A Quantum Memory Qubit in Calcium-43

Experimental Systems and Control

A thesis submitted for the degree of

Doctor of Philosophy

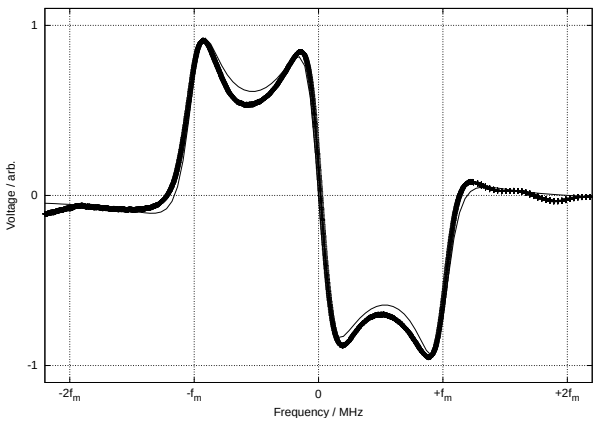

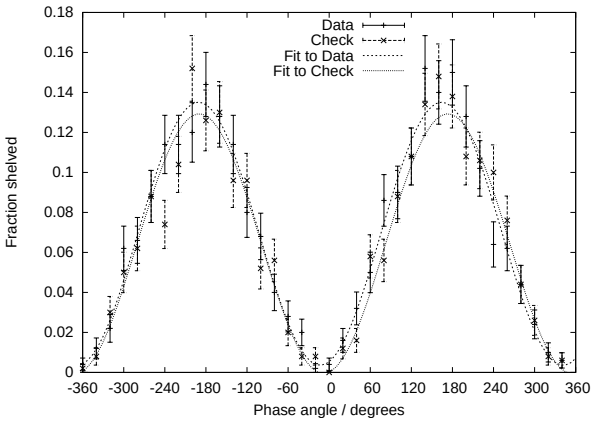

Benjamin Keitch

Trinity Term — Wolfson College

2007 — Oxford

Abstract

A Quantum Memory Qubit in Calcium-43

Benjamin Keitch, Wolfson College

Thesis for Doctor of Philosophy

Trinity 2007

The hyperfine ground states of a $^{43}Ca^{+}$ trapped ion are proposed as a memory qubit. The experimental details to implement such a qubit are presented. These include the design and implementation of a pulse-sequencer to produce pulse sequences for single qubit rotations and other timed experimental sequences. A Raman laser system running at 396nm with a frequency offset of 3.2GHz, based on optically injected violet diode lasers, is presented along with an analysis of its performance. A laser-locking scheme based on the Pound-Drever-Hall method has been built and the design and results are presented here. The dominant limitation on the coherence of the memory qubit, namely fluctuations in the ambient magnetic field, is studied and quantified. The coherence of a superposition of the ground hyperfine clock states of $^{43}Ca^{+}$, $|\downarrow\rangle \equiv S_{1/2}(F=4, M_F=0)$ and $|\uparrow\rangle \equiv S_{1/2}(F=3, M_F=0)$ is measured using the Ramsey technique. The coherence of the same state after implementing a spin-echo pulse sequence is also measured.

The coherence time (T_2) for the memory qubit, using the Ramsey technique is found to be 1.2(2)s. With the implementation of a spin-echo sequence, bounds are put on the spin-echo coherence-time as : $2\,\text{min} \lesssim T_2^{SE} \lesssim 10\,\text{min}$. The difference in these two results leads to the conclusion that the main source of decoherence is due to magnetic field drift. The $^{43}Ca^{+}$ memory qubit has a coherence time that is 10^4 times greater than current gate times, which is sufficient to reach a regime where fault-tolerant quantum information processing should be possible.

Acknowledgements

I wish to thank my supervisor David Lucas for his unerring patience and tolerance over the years.

All the ion-trappers past and present: Andrew Steane, Derek Stacey, Jonathan Home, Matthew McDonnell, Simon Webster, and Michael Curtis. In particular the following people who's work has been included in this thesis (where referenced): Gergely Imreh, David Szwer, and Alice Myerson. To Marek Šašura for helping me understand ion traps and Nicholas Thomas who had to endure sharing an office with me and whose chats were invaluable.

To my parents for their support. Everyone at Wolfson College for sharing the joy and the pain of graduate life.

Finally to İldem Akerman for making me believe it was possible.

Contents

"If you want to make a simulation of nature, you'd better make it quantum mechanical, and by golly it's a wonderful problem, because it doesn't look so easy!"

R P Feynman

Chapter 1

Introduction to Trapped-Ion Quantum Computing with $^{43}Ca^{+}$

1.1 Overview of Computing

Quantum computing has already been well-established theoretically and is a rapidly growing area of physics. The focus is now on realising large-scale quantum computing and the technical difficulties faced in scaling up from a handful of individual quantum systems ("qubits") to a sufficient number to carry out useful computations.

Richard Feynman first proposed quantum computing as an answer to the problem that physicists were facing with modelling quantum systems on conventional computers. Given the computational power available at that time (1980) it seemed that even with coarse approximations, solving problems in physics on a conventional computer would not be efficient. In the same way that a Universal Turing machine can model any classical computer, Feynman envisaged a quantum system that could model any other quantum system [1]. It was Deutsch who first proposed the quantum circuit [2], a way of performing algorithms on a quantum computer with a series of steps, including operations that are analogous to classical logic gates, but use the property of entanglement. At the same time, progress was being made in trapping and cooling single atoms and performing high resolution spectroscopy and metrology experiments with single atoms manipulated by lasers.

However, because the no-cloning theorem forbids the copying of quantum information, should a so-called "qubit" (quantum analogue of a bit) be lost or its state lose coherence (due to measurement, or interaction with its environment) the information is lost. This seems inevitable given any noise in the system. However, due to the independent discovery of quantum error correction by Shor and Steane, it was shown that by having ancillary qubits, these errors could be corrected for [3]. This made the possibility of quantum computers a reality, and in 1994, NIST combined the ideas of trapped single ions and quantum computing and performed the first two qubit gate in an isolated atom [4]. Many other schemes for quantum computing have been proposed and some have been demonstrated. These are discussed below.

A Quantum Information Processor (QIP) can be used for two main tasks. The first, and most highly publicised, is to solve problems that are currently inefficient to solve using ordinary computers. These problems include the factoring of large numbers. The current best classical algorithms for solving these problems are considered inefficient in computing terms, as the time taken to run the algorithm scales exponentially with respect to the size of the problem (it is in class NP). To be efficient they would need to scale as a polynomial function of the size of the problem (in class P). The Shor factoring algorithm [5] utilises the properties of entanglement and superposition in a QIP to reduce this problem from class NP to class P. This means that as the size of the problem grows, the computer can still realistically scale (in size or speed) to be able to solve the problem. It is this difference which would make a QIP profitable in decrypting information that has been encrypted using current public-key cryptography schemes, such as are used on the Internet.

Currently these systems are only secure because as the power of the computer being used to crack the encrypted message increases, it is a simple matter to increase the size of the key and hence the difficulty of decryption. The encryption and permitted decryption remain efficient, whilst unauthorised decryption (cracking) will always be inefficient. It is not the case that the encrypted message can not be decrypted, it is just that it is always made less worthwhile than the reward. For instance, to obtain one credit card number, encrypted using a 640 bit key, it would take 5 months of computing time on a specialist built machine with 80 parallel processors [7]. At today's prices, this would cost approximately £10 000 – far beyond the worth of the credit card

details. Current websites use 1024 bit keys, which are considered safe for the next ten years [8]. If the encryption complexity is increased (the key size increased) then the time taken to encrypt the message will also increase. However, a QIP would shift this balance of power, making decryption without a key a feasible enterprise.

The Grover search algorithm [6] uses a QIP to give an improved efficiency over classic search algorithms for searching in large sets. It finds many applications in searching large databases, such as DNA sequences (see for example [9]), and if implemented will have great commercial and academic benefits.

The second use of a QIP is Feynman's original application – that is to simulate other quantum systems. Currently chemists and biochemists use conventional computers to simulate chemical reactions and the physical aspects of molecular dynamics, which are quantum mechanical in nature. Current modelling schemes rely on making some level of approximation whilst still being very processor-intensive. It was Feynman's original vision that a general purpose quantum system such as a QIP could be used to simulate any given quantum system and produce a result in just a few steps. The requirements to build such a device are less stringent than that for a general purpose QIP, but require more individual elements or qubits [10].

Another motivation for exploring quantum computing exists. As integrated circuits get more complex and the density of transistors increases, the size of these transistors is such that quantum effects are already present. As of 2006, ICs can be produced with features 45nm across (<500 atoms). By 2018, transistors as small as 16nm are expected to be produced [11]. State-of-the-art gate oxide thicknesses are currently between 1.5 and 2.0nm and tunnelling is a substantial source of current-leakage in existing gates. So for technology to progress to smaller and faster devices, quantum effects will have to be considered – either as a disadvantage or, by using QIP, as an advantage.

1.1.1 DiVincenzo Criteria

DiVincenzo [25] summarised the following criteria that a physical system would need to fulfil to be a candidate for a realisable quantum computer:

1. The system is scalable in order to be able to increase the number of qubits.

2. The qubits can be initialised to a simple fiducial state.

3. The quantum gate time is much faster than the qubit decoherence time.

4. A universal gate set exists.

5. A qubit-specific read-out method exists.

Additionally, error correcting algorithms require that the qubit can be re-initialised in real time, and that the fidelity of the gate reaches the fault tolerant threshold. Any realistic computer will also have the ability to hold a quantum state whilst other operations are carried out – this is analogous to a register in a classical computer and is referred to as qubit memory.

To realise the basic unit of a quantum information processor it is necessary to have a single quantum system that can be precisely characterised and to be able to perform entanglement between two or more qubits in a controllable way. Any real system will have some coupling with its environment e.g. the bulk material within which it is contained, which will cause decoherence of this entangled state. Therefore the qubit must be well isolated from its environment. This is the limit imposed by the third DiVincenzo criterion: requiring either long decoherence times, or very fast gates.

Decoherence is the most demanding problem facing quantum computing presently. A lot has been learnt both theoretically and experimentally about decoherence mechanisms in the last ten years and the challenge is to now implement schemes which overcome this problem.

1.1.2 Methods for Quantum Computing

Many methods and systems have been suggested for quantum computing. These include: single photons, ions in electromagnetic traps, neutral atoms, atoms in beams interacting with cavities at optical or microwave frequencies, electronic and spin states in quantum dots, nuclear spins in a molecule in solution or in solid state (NMR), charge states (single Cooper pairs) of nanometre-scale superconductors, flux states of superconducting circuits (SQUIDs), quantum Hall systems, and states of electrons in super-fluid helium.

Most solid-state systems are based on Electron-Spin Resonance (ESR). The electron spin is addressed directly using RF fields rather than optical methods. Neighbouring qubits interact by

spin-spin coupling. Addressing can be achieved, as in NMR, using a field gradient or by varying the spin-site's orientation to the field.

Quantum dots are nano-structures of several hundred atoms either arranged periodically, or embedded in a bulk-substrate. They restrain the motion of charge-carriers in three directions such that they are three-dimensional quantum-wells, and can be used to create coherent fluorescence and laser-light. By performing ESR on electrons trapped at individual sites, it has been proposed that QIP can then be performed on these spin-systems.

Fullerenes, such as the well-known Buckminsterfullerene, or "Bucky ball", are structures made up of carbon atoms which join in a similar way to graphite, except they form a complete three-dimensional structure. By trapping a single atom of e.g. phosphorus or nitrogen inside a Bucky ball it becomes very well isolated from its environment. At cryogenic temperatures decoherence times (spin-relaxation times) of 20μs have been shown [12]. To implement a quantum computer, a row of Bucky balls, each holding one atom (qubit) would be aligned on a substrate. As with quantum dots, ESR is used to manipulate the individual qubits, and spin-spin interactions provide the entanglement. However, aligning the Bucky balls is not easy – one suggestion has been to pack them inside a carbon-nanotube, another type of fullerene.

These solid-state systems, even when cyrogenically cooled, have very short decoherence times and often lack the addressability and ease of manipulation isolated atoms, and trapped ions in particular, can offer.

Current photon experiments are normally non-deterministic – they need to be repeated until an entangled pair is produced, which normally happens with a very low probability. The more entangled photons that are required, the lower the chances of achieving this is, therefore they do not scale well.

There is also an alternative method of realising quantum computing, the measurement model for quantum computation in which only two operations are required, namely, the storage of qubits (quantum memory), and (non-destructive) projective measurements on up to four qubits at a time, in an arbitrary basis. No coherent dynamical operations are involved [28].

Cold Atoms

Cold neutral atoms that are trapped in magneto-optical traps have shown great promise for quantum computing, in particular quantum simulation [29]. Optical lattices (standing-waves of light which can hold atoms in position) can hold a great number of atoms which are isolated from each other. It is possible to entangle them in a massively parallel operation with spin-dependent confinement. By moving one lattice with respect to another, collisions occur and create a resultant entanglement. The interaction energy can be controlled by changing the intensity of the lasers. This can lead to a huge quantum register containing around 10^5 atoms. By using these state-dependent collisions, cluster states can be created.

The greatest advantage of cold atoms is the simulation of bosonic, fermionic and spin many-body systems. It is also possible to do simpler analogue simulations of solid state systems where the interaction energies, doping rates and lattice geometries can all be varied. This is of great importance for studying high-temperature superconductivity. In the near future these simulations will be possible, however, the use of cold atoms for a universal quantum computer such as Deutsch envisaged, is still some way off.

1.2 Ion Trap Quantum Computing

Ion traps have already been shown to be one of the most promising technologies for demonstrating Quantum Information Processing [13].

Ion traps are scalable, unlike for instance, NMR systems which are limited in the number of qubits they can contain. Because the qubit state can be contained within the internal state of an ion, and the ion can then be transported, this means it is only necessary to manipulate a few qubits simultaneously. Ion traps can therefore be easily scaled and multiple-trap architectures based around this idea have been proposed [18]. The necessary stability of the ion's internal state has been demonstrated [20] and is shown in this work.

Due to the separation of the atoms in the trap, it is possible to manipulate and read the internal state of each atom with high fidelity. Entanglement between ions is achieved using the trap's potential well to couple internal states to their external state.

Trapped ion experiments have been able to provide the high degree of isolation from the surrounding environment that is required to reduce decoherence to an acceptable level. As the number of trapped ions is increased and the physical scale of the trap is reduced, problems due to decoherence increase [16]. As experiments are being scaled up to a greater number of qubits, the requirement for long coherence times puts more and more stringent requirements on the technology. Sources of error include heating and field noise and instability of the laser field.

Although eventually solid-state systems may overtake cold-atom and trapped ion technology, currently many of the "firsts" have been achieved using trapped ions. The main reason for this is the high degree of isolation that can be achieved.

1.3 Progress to Date

The Oxford Ion Trap Group is using RF traps and calcium ions for quantum information processing. It is particularly interested in demonstrating Quantum Error Correction (QEC)[14].

Successful implementation of a two qubit gate in an ion trap was first demonstrated at NIST in Boulder, using a single beryllium ion [4]. In subsequent experiments using a pair of beryllium ions, a fidelity of 97% was reported [17].

The same group have also shown entanglement of six particles [32]. Unlike entanglement achieved with photons, entanglement in an ion trap is deterministic. Deterministic in this sense means being able to generate an entangled state with a high degree of certainty, at a specified time [19].

Groups at Michigan [30] and at Innsbruck [31] in 2003, and at Oxford [21] in 2006, have demonstrated gates with two trapped ions. Teleportation experiments have been carried out at in 2004 at Innsbruck [33] and at NIST [34]. Quantum error correction was demonstrated in 2004 at NIST [35] and 8-ion entanglement was achieved at Innsbruck in 2005. Ions have been entangled with photons (to allow for long-distant transportation of quantum states, or "flying qubits") in 2004, also at Michigan [36].

1.3.1 Calcium 40 Ions at Oxford

The Oxford Ion Trap group have achieved a great many milestones, using $^{40}Ca^{+}$ in a linear Paul trap. The group have successfully cooled a single ion to the Doppler limit and below. For $^{40}Ca^{+}$ the Doppler limit is 0.5mK, or an average thermal quantum number ($\bar{n} \approx kT/h\nu$) of around 10, in a trap of strength ν=800kHz. Colder temperatures, using a form of EIT cooling ($1 < \bar{n} < 2$), and pulsed Raman side-band cooling ($\bar{n} < 0.1$) have also been demonstrated [23].

Single ions have been initialised, i.e. prepared into a known state, and read back again. Rabi flopping has been demonstrated, for one and two ions. Strings of up to nine ions in length have been held stably in the trap.

A form of phase-gate using two $^{40}Ca^{+}$ ions, with entanglement being achieved by using a particular motional mode of the ions, has been demonstrated [21].

1.3.2 Calcium 43 Ions at Oxford

As well as working with $^{40}Ca^{+}$, other isotopes of calcium have been investigated. The rare isotopes $^{43}Ca^{+}$ (abundance 0.14%) and $^{46}Ca^{+}$ (abundance 0.004%) have been trapped and laser cooled, following isotope-selective photoionisation. The $^{43}Ca^{+}$ ion has an odd number of nucleons and therefore an odd spin. This means it possesses a hyperfine structure, which makes it a good candidate for use in QIP (see below).

Sympathetic cooling of $^{43}Ca^{+}$ to the ground state using $^{40}Ca^{+}$ has recently been demonstrated and is awaiting publication [43]. The use of a $^{43}Ca^{+}$ ion as a memory qubit using a long-lived internal state is presented later in this thesis.

A pair of hyperfine states in the ground level of $^{43}Ca^{+}$ are used as the qubit levels, e.g. $|\uparrow\rangle$ = $S_{1/2}$(F=3,M_F =0) and $|\downarrow\rangle$ = $S_{1/2}$(F=4,M_F =0). For a picture of the relevant levels see figure 1.1. To manipulate the qubit coherently, a Raman transition or microwaves can be used. A Raman transition transfers population via a virtual level, detuned from the $P_{1/2}$ level. Microwaves can transfer population directly, but they can not impart momentum to the ion, nor are they directional.

Using a Raman transition (a two-photon process) allows reliable coherent qubit manipulation – a one-photon excitation can easily be corrupted by a single spontaneously scattered photon unless it is in a very long-lived state (e.g. [**?**]). Using a two-photon Raman transition, the Rabi

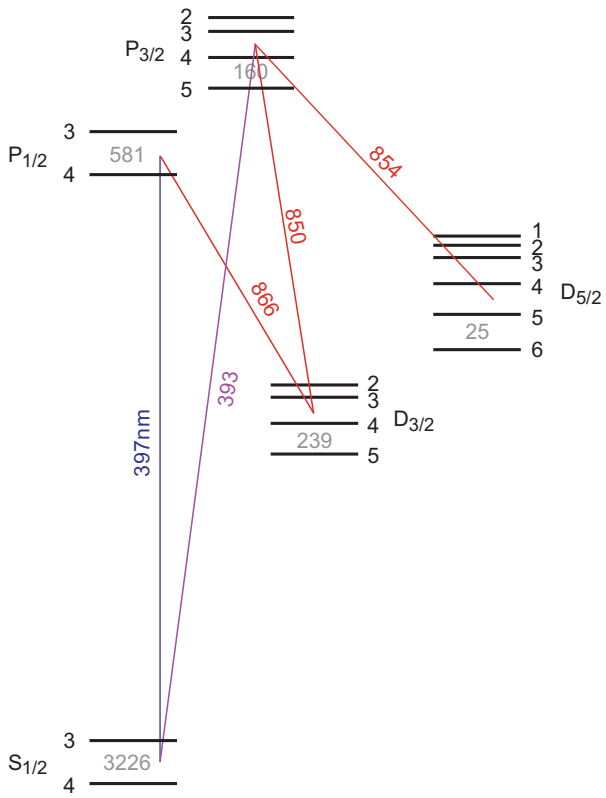

Figure 1.1: Energy level diagram for $^{43}Ca^{+}$. Transitions are shown with wavelengths in nm. Hyperfine splittings are indicated in MHz. Sub-levels are distinguished by their values of F.

frequency is proportional to $1/\Delta$ whereas the single photon scattering rate scales as $1/\Delta^2$, where Δ is the detuning from the upper (4p) level. This is discussed in more detail in chapter 4.

Some of the reasons for choosing $^{43}\mathrm{Ca}^{+}$ are:

1. It can be shown that the ratio of the Raman transition rate to the single-photon scattering rate is proportional to Δ_{fs}/Γ where Δ_{fs} is the fine-structure splitting and Γ is the width of the single-photon transition. In calcium, this ratio is reasonable (approx. 3×10^5) compared to other possible elements. This keeps the possibility of decoherence (due to a scattered photon) low. See chapter 4 for details.

2. Using a hyperfine splitting between the $|\downarrow\rangle$ state and the $|\uparrow\rangle$ state – which are both in the ground level – means that the life-time is effectively infinite.

3. The hyperfine levels are further separated by a Zeeman splitting when a small magnetic field is applied. Two magnetic-field independent states (e.g in the low-field $(\mathrm{F}=3,\mathrm{M_F}=0)$ and

$(F = 4, M_F = 0)$ or for intermediate field $(F = 4, M_F = 2)$ and $(F = 3, M_F = 3)$ can be used, making the states less sensitive to magnetic field noise (see section 1.3.2). This noise would otherwise cause decoherence of the qubit.

4. Using two Raman laser beams instead of, e.g., a microwave source allows the beams to be accurately focused onto an individual ion. The frequency difference can still be created in an equally stable way.

5. Calcium has a $^2D_{5/2}$-level which can be used as a "shelf", allowing reliable (>99%) read-out of the state at the end of an operation.

6. All the required transitions are accessible with currently available diode lasers, without the need for frequency-doubling.

7. Calcium ions can easily be photoionised, and ^{43}Ca can be selectively ionised. Although ^{43}Ca has a low abundance (0.14%), being able to ionise it selectively is still enables reliable loading.

Choice of Qubit in $^{43}Ca^+$

In $^{43}Ca^+$ the $|\downarrow\rangle$ and $|\uparrow\rangle$ states lie within the $4S_{1/2}$ hyperfine doublet. With a large (3.2GHz) separation, these two states are well resolved, long-lived, and robust against environmental noise. In experiments to date the $|\downarrow\rangle$ to $|\uparrow\rangle$ transition is driven directly with a microwave source. In future this can also be done with a Raman laser transition, see chapter 4.

To lift degeneracy, an external magnetic field is applied, further dividing the qubit states into their Zeeman sub-states, and providing a quantisation axis. The choice of sub-states to be used as $|\downarrow\rangle$ and $|\uparrow\rangle$ depends on the application, but three options present themselves.

The stretched states of $S_{1/2}$(F=4,M_F =4) and $S_{1/2}$(F=3,M_F =3) are the easiest to prepare, and the use of selection rules simplifies readout and manipulation (the opposite end of the manifold could also be used). In low field, the energy of these states have an approximately linear dependence on the external field. A small field-fluctuation causes a shift in energy level (first-order Zeeman effect). In the case of an unknown perturbation (noise), the driving field will be slightly off-resonance with this shifted energy separation, and this in turn is one path for decoherence – in

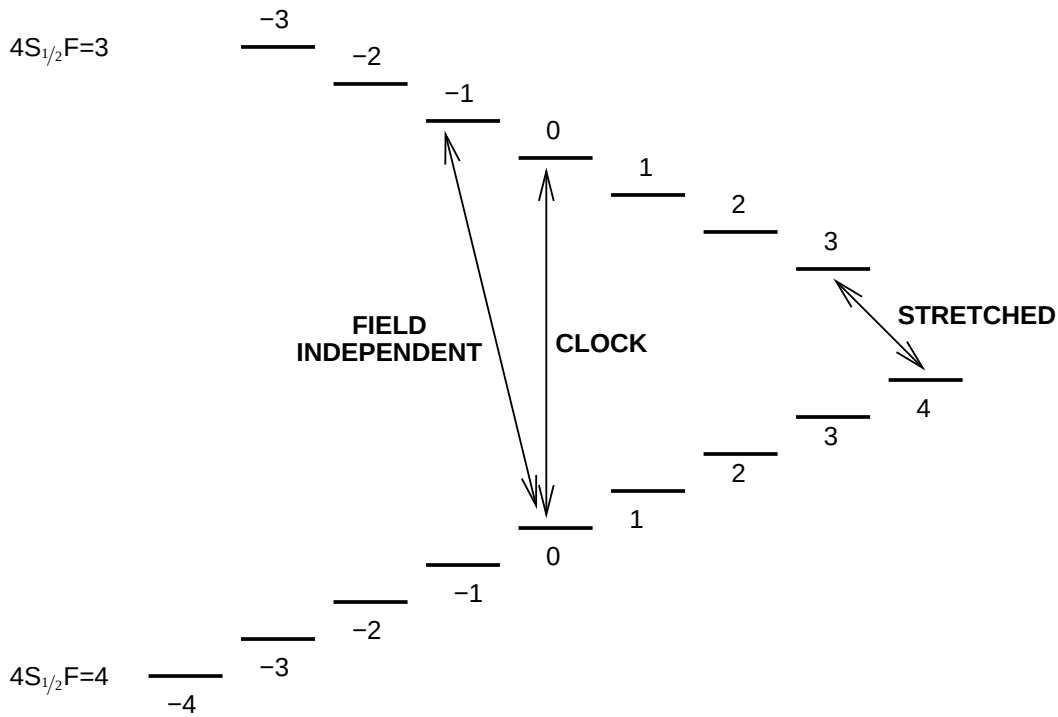

Figure 1.2: Ground state hyperfine transitions in $^{43}Ca^{+}$ showing the M_F states separated by an external magnetic field. The (F=4,M_F =0) ↔ (F=3,M_F =-1) transition is an example of a transition which is field-independent in intermediate field.

the sense of an error caused by an incomplete rotation. For measuring the magnitude of external magnetic fields at the ion, the stretched state is a good choice, and is used in experiments in chapter 6. However, it does not make the best qubit.

Therefore, for computations, it is more desirable to use the "clock" states (F=4,M_F =0) and (F=3,M_F =0). These only have a second-order Zeeman shift at fields close to zero. A non-zero field is still required to lift the degeneracy in M_F so in practice there is still a small linear dependence on external magnetic field. Figure 1.3 shows a plot of energy level versus field strength in the $4S_{1/2}$ states.

A third alternative is to use two states where the transition energy has a zero field dependence. In order to calculate when this occurs, the Breit-Rabi equation is used:

$$E(B) = \frac{-E_{\text{hfs}}}{2(2I+1)} - g_I \mu_N B M_F \pm \frac{E_{\text{hfs}}}{2}\sqrt{1 + \frac{4\alpha B M_F}{2I+1} + \alpha^2 B^2}$$

where E_{hfs} is the hyperfine splitting energy = $h\Delta\nu$, g_I is the Landé g-factor for the nucleus, μ_N

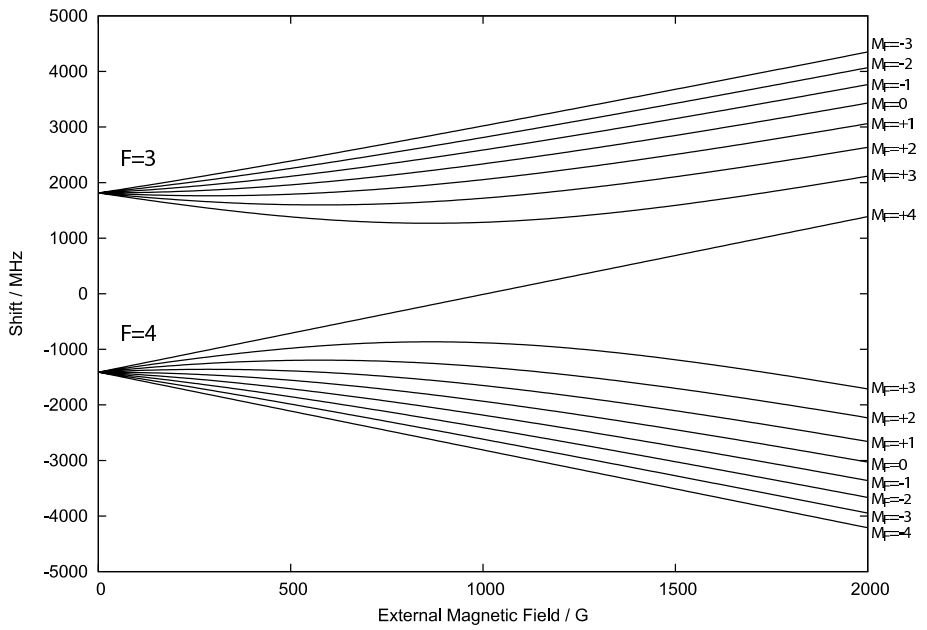

Figure 1.3: Breit-Rabi Plot for $4S_{1/2}$ against magnetic field strength in $^{43}Ca^{+}$ showing Zeeman splitting in increasing field-strength. It can be seen that at certain values of field, the gradient of the energy level for a given pair of sub-states is equal (see also figure 1.4).

is the nuclear magnetron and $\alpha = (g_I\mu_N + g_J\mu_B)/E_{\text{hfs}})$. g_J is the Landé g-factor for the valence electron with total angular momentum $\mathbf{J}$. μ_B is the Bohr magnetron. I=7/2 for $^{43}\text{Ca}^{+}$.

By differentiating the above formula with respect to the field:

$$\frac{\partial E(B)}{\partial B}$$

, the first order dependence on magnetic field can be calculated. By matching this value for the upper and lower states of a transition, a transition that is insensitive to magnetic field (to first order) can be found. Figure 1.4 shows a plot of the first-differential of transition energy w.r.t. B, showing the first occurrences of field-independent transitions in the $4S_{1/2}$ doublet. Only magnetic-dipole allowed transitions (marked with blue circles) are of interest for microwaves. The lowest of these (at non-zero field) occurs at 141.8G, for the (F=4,M_F =-1) to (F=3,M_F =0) transition. This transition is field-independent to first-order. However, these states are more difficult to prepare and read-out than the clock or stretched states, see [37]. No experiments have been done in $^{43}\text{Ca}^{+}$ with these states as yet.

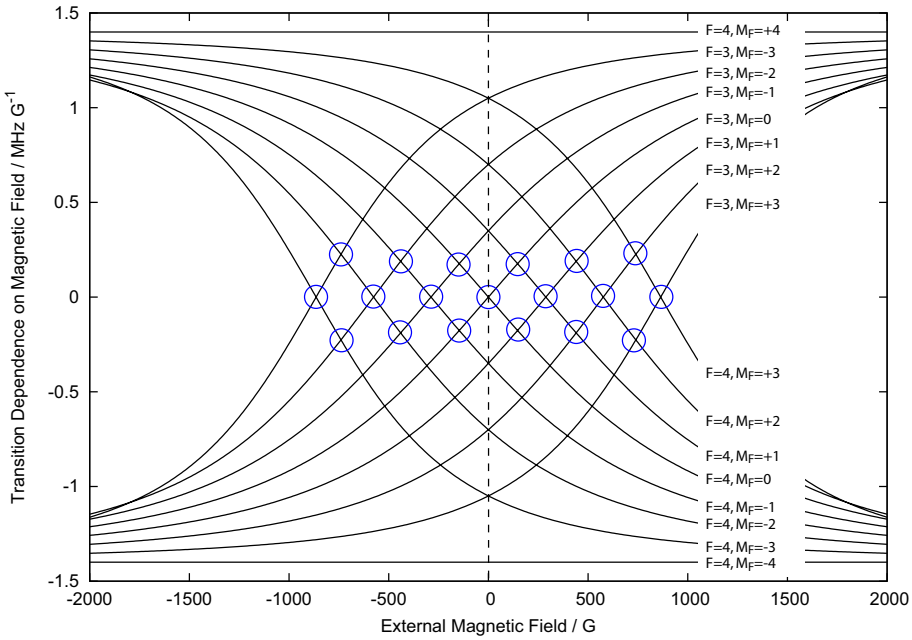

Figure 1.4: Differential of energy $\partial E(B)/\partial B$ against B, for transitions in the $4S_{1/2}$ manifold, in $^{43}Ca^{+}$. Field-independent magnetic dipole-allowed transitions are marked with blue circles. These occur when $\partial E(B)/\partial B$ is equal for both states of the transition. The first of these occurs at zero field, but is not ideal as a field is required to lift degeneracy. Both field directions are shown, where a negative field is taken to be in the opposite direction to σ^{+} polarised light.

1.4 Realisation of a QIP with Trapped Ions

A typical experimental sequence is outlined below:

1.4.1 Loading of Ions

To create ^{43}Ca ions, a beam of calcium neutral atoms is first produced from a heated oven, and then photoionisation is used to generate ions. This is a two step process, with the first transition taking the atom into an excited state ($4s4p^1P_1$), and the second transition removing an electron to the continuum. As the first transition frequency is isotopically-shifted [22], this allows the photoionisation to be selective by isotope. It is also selective by species, so that other impurity ions are not loaded.

These ions are then trapped in a linear Paul trap and cooled.

1.4.2 Doppler Cooling

In the experiments described in this thesis, cooling to the motional ground-state level was not required. It was only necessary to cool to the Doppler limit in order to localise the ion, and ensure that the addressed transitions were not significantly affected by Doppler shifts.

In calcium, the cooling system is a Λ system consisting of the $4S_{1/2}$, $4P_{1/2}$ and $3D_{3/2}$ levels, see figure 1.1. The S and $P_{1/2}$ levels are divided into two hyperfine sub-levels, and the D state is divided into four hyperfine sub-levels. These are further divided into the Zeeman states, giving a total of 64 states. The S-P transition is addressed with a 397nm laser and the D-P repumping transition is addressed with an 866nm laser. Sidebands can be applied using EOMs to address different hyperfine states. The hyperfine splittings are typically around tens to thousands of MHz and are greater than the natural linewidths which are tens of MHz. The laser linewidth of the lasers used was typically < 5MHz.

To cool, it is necessary to drive one pair of the four transitions at 397nm, whilst repumping from the $3D_{3/2}$ manifold. The merits of which transitions to drive have been investigated [37], and it was concluded that the best results could be obtained by using a laser tuned to the $4S_{1/2}(F=3) \rightarrow 4P_{1/2}(F=4)$ transition. This is achieved by using an EOM to generate a sideband at the $4S_{1/2}(F=4) \rightarrow 4P_{1/2}(F=4)$ frequency, 3226MHz to the red. (A symmetrical sideband is produced to the blue, but not used). This is shown in figure 1.5. The natural linewidth is 23MHz which is greater than the Zeeman splitting (0.6MHz in an external field of 1.8G), so all M_F states are driven. Repumping from the 3D level is achieved using an 866nm laser.

Doppler cooling occurs as photons are absorbed from the laser beam, when moving against the beam direction, and the ion absorbs their momentum. The photon is re-emitted in a random direction, causing a net loss of momentum from the ion. The 397nm Doppler beam is at an angle to both the z-axis and the radial axes of the trap. That way cooling can be achieved in all three directions. For maximum cooling efficiency, the Doppler beam is tuned close to the half-maximum point of resonance (for the ion at rest). The 866nm laser is tuned slightly to the blue, to prevent two-photon dark-resonances.

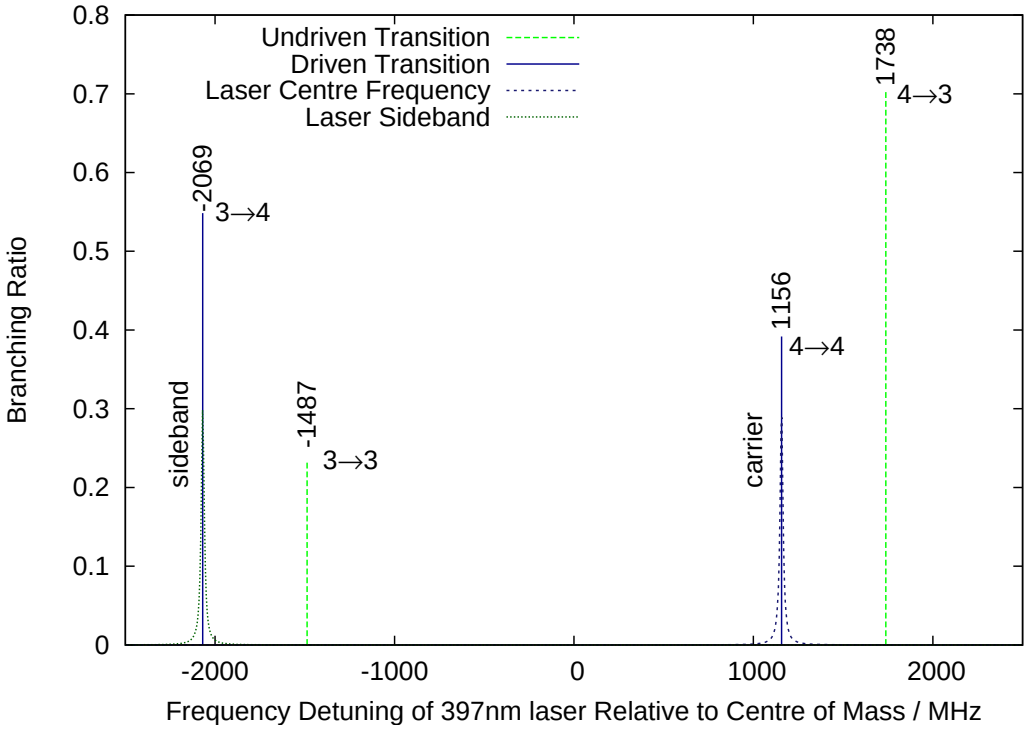

Figure 1.5: Transition energies in $^{43}Ca^{+}$ for the 4S to 4P levels, showing the 397nm beam driving the 4-4 transition and the sideband which drives the 3-4 transition.

1.4.3 State Preparation

DiVincenzo's second criterion requires "the ability to initialise the state of the qubits to a simple fiducial state." To be able to do this with high reliability is essential, as any error in initial state preparation will affect every subsequent qubit operation.

A theoretical study of state preparation has been done by Webster [23], with further work done by Szwer [37]. State preparation is closely connected to cooling, as the S-P transition is involved for both.

As the qubit memory experiments in chapter 7 use the clock transition $|\downarrow\rangle$=(F=4,M_F =0) $\leftrightarrow$ $|\uparrow\rangle$=(F=3,M_F =0), the preparation of this state is described.

By switching off the EOM which generates the sideband for the $S_{1/2}(F=3) \rightarrow P_{1/2}(F=4)$ transition during cooling, it is possible to pump to the $|\uparrow\rangle$ clock state, 4S(F=3,M_F =0), using the same 397nm cooling laser. By allowing the ion to spontaneously decay into the seven M_F sub-states, the clock state is reached statistically $\sim 1/7$ of the time. A rate-equation simulation taking into account Clebsch-Gordan coefficients for different decay routes shows that the probability of

preparing the (F=3,M_F =0) clock state is indeed close to 14%.

It would be possible to prepare just the clock state, by exploiting the selection rule that forbids $M_F = 0 \rightarrow 0$ for $\Delta F = 0$. An applied π polarised beam drives transitions with $\Delta F = 0$, and therefore the lower M_F state goes dark. Experimentally however, this was not possible, as a π polarised beam was not available. Instead, in the experiments described here, the $S_{1/2}$(F=3, M_F =0) clock state is prepared using the statistical method, and the unwanted $6/7$ of the results are discarded.

1.4.4 Qubit Manipulations

Once the ion is cooled and prepared, coherent operations or gate operations can be performed with one or several ions. Entanglement is achieved by using the ion's motional states to couple neighbouring ions. This can be thought of like a bus in a conventional computer, used for conveying information between separate qubits. In the original proposal by Cirac and Zoller it was necessary to first cool the ions to the ground state. However, Sørenson and Mølmer showed that a gate operation could be performed between two ions that were not cooled to the ground state (a so called "warm" gate). In 2005 Oxford demonstrated a form of motional phase gate between two $^{40}Ca^+$ ions, called the "wobble gate" [21].

Other ion-manipulations can be carried out, including very high resolution single-ion spectroscopy and measurements such as the decoherence measurements presented in chapters 6 and 7.

1.4.5 Readout

The fifth DiVincenzo criterion states that a QIP must have "a qubit-specific measurement capability". In $^{43}Ca^+$ no closed cycling transition exists involving the $|\downarrow\rangle$ or $|\uparrow\rangle$ states, such as is used at Boulder [38]. Instead population escapes into the two D-level manifolds.

Assuming that population didn't transfer to the D-states or that it could be pumped back, the off-resonant excitation from the $|\downarrow\rangle$ state to the $|\uparrow\rangle$ state on the S $\rightarrow$ P transition would give a small error rate of:

$$\varepsilon \approx \left(\frac{\Gamma}{\Delta}\right)^2 \approx \left(\frac{22}{3226}\right)^2 = 4.6 \times 10^{-5}$$

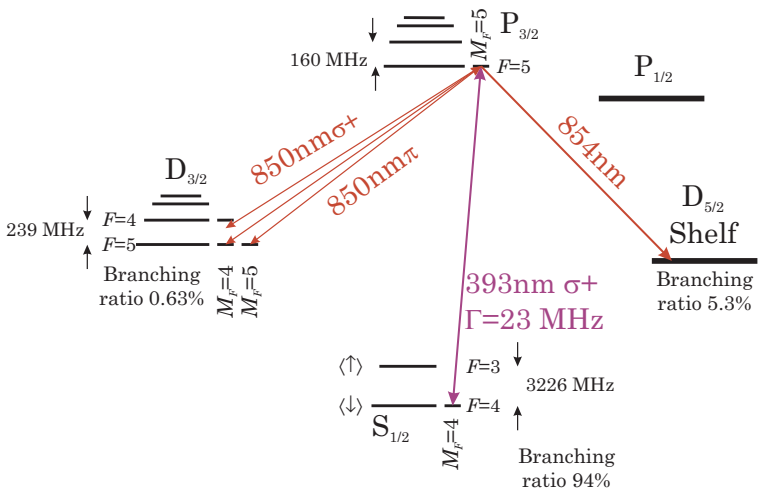

Figure 1.6: Transitions in $^{43}Ca^{+}$ used in read-out. Shown are the relative hyperfine splittings and the long-lived D-level used as a "shelf". The selection of M_F states using σ^{+} light is also shown.

per cycling transition. However, to collect enough fluorescence into the solid-angle available, about 10^4 photons need to be scattered. As soon as an off-resonant transition occurs, a $|\uparrow\rangle$ could become a $|\downarrow\rangle$ and then be trapped in the cycling transition, giving an error rate of 10% or more.

To reduce the number of driven transitions required the D-states can be used advantageously. The $|\downarrow\rangle$ (F=4) state is first "shelved". i.e. the population is transferred to the $D_{5/2}$ with a long-lived lifetime of 1168(7)ms. The $|\uparrow\rangle$ (F=3) state is not shelved, due to being 3.2GHz off-resonance. When the S $\rightarrow$ P cycling transition is driven, the shelf state is a long way from resonance; thus, in principle, errors can be reduced to the 1 in 10^3 level.

Fluorescence corresponds to the ion being in the $|\uparrow\rangle$ (F=3), and fluorescence at the background level corresponds to it being in the $|\downarrow\rangle$ (F=4) state. In order to measure the qubit state with good statistics the experiment is typically repeated 500 times. To differentiate from background light or PMT dark-counts, at least 10 photons need to be collected. For the solid-angle over which photons are collected (0.4 srad – see [39]), and given the quantum efficiency of the PMT (15%), this requires a count time of at least 2ms. As photon counting is a discrete process, the measurement obeys Poissonian statistics. The histogram for a $|\uparrow\rangle$ (F=3) (fluorescence) can be compared to the histogram for a $|\downarrow\rangle$ (F=4) (background count), with a threshold set such that the two qubit states are reliably distinguishable. The net efficiency is about 95%, limited by the decay to the $D_{3/2}$ level.

To obtain typical counting statistics, the ion was prepared, shelved and measured. Results are

shown in figure 1.7. The count-rate differs slightly from the expected Poissonian distribution for the shelved histogram, due to stochastic noise sources. Possible sources of these errors include: dark-counts due to electrical noise or cosmic rays, the shelf spontaneously decaying, and possible errors in state preparation. Each of these has been investigated separately, see [40].

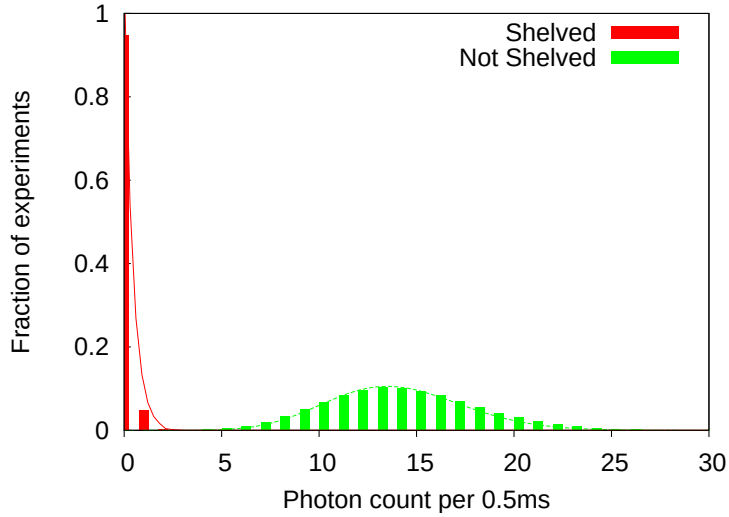

Figure 1.7: Typical photon counts using $^{40}Ca^{+}$ after preparation and readout. The first histogram shows photon count after the ion is shelved. The second histogram shows photon counts for the ion, not shelved. The threshold is set at a count of 2.5. The count-time in $^{40}Ca^{+}$ is shorter due to the higher fluorescence than in $^{43}Ca^{+}$.

1.4.6 Improved Readout Efficiency

The experiments described here do not use a repumping scheme to improve readout efficiency. However, such a scheme would be advantageous and is described below.

Because two D-states exist, population can fall from the $P_{3/2}$ state to the $D_{3/2}$ as well as the $D_{5/2}$ state. The $D_{3/2}$ level is addressed by the cooling and detection lasers, therefore any population here would also fluoresce, creating a readout error. The branching ratio for $(P_{3/2} \rightarrow D_{3/2}) : (P_{3/2} \rightarrow D_{5/2})$ gives approximately a 10% chance of population falling to the $D_{3/2}$ state, giving an unacceptably high error in readout (state discrimination) for actual computation, though this level is tolerable for the diagnostic experiments reported elsewhere in this thesis. The P states are split into their hyperfine components, and further split by the magnetic field into the Zeeman sub-states. These can be used to selectively repump out of the $D_{3/2}$ state, via the $P_{3/2}(F = 5, M_F = +5)$ so that the

only allowed transition is to the $|\downarrow\rangle$ state $S_{1/2}(F = 4, M_F = +4)$ or to the $D_{5/2}$ shelf. (Assuming in this instance the stretched states are being used as the qubit states.) This prevents a $|\downarrow\rangle$ becoming a $|\uparrow\rangle$.

It is possible to make use of selection rules to simplify the number of transitions that need to be addressed. First, it is simplest to consider the stretched state. In this state the $|\downarrow\rangle$ state is $4S_{1/2}$(F=4,M_F =4) and $4S_{1/2}$(F=3,M_F =3) is the $|\uparrow\rangle$ state. i.e. the end of the manifold. These are highlighted in figure 1.6. A 393nm σ^+ pulse will take the population to the $4P_{3/2}$(F=5,M_F =5) state. It will not address the other M_F states or lower F states due to conservation of angular momentum, which requires $\Delta F = \Delta M_F = +1$ for σ^+ light. From this state, i.e. the stretched state of the $P_{3/2}$ manifold, the ion can decay to either the $D_{3/2}$ state or the $D_{5/2}$ state, with the probability of going to the $D_{3/2}$ state being 10.6%. As angular momentum can only change by 0, or ± 1, the $D_{3/2}$(F=5,M_F =5) state will be reached with no change of angular momentum, and the $D_{3/2}$(F=5,M_F =4) or $D_{3/2}$(F=4,M_F =4) states will be reached by a change of -1 in angular momentum. These can be repumped using π light, and σ^+ light at 850nm. However, it is important that the pulses are applied in sequence: σ^+ then π. Otherwise the repumping beams could start to walk population down the P-state manifold.

The other 89.4% of the time the population falls to the $D_{5/2}$ state where it is trapped. The lifetime is long compared to the time taken to perform the fluorescent transitions and repumping ($1/\Gamma$ = 1.2s compared with 1ms for 100 fluorescent transitions) so spontaneous decay from this shelf contributes an error on the level of 1 in 10^3 [37]. By optimising the pulse time in the sequence, Szwer [37] has shown that a theoretical average readout efficiency of 99.96% can be achieved in ten repeats of the pulse sequence, in a total time of 100 µs. In recent measurements, a direction for a π-polarised beam was not available, but this scheme was partially implemented using 850nm σ^+ polarised with a small admixture of π-light, and a net efficiency of 99.7(3)% was achieved. Although readout efficiency only needs to be above a known threshold for state-discrimination, it is important that readout does not significantly increase the overall error-rate. A low readout error rate allows reliable diagnostics and corrections of errors elsewhere in the system. This is a great improvement over a $^{40}Ca^+$ ground-level qubit of $S_{1/2}(M_F = \pm 1/2)$ where readout levels of about 90% were obtained [23]. In most of the experiments presented in this thesis, no

850nm laser was used.

1.4.7 Deshelving

Once readout has been performed, the ion is "deshelved" i.e. population is transferred out of the $D_{5/2}$ shelf level back to the ground state. The ion is then cooled and re-prepared. To deshelve, the $D_{5/2} \rightarrow P_{3/2}$ transition is driven by the 854nm laser. Typically the whole sequence of cool-operation-readout-deshelve is repeated 500–1000 times. The experimental parameters can then be varied (e.g. detuning for a spectroscopic-type measurement) before another set of repetitions is carried out.

1.5 Overview of Thesis

This thesis is organised in the following way:

In chapter 2 the experimental systems including the trap and the lasers are introduced.

In chapter 3, the design, build and testing of a new piece of equipment, the laser control unit, is described. This allows sequences of pulses to be produced that allow the experiments in chapters 6 and 7 to be carried out.

In chapter 4 a laser system based on optical injection is described that would allow Raman transitions to be driven in the qubit states.

In chapter 5 the circuits and setup for PDH locking of lasers is described.

In chapter 6 measurements of the background (lab) magnetic field are made and the effect on the ion is characterised.

Chapter 7 presents results from experiments using $^{43}Ca^{+}$ as a memory qubit. Coherence times are measured for both the Ramsey technique and the spin-echo technique.

Chapter 8 concludes.

Chapter 2

Experimental Setup

2.1 Overview

This section describes the apparatus for experiments with $^{43}Ca^{+}$ presented in this thesis. It also looks ahead at the development of the apparatus which is presented in subsequent chapters.

The apparatus used at Oxford consists of the following: a linear Paul trap to hold the ions, microwaves and diode lasers to manipulate the ions' state, external magnetic fields to lift degeneracy, a CCD or a PMT to collect fluorescence, and a PC to control the experiment and make measurements.

2.2 Paul Traps

The ions are held in a Paul (radio frequency) trap. Currently all work to date, including the results presented in this thesis, has been done in a linear Paul trap with electrodes on the millimetre scale. Further work is being done in the group developing smaller traps, and traps with segmented electrodes. A micro-fabricated trap manufactured by Sandia National Laboratories using silicon based technology with tungsten plated electrodes has recently been used to successfully trap $^{40}Ca^{+}$ ions.

The advantage of the current trap is that the large scale allows easy access with laser beams, and that the heating rate of the the ions from the ground motional state is very low [41]. It has been possible to load and experiment with both $^{40}Ca^{+}$ and $^{43}Ca^{+}$ in the current trap, using an

isotope selective photoionisation process [22].

The current trap does not allow ions to be moved about, nor is it possible to address ions selectively. The latter would be possible with very careful control of a tightly-focused laser beam using a steering device. However, this scheme does not scale well to large numbers of ions and it has not been implemented at Oxford.

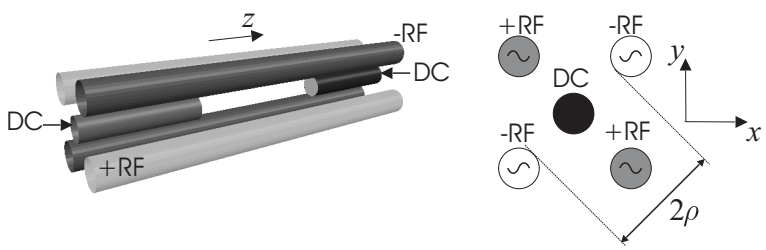

Figure 2.1: A schematic of the current ion trap. This shows the x, y and z directions and the connections of the electrodes to both RF and DC voltages. The +RF and -RF are 180° out of phase. Not shown are DC compensation electrodes which sit further out from the trapping electrodes and allow for compensation of stray fields.

A linear RF trap works on a cylindrical geometry (see figure 2.1). It is conventional to define the trap axis as z, and the radial direction in the x-y plane. Ions are confined in the z direction with two end-cap electrodes. With equal positive voltages on each end, a harmonic potential is formed, along which the positive ions will form a string due to their mutual Coulomb repulsion. The energy levels of the harmonic potential are quantised, and it is along this direction that cooling and coherent motion is directed. In the x-y plane, the ion cannot be held with further sets of DC electrodes. The reason for this is that the Laplace equation, $\nabla^2 V = 0$, must be satisfied and therefore:

$$\frac{\partial^2 V}{\partial^2 x}, \frac{\partial^2 V}{\partial^2 y}, \frac{\partial^2 V}{\partial^2 z} > 0$$

cannot all be satisfied (Earnshaw's Theorem), as is required for trapping. Instead, an RF voltage is applied between opposite pairs of electrodes. This creates a saddle potential which varies in sign, at the frequency of the RF. This in turn creates a pseudo-potential from which the ions cannot escape. However, the ions do undergo micromotion (driven oscillations at the RF frequency) unless the ions are accurately positioned along the z-axis of the trap. For further details of the exact criteria for trapping, and the form of these potentials see [39].

The original Paul trap has been developed into more complex multi-segmented traps [41].

By rearranging the geometry of the electrodes, it is possible to create multiple wells along the z-direction, each holding an ion in a separate potential well. Control of these electrode voltages allows movement of the ion such that ions can be separated (i.e. they no longer interact via the Coulomb interaction) and recombined, or moved to storage areas. Great care must be taken not to impart sufficient momentum to the ion for it to escape [42]. Recent experiments have succeeded in the movement of an ion between two wells without loss of the ion.

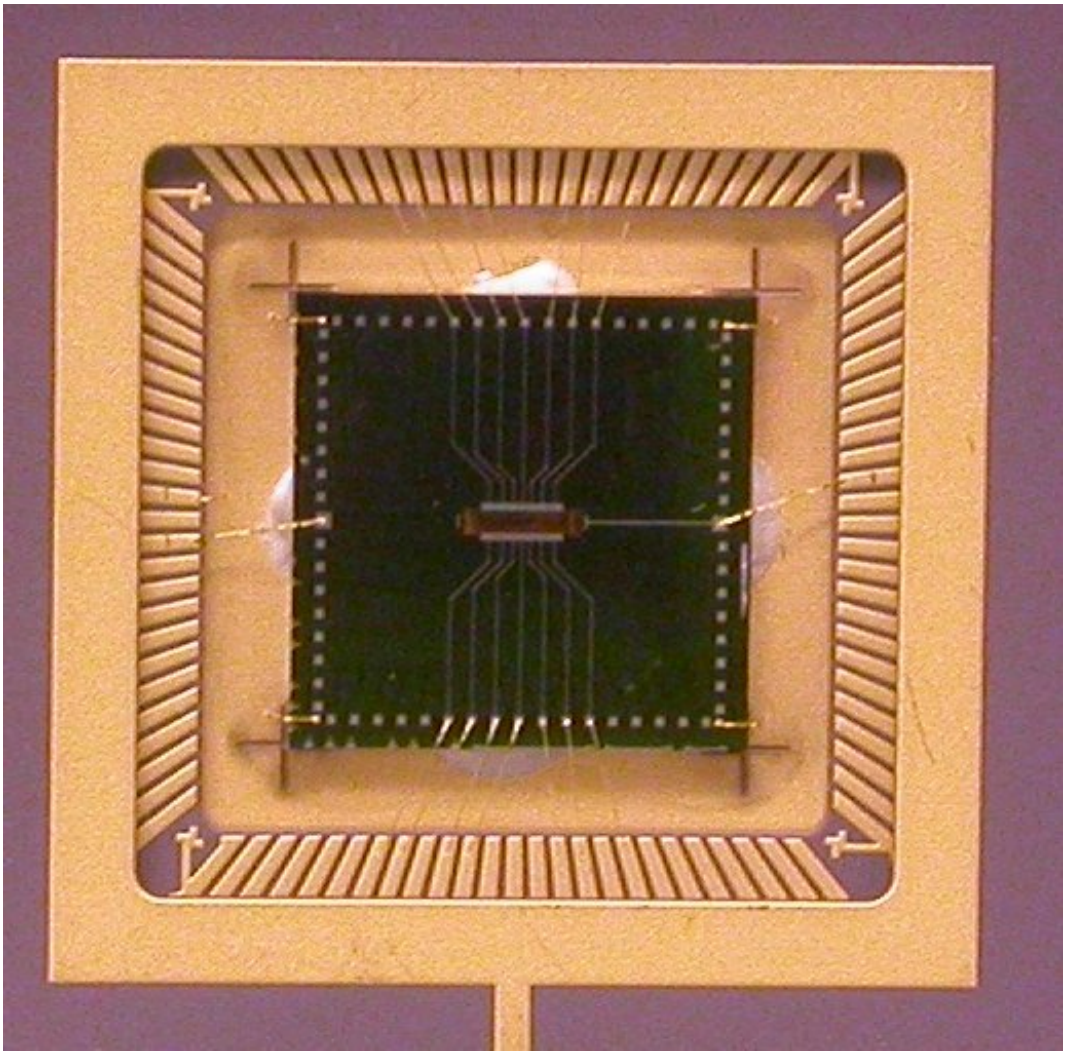

Figure 2.2: Photograph of Sandia trap showing multiple electrodes. It is mounted on a standard chip carrier. The trapping regions is about 2mm long and 0.4mm wide.

Two new traps of different types are being tried at Oxford. One is a millimetre-scale trap built using machine-techniques in conjunction with Liverpool University and has a set of three potential wells. This will allow ion-moving techniques to be investigated. The second trap is manufactured

by Sandia National Laboratories from silicon with gold plating and tungsten plated electrodes. It has multiple trapping areas with electrodes on the one-tenth millimetre scale.

2.3 Laser Systems

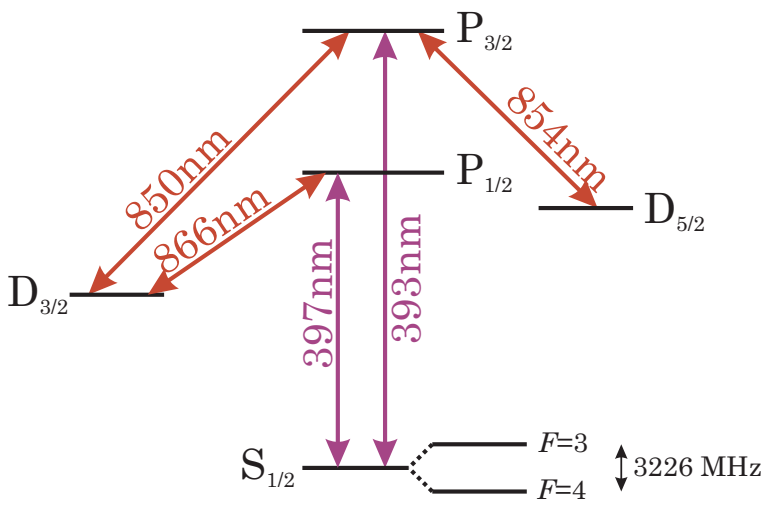

Figure 2.3: Relevant transitions in $^{43}Ca^{+}$ showing the wavelengths of the driving lasers and the frequency of the driving microwaves.

The above figure 2.3 shows the relevant transitions in $^{43}Ca^{+}$. The experimental sequence has been introduced in chapter 1. For the calcium 43 experiments, the lasers that are required are summarised in table 2.1.

Purpose	**λ / nm**	**Notes**
Ionisation of neutral Ca	423 389	Not locked Not locked nor grating-stabilised
Doppler cooling and state-preparation	397 866	Locked using PDH (see chapter 5) Repumping from $D_{3/2}$. Locked.
Coherent state manipulation	396 9.3cm	Raman transition system (see chapter 4) 3.2GHz Microwaves
Read-out:shelving	393 850	Locked using PDH Locked using side-of-fringe lock
Read-out:deshelving	854	Not locked

Table 2.1: Lasers required for $^{43}Ca^{+}$ experiments

The three infrared lasers require 160MHz sidebands in order to efficiently repump from the

complete hyperfine manifolds of the D-levels. It is a great convenience that these sidebands can also be used to lock the laser to a precision reference Fabry-Perot etalon. In future this will be done using the Pound-Drever Hall technique, covered in chapter 5. In the apparatus used for the experiments presented in this thesis, sidebands were not added to the infrared lasers, but reasonable repumping was achieved with sufficiently high laser intensity, see [23]. A Raman beam was not used to manipulate the qubit state – instead the transition was driven directly using microwaves.

2.4 Beam Lines

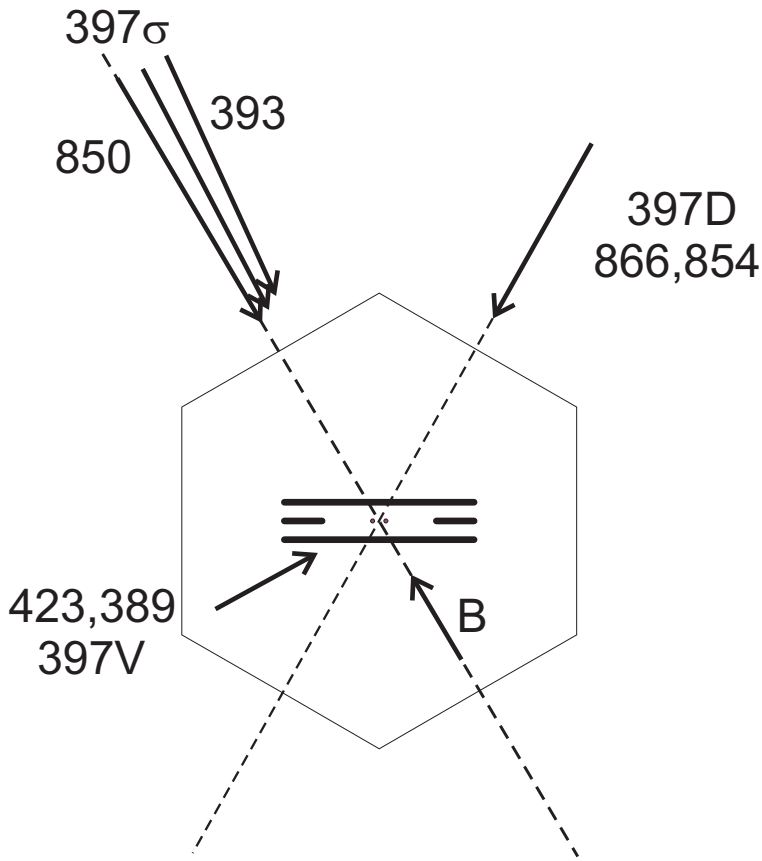

Figure 2.4: Beam directions going into the experimental apparatus. The vertical beams (389, 397V and 423) are actually at 30 ° to the vertical, as shown on the diagram for the imaging system figure 2.16. The 397σ, 393 and 850 beams are circularly polarised. The 397D, 866 and 854 beams are linearly polarised perpendicular to the direction of the magnetic field.

Figure 2.4 shows the direction of each beam going into the apparatus. The vertical beams are

shown more clearly on figure 2.16. Each beam line into the vacuum can is considered in turn.

2.4.1 Photo-ionisation Beams

To ionise neutral calcium, two beams are required with wavelengths of 423nm and 389nm. The 423nm beam excites the atom into the 4s4p$^1P_1$ state and the 389nm beam removes the electron to the continuum. The intensity of the 423nm beam is quite low, ≈10μW. To efficiently load $^{43}Ca^+$, the 389nm needs to run at maximum power, ≈100 μW. No specific polarisation is required. Both beams enter the trap at 30° to the vertical direction via a fibre. Both beams are produced using diode lasers. By tuning the 423nm laser to the correct frequency, isotopic discrimination can be achieved with high precision [22].

2.4.2 397nm Laser

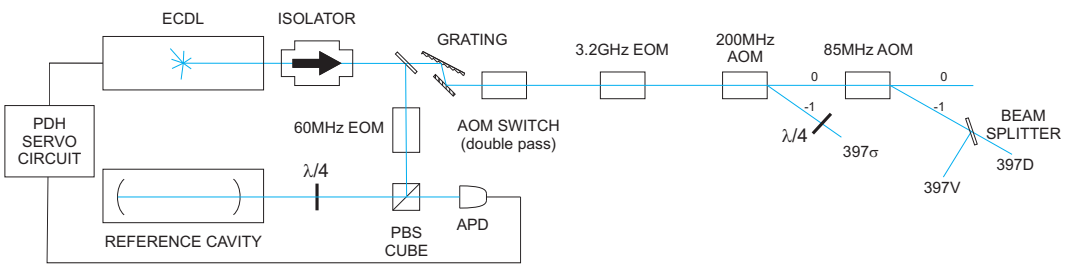

Figure 2.5: Diagram showing the arrangement of the 397nm beam-path. A grating is used to remove any 393nm light present in the ASE. Sidebands are added using an EOM. Multiple paths are created for different uses.

Figure 2.5 shows a schematic representation of the beam-path. The 397nm laser is an External Cavity Diode Laser (ECDL or grating-stabilised). It is PDH-locked to an NPL reference cavity which is an etalon made from zero-expansitivity glass, temperature stabilised and held in a high-vaccuum. The beam is then reflected from a diffraction grating to reject the ASE light which would otherwise excite the 393nm transition. Even a small amount of 393nm light would cause the ion to shelve occasionally and increase the error rate. This grating is 55% effecient at the second-order ([63]).

After this, a double-pass AOM followed by a single-pass AOM is used to switch the beam off and on with good extinction. An EOM is placed in-between to provide 3.2GHz sidebands.

The AOM used is 26% more effecient with one polarisation direction than the other, so a double-pass technique which changes polarisation would not be effecient. Instead the returning beam is separated from the incoming beam in the vertical direction by using a corner-cube see [64] for more details. The AOM allows faster switching times than a mechanical shutter without creating vibrations that may disturb the laser. Each pass achieves an extinction of 10^{-3}, so after three passes only 1 part in 10^9 of the light reaches the ion.

The light is then further divided into two beams to produce a Doppler cooling (397D) beam, and an optical pumping (397σ) beam.

397D Cooling Beam

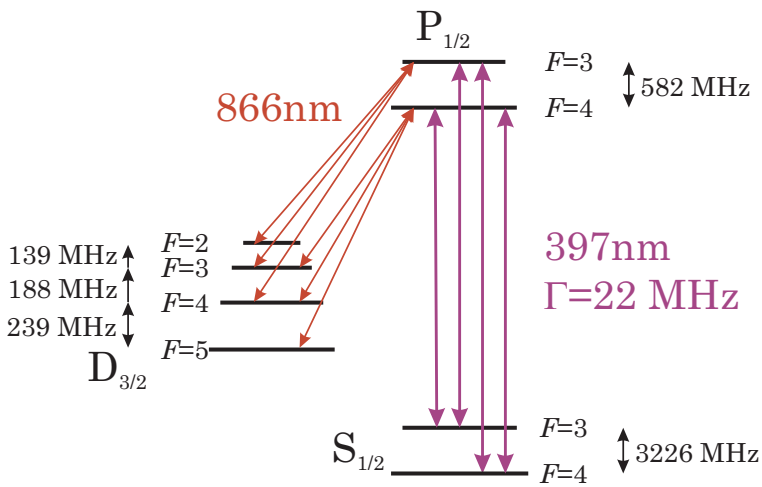

Figure 2.6: Relevant transitions used for Doppler-cooling $^{43}Ca^{+}$.

The 397D beam is used for Doppler cooling. The cooling system is described in section 1.4.2. Initially, two separate laser beams were used to drive the two necessary transitions. In later experiments, an EOM was used to create a sideband at the $4S_{1/2}(F=4) \rightarrow 4P_{1/2}(F=4)$ frequency, 3226MHz to the red. (A symmetrical sideband is produced to the blue, but not used).

This beam-line is taken from the zeroth order of the 397σ AOM and then passed through an 85MHz AOM. This is used as a switch, however the exact frequency was choosen for other experiments, see [23]. The beam is vertically polarised and enters the vacuum can at an angle to both the magnetic field direction and the z-axis of the trap. Therefore it can cool in both the axial and radial directions.

397σ State-preparation Beam

This laser beam is σ^+ polarised, and is introduced to the ion along the magnetic field quantisation direction. This allows transitions to be made which increase the spin quantum number by 1, in order to do state preparation of $S_{1/2}(F = 4, M_F = +4)$ by optical pumping. The required frequency shift for driving the $3 \leftrightarrow 4$ and $4 \leftrightarrow 4$ transitions is present due to the EOM in the beam-path, which produces sidebands at $\pm$ 3.2GHz.

397V Compensation Beam

A further cooling beam may be inserted vertically, along the direction of the ionisation beams. This is a 397nm beam taken from the Doppler beam above and is used to cool and compensate for micromotion in the vertical direction of the trap. Compensation is achieved by minimising correlations between the arrival time of photons and the 6MHz RF drive frequency of the trap. Micromotion causes Doppler shifting of the ion's resonance frequency in sync with the RF drive, so when the 397nm beam is red-detuned on the side of the 397 resonance the fluorescence signal goes up and down in sync with the RF, and the photons arrive in sync with the RF: see [64] for further details.

2.4.3 866nm Repumper Beam

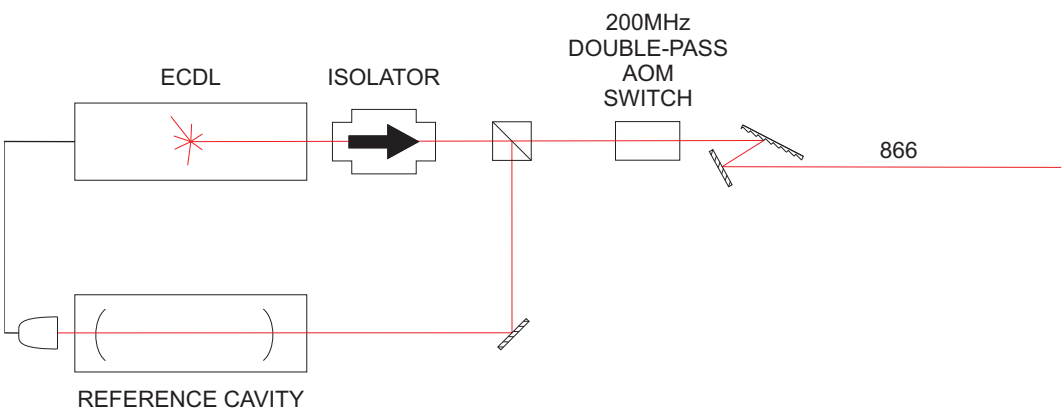

Figure 2.7: Beam path for 866nm laser

The 866nm laser is a 50mW ECDL and is used to repump population from the $D_{3/2}$ level back to the $P_{3/2}$ level during cooling and detection. It is locked to a tunable cavity with a finesse of $\approx$100

and a FSR of 1600MHz (see chapter 5 for further details) using a side-of-fringe lock. Presently there is enough intensity in the beam to drive all the transitions from the $D_{3/2}$ hyperfine manifold, but the addition of 160MHz sidebands would enable the cavity to be PDH-locked and ensure more efficient repumping of all F sub-levels. The exact frequency for maximum effeciency has been calculated by Webster [23]. The beam is switched off and on with a double-pass AOM setup. The beam is reflected from a grating to enable filtering of any 854nm light that would deshelve the ion. The beam is superimposed on to the 854nm beam using a PBS cube. It enters the trap vertically polarised.

2.4.4 Raman Beams

These are not used in the experiments described in this thesis, however the design and implementation of a Raman system for $^{43}Ca^{+}$ is covered in chapter 4.

To drive the $|\downarrow\rangle$ to $|\uparrow\rangle$ transition in $^{43}Ca^{+}$, it is possible to use microwaves, see below. However, this does not allow momentum to be imparted to the ion, or selectivity of position or polarisation. Instead, it is possible to use a two-photon Raman transition to drive the $|\downarrow\rangle$ to $|\uparrow\rangle$ transition. This allows laser beams to be used, which can cool and impart momentum to the ion, for instance in a motional gate. The Raman beams are also useful in temperature diagnostics [23].

To impart momentum to the ion in the z-direction, it is necessary for the difference in the wave-vector between the two Raman beams, $\Delta k = k_1 - k_2$, to have a component along the trap axis.

2.4.5 393nm Laser

393σ+ Shelving Beam

The 393nm beam enters at 3° to the direction of the 397σ. It is used for state-selective shelving, as described above. It is circularly polarised and propagates along the direction of the magnetic field (where the direction can be chosen with the steering coils) such that it drives the $S_{1/2}(F = 4, M_F = +4) \rightarrow P_{1/2}(F = 5, M_F = +5)$ transition. The laser is locked with a Pound-Drever Hall lock, as described in chapter 5.

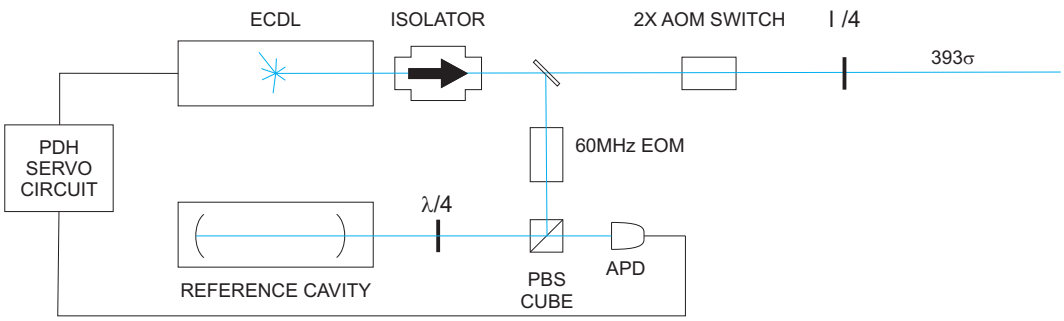

Figure 2.8: Diagram showing the arrangement of the 393nm beam-path including double-pass AOM set-up which switches off the beam, and the PDH lock.

2.4.6 850nm Readout Deshelving Beam

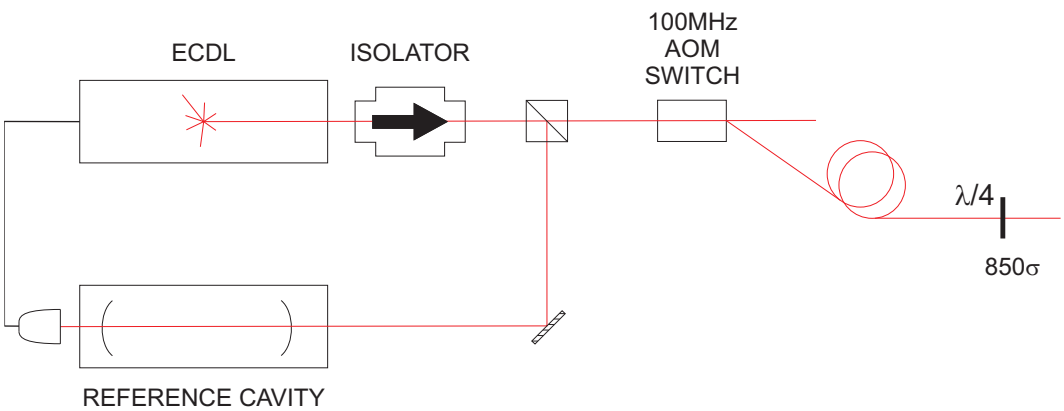

Figure 2.9: Beam path for 850nm laser

The 850nm laser provides repumping from the $D_{3/2}$ levels. It is a 150mW ECDL locked to a cavity using a side-of-fringe lock. By choosing the correct frequency and σ^+ polarisation, pumping occurs from a chosen M_F level. It is switched with a single-pass AOM and transmitted through a fibre before being recollimated and circularly polarised. It then enters the trap, parallel to the magnetic field direction.

2.4.7 854nm Readout Deshelving Beam

To deshelve from the $D_{5/2}$ shelf level after detection, an 854nm laser is used. This consists of a 50mW (ECDL). It is not locked to a cavity. Its width is greater than the spread of the $D_{5/2}$ sub-levels, therefore these transitions are all driven. The addition of 160MHz sidebands would enable this manifold to be deshelved more efficiently and also enable the use of a PDH lock to stabilise

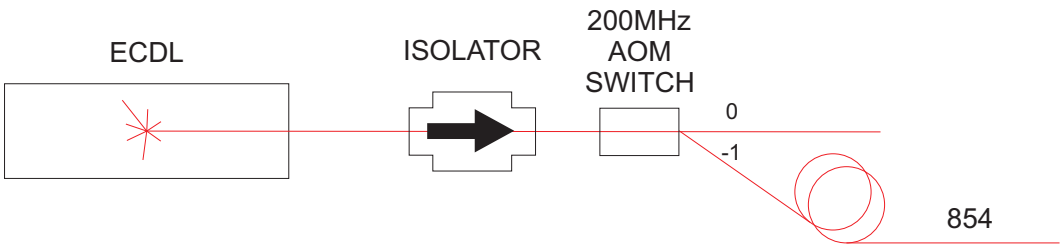

Figure 2.10: Beam path for 854nm laser

the laser from long-term drift. This has not yet been implemented.

2.4.8 Laser Diagnostics

The beam from each laser is first passed through a glass-blank and the reflected beam is sent down a fibre to a fibre-switcher. The output of the switcher goes into an Optical Spectrum Analyser (OSA) and a wavemeter [44].

2.5 Magnetic Field Coils

It is necessary to apply an external magnetic field both to remove Zeeman degeneracy and to prevent optical pumping into dark states.

For this purpose, three pairs of coils are placed around the experiment, as shown in figure 2.11. Greater detail is given in [64].

1. "Parallel field" direction: two coils provide the B-field of 1.8G along the quantization axis i.e. along the 397σ beam direction.

2. "Perpendicular field" direction: two other coils provide a horizontal trimming field perpendicular to the 397σ beam, in the horizontal plane.

3. "Vertical field" direction: two further "high-field" coils provide a vertical trimming field perpendicular to the 397σ beam, in the vertical direction.

The field is steered along the 397σ beam direction such that fluorescence due to optical pumping is minimised.

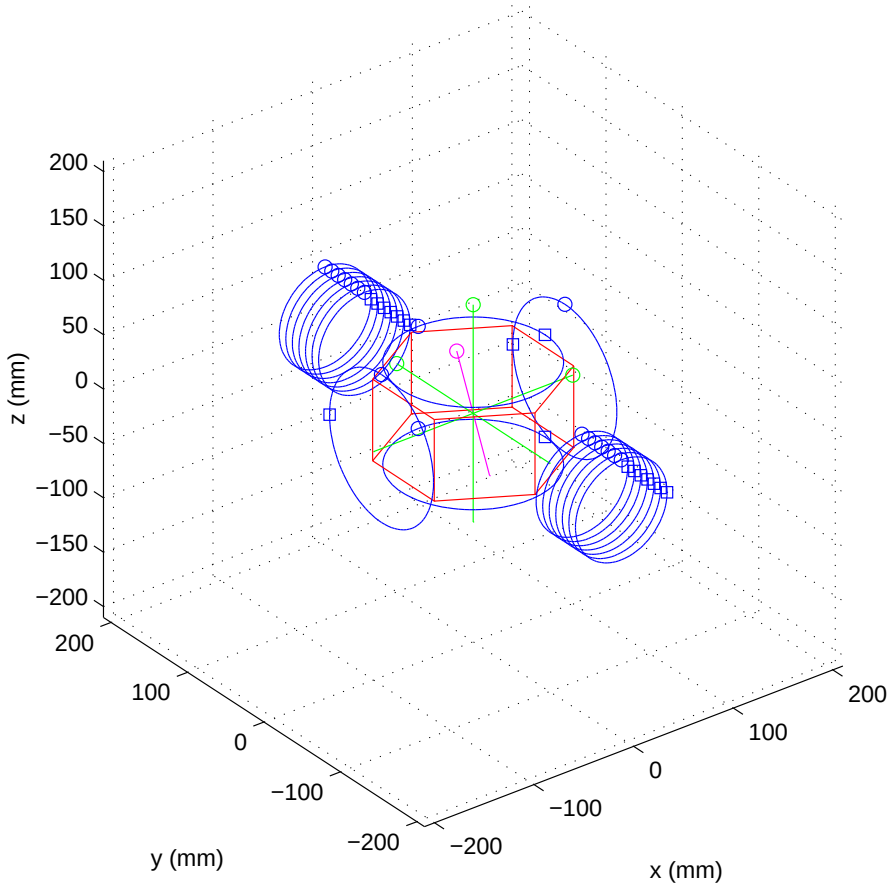

Figure 2.11: Coil-arrangement in the experiment. The x-axis corresponds to the 397σ beam direction. The vacuum can is shown in red.

2.6 Microwave Antenna

In order to drive $|\downarrow\rangle$ to $|\uparrow\rangle$ transitions, it is necessary to excite the ion with 3.2GHz RF radiation, or to use a Raman transition as described above. In these experiments the transition was driven directly using microwaves.

Initially to couple the radiation to the ion an antenna was built. This was driven using an RF amplifier fed from an Agilent E8247C synthesiser, and the amplifier used was an AM38A, on evaluation from Microwave Amps. The AM38 had a gain of 40dB and a maximum output power of 20W over the frequency range 1–4GHz.

Several antennas were tried, see figure 2.12. To compare different antennas, a receiving antenna was used to make measurements on a RF spectrum analyser (Agilent E4405B). In total, four antennas were built:

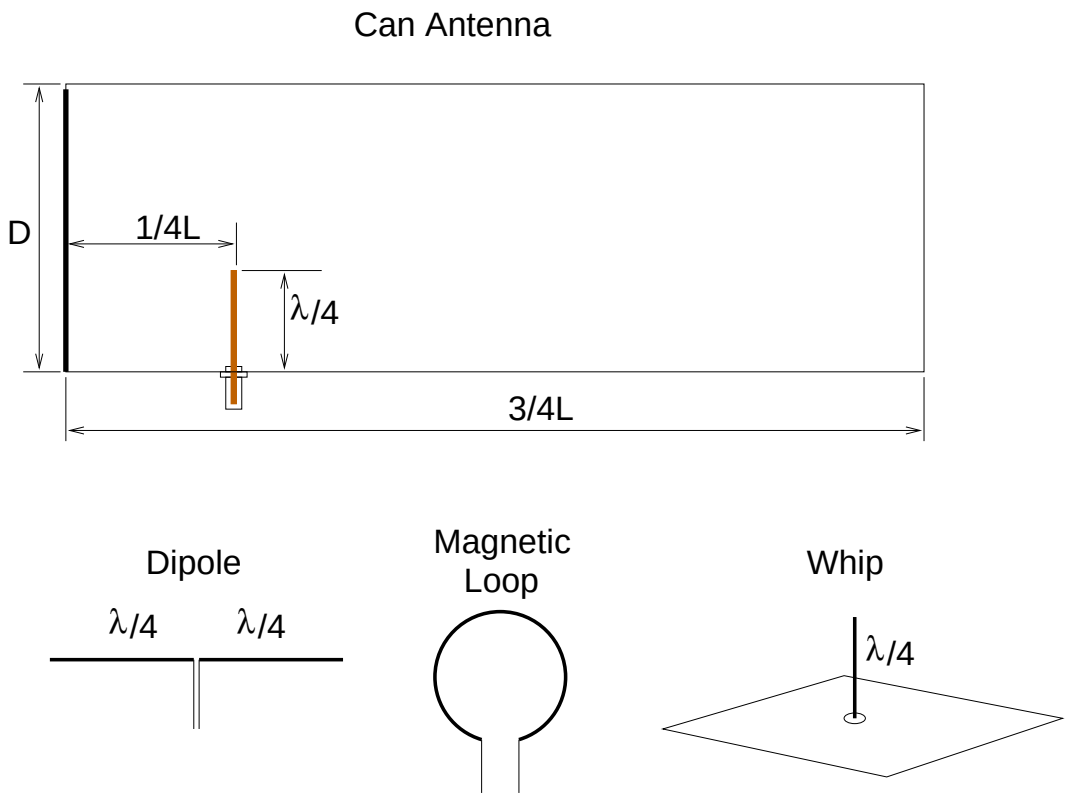

Figure 2.12: Schematic diagram of can antenna showing critical dimensions. One end of the waveguide is blocked with a solid conducting plate. The antenna is fed from a $\lambda/4$ feed-wire connected to a BNC connector. For comparison, a dipole, magnetic loop and whip antenna are also shown.

1. A $\lambda/2$ dipole

2. A magnetic-loop antenna

3. A whip-antenna

4. A horn or can antenna

The $\lambda/2$ dipole is an antenna formed by two conductors whose total length is half the wave length and point in opposite directions. The current is centre-fed, resulting in both conductors being in-phase with each other. The maximum power is transmitted in a direction perpendicular to the conductors.

A magnetic-loop antenna, which is a single loop inductor, was also tried. The loop dimension was smaller than $\lambda/4$ and the antenna has a reactive impedance. It is necessary to tune the antenna

using a capacitor so that resonance is achieved at the driven frequency. The magnetic flux in the antenna generates an RF field in the plane of the loop.

The whip-antenna was constructed from a single conductor perpendicular to a ground plane. This has similar characteristics to a dipole antenna, if the conductor is $\lambda/4$ in length, and the ground plane large by comparison.

Finally a horn or can antenna was built. This is a truncated circular waveguide, fed at one end and open at the other. The dimensions were chosen such that the waveguide supported a single mode, and coupled well to free-space. Schematics of the four antenna types are show in figure 2.12.

The dimensions of the can antenna were calculated using the formula for a cylindrical waveguide.

$$L = \frac{1}{\sqrt{(1/\lambda)^2 - (1/L_c)^2}}$$

L is the standing wavelength inside the can. λ is the wavelength in free-space. L_c is lower dominant mode cut-off wavelength.

The upper and lower cut-off wavelengths can be calculated from D, the interior diameter of the can. L_u is the higher dominant mode cut-off wavelength:

$$L_c = 1.706D$$

$$L_u = 1.306D$$

These are found by numerical solution of the roots of the Bessel function for a TE_{01} mode (see [47]). This gives a limit on D:

$$\lambda/1.706 < D < \lambda/1.306$$

For a given value of D in this range, L and the other dimensions give in figure 2.12 can be calculated.

The antennas were compared under similar conditions. The antennas were mounted on wooden stands, away from metal objects that might have caused reflections. Measurements were made at

different distances and angles using a $\lambda/2$ dipole as the receiving antenna. Different polarisations were also measured.

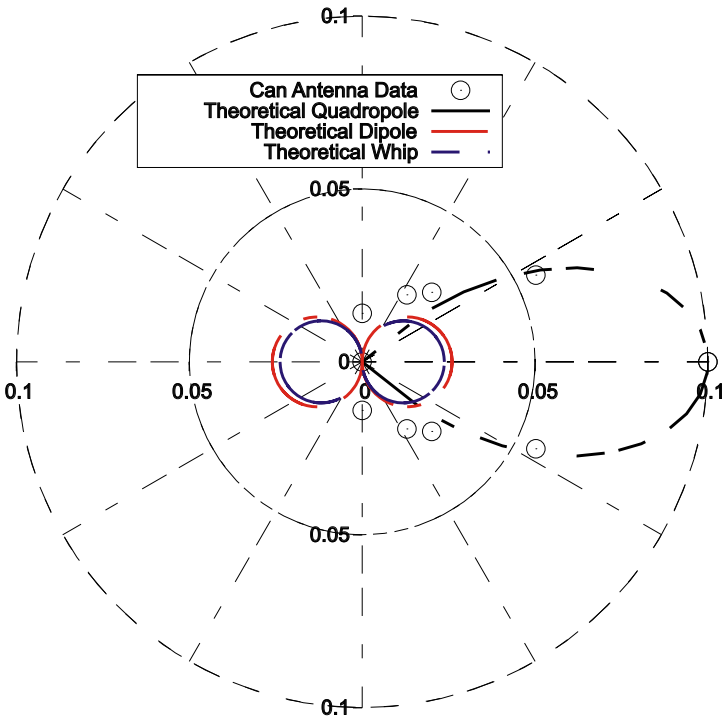

Figure 2.13: Radiation from whip, dipole and can antenna with an input power of 10dBm. Measured power on a receiving dipole at 20cm radial distance is plotted, in mW. It can be seen that the can antenna gives a radiation pattern similar to an quadropole and therefore has better directional gain than the other types of antenna.

The results are plotted in figure 2.13. The antennas were all driven with a power of 10dBm. The theoretical radiation pattern of a can antenna can only be determined numerically, but to guide the eye a quadrapole lobe has been plotted. It can be seen that the can gives a good directional gain of 12.3dB i.e. greater than seventeen-fold. The dipole performs slightly better than the whip antenna, and both perform better than a magnetic loop antenna which could not be made to efficiently radiate energy, due to not matching the resonant condition at 3.22GHz.

The can was initially used to radiate microwaves into the vacuum chamber. The antenna was placed as close as possible to the ion, which, allowing for mechanical practicalities, was about 20cm. To confirm transmission a dipole aerial was used to receive microwaves at a distance from the ion. With microwaves successfully transmitting, a resonance scan was performed on the ion. Resonance was found at 3.221405 GHz (see figure 2.14).

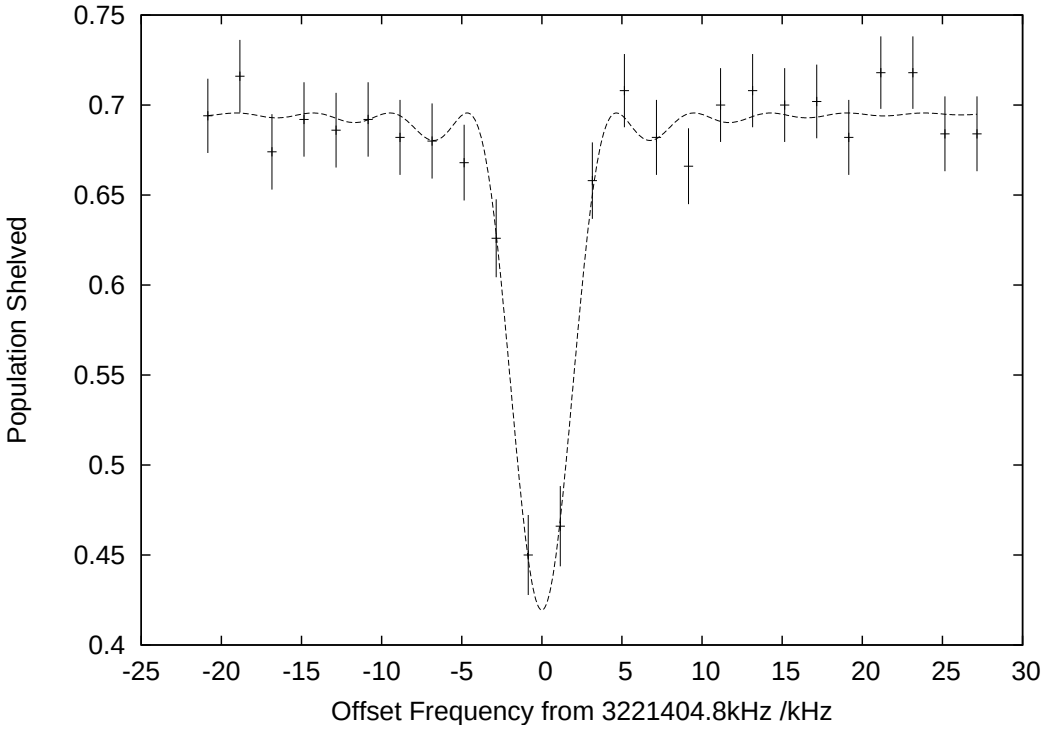

Figure 2.14: Resonance of the $S_{1/2}(F=4, M_F=+4) \leftrightarrow S_{1/2}(F=3, M_F=+3)$ hyperfine transition of $^{43}Ca^{+}$ driven with microwave radiation via a can-antenna. The population shelved is measured as an average of 500 experiments, at each frequency point. In each experiment, the $(F=4, M_F=+4)$ state is prepared, the microwaves are pulsed on, and the F=4 level is shelved to measure the state. The microwave input power was 1W after a 40dB amplifier. A microwave pulse time of 5ms was needed to see this result, showing that the intensity at the ion was very weak. The FWHM is ∼4kHz which would imply a $t_\pi \approx 250\mu s$ but in reality the actual resonance is broadened by stray external fields.

It was found that very little power was entering the vacuum system and reaching the ion. Power was being reflected and absorbed by the surrounding equipment and the vacuum chamber. The vacuum chamber has dimensions of a cylinder approximately 10cm in diameter. The viewports are 6cm in diameter. The wavelength at 3.2GHz is 9.4cm so radiation into the chamber at this frequency is inefficient. To obtain better coupling to the ion, another approach was tried. A co-axial lead that was already connected to a pair of trap electrodes (BL and BR) was used. A microwave signal was applied to this cable. Inside the chamber the cable was unterminated, thus it radiated the RF signal. This was found to give a much better response from the ion even though no attempt to impedance-match the free-end of the cable to free-space was made. Figure 2.15 shows a narrow reasonance, implying a 100-fold increase in Rabi-flopping speed. This technique

was used for the experiments presented in this thesis.

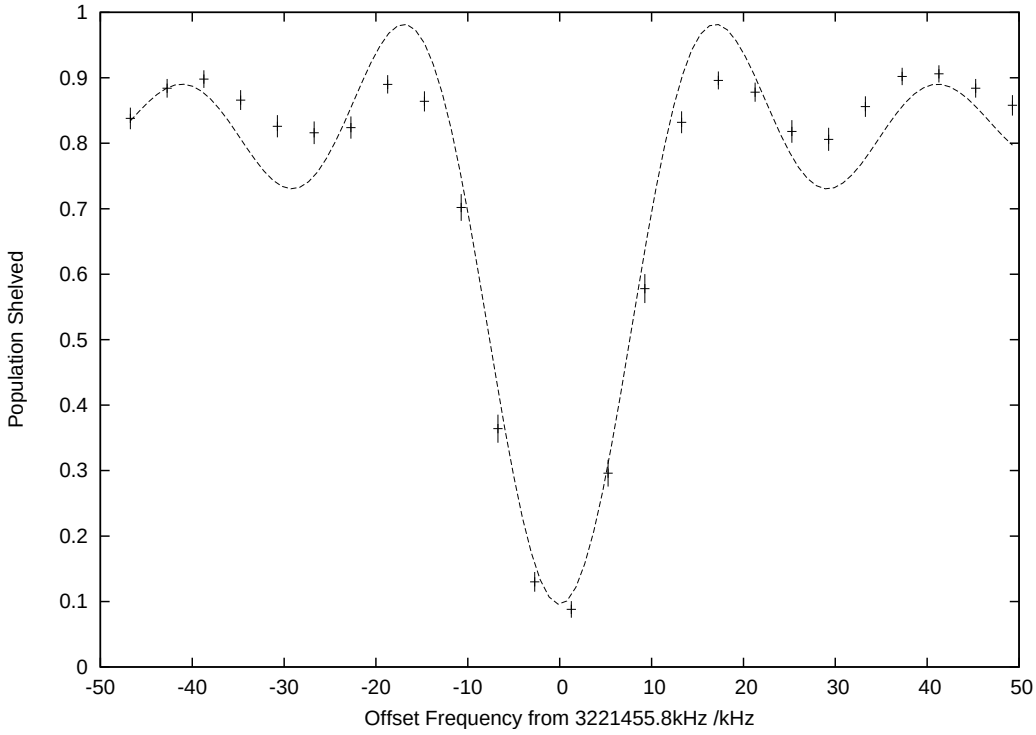

Figure 2.15: Resonance of the ground-state hyperfine transition of $^{43}Ca^{+}$ driven with microwave radiation on the BL/BR DC electrodes. The FWHM is ∼20kHz with a power of +30dBm (0dBm into a 30dB amplifier). t_{π} was found to be 53µs by varying the duration of the microwave pulse while the frequency was tuned to the centre of the resonance, showing the intensity was greater than from the can antenna by a factor of 100.

2.7 Imaging System

To capture the fluorescence as efficiently as possible, the largest solid angle that can be subtended at the ion is required. It would be possible to drive the fluorescence for a good proportion of the shelf's lifetime, however, the longer this time, the more chance there is that the ion might spontaneously deshelve. Ideally the measurement time is kept as short as possible. This is done by being as close to the ion as possible and using a large diameter lens. The vacuum system determines how close to the ion a lens can be placed. A lens with a wide diameter and a large Numerical Aperture (NA) was required. NA is defined as the ratio of the lens's working radius over the distance to the object, $NA = r/d$, i.e. the maximum angle over which light rays can be captured and focused to a point.

For a large NA lens, the paraxial approximation no longer holds, and a simple lens will suffer from chromatic and spherical aberrations. Therefore a compound lens system is required. Typically these systems are designed to work at a particular distance from the object. Placing the object elsewhere will not guarantee removal of aberrations. It would be expensive and difficult to design such a system, therefore a commercial microscope objective was used. Ordinary microscope objectives are designed to work very close to the object, typically 10mm or less. Some even require fluid-coupling to the sample under view. However, due to the vacuum window being greater than 30mm away from the ion, and a requirement to access the ion in the same plane with laser beams, a lens with a large diameter was required. A specialist lens (Nikon ED PLAN 1.5x) was sourced. This lens is designed for use on an inspection microscope which requires good access to the sample and therefore a large diameter (60mm). In normal use it works with a zoom body and an eye piece, to re-image the object to the user's eye. However, in this setup, the lens was used as a single lens.

Information supplied with the Nikon lens was not applicable for the way in which the lens is used in the experiment. It was therefore necessary to measure the position of the principal planes. The focal length was specified to be 65mm and this was confirmed by measurement [40]. The optical arrangement shown in figure 2.16 already existed, see [46] for details. A magnification of 3.2 was achieved by placing the ion at a distance such that the object was re-imaged at $3.2\times$ the object-principal plane distance. This was close to the lens's specified working distance, but not exactly at it. Some aberrations occur because of this difference and because of the vacuum chamber window, but these aberrations do not affect the photon count. A pin-hole was placed at the image plane to block scattered light from the lasers. Additionally, a glass filter is used to filter out any wavelengths other than $\sim$397nm to remove scattered light from the repumping lasers.

The image is collimated using a simple lens and passed through a beam-splitter. The ion is re-imaged onto a CCD array in one direction, and onto a photomultiplier tube in the other direction. To maximise photon-counts, the beam-splitter can be removed.

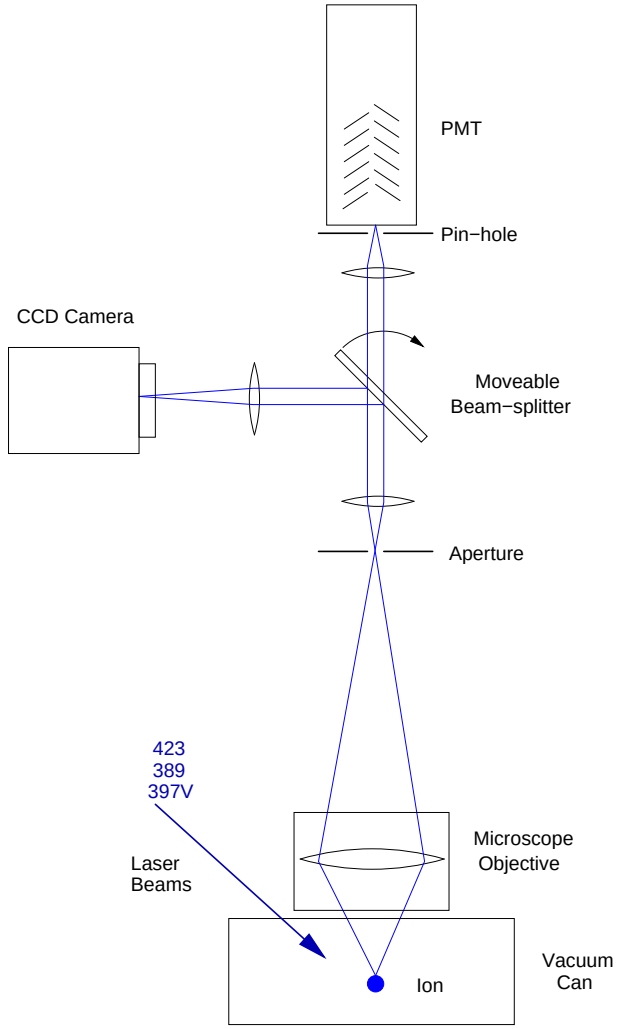

Figure 2.16: Imaging system, showing use of a compound microscope lens as an objective. The second lens recollimates the beam through a movable beam splitter and then the light is focused onto a CCD camera and a PMT simultaneously. To increase PMT counts, the beam-splitter can be rotated out of the way. Apertures, pin-holes and (not shown) chromatic filters reduce scattered laser-light reaching the PMT.

Chapter 3

Laser Control Unit

3.1 Introduction

3.1.1 Purpose

The purpose of the Laser Control Unit (LCU) is to create an accurately timed sequence of pulses to digitally control experimental devices. This allows experiments such as the Ramsey experiments that are described in chapters 6 and 7 to be performed.

The LCU was designed to control 16 devices, such as lasers, Acousto-Optical Modulators (AOMs), RF switches, shutters, and other pieces of optical, control and read-out equipment.

The LCU was designed around a timing card that fits inside a PC and is connected to the PCI bus. This allowed the timing sequences to be programmed from within the existing control environment, which is an IBM-compatible computer running MS-DOS.

3.1.2 Application

An experiment (with typical durations given in parenthesis) might consist of a sequence such as:

1. Switch on cooling laser for a period (1ms).
2. Switch on population preparation laser (10μs).
3. Close shutter (3ms).

4. Switch RF frequency generator to the correct phase and allow to stabilise. This pulse does not have to be accurately synchronised with the other pulses, so long at the RF has settled before the next pulse is fired.

5. Fire a Ramsey pulse for a time $\pi/2\Omega$, where Ω is the Rabi frequency of a particular transition (50μs).

6. Wait a time τ, where τ is the wait-time in a Ramsey experiment (0.1ms to 1s).

7. Fire a Ramsey pulse for a time $\pi/2\Omega$ (50μs).

8. Fire the shelving lasers (100μs).

9. Open a shutter to allow lasers to drive fluorescence (3ms)

10. Count photons with a photomultiplier tube (PMT) for a given time (1ms).

Some devices may need to stay on, such as the cooling laser, whilst other devices might need to be fired for a very short pulse during the "off" interval. To provide for these more complicated situations, an external input allows the device to be triggered from another source (such as a high-resolution pulse generator). An example sequence is shown in chapter 6, figure 6.3. Some parts of the pulse sequence happen independently of the pulse sequence, such as switching on the RF source, before firing the RF pulse. This is an example of an overlapping pulse where synchronisation is not critical.

3.1.3 Specification

The LCU needed to meet the following specifications:

1. Be able to drive up to 16 separate devices. Each device can be switched on or off independently under software control.

2. To provide a variable-duration "on" pulse, followed by an accurately timed "off" period. This sequence needs to be repeated, with the "on" time being continuously varied under software control in real-time.

3. The pulses need to be settable to a resolution of 0.1 µs. This requires a base clock frequency of 10MHz.

4. The switching of devices needs to be accurately synchronised. This requires that the delay between devices being triggered is $\ll$0.1 µs.

5. The transition time i.e. the time to go from a TTL low to a TTL high (0 to 5V) needs to be short compared to the duration of the on and off pulses; about 1/20 of the minimum pulse period. This corresponds to a transition time of 50ns. High-speed TTL components have a typical transition time of 5–10ns.

6. The system needs to be realistically immune from noise and crosstalk errors. This would ideally be a zero error rate over several years of experimental running, however a tolerance of 10^{-6} would be acceptable at current experimental reliability (10^{-3}). For TTL signals a signal to noise ratio of 10^{-2} is sufficient to prevent a zero becoming a one.

7. To provide manual switches which can override any device to be either on or off, regardless of the software state. This allows an experimenter to e.g. switch on a laser whilst aligning the beam, without having to run software or rewire the device.

8. To be able to switch the device from another timing unit such as a high-precision, short pulse generator.

9. To indicate the state of each channel on a 16-channel digital oscilloscope.

10. To be able to drive TTL signals into a 50Ω load with active-high, or active-low logic.

11. To be able to asynchronously switch some devices, independent of the timed-pulse. This allows overlapping pulses that are not accurately synchronised. For example, switching the phase of an RF synthesiser during a pulse sequence, but before the RF pulse is fired.

3.1.4 Rationale

There are commercial cards available that are designed to provide accurately timed pulses; however, these systems do not provide the flexibility required. A device has been built jointly by

NIST, ARDA and MIT [48] which uses a microprocessor system implemented on an Field-Programmable Gate Array (FPGA) – a chip that can be programmed to perform any combinatorial logic. This system was designed to meet the requirements of the quantum information community. However, whilst the system meets all of the above requirements, and provides speeds up to 100MHz, it requires a steep learning curve to program in bespoke machine-code (though tools are provided). The outputs operate at speeds of 100MHz and therefore the signals are routed using a low-voltage differential pair bus. Interfacing this to 50Ω loads would require extra, complex high-speed electronics.

The most advanced commercial systems available, and the NIST system above, allow conditional loops to be programmed into a pulse sequence, but a system built on a PC allows complete programmability in a high-level language such as Pascal. The advantages of a PC-based system over a commercial microprocessor-based system are listed below:

1. There is effectively no limit to the length and complexity of the sequence.

2. The pulse sequence can be dependent on inputs from other devices such as a given count on a PMT. In this way Quantum Error Correction (QEC) can be implemented. QEC requires corrections to be made, based on a PMT reading in real-time. The correction is made by firing additional pulses into the sequence.

3. A microprocessor needs to process instructions during clock-cycles, as well as fire pulses; thus only a few operations can be carried out during short pulses. Using the PC's central-processor and a dedicated timer chip guarantees the length between pulses, and the minimum resolution of pulse-length. Other processing is done at the 1GHz speed, giving time to update the computer display and read inputs, whilst not interrupting the timer-chip. In order to circumvent this problem, the NIST system has two central-processors.

4. A high-level language can be used; therefore, the software is more easily maintained without having future compatibility issues.

5. Cheap generic PC equipment that remains readily available is employed. Electromagnetic Compatibility (EMC) (noise immunity) and low-noise design has already been taken care of by the manufacturer. Processor speeds of 4GHz and higher are available.

6. Other hardware can be connected to the PC via serial ports and GPIB. These can be incorporated into the same control program.

7. Speeds of 100MHz can be achieved using PCI-Express, and future development of the ubiquitous PC will no doubt bring further speed increases.

The advantages of a microprocessor based system are:

1. It can be a stand-alone box that can be unplugged from the computer. However this "black-box" approach means that the user has limited control over what is going on inside.

2. Overlapping pulses can be accurately synchronised possibly with more than one time-base.

3. The speed is not limited by the architecture of the PC (e.g. current PCI bus speeds), currently this means they can go faster.

4. There is not a dependence on a real-time operating system.

3.2 Implementation

The unit is designed around the PCI-DAS1200 from Measurement Computing. This is installed in the host PC used to control the experiment. It is connected to a box called the Laser Control Unit (LCU) which provides synchronisation, buffering and extra logic, terminating in BNC connectors for ease of connection to other devices. This is shown in figure 3.1. A photograph of the finished device in operation is shown in figure 3.2.

The PCI-DAS1200 card provides two counter chips (Intel 82C54) each of which contains three counters. It also provides a 10MHz base-clock, 24 digital I/O lines, 8 differential (or 16 single-ended) analogue inputs, and 3 analogue outputs. The counters are wired to cascade each other, so that the "off" of one pulse triggers the "on" of the next (see figure 3.3). Each counter has a 16 bit resolution and can count for n clock pulses where $1 \leq n \leq 65536$. The base clock was set to its maximum rate of 10MHz (using the internal frequency source), therefore the counter could generate pulses of 0.1µs resolution. A maximum pulse duration of $(2^{16}-1) \times 0.1\mu s \approx 6.5$ms can be achieved. However, it is possible to repeat pulses, with only a small gap between them in order to achieve very long pulse times, such as the 1s times used in Ramsey experiments in chapter 6.

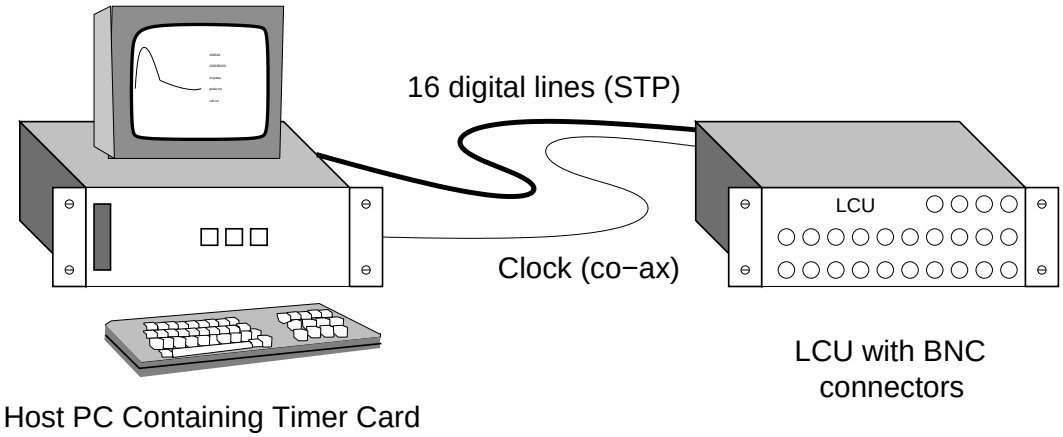

Figure 3.1: The overall system showing timing card, host PC, LCU and interconnections.

Figure 3.2: Photograph of the LCU in use. To the left is the grey cable connecting the unit to the PC. The clock is fed from the PC via co-axial cable. The other cables connect to external devices under control.

This pulse train is synchronised with one of the 16 digital outputs from the PCI-DAS1200 using the combinatorial logic described below. External inputs and manual override are also allowed for, and the output is buffered using a special TTL chip that can drive 50Ω loads. The synchronisation, buffering and extra control circuitry is housed externally, and connected to the PCI-DAS1200 using special wiring also explained below. Devices are connected to the unit using co-axial cables. The output is monitored on a 16 channel scope. The control program is written in Pascal and runs on the host PC.

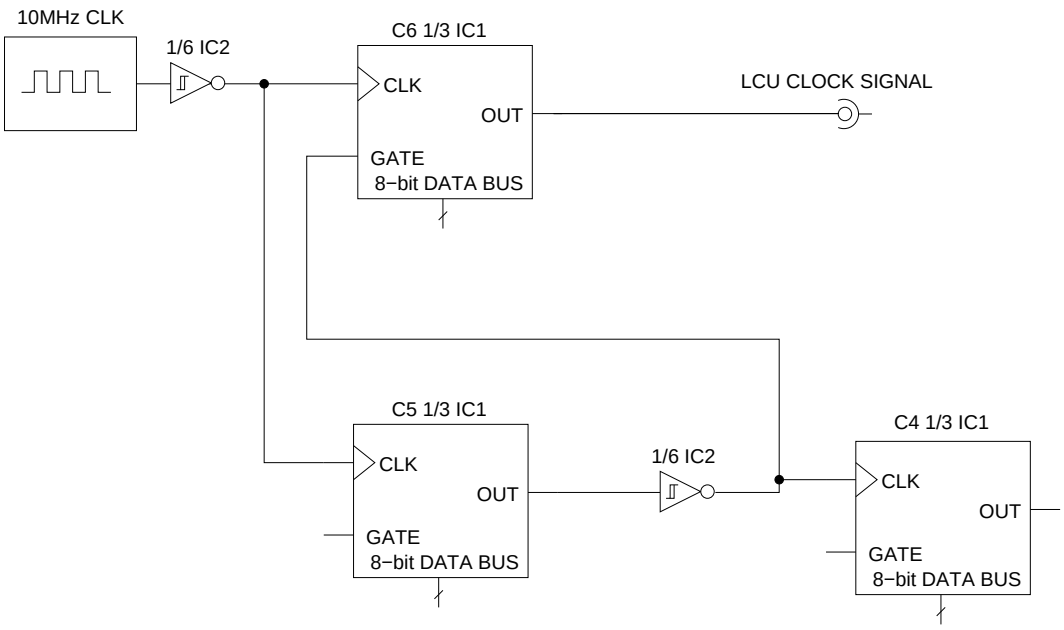

Figure 3.3: Wiring diagram of the second (counters 4, 5 and 6) 82C54 counter IC. C5 is used to gate C6. Schmitt triggers are used to ensure that there are no crosstalk errors. C4 counts the number of pulses to ensure that there are no double firing of pulses. The data-bus of the counters is internally connected to the PCI bus of the PC and can be programmed from there.

3.2.1 Pulse-train Generation

The pulse-train is made up of an on period, called the "mark", and an off period called the "space". To generate the space period, one counter (C6) is used, and a second (C5) is used to generate the outgoing (mark) pulse. The mark counter counts clock pulses supplied from the base clock at 10MHz. When it has counted to the programmed count, n_{mark}, it gates the space counter, and waits until it is reset in software. To ensure that only one pulse is fired each time, C4 is used as a check. The mark counter will count the number of clock pulses, n_{mark} that it has been programmed for. The space time does not change. During this space period, the digital lines are set to be high or low and the next length of pulse is programmed into the mark counter (this will be effective in the next gating of counter C5). It needs to be a minimum, but long enough for the reprogramming to occur. With the total number of commands that need to be sent to the timing card (setting digital lines A,B and C, the next pulse-length, and checking count on C4) a time of 5.4µs is required, in total the space count is 6µs long (n=60) to allow for hold-ups on the PC's bus. Figure 3.4 shows a simplified timing diagram. The mark time can be anywhere between 0 and $(2^{16}-1)\times 0.1\mu s$

less the 6μs gap. Longer pulses are made from two pulses, with a short 6μs gap between. Which device is turned on during the clock mark period is determined by setting the corresponding digital out lines high.

The PCI-DAS1200 supplies an internal clock at 10MHz that was sufficiently stable. However, this could be an external clock if necessary.

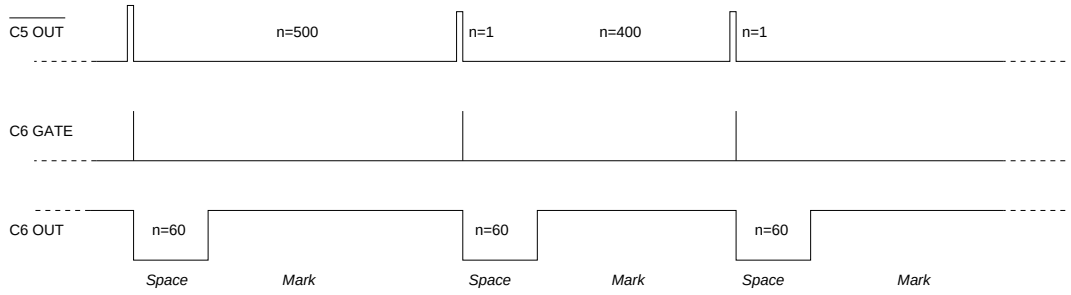

Figure 3.4: Simplified timing diagram showing how the counters are gated to provide a variable length pulse-train. n = number of 0.1μs ticks

To ensure that the correct number of pulses have been fired, C4 is used to count out-going pulses. The count is read from the counter in software. If this does not match the expected number of pulses, an error is displayed on the PC.

It was observed that for pulses over 100μs, an extra clock-period is counted. This occurs randomly, at a level of about 1 extra clock pulse in 10^4. For very long pulses this is quite frequent. For pulses of 100 μs or less it is rarely seen. However, for long pulses such as Ramsey delays, an error of 1 in 10^4 is tolerable. The cause of the error has not been explained.

3.2.2 Synchronisation and Buffering Circuit

To synchronise the pulse-train to the digital output of the PCI-DAS1200 board, and to provide extra functionality, an external circuit was required. Figure 3.5 shows the basic idea of synchronising the data output lines to the clock, using AND gates.

The logic required for synchronising the digital outputs with the clock signal (pulse-train) produced as above was:

```
OUT = (ENABLE AND CLOCK) OR EXTERNAL
```

where ENABLE is one of the digital out lines and CLOCK is the pulse-train (not the 10MHz

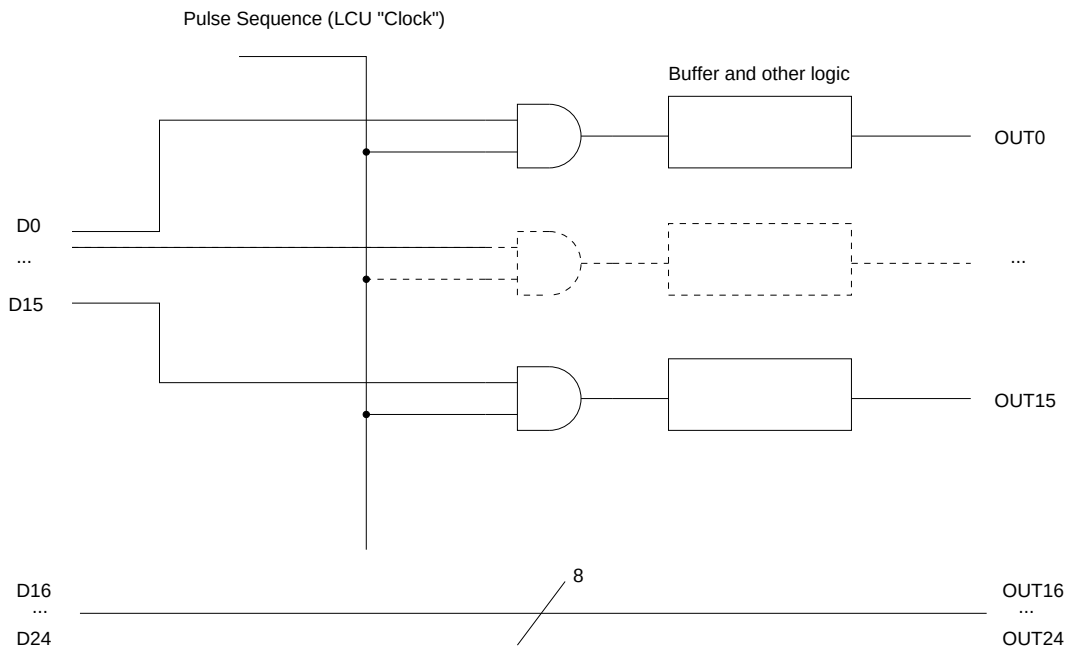

Figure 3.5: An overview of the LCU showing how the 16 digital lines are synchronised with the clock. Also shown are the 8 non-synchronised lines which connect directly to experimental equipment.

base-clock used to generate it). EXTERNAL is an input to the LCU that allows an external pulse generator, or another control card, to trigger the same device, or to hold it permanently on. Additionally it is possible to switch the PC's digital output to control the EXTERNAL line instead; effectively bypassing the CLOCK. This allows the control program to switch devices asynchronously. These pulses will suffer timing inaccuracies due to the PC "multitasking". This allows the asynchronous pulses, such as switching on an RF source, to occur before the pulse is fired. The accuracy of these types of pulse is not important.

The above logic does not allow for a switch to control the output manually, nor does it allow the output to be inverted. A naïve approach to this problem might be to connect the output to an inverter and have a switch to bypass the inverter. However, routing 10MHz signals through a switch requires taking RF signals to the front-panel of the unit and back, which would cause cross-talk and noise. Furthermore, routing the signal either through an inverter, or not through an inverter, would produce a variable propagation delay (each gate generates a propagation delay, typically 30ns for a TTL IC) and the output channels would no longer be synchronised. To avoid both these problems, the manual override and inversion was implemented using the following

logic:

```
OUT = INVERT XOR (((ENABLE AND CLOCK) OR EXTERNAL OR P) AND Q)
```

The P and Q inputs were provided by a three-position, two-gang switch, which sets the manual over-ride to on, off or computer control.

Switch Position	Device Status	LED Colour	P	Q
Up	On	Orange	1	1
Central	Computer Controlled	Green	0	1
Down	Off	Red	1	0

The above combination allows the manual over-ride switch to directly drive a tri-coloured LED, which indicates the state of the device as on, off or under computer control. The last step inverts the output (if desired) so that for a device that is active low, the switch positions, LED colour, and computer output all still have the same meaning, i.e. a red LED would mean the device is off, even though the output is high. The final output is passed to a digital oscilloscope (Agilent 54622D) which can display 16 channels of logic simultaneously. This allows an independent diagnostic of the pulse sequence.

The incoming clock was passed through two inverting Schmitt-triggers. This cleans up the clock pulse, and with the use of a bypass jumper on the second inverter, allows the clock signal to be inverted. The jumper was in place for the final application, as the output from the timing card's clock needed to be inverted to the sense required. A source of pulses other than the PCI-DAS1200 could be used if required.

Implementation

The logic function above was provided by high-speed CMOS TTL (HCT) chips. These chips are high-speed due to the CMOS circuitry, but provide TTL levels, and thus are able to drive TTL devices directly. HCL and LS chips were originally tried, but LS chips were not fast enough – showing considerably slow transition times – and HCL chips were not capable of driving the TTL output stage. Each chip has four gates on a chip, and as there were 16 channels, it was natural to use specific chips for each gate rather than, for instance, building the same logic from either NOR or NAND gates. This reduced the chip count. A circuit-diagram for one channel is shown

in figure 3.6.

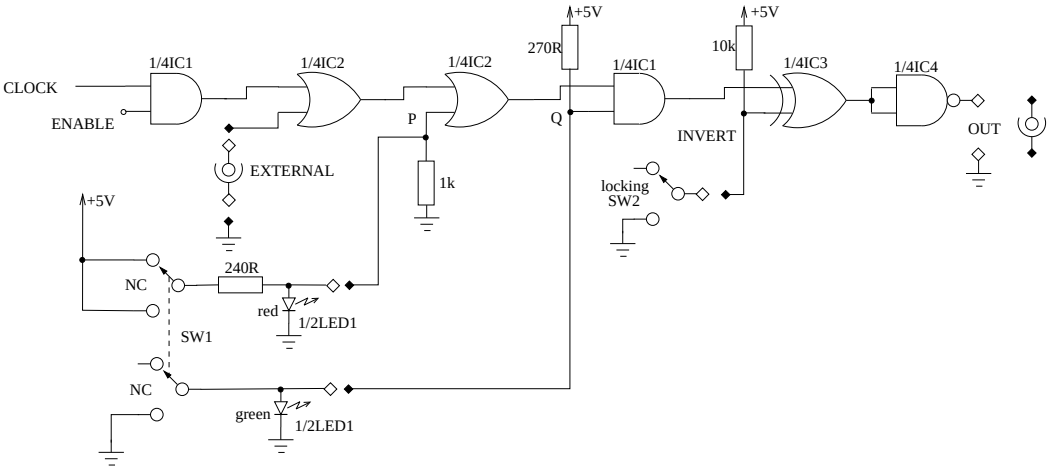

Figure 3.6: Circuit diagram for one channel of LCU. IC1-3 are HCT logic chips. IC4 is a 74F3037 buffer designed to drive 50Ω loads.

Power was provided with a commercial power-supply module (Power-One HB5-3/OVP-AG). The power requirements were calculated on a worst-case scenario – i.e. the outputs would all be driving 5V DC into a 50Ω load at the same time. Assuming $16 \times 5/50 = 1.6$A for the output, and about 50% again to power the ICs, this required a power-supply of 2A. A 3A power-supply was chosen, so that the extra power could be used to drive external circuitry via output connectors on the front-panel.

Construction

All the logic was built onto a PCB. All the switches, inputs, and outputs were built onto the front panel of a 3U 19” relay case.

The PCB was designed using `pcb`, an X-Windows application, available under the GNU Public Licence. The artwork was printed by Photodata, and the circuit board was made by the physics department’s photo-fabrication unit. It was necessary to use a two layer board, due to the complexity of the circuit. This allowed extra grounding to be added to reduce cross-talk between the PCB tracks. A ground plane was not incorporated, due to the potentially high capacitive loading of the logic ICs, which would reduce transition time.

3.2.3 Cabling

All high-speed connections were made from the circuit board to the front-panel using mini-coaxial cable (RG178) to 50Ω BNC connectors. Connections from here are made to the device using standard (RG58) 50Ω coaxial cable, with suitable termination. Termination was necessary to prevent effects, such as end-reflection, which cause loss of power and instability at the device. All other connections were made using twisted-pair cables.

The output from the PCI-DAS1200 was via a 100-way IDC connector designed to take a ribbon cable. The design of the card was such that high-speed digital signals and clock signals were routed next to each other in this connector, as opposed to having alternate grounds. This caused a serious crosstalk problem.

To calculate the size of the crosstalk in a ribbon cable, it is assumed that two parallel high-speed signal lines will be capacitively coupled. Approximating a ribbon cable as a parallel-plate capacitor does not work well, as the wires are longer than they are wide, and they are cylindrical, not flat. A more accurate formula is given by [49]:

$$C = \frac{\pi \varepsilon l}{\cosh^{-1}\left(\frac{d}{b}\right)}$$

where b is the spacing between the wires, l is the length of the ribbon cable and d is the diametter of each conductor.

However, the permittivity is not quite uniform, as it is dependent on the way the insulation is placed between the conductors. It is assumed that $\varepsilon = \varepsilon_0$, but this gives an error of $\sim$2. b is the spacing of conductors (0.05inch or 1.27mm) and d is the diameter of the wire (24AWG or 0.5mm). This gives a capacitance per unit length of 24pF/m. Initially a length of 30cm ribbon cable was used. This gave noticeable crosstalk. The impedance due to the ribbon cables capacitance is:

$$Z_C = \frac{1}{2\pi C f}$$

where f is the switching-speed.

For the clock signal, at 10MHz, $Z_C = 660\Omega$. If the load at the far end of the ribbon cable was 100kΩ (typical for a CMOS gate) then the voltage divider formula can be used (see figure 3.7).

This gives

$$\frac{V_1}{V_2} = \frac{R_g}{R_g + Z_c}$$

where V_1 is the voltage on conductor 1, V_2 is the voltage coupled to conductor 2, R_g is the gate resistance, and Z_c is the effective impedance due to the coupling in the ribbon cable.

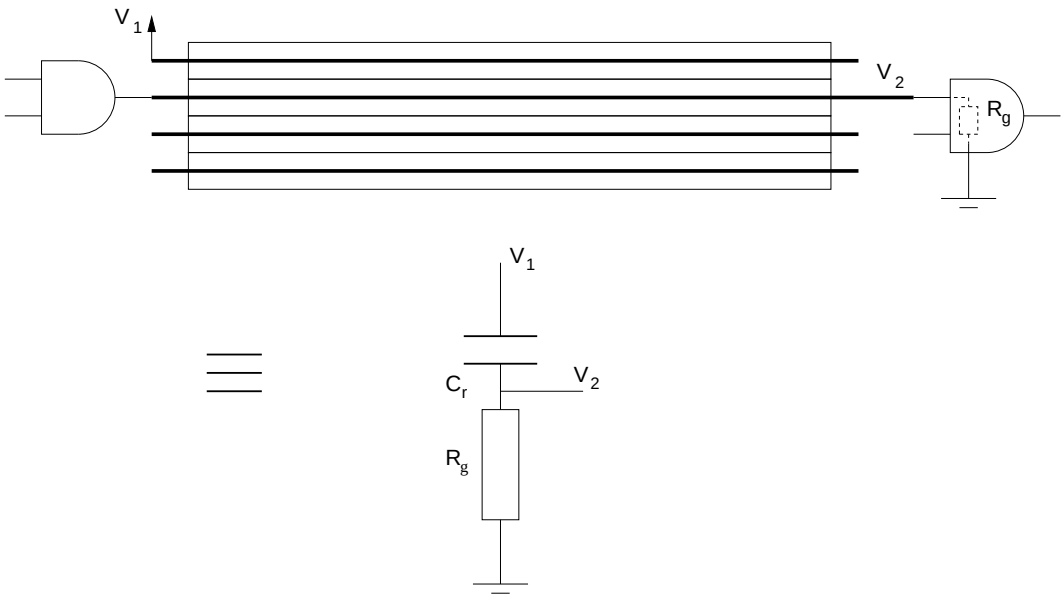

Figure 3.7: Equivalent circuit for a length of ribbon cable used in digital control, showing capacitive coupling between conductors. If conductor 1 is at V_1 and conductor 2 is held low, a voltage, V_2 will be coupled into the second conductor.

For the worst-case scenario, with one conductor high and the other held low by the computer, the voltage on conductor 2 will be $0.98\times$ the voltage on conductor 1. In fact, a value of about 2.5V was seen. The TTL threshold needs to be below 2V to prevent triggering of a gate.

Using capacitance alone, the above calculation gives a good approximation. However, any real cylindrical conductor will have an inductance, and there will be mutual inductance between conductors as well. Additionally, the inductance will cause a voltage drop along the line, which creates a situation closer to a parallel-conductor transmission line. There will also be dissipative loss. These extra effects are much smaller than the capacitive coupling, and the actual voltages measured are in close agreement with the above calculation. In a digital pulse, assuming a square wave, frequencies up to $10\times$ the base clock frequency need to be considered, to give the required transition rate. The first harmonic of a square wave is at $3f$ and at 1/3 the amplitude. The induced

voltage due to this harmonic is therefore:

$$V_{3f} = \frac{RV_f}{3R + Z_c}$$

Assuming $R \gg Z_c$, the induced voltage due to the second harmonic is $< 0.33\times$ the amplitude of the induced voltage, and below the TTL threshold.

It would be possible to effectively increase the gate resistance by adding a resistance in series to the gate input, however this would draw more current from the PC and was not desirable. Initially the ribbon cable length was reduced to 1cm. This reduced the voltage to $0.6\times$ the input, however this was still sufficient to trigger a TTL signal.

One commonly used solution is to alternate grounds between signals in a ribbon cable. This greatly reduces the coupling between conductors, though a second-order effect still exists between the current induced in the ground wire and the next conductor. For ribbon cables under a metre in length, this effect is insignificant. It was not possible to do this, as the wiring of the PCI-DAS1200 did not provide alternate grounds.

Instead, twisted pair cabling was used. Unshielded twisted pair (UTP) cable is very common in telecommunications and computer applications, as it is very cheap, simple to use, and sufficiently immune to noise. Any noise picked up will be equal and opposite on each line, assuming that the period of the twist is smaller than the wavelength of the noise. Typically the twists are 1cm apart, and therefore any signal with a frequency lower than 20GHz is cancelled. This is referred to as common mode rejection (CMR). Crosstalk immunity is provided by adding shielding as well, this is called shielded twisted pair (STP). Therefore, multiple pairs can be bundled together. The characteristic impedance is very high, so unlike co-axial cables, there is no requirement to terminate the cables with a matched impedance. STP has replaced co-axial cables in many applications. Specially designed connectors such as the XLR connector are available for terminating twisted pair (shielded or unshielded) and provide good mechanical performance as well as high noise immunity.

However, the provided connector was an IDC type and could only be connected to ribbon cable. For the timing signals, the ribbon cable had to be cut into, breaking apart the parallel conductors carrying the clock signal, and connected directly to Schmitt triggers (74LS14) which

ensured that the gate signal for each counter was not affected by crosstalk.

For the other digital lines, the signal was sent down 3cm of ribbon cable to a circuit board, which then connected the digital lines of the LCU unit to the PC via a bundle of STP cables, terminated with 25-way D-types. The cable met the IEEE-1284 specification, so that it could easily be replaced. No crosstalk was measurable after the first set of gates inside the LCU. Crosstalk was observed on the 16-channel digital scope but not on the final output. Further investigation showed that this was due to the ribbon cable provided by the oscilloscope manufacturer, which was 2m long, and created crosstalk of its own that was not present on the experimental devices.

Analogue Connections

As well as providing digital outputs, it is also necessary for the control computer to receive analogue inputs, for instance to monitor laser power using a photodiode. The PCI-DAS1200 analogue inputs were used for this purpose. These analogue lines can be used in two modes, either single ended or differential. In differential mode, each line is symmetric. This allows the two signal lines to be set up as a twisted pair using the same cable as the digital signals above and provides very high noise immunity.

However, a lot of equipment provides single-ended outputs connected with BNC connectors. Therefore, the analogue connections from the PCI-DAS1200 analogue inputs are first routed to a patch panel using twisted pair cable. Four inputs terminate in XLR connectors, where they can be used in differential mode. The other four are fed to BNC connectors, with each line being available as the centre pin of the connector, so that one differential pair can be used as two single inputs.

3.2.4 Supply Ripple Reduction

With 16 channels driving 0.1A each, considerable ripple was generated on the supply rails. To prevent this triggering digital inputs, smoothing capacitors were placed close to the ICs, as recommended by the IC manufacturers. The logic ICs each share a 10nF capacitor, which was sited as close to the IC as possible, whilst being connected across low-impedance supply rails. The 50Ω driver chips (74F3037) required the greatest amount of smoothing. Each chip required a 1μF

electrolytic capacitor as well as the 10nF ceramic capacitor that the other ICs required.[1] If the supply rails had a significant intrinsic impedance (e.g. thin PCB tracks, high-resistance material, long wires, or co-axial cable) then parasitic oscillation would have occurred. The cables from the PSU were twisted together to prevent oscillation, and smooth the supply. Even with all these precautions having been implemented, small amounts of transient ripple were seen when several channels switched simultaneously. This ripple was well below the digital switching threshold, and did not cause a problem.

3.2.5 Software

The current control system is based on an IBM-compatible PC running MS-DOS 6.2 and a program written in Turbo Pascal 7.0. This system was chosen because it is fast. As it is only processing the control program, no processor "time-slicing" is occurring, as it does in modern operating system such as UNIX and Windows. MS-DOS can be made to operate in real-time mode by disabling hardware interrupts. It is also non-specialist and affordable. Turbo Pascal will run on MS-DOS, and it provides a library of graphics routines and low-level access to the PC's hardware. Despite being over twenty-years old, there is still an active community providing drivers for current hardware.

However, there are drawbacks to the system. The program is one large monolithic piece of code. It is very unwieldy to manage between several developers. Turbo Pascal and DOS are no longer supported by their production companies. New technologies, like dual-VGA and USB are not officially supported by DOS (though open source drivers are available). Drivers for specialist cards can not be used, and must be implemented in low-level code. Network access and data analysis must be done at a separate time.

Currently the control program provides on-the-fly presentation and statistical analysis of the photon counts. This information can instantly alter the control procedure, for instance to perform QEC. However, graphical presentation is a CPU-intensive task and without time-slice multitask-

[1] Using a small ceramic or tantalum capacitor in parallel with a large electrolytic capacitor is standard practice due to the large series inductance of an electrolytic capacitor. This makes it effective only up to several hundred kilohertz. Above this, the high impedance means that higher frequency signals must be removed using a capacitor with lower series inductance. Ceramic capacitors are good at 10MHz or higher. This was in accord with the manufacturer's recommendations [50].

ing this requires a highly sophisticated level of programming. It is commendable this has been achieved, but as the system increases in complexity, it may be necessary to decouple the presentation and the control aspects of the program. Currently, however, the system works very well and provides the ability to see and edit the experimental sequence, whilst also displaying real-time data on the screen, all from one PC.

To program the PCI-DAS1200, knowledge of the PCI bus is required. The PCI bus is controlled by a bus controller, either from BIOS or the operating system. MS-DOS 6.2 requires the BIOS to setup the PCI peripherals. On boot up, the BIOS queries each physical PCI location (this scheme is referred to as geographical addressing, as opposed to ISA, where each slot is identical). Each device in turn responds with a request to be allocated address space and memory locations. These are assigned by BIOS. If the card moves location or another card is added in a slot that is queried before the timer card, then the addresses will change. The PCI-DAS1200 requires four base-addresses.

Once these base-addresses have been assigned, the control program can access the PCI-DAS1200 card directly. A program is used to discover the base-address, and currently once this is found, it is hard-coded into the control program. It would be possible to use a "plug and play" scheme in the future.

The PCI bus is 32-bits wide and can transfer data at a maximum rate of 33MHz. Faster versions exist, such as the PCI-Express standard, which has data rates of 100MHz and faster. There are two modes: address mode and direct memory addressing (DMA). DMA is faster, as once a connection is established, the card controls data transfer, using the full width and speed of the bus, with no interruptions.

In normal addressing mode, a bus controller is used to queue and send data, and therefore it is inherently slower. The PCI-DAS1200 does not support DMA mode. The bus controller will only send data when it gets a ready signal from the card. Tests carried out showed that the maximum rate the card could be programmed at was 2MHz. The timer chips are 16 bit chips, and need to be loaded with two 8-bit instructions in succession. It requires 6 instructions to the card to set the count of one counter. This takes 3μs. However, some allowance must also be made for the bus controllers queue. A set-up time of 6μs between pulses was used.

MS-DOS achieves a simple form of multi-tasking using interrupts. Interrupts can either be software or hardware, but MS-DOS maps hardware interrupts to software. Hardware interrupts are an electrical signal from a peripheral asking for attention. These go to an interrupt controller, which determines priority and queues interrupts until they are serviced. A software interrupt is also called a system call in other operating systems. It is a quick way to jump out of a user program and to an operating system procedure. The OS ensures that program flow is returned to the program after the interrupt handler, or system routine, has finished. For example, each key-press on the keyboard fires an interrupt. The interrupt handler is called, and, for instance, the character is displayed on the screen. If this occurred during the running of a pulse sequence, a pulse might end and the next one not get programmed, because program flow has been handed over to the interrupt controller. By the time this code returns operation back to the control program, the pulse duration might have extended. Realistically, on a 4GHz processor, the interrupt handler normally returns in under 1µs. However, to ensure reliable operation, it was necessary to disable interrupts during the entire sequence (typically 10ms). Interrupts are turned back on between pulse sequences, e.g. to get input from the keyboard once more.

In order to detect that the pulse has ended, and the card was ready to be programmed with the next pulse length, the third counter IC was used. To start counting, a pulse is sent to this counter. This changes its internal register, which is read in software. This had the added advantage of detecting any double pulses that may have occurred, and the software can flag an error on screen.

An example listing is shown in appendix A.

3.3 Results

Figure 3.8 shows the timing sequence that was used to achieve sideband cooling of both centre-of-mass and stretch modes of two ions, followed by a "wobble gate" operation, and enclosed in a spin-echo sequence – to measure the effect of the gate on the ions' phase.

The figures show the flexibility of the LCU from going from a long-sequence, down to detailed short pulses. Another example is given in chapter 5 for a Ramsey experiment.

The LCU has been proven to work well in a variety of experiments, including recent sympathetic cooling experiments and the Ramsey and spin-echo experiments presented later in this

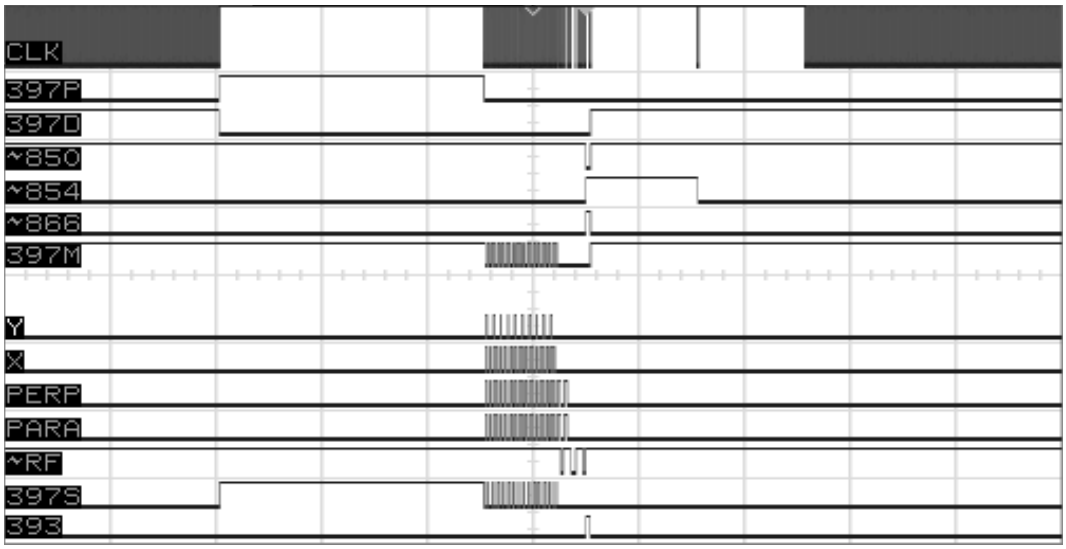

Figure 3.8: Timing Diagram for an Experiment to cool and measure two $^{40}Ca^{+}$ ions. This shows the whole sequence, including the relatively long cooling pulses. The time-scale is 2ms per division.

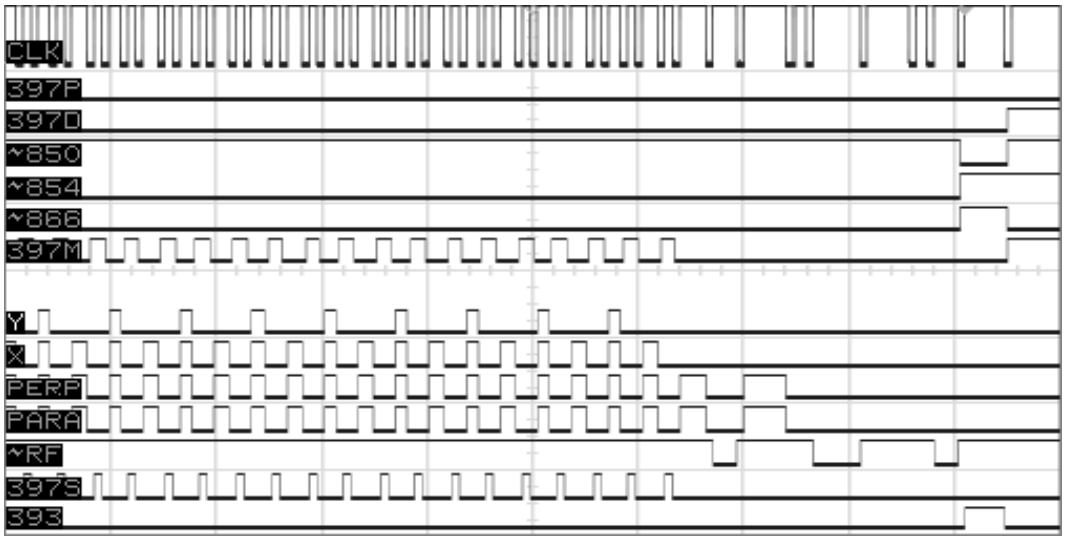

Figure 3.9: A magnified section of figure 3.8 showing the side-band cooling sequence, which consists of much shorter pulses. The Zeeman qubit states are addressed with an RF antenna, labelled RF in this diagram. The time-scale is 200μs per division.

thesis.

The signal-to-noise ratio on the final outputs is better than 10^3. Transition time into an unloaded output is 4ns. Slower transition times are seen on 50Ω loads, but the required switching speed is still achieved for the device. Only one timing error has occurred in three years of sustained use, this corresponds to an estimated error level of 1×10^{-9}, well below that required.

Chapter 4

Master-Slave Laser System

4.1 Introduction

For QIP it is necessary to prepare either the $|\downarrow\rangle$ or $|\uparrow\rangle$ states. As discussed in chapter 2, it is possible to drive coherent transitions between two ground-level hyperfine states, e.g. $|\downarrow\rangle(\mathrm{F}=4,\mathrm{M_F}=0)$ $\leftrightarrow |\uparrow\rangle(\mathrm{F}=3,\mathrm{M_F}=0)$, of $^{43}\mathrm{Ca}^+$ using a Raman transition. This chapter discusses the design and implementation of a laser system for this purpose. It makes use of optical-locking techniques to ensure that the necessary phase-stability and 3.22GHz offset are provided.

4.1.1 Raman Transitions and Spontaneous Scattering

A Raman transition between two levels, $|\downarrow\rangle$ and $|\uparrow\rangle$ involves an intermediate or "virtual" level, that is detuned from an existing atomic level. This is shown in figure 4.1 where the intermediate level is red-detuned from the $\mathrm{P}_{1/2}$ level.

To drive transitions between the ground-level hyperfine states of $^{43}\mathrm{Ca}^+$, the two lasers are far-detuned from the P level but have a difference in frequency of exactly the difference between the chosen $|\downarrow\rangle$ and $|\uparrow\rangle$ states. For the hyperfine states of $^{43}\mathrm{Ca}^+$ this is $\sim$3.23GHz; the exact frequency depending on the choice of $|\downarrow\rangle$ and $|\uparrow\rangle$.

The optimum choice of Raman detuning, Δ, is obtained by considering the spontaneous emission from the virtual level. This needs to be kept to a minimum in order to keep errors of a gate operation to a minimum. Typically 10^{-4} spontaneous photons for every coherent transition would

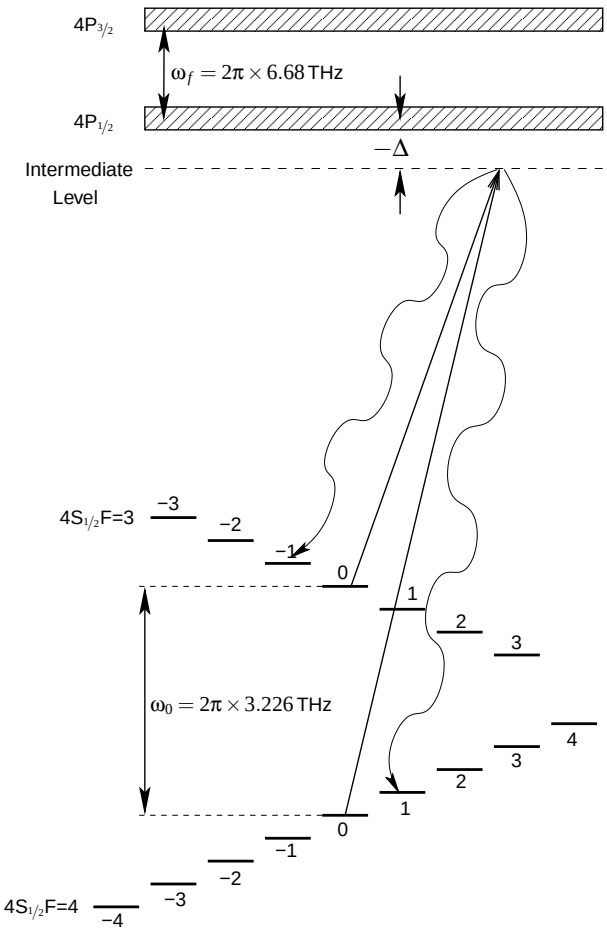

Figure 4.1: A schematic diagram showing the 4S and 4P terms in $^{43}Ca^{+}$ and a virtual level, involved in the Raman transition. The S-level has been expanded into the F and M_F hyperfine states. The P-level hyperfine manifolds are indicated with hatching. The fine-structure splitting is indicated as ω_f. The driven transition shown is for the "clock" qubit states (chapter 2), where ω_0 is the hyperfine splitting. The intermediate level is detuned by Δ from the $P_{1/2}$ level. Spontaneous emission from this level back to the S-level and to the $D_{3/2}$-level (not-shown) is possible.

be required to keep error rates to a level where quantum error correction can be used with success [52]. Additionally the speed of the gate operation needs to be kept high, <1 µs.

Ozeri *et al.* [53] have looked at the optimum detuning for such a Raman transition in trapped ions for both single qubit operations (rotations) and for two qubit gates.

Ozeri shows that where two P-levels exist and the Raman detuning from $P_{1/2}$ is Δ, then the Rabi frequency to drive the clock-transition is:

$$\Omega_R = \frac{g_b g_r}{3} \frac{\omega_f}{\Delta(\Delta - \omega_f)} \tag{4.1}$$

where:

$$g_r = E_0/2\hbar \left\langle \mathrm{P}_{3/2} F = 5, \mathrm{M_F} = 5 \middle| \mathbf{E} \middle| \mathrm{S}_{1/2} F = 3, \mathrm{M_F} = 3 \right\rangle$$

$$g_b = E_0/2\hbar \left\langle \mathrm{P}_{3/2} F = 5, \mathrm{M_F} = 5 \middle| \mathbf{E} \middle| \mathrm{S}_{1/2} F = 4, \mathrm{M_F} = 4 \right\rangle$$

and Δ is the detuning from the lower P-level, $\mathbf{E}$ is the electric dipole operator, and ω_f is the fine-structure splitting between the two P-levels. Equal field-strength (E_0) in both Raman beams and linear polarisation is assumed and also that $g_r \approx g_b = g$.

The rate of total (Raman and Raleigh) spontaneous emission is:

$$P_{\text{total}} = \frac{\pi\gamma}{\omega_f} \left[\frac{2\,\Delta^2 + (\Delta - \omega_f)^2}{\left|\Delta(\Delta - \omega_f)\right|} \right] \tag{4.2}$$

The minimum scattering for a given Rabi frequency can be achieved when:

$$\Delta = (\sqrt{2} - 1)\omega_f$$

where Δ is the Raman detuning, and $\omega_F/2\pi$ is the fine structure splitting.

However, Rayleigh scattering does not collapse the coherence of the ion's state, i.e. a measurement could not determine the state of the ion before emitting the photon and no information is lost. The state can therefore be re-prepared. Raman scattering can be considered to be the only decoherence mechanism from the P-state. Taking the natural line-width of the $\mathrm{P}_{3/2}$ and $\mathrm{P}_{1/2}$ levels to be γ, this scattering rate is shown to be [54] :

$$\Gamma_{\text{raman}} = \frac{4\gamma g^2}{9} \left[\frac{1}{\Delta} - \frac{1}{(\Delta - \omega_f)} \right]^2 \tag{4.3}$$

Therefore, by combining 4.1 and 4.3, it can be shown that the probability to scatter a Raman photon, and therefore cause an error, during the time of a π rotation $t_\pi(\Omega_R t_\pi = \pi/2)$ is:

$$Praman = \varepsilon = \left(\frac{2\pi\gamma}{3}\right)\frac{\omega_f}{\left|\Delta(\Delta-\omega_f)\right|} \tag{4.4}$$

Expression 4.4 is plotted as a function of detuning in figure 4.2.

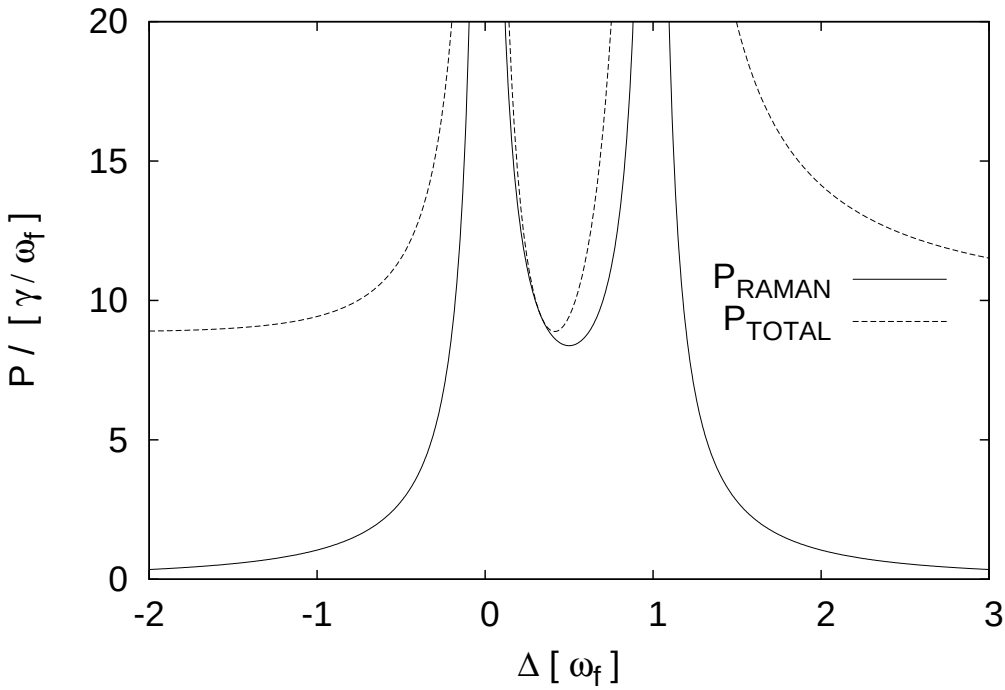

Figure 4.2: Raman scattering (solid line) as a function of detuning (Δ) from $P_{1/2}$, in units of the fine-structure splitting, ω_f. The probability of a scattering event (P_{SE}) is in units of γ/ω_f. For comparison the total scattering (Raman plus Rayleigh) is shown with a dashed line.

It can be seen from figure 4.2 that whilst there is a local minimum at $\Delta = \omega_f/2$, the error rate (or probability of a Raman scattering event) asymptotically approaches very low values for far-detunings on either side of resonance with the two P levels.

The power, $\mathcal{P}$, per Raman beam then becomes a function of the desired error, ε:

$$\mathcal{P} = \frac{2\pi}{3\varepsilon_S}\left(\frac{2\pi w_0}{\lambda}\right)^2 \frac{hc}{\lambda}|\Omega_R| \tag{4.5}$$

where λ and ω are the wavelength and frequency of the resonant transition to the $P_{3/2}$. ε_S is the error rate due to Raman scattering, which is to be chosen. w_0 is the waist-size of the beam at the

ion. This shows that the required power for a given error rate is determined by the cube of the optical transition wavelength, hence ion species with longer wavelengths on the relevant transition (such as calcium) need less laser power (for a given Rabi frequency and error rate). A plot of error rate versus laser power is shown in figure 4.3. There is an additional error due to scattering to the D-levels in calcium, but this is small (due to the low branching ratio) in comparison, and is not shown.

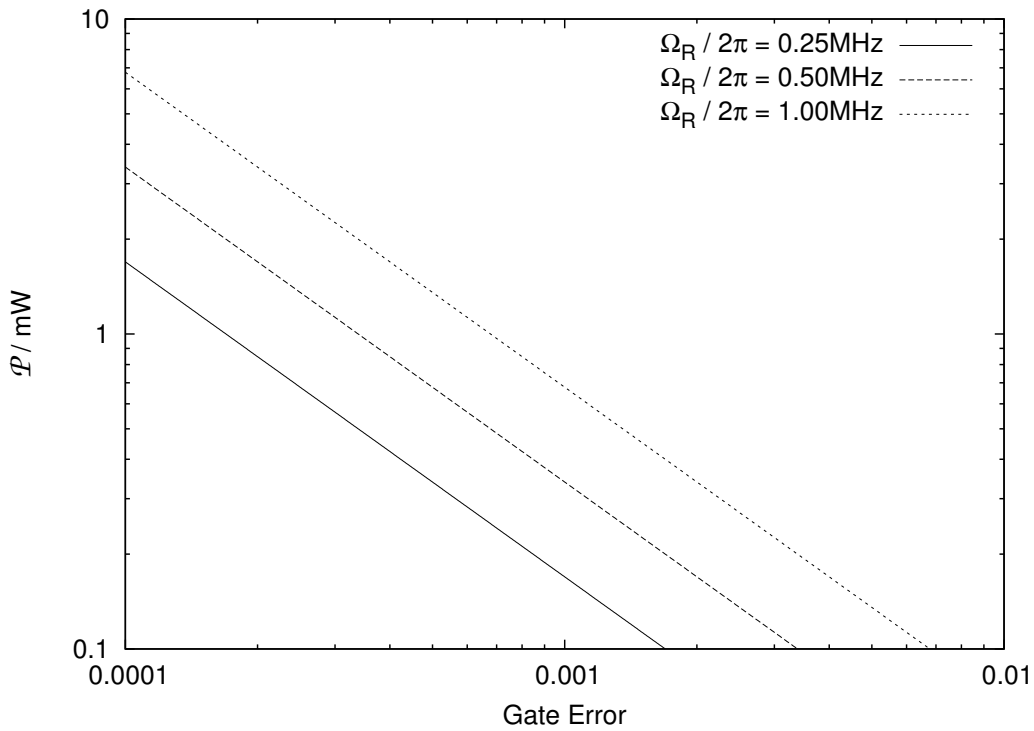

Figure 4.3: Plot of power for a single-qubit operation (π-rotation) versus error rate due to Raman scattering to the S-manifold in calcium for different values of Ω_R. w_0 is chosen to be 20μm as this would be a realisable value in a typical trap.

Figure 4.3 implies that to reduce the error rate, or to increase the single qubit rotation rate for a given detuning, more laser-power is required (the minimum obtainable waist-size is set by the layout of the vacuum system). For this reason, using diode lasers, optical injection techniques are the best solution to provide Raman beams for the experiment. For example, given a beam-waist at the ion of 20μm and a beam-power of 10mW, at a detuning of 2.11 THz below the $P_{1/2}$ level ($\lambda = 398.1$ nm), this would give an error rate of 1.7×10^{-5}.

4.2 Implementation

The Raman laser system is based on semiconductor diode lasers. The required wavelength for the qubit Raman transition, 398nm, is at the very edge of the visible (violet) spectrum. Until recently diode lasers were not available in this region; however, due to a new DVD format which uses 405nm diodes, and due to manufacturing tolerances, some diodes were available in the 395–405nm range. The master-slave setup uses three blue diode lasers. The master is grating stabilised and is used to inject the other two lasers. Whilst being stable, nearly half the master's power is used for feedback and is therefore not available for the experiment. The two slave lasers do not have gratings, but are stabilised by injection of a small fraction of the master power. Maximum power is therefore emitted from the slave lasers as none is used for stabilisation. By tuning the current and operating temperature, the two slave lasers can be locked to the master's wavelength. For a given wavelength it is always possible to find a temperature such that the slave can be injected running close to its maximum operating current, thus providing maximum output power. To provide a difference between the two slaves' frequencies, corresponding to the difference between the two hyperfine levels of $^{43}Ca^{+}$ ($\sim$3.2GHz), a double-pass 1.6GHz AOM was used to shift the frequency of the beam used to inject the second slave laser.

Optical injection techniques have been used before with infrared lasers e.g. for caesium clock experiments [55], but this is the first time it has been achieved using blue-diode lasers and the results are awaiting publication.

The optical injection technique that was employed provides two Raman beams with a frequency difference of 3.2GHz and a measured stability better than 3×10^{-9}. It also allows the slave lasers to be free-running at full power, providing a full 20mW (after the optical isolators) to be available for the experiment. The result is two 20mW beams with stability inherited from the grating stabilised master and a frequency difference of 3.2GHz. The stability of the 3.2GHz offset is set by from the precision microwave synthesiser used to drive the AOM.

4.2.1 Optical Heterodyning of Two Violet Lasers

As a preliminary experiment to the master-slave set-up, the line-width of two independent grating stabilised violet diode lasers was investigated. The master laser (Toptica 1018, nominally 396nm),

and the Doppler cooling laser (Toptica 1016, nominally 397nm) were used. These were set to run at 755 165.5(1)GHz and 755 166.5(1)GHz respectively, measured using a wavemeter (NIST LM10). Both lasers are grating-stabilised in the Littrow configuration (see [46] for details).

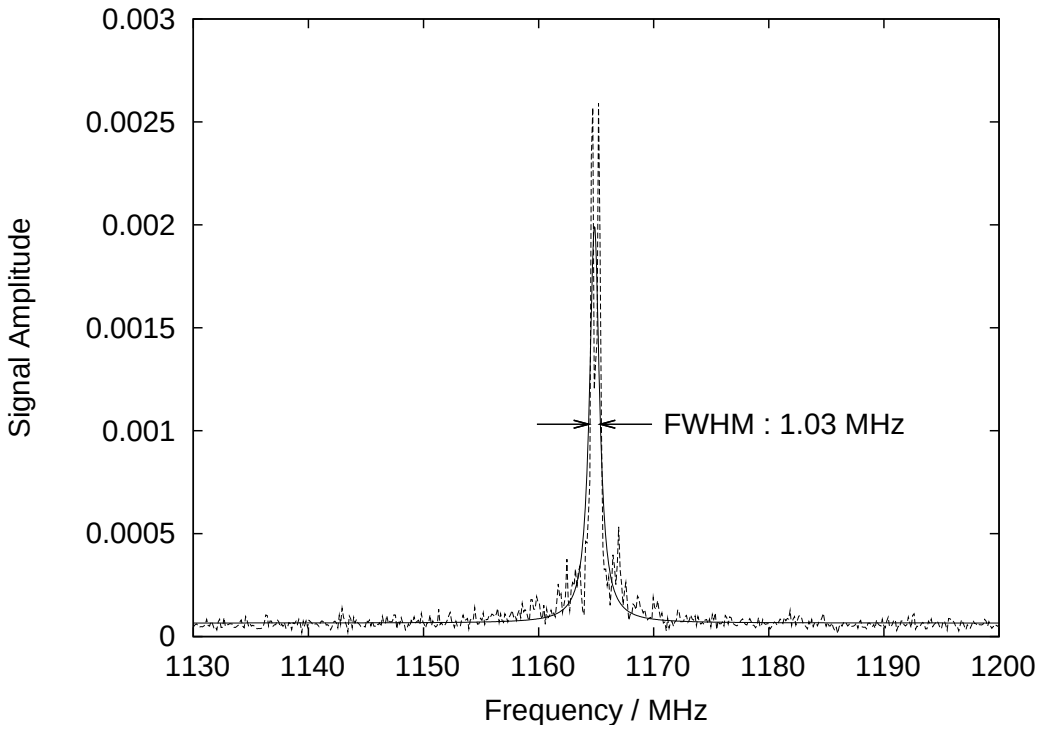

Figure 4.4: Example beat signal for two independent 397nm lasers running at 5.5mW power. The sweep-time was 10ms.

The two beams were combined using a beam splitter. Due to the transmission properties of the beam splitter used, a grazing angle was required to get equal power in the reflected and transmitted beams. The two beams were aligned to co-propagate and then focused onto a fast (25GHz) photodiode (New Focus 1437). The resultant beat signal was observed on an RF frequency analyser (Agilent E4405B), see figure 4.4. The resulting widths are plotted against laser power in figure 4.5. The figure shows the combined laser line-width of two grating stabilised violet diodes with no locking between them. The combined linewidth was 3$\pm$1MHz at a laser power of $\sim$6mW for both lasers measured in a time scan of 10ms, i.e. the analyser scans 7MHz per ms. The width rose to 6$\pm$1MHz at maximum laser power (11.5mW after grating). A full analysis of these results can be found in [56]. This linewidth would make Raman transitions between hyperfine states difficult to drive reliably. Hence the need for offset optical injection locking techniques.

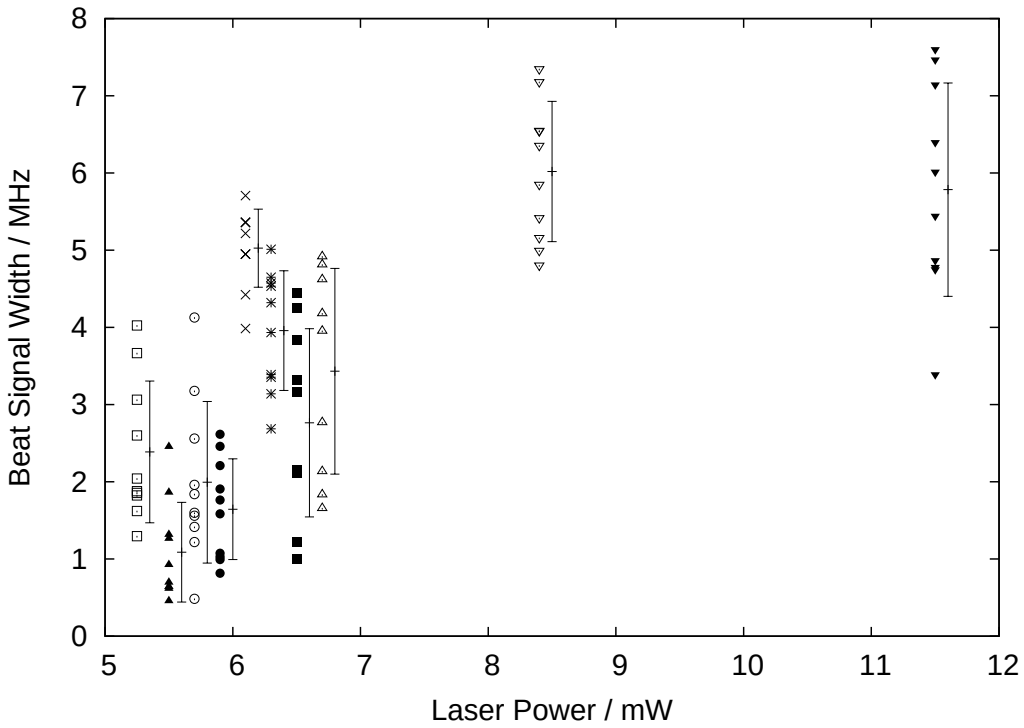

Figure 4.5: Graph of Lorentzian-fitted width of beat signal versus laser power for two unlocked lasers. A different symbol has been used for each power-level. The error bars show the standard deviation at the mean of the sample data and have been offset slightly for clarity. Increasing laser power generally increases the linewidth. Lasing threshold occurs at about 2.5mW

4.2.2 Setup of Optical Injection Scheme

The system is shown in figure 4.6. The lasers each had Faraday rotators to optically isolate the diode from stray reflections back into the laser. These isolators can also be used to inject the slave lasers. A beam leaving the laser is set to pass through the first Polarisation Beam Splitter (PBS). Inside the optically-active medium, due to the strong magnetic field, the polarisation of the beam is rotated through 45°, and exits the second PBS. A reflected beam will only pass the PBS if it is polarised parallel to this beam. This counter-propagating beam will be rotated by 45° in the active medium, and will then be deflected out of the beam path by the PBS nearest to the laser. However, a beam launched into the side port of the second PBS will return through the isolator, if it is polarised orthogonal to the beam leaving the isolator. This is the path that the injection beams take into the slave lasers. By using two mirrors, it was possible to align the injection beam into the slaves successfully. Injection was done without any attempt to change the beam shape. The beam

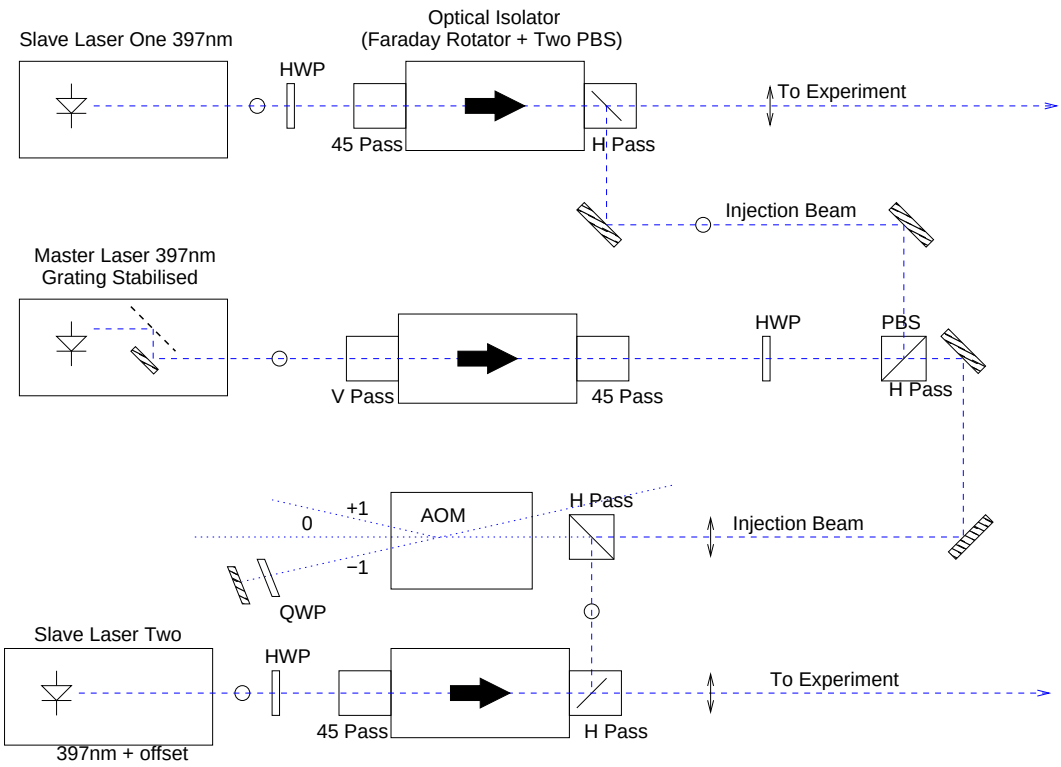

Figure 4.6: Master-slave set-up showing the injection beams into the slaves via the optical isolators. The shifted slave (2) is injected via an AOM driven at 1.6GHz. Polarisations of beams are indicated with arrows for horizontal and circles for vertical. HWP = half-wave plate and QWP = quarter wave plate. Not shown is the lenses required to focus the beam into the AOM and recollimate the beam before reflection.

shape is oval due the semiconductor laser aperture being rectangular. As all three diodes were of nominally identical fabrication, best mode-matching should be achieved by matching similar beams. Successful injection was monitored by scanning the slave's current around the threshold point. With good injection the threshold current drops, and a sharp "step" in output power is observed, see figure 4.7. Without injection, the power varies linearly with current.

This characteristic change is a good diagnostic for successful alignment. Further confirmation was achieved by viewing the injected (slave) beam on an optical spectrum analyser. Scanning the piezo on the master laser's grating caused a corresponding frequency shift on the slave laser.

The second slave was injected via an Acousto Optical Modulator (AOM). The AOM used was based on a crystal of TeO_2 (Brimrose TEF1600-150-395) and had an operating frequency of 1.6GHz $\pm$0.2GHz. By passing the beam twice through the AOM, the injection beam was shifted by a total of 3.2GHz (for a drive frequency of 1.6GHz). An advantage of double-passing the beam

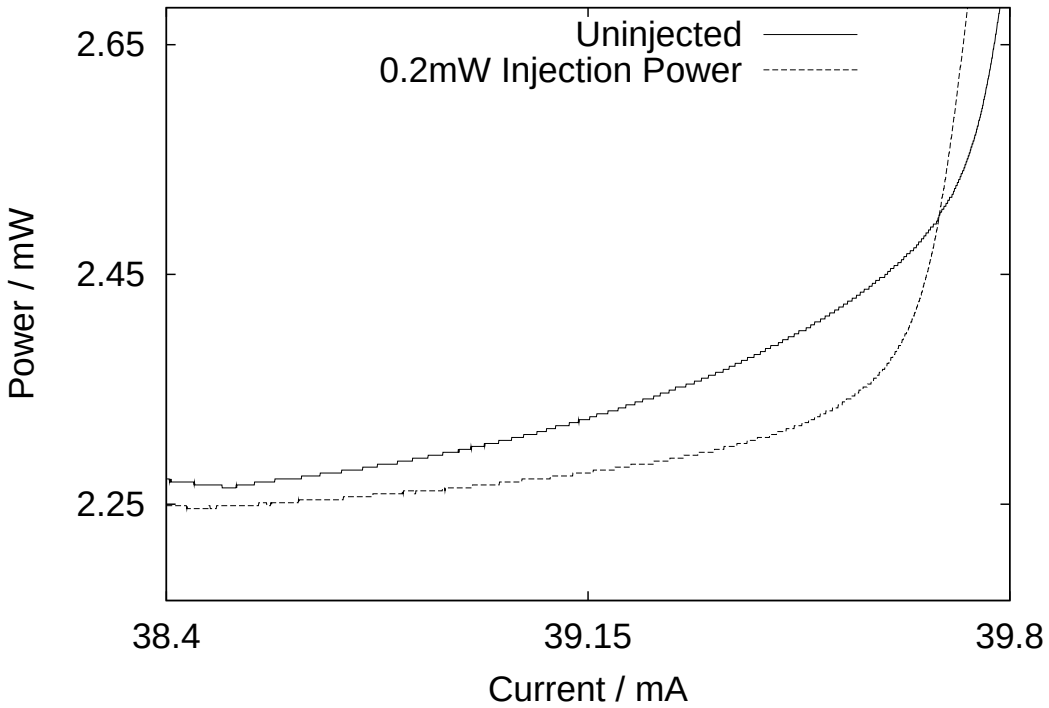

Figure 4.7: Graph of laser power versus input current for a diode laser, showing the effect of injection. The level at which lasing occurs (shown as a sharp increase in power, or "elbow" in the curve) is reduced in the case of injection. This is a useful diagnostic for correct alignment.

is that tuning the AOM drive-frequency does not change the alignment of the beam.

The AOM was found to transmit 25% of the incoming power into the first-order; after two passes 8% of the incident power was available. Starting with a maximum power of 10mW from the master (being grating-stabilised, some of the power was reflected back into the diode), 0.8mW might in principle be available for injection; however, the maximum intensity incident on the AOM was limited by the damage threshold of the device (1W/mm^2). For the beam size used (30 μm by 60μm) the input power was therefore limited to 1mW.

As well as being able to lock to the master frequency, the slave could preferentially lase at its natural modes. The frequency of the mode of a semiconductor laser is a function of two parameters, the length of the internal cavity, and the current at which the diode is running. The length of the cavity is adjusted by changing the temperature of the diode. The required lasing current for stable injection shifts by 5mA per °C. Changing the current changes both the gain of the semiconductor medium and its refractive index. It also causes local heating, therefore a

settling period is required after changing the current in order for the temperature to restabilise.

4.3 Results

4.3.1 Injection without Frequency Shift

Injection "maps" of current versus temperature were obtained for both slave lasers. Neither laser was frequency shifted for this preliminary experiment. The laser was found to lock to the master at a particular current and temperature. There was a small range over which the current could be changed, whilst keeping single-mode injection. The results below are for the minimum current at which stable mode-locking to the master was achieved. There was hysteresis in this current range and it was found that best lasing action was achieved by taking the current beyond the lasing point, and bringing it back to a minimum.

The temperature of the slave was varied between 15 and 30°C. The minimum was 15°C due to the risk of condensation and resulting damage to the laser. The maximum was 30°C as running the laser at higher temperatures was known to shorten its lifetime. The current was varied between 30mA, below which no lasing occurs, and the damage threshold (as set by the manufacturer) of 55mA. Due to the danger of high optical powers burning the laser, output power was also monitored and kept below 30mW. The master wavelength was 397nm (755 168GHz) during these experiments as this gave more power, note however this is not the optimum wavelength for the Raman transition, see above.

Graphs 4.8 and 4.9 show results for slave lasers 1 and 2 under similar conditions. Each set of results forms a straight line, each of which corresponds to one mode of the laser. The injection beam is unshifted, i.e. no AOM was used for either laser whilst obtaining these results.

Power was found to vary approximately linearly with current, at the minimum lasing current [57]. Although both diodes were specified to go to 30mW (20mW after the isolator), Toptica 1019 only produced a maximum power of 10mW. This may be due to the parasitic mode, explained below.

Injection power was varied to find the lowest power at which injection could be obtained. At a temperature of 17.0°C lasing was obtained for as low as 100μW of power, at 397nm. However,

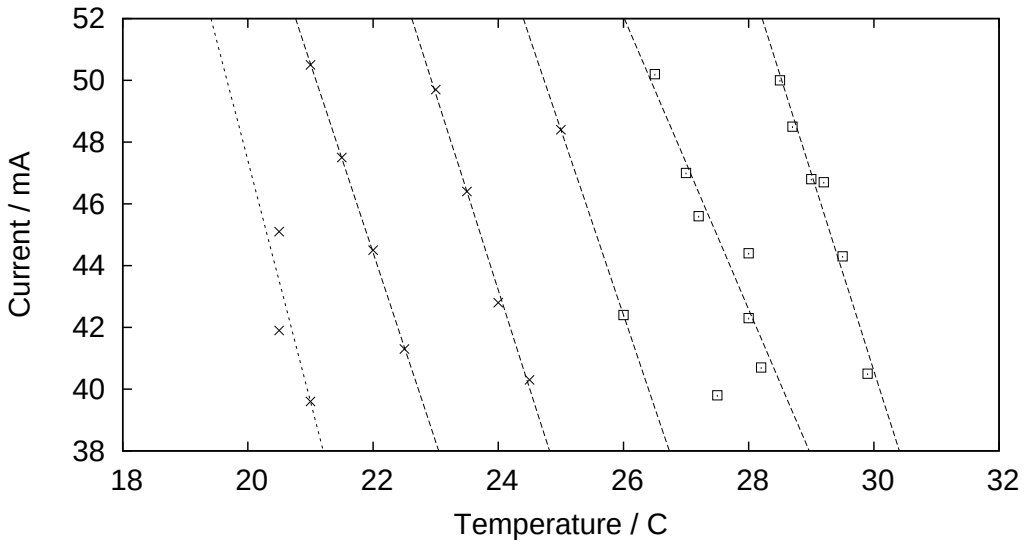

Figure 4.8: Graph of laser threshold current and resultant output power versus temperature for Toptica 1019 (Slave 1) injected at 0.5mW (crosses) and 0.8mW (squares) at the 397nm master wavelength. The power was measured after the Faraday isolator using a glass-pickoff which allowed the power in the beam to be deduced. 52mA corresponds to 12mW output power, this low value is explained in the text.

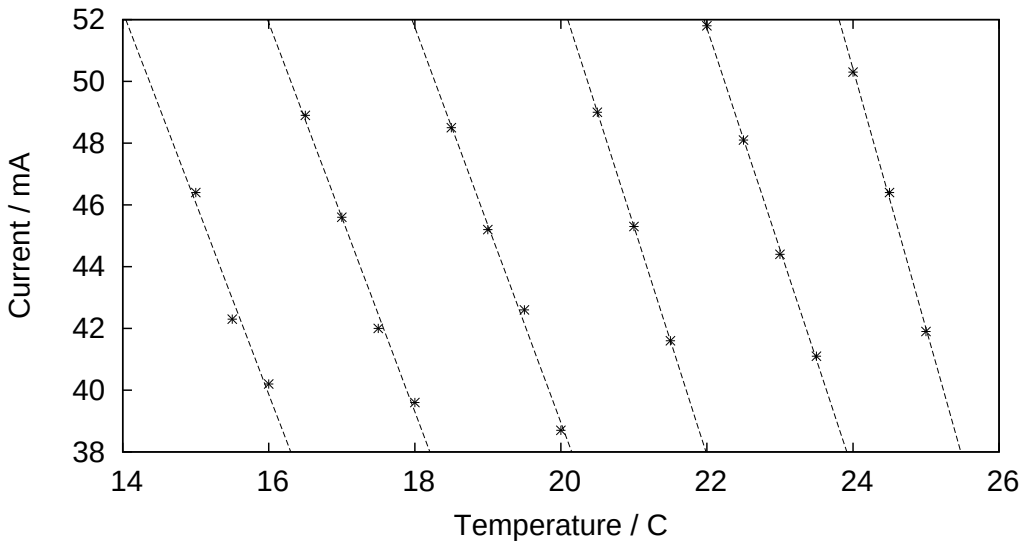

Figure 4.9: Graph of laser threshold current and resultant output power versus temperature for Toptica 1020 (Slave 2) injected at 0.5mW injection power at 397nm master wavelength. The power was measured after the Faraday isolator using a glass-pickoff which allowed the power in the beam to be deduced. 52mA corresponds to about 25mW output power

at other temperatures, whilst injection was achieved, natural modes were also present, which is unsatisfactory for a QIP experiment. Injection was also shown to be achievable with the master

running at 395.5nm for an injection power of 50μW. This shows that large detunings from the lasers nominal wavelength can be achieved, as required for optimum Raman transitions. At 17.4°C and 57.0mA, this gave an output power (after the isolator) of 19.7mW. This was close to the maximum power the slave can run at, which is 30mW before the isolator.

4.3.2 Slave Laser Parasitic Mode

It was observed that a persistent natural longitudinal mode was present in the slave laser, Toptica 1019. This was not observed in the other slave. It was suggested that this is due to a flaw in the structure of the diode, giving rise to a parasitic, longitudinal mode. This was confirmed by the manufacturer's test results.

This will need to be taken into consideration when choosing optimum conditions for the Raman transition and may require operating at a lower power than desired.

4.3.3 Frequency Shifted Injection

The master beam was aligned into the 1.6GHz AOM. A lens with f=250mm was used to reduce the elliptical spot-size to a minimum waist size of 40×53μm. The spot size was measured using a glass-pickoff and a CCD camera (Pulnix PE2015, running in linear response mode) – see figure 4.10. The damage threshold for the AOM is 1W/mm^2 so with this spot-size the maximum incoming power could be $1/0.301 = 3.23$mW. Due to the mode not being TEM$_{00}$ and the possibility of "hot-spots" existing in the mode structure, a maximum incident power of 1mW was used to avoid damage or rapid ageing of the material

Clean injection was achieved using the second-pass return beam from the AOM, which had a power of 80μW. At a temperature of 18.0°C and a current of 47.7mA, a power of 9.0mW was measured after the slave's isolator. A wavemeter was used to confirm the frequency difference between the master and slave was 3.2GHz.

4.3.4 Heterodyning Results

To measure the relative stability (coherence length) of the shifted and unshifted slaves, the heterodyne technique of section 4.2.1 was used. The resultant beat signal of the two superimposed

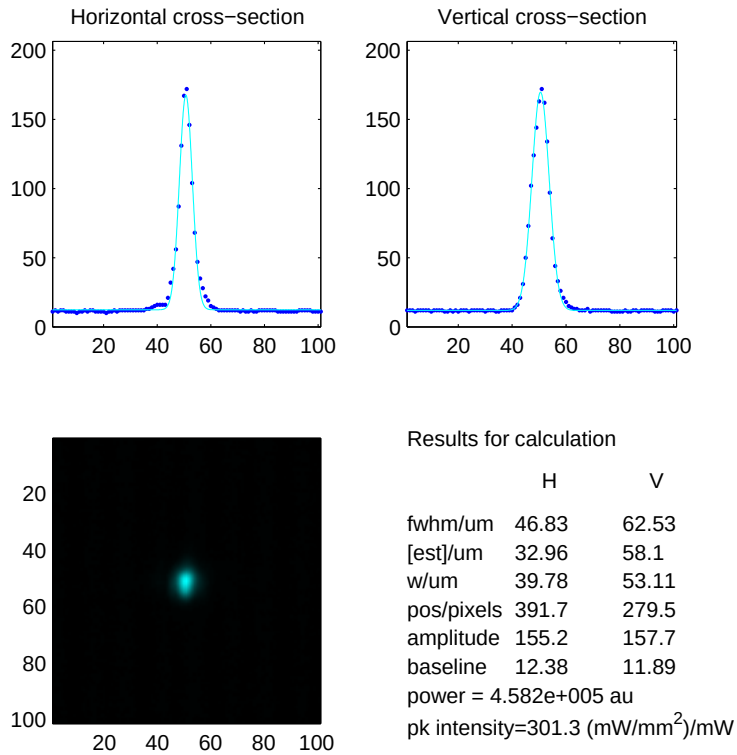

Figure 4.10: CCD image of focused laser beam incident on AOM. The graphs are best-fit Gaussian profiles across the vertical and horizontal axes of the beam. The table shows the fitting parameters.

beams is shown in figure 4.11. Successful injection of the shifted slave was achieved at an input power of 60μW, with a frequency offset of 3.22GHz. This was the lowest power that would result in stable injection at the shifted-frequency though injection was possible at lower powers when unshifted.

The master wavelength was altered by 1.5 nm to reach the required wavelength of 395.5nm, and the experiment was repeated. Both slave lasers were successfully locked to the master, and the frequency offset measured from the beat signal and checked on a wavemeter.

A heterodyne signal with a width of 10Hz at the -3dB points was observed, and this was limited by resolution of the frequency analyser. This remained stable over a period of hours, see figure 4.12, without the master laser being locked to any external frequency reference. Temperature is also plotted as this is varying over time and can affect beam-path stability. This could be improved with temperature control of the lab (air-conditioning). The frequency-offset was tunable via the RF source, and the injection beam remained aligned into the slave, due to the double-pass tech-

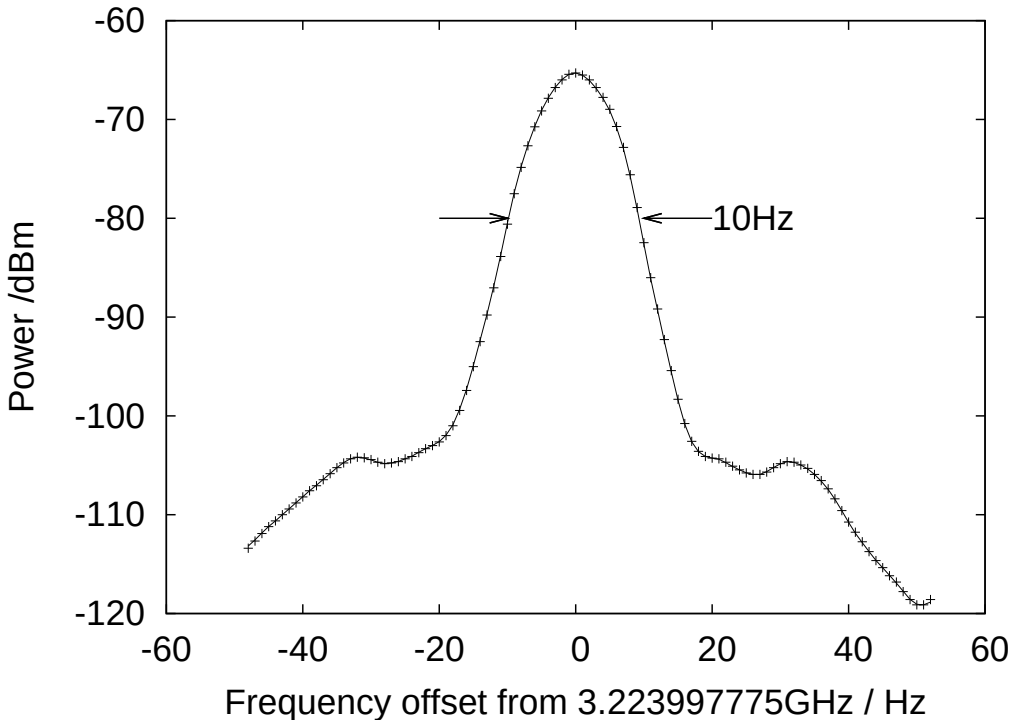

Figure 4.11: The beat signal from two slave lasers both injection locked to the same ECDL running at 397nm. The offset frequency is 3.223997775(1)GHz and the width is 10Hz. The RF source was an Agilent frequency synthesiser (Agilent E8247C). The signal width is as narrow as could be measured by the available RF analyser. There are two harmonics 32dB below the carrier signal.

nique. The absolute frequency was tunable using the external grating of the master laser, provided the slave currents were adjusted accordingly.

4.4 Conclusion

Powers of up to 23.8mW were achieved at injection powers of 50μW for the shifted slave. An output power of 10mW was achieved for a similar injection power from the unshifted slave. Injection was achieved 2.8THz (1.5nm) away from the specified wavelength of the injected laser. The temperature could be adjusted to maximise the output power of the slave for a given wavelength whilst injected. Injection was achieved after a double-pass 1.6GHz AOM setup with 1mW input power before the double-pass, giving the required 3.2GHz offset-lock for $^{43}Ca^{+}$ Raman transitions. This beat-note was measured to be stable in power for more than an hour at a time, without the master being locked to a reference, nor any temperature stabilisation of the lab. The frequency

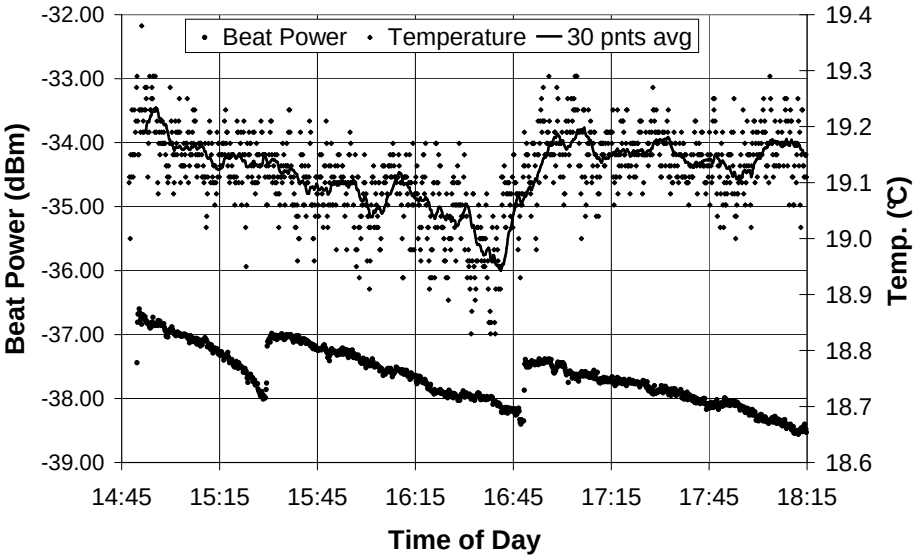

Figure 4.12: Graph of power measured from the beat signal of two injection-locked slave lasers. Also shown is the temperature. The beat-power drops as the alignment changes dues to thermal shift. However, the lasers are staying locked and the difference frequency does not vary in this time-period. At 15:30 and 16:50 the laser mode-hopped and was reset.

stability was only limited by the stability of the RF source driving the AOM ($\lesssim$ 10Hz).

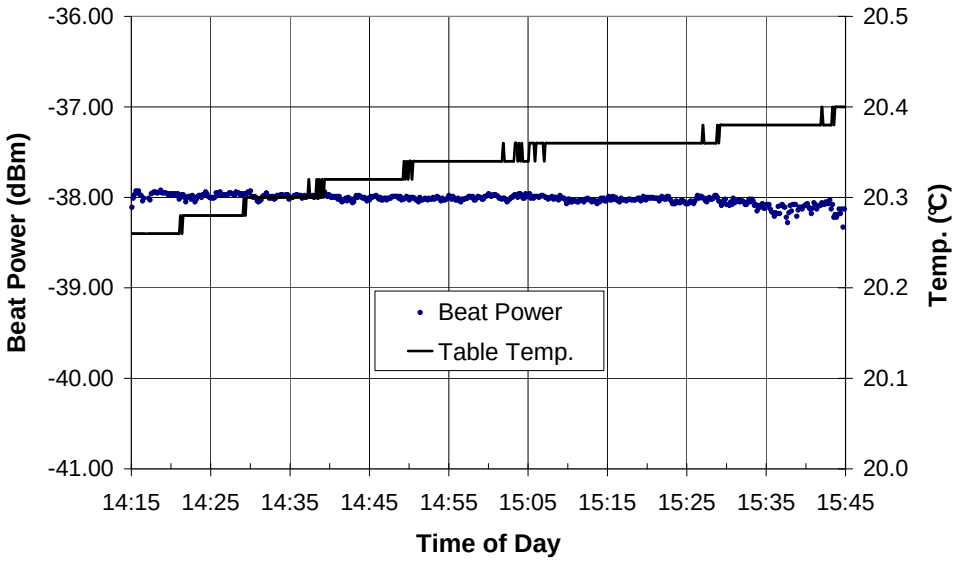

Figure 4.13: A section of results from 4.12 earlier in the day when the laser remained stable for 1.5 hours without any need to relock the laser. The discretisation in the temperature measurement is due to the thermocouple only being precise to 0.1°C.

Chapter 5

Locking of Lasers to a Reference Cavity

5.1 Introduction

For addressing the repumping and cooling transitions in $^{43}Ca^{+}$ it is necessary to have lasers which are stable over the course of the experiment, or sequence of experiments. Typically this could be several hours. The linewidths of these transitions are around 20MHz, so a drift of a few MHz/hour will affect the transition rate and add error to the overall experimental outcome. It is also essential that the laser frequency does not suddenly jump to another mode due to acoustic or mechanical disturbances, as this will cause the experiment to fail. For example, if the 397nm Doppler cooling laser were to jump to the high-frequency side of the cooling transition, heating would occur that might result in the ion escaping from the trap.

This chapter describes stabilisation of lasers using the Pound-Drever-Hall (PDH) technique. The lasers are stabilised against drift over a few hours and made robust against disturbances. No attempt was made to reduce the laser's linewidth. During the generation of the PDH signal, sidebands are generated, which can also be used to address the whole hyperfine manifold of the $D_{3/2} \rightarrow P_{1/2}$ transition, addressed by the 866nm laser [23], see chapter 2 for details.

5.1.1 Pound Drever Hall Theory

Pound-Drever-Hall locking or PDH locking is a technique used to improve a laser's frequency stability. The method uses an external high-finesse cavity (i.e. a Fabry-Perot interferometer) as

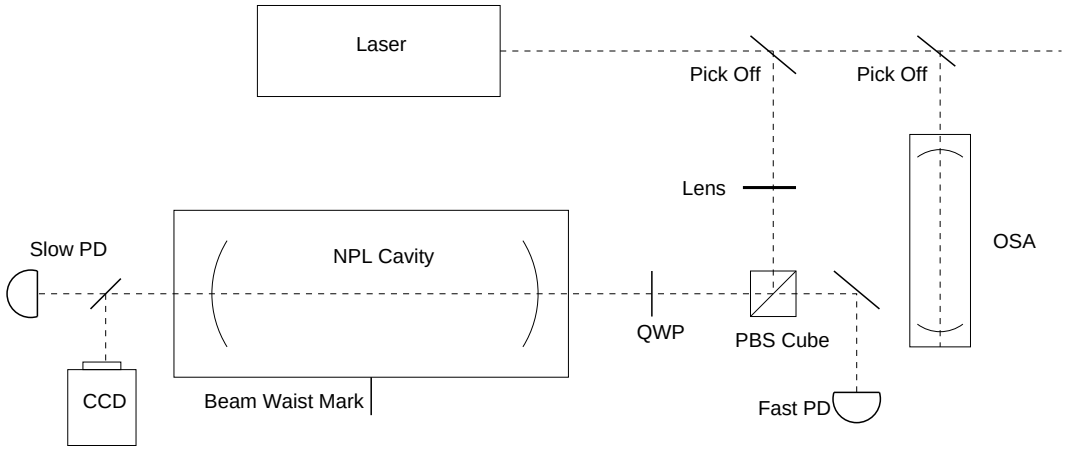

Figure 5.1: Optical arrangement for the PDH lock. The laser was focused to a waist using a lens. Care was taken to ensure that this waist was at the centre of the cavity which was marked on the outside of the cavity housing. A slow photodiode (PD) and a camera were used to help with alignment. An optical spectrum analyser (OSA) was used to ensure the laser is running single-mode and to observe the modulation depth.

the reference, and an electronic servo to control the laser. It was developed by Drever and Hall and is based on the Pound stabiliser [58] – originally used to stabilise microwave cavities. Unlike side-of-fringe locking it is not sensitive to intensity noise and it can have a larger capture range.

An optical cavity has a maximum transmission when on resonance with the input source (see figure 5.3b). The side-of-fringe locking technique is a simple system that uses the transmitted light of the cavity, and locks just off-resonant with the cavity's mode. If the laser is locked to the low-frequency side of the fringe and the frequency of the laser increases (e.g. due to thermal drift), it will come closer to resonance with the cavity. The error signal will increase. If the laser frequency drops, it will move further away from resonance and the error signal will decrease. This error signal can be fed back to the laser controller and used to return the laser's frequency back to the lock-point. There are two problems with this scheme though. Firstly, the capture range is small (comparable to the fringe width of the cavity). It is very easy for a sudden disturbance to change the laser's frequency so far from the lock point that it is on the other side of the reflection peak, and the above situation is reversed. This is sometimes called "anti-lock", and the laser will be rapidly driven to another cavity mode. Secondly, a side-of-fringe lock will respond to intensity noise as well as frequency noise on the laser. This can be corrected for with intensity stabilisation schemes, but it is not ideal.

In PDH locking, the laser is phase-modulated at a frequency, f_m, either directly by the input current (when side-bands are required on the experimental beam), or by using an Electro-Optical Modulator (EOM) when the side-bands are only to be added to the beam used for the lock. Either technique introduces two side-bands onto the carrier frequency of the laser. If the laser's current is modulated directly, sidebands with both a different phase and the same phase are produced. The resultant phase-difference can be seen as a small difference between the power of the two side-bands as seen on an optical spectrum analyser.

Phase-modulation is a form of angular modulation and can be expressed as a Fourier series of Bessel functions (J_n):

$$\sin(\omega_c t + \beta \sin(\omega_m t)) = \sum_{-\infty}^{+\infty} (-1)^n J_n(\beta) \sin(\omega_c t + n\omega_m t)$$

for single-frequency modulation, where the modulation function is $\sin(\omega_m t)$, β is the ratio of the side-band power to the carrier, called the modulation index, and ω_c is the laser's principle frequency (the carrier).

If the modulation index is small ($\beta < 1$) this is equivalent to amplitude modulation [60] – if just the first two terms in the series are taken, and the two side bands, $\omega_c - \omega_m$ and $\omega_c + \omega_m$, have a 180° phase relationship between them. For the purposes of PDH-locking, the only requirement is that two signals either side of the carrier exist, and that they have a distinct phase relationship.

The reflected signal from the cavity is given by [59]:

$$F(\delta) = \frac{r\left(1 - e^{\left(-i\frac{\delta}{\Delta f}\right)}\right)}{1 - r^2 e^{\left(-i\frac{\delta}{\Delta f}\right)}}$$

where δ is the detuning from the cavity's resonance, Δf is the free-spectral range of the cavity and r is the reflectivity of the mirrors (assuming they are equal).

The phase-modulated signal can be expanded using Bessel functions [60]; however, it is noticed that especially for modulation of the laser-current, the sidebands are not of equal amplitude, therefore we replace J_0 with a, and J_1 with b and c respectively,

$$E = E_0 e^{i(\omega_c t + \beta \sin(\omega_m t))}$$

$$E \approx E_0 e^{i\omega_c t}\left(a + be^{i\omega_m t} - ce^{-i\omega_m t}\right) \tag{5.1}$$

When E is incident on the cavity, the reflected signal from the cavity is:

$$E_r = E_0 e^{i\omega_c t}\left(aF(\delta) + bF(\delta+\omega_m)e^{i\omega_m t} - cF(\delta-\omega_m)e^{-i\omega_m t}\right) \tag{5.2}$$

To be able to separate the input signal from the return signal, a form of isolator is used, which makes use of the orthogonal polarisation to direct the input beam on a different path to the output beam. This is done by passing the beam through a polarisation beam splitting cube. This will pass horizontally polarised light into the cavity. The reflected light is rotated by using a quarter-wave-plate. This vertically polarised light returns to a fast photodetector. Good optical isolation is still required to prevent stray reflected light which would cause optical feed-back to the laser and destabilise it; hence the need for the Faraday isolator.

The power received at the photodiode is:

$$\begin{aligned} |E_r|^2 &= E_0^2\left(a^2|F|^2 + b^2|F_u|^2 + c^2|F_l|^2\right) \\ &\quad + aE_0^2 F(bF_u^* e^{-i\omega_m t} - cF_l^* e^{i\omega_m t}) + aE_0^2 F^*(bF_u e^{i\omega_m t} - cF_l e^{-i\omega_m t}) \\ &\quad + 2\omega_m \text{ terms} \end{aligned} \tag{5.3}$$

where $a^2E_0^2$ is the power in the carrier, $F = F(\delta)$, $F_l = F(\delta-\omega_m)$, $F_u = F(\delta+\omega_m)$, and δ is the frequency detuning from the cavity's resonance.

This consists of DC terms, ω_m terms and higher frequency terms. To recover the terms of interest, i.e. the terms containing F_u and F_l, the signal is multiplied by the original modulating signal, $\cos(\omega_m t)$. This results in the error signal, $P_\varepsilon = E_r \times \cos(\omega_m t)$. This has the effect of shifting the terms containing F_u and F_l to DC and the other terms are shifted to frequencies at or above ω_m. These high-frequency signals are rejected by the electronic circuit, leaving a signal involving a convolution of the cavity's response to the carrier and the upper and lower sidebands. After mixing and filtering, the resultant signal is,

$$P_\varepsilon = E_0^2\left[(abF^*F_u - acFF_l^*)e^{i\omega_m t} + (abFF_u^* - acF_lF^*)e^{-i\omega_m t}\right] \times \cos(\omega_m t)$$

To see the frequency dependence of this term, the following identity is used:

$$2\left[\Re\left\{Z\right\}\cos(x)+\Im\left\{Z\right\}\cos(x)\right]\equiv Ze^{ix}+Z^{*}e^{-ix}$$

where Z is complex and x real.

This gives:

$$\begin{aligned} P_{\varepsilon} &= E_0^2 a\left[\Re\left\{bFF_u^{*}-cF^{*}F_l\right\}\cos(\phi)-\Im\left\{bFF_u^{*}-cF^{*}F_l\right\}\sin(\phi)\right] && (5.4)\\ &\quad +(2\omega_m \text{ and higher terms}) && (5.5)\end{aligned}$$

where ϕ is the phase-difference between the two inputs at the mixer. This can be tuned to give a maximum signal by varying the phase-delay between the photodiode and the mixer. The expected error signal from equation 5.5 is plotted in figure 5.2. To see the effect of varying phase and sideband height, see appendix B.

The advantage of the PDH system is that it uses the reflected signal on resonance. The phase of the signal provides an error signal with a sign that depends on whether the laser is above or below resonance with the etalon. That way the laser's frequency can be adjusted in the correct direction.

To understand equation 5.5, suppose the laser is just higher than resonance. The upper sideband will have a lower amplitude than the lower sideband. The received signal would be the same as the case when the laser is lower than resonance, if it were not for the phase-difference between the two sidebands.

The geometry of this error signal is predomintly determined by two things: the fringe width of the cavity which is defined by: $\Gamma = \Delta f/\mathcal{F}$, and the modulation frequency. The width of each peak is determined by a combination of the two. If the modulation frequency is too high, the two "dips" in the curve increase. If these reach zero, the system may lock to them instead of the centre lock-point. The slope in the centre determines the maximum lock-response and is determined by the linewidth of the cavity and the power in the sidebands. Therefore, a careful choice of these values is required. See [61] and [62] for more details.

The etalons used were from NPL and had the following specifications [65]:

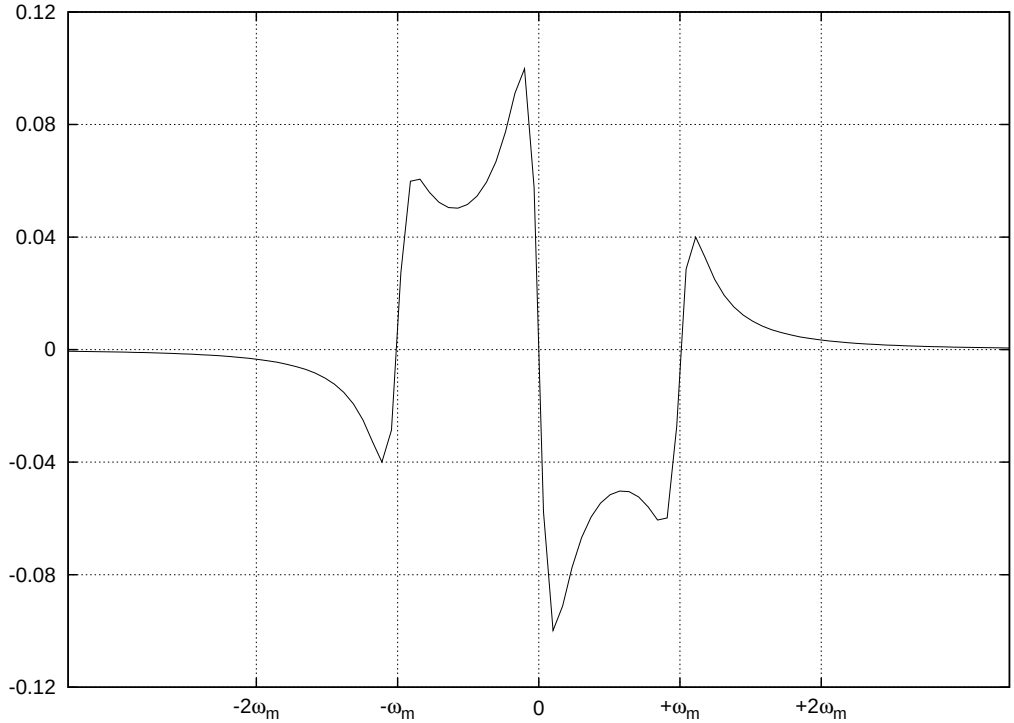

Figure 5.2: A theoretical PDH error signal for a cavity with $\mathcal{F} = 50$, $\Delta f = 1620MHz$. The sideband height is 10dB below the carrier and the modulation frequency is $150MHz$.

1. Wavelength range: 630nm – 1000nm

2. Low frequency drift: $<$ 500kHz / hour (measured to be 270kHz / hour)

3. Scan range: $\lesssim$ 10GHz

4. Free spectral range: 1620 MHz

5. Finesse: ~ 30 ($R \approx 0.90$)

6. The system is evacuated to eliminate pressure effects to a residual pressure $< 10^{-6}$mbar

7. The system temperature is controlled to minimise thermal effects to within 5 mK

The slope at the centre of the signal is given by [63]:

$$D = -\frac{8ab}{\Gamma} \ \ (\mu\text{W/MHz}) \tag{5.6}$$

assuming $\phi = \pi/2$ and $c = b$. This determines the response-speed of the lock. However, in this

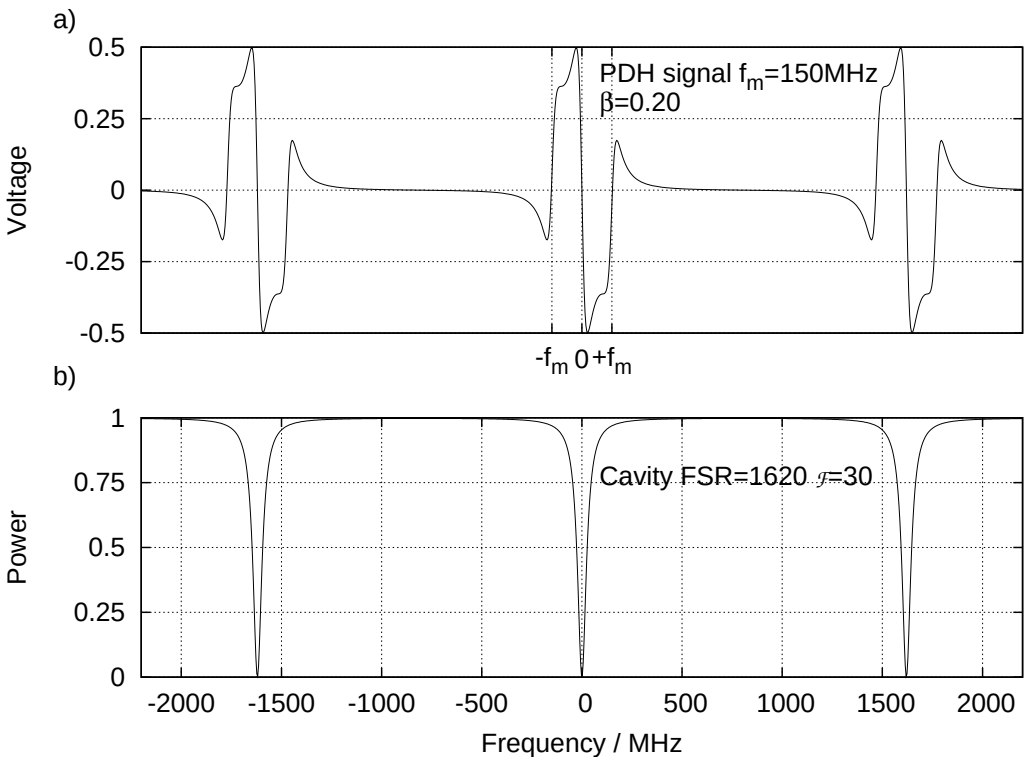

Figure 5.3: a. A theoretical PDH error signal for a cavity with $\mathcal{F} = 50$, $\Delta f = 1620$MHz, side-band height 10dB below carrier and modulation frequency 150MHz. b. The theoretical reflection function for the same cavity.

setup, there is no fast feedback path, and therefore the speed of the lock is determined by the bandwidth of the servo control.

The capture range of the lock is predominately determined by $\omega_m = 2\pi f_m$. For a robust lock this needs to be as high as possible, typically $\Delta f/5$. For the infrared diodes this is determined by the hyperfine splitting in $^{43}Ca^+$. For the violet diodes it is determined by the frequency response of the available EOM.

The modulation depth will change the performance of the lock [61], as will the symmetry of the side-bands and the phase-shift between the two inputs to the mixer. Maximum slope at the centre of the error signal occurs when the modulation is at 42%. Asymmmetric side-bands increase the capture range as long as ϕ is chosen to maximise the signal (see appendix B).

The RF circuit used to recover the error signal is essentially a superheterodyne receiver, as used in high-quality AM radio receivers. See figure 5.8 for a block diagram. It works by having a local oscillator that is running at a given carrier frequency which is mixed with the incoming signal

given by equation 5.5. The mixer generates new frequencies (hence the term heterodyning), which are the product of the two incoming signals. In radio-receivers this is labelled the Intermediate Frequency, as it contains the sums and the differences of the carrier and the signal frequencies. Due to the phase difference introduced by the optical path length and cabling, ϕ may not be at the optimum value to give a good PDH signal (see appendix B). Therefore, it is necessary to adjust the phase difference, or delay, between the photodetector and the mixer. Commercial phase-adjusters were looked at, but were too expensive for this application. Therefore, phase-adjustment was done by simply varying the length of the coaxial cable used. The local oscillator used to drive the laser sidebands is also used at the input to the mixer, resulting in the signal being very phase-stable.

The error signal is fed into a PID circuit (Toptica Module PID110). A PID circuit uses a sum of the Proportional, Integral and Differential of the error signal to control the laser and complete the feedback look. The proportional and differential parts allow for high-speed control, whilst the integral part gives a time-averaged signal. Typically, the three parts are mixed together to allow for the best possible noise and drift cancellation, without causing the servo-loop to oscillate.

In practice, only the integral part of the PID was used to drive the piezo. This was done by setting the proportional and differential contributions to zero. The signal is connected internally within the Toptica system to the piezo controller (SC110) which varies the length of the laser's external (grating) cavity. The maximum bandwidth is determined by the response of the piezo, which depends on several things, including mechanical resonances and the mass of the mirror (see [64] for a full discussion). It is typically in the order of $\lesssim$10kHz. To ensure the laser remains single-mode it is also necessary to adjust the current to keep the laser's gain medium lasing at the adjusted frequency. A signal is "fed-forward" from the piezo-controller to the current controller to ensure this happens.

5.2 Transimpedance Amplifier for Photodiode for use in Lock-circuits

5.2.1 Specifications

The purpose of building this circuit is to be able to detect the reflected etalon signal given by equation 5.5.

It must be able to detect a low power ($\sim$10μW) AM signal at $50-200$MHz against a constant DC background of light, and to amplify it to a sufficient voltage to drive the PDH lock circuit.

Equation 5.6 puts a limit on the total noise budget of the system. The amplitude of the PDH signal is determined by the slope of the error signal around zero, which for a given beam power of 100 μW and the etalons used, gives a response of the PDH of 7.8 μW/MHz. For a change in 1MHz of the laser frequency, the signal on the photodiode (P_ε) will change by 7.8μW on a background signal of 120 μW (carrier + sidebands). The maximum noise that the system can tolerate depends on the open-loop gain of the system including the PID circuit and the HV amplifier which drives the piezo. The contribution of noise from the electronics needs to be much less than the frequency jitter on the laser, which is around 1MHz. Therefore the electronics must contribute $\lesssim$ 100kHz. With the set-up above this corresponds to an overall SNR $\gtrsim 10^3$. The piezo bandwidth limits the noise to $\sim$ 1 kHz.

For the infrared diodes the frequency of the sidebands was 160MHz, so that the hyperfine sub-levels of the $D_{3/2}$ in $^{43}Ca^+$ can be simultaneously addressed. For the violet diodes where the sideband is only used for locking and an EOM must be used, the frequency was between 50 and 85MHz depending on the EOM used. Therefore, the circuit needed to operate between 50 and 200MHz.

5.2.2 Photodiode

For the infrared laser system, the S3883 photodiode from Hamamatsu was chosen because it has a large surface area (a circle of diameter 1.5mm), and its optimum response is in the 850-870nm range. It is a PIN diode which works by having a neutrally doped region called the Intrinsic layer between the P- and N- doped layers. This increases the response speed of the diode. PIN diodes can be operated in two ways. When light is incident on the junction, a voltage is generated which can be amplified, but this is not very fast due to the intrinsic capacitance of the junction. This allows fast voltages to sink to ground. If the diode is reversed biased, the PIN junction acts like a resistor, whose resistance changes with light-intensity. The junction capacitance is reduced by a high reverse bias voltage. For a fixed bias-voltage, the diode generates a current proportional to the light intensity incident on it. The S3883 can withstand a reverse-voltage of 30V, and has a

cut-off frequency of 300MHz. At a reverse voltage of 20V, the capacitance is 6pF. At 860nm the responsitivity is 0.6 A/W. The diode has an area of 1.7mm^2, making alignment very simple.

The other common high-speed photodiode used for low-light levels is an avalanche photodiode, which works by having a very high-field across the junction. When a photon is incident on the junction, a charge carrier (in this case an electron) is released which will cause an avalanche of more carriers, producing a sizable current. Just like a photo-multiplier tube, a large number of electrons are released for each event, but the actual statistics of the detection means that not every photon will generate a current. The signal-to-noise is very good at low-level signals. However, a high-field in the order of a few hundred volts across the junction is required. An avalanche photodiode system provided by NPL was evaluated. It consisted of a diode, a Minicircuits amplifier, and a high-voltage DC supply created by rectifying and filtering the mains (240V) supply. The results were good, but the high-price of the system was not justified, as the signal levels that are used in the PDH system are sufficient to get a good SNR from a cheaper PIN system. However, to achieve this, a PIN photodiode amplifier needed to be designed and built.

For the violet laser systems a commercial APD package from Hamamatsu (C5331) was used. This contains the necessary voltage up-converter to generate the high negative bias, and the first-stage amplification, both running off a single 5V supply. It is cost-effective and easy to use but still more expensive than a home-built system.

5.2.3 Amplifier Circuits

In order to use the S3883 photodiode in a circuit, first-stage amplification was necessary. Five different amplifiers were looked at but in the end a hybrid design was implemented. The five different amplifiers are presented below, along with the final design. The selected photodiode produces about 0.5A/W at the required wavelength of 866nm. For a signal power of $\sim 10\,\mu$W, this gives a photo-current of $\sim 5\,\mu$A. This must be able to drive a 50Ω load to be compatible with standard RF amplifiers and mixers. Therefore, the first-stage amplifier needs a gain of 20dB (a transimpedance of 100). This gives a voltage of 0.5mV (-53dBm) to be connected to the second-stage amplifier and this drives the mixer circuit for the PDH system. Given the background DC signal on the photodiode, a dynamic range of 30dB would be required. To avoid this problem, DC

is blocked from entering the circuit.

Transimpedance Amplifier

The normal circuit for amplifying a low-frequency photodiode signal is the transimpedance circuit shown in figure 5.4. This is a current-to-voltage converter, as a photodiode produces a current proportional to the light-power incident on its surface.

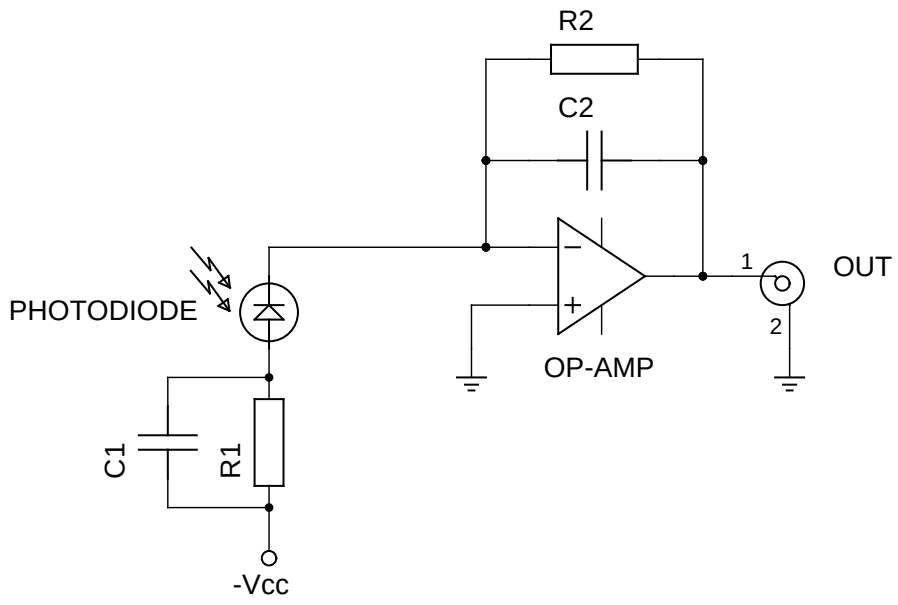

Figure 5.4: Transimpedance amplifier. R2 sets the gain. C2 provides a "speed-up" of the amplifier as discussed in the text. R1 and C1 provide a current limit to the photodiode. Output is via coaxial cable.

This circuit requires an op-amp that is a voltage-amplifying device. Because the input is an effective ground-point, the photo-current is flowing into an almost perfect high-impedance front-end. Negative feedback via R2 and C2 is used to stabilise the amplifier and set the gain.

A bias-capacitor, C1, is added in parallel with the biasing resistor, R1, to ensure that the bias current is not affected by the signal at the photodiode.

Speed-up Capacitor A high-speed amplifier introduces a phase-shift between the input and output terminals. If the output was fed back into the input, out of phase with the signal, oscillation would occur. To remove this oscillation, a capacitor (C2) is introduced in the feedback circuit. At high frequencies the impedance of the capacitor is much lower than the feedback resistor (R2) and

so dominates. The capacitor is chosen such that at its own phase-shift balances the phase-shift of the op-amp. However, op-amps rarely have linear phase-shift over their full working range, and it is necessary to compromise with the best value for the working frequencies [66].

Drawbacks The problem with a transimpedance amplifier at this speed is the gain-bandwidth product (GBP) of the op-amp required is outside the range of available components. The fastest op-amps available have a GBP of about 1600MHz. At 160MHz this gives a gain of 10. This is not large enough for the application, and it would require operating the amplifier at marginal stability (the phase shift may cause unwanted oscillation, even with a speed-up capacitor).

Photodiode Amplifier using a Current Amplifier

Faster op-amps are available; however, they use current-mode amplification. This is not suitable for amplifying a small photo-diode current, as they require a large input current to work, and have low-impedance front ends (they are generally used for video signals).

Discrete Component Amplifiers

Bipolar Transistors A bipolar transistor is effectively a current amplifying device. It is sometimes known as a transresistance amplifier, as it can convert current to voltage, in the same way a transimpedance amplifier does. The ratio of input to output current is characterised by $\beta = {}^{I_{\text{in}}}/_{I_{\text{out}}}$

High-beta transistors ("super-betas") are available. However, the beta value is highly sensitive to thermal fluctuations, and the gain can be highly unstable. Other problems with single transistors include junction capacitance which creates phase-shift (Miller effect), and a fall-off in gain at high frequencies. These can be compensated using other components, e.g. the cascode arrangement below. However, because of the tolerance on single transistors, and thermal noise and effects, discrete components tend to give problems which ICs often solve.

To avoid some of the problems with discrete transistors it is possible to buy an IC containing six transistors on one substrate. These are all in the same thermal environment, and the cascode arrangement can be built using them with improved thermal characteristics. However, all the transistors have the same specifications, and it is not possible, for example, to buy super-beta transistors in one IC.

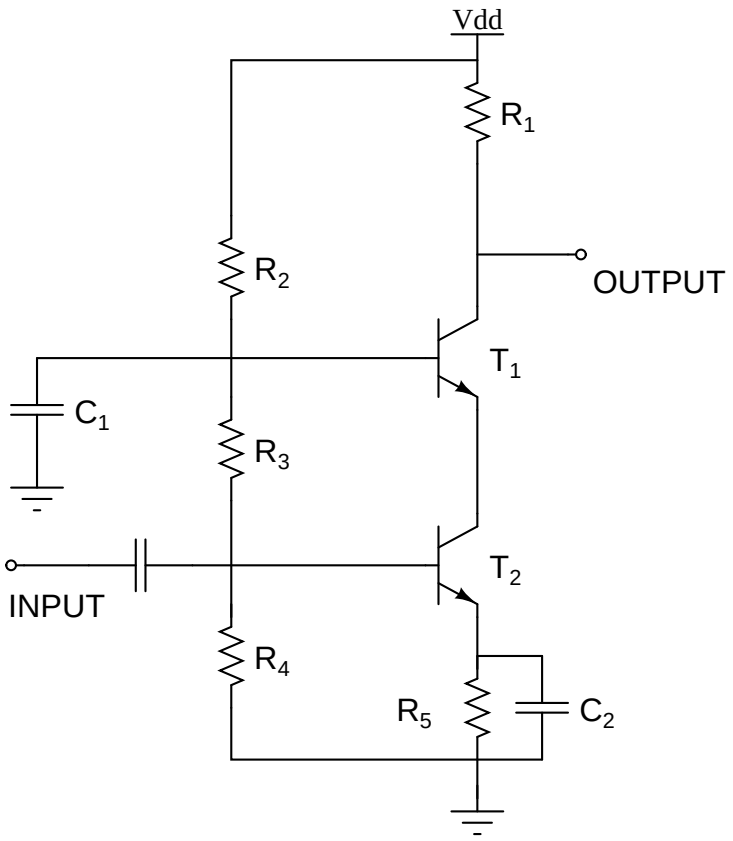

Figure 5.5: Cascode amplifier showing T_1 which compensates for Miller capacitance. Stability is achieved at high-frequencies from capacitors C_1 and C_2. The capacitor on the input is the DC-blocking capacitor. R_2, R_3 and R_4 provide biasing.

Cascode Arrangement The cascode arrangement is used to compensate for Miller capacitance. The Miller effect is due to the junction capacitance between the collector and base. Due to the voltage changing on both sides of the junction, the capacitance appears larger than it would be at DC. This capacitance is in the feedback loop, and therefore a small junction capacitance is amplified, causing the gain to reduce at high-frequency as $1/f$.

The cascode circuit shown in figure 5.5 consists of a common-emitter amplifier driving a common-base stage. This is called a cascode circuit as it resembles the cascaded-anode circuits used with vacuum tubes [71].

Several photodiode amplifiers e.g. [67] have been designed around this arrangement, but these are normally only good to 20MHz or so. This is because the cascode arrangement adds noise to the circuit, due to noise sources in the transistor [68]. These can be overcome at low-speeds, but

are problematic above 20MHz.

The cascode circuit gives a good transimpedance amplification, but not much gain. Super-beta transistors can be used to increase gain as discussed above, but it is more usual to use a second amplification stage, such as an FET, or high-impedance front-end RF amplifier.

FET Amplifiers FETs are essentially voltage amplifiers and are therefore not normally considered suitable for amplifying photo-currents. However, at lower speeds, circuits have been devised that use an FET front-end [69].

Specialist Integrated Circuits

The very popular OPA128 is a transimpedance amplifier using a cascode front end. This is the IC used in the TW board, which is in use in the calcium 40 experiment. However, it is bandwidth limited to <1MHz.

Hamamatsu recommend OPA648/658 in their application notes [70]. Unfortunately this is obsolete. Texas Instrument's fastest photodiode-compatible operational amplifiers are the OPA657 and OPA686 (which is now being replaced by the OPA846). These were not available to buy in the UK.

These and other specialist chips exist which are specified to operate in transimpedance mode up to 200MHz. These devices would be the best option, as anything that can be done in discrete components can be done much better on a single substrate, with extra compensating components and fewer inter-connection problems. However, the cost of obtaining small quantities of these ICs is prohibitively high.

5.2.4 Final Circuit – KPD110

The final circuit is shown in figure 5.6. The circuit consists of an FET follower and, for greater gain, a second stage of amplification was added. This was an RF amp (Minicircuits MAN-1LN) which has a gain of 20dB and can operate up to 500MHz. However, it has matched 50Ω inputs and outputs and therefore the first-stage needs to be able to drive a 50Ω load. The same technique is used in the NPL avalanche detector. An avalanche diode produces sufficient current to drive a MAN-1LN directly. However, a PIN diode cannot drive such a low impedance load directly. The

FET follower provides a high-impedance load on the photodiode whilst being able to drive the 50Ω input impedance of the RF amplifier. A high reverse bias is generated from a 24V source. The FET was biased in its linear region and a DC blocking capacitor used to reduce the required dynamic range. A second DC blocking capacitor between the two amplifier stages is used to obtain best performance from the Minicircuits module. The capacitors in the circuit were chosen for their good RF characteristics and stability. Tantalum capacitors were used for voltage regulation, along with larger electrolytic types. A ceramic capacitor (see chapter 4) was favoured for the DC blocking capacitor as it is less noisy than the cheaper polyester types.

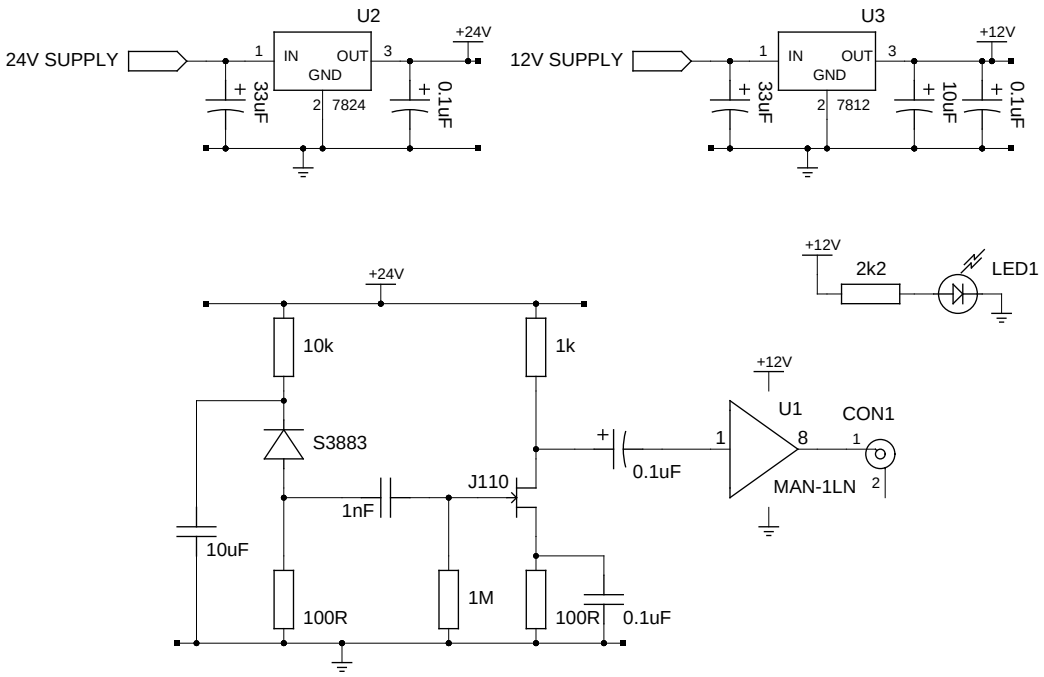

Figure 5.6: Final photodetector circuit showing FET follower, bias arrangement, voltage regulation and second-stage amplification. This circuit gives -6dBm output power at 1mW optical input power at 160MHz.

Construction

The biggest problem with any RF ($f \gg 10$MHz) circuit is unwanted reactances in the circuit. Every piece of wire, whether used for interconnection or as a component lead, will have an intrinsic inductance, and will form a capacitor with neighbouring connections. If these are known, and well characterised, they can be designed into the circuit and compensated for. Often, however,

they are frequency dependent, temperature dependent, and sensitive to mechanical vibrations. An unwanted reactance can cause a phase shift leading to oscillation or reduced gain.

PCB Layout

A badly designed PCB will give the same problems as connections using wires. For this reason, some authors e.g. [68] recommend the "dead-bug" technique where components are glued upside down with their leads in the air, and the components connected using physically separated wires. This reduces inter-connection capacitance, and the wire inductances can be tuned out. It reduces the RF problems associated with stripboard or badly-designed PCBs which often have non-negligible parasitic capacitances. Initially this technique was used with good results. However, a strong "microphonic" effect could be observed, and there were problems with thermal drift. Additionally, keeping the photodiode in place proved difficult and with the additional required voltage regulation, the final system was built on a single board. Modern techniques of multiple layer PCBs avoid traditional PCB design problems. By using a thin PCB track and a ground plane, a transmission line is formed called a "microstrip". Its characteristic impedance can be calculated using this formula [74]:

$$Z = \frac{60}{\sqrt{\varepsilon_{\text{eff}}}} \ln\left(8\frac{H}{W} + 0.25\frac{W}{H}\right) \tag{5.7}$$

where H is the effective height of the track and W is the effective width. This assumes a straight piece of microstrip. For complex board layouts with curves and discontinuities, a finite-element package is used to calculate Z. A transmission line does not create a phase-shift as the inductance balances the capacitance (Z is real). All PCB material now supplied is designed for RF work, with a well-known permittivity. FR4 board was used which has an $\varepsilon_{\text{eff}} = 4.0 \pm 0.5$. It is essential to design the track-width correctly and ensure a low-impedance unbroken ground-plane. There is a more advanced technique for multi-layer board which involves embedding the microstrip between two ground planes (called stripline). This technique gives additional immunity to both noise injection and power-loss but is more complex to design and build.

Without access to a finite element software package, it was important to keep the lines straight where possible, and to avoid edges and acute-angled bends, where formula 5.7 breaks down, and

either stray reactance, or power loss due to radiation will occur.

The ground-plane must be a low impedance path to ground and designed to ensure that the "image" created by the microstrip is not disrupted by holes or breaks in the copper, e.g. component legs or bolt holes.

Because "through-hole" components require leads to connect them, they have an extra inductance due to the lead. It is also necessary to penetrate the ground-plane and leave an insulating gap around the hole. For this reason, in RF circuitry, surface mount components are preferred to through-hole components.

Capacitive Loading

Whilst ground-planes are necessary to provide the transmission line, they can introduce problems of their own. If a component sees a large capacitance to ground, which is not correctly matched to form a transmission line, then this creates greater loading on the component. This means that around component tabs or legs, it is necessary to remove sections of the ground plane to prevent excessive loading. For the Minicircuits parts used, this was specified by the manufacturer in the data sheet. Attempts to deviate from this design resulted in parasitic oscillations. For example, the MAN-1LN requires a generous ground-plane underneath it. Without this, parasitic oscillations and a large loss of gain were observed. Even with good earthing using wires, it did not work correctly but on a ground plane the module performed well. Figure 5.7 shows the extraneous oscillations that occur from poor PCB design for a MAN-1LN amplifier. To prevent power loss, microstrip was used for interconnections on the board, and a PCB mounted SMA connector was used to make connections via 50Ω cable to the rest of the PDH circuit.

Inter-board Connection

Whilst most components can be mounted on the board, some external connections were required. The standard for RF interconnection is 50Ω transmission line, therefore the PCB was designed using 50Ω microstrip, correctly terminating on 50Ω connections, either SMB or SMA type as appropriate.

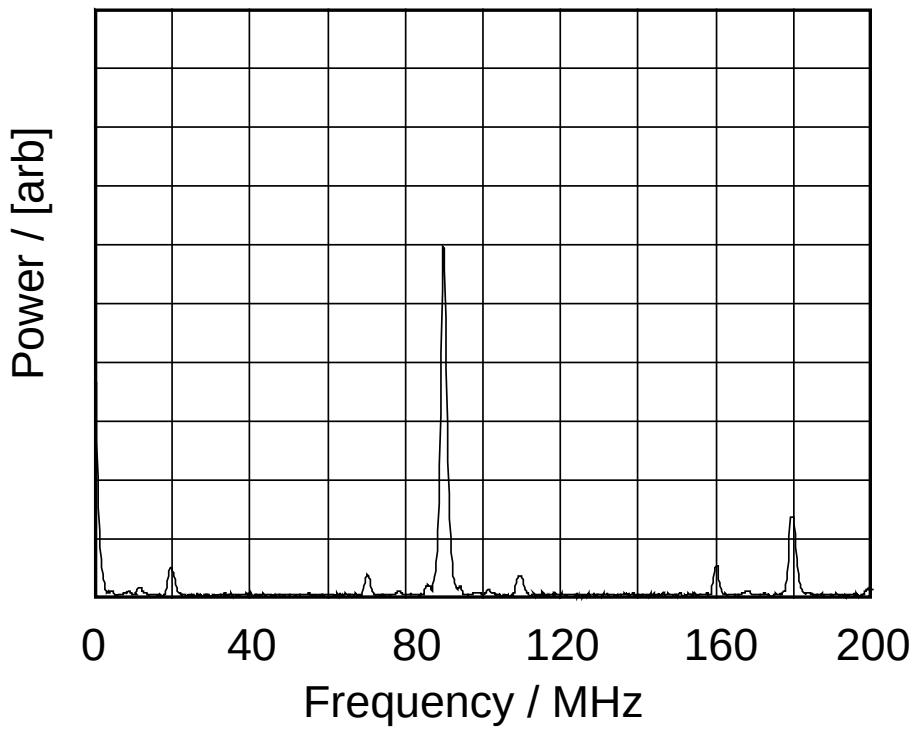

Figure 5.7: Output from a MAN-1LN amplifier for a 160MHz input signal. Several parasitic oscillations including one at 89.8MHz that exists even with no input. This depletes the amplifier so that the input signal is only weakly amplified. This was due to an insufficient ground-plane.

Power supply Ripple

Another important issue is noise injected by the power supply. The application notes for the S3883 [70] recommends using a battery for the power-source in order to reduce noise. Initially a battery was used but this was not very convenient. Therefore, with careful design, voltage regulation was built on to the board.

Where possible all power supply leads were very low impedance (i.e. physically large wires with low resistance) and the connections were designed to give minimum impedance. Supply conditioning using smoothing capacitors and chokes was implemented. The supply leads were twisted to provide CMR and reduced the risk of oscillation with the PSU.

Results

It was difficult to characterise the response, as a well-characterised optical signal was not available. The laser was directly modulated via a bias-T using a 160MHz RF source. This created phase

modulation at 160MHz, the level of which was estimated by comparing the heights of the optical sidebands to the carrier on an optical spectrum analyser. This can be taken as a percentage of the average optical power in the beam. The beam used was 1mW and the percentage of AM modulation estimated to be about 5% (a modulation index of $\beta = 0.1$). This gives a signal of 50μW. This should give a current of 30μA. This was fed into the above KPD110 circuit. At 160MHz modulation, this gave a signal of -6dBm into a 50Ω load, measured using an Agilent spectrum analyser. This gives an overall gain of the two-stage amplifier of 94dB. To characterise the SNR the PDH signal from the cavity was measured. The resulting signal was 0.5V peak-to-peak. The noise was less than 0.1mV RMS (measured on a digital oscilloscope). This gives an SNR of about 34dB. This compares well with the required level of 30dB, and it is noted that the level of the noise is about the same as the oscilloscope when grounded. The resulting frequency noise on the laser is therefore <175kHz. This compares favourably with the frequency noise of the locked laser which is 1.1MHz RMS.

5.2.5 Conclusion

The KPD-110 is an inexpensive photodiode module suitable for high-speed modulated laser beams (up to 200MHz) that is fully integrated with on-board supply regulation. It is fully compatible with the PDH system described later in this chapter, for wavelengths of 850±20 nm. The SNR of the overall system was 34dB.

5.3 Integrated Laser Lock System

The integrated laser lock system (KILL-110) is the combination of all the components required for a PDH lock (except the photodiode amplifier) on a single PCB, fitted into a Eurocard format. It uses the back-plane of the Toptica rack (DL-110) for power, and supplies a modulation source to either drive the laser's current via a bias-tee or FET, or to drive an EOM to modulate the light (after the required amplification). The PDH signal comes from either a high-speed photodetector such as the Hamamatsu APD module (C5331) or the KPD-110 which is a high-speed infrared photodetector, based on the Hamamatsu S3883 PIN photodiode (described previously).

5.3.1 Overview

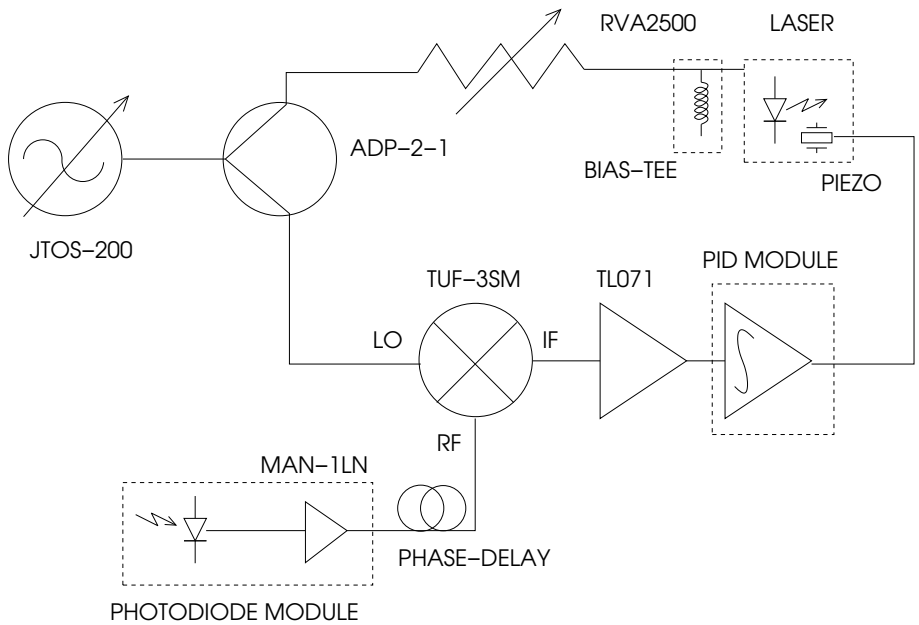

Figure 5.8: Block-diagram of RF Section of PDH circuit implemented in the KILL-110. Part numbers are Minicircuits parts, which are all surface-mount type.

The KILL-110 consists of an RF circuit which implements the PDH scheme, and some associated circuitry which provides ancillary functions. The block-diagram for the RF circuitry is shown in figure 5.8. The source for both the modulation of the laser, and the Local Oscillator (LO) is the same Voltage Controlled Oscillator (VCO). This gives excellent stability compared to using two separate oscillators. The oscillator output is fed into an RF power-splitter(ADP-2-1) which is a carefully made transformer that ensures there is a good impedance match at both the input and output of the splitter. This is more efficient than, for instance, splitting the output of the VCO using a tee, which would cause lost power due to impedance mismatch. One output goes to the laser via a voltage controlled attenuator (RVA-2500). This allows the level of the signal to be adjusted before going to the bias-tee on the laser, or the amplifier driving an EOM. A mechanism to adjust sideband power is required, as the exact level of the sidebands for repumping out of the hyperfine sub-levels of $^{43}Ca^{+}$ is a trade-off against side-band power and the total laser power that can be achieved [23] and is different under different conditions. The amplifier used with the EOM is of a fixed-gain, so being able to attenuate the signal for optimum EOM performance is also useful. The frequency of the oscillator is also tunable for the same reasons.

The other output of the power-splitter feeds the Local Oscillator input of a Diode Ring Mixer (DRM). A DRM works by using a diode as a linear mixing element. If a voltage is applied at both ends of a component, the component is subject to the sum of those voltages. Isolation is required between the two input ports. At DC a diode would provide this, but would not allow mixing of the negative-half of the signal. By using a ring of diodes and two matching transformers, good mixing and good isolation can be achieved. To ensure correct operation, the PN junctions of the diodes must be saturated, otherwise the diode will create harmonics of the input signal. Mixers are rated by the minimum level of LO power required to saturate the PN junctions. Driving at a lower or higher level will create distortion of the mixed signal. A Minicircuits mixer (TUF3-SM) was chosen. This device requires 7dBm of input power to correctly saturate the diodes (level-7). The power from the splitter was chosen to match the required power of the mixer before losses. Some loss was inevitable; however, the manufacturer's data shows that this loss has a negligible effect on the performance of the chosen mixer. The mixer worked satisfactorily at a much lower power of only 0dBm.

The RF input of the mixer is connected to the signal from the photodetector. This has already been amplified at the source (to keep a good SNR) so it is sufficient to drive the mixer (mixers can detect very weak signals at the RF input, down to -30dBm). The resultant IF will only be as strong as the RF input signal, with considerable loss in the mixer, due to inefficiency of the diodes. The IF signal is therefore further amplified by a post-amplifier circuit. The amplifier chosen was a low bandwidth (2MHz) operational amplifier (TL071) as it is only the error signal that is of interest in the PDH scheme. The modulation signal is not explicitly filtered away, and can still be detected at the output. This does not cause a problem as the rest of the lock-circuit has a bandwidth of $<$ 1MHz. The post-amplifier needs a bandwidth that is as fast as the fastest error signal being corrected for. This is limited by the response of the piezo, which is 600Hz for the 866nm laser. A faster path was not provided to the current control; however, should this be added at a later date signals up to 2MHz are available at the output. To reduce noise, the signal is integrated before being amplified and fed to the piezo, using a Toptica PID-110 module.

Between the laser being modulated and the signal being detected and returned to the mixer, a phase-delay is introduced by the effective optical path and interconnecting cables. To ensure

maximum error signal, this phase needs to be adjusted. To do this a variable phase delay device could have been introduced into the circuit. However, these devices were not readily available at the required frequency. Instead, the length of co-axial cable between the photodetector and the KILL box was adjusted to provide a phase shift to a resolution of $1\text{m} \approx \lambda/3 \approx 2\pi/3$. This is not the most convenient situation but suffices.

The rest of the KILL-110 circuit is the required voltage regulation for the individual RF parts, and a slow-switch on function. When the power is applied to the VCO, an initial transient response was observed. These transients could potentially damage the laser, although their power is low. To prevent the switching of the VCO causing damage, and to provide a convenient way of switching the RF supply to the laser or EOM off and on, a soft-switch function was provided. This makes use of the excellent isolation provided by the adjustable attenuator. This is comparable with the best RF switches, and so a separate RF switch was not employed. This also prevented further loss of RF power. The voltage control pin of the attenuator is ramped up from zero to the preset power level, and back down again at the flick of a switch. The ramping is generated by charging and discharging a capacitor. This is not very linear, but for this application, linearity was not required. The charge/discharge voltage level is amplified by an op-amp and then by a simple (class B) push-pull transistor amplifier. This transistor amplifier was required as the attenuator draws 30mA on its control pin, which is more than commercially available op-amps can normally provide (the 741 can give about 20mA) and for the cost, two transistors do an effective job. It was not necessary to cancel the distortion caused by this arrangement. The complementary pair BC108 and BC178 general purpose small signal transistors were chosen for their price and availability. The switch was not debounced, as the large capacitor effectively filters out the transients created in the switch. For results of this circuit, see figure 5.9.

5.4 Results

Lock-signals (P_ε) from an infrared and violet laser are shown in figures 5.10 and 5.11 respectively. These two signals have a SNR better than $1 : 10^6$ as was required. The theoretical fits highlight both the asymmetry of the sidebands and the phase-difference introduced in the system. This results in a DC offset of the signal which does not effect the lock in anyway.

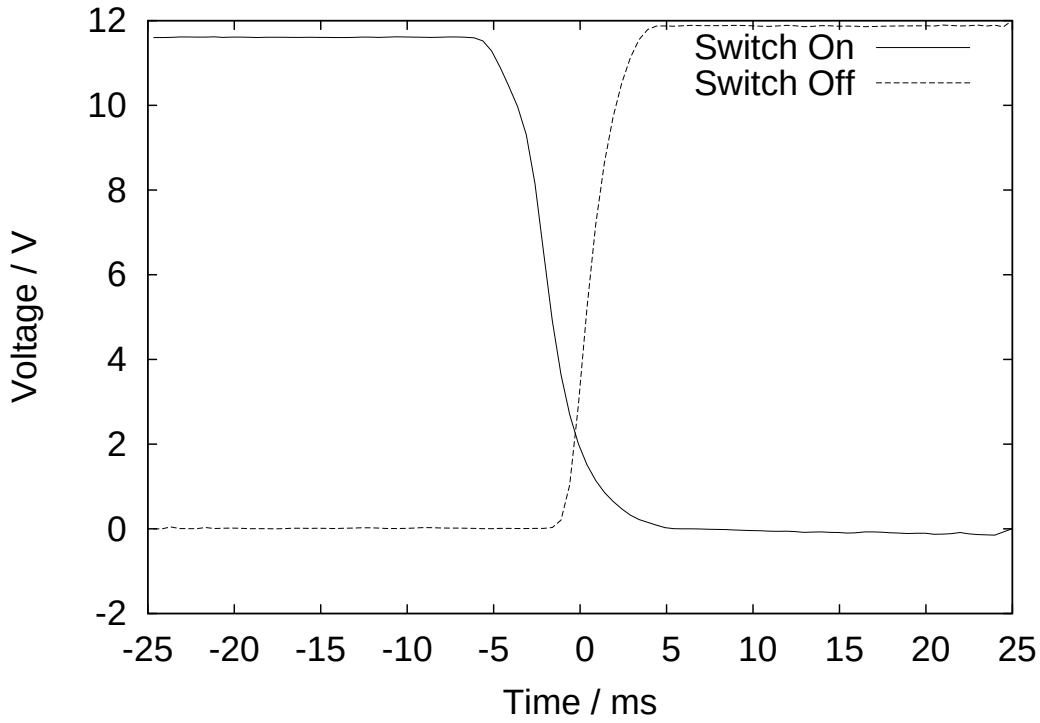

Figure 5.9: Output from the slow switch-on (and switch-off) circuit. The voltage ramps slowly to ensure that no damage occurs to the laser when the KILL unit is powered up.

Long-term drift of a laser locked to a reference cavity at 854nm was investigated. The results are presented in figure 5.12. This shows results taken on a wavemeter (LM10) over a period of 23 hours. The large scatter of results is partially due to discretisation errors in the LM10 wavemeter [72]. Over longer timescales, the laser was seen to occasionally mode-hop. When the laser hops to another mode, it can lock to another cavity mode. However, this was a very rare occurrence over the a timescale of days, and the laboratory was not temperature stabilised which may cause a problem with the cavity.

The drift is 0.27MHz per hour which is within the specified drift of the NPL cavity of 0.5MHz per hour. Lab temperature fluctuations may affect the cavity temperature stabilisation and the accuracy of the LM10 wavemeter, which relies on a HeNe reference laser which is not temperature stabilised. It is a 100-fold improvement on the frequency drift of an unlocked laser which was measured to be 26MHz per hour.

Tests involving mechanical vibration have shown this implementation of the PDH-locking scheme to be very robust. Experiments have been performed with trapped ions using locked lasers

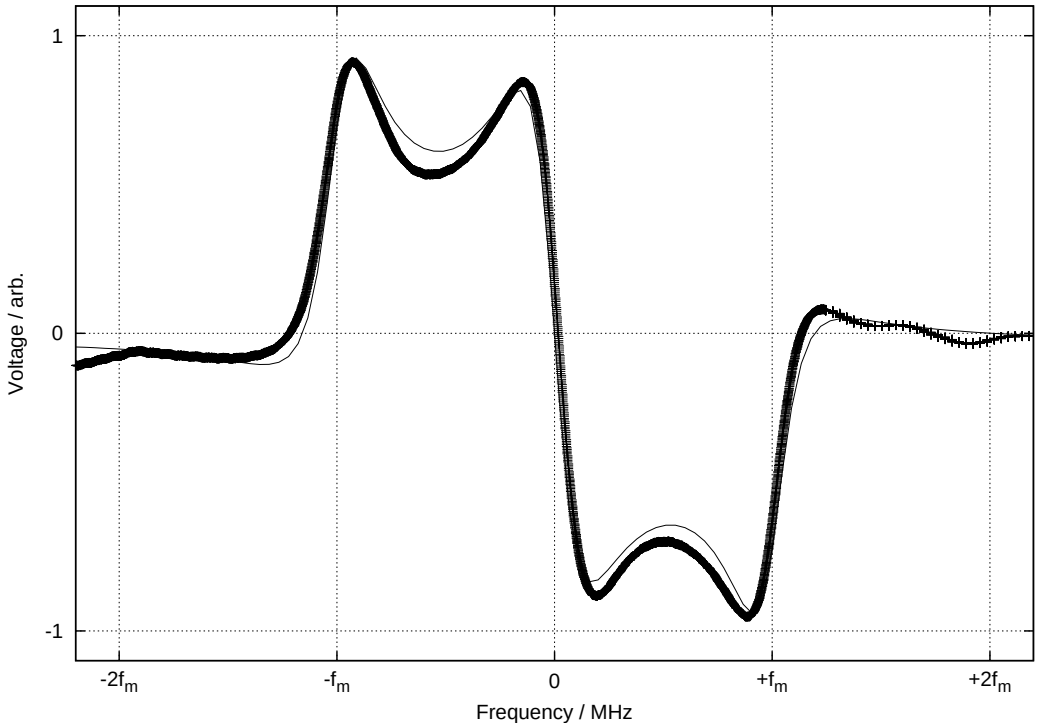

Figure 5.10: Signal from the KILL circuit after mixing and amplification for the 866nm laser. The crosses are data and the continuous (thin) line is the theory. The current is directly modulated at 160MHz and the signal reflected is detected on the S3883 circuit described above. The etalon is an NPL low drift etalon with $\mathcal{F} = 30$ and FSR=1620MHz. The fitted parameters of the plotted theory curve suggest the sidebands are at 5% of the carrier and are asymmetric by a factor of 1%. This can be seen on a wavemeter.

and experiments greater than 12 hours in duration have been run, without any need to relock the laser.

5.5 Conclusion

PDH locking to a cavity with a finesse of $30-50$ allows a laser to stay at a chosen frequency for several hours. It is robust against thermal fluctuations and mechanical vibrations at low frequency. This type of lock does not attempt to narrow the laser linewidth or to stabilise the laser against high-frequency noise such as electrical pickup as it is limited by the piezo. Its success is due to the high-modulation frequency resulting in a wide capture range and good SNR in the electronics resulting in low frequency noise in the laser. The low-drift etalons provide good long-term stability. The KILL-110 unit is easy to use and has been successfully used in many trapped ion

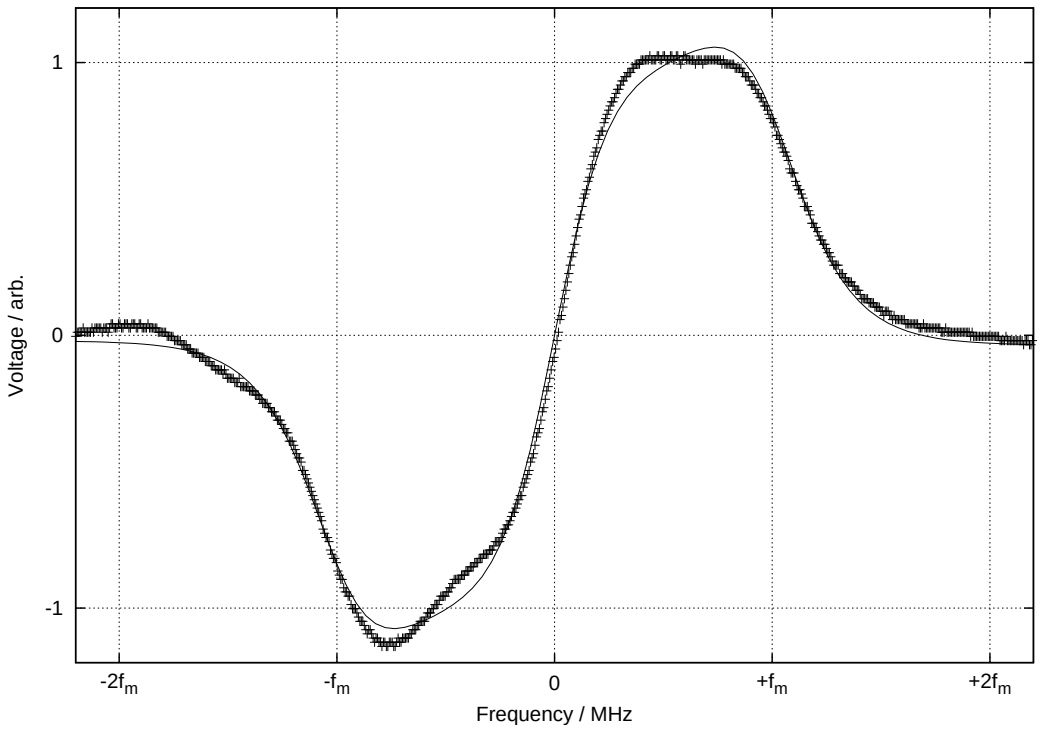

Figure 5.11: Similar results to 5.10 are shown for the 397nm laser with a modulation of 85MHz using an EOM. A Hamamatsu APD is used. The modulation depth is nearly 30% and the cavity has a lower finesse. The fitted parameters give almost perfectly symmetrical sidebands as would be expected for an EOM.

experiments, including the experiments presented later in this thesis. The frequency jitter on the laser was measured to be 1.1MHz with $\lesssim$100kHz being due to noise due to the locking electronics.

By using a frequency of $\approx$160MHz for the lock-servo, sidebands can be generated to drive the complete hyperfine manifold of the infrared D-P transitions in $^{43}Ca^{+}$.

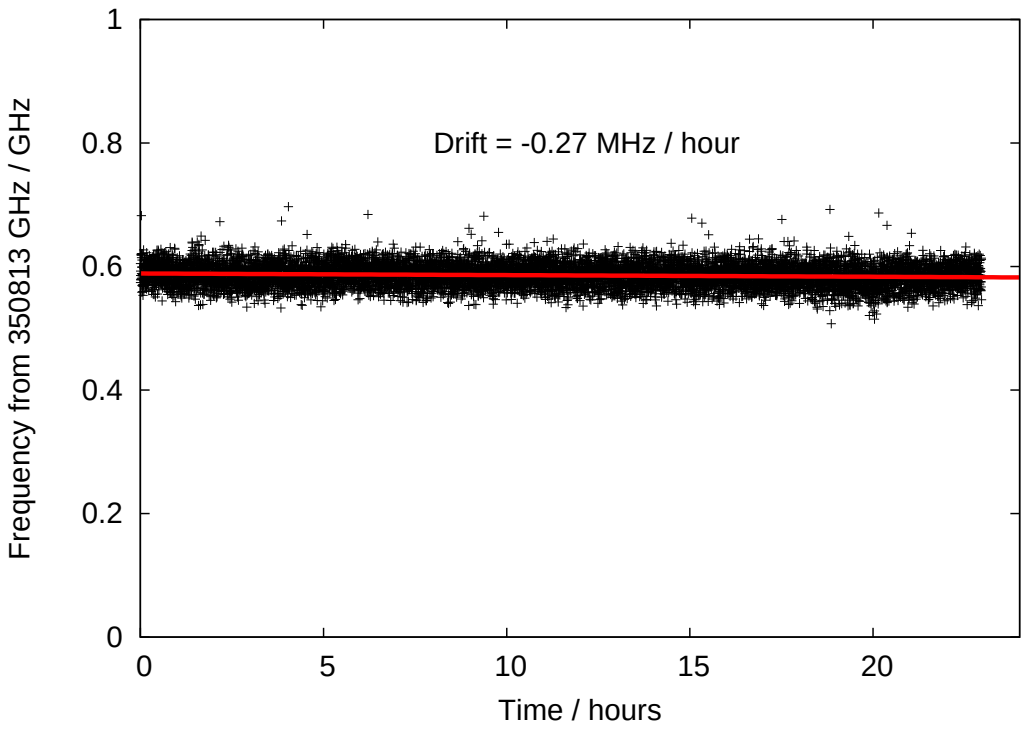

Figure 5.12: Results from LM10 wavemeter. The scatter of points above the mean is due to the wavemeter miscounting fringes. Due to the modulation on the laser, the interference fringe amplitude is modulated, which may cause the wavemeter to occasionally miss one or more fringe-counts.

Chapter 6

Investigation of Magnetic Field Fluctuations

6.1 Introduction

The main source of decoherence in the qubit states in the experiments presented in this thesis, is understood to be magnetic field noise and drift. For the T_2 measurement in chapter 7, the clock transition, $S_{1/2}$(F=4,M_F =0)$\rightarrow$$S_{1/2}$(F=3,$M_F$ =0) was used which has a field-dependence of $\sim$4.0Hz/mG at an applied field of 1.8G. For the required fidelity as discussed in chapter 7, this requires the transition frequency to be stable to $\sim$1Hz, or a field stability of 0.25mG.

The splitting between energy levels in $^{43}Ca^{+}$ is influenced by external magnetic and electric fields, their directions, and their amplitudes. To lift the degeneracy of the ion's states, an external Zeeman field is applied. Fluctuations in this field, due to the current supplied to the coil fluctuating, can also lead to energy-shifts at the ion. In the following experiments, an external field of around 1.8G was used, as at lower fields it is more difficult to drive fluorescence from the ion, because optical pumping into dark states occurs [63].

The electric fields shifts are much smaller than those due to the magnetic field. For example, comparing the Stark effect from the electrodes to the magnetic effect, the energy shift is < 0.1Hz [75], compared to magnetic shifts on the order of 10Hz on the clock-transition.

The main source of magnetic field noise is likely to be due to the electrical supply in the lab,

and variation in the local magnetic field due to movement of large ferrous objects. Changes of the local magnetic field by several mG due to vehicles, lifts, and workshop activities have been measured [76]. A further source of high-strength DC magnetic fields are the superconducting magnets used by other research groups in the vicinity. These are on the order of several Tesla and are no more than 100m away from the experiment. It is possible that ferromagnetic materials both in the building, and in other equipment, may be temporarily magnetised by a field of this magnitude.

As well as mains supply noise, circuitry in the lab generates other fields at higher frequencies. One common source is switch-mode power supplies, which use a high frequency oscillator to generate a lower voltage, and then regulate the voltage by varying the frequency of this oscillator. Commonly this will generate 100kHz and higher harmonics. Other sources of fields also include distant radio transmitters and cosmic sources which are of even higher frequencies. The experiment is situated in a basement, which means these sources are much weaker, and of less significance than local sources. Most well-designed lab equipment, especially that which meets EMC regulations, will be well-shielded. High-fidelity (>98%) π-pulse experiments have shown that these fields are not a significant problem (see figure 6.1).

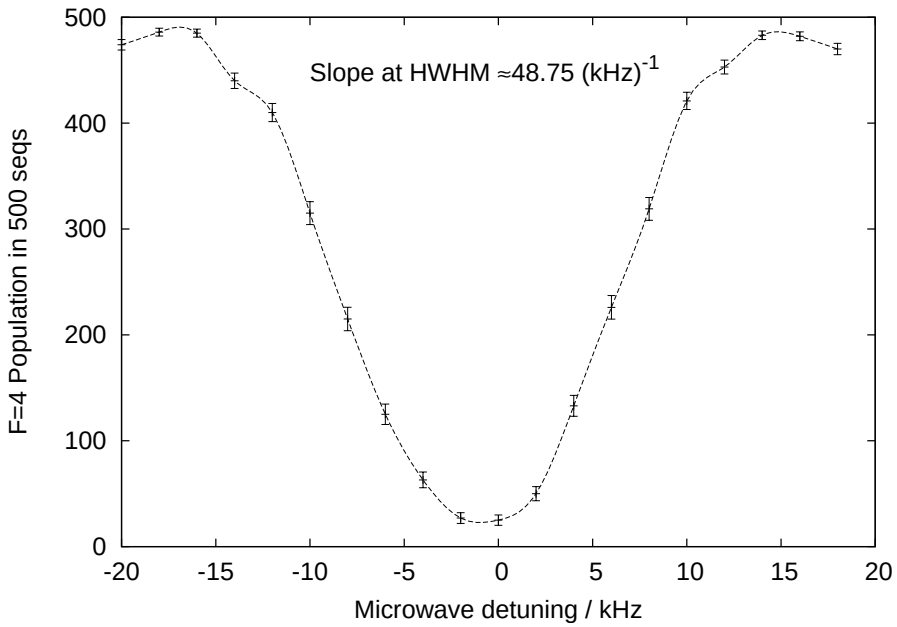

Figure 6.1: A π-pulse (Rabi flop) on the stretched-state hyperfine transition of $^{43}Ca^{+}$ showing a 96% contrast relative to the available read-out efficency. Detuning is relative to 3221344kHz.

The process of heating calcium metal by passing a current through a thin tube to generate the atomic beam used to load the ion-trap, is known to generate a local field inside the vacuum apparatus [41], and is compounded by the local heating which creates a temporary magnetisation of the surrounding metal (see figure 6.6).

The aim of these experiments was to measure the stray fields in the lab under different conditions and to try and correlate them with possible sources. Previous measurements [76] were made using a field probe. However, the most sensitive test, and the most accurate in terms of location, was to measure the effect on the ion itself by probing a magnetic field-sensitive transition.

6.1.1 Magnetic Sensitive Transition

To characterise the effect of magnetic field fluctuations on the ion, both Rabi-flopping and Ramsey experiments [77] were performed. The $^{43}Ca^{+}$ stretched state $S_{1/2}$(F=4,M_F =+4)$\rightarrow$$S_{1/2}$(F=3,$M_F$ =+3) qubit transition was used (see chapter 2) as this is sensitive to the field by -2.45kHz/mG in a nominal applied field of 1.8G, and the easiest state to prepare.

6.1.2 Theory

The effect of radiation on a two-state system is considered by treating the radiation field classically. Here we follow the derivation in Foot [78]. In the dipole approximation, when the wavelength of light (or microwaves) is greater than the size of the atom, the radiation can be considered as a perturbation on the atom.

For a hyperfine transition that is not far from resonance, the atomic system is assumed to be two-level. Because the detuning is small compared to the frequencies of the transition, the rotating wave approximation can be made.

The time-dependent wavefunction for a two state system can be written:

$$\Psi(\mathbf{r},t) = c_1 \left|1\right\rangle e^{-i\omega_1 t} + c_2 \left|2\right\rangle e^{-i\omega_2 t} \tag{6.1}$$

where $|i\rangle$ are the stationary solutions to the wavefunctions for the atom's levels 1 and 2. $\omega_i = E_i/\hbar$ is the frequency of that level and c_i is the corresponding time-varying amplitude.

It can be shown [78] that a perturbing oscillating electric field of the form,

$$\mathbf{E} = \mathbf{E_0} \cos \omega t$$

causes the two atomic states to mix, driving population between the two states. This gives rise to Rabi flopping, where the Rabi frequency is defined as:

$$\Omega_R = \frac{\langle 1 | e\mathbf{r} \cdot \mathbf{E_0} | 2 \rangle}{\hbar} \tag{6.2}$$

The population of each state changes in a time-dependent way:

$$\begin{aligned} i\dot{c}_1 &= \Omega_R \cos \omega t e^{-i\omega_0 t} c_2 \\ i\dot{c}_2 &= \Omega_R^* \cos \omega t e^{i\omega_0 t} c_1 \end{aligned} \tag{6.3}$$

where $\omega_0 = (E_2 - E_1)/\hbar$

In a Ramsey experiment, the state is prepared in a linear superposition of states, i.e. $\frac{1}{\sqrt{2}}(|1\rangle + |2\rangle)$. This is done with a $\pi/2$ pulse i.e. a pulse of duration, $t = \frac{\pi}{2\Omega_R}$ on resonance with the transition. By comparison, a π pulse would take population from $|1\rangle$ to $|2\rangle$.

By making the rotating-wave approximation, (i.e. the radiation frequency is close to resonance with the transition frequency, ω_0, and therefore terms that contain $\exp(i(\omega + \omega_0)t/2)$ average to zero when integrated w.r.t. time) the coupled differential equations in equation 6.3 can be solved for monochromatic radiation to give:

$$|c_2|^2 = \frac{\Omega_R^2}{W^2} \sin^2 \left(\frac{Wt}{2} \right) \tag{6.4}$$

where

$$W^2 = \Omega_R^2 + \delta^2$$

and $\delta = (\omega - \omega_0)$ is the detuning of the radiation from resonance.

This can be rearranged in terms of density matrix elements. The density matrix of a two level

system can be expressed as:

$$|\Psi\rangle\langle\Psi| = \begin{pmatrix} |c_1|^2 & c_1c_2^* \\ c_2c_1^* & |c_2|^2 \end{pmatrix} = \begin{pmatrix} \rho_{11} & \rho_{12} \\ \rho_{21} & \rho_{22} \end{pmatrix} \tag{6.5}$$

Equation 6.3 can be written:

$$\begin{aligned} i\dot{c}_1 &= \frac{\Omega_R}{2}e^{i\delta t}c_2 \\ i\dot{c}_2 &= \frac{\Omega_R}{2}e^{i\delta t}c_1 \end{aligned} \tag{6.6}$$

By making the substitution:

$$\tilde{c}_i = c_i e^{-i\delta t/2}$$

in equation 6.6 this gives:

$$\begin{aligned} i\dot{\tilde{c}}_1 &= \frac{1}{2}(\Omega_R\tilde{c}_2 + \delta\tilde{c}_1) \\ i\dot{\tilde{c}}_2 &= \frac{1}{2}(\Omega_R\tilde{c}_1 + \delta\tilde{c}_2) \end{aligned} \tag{6.7}$$

By making the following substitutions and by considering a reference frame rotating at ω_0,

$$\begin{aligned} u &= \tilde{\rho}_{12} + \tilde{\rho}_{21} \\ v &= -i(\tilde{\rho}_{12} - \tilde{\rho}_{21}) \\ w &= |c_1|^2 - |c_2|^2 \end{aligned} \tag{6.8}$$

the following set of equations is arrived at:

$$\begin{aligned} \dot{u} &= \delta v \\ \dot{v} &= -\delta u + w\Omega_R \\ \dot{w} &= -v\Omega_R \end{aligned} \tag{6.9}$$

This can be expressed as a vector product,

$$\mathbf{R} \times \mathbf{W}$$

where

$$\mathbf{R} = u\hat{\mathbf{e}}_1 + v\hat{\mathbf{e}}_2 + w\hat{\mathbf{e}}_3$$

and

$$\mathbf{W} = \Omega_R\hat{\mathbf{e}}_1 + \delta\hat{\mathbf{e}}_3$$

therefore a state-vector can be mapped on to a sphere, with axes in Hilbert space of $\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3$.

Any normalised state vector for a two-state system lies on a unit sphere, the Bloch sphere (in the rotating frame). The "north pole" corresponds to the state $|\uparrow\rangle$ and the "south pole" is the state $|\downarrow\rangle$. The equatorial plane corresponds to an equal linear superposition of $|\downarrow\rangle$ and $|\uparrow\rangle$ with some phase-difference relating to the position around the circumference (see figure 6.2).

A $\pi/2$ pulse drives the population, for example, from the "south pole" of the Bloch sphere to the equatorial plane. If the non-rotating frame is considered, the state of the ion then precesses around the magnetic field axis at ω_0, so that the state acquires phase at the frequency given by the energy of the state separation. On the Bloch sphere this is a continuous rotation along the equator. If this rotation is at a constant rate, and in-phase with the local oscillator (the microwaves or Raman difference frequency) which was used for the $\pi/2$ pulse, then a second $\pi/2$ pulse would put the ion in to the upper state, $|\uparrow\rangle$. However, fluctuations in the magnetic field strength and direction cause the precession frequency to change. This causes a phase difference between the original driving source, and the ion's internal state. This could be compensated for, if it was well-known. However, it is not possible to measure the exact field level and direction at high levels of precision at the ion's position, except by using the ion itself. Therefore a second $\pi/2$ pulse will not necessarily leave the ion in the $|\uparrow\rangle$ state. The fidelity of measuring the $|\uparrow\rangle$ state is directly correlated to the acquired phase difference between the ion and the driving source, and therefore the field fluctuation. Instability in the local oscillator will also create the same effect.

In order to observe a change in the ion's frequency with respect to the local oscillator, either the detuning of the local oscillator or the phase difference between the two $\pi/2$ pulses can be

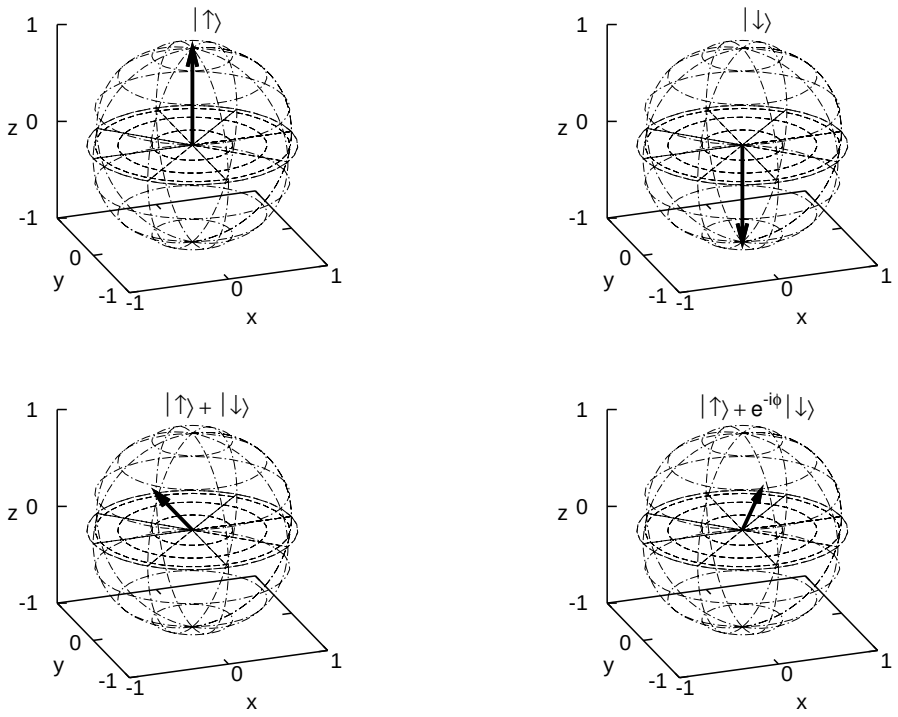

Figure 6.2: Bloch-sphere representations of two-level systems. Note that wave-functions are taken to be normalised.

varied. This leads to Ramsey fringes, similar in nature to the interference pattern generated by light passing through two slits. The height and phase of these fringes directly relates to the magnitude in the direction of changes the external field.

This gives rise to the Ramsey formula [78]:

$$|c_2|^2 = \left|\frac{\Omega_R \tau_P}{2}\right|^2 \left[\frac{\sin(\delta\tau_P/2)}{\delta\tau_P/2}\right]^2 \cos^2\left(\frac{\delta\tau_L}{2}\right) \tag{6.10}$$

where: $\delta = \omega - \omega_0$ is the detuning; τ_P is the pulse-time (chosen such that a $\pi/2$ pulse is realised); τ_L is the delay-time; and Ω_R is the Rabi frequency.

In order to generate the fringes, population can be measured whilst δ is varied. Alternatively, as was done in some experiments, the phase of the second $\pi/2$ pulses, ϕ, was varied, whilst the

detuning was held constant. This removes the sinc envelope.

6.2 Setup

To measure field fluctuations, a microwave source was used to drive the stretched-state transition, $|\downarrow\rangle$ = $S_{1/2}$(F=4,M_F =+4) $\rightarrow$ $|\uparrow\rangle$ = $S_{1/2}$(F=3,M_F =+3) (see figure 1.2). At low field, this transition is the most sensitive to field changes.

The microwave source was an oven-stabilised synthesiser (Agilent E4426B), fed by a Stanford Instruments Rubidium clock source (FRS725). This is specified to be stable to $\sigma_y(10s) = 3 \times 10^{-12}$ where $\sigma_y(10s)$ is the Allan deviation for a 10s period (see [37]). This creates a negligible error compared to field-fluctuations at the ion.

Conveniently, it is possible to prepare the $|\downarrow\rangle$ state by optically pumping the population using σ-polarised light, i.e. photons can only increase M_F until the end of the manifold is reached. $S_{1/2}$(F=4,M_F =+4) was prepared using the 397nm $\sigma+$ beam on the $S_{1/2}$(F=4) $\rightarrow$ $P_{1/2}$(F=4) and $S_{1/2}$(F=3) $\rightarrow$ $P_{1/2}$(F=4) transitions, (with M_F always increasing by one).

The experimental sequence consisted of Doppler cooling (no further cooling was necessary) to localise the ion, followed by state preparation to the $|\downarrow\rangle$ = $S_{1/2}$(F=4,M_F =+4) state. A $\pi/2$ pulse was applied, followed by the Ramsey delay-time, and then the second $\pi/2$ pulse. Readout of the state was carried out, as described in chapter 1. The whole pulse sequence is shown in figure 6.3

Two different experimental setups were used. By triggering the experiment's commencement from the mains power supply (so called "line-triggering"), effects generated by the 50Hz mains supply could be compensated for. This allowed the measurement of slower changes in the field. Drift in magnetic field on a slower timescale was more likely due to moving objects distorting the Earth's magnetic field. Likely sources were discovered to be: the building's lift, forklift truck operation, delivery vehicles moving, movement of metal doors and movement of ferrous stock on the floor above. By line-triggering, these effects could be seen without the additional effect of the mains-supply.

In the second setup, the phase difference between the line source and the start of the experiment was varied, in order to show the effect of the mains supply field.

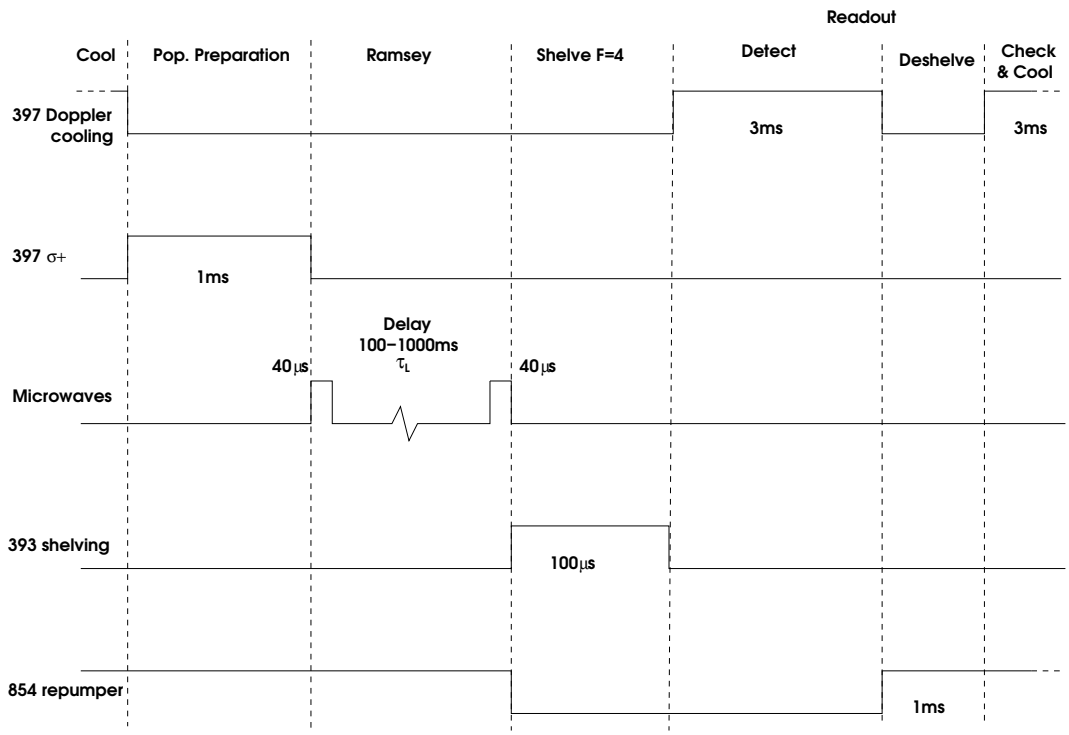

Figure 6.3: Simplified Ramsey pulse sequence. τ_L is the time between the leading edge of each $\pi/2$ pulse.

6.2.1 Field Cancellation

A commercial field-cancelling system (Spicer SC20) was tested. This device works by measuring the field in three orthogonal directions at one or two locations, and applying current to three corresponding orthogonal coils in an attempt to cancel the field. It can work from DC up to 2.5kHz. The coils are relatively large (having a radius of 1-2m) and were placed outside of the experimental apparatus (see figure 6.4). Electrical equipment close to the SC20's sensors generates a field that can not necessarily be cancelled at the ion, and can therefore contribute to field noise. The two detectors were placed as close to the ion as was possible without saturating the sensors (which saturate at 2G). Due to the high field around the experiment's Zeeman coils this needed to be ∼23cm either side of the coils (see figure 6.4). One detector was able to measure the field at its location. The way in which the sensor works has not been disclosed by the manufacturer. With two detectors the gradient of the field can be measured, allowing the field in the area between the sensors to be compensated. This is done by mixing proportions of each sensor's output together, for each orthogonal direction separately. This was tried, with different mixer settings, however, no

improvement over using one detector was discernible, so only one detector was used for further investigations. The three directional inputs are integrated before being fed to the cancelling coils. The integrator gain for each direction was varied and measurements were taken.

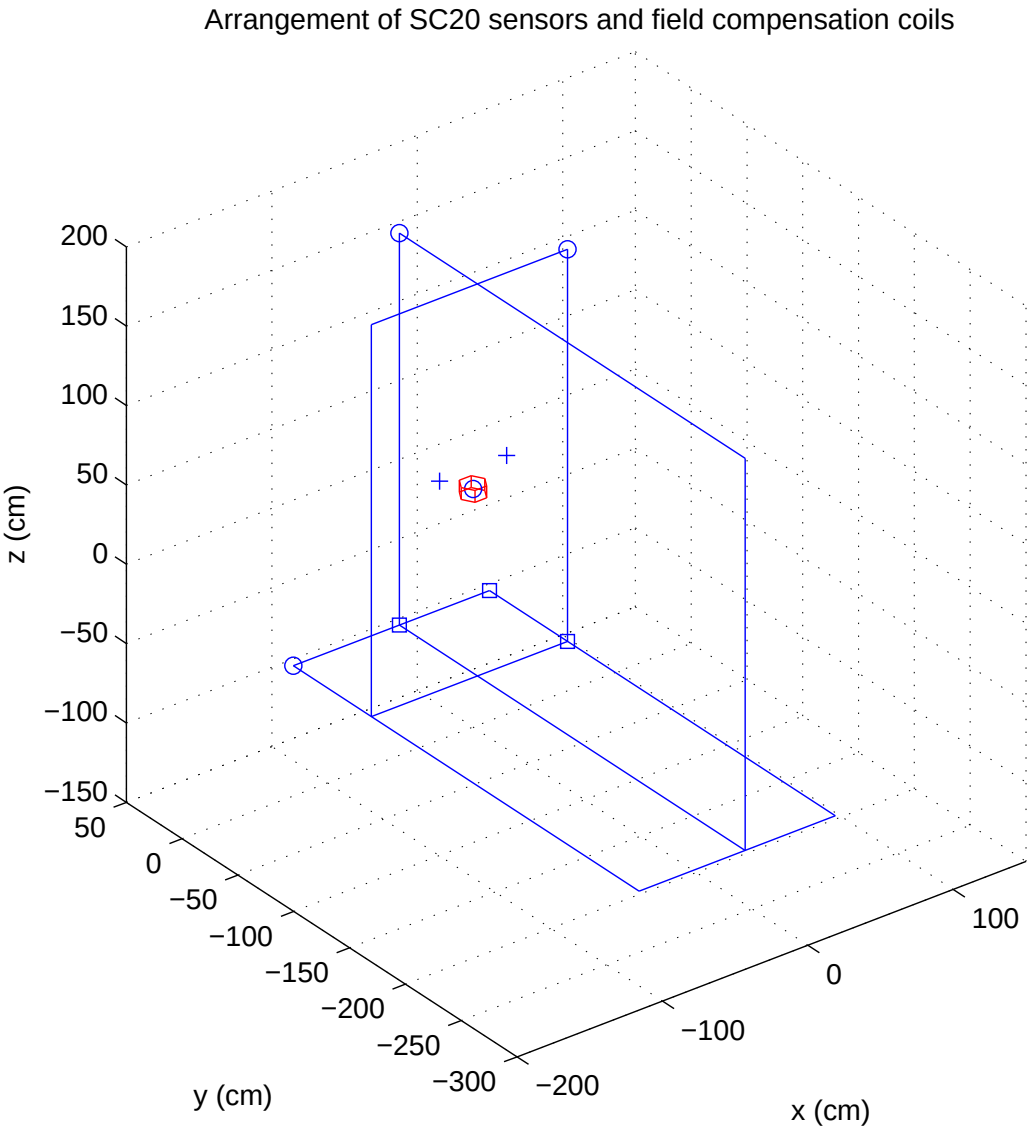

Figure 6.4: Location of the SC20 field compensation coils and sensors. There are three orthogonal compensation coils, each roughly rectangular in shape. The trap is at the co-ordinate origin, with the hexagonal vacuum can (red) shown to scale. The "left" and "right" sensors are at $(\pm 23, 0, 13)$*cm, indicated by* +.

Measurements were made with and without field cancellation operational. The SC20 was also used as a probe to take measurement of the field (by opening the feedback loop), allowing fluctuations in the stretched-state transition frequency to be compared with changes in the field,

as measured by the field cancellation system. This proved to be a useful way of correlating large, slow changes to the static field, but it was not a reliable measure of the field at the ion.

6.3 Results

6.3.1 Fitting

The result of a typical Ramsey experiment, where the timescale is shorter than the slow drifts of the field (120μs), is shown in figure 6.5.

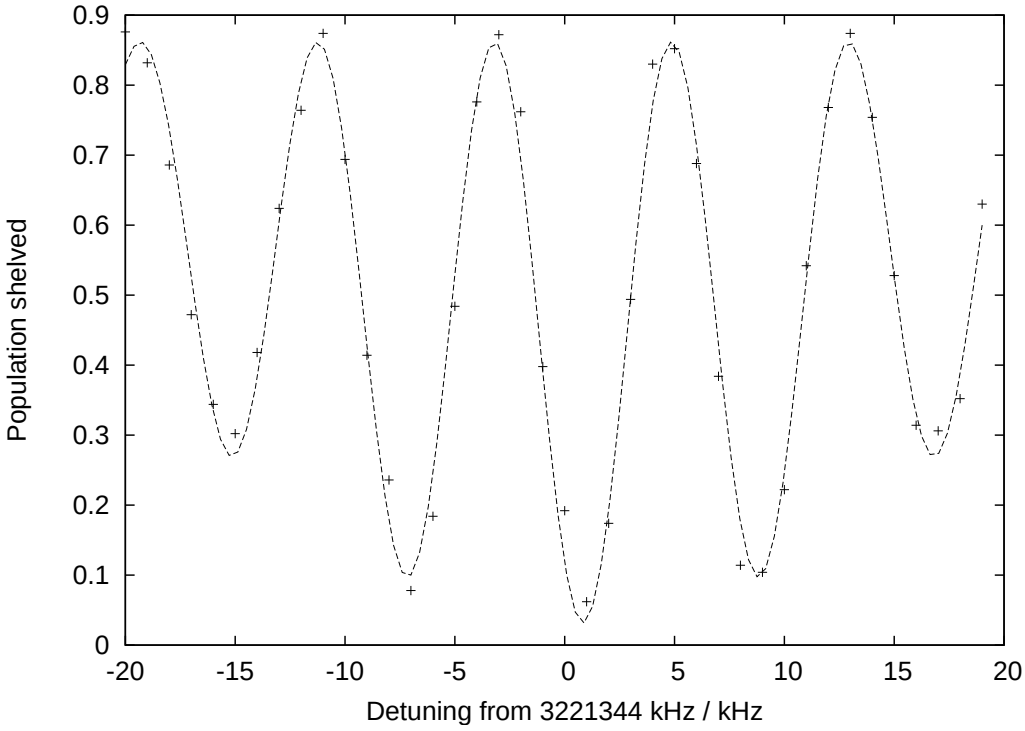

Figure 6.5: Results from a Ramsey experiment with a delay (τ_L) of 120μs. From the fitted function, the detuning, δ, was found to be 0.83(3)kHz and the amplitude 0.83(2). In this experiment line-triggering was used (see section 6.3.4) and the total run time was 400s.

This allows the detuning from the expected resonant transition frequency(f_0) to be determined, by fitting each Ramsey fringe with the formula:

$$P_1 - P\left[\mathrm{sinc}(\pi(f-f_0)\tau_p)\right]^2 \cos^2(\pi(f-f_0)\tau_L)$$

P_1 is an offset allowing for preparation or readout imperfection. P_1, P and f_0 were floated to

minimise the residual, whilst τ_L and τ_p were known to a precision of 0.1µs.

Relating this to the Ramsey formula, 6.10 it can be seen that:

$$P \propto \left| \frac{\Omega_R \tau_P}{2} \right|^2$$

and:

$$\delta = 2\pi(f - f_0)$$

Experimental data points were weighted using the binomial error formula: $\sqrt{np(1-p)}$, as the population is derived from a discrete count of either shelved or not shelved (a Bernoulli Trial). The frequency offset, f_0, was used as a measure of magnetic field. Four different sets of experimental results were taken. Initially the effect of the oven's supply current and associated thermal change were investigated. The second set of experiments were to investigate general background fields due to machinery, etc. The third set were to investigate the effect of the mains electricity supply on the ion, and the fourth set were to investigate the use of the active field-compensation unit (Spicer SC20), to cancel the 50Hz supply field and slow drift.

6.3.2 Oven Current

To vaporise calcium metal and generate calcium ions, a current of typically 4.25A is passed through a thin metal tube filled with calcium, situated inside the vacuum can. See [39] for details. This current produces a local magnetic field near the ion, and the heating appears to create a thermal effect that sustains that field, before decaying exponentially back to equilibrium. This temporary magnetisation could be due to ferrous materials in the vicinity of the ion trap. The long time constant indicates this is probably due to a thermal effect. The results are shown in figure 6.6. The fitted expression implies that any effect is below the 1mG level after approximately 80mins. Therefore, as a precaution, all magnetic field sensitive experiments are done at least one hour after loading the ion and switching off the oven.

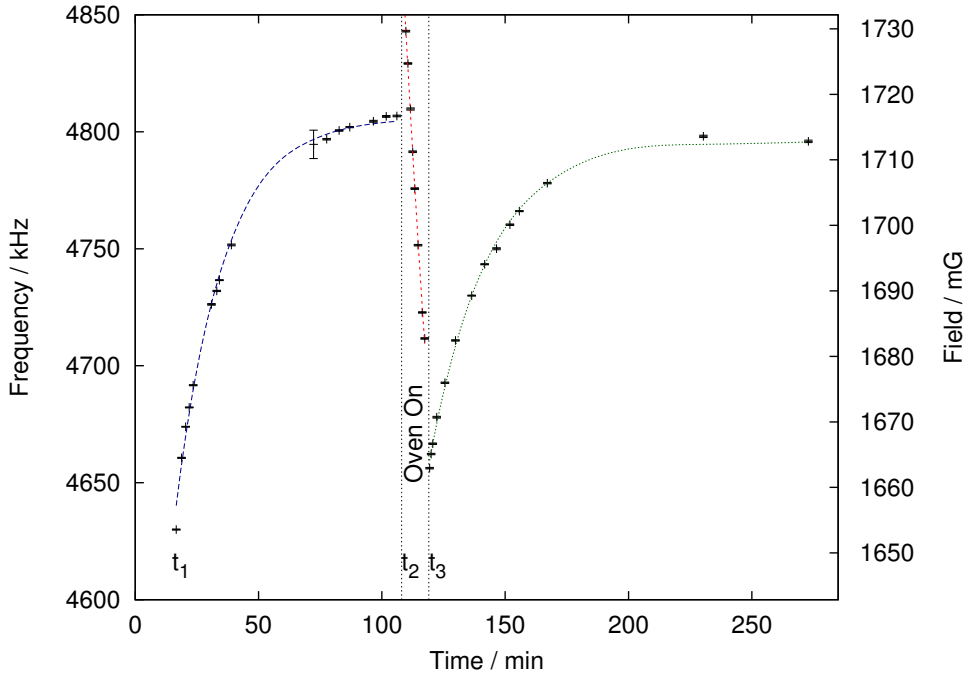

Figure 6.6: Graph of detuning vs time after switching on and off the oven's current supply. After loading an ion at t_1, the oven was switched off and the field decays exponentially ($\tau = 19min$). At t_2 the oven was switched on again and a field of $\approx 50kHz \approx 18mG$ due to the DC oven current (4.25A) is immediately visible, followed by an increase in field at ~5mG/min while the oven is on. When the oven is turned off at t_3 the 18mG field is removed and the field decays again with $\tau = 23min$. This data was taken using a $^{40}Ca^+$ ion; similar results are obtained for $^{43}Ca^+$.

6.3.3 Background Field Drift

Preliminary measurements using a field-probe show that movement of ferrous objects create a change in the field of several mG. Figure 6.7 shows these results. Further investigation using the ion and the Ramsey technique described above, produced the results in figure 6.8. Notes on the figure indicate correlations with activity in the building. Use of high-strength ~1T superconducting magnets in neighbouring labs (<100m away) also affected the results. Data collected on the Christmas and Easter holiday showed less drift. This implies that activity in the building contributes a background field of several milligauss that changes on the timescale of 10 mins – 1 hour.

To measure drift on a faster timescale, the centre of a resonant scan was measured. The results are shown in figure 6.9. It can be seen that as well as a slow background drift, the field is fluctuating on a timescale of minutes as well.

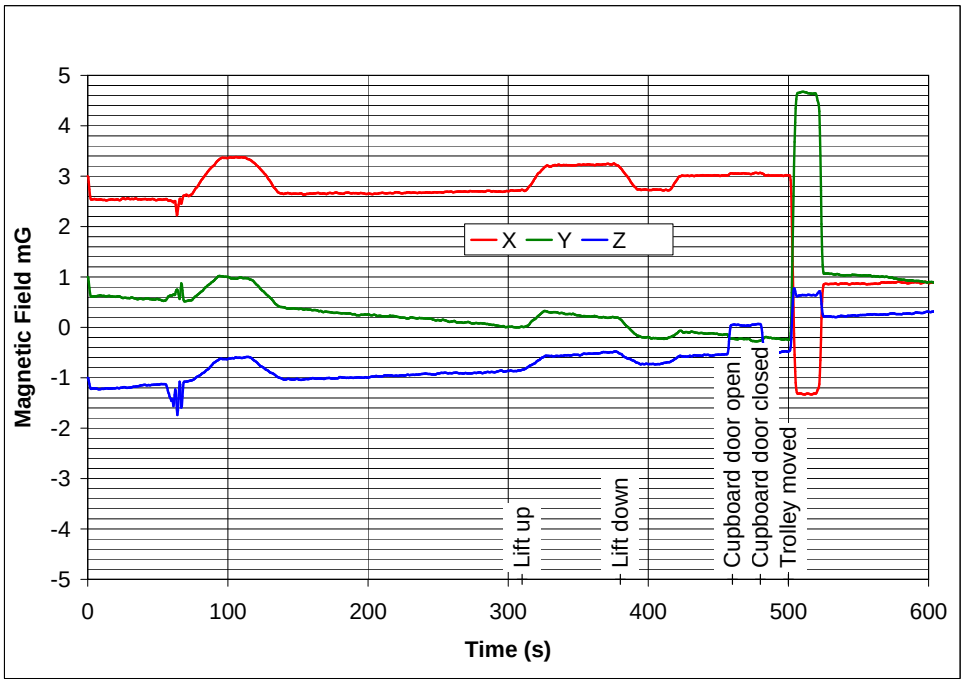

Figure 6.7: Measurements made with a field probe in the lab using an SC10 sensor. The labels indicate different activities in the room.

6.3.4 Mains Supply

The mains electricity supply in the lab is three-phase, where each of the three supplies or "phases", is at 240V, 50Hz and separated by a phase shift of 120°. If equal current were to flow in each phase, then there would be no net field, which is the desired situation. However, there is a residual field which is a combination of the three phases, and depends on the current being drawn on each circuit.

During any experimental run, the load on each phase doesn't vary significantly, but the total phase imbalance will leave a resulting 50Hz magnetic oscillation. The exact magnitude, phase and direction of this field will vary as the current drawn in the lab, and places nearby, varies. Additionally there may be mains supply ripple on the linear power supplies that drive the Zeeman coils.

A set of experiments was done using a line-triggering technique. The delay from a fixed moment in the mains cycle was varied, and the Ramsey experiment was performed. The data

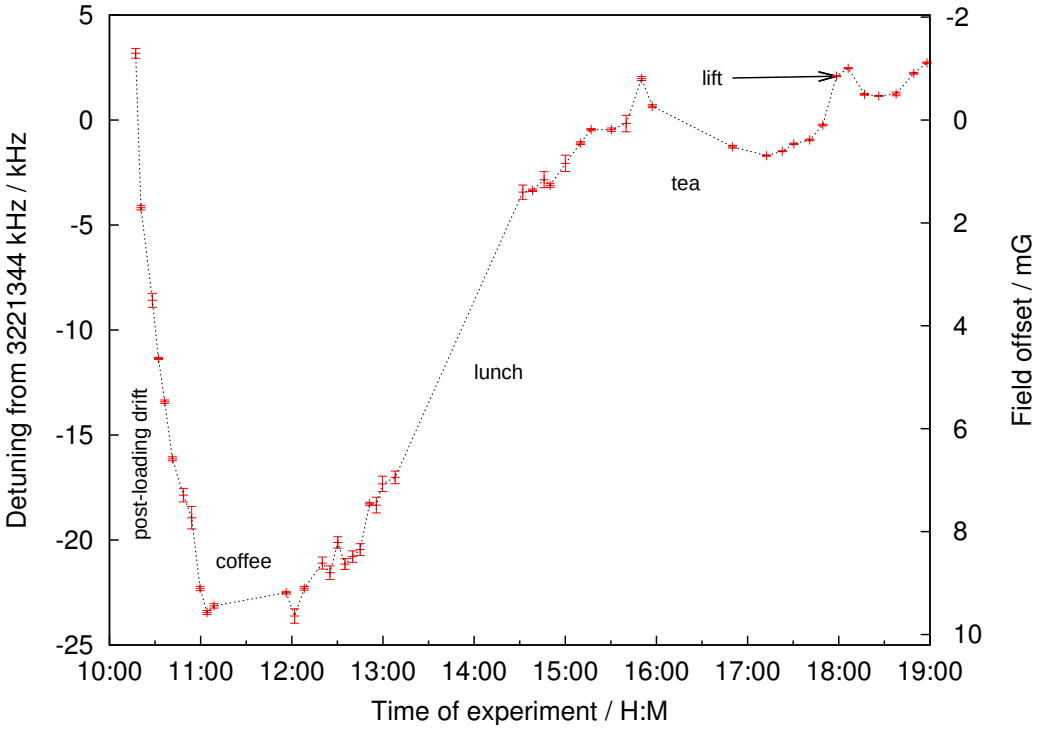

Figure 6.8: Measurements of field-fluctuation at the ion, made using the Ramsey technique. The labels indicate activity that was occurring in the same building. The small peak at 18:00 was due to the lift being operated.

were fitted, and the detuning from the ion's expected resonant frequency was calculated. This was then plotted as a function of delay from the mains cycle time. A sinusoidal curve was fitted to this data with a frequency of 50Hz (the mains frequency). It can be seen that at the ion, a field of about 4mG (peak-peak) from the mains supply is detectable. This is shown in figure 6.10.

Included in this set, slow "DC" field drifts were still present. For comparison, a set of data has been selected when the DC field was relatively stable, labelled "best data". The 50Hz residual field is still evident.

6.3.5 Spicer Field Cancellation

DC Cancellation

Experiments were run with the field cancellation, Spicer SC20, on and off. The results are shown in figure 6.11. The unit creates a small DC field offset, however it can be seen that the field is changing whether the unit is off or on, i.e. it fails to track and null the external field.

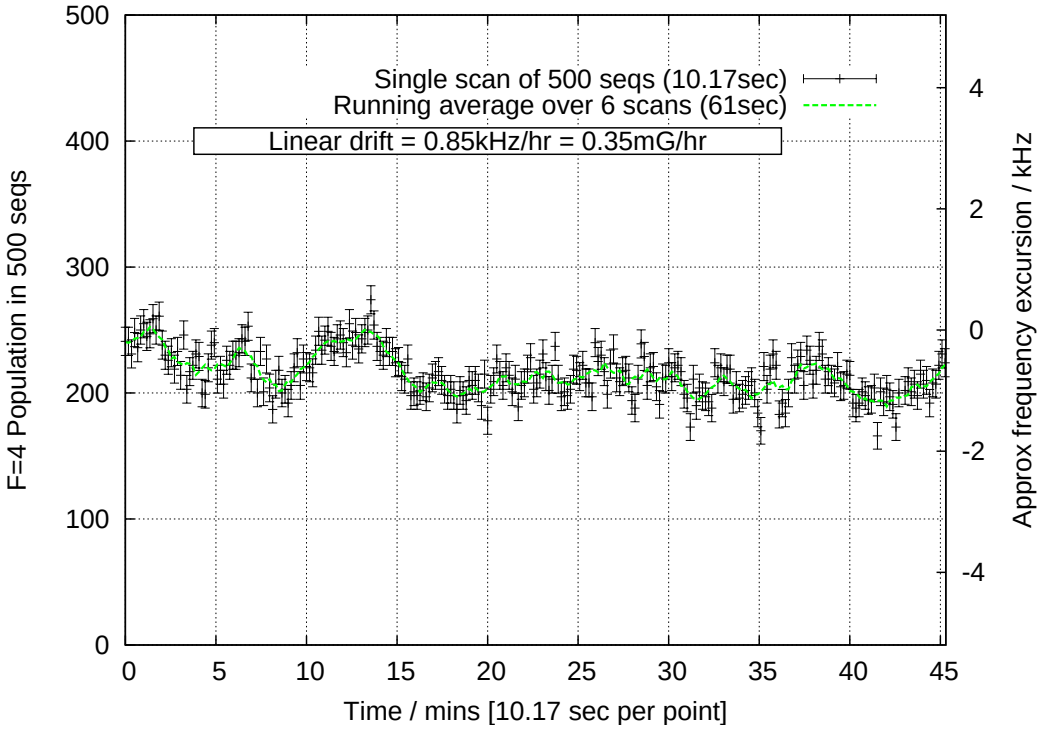

Figure 6.9: Fluctuations of the magnetic field observed by sitting on the side of the $(F=4, M_F=+4) \leftrightarrow (F=3, M_F=+3)$ *resonance and measuring the shelving probability.*

Mains-supply (50Hz) Cancellation

The same experiment performed in section 6.3.4 was carried out with the SC20 enabled. The results are shown in figure 6.12. It can be seen that there is a slightly lower residual 50Hz. This maybe due to ripple on the power supplies which drive the Zeeman coils, and would be closer to the ion than the SC20's sensors (and therefore could not be compensated for).

To see if this residual field could be reduced further, the settings available on the SC20 were investigated. To do this, a faster experiment was devised. The microwaves were detuned to the side of the π-pulse resonance. This allowed the unit's setting to be adjusted whilst monitoring the field. At the HWHM point (see figure 6.1), the slope of the fringe is nearly linear with a normalised value of -0.10(1)kHz^{-1}. This allows a shelving-fraction to be converted to a magnetic field, by multiplying by -2.45kHz/mG which gives 0.25mG^{-1}. Different combinations of settings were tried, and the results are summarised in figure 6.13. The combinations of settings are summarised in table 6.3.5. It can be seen that the best combination is achieved in run 40. However, this still gives a residual field amplitude of 0.46(8)mG. The problem with changing the mixer settings was

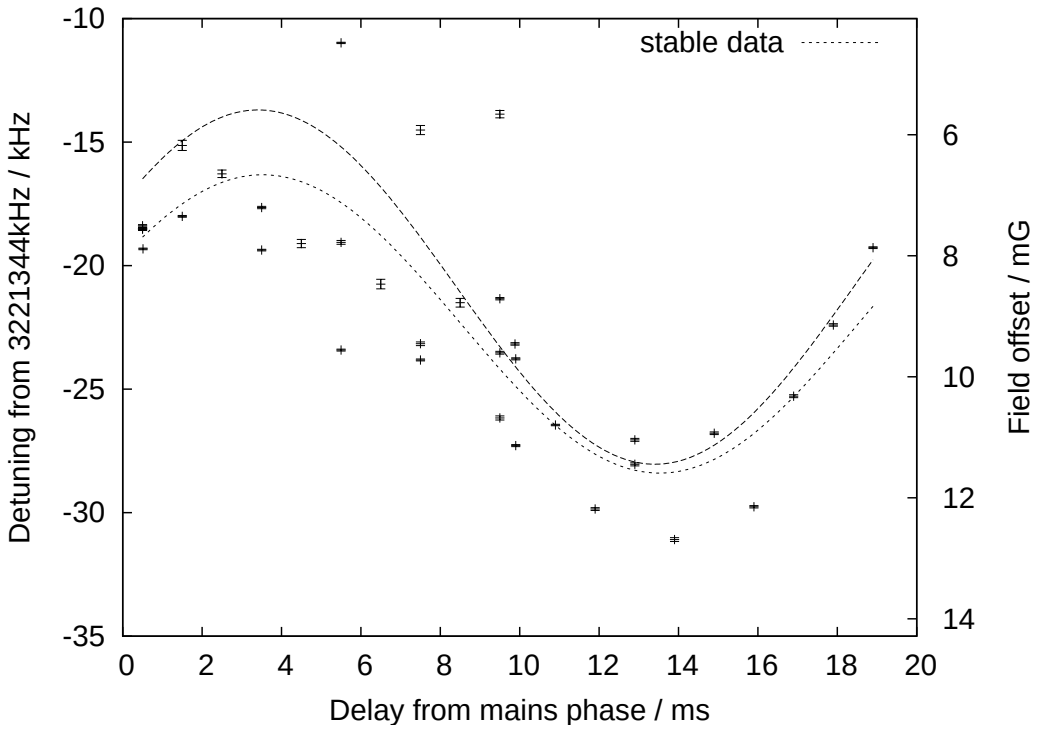

Figure 6.10: Detuning as a function of delay from line-trigger point. Each point is obtained from a fitted Ramsey experiment. The complete set is shown, along with a sinusoidal fit (f=50Hz) as well as data selected at times when the DC field was not varying ("stable data"). The residual mains field is about 5mG peak-peak. No field cancellation unit was used.

that there was no control of the external field. It was assumed that there was a constant 50Hz background and that extra DC fields were steady over the course of changing the mixer settings. To ensure this was the case, the central frequency of the resonance was checked and accounted for periodically.

Spicer Mixer Settings

		Mixers L=-ve R=+ve			Loop Gain		
Run Numbers	Delay /ms	X	Y	Z	X	Y	Z
12 15 19 28 30	0-10	0	0	0	3	-3	12
13	0-10	0	0	-5	3	-3	12
14	0-10	0	0	5	3	-3	12
16	0-10	-5	0	0	3	-3	12
17 18 21 24 44 46 48	0-10	5	0	0	3	-3	12
20	0-10	3	0	0	3	-3	12
22	0-10	5	-5	0	3	-3	12

23	0-10	5	5	0	3	-3	12
25 26 27 31	0-10	5	0	-5	3	-3	12
33 47	0-10	5	0	5	3	-3	12
34	0-10	5	0	5	0	0	0
35	0-10	0	0	0	0	0	0
36	0-10	5	0	0	0	0	0
37	0-10	5	0	5	0	0	0
38 42 43 45	0-10	5	0	0	0	-3	0
39	0-10	5	0	0	0	-3	6
40	0-10	5	0	0	0	-3	-6
41	0-10	5	0	0	0	-3	-12

6.4 Conclusion

Variations of $\pm$5mG in the background field at the position of the ion were observed. The dominant effect was the 50Hz mains field, but this can be reliably allowed for using line-triggering. This corresponds to an energy shift on the stretched state transition of $\pm$12kHz and on the clock-state transition of $\pm$20Hz. With line-triggering this could be reduced to a change of $<$1Hz over 1 hour. The effect of the oven was shown to be negligible after a 1 hour wait.

Use of the Spicer unit did not help cancel this field. This could be due to the detector being to one side of the vacuum can hosting the ion, and therefore it was unable to detect or respond to a field at the ion's position that has diverged from the field at the detector's position. Two detectors did not improve the situation. The field-cancellation unit did not reliably compensate for the mains electricity supply even after trying to optimise the unit's settings, nor did the unit successfully cancel DC field drift. This maybe due to the field having a gradient across the experimental apparatus, and therefore the placement of the detectors would have to be optimised. Also, if the power supplies on the experimental coils were injecting mains ripple or drifting, the SC20 would be unable to compensate this variation. Drifting DC fields were still present, and therefore, the best experimental results were achieved at night when the building was not in use by other people.

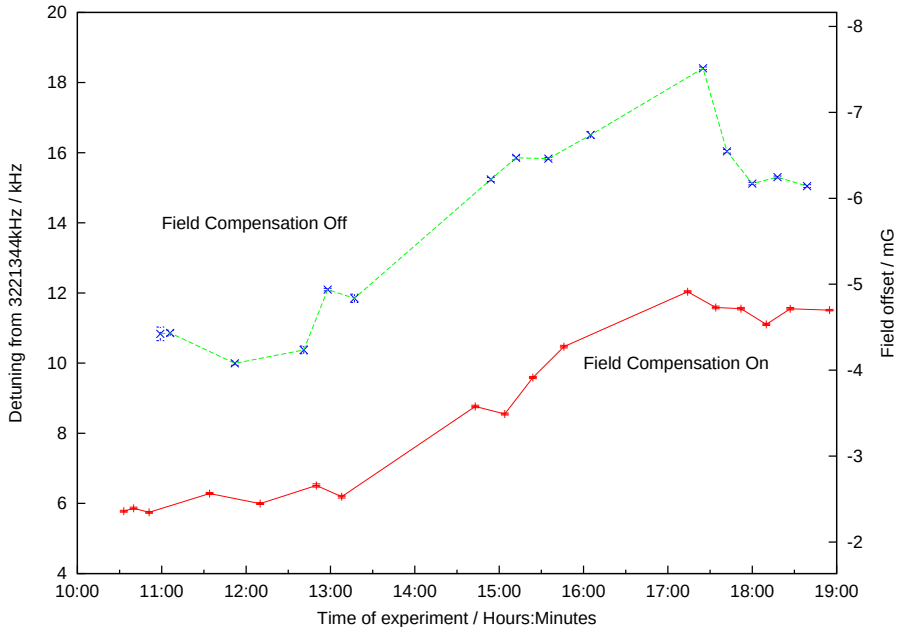

Figure 6.11: Effect of switching on and off active field cancellation. It can be seen that the cancellation unit is not compensating for field-drift but it is creating a local field of about 2mG at the ion. The increase in field between 13:30 and 17:30 was possibly due to a superconducting magnet that was steadily ramped up during this period. The same magnet was immediately ramped down again.

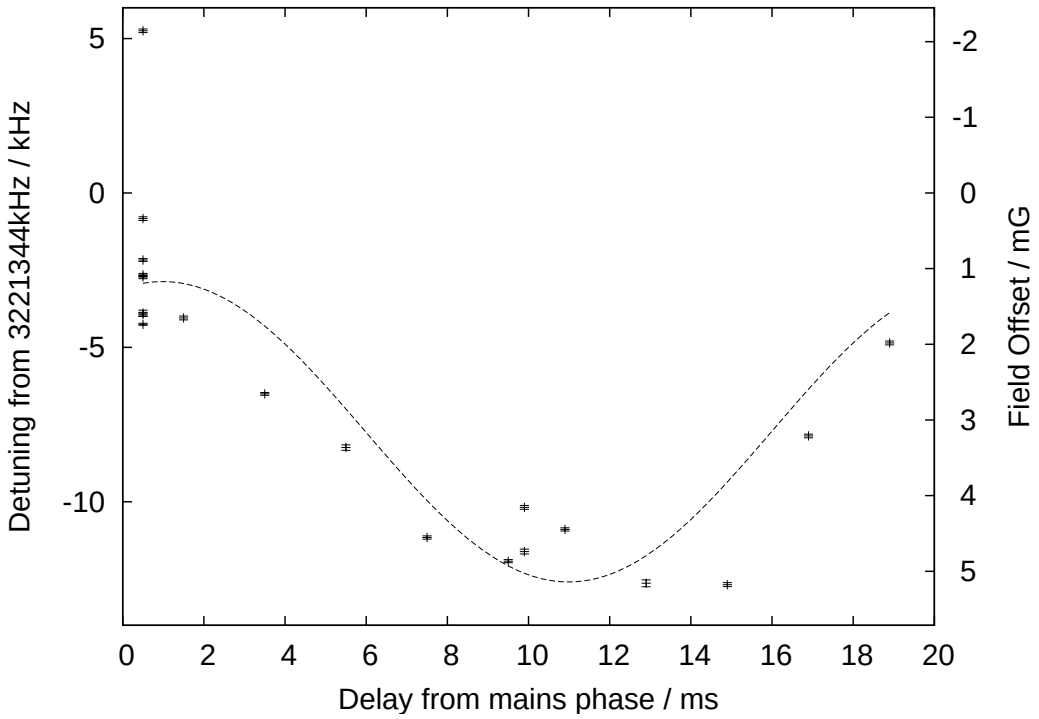

Figure 6.12: Detuning as a function of delay from line-trigger point. Each point is obtained from a fitted Ramsey experiment. A sine curve with f=50Hz, is fitted to the residual field. The SC20 is enabled. The residual mains field is about 4mG peak-peak, slightly lower than when the SC20 was disabled, see figure 6.10.

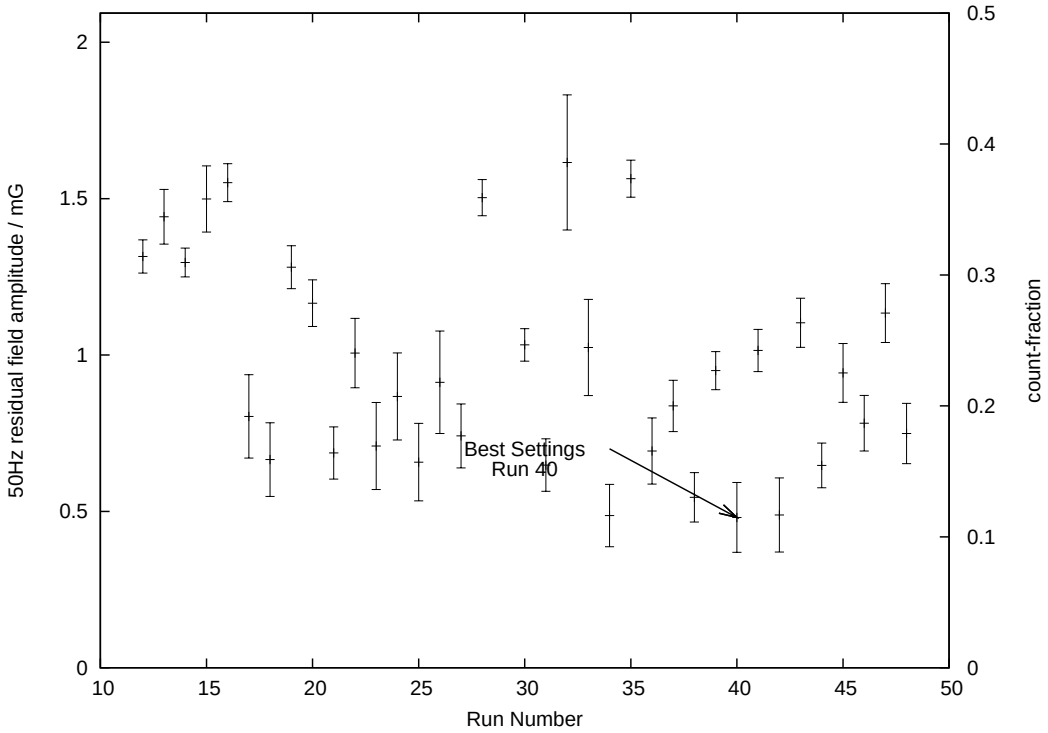

Figure 6.13: Different configurations of active field cancelling unit. To see individual settings for the unit, see table 6.3.5. To allow the effect of settings on the SC20 to be seen, the detuning was held to the left of the atomic resonance, see figure 6.1. This allowed a change in field to be seen as a change in shelving probability.

Chapter 7

A Long-lived Memory Qubit Using The Hyperfine Ground States of $^{43}Ca^{+}$

7.1 Introduction

Current QIP schemes [80] require ions for both processing information, (i.e. involved in gates), and for storage of information between processing i.e. memory qubits. These qubits would play a similar role to registers in a conventional processor. They would hold information temporarily, whilst gate operations on other qubits are carried out. However, they also play another very important role. The no cloning theorem forbids the creation of identical copies of an arbitrary unknown quantum state [79].

The state of one system can be entangled with the state of another system. This is not cloning. No well-defined state can be attributed to a subsystem of an entangled state. Cloning is a process whose end result is a separable state with identical factors. Due to the no-cloning theorem, it is neccessary to transport the physical system from one place to another, without collapsing the encoded state. Photons are one such quantum system, and would enable easy transmission over long distances. However, coupling ions or atoms to photons efficiently is still a difficult task [81], and for the short distances inside the processor, it is easier to move the ions themselves. Therefore memory ions serve two purposes. To hold information between operations, and to transport information from one interaction area to another, within the processor. In a complex

algorithm, requiring thousands of gate operations, the required storage time is several orders of magnitude greater than the gate time. Typical gate-times are currently about 100μs [21]. To reach a required fidelity of 0.9999 this requires a memory coherence time on the order of 1s.

By encoding the information in the internal state of the ion using a long-lived (metastable) state, the information will be robust to movement of the ion. The hyperfine ground state of $^{43}Ca^{+}$ is a very good candidate for a memory qubit.

In this experiment the coherence lifetime of that memory is measured. By investigating what the limiting factors are on the lifetime and coherence, better understanding of noise in other operations, such as gate operations, is also obtained.

7.2 Experimental Detail

The following experiments were performed on the clock-states of $^{43}Ca^{+}$. i.e. $|\uparrow\rangle$ = $S_{1/2}$(F=3,M_F =0) and $|\downarrow\rangle$ = $S_{1/2}$(F=4,M_F =0). These states have no first-order dependence on magnetic field. As was shown in the previous chapter, magnetic field drift is still the greatest source of error in measuring the qubit state, and therefore the largest source of decoherence. The linear field-dependence due to the quadratic Zeeman effect in the clock transition is $\partial f/\partial B \approx$ 4Hz/mG in an external field of 1.8G, which was used in these experiments. This is much lower than the qubit state ($S_{1/2}$, $M_s = \pm 1/2$) in $^{40}Ca^{+}$ which has a Zeeman-splitting of $\sim$ 2.8kHz/mG.

Initially the ion is loaded and Doppler-cooled on the $S_{1/2}$(F=4)$\rightarrow$$P_{1/2}$(F=4) and $S_{1/2}$(F=3)$\rightarrow$$P_{1/2}$(F=4) transitions. These are driven by a 397nm laser with a 3.22GHz sideband that is produced using an EOM (as described in chapter 2). A single-frequency 866nm beam is used to repump from the $D_{3/2}$ levels. It would be possible to prepare the ground clock state, $S_{1/2}$(F=4,M_F =0) by making use of the selection rule that forbids M_F =0 $\rightarrow$ M'_F=0 if ΔF=0. Two π-polarised beams could be used to drive the $S_{1/2}$(F=4)$\rightarrow$$P_{1/2}$(F=4) and $S_{1/2}$(F=3)$\rightarrow$$P_{1/2}$(F=4) transitions until all the population is in $S_{1/2}$(F=4,M_F =0). This is limited only by off-resonant excitation of the $S_{1/2}$(F=4) $\rightarrow$ $P_{1/2}$(F=3) transition. The natural line width of the $P_{1/2}$(F=3) is 22MHz, as opposed to the energy difference of 581MHz between the two F states. Theoretical state-preparation efficiencies of 99.0% can be achieved [37]. However, the required beam direction for the two π-polarised beams was not available, so instead the EOM was switched off, and each of the M_F states were prepared with

$\approx 1/7 \approx 14\%$ probability. The non-clock state results were discarded.

The most sensitive way of measuring the coherence lifetime of the ion is to use a Ramsey experiment (see chapter 6). Starting from the $|\uparrow\rangle$ state, a $\pi/2$ pulse puts the ion into a superposition of $|\downarrow\rangle$ and $|\uparrow\rangle$ states and a second $\pi/2$ pulse should transfer the ion to the $|\downarrow\rangle$ state, assuming zero detuning and zero phase-change between the ion and the local oscillator. The $\pi/2$ pulses are typically 35μs long.

The setup used microwaves introduced at a spare electrode inside the trap, as before. This was driven using an RF synthesiser (Agilent E4426B) which, in the majority of experiments, was locked to a rubidium-referenced quartz oscillator (Stanford FS725) with stability (Allan deviation) below 2×10^{-11} over the timescale of the experiment. The synthesiser was set to output 0dBm then amplified by 30dB and fed via co-axial cable to the feed-through and into the vacuum can. Typically experiments were repeated 500 times for each setting of either phase or detuning, and twenty phase or detuning points were taken for each setting of Ramsey delay time.

During the Ramsey delay time, τ_L, any drift of the resonant frequency will result in the ion undergoing an incomplete rotation and therefore the probability of reaching the final state being less than one. As was shown in chapter 6, the largest cause of change in the ion's resonant frequency is magnetic field drift. The experiment was repeated whilst the microwave detuning was varied, over an amount $\delta \sim \pm 1/\tau_L$ Alternatively, by varying the phase of the second $\pi/2$ pulse relative to the first (in this case using a second synthesiser with a different phase to generate the microwaves for the second $\pi/2$ pulse) a Ramsey fringe was also observed. If the delay time, τ_L, was changed and there was no detuning, no fringes would be apparent. Changing the phase ensures that fringes are always visible even at zero detuning. Over the timescale of the experiment, the detuning changes due to the drift in stray magnetic fields. This results in a reduction in fringe amplitude if the drift in detuning is comparable with $1/\tau_L$.

7.2.1 Check Experiment

To control for drifts in, for example, readout efficiency, each Ramsey sequence was interleaved with a check sequence. This was a very short Ramsey sequence, with the time between the two $\pi/2$ pulses being only τ_L=0.145ms, which meant that the short Ramsey sequence effectively con-

stitutes a single π pulse. This gives a fringe amplitude against which the longer experiments can be normalised. The check experiment was executed immediately before data from the long Ramsey sequence was collected, and was subject to the same environmental conditions, e.g. background magnetic field. These control experiments showed no detectable drift even during the longest experiments whose total run time was 54min.

7.2.2 Spin-echo Technique

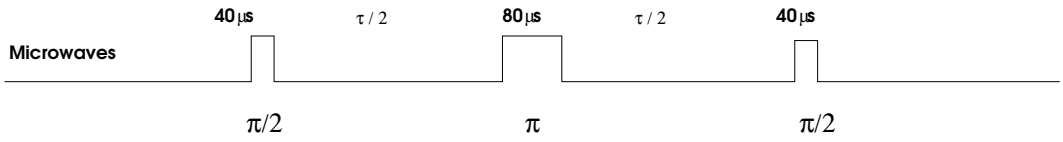

Figure 7.1: Spin-echo pulse sequence

The coherence of the qubit was also measured using the spin-echo technique. It is possible to undo the effect of a constant non-zero detuning, using the spin-echo technique. By inserting a π pulse in the middle of the two $\pi/2$ pulses (see figure 7.1), the ion's state is rotated 180° about a horizontal axis on the Bloch sphere and the effect of the detuning is reversed (see figure 7.2). After an equal amount of time ($\tau_L/2$), the effect of the detuning is removed. This process re-phases the qubit, i.e. removes this phase-change. After the second $\pi/2$ pulse the ion will be in the same state in which it started. This works well for slow drifts that are longer than the duration of a single experimental sequence, but can not compensate for changes of magnetic field during the delay time, τ_L. In these experiments, after the work in chapter 6, it was expected that the slow drift in the lab's magnetic field would be the main source of decoherence, once mains supply effects had been removed using line-triggering.

7.3 Results

7.3.1 Data Analysis

The result of a typical Ramsey experiment is shown in figure 7.3. The check experiment results are also shown, the average of which is used to normalise the 200ms fringe amplitude. For each experiment, the results were fitted with the following equation using the Levenberg-Marquardt

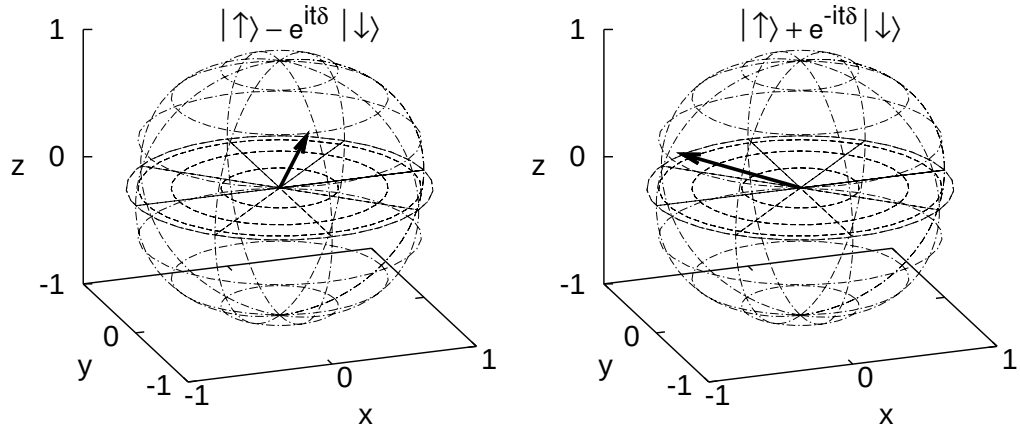

Figure 7.2: Bloch sphere showing a superposition state that is detuned by δ from the local oscillator (the rotating frame) after some time, $\tau_L/2$. The second sphere shows the effect of applying a π pulse, i.e. a 180° rotation around the y-axis. If the detuning remains constant, this phase will have returned to zero after the same time, $\tau_L/2$ and the second $\pi/2$ pulse returns the qubit to its original state.

method as implemented in `gnuplot`:

$$P = P_0 + A\,(1 - \cos[\tau(2\pi f - \phi)])$$

The value P_0 allows for a baseline offset, as explained below. The amplitude, A, is normalised to the check experiment and is expected to be 1, unless decoherence effects have reduced it. τ is the period of the Ramsey fringes and is expected to equal τ_L. However it was allowed to float, and found not to be equal, also explained below. f is the microwave offset frequency and ϕ is the phase-difference between the ion's resonance and the microwave source.

The points were weighted with a value, $w = 1/\sigma^2$. The standard deviation, σ, is calculated using the binomial formula:

$$\sigma = \sqrt{NP(1-P)}/N$$

where N is the number of trials (typically 500) and P is the shelving probability.

However, for a P of zero, this gives an infinite weight, which would cause the Levenberg-Marquardt algorithm to fail. To circumvent this problem three approaches were tried, and the results were found to be in good agreement. The first approach is to ignore those data points with

infinite w. However, they are statistically significant, and this gives a larger final error. The second method was to assign an error of less than σ for $P = 1/N$, but large enough to be significant for those points that should have zero error. A value of $P = 1/(2N)$ was chosen. Thirdly, a statistical method of modifying the error was used. A modifier was deliberately introduced that skews all the results, but keeps the relative weights the same. This has to be done with some care. The addition of a constant, k, to the standard deviation was used. This constant was chosen so as to be large enough to be significant, but not to be so large as to dominate the statistical results. This was done by minimising the coefficient of variation, $C_V = \frac{\sigma}{\mu}$ where μ is the mean, which for a normalised binomial distribution is equal to P. The value was taken from the domain of possible values for $k \in \{P\}$. The value of k which minimises C_V was found using numerical methods. The modified function for the error used was:

$$\sigma' = \frac{\sqrt{NP(1-P)} + 0.18}{N}$$

This approach gave good agreement with the other two methods, see section 7.3.3.

7.3.2 Ramsey Experiments

Figure 7.3 shows a typical Ramsey fringe, with no spin-echo pulse and τ_L=200ms. The fitted amplitude relative to the mean of the check experiment is 0.83(1). The fitted value for $1/\tau$ = 6.13(2)Hz, differing from the expected value of $1/\tau_L$ = 5.00Hz. This can be explained due to the second-order Zeeman shift on the ion of 4.33Hz/mG at the 1.78G field used. A linear field drift of +1.3mG/ hour over the 37 minute duration of the experiment would account for the incorrect fringe spacing, but simulations show that the amplitude would be reduced by a negligible amount. However, the reduction in amplitude is consistent with faster field fluctuations at the level of 0.25mG over the timescale of a few minutes. Both drift and fluctuations at these levels were demonstrated in chapter 6, see figure 6.9.

In the first set of experiments, Ramsey sequences were performed with no spin-echo pulse, and τ_L was varied. These results are shown in figure 7.4. These data show several different datasets taken on different dates, including over a holiday, and with and without an active field compensation unit, the SC20.

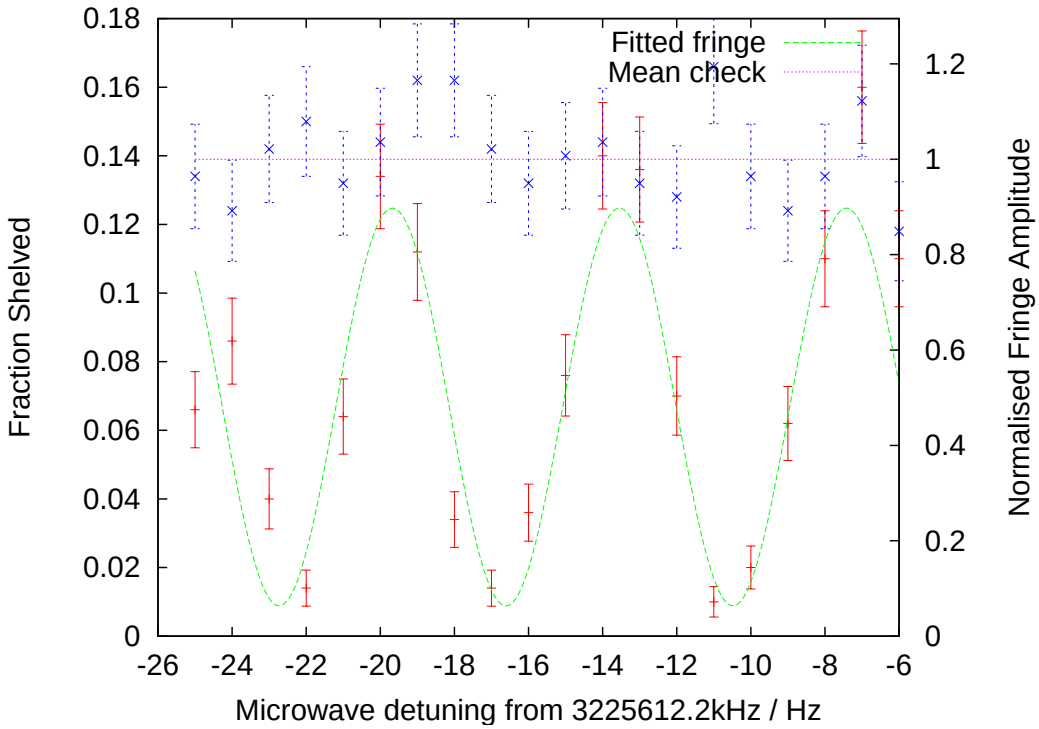

Figure 7.3: A typical fringe from a Ramsey experiment, with τ_L=200ms, performed on the clock transition in $^{43}Ca^{+}$. The maximum amplitude only reaches $1/7$ as the preperation method only prepares the desired state $1/7$ of the time.

Figure 7.5 shows the various conditions for each day. These were as follows:

1. The first data taken in June did not use any of the improvements made in subsequent experiments.

2. In December, shutters were added to both the 393 and 397 beams. The rubidium clock was used to provide better stability to the synthesiser. Field compensation was used for some of the runs but did not make a noticeable improvement.

3. In February, the 3.2GHz EOM was added, as described in chapter 2. Previously two separate 397nm lasers had been used for hyperfine repumping. The switching between two synthesisers ("Multiplexer") was added to allow the phase to be varied instead of the detuning. However, in February it was discovered that the 393nm double-pass AOM was not completely extinguishing the beam and this accounts for the two low values for T_2 times. On the 15th February the SC20 was switched on and off with no result.

4. In April the coil power-supplies were replaced with better regulated supplies. Again no discernible difference was measured. On the 17th April, the shutters were again placed on the 393 and 397 beams.

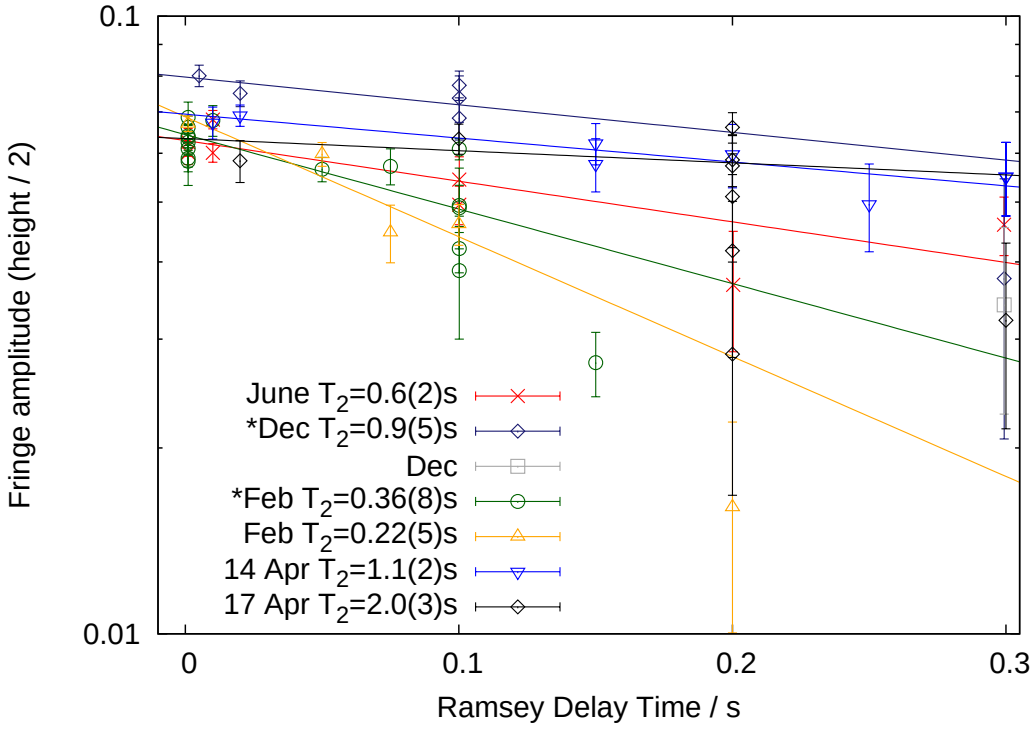

Figure 7.4: Summary of decoherence rates for a $^{43}Ca^{+}$ memory qubit on different days. The field compensation, SC20, was off except for those runs marked with an asterisk, which were analysed separately. The data for each day is fitted with an exponential decay, with time-constant (T_2) and error given in the key, in seconds.

The data was fitted with an exponential decay, assuming that the fringe contrast decays with increasing decoherence. The fitted value for the function, $\alpha\exp(-\tau_L/T_2)$ gives a value of T_2 = 1.2(2)s for the internal state coherence time for results obtained from all runs (excluding February, for the reasons given above). This compares favourably to gate operations which take on the order of 100µs, giving a ratio of the two of 10^4. This is comparable to other techniques such as using decoherence-free subspaces, which are not as advantageous, due to the necessary coding and decoding overhead [82].

The reduction in fringe amplitude can be accounted for by fluctuations in the local magnetic field, or other processes such as photon-scattering from imperfectly extinguished laser beams.

These project the qubit, making a measurement and collapsing the superposition. For some experiments shutters were used on the 397 and 393 beams. Taking just data with the shutters present (see figure 7.6) showed an increase in T_2 to 1.6(4) s compared to the data without shutters, T_2 = 0.8(2) s (excluding February). This could be explained due to imperfect extinction using AOMs, but is of marginal significance.

The spin-echo technique will cancel de-phasing due to field-drift whilst not compensating for other effects such as photon-scattering. Therefore this method was tried next.

7.3.3 Spin-echo

In order to see if the decoherence of the ion was due to magnetic field drift or some other cause, the spin-echo technique was used. Figure 7.7 shows a typical fringe for the spin-echo sequence as described in 7.2.2. The phase between the two $\pi/2$ pulses is being varied, therefore both the check

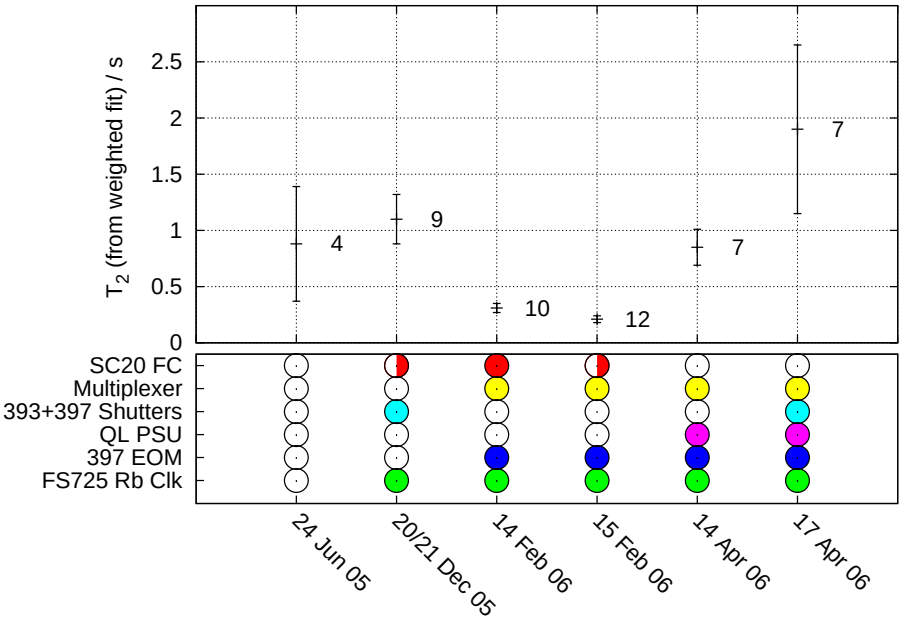

Figure 7.5: Decoherence rate calculated from fitted exponential decays, as shown in figure 7.4. The lower plot shows the different conditions for that day's run, where a filled circle indicates the condition was true and an empty circle indicates the condition was false. A half-filled circle indicates the condition was true for some of the runs on that day. The number to the right of the data-point is the number of experimental runs.

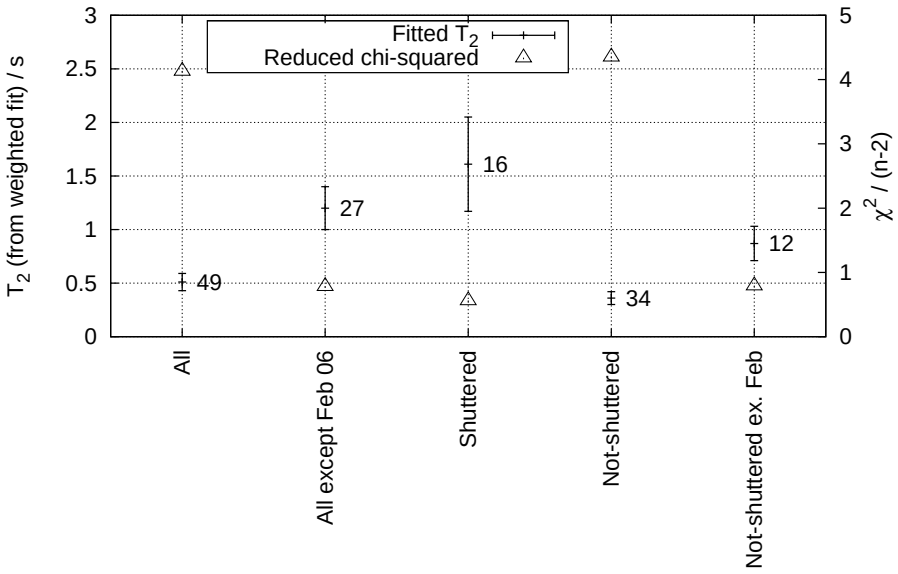

Figure 7.6: The same data as shown in figure 7.5 grouped to show the effect of shutters and of excluding the February data. The reduced χ^2 values are shown to give a measure of fit. The χ^2 values for data that includes February data is suspiciously large. When the February runs are excluded the difference between shuttered and non-shuttered T_2 values is barely significant.

experiment and the actual experiment with a delay time of τ_L are fitted with sinusoidal functions. As before, the fringe amplitude is found by fitting and this is used as a measure of decoherence after being normalised to the check-sequence result.

The results of the spin-echo experiments are summarised in figure 7.8. This gives longer coherence times than for the Ramsey experiment, as the detuning is compensated for by the π pulse. This implies that the cause of the decoherence is due to a detuning at the ion (due to stray fields) and not due to another factor such as scattered light. For the longest pulse sequences tried, τ_L=1s, no measurable decoherence can be observed (see figure 7.7). A straight line-fit to the data gives a value of 0.98(2) for the fringe amplitude, with a gradient of $-(0.04 \pm 3) \times 10^{-5}$ (ms)$^{-1}$ which implies the rate of amplitude and phase decoherence of the qubit, due to photon scattering or background collisions, is negligible on the 1s timescale. If, as above, it is assumed that decoherence can be modelled as an exponential decay, this implies an effective coherence time of $T_2 > 45$s.

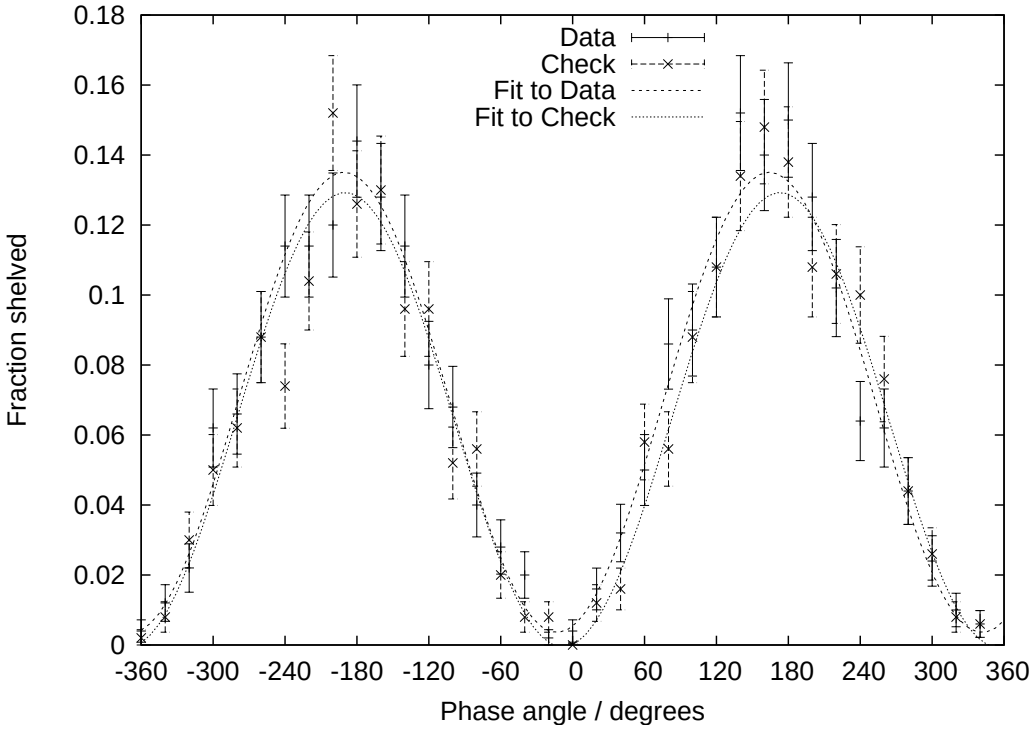

Figure 7.7: Fringe from spin-echo sequence with a delay time of τ_L = 1.0s. The phase between the two $\pi/2$ pulses is being varied instead of detuning. The fitted check-fringe with a delay time of only 0.30 ms is also shown. The ratio of the two fringe heights is 1.01(4). The two amplitudes agree within errors showing that the decoherence after 1s is not detectable. The baseline, however, shows a slight increase. The total length of the run was over six hours.

To compare the different fitting methods discussed in section 7.3.1, the data was fitted with a sinusoidal function and weighted using the three different methods, and the data fitted with a straight line fit of $A = m\tau_L + c$. All three methods give good agreement.

With zero counts:	$m = (-0.6 \pm 3.1) \times 10^{-5}\,(\mathrm{ms})^{-1}$	$c = 0.98(2)$
Without zero counts:	$m = (0.2 \pm 3.3) \times 10^{-5}\,(\mathrm{ms})^{-1}$	$c = 0.98(2)$
Using modifier method:	$m = (-0.03 \pm 3.1) \times 10^{-5}\,(\mathrm{ms})^{-1}$	$c = 0.97(2)$

Although the amplitude does not show any drop, up to a delay time of 1s (the amplitude is >97%), the fringe baseline is increasing as the delay time increases. The baseline increase is $0.03(2)\,\mathrm{s}^{-1}$ as a fraction of the amplitude. This can be seen in figure 7.9. This is conceivably due to the contribution from all seven of the $S_{1/2}$(F=3) states to the signal. It could be a loss in ion fluorescence during the long delay-time, which would be recorded as an ion in the $|\downarrow\rangle$ state. This could be due to the ion heating up during the long delay-time. Alternatively, if due to real

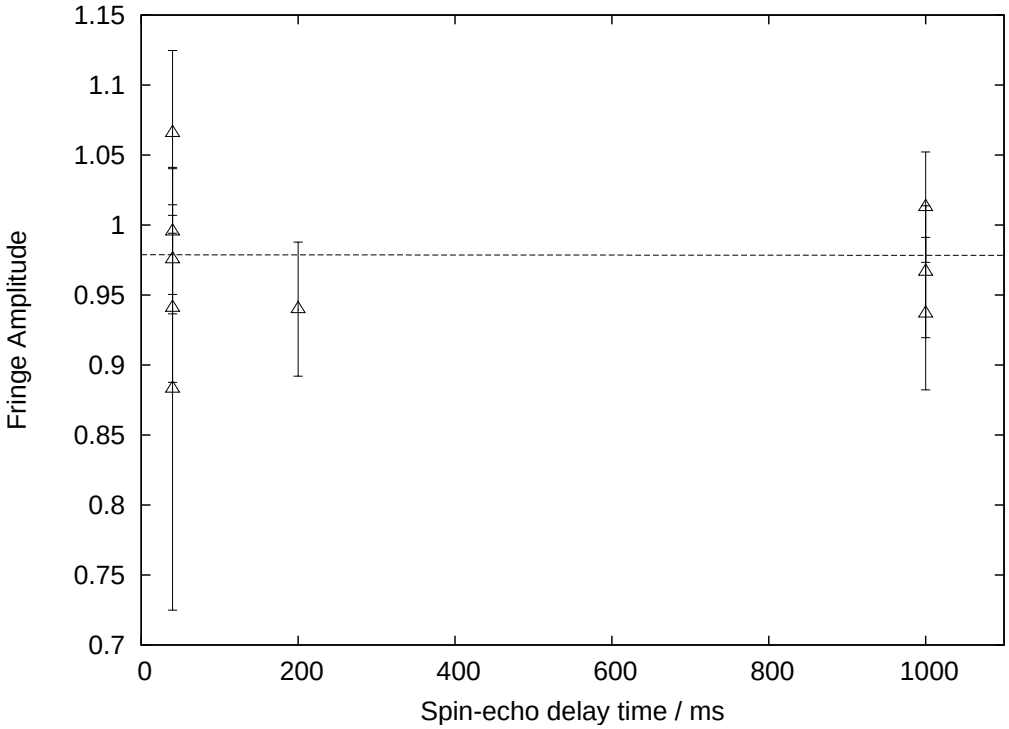

Figure 7.8: Normalised amplitudes for spin-echo experiments with differing delay-times. The straight-line fit shows no measurable decoherence on this time-scale.

decoherence effects, the decoherence would be in the order of 0.5% in 1s, which is smaller than the limit on the decoherence discussed above, implying $2\,\text{min} \lesssim T_2 \lesssim 10\,\text{min}$ (assuming exponential decay).

7.4 Conclusion

The hyperfine states of $^{43}Ca^{+}$ make an exceptionally good memory qubit. With the spin-echo technique, it is possible to obtain a coherence lifetime of order one million times higher than gate times. The spin-echo technique is expected to work well with slowly varying fields, such as a back-ground magnetic field. Because this technique works so well, it confirms the results in chapter 6 that the main source of decoherence on the ion is the local magnetic field. It is concluded that other possible sources of decoherence such as photon-scattering are negligible for time-scales $\lesssim$1s.

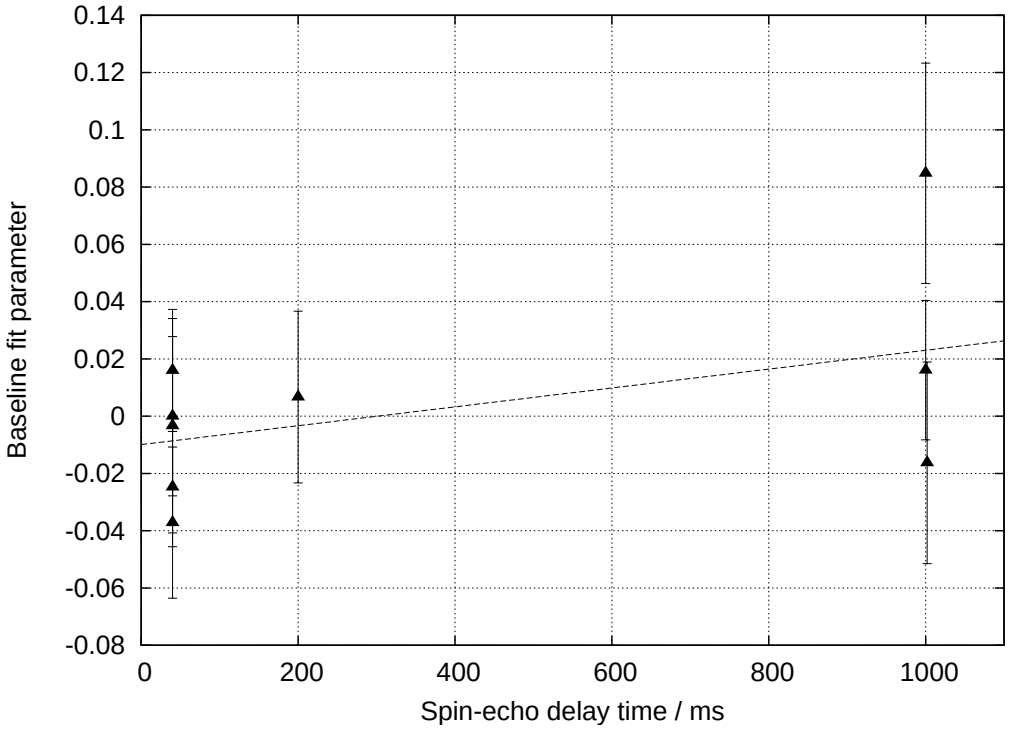

Figure 7.9: Fitted baseline parameter, P_0, against spin-echo delay time. A straight line fit gives an intercept value of -0.01(1) and a gradient of 0.03(2) s^{-1}.

Chapter 8

Conclusion

This thesis has looked at the systems and experimental techniques required to implement QIP in $^{43}Ca^{+}$, specifically the requirements and limitations to implement a long-lived memory qubit using the hyperfine structure.

Chapter one presented an overview of QIP in general. Cold atoms and trapped ions are proving to be the best technology for realising scalable quantum computing. The atomic system of $^{43}Ca^{+}$ was looked at and compared to previous work done at Oxford with $^{40}Ca^{+}$. Using the hyperfine ground states was presented as the best way to store quantum information due to their long lifetime and robustness. The problem of decoherence due to external fields was discussed.

Chapter two presented the current experimental setup including the electrode system that keeps the ion in place as well as allowing coupling between ions. The atomic system in $^{43}Ca^{+}$ was described. The current laser systems were presented along with the microwave system that allow direct manipulation of the ion's hyperfine state.

Chapter three described the design and construction of a new piece of equipment, the LCU, which allows experiments to be performed that required complex, precisely timed pulse sequences. This has allowed the experiments presented in chapters 6 and 7 to be carried out. It has also been essential for other work at Oxford, such as the motional-gate and sympathetic cooling. The LCU met the necessary timing requirements by providing a pulse resolution of 0.1 μs and effectively an unlimited length of pulse sequence. It also provides a high-degree of flexibility via a high-level programming language and has proven to robust and reliable with only one timing error detected

in three years (in over one billion pulses).

Chapter four was concerned with the building of an optical-injection system which allows the Raman transitions between two hyperfine states to be driven as a better alternative to the microwaves used presently. Optical injection was achieved by using the side-port of a Faraday isolator and the offset was generated using a double-pass AOM at 1.6GHz. The laser system was built and thoroughly characterised. It was shown to yield a power of 20mW in the frequency-shifted beam and 12mW in the other beam, (measured after the Faraday isolator). This provides ample power to drive transitions at a fast-enough rate. The frequency offset showed a stability of better than 10Hz in 3.2GHz, the limit of this measurement being the measuring equipment. Wavelength tunability of over 1.5nm was obtained and the injection remained stable for more than a day of operation.

Chapter five contained information about the design and implementation of a Pound-Drever-Hall locking system for the lasers used in the experiment. The system relied on a stabilised reference cavity provided by NPL (which had a specified stability of $\lesssim$500kHz/hr) and high-speed RF electronics. The system worked at a modulation of up to 200MHz and used cavities with a FSR of 1620MHz and low finesse of $\sim$ 30. This gave a very broad capture range, and a sufficiently fast response. This resulted in locks that were very robust against mechanical and thermal changes, but responded quickly to frequency fluctuations, resulting in a laser than could remain locked for over a day in a typical lab environment. This was essential for the very long coherence time experiments presented in chapter 7 allowing unattended run as long as 6 hours. To implement the PDH scheme a photodiode circuit working at speeds up to 200MHz was required. This was designed and built. The electronics gave an overall SNR of 10^4 which is sufficient to contribute negligible ($\ll$100 kHz) frequency noise to the laser. When locked, the laser's frequency jitter was $\sim$1.1MHz RMS. These laser-locks have been implemented on the infrared lasers in the experiment, and have been used for some single qubit experiments. The main benefit has been to vastly improve the reliability of the experiment, reducing error counts, time spend re-locking lasers, and ion loss due to the Doppler laser losing lock.

Chapter six investigated the effect of stray fields on a $^{43}Ca^{+}$ trapped ion. The stretched state transition $S_{1/2}(F=3, M_F=3) \leftrightarrow S_{1/2}(F=4, M_F=4)$ was used as it has a relatively large field-

dependence of -2.45kHz/mG. By using Ramsey pulse sequences, produced with the LCU described in chapter three, the change in transition frequency due to the field at the ion could be accurately measured. Both slow drifts (on the timescale on minutes) and fluctuations at the mains frequency (50Hz) were investigated. It was found that the mains supply was a source of significant fluctuations, giving about 4mG peak to peak in a 50Hz cycle, but these could be avoided by synchronising the experiment with the same point in the mains cycle (line-triggering). Slow drifts in the field were attributed to large ferrous objects being moved and nearby superconducting magnets, as well as electrical equipment. An attempt to use a commercial field cancellation system (Spicer SC20) did not result in any consistent improvement. This may be due to the inability of the system to accurately monitor the field at the position of the ion. The field was seen to drift by 4-8mG over the day. Changes on the 1mG level over shorter timescales were also observed. Much larger field-changes ($\sim$50mG) due to the oven used to load the trap were discovered but this was shown to decay to a negligible level after one hour.

Chapter seven drew on work in the previous chapters. Measurements of the decoherence time of a qubit stored in the hyperfine ground state of $^{43}Ca^{+}$, using the clock state transition $S_{1/2}(F=3, M_F=0) \leftrightarrow S_{1/2}(F=4, M_F=0)$, were presented. The decoherence time, T_2, of this qubit was measured and was found to be 1.2(2)s. The main source of error was expected to be the magnetic field drift, based on the results in chapter 6. This was factored out of the experiment (to first order) by performing spin-echo pulses. With these pulses in place, the decoherence was undetectable given the error bars in the experiment and led to the conclusion that the T_2 time (spin-echo) is better than 45s. Spin-echo experiments with a delay time of over 1s were performed that showed the ion was still in a coherent state, with a fringe contrast $>$98%. This also showed that the decoherence (to this level) was not due to other factors such as incomplete extinction of laser beams, or heating of the ion. The Ramsey and spin-echo results are summarised in figure 8.1. This is the lowest decoherence ever measured for a single physical qubit, as far as the experimenters are aware.

In conclusion, the necessary lasers, timing and control system for performing experiments in $^{43}Ca^{+}$ trapped-ions have been realised, and single qubit experiments have been performed, for the first time in this ion. The main limitation currently is decoherence due to stray magnetic fields,

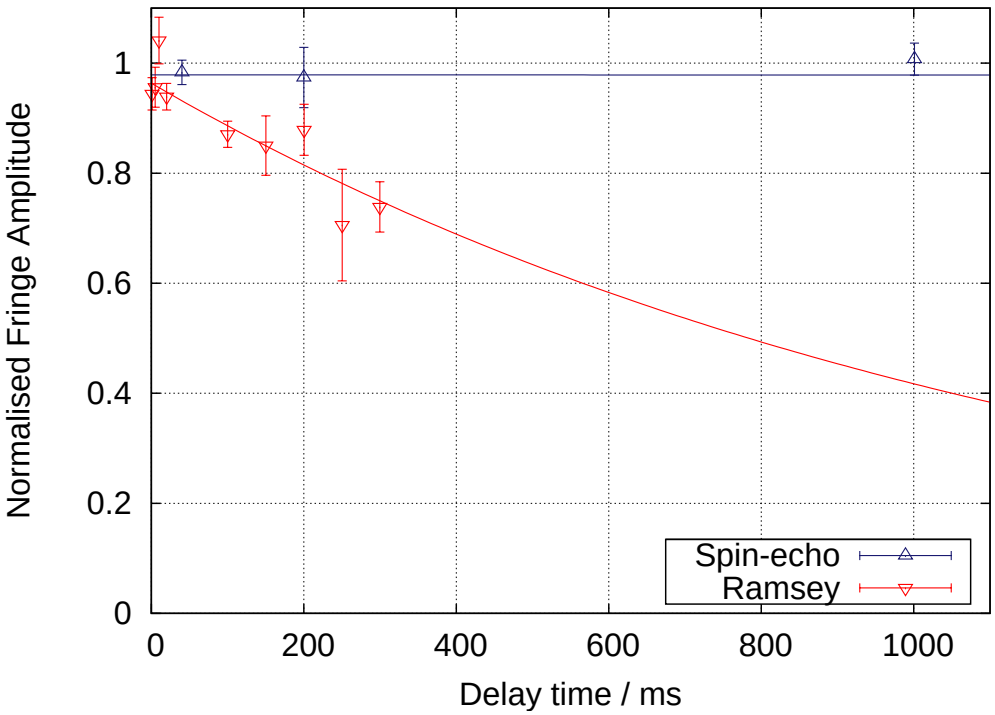

Figure 8.1: Comparison between T_2 times from Ramsey and spin-echo sequences. For clarity, a weighted average at each time-point has been taken, with the error bar derived from the individual errors. The spin-echo data is fitted with a straight line fit that shows no decoherence out to 1s (within errors). The Ramsey experiment is fitted with an exponential decay with a decay-constant of $1.2(2)s^{-1}$, this has been extrapolated out to 1s for comparison with the spin-echo result. The spin-echo result is markedly better, implying that the main source of decoherence is magnetic-field drift.

but this can be compensated for using the spin-echo technique and line-triggering, and is already several orders of magnitude better than for other operations such as two-qubit gates or movement of ions. To reach the necessary fidelity of 0.9999 a $T_2 > 10^4\times$ the gate time is required, and this has now been demonstrated.

8.1 Future Directions

Quantum memory is a fundamental building block in multi-particle QIP that relies on the movement of ions to isolate qubits during gate operations. With high-fidelity memory now realised, the next steps are to perform QIP in multiple well traps such as those discussed in chapter 2. Already ions have been trapped, held and moved around in the Sandia trap (see figure 2.2). The next stages

are:

1. The movement and control of ions around a multi-well trap (processor)
2. The implementation of gates in such a processor
3. Measurements of coherence of a memory qubit after such movement

Once these milestones have been achieved, scalable quantum processors with trapped-ions will have become a reality and algorithms that out-compete the best conventional computers can be implemented. Feynman and others' dream from the 1980's will have become a reality in less than two decades. "The discovery of the computer and the thinking about computers has turned out to be extremely useful in many branches of human reasoning" [83]. No doubt the same will be said for quantum computers in the near future.

Glossary

AOM	Acousto-Optic Modulator
APD	Avalanche Photodiode
BIOS	Basic Input/Ouput System
BNC	Bayonet Neill-Concelman
DMA	Direct Memory Access
DOS	Disk-Operating System
ECDL	Extended Cavity Diode Laser
ECP	Extended Capability Port
EMC	ElectroMagnetic Compatibility
EOM	Electro-Optic Modulator
EPP	Enhanced Parallel Port
GBW	Gain-Bandwidth product
GNU	Gnu's Not Unix
HCL	High-speed CMOS Logic
HWP	Half-wave Plate
IBM	International Business Machines

IDC	Insulation Displacement Connector
ISA	Industry Standard Architecture
LED	Light Emitting Diode
PCB	Printed Circuit Board
PCI	Peripheral Component Interconnect
PIN	Positive-Intrinsic-Negative
PMT	Photomultiplier Tube
PSU	Power Supply Unit
QEC	Quantum Error Correction
QWP	Quater-wave Plate
SMA	SubMinature connector-A
TTL	Transistor-Transistor Logic
UTP	Unshielded Twisted Pair
VGA	Versatile Graphics Array
XLR	X-series Latch (Rubberised)

The following citations are followed by the page number on which they occur.

Bibliography

[1] "Feynman Lectures on Computation", Richard P. Feynman, Perseus Books, ISBN-13: 978-0738202969 (2000) 7

[2] "Quantum Computational Networks", D. Deutsch, Proc. R. Soc. A 425 73-90 (1989) 7

[3] ""Error correcting codes in quantum theory" A. M. Steane, Phys. Rev. Lett. 77, 793-797 (1996) 8

[4] "Demonstration of a Fundamental Quantum Logic Gate", Monroe et. al. Phys. Rev. Lett. 75 4714-4717 (1995) 8, 13

[5] "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer." Peter W. Shor. arXiv:quant-ph/9508027v2 (Jan 1996) 8

[6] "A fast quantum mechanical algorithm for database search", L.K. Grover, Proceedings, 28th Annual ACM Symposium on the Theory of Computing, p. 212, (May 1996) 9

[7] "Factored RSA640 by GNFS", F. Bahr, M. Boehm, J. Franke, T. Kleinjung, http://www.crypto-world.com/announcements/rsa640.txt, 04 Nov 2005 8

[8] "On the cost of factoring RSA-1024", Adi Shamir, Eran Tromer, RSA CryptoBytes, vol. 6 no. 2, 10-19, 2003 9

[9] "Fast quantum search algorithms in protein sequence comparisons: Quantum bioinformatics", Hollenberg, Lloyd C. L., Phys. Rev. E, 62,7532–7535, Nov 2000 9

[10] "Cold Atom Quantum Simulators", Jim Gubernatis and Gerardo Ortix, LANL Research Higlights (2005) 9

[11] “Pushing the Limits”, Paul A. Packan, Science 285 (5436), 2079, (24 September 1999) 9

[12] “Electron spin relaxation of N@C_{60} in CS_2” J. Morton et al. J. Chem. Phys. 124, 014508 (2006) 11

[13] “The Ion Trap Quantum Information Processor”, Andrew Steane, arXiv:quant-ph/9608011, 9 Aug 1996 12

[14] “Quantum Computing”, Andrew Steane, Rep. Prog. Phys. 61 117-173, (1998) 13

[15] “Fullerene-based electron-spin quantum computer”, W Harneit, Phys. Rev. A 65, 032322 (2002)

[16] “ Scaling and Suppression of Anomalous Heating in Ion Traps”, L. Deslauriers et al. Phys. Rev. Lett. 97, 103007 (2006) 13

[17] “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate”, Wineland et al, Letters to Nature, 422, 27 March 2003 13

[18] “Architecture for a large-scale ion-trap quantum computer.”, Wineland, Monroe and Kielpinski, Nature, 417, 13 June 2002 12

[19] “Experimental entanglement of four particles”, Monroe et al, Letters to Nature, 404, 16 March 2000 13

[20] “Transport of quantum states and separation of ions in a dual RF ion trap”, Rowe et al, Quantum Inform. Comput. July 2002, Vol. 2 Ed. 4 12

[21] “”Deterministic entanglement and tomography of ion spin qubits” J. P. Home, et al. New J. Phys. 8, 188 (2006) 13, 14, 22, 131

[22] “Isotope-selective photoionization for Calcium ion trapping”, Lucas et al. PRA 69, (2004) 19, 28, 32

[23] “Raman Sideband Cooling and Coherent Manipulation of Trapped Ions”, S. Webster, PhD Thesis, University of Oxford, (2005) 14, 21, 25, 31, 33, 35, 83, 102

[24] “Quantum computations with cold trapped ions”, Cirac, J. I. & Zoller, P. Phys. Rev. Lett. 74, 4091-4094, (1995)

[25] “The Physical Implementation of Quantum Computation”, David P. DiVincenzo, IBM Publication, 23 May 2006 9

[26] “Quantum Simulations of Physics Problems”, Rolando Somma, Gerardo Ortiz, Emanuel Knill, James Gubernatis, arXiv.org:quant-ph/0304063 (2003)

[27] “Quantum Simulations on a Quantum Computer” S. S. Somaroo and C. H. Tseng and T. F. Havel and R. Laflamme and D. G. Cory, Physical Review Letters, 82, 5381 (1999)

[28] “Quantum computation by measurement and quantum memory”, M. A. Nielsen, Phys. Let. A Vol. 308, 96-100 Issues 2-3, (2003) 11

[29] “Quantum information processing with cold atoms and trapped ions”, J Garca-Ripoll, Zoller and Cirac, J. Phys. B: At. Mol. Opt. Phys. 38 S567-S578 (2005) 12

[30] “Entanglement of Trapped-Ion Clock States”, P. C. Haljan, P. J. Lee, K. A. Brickman, M. Acton, L. Deslauriers, and C. Monroe, Phys. Rev. A 72, 062316 (2005) 13

[31] “Realization of the Cirac-Zoller controlled-NOT quantum gate”, F. Schmidt-Kaler, H. Haffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner and R. Blatt, Nature 422, 408-411 (2003) 13

[32] “Creation of a six-atom ’Schroedinger cat’ state”, D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, R. Reichle, and D. J. Wineland, Nature. 438, 639-642 (2005) 13

[33] “Deterministic quantum teleportation with atoms”, M. Riebe, H. Haffner, C. F. Roos, W. Hansel, J. Benhelm, G. P. T. Lancaster, T. W. Karber, C. Becher, F. Schmidt-Kaler, D. F. V. James, R. Blatt, Nature 429, 734 (2004) 13

[34] “Deterministic quantum teleportation of atomic qubits”, M.D. Barrett et. al., Nature 429 737-739 (2004) 13

[35] "Realization of quantum error correction", J. Chiaverini et. al., Nature 43 602-605 (2004) 13

[36] "Observation of entanglement between a single trapped atom and a single photon", B. B. Blinov, D. L. Moehring, L. M. Duan and C. Monroe, Nature 428, 153-157 (2004) 13

[37] "First Year Report: A $^{43}Ca^{+}$ Trapped Ion Qubit", D. Szwer, University of Oxford, 31 Aug 2006 18, 20, 21, 25, 116, 131

[38] "Demonstration of a fundamental quantum logic gate", Monroe, C. Meekhof, D. M. King, B. E. Itano, W. M. Wineland, Phys. Rev. Lett. 75, 4714-4717 (1995) 22

[39] "The Quantum Manipulation of Ions", D. Stevens, University of Oxford, (1999) 23, 28, 120

[40] "High-precision Qubit Readout and Microfabricated Ion Traps for QIP", A. Myerson, First Year Report, University of Oxford, 15 August 2007. 24, 44

[41] "Entanglement of Two Trapped-Ion Spin Qubits", J. Home, PhD Thesis, University of Oxford, (2006) 27, 28, 111

[42] "First Year Report", G. Imreh, First Year Report, University of Oxford, 31 August 2005 29

[43] "Coherence of a sympathetically cooled memory qubit", Home et al. Awaiting publication. 14

[44] "Laser Diagnostics", D. Allcock, Univeristy of Oxford, MPhys report (2007) 37

[45] "Isotope shifts and hyperfine structure in calcium", A Ramas. University of Oxford, 28 Nov 2006

[46] "Development of an Ion Trap Quantum Information Processor", C. Donald, PhD Thesis, University of Oxford, (2000) 44, 72

[47] "Electromagnetics", Kraus and Carver, McGraw-Hill, ISBN 0070353964 (1973) 40

[48] "A general-purpose pulse sequencer for quantum computing", P. Pham, PhD Thesis, Massachusetts Institute of Technology, February 2005 49

[49] “Inductance and Capacitance”, Richard E. DuBroff and James L. Drewniak EMC Laboratory, University of Missouri Rolla, IEEE Soc. Electromagnetic Compatibility Educational Paper 57

[50] “74F3037 Quad 2-input NAND 30Ωdriver”, NXP Semiconductors (Philips) http://www.nxp.com/acrobat/datasheets/74F3037_2.pdf 61

[51] “Quantum information processing with trapped ions”, Wineland, M. Barrett, J. Britton, J. Chiaverini, B. DeMarco, W. M. Itano, B. Jelenkovicy, C. Langer, D. Leibfried, V. Meyer, T. Rosenband, and T. Schatz. arXiv:quant-ph/0212079 v1 12 Dec 2002.

[52] “How to build a 300 bit, 1 Giga-operation quantum computer”, Andrew M. Steane arXiv:quant-ph/0412165v2 (2004). 67

[53] “Errors in trapped-ion quantum gates due to spontaneous photon scattering”, R. Ozeri et al. Physical Review A 75, 042329 (2007). 67

[54] “Hyperfine Coherence in the Presence of Spontaneous Photon Scattering”, R. Ozeri et al. Phys. Rev. Lett. 95, 030403 (2005) 68

[55] “Microwave signal generation with optical injection locking”, P. Bouyer et al. Opt Lett 18, 1502 (1996). 71

[56] “Diagnostics of DL100 Diode Laser System”, Patrick Parkinson, Internal Report, University of Oxford, 22 April 2004. 72

[57] “Master Slave Laser Injection Locking”, Ursula Pavlish, University of Oxford, 14 Aug 2004. 76

[58] “Electronic Frequency Stabilization of Microwave Oscillators”, RV Pound, Review of Scientific Instruments 17, 490 (1946) 84

[59] “Optics”, Hecht, Addison-Wesly ISBN:0201116111 (1987) 85

[60] “Electronic Communications Handbook”, Andrew F. Inglis, McGraw-Hill Education (Aug 1988) ISBN-13: 978-0070317116 85

[61] “An introduction to Pound-Drever-Hall laser frequency stabilization”, Am. J. Phys. 69 (2001) Eric D. Black DOI: 10.1119/1.1286663 87, 89

[62] “Laser Phase and Frequency Stabilization Using an Optical Resonator” Drever, Hall et al PACS: 06, 07.60, 07.65 87

[63] “Two-Photon Readout Methods for an Ion Trap Quantum Information Processor”, Matthew McDonnell, PhD Thesis, University of Oxford, (2003) 32, 88, 109

[64] “Stabilization and Control In a Linear Ion Trap”, John Patrick Stacey, Oxford Thesis 33, 34, 37, 90

[65] “NPL Etalon Specifications”, http://www.npl.co.uk/length/wss/services/equipment/etalons.html, retrieved : 2007-01-01 87

[66] “Photodiode Technical Guide”, http://sales.hamamatsu.com/assets/applications/SSD/photodiode_technical_information.pdf, retrieved: 2006-03-01 94

[67] “High Speed Light Detector with FET Input”, David Johnson and Associates, FE-TAMP0.DSN 95

[68] “Photodiode Front Ends The REAL Story”, Philip C.D.Hobbs, April 2001 Optics & Photonics News 42-45 95, 98

[69] “20MHz Light Receiver Front-End Circuit”, David Johnson and Associates, 20MRVR2.DSN 96

[70] “Application circuit examples of Si Photodiode”, Hamamatsu Application Notes 96, 100

[71] “Handbook of Basic Transistor Circuits and Measurements”, SEEC, Wiley (1967) 95

[72] “Transfer Report”, B. Keitch, University of Oxford, (2004) 105

[73] “The Art of Electronics, Horowitz and Hill”, Cambridge, ISBN 052137095

[74] “Encyclopedia of RF and Microwave Engineering”, Kai Chang (Editor), ISBN-13: 978-0471270539, John Wiley & Sons Inc (2005) 98

[75] “Precise determination of the ground state hyperfine structure splitting of 43Ca II” Arbes et al. Z. Phys. D 31, 27-30 (1994) 109

[76] “Magnetic field measurement around the ion trap”, Gergely Imreh, Internal Report, University of Oxford, (2005) 110, 111

[77] “Experiments with seperated oscillatory fields and hydrogen masers” Norman F Ramsey. Rev. Mod. Phys. Vol. 62 No. 3 (1990) 111

[78] “Atomic Physics” Christopher J Foot, OUP, ISBN 0198506961 (2004) 111, 112, 115

[79] “A Single Quantum Cannot be Cloned”, W.K. Wootters and W.H. Zurek, Nature 299 (1982), pp. 802-803. 130

[80] “Architecture for a Large-Scale Ion-Trap Quantum Computer”, D. Kielpinski, C.R. Monroe, and D.J. Wineland, Nature 417, 709-711 (2002) 130

[81] “A Single-Photon Server with Just One Atom”, Markus Hijlkema, Bernhard Weber, Holger P. Specht, Simon C. Webster, Axel Kuhn, and Gerhard Rempe, Nature Physics 3, 253-255 (2007) arXiv:quant-ph/0702034v1 130

[82] “Long-Lived Qubit Memory Using Atomic Ions”, C. Langer et al., Phys. Rev. Lett. 95, 060502 (2005) 137

[83] R. Feynman in Keynote Speech: “Simulating Physics with Computers” Int. J Theoretical Phys. 21 6-7 (1982) 147

Appendix A

Example Listing for LCU

The following sample Pascal code shows how to set-up and program the PCI-DAS1200 Card for the LCU.

```
program clktest;
{Test of clock on LCU and LCU break-out board
Uses PCI-DAS1200 on-board counters to generate pulse-sequence.
Version 1.0.0 BCK 22/05/2006
}

{The following two constants are the base-addresses of the PCI
card
pertaining to the counters.
On boot-up these are assigned by BIOS on a ''geographical'' bases
i.e. they don't change unless cards are added, removed or moved.
They are discovered using a probe program, PCI.pas }

const DAS1200badr3 = $D400;
const DAS1200badr1 = $CC00;
{PWMgap is set to 6us as this is the minimum time required
to repogram the counters.}
const pwmgap = $1E;
const tickover = $5000;

procedure DisableInterrupts; inline($FA); {CLI}
procedure EnableInterrupts;  inline($FB); {STI}

procedure InitDAS1200counters;
{Initialize DAS1200 counters C5, C6 for "PWM" mode, and start
ticking over.}
{C5, C6clk assumed are fed with same 10MHz clock signal (from
DAS1200 pin 77).}
{Note that all addresses requre two words to program hence 2
'loads' per addr}
begin
{Set up C5 as rate generator: sets (reprogrammable) time between
gap pulses}
port[DAS1200badr3+11] := $74;            {C5: mode 2=rate-gen,
lsb:msb, binary}
port[DAS1200badr3+ 9] := lo(tickover);   {load C5 with tickover
period}
port[DAS1200badr3+ 9] := hi(tickover);

{Set up C6 as one-shot: counts the (fixed length) gap pulse}
port[DAS1200badr3+11] := $B2;            {C6: mode 1=one-shot,
lsb:msb, binary}
```

```
port[DAS1200badr3+10] := pwmgap mod 256; {load with length of
gap pulse}
port[DAS1200badr3+10] := pwmgap div 256; {load with length of gap
pulse}

{Set up C4 to count C5out so we know when the next gap pulse is
starting}
portw[DAS1200badr1+4] := $0000; {select external clock for C4;
ADC pretrig off}
port[DAS1200badr3+11] := $10;              {C4: mode 0=potc, lsb
only, binary}
port[DAS1200badr3+ 8] := $01;              {load with £01 initial
count}
{could replace C4 with a simple flip/flop which would only
require a
single port[] instruction to read (via a digital input).}
{or could set up C4 as a "1-bit counter" and read C4out.}

{wait until first pulse from C5 has loaded new count into C4}
repeat
  port[DAS1200badr3+11] := $E2;            {read-back C4 status
byte}
until (port[DAS1200badr3+8] and $40)=0;  {wait for "null count"
bit to go low}
end;


procedure RunSequence(m: integer);
{This generates a simple sequence (counting up and down) to test
each
channel of the LCU. In reality a sequence is loaded from a
file.}

var
  i,extrapulses: integer;
  tout,numpulses: longint;
  c4,c4inc,prevc4: byte;
  pulselen,L,pmtcounts,D: word;
  pmtpulse,dacpulse,timeout: boolean;

  function Read1200counter4: byte;
    begin
    port[DAS1200badr3+11] := $00; {latch C4}
    Read1200counter4 := port[DAS1200badr3+8]; {read C4 (lsb-only
mode)}
    end;
```

```
  procedure WaitForGapStart;
   const toutmax = 100000; {gives 0.1sec timeout}
    begin
    {wait for start of next gap :
(C4 counts trigger just before start of gap pulse)}
    tout := 0; {timeout counter to avoid crashing}
    repeat
      port[DAS1200badr3+11] := $00; {latch C4}
      c4 := port[DAS1200badr3+8]; {read C4 (lsb-only mode)}
      inc(tout);
    until (c4<>prevc4) or (tout>toutmax);
    if (tout>toutmax) then timeout:=true;
    c4inc := prevc4-c4; {auto-wraparound for byte type}
    prevc4 := c4;
    extrapulses := extrapulses + c4inc - 1; {one pulse should have
occurred}
    end;
begin

DisableInterrupts; {interrupts off throughout sequence for reliable
timing}

{re-initializing counters every sequence seems to make timing
glitches rarer}
InitDAS1200counters;

extrapulses := 0; {counts any extra pulses detected by C4}
timeout := false; {set to true if timeout while waiting for next
gap}
prevc4 := Read1200counter4; {read initial C4 count}
WaitForGapStart;
numpulses := 10000;
i := 0;
L := $5000; {2 ms}
D := $1;
pulselen := $FF00;
while (i<=numpulses) do
  begin
  pulselen := L+pwmgap;
  {reload C5 for next pulse}
  port[DAS1200badr3+9] := pulselen and $FF;
  port[DAS1200badr3+9] := pulselen shr 8;
  {wait for start of next gap :
(C4 counts trigger just before start of gap pulse) }
  WaitForGapStart;
```

```
  {set external (bypassed) and multiplexer outputs for this
pulse}
  port[DAS1200badr3+4] := D and $FF;            {set port A
outputs}
  port[DAS1200badr3+5] := (D shr 8) and $FF;  {set port B
outputs}
  port[DAS1200badr3+6] := (D shr 16) and $FF; {set port C
outputs}
  inc(i);
  end;

EnableInterrupts;
if (extrapulses<>0) then
  begin
  write('Sequence error: counted extra pulses');
  end;
if (timeout) then
  begin
  write('Sequence error: timeout');
  end;
end;


begin
RunSequence (1);
END.
```

Appendix B

Effect of changing sideband height and detuning in PDH signal

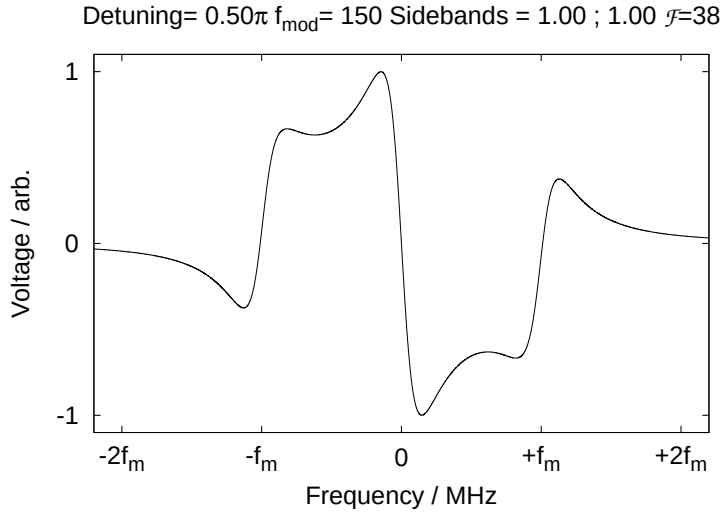

Figure B.1:

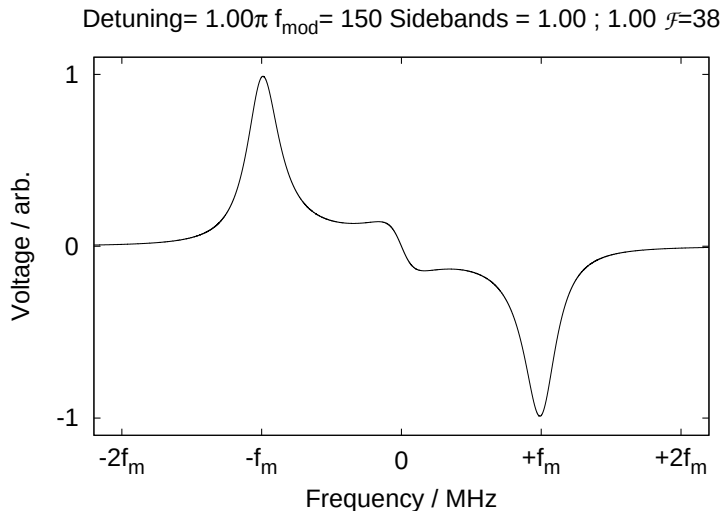

Figure B.2:

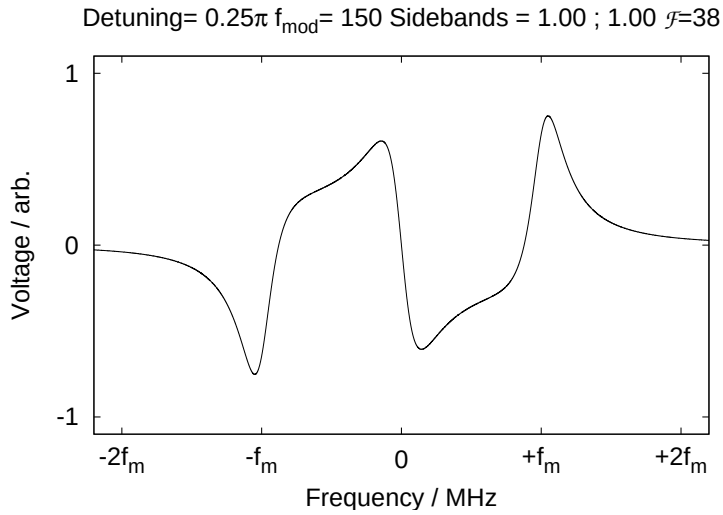

Figure B.3:

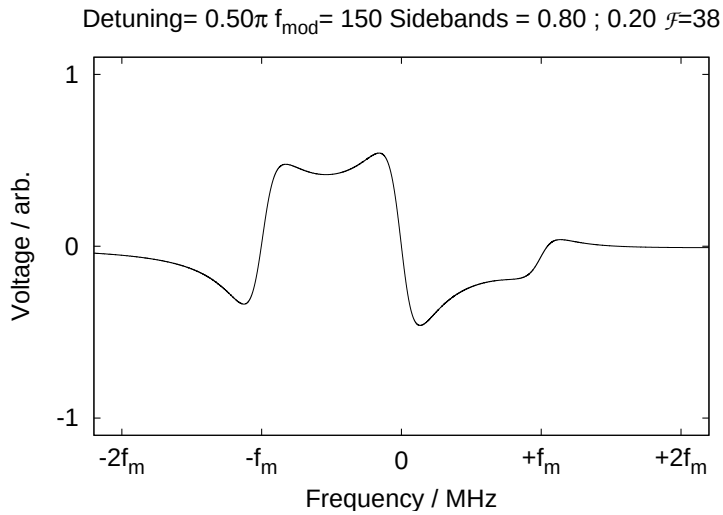

Figure B.4:

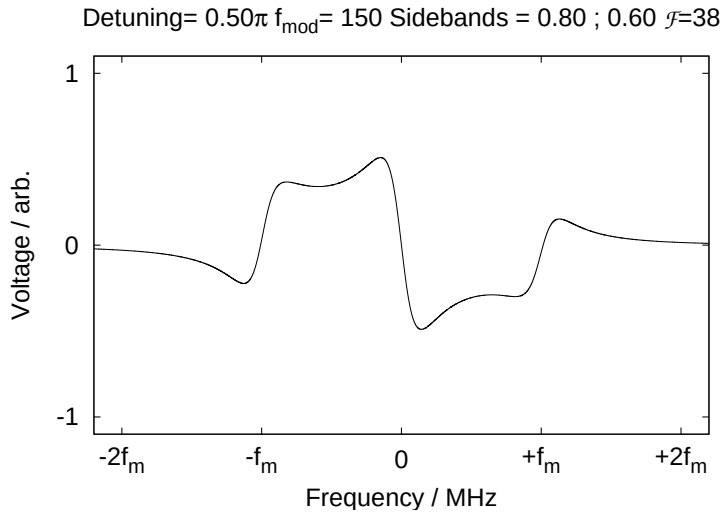

Figure B.5:

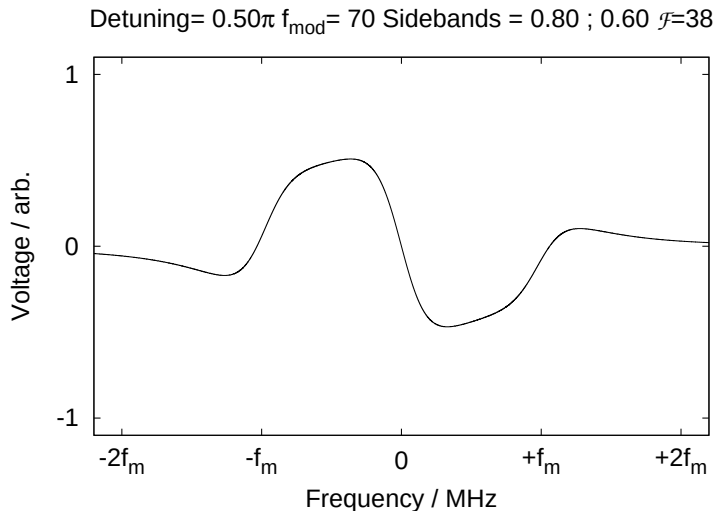

Figure B.6:

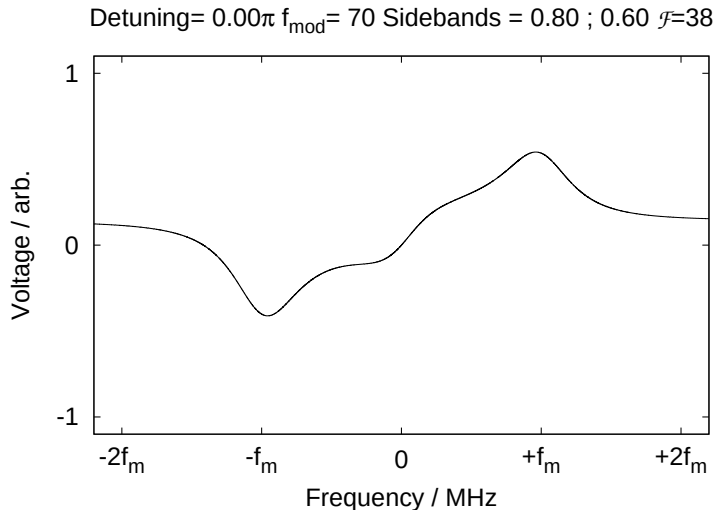

Figure B.7: