in any of the other voltages does. It can be seen that 1 volt of error in V_E causes an ion to miss its target by ~ 0.5 mm.

*By keeping the voltage on C unchanged between the two modes, only four switching channels are needed instead of five.



(a) Trapping potential



(b) Shuttling potential

Figure 4.9: Simulated electrostatic potentials at the mid-plane between the two trap boards. Note the change in scale of the vertical axis.




Figure 4.10: Example of an ion trajectory while moving from one trap to the next. Lines of equipotential in the mid-plane between the two boards are shown in blue. The x axis is horizontal while the y axis is vertical.




Figure 4.11: Residual position of an ion relative to the centre of the second trap, immediately after shuttling. $B = 0.9$ tesla.




Figure 4.12: Residual kinetic energy of an ion, immediately after shuttling an ion in a field of 0.9 tesla.

Figure 4.12 shows how an ion gains kinetic energy if the shuttling voltages are applied for too long or too short a time. The voltages were set appropriately to move the ion to the centre of the second trap in a single cycloid loop. It can be seen that the kinetic energy acquired by an ion varies quadratically with $t - t_{\text{ideal}}$, and reaches roughly 100 meV when the duration is wrong by 100 ns ($\sim 4\%$ of the pulse duration).

There is a range of possible sets of voltages which will cause an ion to move between the centres of adjacent traps. For example, at a magnetic field of 0.9 tesla, $(V_A, V_B, V_C, V_D, V_E) = (16.3, 16.3, 2.5, -13.8, -13.8)$ will produce a good trajectory, as will $(V_A, V_B, V_C, V_D, V_E) = (18.7, 2.5, 2.5, 2.5, -16.2)$, and even $(V_A, V_B, V_C, V_D, V_E) = (2.5, 2.5, 2.5, 2.5, -19.7)$. When the ion is initially at rest and perfectly centred in the first trapping zone, all of these different voltage sets, and everything in between will produce good results. However, if the ion begins at a position slightly offset from the trap centre – or equivalently, the trap is not perfectly aligned to the B field – then this initial imperfection is amplified by different amounts depending on the specific set of voltages used. Finding the best possible parameters is a non-trivial problem. It would involve finding a line in four dimensional voltage-space (V_A, V_B, V_D, V_E) for which an ion with ideal initial conditions reaches the centre of the target trap, and then finding the point on this line which produces an optimal final condition given a range of realistic initial conditions. To simplify matters, we consider a set

of voltages which are symmetric about electrode C

$$\begin{pmatrix} V_A \\ V_B \\ V_C \\ V_D \\ V_E \end{pmatrix} = \begin{pmatrix} 2.5 + a \\ 2.5 + \epsilon a \\ 2.5 \\ 2.5 - \epsilon a \\ 2.5 - a \end{pmatrix} \quad (4.2)$$

and also a set which is not symmetric

$$\begin{pmatrix} V_A \\ V_B \\ V_C \\ V_D \\ V_E \end{pmatrix} = \begin{pmatrix} 0.0 \\ 0.0 \\ 2.5 \\ 2.5 - \epsilon a \\ 2.5 - a \end{pmatrix}. \quad (4.3)$$

The asymmetric set has the advantage that only two electrodes need to be switched instead of four in order to perform a single jump.

The simulation is scanned over a range of values of ϵ . For each ϵ , an optimal value for a is found (along with an optimal pulse duration), and then some measure of the final condition of the ion can be found, for a particular initial condition. The best values of a and of the pulse duration are shown in figure 4.13. The duration of a hop is defined as the time between the application of the shuttling voltages and the next minimum in the ion's kinetic energy. The best value of a is found by minimising the distance along x between the ion and the centre of the target trap, immediately after a hop.

Figure 4.14 show how different values of ϵ lead to different final conditions of an ion. The initial condition is the ion at rest but displaced by 0.1 mm along the axis of the first trap. The final position of the ion relative to the centre of the destination trap is plotted, along with the final velocity, the final kinetic energy, and final total energy. All of these measures should ideally be as low as possible.

It was found that the symmetric voltage set, (4.2), produced better results (lower final KE etc.) than the asymmetric set, (4.3). The asymmetric voltage sets cause the ion to gain significantly more kinetic energy. It was also found that the best results were obtained when $\epsilon \approx 0.5$ (or slightly higher) – i.e. a linear step down of the five electrode voltages.




Figure 4.13: Simulated optimal voltage and timing parameters for $B = 0.9$ tesla. The true cyclotron period is also shown as a red line in (c) and (d).




Figure 4.14: Final conditions of an ion after shuttling for various voltage sets. The ion was initially at rest and displaced 0.1 mm along z . $B = 0.9$ tesla.

Another asymmetric voltage set

$$\begin{pmatrix} V_A \\ V_B \\ V_C \\ V_D \\ V_E \end{pmatrix} = \begin{pmatrix} 2.5 + a \\ 2.5 + \epsilon a \\ 2.5 \\ 0.0 \\ 0.0 \end{pmatrix} \quad (4.4)$$

was also simulated. This set gave worse results (much larger residual KE and position) than either of the two sets shown above. It also requires much larger voltages (50-100 volt) to be applied to the electrodes*.

An initial displacement of the ion along x or y has less effect on the outcome than an initial displacement along z . It can be seen that for typical voltage sets there is an overall restoring force pushing ions towards $z = 0$. The z component of the force on an ion integrated over a full trajectory is shown in figure 4.15. Note that the stronger force at low ϵ does not indicate a more stable potential, but that there is a larger z displacement along the ion trajectory. Unfortunately, even with this restoring force present, any initial displacement or initial velocity will cause an ion to gain energy when moving between traps.




Figure 4.15: Force on an ion along z , integrated over a shuttling trajectory, for various symmetric voltage sets. The ion was initially at rest and displaced 0.1 mm along z .

Since the ion moves perpendicular to the magnetic field, a displacement

*The ion moves closer to the negative electrodes as it is shuttling, so electrodes with positive voltages produce a relatively small effect on the ion's trajectory.

of the ion relative to the trap centre along z is similar to a misalignment of the trap with the magnetic field. An initial displacement of 0.1 mm along z is roughly equivalent to a misalignment of just half a degree. Clearly this misalignment must be reduced as much as possible when attempting to shuttle ions in the real trap.

4.3 Electronics

In order to produce the voltage pulses needed to shuttle ions between traps, a system of high speed switching electronics was developed. Each electrode (except for electrode C) needs to be switchable between 0 V (or close to 0 V) for trapping, a positive voltage of up to around 20 V for shuttling ions in one direction, and a negative voltage of similar magnitude for shuttling in the other direction. The voltages need to be set precisely, and remain stable over time. The switching needs to occur on a timescale which is short compared to the length of the pulse. The pulse length must be set precisely and must also remain stable over time. There should also be a negligible relative time delay between the pulses on the various electrodes.

To generate pulses of a precise duration, a commercial digital pulse generator (Stanford Research Systems DG535) is used to produce TTL pulses. This pulse generator has a resolution of 5 ps, a jitter of ~ 50 ps, and an accuracy of ~ 500 ps (500 ps typical, 1500 ps max). The TTL pulses are then converted into pulses of the required voltages. The digital pulse generator is triggered from a LabView program via a National Instruments card. The generalised layout is shown in figure 4.16.




Figure 4.16: Generalised layout of the system used for shuttling ions.

4.3.1 Output Stage

Figure 4.17 shows one of the output stages of the pulse converter. There are four of these in total. The OUTPUT PULSE is connected straight to one of the trap electrodes. The POS V and NEG V are connected to positive and negative voltage sources, discussed below. The capacitors are $3.3\ \mu\text{F}$ ceramic chip capacitors in parallel with larger $100\ \mu\text{F}$ electrolytics. The mosfet transistors are BSP318 (N type) and BSP315 (P type). RF chokes (L1 and L2) are used to protect the voltage sources from the capacitive loads (which were found to cause oscillations and overheating) and also from the transient loads.




Figure 4.17: Simplified schematic of the output stage of the pulse converter.

In ‘trap mode’, both the POS CONTROL and NEG CONTROL channels are low (close to 0 V), and all the transistors are switched off. When the POS CONTROL goes high ($\sim 5\text{ V}$), Q1 is switched on. This brings down the gate of Q2 which then also switches on, and the output is pulled up to POS V (assuming the on-resistance of Q2 is much less than R2). When POS

CONTROL drops back down to 0 V, the transistors turn off and the output is pulled back to ground. The turn off time is slower than the turn on time, and is limited by current flowing through R1 to charge up the gate capacitance of Q2, and also through R2 to charge up the capacitance of the trap (and cable and feedthrough). Figure 4.18 shows the output voltage for various values of R2 when a 2 μ s TTL pulse is applied to POS CONTROL. It can be seen that when R2 is large, the fall time of the output pulse is too long. When R2 is very small, C1 begins to significantly discharge before the pulse has finished and the voltage sags*. Also, for very small values of R2 the wasted power is high and there are errors due to the finite on-resistance of Q2. Resistances of around 20 to 100 Ω lead to good output pulses with rise and fall times below 50 ns. Metal film 1% resistors of 47 Ω were used for R1, R2, R3 and R4.




Figure 4.18: Example 20 volt pulses for various resistances, R2. Graphs are offset for clarity. The apparent voltage noise appears similar when measuring a clean 20 volt DC source, so the noise is most likely due to RF pickup and/or imperfections in the oscilloscope.

The R3-Q3 combination acts as an nMOS logic style NOT gate. So when the NEG CONTROL input is pulled up to 5 V, Q3 is switched on, causing the switch on of Q4 and then Q5. This pulls the output down to NEG V.

It should be noted that POS CONTROL and NEG CONTROL should never be set high at the same time. This would short POS V to NEG V, and possibly destroy Q2 and Q5. Also, each of the trigger inputs should not be

*Even with R2 as low as 5 Ω , the 100 μ F electrolytic capacitor should be large enough to prevent significant sag. However the equivalent series resistance and inductance of the electrolytic capacitor prevent it from responding well at high frequencies.

high for an extended amount of time*. For this reason, the gates of Q1 and Q3 are tied to ground with large resistors (not shown in figure 4.17).

Another addition, not shown in figure 4.17, is that one side of R2 can be connected to the output of a high speed analogue buffer, instead of to ground. The input of the buffer is connected to the small AC voltage of the axialisation drive mentioned in chapter 3.

4.3.2 Voltage Sources

There are eight different DC voltages required[†]: POS V and NEG V for each of the electrodes *A*, *B*, *D* and *E*. When hopping to the left, *A* and *B* are given positive voltage pulses while *D* and *E* are given negative pulses, and vice versa for hopping to the right. Assuming everything in the trap is perfect, POS V *A* = POS V *E* etc., and only four voltages are really required. However, since a real trap is never perfect, it is desirable to have the capability of adjusting the similar voltages separately.

Figure A.3 in the appendix shows the complete schematic of the DC voltage control board. Figure 4.19 shows a simplified picture of the circuit, containing only the section used to control NEG V *A* and NEG V *E*. A 128 step, 10 k Ω digital potentiometer (IC1) is connected to a stable 5 volt reference (REF02) to produce a tunable voltage. The wiper of IC1 is loaded by another pair of similar potentiometers, with resistors R1 and R2 such that a fine ($\div 10$) voltage control is provided for *A* and *E* separately. The voltages are buffered and then inverted, before being amplified by a factor of 6. This provides linear control between 0 and -30 volt, with a precision of ~ 20 mV. On the positive voltage side, the buffer and inverter are simply not used. The output op amps (IC9) are powered from single ended supply rails: 0 and +30 V for the positive, 0 and -30 V for the negative.

The 12 digital potentiometers are set from a LabView program via a National Instruments card. They share common up/down (U/D) and increment (INC) signals. The chip select (CS) signal is demultiplexed on the board[‡] so that the computer need only provide 6 signal lines instead of 14.

The results presented in §4.4 used an earlier version of this voltage control board. The earlier version did not have computer controlled potentiometers,

*Trying to produce pulses with durations on the order of milliseconds to seconds causes the voltage sources to shut down from overheating or too high a current draw.

[†]Excluding the trapping bias on electrode *C* and the filament bias.

[‡]A 74154 4-to-16 TTL demultiplexer was used.




Figure 4.19: Partial schematic of voltage source generator.

and used a less stable 78L05 voltage regulator instead of the REF02. Otherwise both versions were very similar.

4.3.3 Control Pulse Demultiplexer

Referring back to figure 4.17, a TTL pulse on the NEG CONTROL line causes three transistors to turn on, while a pulse on the POS CONTROL line turns on only two transistors. It was found that this inconsistency* causes the negative pulses to be roughly 100 ns longer on average than the positive pulses. The negative pulses also have a slightly longer delay between the rising edge of the control pulse and the rising edge of output pulse.

The digital pulse generator used to generate the precise TTL pulses has two output channels which can be set independently. Utilising this feature, the positive voltage and negative voltage pulses are separately triggered from the two different channels. An adjustment can be made to the relative timing of the two TTL pulses, to compensate for the greater delay times of the negative output pulses. So for example, TTL pulses are sent to POS CONTROL A, POS CONTROL B, NEG CONTROL D and NEG CONTROL E all at roughly the same time, but the NEG CONTROL pulses are set to come

*Another inconsistency between the positive and negative pulses is caused by the physical differences between the N type and P type mosfets.

slightly earlier and be slightly shorter than the POS CONTROL ones.

To trigger an ion hop, a digital output line from a National Instruments card first selects the direction of the hop (LEFT/RIGHT). The same card then sends a trigger pulse (of any duration) to the digital pulse generator, which sends two pulses (CLK POS and CLK NEG) to the pulse converter electronics. These three signals are processed, using fast TTL logic gates, into the eight signals required for the POS CONTROL and NEG CONTROL of the four different channels (A , B , D , E). The circuit to perform this is shown in figure 4.20. The AND gates are contained in a 74F08 fast TTL logic IC. The NOT gate does not need to be high speed, and is a simple nMOS type gate consisting of a 2N7000 transistor and a 10 k Ω resistor.




Figure 4.20: Schematic of the control pulse demultiplexer.

As well as the small unwanted time differences between the positive and negative pulses, there are also small differences between the four (nominally identical) channels. There is approximately 50 ns difference between the shortest and longest pulses. The magnitude of this timing inconsistency is similar to the rise and fall times of the pulses, and was not corrected for. Ideally, a precise TTL pulse generator with four separate channels would be used to compensate for all of these timing errors. Alternatively, many transistors could be individually tested and matched before being used in the circuit.

The real pulses, as measured using a voltage probe on the pins of the feedthrough plug, were loaded into SIMION. Even with the imperfections of

reality – such as finite rise and fall times, overshoot, pulse duration inconsistencies and voltage noise – virtual ions were still shuttled reliably between traps in the simulation.

4.4 Results

Figure 4.21 shows the fluorescence from a large cloud of ions in the central trapping zone. After a few seconds, $2.59 \mu\text{s}$ pulses are applied and the ion cloud is moved into a different trap. The signal level drops to the background level because fluorescence is only collected from the central trap. After a few more seconds, the reverse set of voltage pulses are applied, and most of the signal returns. As the process is repeated, ions are lost on each journey, but some ions still remain after 20 hops (10 return trips). The voltages used were $(V_A, V_B, V_C, V_D, V_E) = (20.0, 12.5, 2.5, -7.5, -15.7)$ volt, and $(V_A, V_B, V_C, V_D, V_E) = (-15.7, -7.5, 2.5, 12.5, 20.0)$ volt.




Figure 4.21: Fluorescence from a cloud of ions, repeatedly shuttled away from the central trap and back again.

An interesting feature can be seen each time a medium sized cloud is moved back into the central trap. First a very small amount of signal returns instantly (in less time than one bin width). The signal level then rises fairly sharply (within around 1-3 seconds), but plateaus out and remains quite low for up to 20 seconds, after which the signal level sharply rises again to its maximum value. This behaviour does not seem to occur for very large clouds ($\gtrsim 10^4$ ions) or for very small clouds ($\lesssim 10$ ions). A possible explanation

for this is as follows: When an ion cloud is moved between traps, the cloud expands* and the ions become much hotter. Some small fraction of the cloud however, will land within the laser beam waist and be cool enough to produce fluorescence straight away. The cyclotron and axial† motions of the ions are strongly cooled, and so the signal level rapidly increases. The cloud then gradually shrinks as energy is pumped into the magnetron motion (the laser beam was offset to facilitate this effect, but no axialisation was applied). The smaller and denser the cloud becomes, the more strongly it interacts with the laser beam, so there is a runaway effect causing the final signal increase to be rapid.

The final signal level after each hop, plotted as a function of the number of hops, gives a remarkably good fit to a decaying exponential. The free parameter of this fit gives the efficiency of the shuttling procedure. Referring back to figure 4.11, changing one of the voltages away from its ideal value will cause an ion to land away from the centre of the target trap. The same experiment was performed in reality, altering one of the voltages and observing the efficiency. The results are shown in figure 4.22, where the simulated residual displacement is also shown. Two shuttling sequences were performed for each data point. The error bar is simply the difference between the two. It can be seen that the measured efficiency does indeed peak where the simulation predicts the optimal voltage.

Although single ions have been trapped and laser cooled for extended periods of time, single ions have not yet been reliably shuttled. To attempt to push the shuttling efficiency up above 75%, two improvements were made: A new electronics system was built, with faster rise and fall times of ~ 10 ns. Secondly, coils were wound to produce a small magnetic field perpendicular to the main field, thus allowing the angle of the field to be finely adjustable (§4.5.1). Unfortunately, the electron beam filament failed before we were able to test the new improvements. They will be tested in the near future, but time constraints prevent those results from being reported here.

*The cloud expands due to Coulomb repulsion, but also because ions which are initially offset from the centre of the first trap will gain a larger offset at the destination trap just due to the imperfect shuttling field.

†The laser is perpendicular to the axis of the trap, so the axial motion will most likely cool slower than the cyclotron. The relative cooling rates of the axial and cyclotron modes have not been measured, so it could be that in fact the hot axial motion is limiting the rate of signal increase.




Figure 4.22: Efficiency of the return trip shuttling procedure, for various voltages.

4.5 Future Improvements

4.5.1 Shim Coils

Recalling from §4.2, a small angle between the trap axis and the magnetic field direction causes an ion to gain energy as it is shuttled between traps. The vacuum chamber was physically adjusted between the pole pieces of the magnet such that the electron beam current passing through the trap from the filament to the oven was maximised. To reach the oven, the electron beam passes through several concentric holes of 1 mm diameter. The holes closest and furthest from the filament are 10 mm apart. Thus in the worst case scenario, the angle between the trap and the B field could be up to 0.1 radian. An angle this large causes an ion to gain at least 0.1 eV after one hop. If the ion is not cooled effectively in the second trap*, then this error can be amplified further when it is shuttled back to the first trap.

To fine tune the orientation of the field, coils have been added outside the vacuum chamber such that a small field perpendicular to the main trapping field could be produced. The shim coils are connected to relays so that they can be switched on and off with computer control.

*Although the Raleigh range of the laser cooling beam is similar to the inter-trap spacing, the cooling parameters are only optimised for the central trap. Any misalignments or magnetic field inhomogeneities will lead to poor heating in the outer two trapping zones.

4.5.2 Electronics

A second pulse converter board was made, to produce pulses with faster and more consistent rise and fall times. The schematic of the second generation pulse converter is shown in figure A.4 in the appendix. A simplified schematic of one of the output channels is shown in figure 4.23. Instead of using a single ended nMOS style output (one resistor and one mosfet), a CMOS style output is used (one N-type mosfet and one P-type) to provide a lower output impedance. Smaller transistors (2N7002 and BSS84) were used, with larger on-resistances, but smaller gate capacitances and hence faster turn on and turn off times. Relays are used to switch between the positive and negative outputs. Relays are also used to switch between the shuttling pulse voltage and the axialisation voltage.




Figure 4.23: Simplified schematic of one channel of the second generation pulse converter.

The pulses produced by this circuit have rise and fall times of ~ 10 ns, compared to ~ 50 ns. The timing inconsistencies between the different channels are smaller, as is the timing inconsistency between the positive and negative voltage pulses. On the other hand, the faster pulses have significantly greater overshoot and ringing than the old pulses. Hopefully in the near future this new system will be tested and compared to the original prototype.

If a new PCB based trap is to be built in the future, the switching elec-

tronics could be soldered directly onto the trap board. If this is done then the stray inductance and capacitance would be very low indeed. There would be no imperfections due to impedance mismatch between the electronics, the cable, the feedthrough, and the trap. There may be UHV issues if there are electronic components inside the vacuum chamber, but they could perhaps be coated with UHV compatible glue if the outgassing was found to be too problematic.

FR4 board has a dielectric constant in the range 4.4 to 5.2. At microwave frequencies, dielectric losses make FR4 unsuitable [83]. If sub-nanosecond switching is required, then other materials such as Rogers RO4000 board* may need to be used.

4.5.3 Prototype Design

Alumina is a good insulator, UHV compatible, high-temperature compatible, and harder than some ceramics such as Macor[†]. Even though alumina is a fairly hard ceramic, the washers used to insulate the oven and filament holders are only 0.3 mm thick and can very easily crack if too much pressure is applied. The method used for holding the oven and filament is very good because each is held firmly in place and will not move significantly when heated (unlike in older traps such as the split-ring trap). However if a screw experiences too much force, then a washer will crack – most likely causing a short to ground as well as catastrophic misalignment. A suggested improvement would be to simply use larger ceramic washers. Size M3 or even M2.5 washers are significantly thicker, and the larger surface area would allow the same force to be applied with less pressure.

As mentioned in §4.1.2, the trap was originally designed with three separate lenses in the vacuum system, one for each trapping zone. The ability to collect fluorescence from all of the traps has many advantages. One important advantage is the possibility of optimising the shuttling parameters for a single jump separately in each direction. At the moment this cannot be done straightforwardly because fluorescence returns only after two shuttling procedures. A suggestion for a similar trap built in the future would be to

*Rogers RO4000[®] series high-frequency circuit materials are glass reinforced hydrocarbon/ceramic thermoset laminates designed to offer superior high-frequency performance. RO4000 also has a higher glass transition temperature ($\gtrsim 280^{\circ}\text{ircC}$) than FR4.

[†]MACOR is a machineable glass-ceramic developed and by Corning Incorporated. Unlike alumina it can be machined into any desired shape using standard metalworking bits and tools.

have two separate fluorescence collection systems available at the same time: one similar to an old tried and tested system, and another new system with the capability of imaging multiple trapping regions. On the other hand, if the size of the trap is reduced, it should be possible to collect light from several trapping zones using just one lens.

4.6 Outlook and Comparison With Other Work

This work can be compared to similar work performed using segmented RF traps. Huber *et al.* reports an efficiency of 99.0% for transporting single ions over a distance of 2 mm and back in a round trip time of 20 μs , using a segmented PCB based linear Paul trap [46]. Hensinger *et al.* reports 100% efficiency (881 out 882 attempts) for transporting ions round a T-junction corner in 30 μs in a trap made from gold and alumina [50].

One advantage of our system (if the technical difficulties of reliably transporting single ions can be overcome) is that the shuttling time is lower by about an order of magnitude (using a 1 tesla field). If a higher field (or a lighter ion) is used, sub μs transport times would be possible*.

Another clear advantage of our system is that the switching of electrode voltages is relatively simple. Transport of ions between trapping zones in an RF trap has so far required a much more complicated set of precisely controlled analogue voltage ramps to be applied to the various electrodes.

For the trap described in this chapter, the best set of shuttling voltages was an ascending series over all five electrodes. However it may be possible to build an even simpler trap with just three electrodes: a row of endcaps, and two ring type electrodes. The shuttling could be performed by pulsing just one of the ring electrodes, and pulsing the other one to go in the other direction. Although this goes against the simulation results of §4.2, the electronics would be greatly simplified. No negative voltages or tri-state switches would be required. If a similar but smaller trap is built, then the required shuttling voltages would also be smaller. If the trap was designed such that the shuttling voltages were ~ 5 V, then the pulses could be produced directly by high performance TTL, CMOS or ECL devices. The shuttling voltage would be tuned by tuning the power supply of the fast logic device itself. This would not

*On the other hand, future generations of smaller, tighter RF traps would also allow reliable shuttling at faster speed. An 80% efficiency for a single trip performed in just 3 μs was reported in [46].

only simplify the electronics, it would also allow the incredible performance of modern digital ICs to be directly utilised.

A further advantage of the scalable Penning trap concept is seen when considering the transportation of ions round corners. Although a near perfect efficiency is reported by Hensinger *et al.*, they state that in their simulations, an ion acquires about 1.0 eV of kinetic energy during a corner-turning protocol [50]. Real ions in our system certainly gain some energy after shuttling, but at least in theory the gain in energy can be extremely small. Moving an ion round a corner in a multiple Penning trap would involve nothing more than performing two regular cycloid loops at right angles to each other. Alternatively, a corner could be a movement along the trap axis between two axially aligned traps, followed by a cycloid loop perpendicular to the magnetic field. A next generation prototype of the PCB Penning trap should certainly make use of this possibility.

The electrodes of a future trap should be designed in such a way as to facilitate ion transport round a corner in the radial plane, and possibly also allow ion transport along the B field direction*. Building a three dimensional (or even just two dimensional) trap array, with the ability to move ions in a controlled manner between each trapping region, would certainly be a challenging technical feat, but the results of this chapter have demonstrated the first proof of principle of the scalable Penning trap. It may be the case that a large scalable Penning trap would be simpler than an equivalent RF microtrap, with faster shuttling and lower heating rates. Any of these possible advantages could become essential in building a real quantum computer.

*Moving ions along the magnetic field of a Penning trap is a fairly standard technique of many groups using cylindrical Penning traps [81].

Chapter 5

Quantum Jumps and J -Mixing

“It was on March 12, 1862, in the laboratory of the Royal Institution that Faraday carried out [his last recorded] experiment. The notes in his notebook, although not quite clear, leave no doubt that he was attempting to demonstrate by means of a spectroscope that magnetism has a direct effect on a light source. The result was however absolutely negative, and Faraday writes in his notebook ‘not the slightest effect demonstrable either with polarised or unpolarised light.’ ” – P. Zeeman. Nobel Prize Lecture (1903) [84].




Figure 5.1: Fluorescence from a single Ca^+ ion at 0.9 tesla.

A typical fluorescence signal from a single Ca^+ ion in the PCB trap is

shown in figure 5.1. Two distinct levels can clearly be seen. The lower level is the background from laser scatter, PMT dark current etc., and the upper level is the fluorescence. When the ion is in the metastable $D_{5/2}$ state, it does not fluoresce. A trace such as the one shown in figure 5.1 is clear evidence that a single fluorescing ion is present in the trap. The trace shown in figure 5.1 has a number of other interesting features:

1. The signal level is low compared to when individual ions are observed in the same system operated as an RF trap (where ~ 7000 counts s^{-1} are seen).
2. The average length of time for which the ion is dark is significantly lower than the accepted lifetime of the $D_{5/2}$ state (1.168 ± 0.007 s [85]).
3. The ion frequently goes into the $D_{5/2}$ state even though $P_{1/2} \rightarrow D_{5/2}$ decay is very strongly forbidden and there should be no light other than at 397 nm and 866 nm.

One possible explanation for the first point could be coherent population trapping. The many different laser wavelengths interacting with the ion can put it into a superposition of energy eigenstates which does not fluoresce, known as a dark state. Figure 5.2 shows an example of this for a small cloud of Ca^+ ions in a Penning trap. One of the four 866 nm lasers is slowly scanned over a few hundred MHz as all the other lasers frequencies are held constant. A dip in the signal can be seen. The effect has been theoretically studied in our system, and is described in more detail in [52]. The overall conclusion of this study was that although a large number of dark states do exist, perfect optimisation of laser tunings should in principle allow operation without reduced fluorescence. This requires the simultaneous optimisation of six laser wavelengths, but the wavemeter is not accurate enough for the laser frequencies to be set at the calculated optimal values. Instead, the laser frequencies are approximately set and then optimised by maximising the ion signal itself. Given these experimental uncertainties, coherent population trapping still seems to be a likely mechanism leading to the lower signal seen in the Penning trap. It has also been reported that in the rare isotope of $^{43}Ca^+$, which has a hyperfine splitting comparable to our Zeeman splitting, the fluorescence rate is similarly reduced in comparison to the simpler $^{40}Ca^+$ system [55].

The simplest explanation of the second and third points is that light at 854 nm and 850 nm or 393 nm is interacting with the ions. Referring to




Figure 5.2: Fluorescence trace showing a dark state in a small cloud of Ca^+ ions in a Penning trap as one 866 nm laser is scanned in a period of 10 seconds. The horizontal scale is approximate.

figure 2.2 in chapter 2, light at 850 nm and 393 nm will excite to the $\text{P}_{3/2}$ state, which can then decay to $\text{D}_{5/2}$, causing the ion to turn dark. Light at 854 nm will repump the $\text{D}_{5/2}$ state, shortening the duration of dark times. This effect was mentioned as significant in the precise $\text{D}_{5/2}$ lifetime measurement experiment of Barton *et al.* [85]. Quantum jumps due to unwanted light at 393 nm from diode lasers have also been reported elsewhere [86].

It was found that there was indeed light present at 850 nm and 854 nm. However, another more subtle effect was also afoot, shelving ions into the $\text{D}_{5/2}$ state even in the absence of 850 nm and 393 nm light. This effect is caused by the strong magnetic field of the Penning trap. As well as producing a significant Zeeman splitting, the magnetic field mixes the different J levels of the ion, adding a small admixture of $\text{D}_{3/2}$ to the $\text{D}_{5/2}$ levels, and a small admixture of $\text{P}_{3/2}$ to the $\text{P}_{1/2}$ levels. Thus a decay from $\text{P}_{1/2}$ to $\text{D}_{5/2}$ will (with a very small branching ratio) become an allowed transition. This effect is very relevant because reliable internal state readout is required for quantum computation. The amount of mixing, as explained in §5.2, is (to first order) proportional to the square of the magnetic field. The use of a very similar effect in clouds of neutral atoms has been proposed as a tool for atomic clocks: A magnetic field can be used to tune the linewidth of a very narrow linewidth atomic transition [87, 88]. However, the effect of magnetic J -mixing has never been observed before in individual atomic particles.

5.1 Notes on the Experimental Setup

5.1.1 Amplified Spontaneous Emission in Diode Lasers


Figure 2.4 (in chapter 2) shows a typical gain curve of one of the IR laser diodes used for repumping the $D_{3/2} \rightarrow P_{1/2}$ transition at 866 nm. The diodes were used in the Littrow configuration with external gratings, as described in §2.3.1. Thus only the selected wavelength oscillates inside the diode and lases. However, any light between about 800 and 890 nm can undergo gain if it passes through the diode. Spontaneous emission inside the diode at 850 nm and 854 nm can be amplified even if it just passes through the diode once*, leading to amplified spontaneous emission (ASE). This effect is present to some degree in all lasers. The actual amount of ASE is dependent on the details of the gain medium and the laser cavity. The laser output is therefore made up of a narrow band of true laser emission sitting on a broad spectrum of ASE which covers the entire gain bandwidth of the lasing medium.

To remove unwanted 850 and 854 nm light, a filter was placed in the path of the IR beam. A Thorlabs FB870-10 bandpass filter was used. The transmission at 854 nm is significantly below 1%, and the transmission at 850 nm is even smaller. Transmission at 866 nm is $\sim 50\%$.


Light at 393 nm is also undesirable. The gain curve of UV diode lasers is typically much narrower than that of IR diode lasers. The gain of the diodes used at 397 nm does not extend down to 393 nm, so there should be no ASE at this wavelength in these diodes. Even so, to be sure of removing any 393 nm light, a filter was also placed in the path of the blue laser beams. This filter, a Semrock FF01-406/15-25 bandpass filter, has less than 1% transmission at 393.4 nm, while retaining 90% transmission at 396.8 nm [89].

Fluorescence from an ion in a Penning trap at 0.9 tesla with both of the laser beams filtered is shown in figure 5.3(a). The average duration of dark periods is now significantly longer than in figure 5.1. With the lasers filtered, a histogram of the dark time durations gives a good fit to the expected lifetime (figure 5.4). However, the ion still goes into the shelved state so often that it now spends more time dark than bright.

*Also, even though the output surface of the diode is anti-reflection coated, some fraction of the light could internally reflect and pass through the diode more than once.



(a) 0.9 tesla



(b) 0.2 tesla

Figure 5.3: Fluorescence from a single Ca^+ ion with filtered laser beams. (a): Penning trap with a magnetic field of 0.9 tesla. (b): Combined trap with a field of 0.2 tesla.




Figure 5.4: Histogram of ~ 500 dark (shelved) time durations. The accepted $D_{5/2}$ lifetime is 1.168 s.

5.1.2 Combined Trap

When the lasers were filtered and an ion was trapped in an RF trap at low magnetic field*, quantum jumps were almost never observed. This implies that the magnetic field has a significant effect on the rate of shelving to the $D_{5/2}$ level. In order to investigate this effect in detail it was necessary to study ions trapped over a wide range of magnetic fields.

A Penning trap requires a large magnetic field to work reliably. It was found that clouds of ions could be trapped at a field of no lower than 0.6 tesla. As the field was reduced, trapping became more difficult and less reliable. On the other hand, 0.6 tesla is a very large magnetic field compared to that in an RF trap, which is typically on the order of a few gauss to a few hundred gauss. To investigate the dependence of the shelving rate on magnetic field strength, the trap was operated as a combined trap with both a trapping RF potential and a magnetic field. An example of fluorescence from a single ion in a combined trap with a magnetic field of $B = 0.2$ tesla is shown in figure 5.3(b). It can be seen that the strength of the magnetic field makes a remarkable difference to the shelving rate.

The PCB trap described in chapter 4 was used for all the measurements in this chapter. An RF voltage of 200 V peak to peak, at 2.27 MHz, was applied to the endcap electrodes. This was produced using a function generator, an

* $B \ll 0.01$ tesla.

RF amplifier and a simple step-up transformer, resonant at 2.27 MHz. The ring electrodes were set to a bias of +2.4 V. A 200 μm pinhole was used in front of the PMT.

The scaling of B with current through the magnet coils was calibrated using a Hall probe outside the vacuum chamber. As the Hall probe was not in the centre of the trap, the field measured does not correspond perfectly to the real field at the position of an ion. To compensate for this, an accurate measurement of B was made by measuring the Penning trap motional frequencies of a small ion cloud at a magnet coil current of 30 amp. The magnetron and modified cyclotron frequencies were found to be 50.7 kHz and 291 kHz respectively, using the technique described in [51]. This implies a true cyclotron frequency of 342 kHz and hence a field of 0.898 tesla. The Hall probe was placed close to the trap, in such a position as to measure 0.898 ± 0.005 tesla at 30 amp. The strength of B was then measured with the probe at various currents. The results of this calibration were shown earlier in figure 2.15. By repeatedly ramping the current up and down, and taking measurements with the probe each time, an approximate uncertainty on B was estimated to be ~ 0.005 tesla.

5.2 Theory

“I think I can safely say that nobody understands quantum mechanics.” – Richard Feynman

5.2.1 Transition Rates

Experimentally we see that the ion goes dark ($D_{5/2}$ shelving) after emitting a (large) number of 397 nm photons. The average number of blue photons scattered before shelving occurs is very high at low magnetic fields. It then decreases as B is increased. This number is equivalent to the ratio of the transition rates, $\Gamma(i, f)$.

$$n = \frac{\Gamma(P_{\frac{1}{2}}, S_{\frac{1}{2}})}{\Gamma(P_{\frac{1}{2}}, D_{\frac{5}{2}})} \rightarrow \infty \text{ at low } B. \quad (5.1)$$

The inverse of this will be considered from now on, since this scales posi-

tively with B .

$$n^{-1} = \frac{\Gamma(P_{\frac{1}{2}}, D_{\frac{5}{2}})}{\Gamma(P_{\frac{1}{2}}, S_{\frac{1}{2}})} \rightarrow 0 \text{ at low } B. \quad (5.2)$$

From Fermi's Golden Rule, a decay rate can be expressed in terms of a matrix element and a density of final states. The spontaneous emission rate from $|e, J'\rangle$ to $|g, J\rangle$ (with an energy difference of $\hbar\omega_g$) is

$$\Gamma = \frac{4\omega_g^3}{3\hbar c^3} \frac{|\langle g, J | \mathbf{d} | e, J' \rangle|^2}{2J' + 1} \quad (5.3)$$

where $\langle g, J | \mathbf{d} | e, J' \rangle$ is the reduced matrix element*, and $\mathbf{d}$ is the electric dipole operator [90].

Combining (5.2) and (5.3) leads to

$$n^{-1} = \frac{\omega_{D_{\frac{5}{2}}}^3 \left| \langle D_{\frac{5}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle \right|^2}{\omega_{S_{\frac{1}{2}}}^3 \left| \langle S_{\frac{1}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle \right|^2}. \quad (5.4)$$

Now we divide the top and bottom by the transition rate of the 866 nm allowed transition, $\Gamma(P_{\frac{1}{2}}, D_{\frac{3}{2}})$. The branching ratio between the 397 nm and 866 nm transitions, R , is known† and can be inserted. Also, since the fine structure splitting is so small compared to the photon energies, $(\omega_{D_{\frac{5}{2}}}/\omega_{D_{\frac{3}{2}}})^3 = 0.984 \approx 1$.

$$n^{-1} = \frac{\omega_{D_{\frac{5}{2}}}^3 \left| \langle D_{\frac{5}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle \right|^2}{\omega_{S_{\frac{1}{2}}}^3 \left| \langle S_{\frac{1}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle \right|^2} \times \frac{\omega_{D_{\frac{5}{2}}}^3 \left| \langle D_{\frac{5}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle \right|^2}{\omega_{D_{\frac{3}{2}}}^3 \left| \langle D_{\frac{3}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle \right|^2} \quad (5.5)$$

$$= \frac{\Gamma(P_{\frac{1}{2}}, D_{\frac{3}{2}})}{\Gamma(P_{\frac{1}{2}}, S_{\frac{1}{2}})} \times 0.984 \left| \frac{\langle D_{\frac{5}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle}{\langle D_{\frac{3}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle} \right|^2 \quad (5.6)$$

$$\approx R \times \left| \frac{\langle D_{\frac{5}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle}{\langle D_{\frac{3}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle} \right|^2. \quad (5.7)$$

*While a regular matrix element between an initial state and a final state, each containing several sub-states, would be computed by averaging over the initial sub-states and summing over the final sub-states, the reduced matrix element is computed by summing over both the initial and final sub-states.

† $R \equiv \frac{\Gamma(P_{\frac{1}{2}}, D_{\frac{3}{2}})}{\Gamma(P_{\frac{1}{2}}, S_{\frac{1}{2}})} = \frac{1}{15.4}$ as calculated in [91].

Using the identity*

$$|\langle \xi, j \| \mathbf{d} \| \xi', j' \rangle|^2 \equiv \sum_{m, m'} |\langle \xi, j, m | \mathbf{d} | \xi', j', m' \rangle|^2 \quad (5.8)$$

where the label ξ accounts for all the quantum numbers apart from j and m , the numerator of (5.7) can be expanded into all the possible $P_{1/2}$ to $D_{5/2}$ transition probabilities.

$$n^{-1} = 2R \sum_{m_J = -\frac{1}{2}}^{+\frac{3}{2}} \left| \frac{\langle D_{\frac{5}{2}}^{m_J} | \mathbf{d} | P_{\frac{1}{2}}^{\frac{1}{2}} \rangle}{\langle D_{\frac{3}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle} \right|^2. \quad (5.9)$$

The factor of 2 comes from the sum over m of the $P_{1/2}$ states.

5.2.2 Magnetic Perturbation

A magnetic field will produce a small perturbation to the standard Hamiltonian of the system. The eigenstates are usually defined by n (principal), L (orbital angular momentum), J (total angular momentum) and m_J (Zeeman). However, since a magnetic field interacts differently with spin and orbital angular momentum, J is no longer a good quantum number to use when considering the perturbation caused by the magnetic field. To label the eigenstates of the magnetic interaction operator, the quantum numbers n , l , m_l and m_s are used instead.

From first order perturbation theory, for a perturbed Hamiltonian $\mathbf{H} = \mathbf{H}_0 + \mathbf{H}'$, the perturbed eigenfunctions are

$$|\psi'_a\rangle = |\psi_a\rangle + \sum_{b \neq a} \frac{\langle \psi_b | \mathbf{H}' | \psi_a \rangle}{E(a) - E(b)} |\psi_b\rangle + \dots \quad (5.10)$$

where $E(a)$ is the energy of state $|\psi_a\rangle$ etc. This can be applied to (5.9).

$$n^{-1} = 2R \sum_{m_J = -\frac{1}{2}}^{+\frac{3}{2}} \left| \frac{\left(\langle D_{\frac{5}{2}}^{m_J} | + \delta_{m_J} \langle D_{\frac{3}{2}}^{m_J} | \right) \mathbf{d} \left(| P_{\frac{1}{2}}^{\frac{1}{2}} \rangle + \delta_P | P_{\frac{3}{2}}^{\frac{1}{2}} \rangle \right)}{\langle D_{\frac{3}{2}} | \mathbf{d} | P_{\frac{1}{2}} \rangle} \right|^2 \quad (5.11)$$

*This can be proved by considering the Wigner-Eckart Theorem (5.24) along with the general properties of Clebsch-Gordan coefficients.

where

$$\delta_{m_J} = \frac{\left\langle D_{\frac{3}{2}}^{m_J} \left| -\mu \cdot \mathbf{B} \right| D_{\frac{5}{2}}^{m_J} \right\rangle}{E(D_{\frac{5}{2}}) - E(D_{\frac{3}{2}})}, \quad \delta_P = \frac{\left\langle P_{\frac{3}{2}}^{\frac{1}{2}} \left| -\mu \cdot \mathbf{B} \right| P_{\frac{1}{2}}^{\frac{1}{2}} \right\rangle}{E(P_{\frac{1}{2}}) - E(P_{\frac{3}{2}})}. \quad (5.12)$$

Multiplying out equation (5.11) gives four terms. One term contains the transition matrix $\left\langle D_{\frac{5}{2}}^{m_J} \left| \mathbf{d} \right| P_{\frac{1}{2}}^{\frac{1}{2}} \right\rangle$, which is a forbidden ($\Delta J = 2$) transition and hence zero. Another term contains a second order mixing, $\delta_{m_J} \delta_P$, which gives a negligible contribution since $\delta_{m_J}, \delta_P \ll 1$. This leaves the two allowed transitions with first order coefficients.

$$n^{-1} = 2R \sum_{m_J=-\frac{1}{2}}^{+\frac{3}{2}} \left| \frac{\delta_{m_J} \left\langle D_{\frac{3}{2}}^{m_J} \left| \mathbf{d} \right| P_{\frac{1}{2}}^{\frac{1}{2}} \right\rangle + \delta_P \left\langle D_{\frac{5}{2}}^{m_J} \left| \mathbf{d} \right| P_{\frac{3}{2}}^{\frac{1}{2}} \right\rangle}{\left\langle D_{\frac{3}{2}} \left\| \mathbf{d} \right\| P_{\frac{1}{2}} \right\rangle} \right|^2. \quad (5.13)$$

The Ca^+ electric dipole transition amplitudes have been calculated using relativistic many-body theory by C. Guet and W. R. Johnson [91]. They quote

$$\left\langle D_{\frac{3}{2}} \left\| \mathbf{d} \right\| P_{\frac{1}{2}} \right\rangle = 2.373 \text{ and } \left\langle D_{\frac{5}{2}} \left\| \mathbf{d} \right\| P_{\frac{3}{2}} \right\rangle = 3.186 \quad (5.14)$$

in arbitrary units*. Therefore,

$$n^{-1} = 2R \sum_{m_J=-\frac{1}{2}}^{+\frac{3}{2}} \left| \frac{\delta_{m_J} \left\langle D_{\frac{3}{2}}^{m_J} \left| \mathbf{d} \right| P_{\frac{1}{2}}^{\frac{1}{2}} \right\rangle}{\left\langle D_{\frac{3}{2}} \left\| \mathbf{d} \right\| P_{\frac{1}{2}} \right\rangle} + \frac{3.186}{2.373} \frac{\delta_P \left\langle D_{\frac{5}{2}}^{m_J} \left| \mathbf{d} \right| P_{\frac{3}{2}}^{\frac{1}{2}} \right\rangle}{\left\langle D_{\frac{5}{2}} \left\| \mathbf{d} \right\| P_{\frac{3}{2}} \right\rangle} \right|^2. \quad (5.15)$$

To compute the mixing coefficients, δ_{m_J} , δ_P , the $|J, m_J\rangle$ states are written in the $|l, m_l; s, m_s\rangle$ basis. This can be done straightforwardly using the Clebsch-Gordan coefficients (C_i).

$$|J, m_J\rangle = \sum_i C_i |l, m_l^i; s, m_s^i\rangle \quad (5.16)$$

*No explicit errors are given, but the largest disparity between the calculated and measured lifetimes in Ca^+ is 1.8% [91].

P states:

$$\left| \frac{1}{2}, \pm \frac{1}{2} \right\rangle = \pm \sqrt{\frac{2}{3}} \left| 1, \pm 1; \frac{1}{2}, \mp \frac{1}{2} \right\rangle \mp \sqrt{\frac{1}{3}} \left| 1, 0; \frac{1}{2}, \pm \frac{1}{2} \right\rangle \quad (5.17a)$$

$$\left| \frac{3}{2}, \pm \frac{1}{2} \right\rangle = \sqrt{\frac{1}{3}} \left| 1, \pm 1; \frac{1}{2}, \mp \frac{1}{2} \right\rangle + \sqrt{\frac{2}{3}} \left| 1, 0; \frac{1}{2}, \pm \frac{1}{2} \right\rangle. \quad (5.17b)$$

D states:

$$\left| \frac{3}{2}, \pm \frac{1}{2} \right\rangle = \pm \sqrt{\frac{3}{5}} \left| 2, \pm 1; \frac{1}{2}, \mp \frac{1}{2} \right\rangle \mp \sqrt{\frac{2}{5}} \left| 2, 0; \frac{1}{2}, \pm \frac{1}{2} \right\rangle \quad (5.18a)$$

$$\left| \frac{5}{2}, \pm \frac{1}{2} \right\rangle = \sqrt{\frac{2}{5}} \left| 2, \pm 1; \frac{1}{2}, \mp \frac{1}{2} \right\rangle + \sqrt{\frac{3}{5}} \left| 2, 0; \frac{1}{2}, \pm \frac{1}{2} \right\rangle \quad (5.18b)$$

$$\left| \frac{3}{2}, \pm \frac{3}{2} \right\rangle = \pm \sqrt{\frac{4}{5}} \left| 2, \pm 2; \frac{1}{2}, \mp \frac{1}{2} \right\rangle \mp \sqrt{\frac{1}{5}} \left| 2, \pm 1; \frac{1}{2}, \pm \frac{1}{2} \right\rangle \quad (5.18c)$$

$$\left| \frac{5}{2}, \pm \frac{3}{2} \right\rangle = \sqrt{\frac{1}{5}} \left| 2, \pm 2; \frac{1}{2}, \mp \frac{1}{2} \right\rangle + \sqrt{\frac{4}{5}} \left| 2, \pm 1; \frac{1}{2}, \pm \frac{1}{2} \right\rangle. \quad (5.18d)$$

Writing (5.12) in the new basis, we have for $m_J = \frac{3}{2}$,

$$\delta_{\frac{3}{2}} = \frac{\langle \frac{3}{2}, \frac{3}{2} | -\mu \cdot \mathbf{B} | \frac{5}{2}, \frac{3}{2} \rangle}{E(D_{\frac{5}{2}}) - E(D_{\frac{3}{2}})} \quad (5.19)$$

$$= \frac{\left(\sqrt{\frac{4}{5}} \langle 2; -\frac{1}{2} | -\sqrt{\frac{1}{5}} \langle 1; \frac{1}{2} | \right) (-\mu_B B (L_z + 2S_z)) \left(\sqrt{\frac{1}{5}} | 2; -\frac{1}{2} \rangle + \sqrt{\frac{4}{5}} | 1; \frac{1}{2} \rangle \right)}{E(D_{\frac{5}{2}}) - E(D_{\frac{3}{2}})} \quad (5.20)$$

$$= \frac{-\mu_B B \left(\sqrt{\frac{4}{5}} \sqrt{\frac{1}{5}} \langle 2; -\frac{1}{2} | L_z + 2S_z | 2; -\frac{1}{2} \rangle - \sqrt{\frac{1}{5}} \sqrt{\frac{4}{5}} \langle 1; \frac{1}{2} | L_z + 2S_z | 1; \frac{1}{2} \rangle \right)}{E(D_{\frac{5}{2}}) - E(D_{\frac{3}{2}})} \quad (5.21)$$

$$= \frac{2 \mu_B B}{5 \Delta E_D} \quad (5.22)$$

where the l and s labels have been dropped, ΔE_D is been defined as $\Delta E_D \equiv E(D_{\frac{5}{2}}) - E(D_{\frac{3}{2}})$, and the magnetic Hamiltonian $-\mu \cdot \mathbf{B} = -\mu_B B (L_z + 2S_z)$ has been used.

The other mixing coefficients can be found similarly.

$$\delta_{\frac{3}{2}} = \frac{2}{5} \frac{\mu_B B}{\Delta E_D} \quad (5.23a)$$

$$\delta_{\frac{1}{2}} = \frac{\sqrt{6}}{5} \frac{\mu_B B}{\Delta E_D} \quad (5.23b)$$

$$\delta_{-\frac{1}{2}} = \frac{\sqrt{6}}{5} \frac{\mu_B B}{\Delta E_D} \quad (5.23c)$$

$$\delta_P = \frac{\sqrt{2}}{3} \frac{\mu_B B}{\Delta E_P}. \quad (5.23d)$$

Note that $\Delta E_P \equiv E(P_{\frac{1}{2}}) - E(P_{\frac{3}{2}}) < 0$, while $\Delta E_D > 0$. So δ_{m_J} is positive but δ_P is negative.

5.2.3 Wigner-Eckart Theorem

Using the Wigner-Eckart theorem, the matrix elements for the various m_J can be written in terms of the reduced matrix element and the appropriate Clebsch-Gordan coefficients.

$$\langle \xi, j, m | T_q^\kappa | \xi', j', m' \rangle = \frac{\langle \xi, j || T^\kappa || \xi', j' \rangle}{\sqrt{2j+1}} \langle j', m'; \kappa, q | j, m \rangle. \quad (5.24)$$

The irreducible tensor operator T_q^κ can be used to describe the electric dipole decay operator (written as $\mathbf{d}$ in the equations above): T_0^1 for a π polarised photon, $T_{\pm 1}^1$ for $\sigma_{\mp}$ polarisation.

Applying this to equation (5.15) gives

$$n^{-1} = 2R \sum_{m_J = -\frac{1}{2}}^{+\frac{3}{2}} \left| \frac{\delta_{m_J} \langle \frac{1}{2}, \frac{1}{2}; 1, m_J - \frac{1}{2} | \frac{3}{2}, m_J \rangle}{\sqrt{4}} + \frac{3.186}{2.373} \frac{\delta_P \langle \frac{3}{2}, \frac{1}{2}; 1, m_J - \frac{1}{2} | \frac{5}{2}, m_J \rangle}{\sqrt{6}} \right|^2. \quad (5.25)$$

Finally, putting in the Clebsch-Gordan coefficients and the mixing coeffi-

cients from (5.23) yields

$$n^{-1} = 2R \left(\left| \frac{1}{2} \frac{\sqrt{6}}{5} \frac{\mu_B B}{\Delta E_D} \sqrt{\frac{1}{3}} + \frac{3.186}{2.373} \frac{1}{\sqrt{6}} \frac{\sqrt{2}}{3} \frac{\mu_B B}{\Delta E_P} \sqrt{\frac{3}{10}} \right|^2 + \left| \frac{1}{2} \frac{\sqrt{6}}{5} \frac{\mu_B B}{\Delta E_D} \sqrt{\frac{2}{3}} + \frac{3.186}{2.373} \frac{1}{\sqrt{6}} \frac{\sqrt{2}}{3} \frac{\mu_B B}{\Delta E_P} \sqrt{\frac{3}{5}} \right|^2 \right) \quad (5.26)$$

$$+ \left| \frac{1}{2} \frac{2}{5} \frac{\mu_B B}{\Delta E_D} 1 + \frac{3.186}{2.373} \frac{1}{\sqrt{6}} \frac{\sqrt{2}}{3} \frac{\mu_B B}{\Delta E_P} \sqrt{\frac{3}{5}} \right|^2 \Bigg) \\ = 2 \frac{1}{15.4} \left(|0.00110 - 0.00030|^2 + |0.00156 - 0.00042|^2 + |0.00156 - 0.00042|^2 \right) B^2 \text{ tesla}^{-2} \quad (5.27)$$

$$= 4.2 \times 10^{-7} B^2 \text{ tesla}^{-2}. \quad (5.28)$$

Dividing this by the detection efficiency gives a prediction for the number of photons detected between each shelving event.

$$n_{\text{detected}}^{-1} = \frac{n^{-1}}{\eta} \quad (5.29)$$

$$\approx 0.001 B^2 \text{ tesla}^{-2} \quad (5.30)$$

for $\eta = 4 \times 10^{-4}$.

5.3 Data Analysis

For each magnetic field, ~ 1000 seconds of data were recorded. The data were saved as files containing the time and the number of photons detected per time bin. The time bins were 10 ms wide. The files were processed using the Perl script listed in appendix C. The outline of the analysis algorithm is as follows:

For each I_{mag} :

1. Find the best value to use as the threshold signal level between a bright ion and a dark ion.
2. Create three sets of data: The duration of each bright period*; the number of photons detected (minus the background level) during each

*For figure 5.6.

bright period^{*}; and the duration of each dark period[†].

3. Compute the mean of each of the three data sets. This is a simpler method than fitting a decaying exponential to a histogram (i.e. figure 5.4), and gives similar results.
4. Estimate the standard deviations of the various durations.
5. Look up the value of B as a function of I_{mag} (figure 2.15) and print out the data points as shown in figures 5.6, 5.7 and 5.8.

The script scans through the data and works out the mean signal level of each bin with a level below a certain (guessed) threshold[‡], S_{dark} . Similarly, the average level of the signal above the threshold guess, S_{bright} , is computed. Now an improved estimate of the best threshold level is set as

$$S_{\text{threshold}} = \frac{2S_{\text{dark}} + S_{\text{bright}}}{3}. \quad (5.31)$$

This formula follows the advice given in [85]. The best threshold is not halfway between S_{bright} and S_{dark} because the higher level is more noisy. Assuming the count rate from the ion fluorescence forms a Poisson distribution, the standard deviation of the signal when the ion is bright is $\sqrt{S_{\text{bright}}}$, while the standard deviation of the background signal is a much smaller. This is clearly seen in figure 5.5.

When an ion is in the metastable state and not fluorescing, the PMT signal mostly lies below the threshold level. Occasionally however, noise in the dark current, laser scatter, cosmic rays etc., causes the signal to go briefly above $S_{\text{threshold}}$. Similarly, the Poisson distributed count rate when the ion is fluorescing has a tail which extends below the threshold level. As mentioned above, for higher magnetic fields the bright signal level tends to be lower. This causes the tail of the distribution to drop significantly below the threshold level, as shown in figure 5.5(b). Clearly for large fields, where the dark and bright peaks are not perfectly resolvable, a bright ion will sometimes be wrongly detected as a dark ion.

The analysis program attempts to reject these spurious events by looking ahead to see if several bins are also above/below threshold, i.e.,

^{*}For figure 5.7.

[†]For figure 5.8.

[‡]The initial guess of the threshold level was 5 counts per 10 ms.





Figure 5.5: Distribution of fluorescence levels from a single Ca^+ ion. At high mean values, a Poisson distribution converges to a Gaussian distribution. The data has been fit to $f(S) = f(S)_{\text{dark}} + f(S)_{\text{bright}}$, where $f(S)_{\text{dark}}$ and $f(S)_{\text{bright}}$ are separate Gaussian functions.

```

IF (ion_was_bright AND S[i] < threshold
    AND S[i+1] < threshold AND S[i+2] < threshold
    AND S[i+3] < threshold AND S[i+4] < threshold)
THEN
    Ion_has_just_turned_dark;

```

where $S[i]$ is the signal level in the bin of index i . The dark→bright events are detected similarly.

This method rejects noise, but it also rejects any genuine durations which are shorter than δt (5 time bins in this case). This will have a completely negligible effect only if $\delta t \ll \tau$, the lifetime. We wish to find τ , but have only the average measured lifetime, $\langle t \rangle$. If $\delta t \rightarrow 0$ and the measurement is perfect then $\tau = \langle t \rangle$, but since this is not the case then there is a small shift. The effect is to make $\langle t \rangle = \tau + \delta t$. To see this, consider an ion in a metastable excited state. There is a certain probability that it will decay within a certain time. The probability, $p(t)$, of decaying within a time t is proportional to $\exp(-\frac{t}{\tau})$. Since an excited ion always decays eventually, the sum of the probabilities for all possible durations $0 < t < \infty$ must be 100%.

$$p(t) = \alpha \exp\left(-\frac{t}{\tau}\right) dt \quad (5.32)$$

$$1 = \int_0^\infty \alpha \exp\left(-\frac{t}{\tau}\right) dt \quad (5.33)$$

$$= \alpha \tau \quad (5.34)$$

$$\therefore p(t) = \frac{1}{\tau} \exp\left(-\frac{t}{\tau}\right) dt. \quad (5.35)$$

The average measured time duration is the mean of all the durations not between 0 and ∞ , but between δt and ∞ (normalised such that the total probability is 1):

$$\langle t \rangle = \frac{\int_{\delta t}^\infty p(t) t}{\int_{\delta t}^\infty p(t)} \quad (5.36)$$

$$= \frac{\int_{\delta t}^\infty \frac{t}{\tau} \exp\left(-\frac{t}{\tau}\right) dt}{\int_{\delta t}^\infty \frac{1}{\tau} \exp\left(-\frac{t}{\tau}\right) dt} \quad (5.37)$$

$$= \frac{(\tau + \delta t) \exp\left(-\frac{\delta t}{\tau}\right)}{\exp\left(-\frac{\delta t}{\tau}\right)} \quad (5.38)$$

$$= \tau + \delta t. \quad (5.39)$$

For $\delta t = 50$ ms the shift is not much smaller than the standard error and so must be taken into account.

Any sub-50 ms dark durations will also cause a shift in the measured bright time duration (and vice versa). Two bright periods, separated by a very short dark period, will be detected as one long bright period. This effect is assumed to have only a small effect on the final lifetime measurements, but it becomes much more significant at very high B where there are many short bright periods.

There is also another shift, in the opposite direction, because it is possible (though unlikely) for five bins in a row to be below threshold even when the ion is still fluorescing. The probability, p_1 , of wrongly measuring an ion (in one time bin) to be dark when it is actually still fluorescing is

$$p_1 \approx \frac{\int_0^{S_{\text{threshold}}} f(S)_{\text{bright}} dS}{\int_0^{\infty} f(S)_{\text{bright}} dS} \quad (5.40)$$

where $f(S)_{\text{bright}}$ is the fit to the count rate distribution of the bright level (shown in figure 5.5). Considering several bins in a row must be below threshold, the effect causes $\langle t \rangle$ to be smaller than τ by a factor:

$$\frac{\langle t \rangle}{\tau} \approx (1 - p_1^n)^N \quad (5.41)$$

where n is the number of bins which must be above threshold*, and N is the average number of bins in the period being measured†. For the case with the poorest resolution ($p_1 \approx 0.1$), using $n = 5$ gives $\langle t \rangle \approx 0.998\tau$. This is a very small offset and so can be ignored. Note that these equations are only true assuming a very large number of measurements have been taken.

For each of the three data sets, at each magnetic field, a standard error was estimated using the bootstrap method. For each set, containing n_{bs} time durations, with average duration $\langle t \rangle_0$, a new data set is created by randomly picking n_{bs} values of t from the original set. The mean duration of this new set, $\langle t \rangle_i$, is then computed. This is done again for many ($N_{\text{bs}} = 10^4$) randomly generated sets. The standard deviation is the square root of the variance of

* $n = 5$.

† $N = \tau / 10$ ms.

these mean durations.

$$\sigma = \sqrt{\sum_{i=1}^{N_{\text{bs}}} \frac{(\langle t \rangle_i - \langle t \rangle_0)^2}{N_{\text{bs}}}}. \quad (5.42)$$

The detection efficiency fluctuates / drifts on a day to day timescale. This effect has not been quantitatively measured, but has been noticed to be on the order of 10-20%. The detection efficiency affects the normalised quantum jump rate (figure 5.7), but was not taken into account by the bootstrap algorithm. Thus we expect that many points will lie further than a standard deviation away from the line of best fit.

5.4 Results




Figure 5.6: Shelving rate as a function of magnetic field.

As estimated in §2.4.3, the expected scattering rate of 397 nm photons is roughly $\Gamma/8 = 1.8 \times 10^7 \text{ s}^{-1}$. The expected rate of shelving to $D_{5/2}$ is then this rate times the branching ratio to $D_{5/2}$.

The shelving rate is inversely proportional to the mean duration of the bright periods in the fluorescence trace. This rate, measured at various magnetic fields, is plotted in figure 5.6. The increase in rate with magnetic field is clear, but it does not appear to scale quadratically with B as predicted. The deviation can be explained by the fact that the average fluorescence sig-

nal level is different at different magnetic fields. Each field strength requires re-optimisation of the six laser frequencies, so maintaining the same signal level at each field strength is very difficult. There is also a general trend of decreasing signal as the magnetic field is raised. This means there are less photons scattered per second, less decays from $P_{1/2}$, and hence less decays to $D_{5/2}$. As mentioned above, the reason for this decrease in signal is still not totally clear. It could be due to stronger coherent population trapping in the Zeeman split levels at large B . Other possible causes are:

1. Weaker cooling, since the ion spends more time in the dark state as the field increases.
2. More possibility of trapping and sympathetically cooling ions other than $^{40}\text{Ca}^+$, since the additional B field can increase the range of stable trapping parameters.
3. Difficulty in optimising the laser wavelengths, since improvements cause an ion to rapidly turn dark.
4. Chaotic motion in the combined trap causing delocalisation of the ion, which could lead to stronger RF heating, and perhaps a larger required red-detuning*.

Or several of the above.

Allowance for this effect can be made if the shelving rate is adjusted (normalised) by dividing by the absolute signal level. In other words, the number of photons detected between shelving events, n_{detected} , is considered instead of the time between events. This is shown in figure 5.7.

The result is now independent of the scattering rate, but does depend on the detection efficiency. The solid lines in figure 5.7 show the theoretical result assuming a detection efficiency of $0.04 \pm 0.014\%$. Note that the assumed detection efficiency alters the slope of the straight line (equation (5.29)). Although there is a very large uncertainty in the detection efficiency, the B^2 trend can clearly be seen.

The same data were also analysed to find the lifetime of the $D_{5/2}$ state at each magnetic field. The results, shown in figure 5.8, are in reasonable agreement with the best published results [85, 92]. This indicates that 854 nm

*The best signal seen in the combined trap at high fields was somewhat lower than the best signal seen in a pure Penning trap. The trap was not compensated to reduce micromotion, so RF heating certainly does occur.




Figure 5.7: Normalised shelving rate. The theoretical prediction assuming a detection efficiency of $0.04 \pm 0.014\%$ is shown in red. Both linear and log-log scales are shown for clarity.

light has been mostly eliminated, and that the data analysis script is working correctly. The average measured lifetime lies slightly below the accepted value, most likely because of a small residual amount of 854 nm light. The error on the accepted value is negligible in comparison with errors on our data.




Figure 5.8: Mean lifetime of the $D_{5/2}$ state.

Chapter 6

Discussion

The main motivation of this work was to assess the possibility of using ions in Penning traps as the qubits of a quantum computer. As discussed in chapter 1, there are many groups around the world working towards similar goals with small numbers of laser cooled ions in RF traps. As outlined below, the RF trap obviously has several advantages over the Penning trap, otherwise Penning traps would be more common in quantum computing studies. There are however some possible benefits of using Penning traps, and there may be more, as yet undiscovered.

One disadvantage of Penning traps is simply that generating the $\gtrsim 1$ tesla trapping field is more difficult than generating an RF field that produces a similar trap depth. An RF trapping potential can be created using a simple RF oscillator, a few-watt RF amplifier, and a resonant transformer. On the other hand, generating a large, stable magnetic field typically requires a large, heavy, expensive, power-hungry electromagnet and power supply. Superconducting magnets require less electrical power, but they do consume cryogenic liquids and are also large, expensive and can severely limit optical access to the trap. In the future, permanent magnets may offer a good alternative solution. When temperature stabilised, permanent magnets produce a very stable magnetic field and do not consume power or cryogens. They also provide a method of magnetic field generation which is much more scalable than electromagnetic systems. In recent years, magnets made from Neodymium and other materials have become stronger and cheaper. A Penning trap based on permanent magnets is currently being developed by the group, and could in principle improve on the RF trap technique when it comes to the ease and

simplicity of generating trapping fields.

The unstable magnetron motion makes cooling and localising ions in a Penning trap more difficult than in an RF trap. However, as shown in chapter 3, this difficulty can be overcome. We are able to cool single ions and strongly align pairs of ions along the axis of the trap.

Another obvious disadvantage of laser cooling Ca^+ in a Penning trap is the large Zeeman splitting, creating the requirement for many more laser frequencies. We have shown however, that even with the more complicated system of lasers, single Ca^+ ions can be trapped and cooled for extended periods of time in a Penning trap. The large Zeeman splitting may in fact be turned into an advantage if the sublevels of the $S_{1/2}$ state are used as qubit states. The current aim of the group is to use the metastable $D_{5/2}$ state, along with the ground state as a qubit, but there would be several advantages, not available at low magnetic fields, of using the Zeeman sublevels of the ground state instead*.

The J -mixing effect presented in chapter 5, and also the fluorescence rate reduction discussed in the same chapter, will reduce the qubit readout fidelity. The rate of shelving to the $D_{5/2}$ level was seen to be up to a few events per second at high magnetic field. This is of a similar timescale to the lifetime of the $D_{5/2}$ state, so the J -mixing effect should not limit fidelities much more than they are limited by spontaneous decay. Similarly, while the main cause of the signal level reduction is not fully clear, the level is only reduced by a factor of ~ 2 compared to single ions at low field. Thus the readout fidelity will be somewhat reduced, but not by orders of magnitude. The J -mixing problem could be greatly reduced if a different ion (with larger fine structure splitting or smaller Landé g factors) were to be used. This may be necessary if significantly larger magnetic fields are used in future experiments. If the fluorescence detection efficiency was well calibrated, then the J -mixing effect could perhaps be useful in measuring, for instance, the branching ratios of the various P state decay channels.

An obvious advantage of the Penning trap over the RF trap is the lack of micromotion and RF heating. Although ions in a well compensated RF

*The lifetime of the upper $S_{1/2}$ sublevel is much longer than the $D_{5/2}$ lifetime, so there would be no decoherence caused by spontaneous decay. Also, generating a pair of Raman laser beams or a microwave source with a narrow linewidth is often much simpler than building an ultra-narrow linewidth laser for addressing the quadrupole transition. The main difficulty lies in producing Raman beams with such a large (28 GHz at 1 tesla) frequency difference.

trap show little micromotion when they are very close to the trap centre/axis, a small displacement of the ion causes a significant increase in micromotion. The effect of RF heating on qubit decoherence can only be fully studied by performing coherence measurements on ions in a trap with no RF heating.

The recent trend in RF trap development is the manufacturing of very small traps. One reason for this is that steeper trapping potentials (and hence higher motional frequencies and smaller Lamb-Dicke parameters) can be produced more easily*. In a Penning trap however, the motional frequencies are limited by the strength of the magnetic field. This could be a disadvantage or an advantage. The strongest magnets available cannot produce fields beyond ~ 20 tesla, so the Penning Lamb-Dicke parameter is fundamentally limited by magnet technology. On the other hand, if motional frequencies of around 1 MHz are enough, a few-tesla superconducting magnet could produce these frequencies in a very large trap.

In small RF traps, there is a large (and not fully understood) heating rate due to electrically fluctuating patches on the surface of the trap electrodes. The heating rate of an ion at a distance R from the trap electrodes is found to scale approximately as R^{-4} [93]. Patch potential heating is thus very large in small microtraps. Since Penning traps could be made much larger, this heating effect (along with any other causes of decoherence caused by ions being in close proximity to electrodes), could be greatly eliminated. Also, since heating rates have not been measured in very small Penning traps, it may be that the anomalous heating effect is different for electrodes which only support DC voltages.

The problems of laser scatter can be much more difficult to eliminate in very small traps. Using a large Penning trap instead of a small RF trap could reduce this technical difficulty.

Another problem encountered with RF microtraps is excessive heating of the trap structure itself. Even if most of the applied RF power is extracted from the trap, some of the power will always be absorbed. This effect physically damaged many of the first microtraps, and has even caused UHV problems in some experiments. Making use of a Penning trap eliminates this problem entirely.

As discussed at the end of chapter 4, the scalable Penning trap has some possible advantages over RF trap arrays when considering the shuttling of ions

*A voltage applied to the electrodes of a small trap will produce larger electric field gradients compared to when the same voltage is applied to the electrodes of a large trap.

between sub-traps. The time required to shuttle an ion in our trap is to a first approximation that of a single cyclotron period ($t \approx \frac{2\pi m}{eB}$). This duration is favourably short when compared to similar shuttling experiments in RF traps. When an ion is shuttled between a pair of Penning traps, it automatically comes to rest in the centre of the second trap*, unlike in RF trap arrays. Also unlike shuttling ions in RF traps, the Penning trap shuttling scheme can be extended relatively straightforwardly to the important problem of moving ions round corners. The scheme cannot be simply applied to the splitting and joining of pairs of ions, but ion pairs could be split and joined along the axis of a trap in a similar way to ions in linear RF traps. Finally, as multiple trap structures become more and more complicated, the relative simplicity of the voltage switching system could be the most important advantage of Penning trap arrays.

This work has shown that Ca^+ ions in Penning traps still seem to be a viable approach for quantum computing. There are many difficulties associated with using Penning traps instead of RF traps, but these difficulties can be overcome. It is hoped that, with the improvements described in §4.5, single ions can be reliably shuttled between different trapping zones. It is also hoped that the decoherence rates of ionic qubits in Penning traps will be studied in the near future. Soon after completing these two goals, they will be combined and the effects of fast shuttling on qubit decoherence will be measured.

*Assuming good control voltages, timing, and alignment.

Appendix A

Miscellaneous Electronic Schematics




Figure A.1: HeNe stabilisation circuit schematic.




Figure A.2: Hot cathode filament and oven controller schematic.

Chapter A




Figure A.3: Voltage source generator schematic.

Appendix B

Script to Generate SIMION Geometry Files

```
#!/usr/bin/perl -w

# Script to create SIMION geometry file for the PCB trap

$R = 2.78; # Radius of pad (centre to point)
$D = 0.2; # Track width
$Z = $ARGV[1]; # Distance between pads in z
$HOLE2 = 2.0 / 2; # Radius of hole
$HOLE1 = 1.0 / 2;
$HOLE0 = 0.5 / 2;
@SIZE = (40.0, 20.0, 10.0); # Size of universe
$GRID_UNITS_PER_MM = 10;
$thickness = $delta = 1 / $GRID_UNITS_PER_MM;
$thick = 2*$thickness;

$PI = atan2(1, 1) * 4;

@e = (0, 0, 0, 1, 0, 0);

#####

sub Write_Header {
    @grid_size = ($SIZE[0] * $GRID_UNITS_PER_MM, $SIZE[1] * $GRID_UNITS_PER_MM, $SIZE[2] * $GRID_UNITS_PER_MM);
    @grid_centre = ($SIZE[0] * $GRID_UNITS_PER_MM / 2, $SIZE[1] * $GRID_UNITS_PER_MM / 2,
        $SIZE[2] * $GRID_UNITS_PER_MM / 2);
    @size = ($SIZE[0] * $GRID_UNITS_PER_MM - 1, $SIZE[1] * $GRID_UNITS_PER_MM - 1, $SIZE[2] * $GRID_UNITS_PER_MM - 1);
    print FILE "; $file\n; Created using makeGeo.perl, by Dan Crick.\n\n";
    print FILE "pa_define($grid_size[0], $grid_size[1], $grid_size[2], planar, non-mirrored, Electrostatic)\n";
    print FILE "locate($grid_centre[0], $grid_centre[1], $grid_centre[2], $GRID_UNITS_PER_MM) {
        ; We have $GRID_UNITS_PER_MM grid points per mm, and have shifted to the centre.\n\n";
    }
}

sub Print_File {
    print "\n";
    system "cat $file";
    print "\n\n";
    system "cat M_$file";
    print "\n";
}

sub Put { # Coords should be passed here in mm
    ($eID, $x, $y) = @_;
    $width = 2 * $R * cos($PI/6);
    $height = $R;
    print FILE "locate($x,$y,$z, 1, 0, 90) {\n";
    print FILE "e($eID) {fill {within {centered_box3d(0,0,0, $width, $height, $thickness)}}}\n\n";
}
```

Chapter B

```
print FILE "locate($x,$y,$z, 1, 0, 150) {\n";
print FILE "e($eID) {fill {within {centered_box3d(0,0,0, $width, $height, $thickness)}}}\n\n";
print FILE "locate($x,$y,$z, 1, 0, 210) {\n";
print FILE "e($eID) {fill {within {centered_box3d(0,0,0, $width, $height, $thickness)}}}\n\n";
}

sub Drill {
  ($x, $y) = @_;
  $holelength = $SIZE[2];
  print FILE "locate($x,$y,$z, 1, 0, 0) {\n";
  print FILE "non_electrode() {fill {within {cylinder(0,0,0, $holesize, $holesize, $holelength)}}}\n\n";
  print FILE "locate($x,$y,$z, 1, 180, 0) {\n";
  print FILE "non_electrode() {fill {within {cylinder(0,0,0, $holesize, $holesize, $holelength)}}}\n\n";
}

sub Put_Single_Array {

#####
## Electrode 1 #####
#####

Put $e[1], 0.0, (2*$R*cos($PI/6) + $D);
Put $e[1], (3*$R + 2*$D*cos($PI/6)), (2*$R*cos($PI/6) + $D);
Put $e[1], -(3*$R + 2*$D*cos($PI/6)), (2*$R*cos($PI/6) + $D);
Put $e[1], (6*$R + 4*$D*cos($PI/6)), (2*$R*cos($PI/6) + $D);
Put $e[1], -(6*$R + 4*$D*cos($PI/6)), (2*$R*cos($PI/6) + $D);

# Connect the top row:
$height = $R * cos($PI/6);
$width = 33.5;
$bottom = (2*$R*cos($PI/6) + $D) + $D*0.5;
$left = -33.5/2;
print FILE "e($e[1]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$height = 7.5;
$width = 0.7;
$left = -33.5/2;
$bottom = $D/2;
print FILE "e($e[1]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$height = 7.5;
$width = 0.7;
$left = 33.5/2 - $width;
$bottom = $D/2;
print FILE "e($e[1]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$left = -($R + $D*cos($PI/6) + 0.5*$R) - (3*$R + 2*$D*cos($PI/6)) - $R - $D;
$width = 0.5*$R;
$bottom = $D/2;
$height = 0.866*$R;
print FILE "e($e[1]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$width = 0.5*$R;
$left = ($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6)) + $R + $D - $width;
$bottom = $D/2;
$height = 0.866*$R;
print FILE "e($e[1]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

# Cut away the bits we don't want.
$x = -($R + $D*cos($PI/6) + 0.5*$R) - (3*$R + 2*$D*cos($PI/6));
$y = 0.5*$D + $R*cos($PI/6);
$width = 2 * $R * cos($PI/6) + $D;
$height = $R;
print FILE "locate($x,$y,$z, 1, 0, 210) {\n";
print FILE "non_electrode() {fill {within {centered_box3d(0,0,0, $width, $height, $thick)}}}\n\n";

$x = ($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6));
$y = 0.5*$D + $R*cos($PI/6);
$width = 2 * $R * cos($PI/6) + $D;
$height = $R;
print FILE "locate($x,$y,$z, 1, 0, 150) {\n";
print FILE "non_electrode() {fill {within {centered_box3d(0,0,0, $width, $height, $thick)}}}\n\n";
```

Chapter B

```
#####
## Electrode 5 #####
#####

Put $e[5], 0.0, -(2*$R*cos($PI/6) + $D);
Put $e[5], (3*$R + 2*$D*cos($PI/6)), -(2*$R*cos($PI/6) + $D);
Put $e[5], -(3*$R + 2*$D*cos($PI/6)), -(2*$R*cos($PI/6) + $D);
Put $e[5], (6*$R + 4*$D*cos($PI/6)), -(2*$R*cos($PI/6) + $D);
Put $e[5], -(6*$R + 4*$D*cos($PI/6)), -(2*$R*cos($PI/6) + $D);

# Connect the bottom row:
$height = ($R * cos($PI/6));
#$width = (($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6)) + $R) * 2;
$width = 33.5;
$bottom = -(3*$R*cos($PI/6) + $D) - $D*0.5;
#$left = -(($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6)) + $R);
$left = -33.5/2;
print FILE "e($e[5]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$height = 7.5;
$width = 0.7;
$left = -33.5/2;
$bottom = -7.5 - $D/2;
print FILE "e($e[5]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$height = 7.5;
$width = 0.7;
$left = 33.5/2 - $width;
$bottom = -7.5 - $D/2;
print FILE "e($e[5]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$left = -($R + $D*cos($PI/6) + 0.5*$R) - (3*$R + 2*$D*cos($PI/6)) - $R - $D;
$width = 0.5*$R;
$height = 0.866*$R;
$bottom = - $height - $D/2;
print FILE "e($e[5]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

$width = 0.5*$R;
$left = ($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6)) + $R + $D - $width;
$height = 0.866*$R;
$bottom = -$height - $D/2;
print FILE "e($e[5]) {fill {within {corner_box3d($left, $bottom, $z, $width, $height, $thickness)}}}\n";

# Cut away the bits we don't want.
$x = -($R + $D*cos($PI/6) + 0.5*$R) - (3*$R + 2*$D*cos($PI/6));
$width = 2 * $R * cos($PI/6) + $D;
$height = $R;
$y = -(0.5*$D + $R*cos($PI/6));
print FILE "locate($x,$y,$z, 1, 0, 150) {\n";
print FILE "non_electrode() {fill {within {centered_box3d(0,0,0, $width, $height, $thick)}}}\n\n";

$x = ($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6));
$width = 2 * $R * cos($PI/6) + $D;
$height = $R;
$y = -(0.5*$D + $R*cos($PI/6));
print FILE "locate($x,$y,$z, 1, 0, 210) {\n";
print FILE "non_electrode() {fill {within {centered_box3d(0,0,0, $width, $height, $thick)}}}\n\n";

#####

# Now the rest of the electrodes...

Put $e[3], 0.0, 0.0;
Put $e[3], (3*$R + 2*$D*cos($PI/6)), 0.0;
Put $e[3], -(3*$R + 2*$D*cos($PI/6)), 0.0;

Put $e[2], $R + $D*cos($PI/6) + 0.5*$R, 0.5*$D + $R*cos($PI/6);
Put $e[2], -($R + $D*cos($PI/6) + 0.5*$R), 0.5*$D + $R*cos($PI/6);
Put $e[2], ($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6)) , 0.5*$D + $R*cos($PI/6);
Put $e[2], -($R + $D*cos($PI/6) + 0.5*$R) - (3*$R + 2*$D*cos($PI/6)) , 0.5*$D + $R*cos($PI/6);

Put $e[4], $R + $D*cos($PI/6) + 0.5*$R, -(0.5*$D + $R*cos($PI/6));
Put $e[4], -($R + $D*cos($PI/6) + 0.5*$R), -(0.5*$D + $R*cos($PI/6));
Put $e[4], ($R + $D*cos($PI/6) + 0.5*$R) + (3*$R + 2*$D*cos($PI/6)) , -(0.5*$D + $R*cos($PI/6));
```

Chapter B

```
Put $e[4], -($R + $D*cos($PI/6) + 0.5*$R) - (3*$R + 2*$D*cos($PI/6)) , -(0.5*$D + $R*cos($PI/6));

# Now make the 1mm holes:
$holesize = $HOLE1;
Drill 0, 0;
Drill (3*$R + 2*$D*cos($PI/6)), 0;
Drill -(3*$R + 2*$D*cos($PI/6)), 0;

# And the 0.5mm holes:
$holesize = $HOLE0;
Drill 0, 6.25;
Drill -12.35, 3.8;
Drill -4.35, 3.8;
Drill 4.35, 3.8;
Drill 13.25, 3.8;
Drill -12.1, -3.7;
Drill -4.35, -3.75;
Drill 4.35, -3.65;
Drill 12.1, -3.7;
Drill 0, -6.35;
}

sub Write_Magnet {
    open FILE, "> M_$file" or die "Can't open M_$file : $!";
    print FILE "; M_$file\n; Created using makeGeo.perl, by Dan Crick.\n\n";
    $z1 = $SIZE[2] * $GRID_UNITS_PER_MM - 1;
    $ng = $SIZE[2] * $GRID_UNITS_PER_MM - 1;
    print FILE "pa_define($grid_size[0], $grid_size[1], $grid_size[2], planar, non-mirrored, Magnetic, $ng)\n";
    print FILE "e(0){ fill{ within{corner_box3d(0, 0, 0, $grid_size[0], $grid_size[1], 0)}}}\n";
    print FILE "e(10000){ fill{ within{corner_box3d(0, 0, $z1, $grid_size[0], $grid_size[1], $z1)}}}\n";
    close FILE;
}

#####

$file = $ARGV[0];
open FILE, "> $file" or die "Can't open $file : $!";

Write_Header;

$z = -$Z/2;
Put_Single_Array;
#Put 1, 0, 0;

$z = $Z/2 + $thickness;
Put_Single_Array;

# Cut away the bits which go outside the board:
$x = -$SIZE[0]/2;
$y = -$SIZE[1]/2;
$z = -$SIZE[2]/2;
$width = ($SIZE[0]/2) - (33.5/2);
$height = $SIZE[1];
$depth = $SIZE[2];
print FILE "locate($x,$y,$z, 1, 0, 0) {\n";
print FILE "non_electrode() {fill {within {corner_box3d(0,0,0, $width, $height, $depth)}}}\n\n";

$width = ($SIZE[0]/2) - (33.5/2);
$height = $SIZE[1];
$depth = $SIZE[2];
$x = $SIZE[0]/2 - $width;
$y = $SIZE[1]/2 - $height;
$z = $SIZE[2]/2 - $depth;
print FILE "locate($x,$y,$z, 1, 0, 0) {\n";
print FILE "non_electrode() {fill {within {corner_box3d(0,0,0, $width, $height, $depth)}}}\n\n";

$x = -$SIZE[0]/2;
$y = -$SIZE[1]/2;
$z = -$SIZE[2]/2;
$width = $SIZE[0];
$height = ($SIZE[1]/2) - (15/2);
$depth = $SIZE[2];
print FILE "locate($x,$y,$z, 1, 0, 0) {\n";
```

Chapter B

```
print FILE "non_electrode() {fill {within {corner_box3d(0,0,0, $width, $height, $depth)}}}\n\n";

$width = $SIZE[0];
$height = ($SIZE[1]/2) - (15/2);
$depth = $SIZE[2];
$x = $SIZE[0]/2 - $width;
$y = $SIZE[1]/2 - $height;
$z = $SIZE[2]/2 - $depth;
print FILE "locate($x,$y,$z, 1, 0, 0) {\n";
print FILE "non_electrode() {fill {within {corner_box3d(0,0,0, $width, $height, $depth)}}}\n\n";

# End the locate clause
print FILE "}\n";

# Draw the grounded box:
print FILE "e(0) {fill {within {corner_box3d(0, 0, 0, $grid_size[0], $grid_size[1], 0)}}}\n";
print FILE "e(0) {fill {within {corner_box3d(0, 0, $size[2], $grid_size[0], $grid_size[1], $size[2])}}}\n";
print FILE "e(0) {fill {within {corner_box3d(0, 0, 0, $grid_size[0], 0, $grid_size[2])}}}\n";
print FILE "e(0) {fill {within {corner_box3d(0, $size[1], 0, $grid_size[0], $size[1], $grid_size[2])}}}\n";
print FILE "e(0) {fill {within {corner_box3d(0, 0, 0, 0, $grid_size[1], $grid_size[2])}}}\n";
print FILE "e(0) {fill {within {corner_box3d($size[0], 0, 0, $size[0], $grid_size[1], $grid_size[2])}}}\n";

close FILE;

Write_Magnet;
Print_File;
```

Appendix C

Data Analysis Script

```
#!/usr/bin/perl -w

if ($ARGV[0] && $ARGV[0] eq "-v") {$verbose = 1;} else {$verbose = 0;}

$Delta_t = 0.05; # Signal must be above / below threshold for at least this amount of time to register a change

sub Main_Loop
{
    $file = "$magnet_current.complete";
    @S = (); # Signal
    @t = (); # Time
    $threshold_estimate = 5;

    # Load data
    if ($verbose) {print "\nLoading $file";}
    open FIN, "<$file" or die "Couldn't open $file\n";
    $i = 0;
    while (<FIN>)
    {
        if (/(\S+)\s(\S+)/)
        {
            $t[$i] = $1;
            $S[$i] = $2;
            $raw_hist[$S[$i]]++;
            $i++;
        }
    }
    close FIN;
    $i_max = $i;

    # Set the threshold level
    $sum_S = 0;
    $sum_bins = 0;
    for ($i = 0; $i < $i_max; $i++)
    {
        if ($S[$i] < $threshold_estimate)
        {
            $sum_S += $S[$i];
            $sum_bins++;
        }
    }
    if ($sum_bins != 0) {
        $S_dark = $sum_S / $sum_bins;
    }
    else {
        $S_dark = 0;
    }

    $sum_S = 0;
    $sum_bins = 0;
}
```


Chapter C

```
for ($i = 0; $i < $i_max; $i++)
{
    if ($S[$i] >= $threshold_estimate)
    {
        $sum_S += $S[$i];
        $sum_bins++;
    }
}
$S_bright = $sum_S / $sum_bins;
$average_bright_level += $S_bright;
$S_threshold = (2*$S_dark + $S_bright) / 3.0;
if ($verbose) {print "Threshold set at $S_threshold\n";}

# Run through and collect the data
$darkcount = $brightcount = 0;
if ($S[0] <= $S_threshold)
{
    $ion_on = 0;
    $darkcount++;
}
else
{
    $ion_on = 1;
    $brightcount++;
}
$dark_timer_start = 0;
$dark_timer_finish = 0;
$bright_timer_start = 0;
$bright_timer_finish = 0;
$signal_total = 0;
for ($i = 1; $i < $i_max - 4; $i++)
{
    if ($ion_on == 0 && $S[$i] <= $S_threshold)
    { # Still dark. Do nothing yet.
        $darkcount++;
    }
    elseif ($ion_on == 1 && $S[$i] > $S_threshold)
    { # Still bright.
        $signal_total += ($S[$i] - $S_dark);
        $brightcount++;
    }
    elseif ($ion_on == 1
        && $S[$i] <= $S_threshold
        && $S[$i+1] <= $S_threshold
        && $S[$i+2] <= $S_threshold
        && $S[$i+3] <= $S_threshold
        && $S[$i+4] <= $S_threshold)
    { # Ion has just gone dark.
        $darkcount++;
        $ion_on = 0;
        $dark_timer_start = $t[$i];
        $bright_timer_finish = $t[$i];
        $delta_t = $bright_timer_finish - $bright_timer_start;
        $total_brighttime += $delta_t;
        $total_compensated_brighttime += $signal_total;
        push (@brighttimes, $delta_t);
        push (@compensated_brighttimes, $signal_total);
        $signal_total = 0;
    }
    elseif ($ion_on == 0
        && $S[$i] > $S_threshold
        && $S[$i+1] > $S_threshold
        && $S[$i+2] > $S_threshold
        && $S[$i+3] > $S_threshold
        && $S[$i+4] > $S_threshold)
    { # Ion has just gone bright.
        $brightcount++;
        $ion_on = 1;
        $bright_timer_start = $t[$i];
        $dark_timer_finish = $t[$i];
        $delta_t = $dark_timer_finish - $dark_timer_start;
        $total_darktime += $delta_t;
        push (@darktimes, $delta_t);
        $signal_total += ($S[$i] - $S_dark);
    }
}
```

Chapter C

```
}
elseif ($ion_on == 1)
{ # Still bright but fell below threshold on this bin
  $signal_total += ($S[$i] - $S_dark);
}
elseif ($ion_on == 0)
{ # Still dark but went above threshold on this bin
}
else
{
  print "Error!\n";
}
}

# Work out the mean and variance:
if ($#darktimes > 0 && $#brighttimes > 0)
{
  $average_darktime = $total_darktime / ($#darktimes+1);
  $average_brighttime = $total_brighttime / ($#brighttimes+1);
  $average_compensated_brighttime =
    $total_compensated_brighttime / ($#brighttimes+1);

  if ($verbose)
  {
    print "Average darktime = $average_darktime, Average brighttime = $average_brighttime,
      Average compensated brighttime = $average_compensated_brighttime\n";
  }

  # Now bootstrap to get the error estimate:
  if ($verbose) {print "Bootstrapping...\n";}
  $bootstrap_size = 10000;
  $bs_average_darktime = $bs_average_brighttime = $bs_average_compensated_brighttime = 0;
  $mean_bootstrapped_darktime = 0;
  $mean_bootstrapped_brighttime = 0;
  $mean_bootstrapped_compensated_brighttime = 0;
  $variance_bootstrapped_darktime = 0;
  $variance_bootstrapped_brighttime = 0;
  $variance_bootstrapped_compensated_brighttime = 0;
  for ($bootstrap_counter = 0; $bootstrap_counter < $bootstrap_size; $bootstrap_counter++)
  {
    # Dark periods:
    $bs_total_darktime = 0;
    for ($i = 0; $i <= $#darktimes; $i++)
    {
      $bs_total_darktime += $darktimes[int(rand($#darktimes+1))];
    }
    $bs_average_darktime = $bs_total_darktime / ($#darktimes+1);
    $mean_bootstrapped_darktime += $bs_average_darktime / $bootstrap_size;
    $variance_bootstrapped_darktime += ($bs_average_darktime - $average_darktime)**2 / $bootstrap_size;

    # Bright periods:
    $bs_total_brighttime = 0;
    $bs_total_compensated_brighttime = 0;
    for ($i = 0; $i <= $#brighttimes; $i++)
    {
      $n = int(rand($#brighttimes+1));
      $bs_total_brighttime += $brighttimes[$n];
      $bs_total_compensated_brighttime += $compensated_brighttimes[$n];
    }
    $bs_average_brighttime = $bs_total_brighttime / ($#brighttimes+1);
    $mean_bootstrapped_brighttime += $bs_average_brighttime / $bootstrap_size;
    $variance_bootstrapped_brighttime += ($bs_average_brighttime - $average_brighttime)**2 / $bootstrap_size;
    $bs_average_compensated_brighttime = $bs_total_compensated_brighttime / ($#brighttimes+1);
    $mean_bootstrapped_compensated_brighttime += $bs_average_compensated_brighttime / $bootstrap_size;
    $variance_bootstrapped_compensated_brighttime +=
      ($bs_average_compensated_brighttime - $average_compensated_brighttime)**2 / $bootstrap_size;
  }

  $sigma_dark = sqrt($variance_bootstrapped_darktime);
  $sigma_bright = sqrt($variance_bootstrapped_brighttime);
  $sigma_compensated_bright = sqrt($variance_bootstrapped_compensated_brighttime);
  if ($verbose)
  {
```

Chapter C

```
    print "$mean_bootstrapped_darktime, $sigma_dark\n";
    print "$mean_bootstrapped_brighttime, $sigma_bright\n";
    print "$mean_bootstrapped_compensated_brighttime, $sigma_compensated_bright\n";
  }
}

open DARKTIME, ">./darktime_vs_B";
open JUMPFREQ, ">./jump_frequency_vs_B";
open COMPFREQ, ">./compensated_jump_frequency_vs_B";

for ($magnet_current = 0; $magnet_current <= 30; $magnet_current++)
{
  if (-d ".$magnet_current")
  {
    $average_darktime = 0;
    $total_darktime = 0;
    @darktimes = ();
    $average_brighttime = 0;
    $total_brighttime = 0;
    @brighttimes = ();
    $average_compensated_brighttime = 0; # Here we multiply by the signal level
    $total_compensated_brighttime = 0;
    @compensated_brighttimes = ();
    $average_bright_level = 0;
    @raw_hist = ();
    print "$magnet_current";
    Main_Loop();
    print "\n";

    # Make histogram files:
    open HISTFILE, ">$magnet_current.hist";
    for ($i = 0; $i < $#raw_hist; $i++) {
      if ($raw_hist[$i]) {
        $y = $raw_hist[$i];
      }
      else {
        $y = 0;
      }
      print HISTFILE "$i $y\n";
    }
    close HISTFILE;

    # Find the B field value
    open MAGVSCURRENT, "<./magnetic_field_vs_current";
    while (<MAGVSCURRENT>)
    {
      if (/(\S+)\s+(\S+)/)
      {
        $I = $1;
        $B = $2;
        if ($I == $magnet_current)
        {
          $x = $B**2;
        }
      }
    }
    close MAGVSCURRENT;

    $B_error = 0.005 * 2 * sqrt($x);
    $s = $sigma_dark;
    # Delta_t fixes the shift caused by the 50ms bin period.
    $dark_tau = ($average_darktime - $Delta_t);
    $bright_tau = ($average_brighttime - $Delta_t);
    $comp_tau = $average_compensated_brighttime - $Delta_t*($s_bright-$s_dark);
    print DARKTIME "$x $dark_tau $B_error $s\n";
    $jumpfreq = 1.0 / $bright_tau;
    $s = ($sigma_bright / $average_brighttime) * $jumpfreq;
    print JUMPFREQ "$x $jumpfreq $B_error $s\n";
    $jumpfreq = 1.0 / $comp_tau;
    $s = sqrt($sigma_compensated_bright / $average_compensated_brighttime);
    print COMPFREQ "$x $jumpfreq $B_error $s\n";
  }
}
```

Chapter C

```
}

close DARKTIME;
close JUMPFREQ;
close COMPFREQ;

if($verbose) {print "Plotting results...\n";}
system "pyxplot plot_data.pyxplot";
system "acroread preliminary_results.pdf";
```

Bibliography

- [1] P. K. Ghosh. Ion Traps (Oxford University Press, 1995).
- [2] F. Major, V. Gheorghe, and G. Werth. Charged Particle Traps: Physics and Techniques of Charged Particle Field Confinement (Springer, 2004).
- [3] S. Earnshaw. On the nature of the molecular forces which regulate the constitution of the luminiferous ether. *Trans. Camb. Phil. Soc.* **7**, 97 (1842).
- [4] D. J. Bate, K. Dholakia, R. C. Thompson, and D. C. Wilson. Ion oscillation frequencies in a combined trap. *J. Mod. Opt.* **39**, 305 (1992).
- [5] E. Fischer. Die dreidimensionale Stabilisierung von Ladungsträgern in einem Vierpolfeld. *Z. Phys.* **156**, 1 (1959).
- [6] T. W. Hansch and A. L. Schawlow. Cooling of gases by laser radiation. *Optics Communications* **13**, 68 (1975).
- [7] D. J. Wineland and H. Dehmelt. Proposed $10^{14}\delta\nu/\nu$ laser fluorescence spectroscopy on Tl^+ mono-ion oscillator. *Bull. Am. Phys. Soc.* **20**, 657 (1975).
- [8] Encyclopedia of Chemical Physics and Physical Chemistry, volume 3 (Taylor and Francis, 2001).
- [9] I. Marzoli, J. I. Cirac, R. Blatt, and P. Zoller. Laser cooling of trapped three-level ions: Designing two-level systems for sideband cooling. *Phys. Rev. A* **49**, 2771 (1994).

- [10] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Laser cooling to the zero-point energy of motion. *Phys. Rev. Lett.* **62**, 403 (1989).
- [11] F. Schmidt-Kaler *et al.* Realization of the Cirac-Zoller controlled-NOT quantum gate. *Nature* **422**, 408 (2003).
- [12] D. R. Crick *et al.* Two-ion Coulomb crystals of Ca^+ in a Penning trap. *Optics Express* **16**, 2351 (2008).
- [13] W. M. Itano and D. J. Wineland. Laser cooling of ions stored in harmonic and Penning traps. *Phys. Rev. A* **25**, 35 (1982).
- [14] R. J. Hendricks, E. S. Phillips, D. M. Segal, and R. C. Thompson. Laser cooling in the Penning trap: an analytical model for cooling rates in the presence of an axializing field. *J. Phys. B* **41**, 035301 (2008).
- [15] R. Van Dyck Jr, P. Ekstrom, and H. Dehmelt. Axial, magnetron, cyclotron and spin-cyclotron-beat frequencies measured on single electron almost at rest in free space (geonium). *Nature* **262**, 776 (1976).
- [16] E. S. Phillips *et al.* Dynamics of axialized laser cooled ions in a Penning trap. *Phys. Rev. A* **78**, 032307 (2008).
- [17] H. F. Powell *et al.* Improvement of laser cooling of ions in a Penning trap by use of the axialization technique. *J. Phys. B* **36**, 961 (2003).
- [18] G. Taubes. An interview with Dr. David Deutsch. *ESI Special Topics* (2002).
- [19] C. Bennett. Logical reversibility of computation. *IBM J. Res. Develop.* **17**, 525 (1980).
- [20] P. Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *J. Stat. Phys.* **22**, 563 (1980).
- [21] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proc. R. Soc. Lond.* **400**, 97 (1985).
- [22] R. Feynman. Quantum mechanical computers. *Opt. News*. **11**, 11 (1985).
- [23] A. White. Quote from lecture given during SCALA summer school (2006).

- [24] Nielson and Chuang. Quantum Computation and Quantum Information (Cambridge University Press, 2000).
- [25] P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. *Proc. 35nd Annual Symposium on Foundations of Computer Science* page 124 (1994).
- [26] L. M. K. Vandersypen *et al.* Experimental realization of Shor’s quantum factoring algorithm using nuclear magnetic resonance. *Nature* **414**, 883 (2001).
- [27] D. Porras and J. I. Cirac. Quantum manipulation of trapped ions in two dimensional Coulomb crystals. *Phys. Rev. Lett.* **96**, 250501 (2006).
- [28] M. A. Nielsen. Cluster-state quantum computation. *arXiv:quant-ph/0504097* (2005).
- [29] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.* **74**, 4091 (1995).
- [30] A. Sørensen and K. Mølmer. Quantum computation with ions in thermal motion. *Phys. Rev. Lett.* **82**, 1971 (1999).
- [31] J. I. Cirac and P. Zoller. A scalable quantum computer with ions in an array of microtraps. *Nature* **404**, 579 (2000).
- [32] G. J. Milburn, S. Schneider, and D. F. V. James. Ion trap quantum computing with warm ions. *Fortschr. Phys.* **48**, 801 (2000).
- [33] C. Monroe *et al.* Demonstration of a Fundamental Quantum Logic Gate. *Phys. Rev. Lett.* **75**, 4714 (1995).
- [34] D. Leibfried *et al.* Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature* **422**, 412 (2003).
- [35] D. Leibfried *et al.* Creation of a six-atom ‘Schrödinger cat’ state. *Nature* **438**, 639 (2005).
- [36] H. Häffner *et al.* Scalable multiparticle entanglement of trapped ions. *Nature* **438**, 643 (2005).
- [37] A. Barenco *et al.* Elementary gates for quantum computation. *Phys. Rev. A* **52**, 3457 (1995).

- [38] D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik* **48**, 771 (2000).
- [39] ARDA. Quantum information science and technology roadmap. (2004).
- [40] W. S. Warren, N. Gershenfeld, and I. Chuang. The usefulness of NMR quantum computing. *Science* **277**, 1688 (1997).
- [41] D. Beckman, A. N. Chari, S. Devabhaktuni, and J. Preskill. Efficient networks for quantum factoring. *Phys. Rev. A* **54**, 1034 (1996).
- [42] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature* **417**, 709 (2002).
- [43] S. Seidelin *et al.* Microfabricated surface-electrode ion trap for scalable quantum information processing. *Phys. Rev. Lett.* **96**, 253003 (2006).
- [44] J. Britton *et al.* A microfabricated surface-electrode ion trap in silicon. *arXiv:quant-ph/0605170* (2006).
- [45] D. Stick *et al.* Ion trap in a semiconductor chip. *Nature Physics* **2**, 36 (2006).
- [46] G. Huber *et al.* Transport of ions in a segmented linear Paul trap in printed-circuit-board technology. *New J. Phys.* **10**, 013004 (2008).
- [47] C. E. Pearson *et al.* Experimental investigation of planar ion traps. *Phys. Rev. A* **73**, 032307 (2006).
- [48] D. L. Moehring *et al.* Entanglement of single-atom quantum bits at a distance. *Nature* **449**, 68 (2007).
- [49] D. J. Wineland *et al.* Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *J. Res. Natl. Inst. Stand. Tech.* **103**, 259 (1998).
- [50] W. K. Hensinger *et al.* T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation. *Appl. Phys. Lett.* **88**, 034101 (2006).
- [51] R. Hendricks. Spectroscopy and dynamics of laser-cooled Ca^+ ions in a Penning trap. Ph.D. thesis, Imperial College London (2006).

- [52] H. Ohadi. Single Ca^+ ions in a Penning trap for applications in quantum information processing. Ph.D. thesis, Imperial College London (2008).
- [53] E. S. G. Phillips. Controlled dynamics of laser-cooled ions in a Penning trap. Ph.D. thesis, Imperial College London (2004).
- [54] D. M. Lucas *et al.* Isotope-selective photoionization for calcium ion trapping. *Phys. Rev. A* **69**, 012711 (2004).
- [55] A. H. Myerson *et al.* High-fidelity readout of trapped-ion qubits. *Phys. Rev. Lett.* **100**, 200502 (2008).
- [56] NIST. Atomic spectra database. (2008).
- [57] N. Kjaergaard *et al.* Isotope selective loading of an ion trap using resonance-enhanced two-photon ionization. *Appl. Phys. B* **71**, 207 (2000).
- [58] S. Gulde *et al.* Simple and efficient photoionization loading of ions for precision ion-trapping experiments. *Appl. Phys. B* **73**, 861 (2001).
- [59] E. Photonics. Ridge waveguide laser for external cavity setups (datasheet). (2008).
- [60] A. S. Arnold, J. S. Wilson, and M. G. Boshier. A simple extended-cavity diode laser. *Rev. Sci. Instrum.* **69** (1998).
- [61] Melles Griot Inc. Optics guide. (2002).
- [62] NIST. Engineering metrology toolbox. (2008).
- [63] A. Devices. REF01 (datasheet). (2005).
- [64] NPL. Low drift Fabry-Perot etalon (technical specifications). (2008).
- [65] W. Z. Zhao, J. E. Simsarian, L. A. Orozco, and G. D. Sprouse. A computer-based digital feedback control of frequency drift of multiple lasers. *Rev. Sci. Instrum.* **69**, 3737 (1998).
- [66] B. M. A. Rossi, V. Biancalana and L. Tomassetti. Long-term drift laser frequency stabilization using purely optical reference. *Rev. Sci. Instrum.* **73**, 2544 (2002).
- [67] I. Bhatti. Computerised stabilisation of a multiple laser system. Master's thesis, Imperial College London (2007).

- [68] M. Storkey. Studies of laser-cooled trapped ions. Ph.D. thesis, Imperial College London (2001).
- [69] M. van Eijkelenborg *et al.* A driven, trapped, laser cooled ion cloud: a forced damped oscillator. *Optics Communications* **159**, 169 (1999).
- [70] Danfysik. System 8000 Model 853 manual.
- [71] P. Ranin *et al.* Study of the influence of the Earth's magnetic field on PMTs for Pierre-Auger surface detectors. *UCLA-Cosmic* (2001).
- [72] F. Diedrich *et al.* Observation of a phase transition of stored laser-cooled ions. *Phys. Rev. Lett.* **59**, 2931 (1987).
- [73] S. Fishman, G. de Chiara, T. Calarco, and G. Morigi. Structural phase transitions in low-dimensional ion crystals. *Phys. Rev. B* **77**, 064111 (2008).
- [74] R. C. Thompson and D. C. Wilson. The motion of small numbers of ions in a Penning trap. *Z Phys. D* **42**, 271 (1997).
- [75] D. G. Enzer *et al.* Observation of power-law scaling for phase transitions in linear trapped ion crystals. *Phys. Rev. Lett.* **85**, 2466 (2000).
- [76] D. Reiss *et al.* Raman cooling and heating of two trapped Ba^+ ions. *Phys. Rev. A* **65**, 053401 (2002).
- [77] G. Birkel, S. Kassner, and H. Walther. Multiple-shell structures of laser-cooled $^{24}\text{Mg}^+$ ions in a quadrupole storage ring. *Nature* **357**, 310 (1992).
- [78] T. B. Mitchell *et al.* Direct observations of structural phase transitions in planar crystallized ion plasmas. *Science* **282**, 1290 (1998).
- [79] J. Chiaverini *et al.* Realization of quantum error correction. *Nature* **432**, 602 (2004).
- [80] J. R. Castrejón-Pita *et al.* Novel designs for Penning ion traps. *J. Mod. Opt.* **54**, 1581 (2007).
- [81] H. Häffner *et al.* High-accuracy measurement of the magnetic moment anomaly of the electron bound in hydrogenlike carbon. *Phys. Rev. Lett.* **85**, 5308 (2000).

- [82] C. Rouki, L. Westerberg, and the CHICSi development group. Ultra-high vacuum compatibility measurements of materials for the CHICSi detector system. *Physica Scripta* **T104**, 107 (2003).
- [83] SMART Group. FR4 laminate. Technical report, EMT Worldwide (2002).
- [84] P. Zeeman. Light radiation in a magnetic field. *Nobel Prize Lecture* (1903).
- [85] P. A. Barton *et al.* Measurement of the lifetime of the $3d\ ^2D_{5/2}$ state in $^{40}\text{Ca}^+$. *Phys. Rev. A* **62**, 032503 (2000).
- [86] K. Toyoda *et al.* Laser cooling of calcium ions by use of ultraviolet laser diodes: significant induction of electron-shelving transitions. *Opt. Lett.* **26**, 1897 (2001).
- [87] A. V. Taichenachev *et al.* Magnetic field-induced spectroscopy of forbidden optical transitions with application to lattice-based optical atomic clocks. *Phys. Rev. Lett.* **96**, 083001 (2006).
- [88] Z. W. Barber *et al.* Direct excitation of the forbidden clock transition in neutral ^{174}Yb atoms confined to an optical lattice. *Phys. Rev. Lett.* **96**, 083002 (2006).
- [89] Semrock. Brightline single bandpass filter catalog (2008).
- [90] D. Budker, D. F. Kimball, and D. P. DeMille. *Atomic Physics* (Oxford University Press, 2004).
- [91] C. Guet and W. R. Johnson. Relativistic many-body calculations of transition rates for Ca^+ , Sr^+ and Ba^+ . *Phys. Rev. A* **44**, 1531 (1991).
- [92] A. Kreuter *et al.* New experimental and theoretical approach to the $3d\ ^2D$ -level lifetimes of $^{40}\text{Ca}^+$. *Phys. Rev. A* **71**, 032504 (2005).
- [93] L. Deslauriers *et al.* Scaling and suppression of anomalous heating in ion traps. *Phys. Rev. Lett.* **97**, 103007 (2006).