

Measurement-Selected Ensembles in Trapped-Ion Qubits

Michael J. Curtis

A thesis submitted for the degree of Doctor of Philosophy

University of Oxford, Department of Physics

Hilary Term, 2010

Contents

Contents	i
1 Introduction	5
1.1 Segmented ion traps	6
1.1.1 Degrees of freedom	6
1.1.2 Trapping fields	7
1.2 Quantum measurements	8
1.3 Structure of this thesis	10
2 Segmented Ion Trap Modelling	13
2.1 Radio-frequency trapping	13
2.2 The RF pseudo-potential	14
2.3 Ideal potentials	15
2.4 Electrode configuration space	16
2.4.1 Numerical modelling	17
2.4.2 Parametrisation	17
2.4.3 Symmetry	19
3 Measurement-Selected Ensembles	21
3.1 Two-State Vector Formalism	22
3.2 Weak measurement	24
3.2.1 Review of continuous pointer measurements	25
3.2.2 Qubit pointer measurements	27
3.3 Bloch vector representation of the density operator	30
3.4 Partial collapse and uncollapse of the wavefunction	33
3.4.1 Partial collapse	34
3.4.2 Uncollapse	35
3.4.3 Asymmetric case	38
3.4.4 Practical value	42
3.4.5 Leaky transfer	44
3.4.6 Implementation with a trapped ion qubit	44
4 Experimental Methods	47
4.1 Calcium atomic structure	47
4.2 Stabilised diode lasers	49
4.2.1 Photoionisation lasers	51

4.2.2	Doppler laser	51
4.2.3	Shelving laser	51
4.2.4	Deshelving lasers	52
4.2.5	Light-shift laser	52
4.3	Laser Control Unit	52
4.4	Photon counting	53
5	Lucent Trap	55
5.1	Apparatus	55
5.1.1	Design and fabrication	55
5.1.2	Electrical properties and faults	56
5.1.3	Vacuum system	57
5.1.4	Optical setup	59
5.1.5	Generation of DC trapping potentials	60
5.1.6	Generation of RF trapping potential	63
5.1.7	Magnetic field	63
5.2	Neutral calcium fluorescence	63
5.3	Loading attempts	64
5.3.1	The Dust	64
5.3.2	Background scatter	66
5.3.3	Procedure	66
5.4	Comments	68
6	Liverpool Trap	69
6.1	Apparatus	69
6.1.1	Design and fabrication	69
6.1.2	Vacuum system	71
6.1.3	Generation of RF trapping potential	72
6.1.4	Generation of DC trapping potentials	73
6.1.5	Magnetic field	74
6.2	Loading	74
6.2.1	Beam alignment	74
6.2.2	Neutral calcium fluorescence	75
6.2.3	Ion trapping	76
6.2.4	Micromotion compensation	76
6.3	Characterisation	78
7	$D_{5/2}$ Spin Experiments	83
7.1	Theory	83
7.1.1	State preparation	84
7.1.2	Coherent manipulation	87
7.1.3	Readout	97

7.2	Polarization	99
7.2.1	Optical setup	99
7.2.2	Alignment	101
7.2.3	Results	103
7.3	Magnetic driving	107
7.3.1	Magnetic coil and control electronics	107
7.3.2	Results	109
7.4	The $D_{5/2}$ qubit	111
7.4.1	Application of light shifts	111
7.4.2	Qubit manipulations	114
7.4.3	Decoherence and protection sequences	116
8	Uncollapse	119
8.1	Pulse sequence	119
8.2	Process tomography	121
8.3	Results	122
9	Conclusions	133
9.1	Comments on ion trapping	133
9.2	Comments on measurement-selected ensembles	134
9.2.1	Higher fidelity	135
9.2.2	Symmetric uncollapse	136
9.2.3	Weak measurements	136
	Bibliography	139
A	Details of $^{40}\text{Ca}^+$ Calculations	145
A.1	Transforming the magnetic dipole Hamiltonian	145
A.2	<i>Mathematica</i> code for rotation matrices	146
A.3	Time-independent light shift Hamiltonian	147
A.4	Eigenvalues of light shift Hamiltonian	149
A.5	Light shift Rabi frequency	149
B	Partial and Uncollapse Experiment Detailed Data Tables	151
B.1	Pulse Timings	152
B.2	Partial and uncollapse experiment ensemble sizes	153
	List of Symbols and Abbreviations	155
	List of Figures	157
	List of Tables	159

Abstract

This thesis describes the development of two segmented ion trap systems for quantum information experiments and investigations of novel quantum measurements. We report successful loading of $^{40}\text{Ca}^+$ ions in one of the trap systems and the successful demonstration of the partial collapse and revival of the wavefunction of an ion qubit using the system.

We describe the operation of ion trap systems with segmented electrodes and develop a practical strategy for choosing the voltages to apply to the electrode segments to obtain desired trapping potentials and correct for stray fields. The linearity of the Laplace equation results in a linear map from the configuration space of electrode voltages to the function space of electric potentials in the trap. We decompose the electrode configuration space into basis vectors whose corresponding potentials have intuitive meaning.

We describe attempts at loading Ca^+ ions in a micro-scale segmented ion trap built at Lucent Technologies. The Lucent Trap is a planar ion trap fabricated using lithographic techniques on a silicon substrate. We experienced several problems with this trap and were unable to successfully load ions in it.

A second ion trap, built at the University of Liverpool, is described. The Liverpool trap has miniature segmented electrodes constructed using conventional machining. We describe loading of Ca^+ ions, and the demonstration of stable trapping with long ion lifetimes and stable electric and magnetic fields.

The Liverpool trap is used to demonstrate non-projective “partial” measurements of a qubit wavefunction, which result in non-unitary, non-projective evolution of the wavefunction dubbed *partial collapse*. This effect can be reversed by performing another partial measurement, but only if the second measurement results in a particular outcome. When this outcome is observed, the qubit wavefunction is restored to the original state. An exact restoration occurs when both measurements are ideal, and we quantify the success of our demonstration by quantum process tomography. We obtain process fidelity > 0.65 for a wide variation in strength of the partial collapse from 0–0.94 (where 0 leaves the system unaffected and 1 is the limit of a complete projective measurement).

Acknowledgements

I thank Andrew Steane, my supervisor, for his support, guidance, and for my first introduction to both ion trapping and weak measurements. I thank Derek Stacey for his encouragement, his enthusiasm for interesting problems, and the many discussions which shaped my understanding of physics. I thank David Lucas for sharing his expertise and solving problems in the lab, for the early discussions which led to the design of the uncollapse experiment, and his advice and guidance throughout my time in Oxford.

All the Oxford Ion Trap Group members, past and present, have had some hand in various components of the experimental setup, and I also thank all those whose tenure was concurrent with mine for useful discussions. I wish to mention a few in particular. Nick Thomas and Gergely Imreh built some of the laser systems and vacuum systems used in the experiments in this thesis, and generally showed me the ropes when I first arrived. Alice Burrell built the imaging systems used in the experiments. Ben Keitch and David Allcock built the laser control units, and David Lucas programmed the experiment control computer. Jeff Sherman built many various pieces of electronics kit, and was my most significant partner in the laboratory for the experiments presented here. The implementation of coherence-preserving pulse sequences follows directly from the work of David Szwer.

This work is also based on and influenced by the intellectual contributions of a great number of people. It is my hope that the majority of those most relevant have been included in the bibliography. However, I wish to mention two people who are not referenced, but who nonetheless have influenced this work in a significant way. L. Nick Trefethen's numerical computing course at Oxford has led to a deeper understanding of linear transformations and numerical approximations. The former in particular is a pervasive mathematical foundation in this work. Finally, the writings and examples of Edward Tufte have significantly informed the design of this thesis.

1

Introduction

Trapped ions constitute a system in which coherent control over both motional and internal degrees of freedom can be attained. This makes them an attractive candidate for quantum information processing, using relatively robust internal degrees of freedom as qubits and motional degrees to achieve entanglement between qubits. The first schemes for coupling internal and motional states were conceived by Cirac & Zoller [12], and recent experiments [39] [21] have demonstrated simultaneous coherent control over several ions. In addition, several schemes for coupling internal states to macroscopic observables, such as scattered photons from an illuminating laser [40], have been demonstrated.

The science of constructing controllable, interacting, quantum systems which are measurable at an extremely detailed level will be of major importance in a large number of disciplines—not just those interested in doing computations. For example, the ion-ion entanglement and high fidelity readout techniques of trapped ion qubits [47] have been applied to create an extremely stable (i.e. a fractional inaccuracy of 8.6×10^{-18}) atomic clock [11].

One area in which ion trapping techniques can be fruitfully exploited is the study of ensembles of quantum systems when they are subject to non-projective measurements. In particular, here we are concerned with two classes of measurement. The first, called *weak measurement*, attempts to minimise the disturbance of the system being measured, and leads to interesting correlations with the results of projective measurements made before and after the weak measurement. The second is partial measurement, in which the process gives incomplete information about the state of the system. We formally consider the state of the system following a measurement of this type, and this leads to the concept of the *partial collapse* of the wavefunction, in which the system state is altered by the measurement. The reversal of this partial collapse is called the *uncollapse* of the wavefunction.

In this thesis we develop two related, but distinct avenues of research in quantum information. First we describe progress in the development of ion trap systems with segmented electrodes to allow simultaneous control over a larger

number of ions. This includes the discussion of a general strategy for generating the desired electric potentials in these traps, and our attempts at loading two different segmented ion trap systems.

The second is a presentation of experimental results showing the partial collapse and uncollapse of a qubit wavefunction, using one of the ion trap systems developed. We also discuss some general aspects of the theory of measurement-selected ensembles and give examples of weak measurements and partial measurements in qubits.

1.1 Segmented ion traps

As ions are added successively to a single trap, it becomes more difficult to coherently control their collective motional states. A scalable system would have a large number of ions, divided into an array of traps, and some method for moving information between them. One method is to move the ions themselves—shuttling them from trap to trap [30]. Many novel trap designs have been proposed, termed *segmented* ion traps, due to the electrodes used to create DC fields being divided into segments. Manipulation of the voltages applied to these segments creates multiple trapping regions and time-varying potentials which move and separate ions in the trap array.

A general trend has been toward smaller traps and more densely segmented electrode configurations. The fabrication techniques required to construct reliable trap arrays at the 10-100 μm scale are not well established. Our research group has obtained several segmented traps with which to begin enter this territory. We are concerned here with two: the Liverpool trap is a sub-millimetre scale trap designed as an intermediate size trap in which controlling a multi-segment trap can be tested without the fast heating [13] and short ion lifetime problems associated with microscopic traps. The Lucent trap is micro-scale, designed to test microfabrication techniques.

1.1.1 Degrees of freedom

An ion is a strongly interacting system of many particles, and correspondingly, has a large number of degrees of freedom: 3 continuous degrees and a finite set of spin states for each proton, neutron*, and electron. However, the particles are all strongly bound together, so we may consider the quantised internal structure separately from 3 external, centre-of-mass degrees of freedom. Fortunately, the configuration of the protons and neutrons in the nucleus is stable, and on account of the large energy scales involved in these configurations, we shall never appreciably couple to the degrees of freedom therein. We simplify the electronic

*of course, each proton and neutron also has its own internal degrees of freedom!

structure by only considering the ground state and a few low-lying energy levels. We might then classify interactions with ions as those operating on external degrees of freedom, those operating on the internal degrees of freedom, and those which couple the two.

It is easy to imagine an example of the first category: static electric fields affect the motion of ions, but couple much more weakly to internal degrees of freedom. An example of an interaction which couples both is an oscillating electric field near an optical resonance—which for an unbound ion also changes the state of motion because of the momentum associated with absorbing or emitting a photon.

Are there examples of interactions which only couple significantly to internal degrees? For small energy gaps such as Zeeman or hyperfine splittings, microwave or Raman excitation will have negligible influence on the motion. On the other hand, it would seem that any single-photon interaction which has enough energy to couple to optical transitions will always change the motion, and for unbound ions, this will be the case. The situation changes for bound ions, however, since their energy spectrum will be quantised. The coupling between internal and external degrees of freedom can be frozen out if the spacing of the ions' motional energy levels is made large compared with the kinetic energy change associated with absorbing or emitting a photon [58]. A large splitting of the ions' motional energy levels corresponds to tight confinement.

Detuning from the internal resonance reintroduces coupling between motional and electronic degrees of freedom, allowing for Doppler and sideband cooling of ions, and controlled coupling between ions in the same trap via shared states of motion.

Thus, precise control over the electromagnetic field—its frequency, amplitude, polarisation, and direction—will translate to precise quantum control over both the internal and external degrees of freedom of an ion-qubit. Static and radio-frequency (RF) electric fields are used to tightly confine ions in traps, and optical electric fields are used both to manipulate the internal degrees of freedom and to couple internal and external degrees for cooling processes and multi-qubit gates. A static magnetic field is used to define a quantisation axis and prevent optical pumping. An RF magnetic field is used to manipulate the spin state of the ions.

1.1.2 Trapping fields

It is impossible to trap a charged particle in 3-dimensional free space using static electric fields. For a charged particle to feel no force at a point in space we must have $\nabla V = 0$ at the point, and in order for this to be a stable equilibrium we require (for a positive charge) $\nabla^2 V > 0$. Of course, the second of these requirements is forbidden in free space by Laplace's equation.

The innovation pioneered by Paul [43] was the use of an oscillating electric field to trap charges. When a charged particle is in an oscillating electric field, it oscillates at the same frequency along the polarisation of the field, and will feel a time-averaged force along any gradients in the strength of the field. This is called the *ponderomotive force*. If there is a gradient in field strength, then the ion gains more momentum during one half of its oscillation than the other. Averaged over oscillation cycles, this change in momentum may be regarded as the action of a force-field proportional to $\nabla|\mathbf{E}|^2$, or equivalently as a potential proportional to $(\nabla V)^2$. This is the key: while it is impossible to have a minimum of V in free space, it is possible to have a minimum of $(\nabla V)^2$.

Consider, for example, a quadrupole:

$$V = \alpha x^2 + \beta y^2 + \gamma z^2 \tag{1.1}$$

and,

$$(\nabla V)^2 = 4(\alpha^2 x^2 + \beta^2 y^2 + \gamma^2 z^2). \tag{1.2}$$

Laplace's equation prevents the coefficients α , β , γ from all being positive, as required to have a minimum of V , but regardless of the signs, $(\nabla V)^2$ has a minimum.

While RF ion traps vary widely in their electrode configurations, they all have RF electrodes designed to create a minimum of $(\nabla V)^2$ in the trapping region.

1.2 Quantum measurements

In quantum mechanics, information and measurement are closely linked. Knowing the state of a physical system represents the most complete information we may have about it. The quantum state tells us, among other things, the possible outcomes of measurements on the system. The inverse is also the case: knowing the outcome of a measurement made on the system yields information about its quantum state prior to (and following) the measurement. However, quantum mechanics' predictive capability is, in general, limited to probability: the quantum state gives the statistics of measurements on an ensemble of systems[†], and (appropriately chosen) measurements on an ensemble of systems whose state is initially unknown allow the determination of that state up to statistical uncertainty.

The consequence of the Measurement Postulate of quantum mechanics is that it is not always possible to measure a system and leave its state unchanged. When the system is in a state which a measurement will leave unchanged, the

[†]an ensemble of one presents no difficulty

state is said to be *sharp* with respect to the measurement. There are no states which are sharp with respect to all possible measurements (a consequence of the Uncertainty Principle), and we may in general regard a measurement as a process which leaves the system altered. We gain information about a system with the consequence that in so doing we change it.

Qubits in a quantum information processor must be well-isolated from the environment. It is insufficient for us to merely ensure that *we* have not measured the system—it must be the case that no measurement on the environment, no matter how precise, can yield information about the state of the system. Each qubit must be initialised to a known input state, then allowed to interact according to a set of rules which forms the program, and finally the qubits are measured to form the output. During the progress of the computation, we must take care that no information about the states of the qubits gets out, for this would spoil the computation. Measuring the system before the program has finished would be disastrous!

Of course, this view is too simplistic. The only realistic paradigm for a large scale quantum information processor involves protecting the coherence of the “logical” qubits by encoding them in several “physical” qubits with an error-correcting code [57][52], and designing the program to be fault-tolerant [53][44]. In this paradigm, measurements are constantly being performed throughout the computation to gain information about errors (which allows for their correction). These measurements, called the *syndrome extraction*, are designed such that the results reveal no information about the state of the encoded “logical” qubit. No logical information is extracted, and no harm is done to the computation in progress.

As it turns out, there are other measurement processes which will leave the system unaltered, or approximately so, even if it is not sharp with respect to the observable being measured. These novel processes, which all involve non-projective measurements, are a major subject of this thesis. We will examine one type of measurement in which we reduce the amount of information gained to nearly zero. In this “weak measurement,” the effect on the system being measured is reduced so that the statistics of the outcomes of subsequent measurements are changed only negligibly. We will find, however, that although the correlation between the initial state and final state is unaffected, very interesting correlations arise in the results of the weak measurement with respect to the initial and final states.

In a second scenario, we will show that a measurement which is not projective does not completely destroy the coherence of a quantum system. Partial measurements extract incomplete information, and the subsequent state of the system cannot necessarily be determined from the partial measurement alone. We will also show that if a second partial measurement is performed and a par-

ticular outcome obtained (which, in general, will not occur with certainty), then the system can be restored back to its original state—the uncollapse of the wavefunction.

In the case of uncollapse, when the information gained from the two partial measurements are examined together, they yield zero net information about the system being measured. This indicates that we cannot regard the effect of measurement as fundamentally irreversible with respect to the system being measured. However, when the experiment is repeated many times, the partial measurement results allow the state of the system to be inferred.

1.3 Structure of this thesis

Chapters 2 and 3 are primarily theoretical and introduce the topics discussed. They give the background necessary for the experimental Chapters 4–8 that follow. The two themes of segmented ion traps (Chapters 2, 5, 6) and measurement-selected ensembles (Chapters 3, 7, 8) are interspersed, rather than developed sequentially, because experiments on loading ions, cooling them, and compensating for stray fields are the foundation for the demonstrations of the partial collapse and uncollapse of the qubit wavefunction. Unfortunately this means the transition between the two theory sections, from segmented traps to measurement-selected ensembles, will not be as natural.

Chapter 2 gives a theoretical treatment of the task of choosing the voltages to apply to the electrodes in a segmented ion trap, discussing along the way the essential physics of operation of traps. This serves as a more complete introduction to ion traps in general, and details the techniques used in the trap loading sections.

We then turn to the treatment of measurement-selected ensembles. In Chapter 3 we introduce and define the measurement-selected ensemble, then discuss two interesting examples which the techniques of quantum information processing have put within reach of experimental tests. We review the relevant literature on each, and make explicit the experimental demonstrations which could be accomplished using qubits. This gives the necessary theory to discuss the results of our demonstration of the second of the examples: the partial collapse and restoration of a qubit wavefunction. These results are discussed in Chapter 8.

However, before examining those results, we present the apparatus used in the experiments. This includes laser systems, photon counting, experimental control, and of course the trap system itself. The first three are discussed in Chapter 4.

This leads into a discussion of two different segmented trap systems which we attempted to load with Ca^+ . We present results with two different trap sys-

tems: the Lucent Trap and the Liverpool Trap. The Lucent Trap is a micro-scale planar surface trap fabricated on a silicon substrate. It was built by Lucent Technologies (formerly Bell Labs). We failed to load ions in the Lucent Trap. In Chapter 5 we discuss our attempt and comment on the difficulties we encountered.

The Liverpool Trap is a mesoscopic scale ion trap constructed from miniature stainless steel electrodes and MACOR[‡] insulator. It was built at the University of Liverpool and is discussed in Chapter 6. We successfully loaded $^{40}\text{Ca}^+$ ions into this trap, and it is the trap system used for the subsequent experiments presented in this thesis.

The demonstration of partial collapse and recovery requires the experimental ability to implement a partial measurement without completely destroying the qubit coherence. This is experimentally non-trivial, and required us to develop a “new qubit” in $^{40}\text{Ca}^+$. The qubit levels are two Zeeman states of the metastable $3D_{5/2}$ level of $^{40}\text{Ca}^+$. Chapter 7 gives the details of the operation (initialisation, coherent dynamics, and readout) of this qubit, and gives the results of the diagnostic experiments performed on it.

Chapter 8 then presents and comments on the results of the partial collapse and uncollapse experiment on the qubit. We demonstrate a fidelity of uncollapsing greater than 0.65 for partial collapse strengths in the range 0.01–0.94.

We conclude the thesis with some comments and ideas for future work.

[‡]MACOR is a machinable glass ceramic manufactured by Corning Inc.

2

Segmented Ion Trap Modelling

Segmented ion traps, by their very nature, have a large number of degrees of freedom in their configuration—of order one for every electrode segment. In this chapter we present a general strategy to choose these parameters, and reduce the complexity of the configuration to a manageable level. It is to our great advantage that electric potentials obey the principle of superposition; this allows the mathematics of abstract vector spaces to be employed to manipulate the configuration space.

We first review the basics of trapping ions with static and oscillating fields, leading to a functional form of the trapping potentials. This will lead naturally into a discussion of how to generate such potentials by manipulating the voltages applied to trapping electrodes. The basic trapping potentials shall serve as an example throughout, but we shall keep the treatment general enough that it may be applied to more complex situations, such as the separation of ions into two trapping regions, which arise in the evolution of the ion trap as a tool in quantum information science.

2.1 Radio-frequency trapping

It is impossible to confine charged particles in free space using only static electric fields, but by adding an oscillating component it will be possible to confine in all three dimensions. An oscillating electric field causes a charged particle to undergo a driven motion at the oscillating frequency, and it feels a time-averaged force along gradients in the strength of the oscillating field. This ponderomotive force can be used to confine charged particles.

For ions, the oscillating component will be in the radio-frequency domain (5-50 MHz is typical for traps of this type). The wavelengths at this frequency are much larger than the size scale of the trap, and thus it is appropriate to consider only a quasi-static electric potential in the treatment of ion dynamics: $V(\mathbf{r}, t)$, and the Maxwell equations reduce to Poisson's equation

$$\nabla^2 V(\mathbf{r}, t) = \frac{\rho(\mathbf{r}, t)}{\epsilon} \quad (2.1)$$

with ρ the charge density, and ϵ the permittivity.

We may further decompose the electric potential into a static component and a single-frequency oscillating component

$$V(\mathbf{r}, t) = V_{\text{dc}}(\mathbf{r}) + V_{\text{rf}}(\mathbf{r}) \cos(\Omega_{\text{rf}} t) \quad (2.2)$$

which allows our analysis to focus now on two time-independent functions $V_{\text{dc}}(\mathbf{r})$ and $V_{\text{rf}}(\mathbf{r})$. Our general procedure is to consider what these functions should be in the ideal case, and then determine what voltages to apply to the electrodes to realise them as nearly as possible. (2.1) has no differential dependence on time, and we may simply drop the time coordinate when using it to find time-independent potentials.

2.2 The RF pseudo-potential

The form of $V_{\text{rf}}(\mathbf{r})$ is typically determined by the trap geometry, with the experimenter having control over a single parameter: the amplitude

$$V_{\text{rf}}(\mathbf{r}) = Q\phi(\mathbf{r}) \quad (2.3)$$

with $\phi(\mathbf{r})$ a function determined by the geometry. In the regime where Q is not too large, a simple, approximate description of the ion's motion is appropriate. It involves deriving a *pseudo-potential* due to the oscillating field, which describes the secular response of the ion, neglecting the small micromotion at the RF driving frequency.

The RF pseudo-potential is given by

$$V_{\text{pseu}}(\mathbf{r}) = \frac{eQ^2(\nabla\phi) \cdot (\nabla\phi)}{4m\Omega_{\text{rf}}^2} \quad (2.4)$$

with e and m the charge and mass of the ion.

The pseudo-potential is advantageous because it can be calculated directly from the RF potential, and is particularly useful in calculating the depth of a real trapping potential. Neglecting micromotion, the ion behaves as if it is in an effective potential

$$V_{\text{eff}} = V_{\text{pseu}} + V_{\text{dc}} \quad (2.5)$$

However, the condition that the effective potential has a local minimum is not sufficient to guarantee that stable trapping is possible. The ion's dynamics

at the effective potential minimum are governed by a set of differential equations called the Mathieu equations, and stable solutions are possible for a certain range of values for the confinement due to RF and DC components. We shall return to this issue after we examine the functional form of the potentials.

2.3 Ideal potentials

Our ideal effective potential is an anisotropic harmonic oscillator

$$V = \frac{m}{2e} (\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2) \quad (2.6)$$

To create an *axial* trap, we will arrange for the trapping strength along one of the axes of the trap to be significantly lower than the other two, $\omega_z < \omega_x, \omega_y$. In this regime, when ions are crystallized in the trap they line up along the axis with weak confinement. For this reason, we will refer to the z -axis as the *trapping axis*, or simply *the axis* of the arrangement.

A 2-dimensional quadrupole potential oscillating at RF is used to tightly confine the ions in the xy (radial) plane. The form of this potential is

$$V_{\text{rf}} = Q_{\text{rf}}(x^2 - y^2) \quad (2.7)$$

Using (2.4) we obtain the RF pseudo-potential

$$V_{\text{pseu}} = \frac{eQ_{\text{rf}}^2}{4m\Omega_{\text{rf}}^2} (x^2 + y^2) \quad (2.8)$$

Confinement along the trapping axis is created with a static potential. We assume that in the region of interest the form of this potential is quadratic:

$$V_{\text{dc}}(x, y, z, t) = \lambda x^2 + \sigma y^2 + \gamma z^2 \quad (2.9)$$

Since there are no charges in this region, we must satisfy $\nabla^2 V = 0$,

$$\lambda + \sigma + \gamma = 0 \quad (2.10)$$

This means that in order for the static potential to be confining in z (i.e. $\gamma > 0$), it will need to be expelling in either x or y or both. Since we will depend on the RF potential to confine in the xy plane, we choose $\lambda = \sigma = -\gamma/2$ as the best-case scenario. The effective potential is

$$V_{\text{eff}} = \left(\frac{eQ_{\text{rf}}^2}{m\Omega_{\text{rf}}^2} - \frac{Q_{\text{dc}}}{2} \right) (x^2 + y^2) + Q_{\text{dc}} z^2 \quad (2.11)$$

where we have replaced γ with Q_{dc} for consistency in notation. We may read off the trapping frequencies from this expression.

As we mentioned briefly in our discussion of the RF pseudo-potential, stable trapping will only be possible with some values of Q_{rf} and Q_{dc} , predicted by solutions to the Mathieu equation. See [19] or [14] for details. Practically speaking, we may extract the Mathieu parameters from the form of our potentials, and compare with a stability diagram, such as [20]. There is somewhat wide variation in notation for these potentials. The canonical Mathieu parameters for the potentials we have just described are:

$$q = \frac{4eQ_{\text{rf}}}{m\Omega_{\text{rf}}^2} \quad a = -\frac{4eQ_{\text{dc}}}{m\Omega_{\text{rf}}^2} \quad (2.12)$$

We typically operate traps with these parameters in the range 0.1 to 0.5.

2.4 Electrode configuration space

Having established what potentials are needed to create a trap, we now turn to the more difficult question of what voltages to apply to the electrodes to generate the appropriate potentials. Here we take advantage of the linearity of Laplace's equation. The potential due to a set of voltages $\{\psi_n\}$ applied to electrodes 1 to N can be written

$$V(\mathbf{r}) = \sum_{n=1}^N \psi_n k_n(\mathbf{r}) \quad (2.13)$$

where $k_n(\mathbf{r})$ are the potentials which result when unit voltage is applied to electrode n and all others are at ground. (2.13) can be written as a linear algebra equation with a pseudo-matrix K whose columns are the k_n

$$\begin{bmatrix} k_1, \dots, k_n, \dots, k_N \end{bmatrix} \begin{bmatrix} \psi_1 \\ \vdots \\ \psi_n \\ \vdots \\ \psi_N \end{bmatrix} = V \quad (2.14)$$

$$K\boldsymbol{\psi} = V \quad (2.15)$$

We have two problems: firstly, we need to find the functions k_n , and secondly, a continuous function has an infinite number of degrees of freedom, and our electrode voltage vector $\boldsymbol{\psi}$ is finite. The electrode configuration $\boldsymbol{\psi}$ is a real vector in the \mathbb{R}^N configuration space for the trap. K is a linear map from the N -dimensional configuration space to the infinite-dimensional space of scalar fields. We cannot invert a matrix which has infinite rows but finite columns.

2.4.1 Numerical modelling

To find the functions k_n , we need to solve Poisson's equation, (2.1), for the N cases of unit voltage on each electrode in turn with all others held at ground. For most segmented trap geometries, we will be unable to find closed-form solutions to these equations*. Therefore, we turn to numerical modelling to obtain approximations to the k_n .

The most common method in modelling ion traps is the boundary element method (BEM). BEM divides the boundaries, i.e. the surfaces of the electrodes, into elements, and solves for the electric charge on the elements. Given j surface elements, the potentials are related by a set of j linear equations to the charges on each of the elements via Coulomb's Law. Cast in this light, solving for a set of charges from an input set of potentials involves inverting a $j \times j$ matrix. Once the charges have been calculated, the electric potential anywhere in space can be calculated by summing the contributions from each surface element via Coulomb's Law. BEM solves the special case of Poisson's equation when there are no dielectrics near the trapping region, i.e. Laplace's equation $\nabla^2 V = 0$. This is not a major limitation for ion traps because dielectrics cause numerous problems like creation of stray fields and distortion of RF potentials and so are kept well away from the trapping region.

Calculations used in this work were made with a BEM package called *Charged Particle Optics* (CPO).

We must choose the region in which to evaluate the k_n with some care. Real trapping potentials have a finite depth, and therefore must eventually turn over and decrease as we move away from the trap centre—so the target harmonic potential will only hold in a region smaller than the electrode-electrode spacing in the trap. If we make the region too small, however, numerical or rounding errors in the simulation may make it impossible to extract accurate values for the derivatives of V which are our chief concern.

2.4.2 Parametrisation

Once discrete representations of k_n have been calculated, we may calculate an equivalent representation of our target V (i.e. evaluating it at the same discrete points). In this way, we convert the pseudo-matrix equation (2.14) into a real matrix equation. Typically we will want to evaluate the k_n at many more points than there are electrode degrees of freedom, and thus we will be unable to solve the equation exactly. We may calculate a least-squares solution in the usual way to get a best approximation of the target potential.

*The major exception are surface traps whose electrodes are confined to a single plane, provided the gaps between conductors are sufficiently narrow [66]

However, calculating the least-squares solution does not help us understand the structure of the configuration space. We will need a deeper analysis in order to understand how to compensate for stray fields and how to trap in more than one region at a time. A better approach is to express the V and k_n functions as linear combinations of simpler functions which capture some particular aspect of the trap. For example, the function $q(\mathbf{r}) = z^2$ represents the confinement in the z -direction. The coefficient on q in an expansion of k_n is the n^{th} electrode's contribution to confinement in the z , and the coefficient in an expansion of V is the target confinement. This example suggests that a Taylor expansion is a good basis in which to express the V and k_n functions.

The choice of what terms to include will depend on the specific trap being modelled. We may truncate the expansion when the coefficients become negligible, and this will depend on the size of the regions for which we calculate the k_n , and on the length scales of the surrounding electrodes. For traps where the electrode-electrode distance is of the same order as the electrode-ion distance, expanding up to quadratic terms will usually suffice. Terms like x^2 and y^2 are sufficient to describe harmonic potentials whose axes lie along the coordinate axes, but cross terms like xz will be required to describe general harmonic potentials.

A concrete example is in order. Say we choose to expand with the functions $\{x, x^2, y, y^2, z, z^2\}$. Then each discrete k_n becomes a column vector $\mathbf{c}^{(n)}$ with

$$k_n(\mathbf{r}) \approx c_1^{(n)}x + c_2^{(n)}x^2 + c_3^{(n)}y + c_4^{(n)}y^2 + c_5^{(n)}z + c_6^{(n)}z^2 \quad (2.16)$$

and (2.14) becomes

$$\begin{bmatrix} c_1^{(1)} & \dots & c_1^{(n)} & \dots & c_1^{(N)} \\ \vdots & & \vdots & & \vdots \\ c_m^{(1)} & \dots & c_m^{(n)} & \dots & c_m^{(N)} \\ \vdots & & \vdots & & \vdots \\ c_6^{(1)} & \dots & c_6^{(n)} & \dots & c_6^{(N)} \end{bmatrix} \begin{bmatrix} \psi_1 \\ \vdots \\ \psi_n \\ \vdots \\ \psi_N \end{bmatrix} = \begin{bmatrix} 0 \\ -\frac{Q_{\text{dc}}}{2} \\ 0 \\ -\frac{Q_{\text{dc}}}{2} \\ 0 \\ Q_{\text{dc}} \end{bmatrix} \quad (2.17)$$

for V_{dc} in (2.11).

The $c_m^{(n)}$ are computed by least squares fits to the discrete k_n functions. This coefficient matrix C with columns $\mathbf{c}^{(n)}$ contains the primary features of interest about the trapping geometry. The solution to (2.17) is the basic DC confinement, but using C , we can compute many other interesting electrode configurations. For example, we often want to be able to compensate errors in the DC potential due to charging of insulators or stray fields from outside the trap system. Since these fields typically originate far from the trapping region, they can be well approximated as uniform fields—which are the linear function components in the potential.

The solution to

$$C\mathbf{e}_x = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (2.18)$$

gives a set of electrode voltages which create a uniform field in the x -direction. We may compute similar voltages configurations for y and z directed fields. We can build up a total electrode configuration as a linear combination of the various solutions we compute in this way. For example

$$\psi_{\text{tot}} = Q_{\text{dc}}\psi_t + E_x\mathbf{e}_x + E_y\mathbf{e}_y + E_z\mathbf{e}_z \quad (2.19)$$

is an electrode configuration which has four parameters, each with a clear purpose: Q_{dc} sets the z -axis confinement, and the E_i are field compensation parameters along the principal axes.

Let us examine one more example of practical importance. Suppose we wish to lift the degeneracy between x - and y -oscillations in the harmonic potential, for example, to be able to address these modes individually in frequency space. Solving the equation

$$C\mathbf{q}_{xy} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ -1 \\ 0 \\ 0 \end{bmatrix} \quad (2.20)$$

yields a 2-d quadrupole which breaks the symmetry in the xy plane. We may include it in our linear combination to split the x and y radial modes multiplied by a strength parameter to determine the size of the splitting.

2.4.3 Symmetry

Geometric symmetries also fit nicely into this framework. If the electrode structure is symmetric and voltages are applied symmetrically about the $z = 0$ plane, for example, then the resulting potentials will be symmetric too. Functions which are odd in the z parameter will have coefficients of zero in expansions. These symmetries can be used to cut down the configuration space while at the same time cutting out some of the linear constraints. If all electrodes were in symmetric pairs about z , then we could eliminate the requirement that the coefficient of z vanish by deleting that row from C and the target V .

2. SEGMENTED ION TRAP MODELLING

We account for fewer degrees of freedom in this arrangement by generating a new C with columns corresponding to the new degrees of freedom. Each column in the new arrangement is simply the sum of the columns in the old arrangement which are held at the same voltage for that degree of freedom[†].

[†]Of course, such reductions of the configuration space do not need to be motivated by geometric symmetries. For example, it may be simple laboratory convenience to drive several electrodes from the same source.

3

Measurement-Selected Ensembles

The typical approach in describing a quantum mechanical system is to describe the initial conditions, follow the evolution of the system according to the Schrödinger Equation, and then consider some measurement process on the system. In a complete (von Neumann projective) measurement, the associated complete set of measurement outcomes and corresponding final states of the system do not depend on the initial conditions, evolution, or anything that precedes the measurement—only the probability of outcomes can be affected. In this sense, a complete measurement represents a break with the past, and we may safely ignore measurement outcomes prior to the most recent complete measurement in a description of the state of the system. Thus, complete measurements of a system provide a natural segmentation to the description of quantum systems—outcomes of complete measurements provide the initial conditions to the next evolution-measurement cycle.

However, this usual prescription of considering a single iteration of the cycle

Initial Conditions → Evolution → Complete Measurement

is only one perspective. The usual state vector which represents the initial conditions is an accurate description of an ensemble of systems which are selected only based on the initial conditions. When we instead consider ensembles of systems which are selected based not only on initial conditions, but also on the results of one or more later measurements, we form what we shall denote in this work a *measurement-selected ensemble*. We shall consider several scenarios which, in the context of a measurement-selected ensemble, are described by physics that is quite different from the “usual” Hamiltonian evolution. These scenarios are unified under the theme of non-projective measurements, and may be divided into two categories: weak measurements and partial measurements.

What follows is a literature review, with brief extensions, of the two novel measurement scenarios. We shall introduce a useful formalism for describing

measurement-selected ensembles which are selected by initial conditions *and* final conditions, and then use it to review the weak measurement of an observable. We then extend this idea of weak measurement explicitly to qubit systems and discuss the outlook for implementation in trapped ion qubits. Next, we introduce another useful mathematical formalism which gives a clean geometric interpretation of the density operator for a single qubit. This interpretation will aid in our review of the partial collapse and so called “uncollapse” of a qubit wavefunction. We conclude with a description of an implementation of this scenario in a single ion of $^{40}\text{Ca}^+$.

3.1 Two-State Vector Formalism

We shall begin with a brief review of the Two-State Vector Formalism (TSVF) for measurements, originally considered by [1] to describe quantum measurement in a formalism which is symmetric with respect to time-reversal. We shall not consider its utility for that purpose here; [3] should be consulted if this is of interest. For our purpose, the TSVF leads naturally to the consideration of so-called “weak measurements,” in which we attempt to minimise the back-action of the measuring device on the quantum system, with interesting consequences.

The TSVF describes the situation in which a series of measurements are performed on a quantum system, say $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \mathcal{M}_4, \dots$, etc. Standard quantum mechanics uses the state vector $|\Psi\rangle$ to keep track of information about the possible outcomes of a measurement in the sequence, and it is only affected by results of measurements in the past. For example, the state just before \mathcal{M}_2 is affected by the outcome of \mathcal{M}_1 , but not by $\mathcal{M}_3, \mathcal{M}_4, \dots$, etc. However, the results of \mathcal{M}_2 will, in general, be correlated with the outcomes of future measurements. The standard interpretation is to consider the results of \mathcal{M}_2 propagating forward in time to affect \mathcal{M}_3 , etc., but it need not be so. Suppose we consider a question of the following form: given that we measure outcome a in \mathcal{M}_1 , and outcome c in \mathcal{M}_3 , what probabilities do we expect to observe for outcomes $\{b_n\}$ in \mathcal{M}_2 ?

Using the standard formalism, we could work this all out by considering the probabilities of \mathcal{M}_2 given a , $P(b_n|a)$, and then the probabilities of c given the state vectors resulting from outcomes of \mathcal{M}_2 , $P(c|b_n)$, and work back using Bayes Theorem. This method is explicitly tied to the form of \mathcal{M}_2 and its associated outcomes.

Alternatively, we may consider a state vector evolving forward in time from $\mathcal{M}_1, |\Psi\rangle$, and a state vector evolving backward in time from $\mathcal{M}_3, \langle\Phi|$.

$$\begin{array}{ccccccc} \text{Initial Conditions} & \rightarrow & \text{Evolution} & \rightarrow & \text{Measurement} & \leftarrow & \text{Evolution} & \leftarrow & \text{Measurement} \\ \mathcal{M}_1 \rightarrow |a\rangle & & \rightarrow & & |\Psi\rangle & & \rightarrow & & \mathcal{M}_2 & & \leftarrow & & \langle\Phi| & & \leftarrow & & \langle c| \leftarrow \mathcal{M}_3 \end{array}$$

The *two-state vector* $\langle \Phi | | \Psi \rangle$ then gives the complete description of the statistics for \mathcal{M}_2 via the Aharonov-Bergmann-Lebowitz rule [1]

$$P(b_n) = \frac{|\langle \Phi | \hat{B}_n | \Psi \rangle|^2}{\sum_j |\langle \Phi | \hat{B}_j | \Psi \rangle|^2} \quad (3.1)$$

where \hat{B}_n are projector operators for outcomes b_n . In the case where $\hat{H} = 0$ for the system, $|\Psi\rangle = |a\rangle$ and $\langle \Phi| = \langle c|$ for the above situation. In the case $\hat{H} \neq 0$, the states evolve as normal according to the Schrödinger Equation forward (backward) from the instant of \mathcal{M}_1 (\mathcal{M}_3). The two-state vector $\langle \Phi | | \Psi \rangle$ is a complete description that does not depend on an explicit form for \mathcal{M}_2 , subject to the constraints of past *and future* measurements, in the same sense as the state vector describes possible outcomes subject to the constraints of past measurements.

How do we interpret the two-state vector? In this work, which is primarily experimental, we can straightforwardly see the two-state vector as a compact representation of an ensemble of experiments. If we repeat a sequence of measurements over many experimental runs, we may select only those experiments with particular outcomes in particular measurements, e.g. the results of \mathcal{M}_1 and \mathcal{M}_3 being a and c . The two-state vector then should match the results of \mathcal{M}_2 within that ensemble.

We now introduce a simple symbolic language which will allow us to discuss these and other measurement scenarios with precision and brevity. As we mentioned in the introduction, complete projective measurements provide a natural segmentation in our description of quantum systems, each of which we denote with a \bullet . In the previous discussion, we made a distinction between measurements whose results we use to select our ensemble and those for which we wish to calculate probabilities of results. The former were measurements like \mathcal{M}_1 and \mathcal{M}_3 , which we might call the “selection” measurements. We will denote these with a bracket underneath: \bullet . The measurements for which we calculate probabilities, like \mathcal{M}_2 , are the ones for which some notion of a “quantum state” can be defined, such as a state vector or two-state vector. We denote these measurements with angled brackets: \diamond . Time evolution will be represented with horizontal lines and time progresses from left to right. The typical quantum mechanical description which takes into account measurements in the past is $\bullet \text{---} \diamond$, which is called a *pre-selected* ensemble by Aharonov et al. An ensemble which takes into account results of both past and future measurements is $\bullet \text{---} \diamond \text{---} \bullet$, which is called a *pre- and post-selected* ensemble. In this thesis we take the general term *measurement-selected ensemble* to include these and other (Section 3.4) ensembles.

For completeness we mention that Aharonov and Vaidman also describe scenarios which are post-selected only, $\diamond \text{---} \bullet$, corresponding to a single backward-

evolving state vector $\langle \Phi |$. The physical realisation of such a system involves performing a Bell-type measurement on the system with an ancilla and guarding the ancilla from possible measurements until the post-selection is complete [3], but we will not consider this in any further detail in this work.

3.2 Weak measurement

So far, we have implicitly considered the situation where each measurement considered was complete in the sense that it definitely projected the system into an exactly known quantum state. In the case where the measured operators do not commute, such measurements irrevocably alter the system. The opposite limit is to consider a measurement designed to minimise the disturbance of the two-state vector by weakening the interaction between the system and measuring device. Given sufficiently weak coupling, the effect of the weak measurement negligibly alters the statistics of the pre- and post-selection measurements which establish the ensemble. In accordance with the Uncertainty Principle, we also gain negligible information about the state of a system from a single shot of a weak measurement experiment. However, repeating the experiment many times will allow accurate measurement of the small change in the measuring device—yielding the *weak value* [3]

$$B_w = \frac{\langle \Phi | \hat{B} | \Psi \rangle}{\langle \Phi | \Psi \rangle} \quad (3.2)$$

of the observable \hat{B} for the pre- and post-selected ensemble. This gives the expectation value of observations of the measuring device, suitably scaled by the coupling strength. Its derivation involves considering the interaction of the system with a measuring device, an example of which is given in (3.17).

Symbolically, we denote a weak measurement \circ .

In the case of a quantum system which is pre-selected only, $\bullet \rightarrow \circ$, the weak value reduces to the expectation value of the system [3], but for some cases of pre- and post-selected ensembles, $\bullet \rightarrow \circ \rightarrow \bullet$, the weak value can take a value which lies outside the range of the eigenvalue spectrum of \hat{B} .

Let us consider an example where this is the case. Imagine making measurements on a spin- $\frac{1}{2}$ particle. We may measure the angular momentum along an axis defined by \mathbf{r} , with corresponding observable $\hat{B} = \boldsymbol{\sigma} \cdot \mathbf{r}$, and we take the eigenvalues of the observable to be ± 1 , as usual. Say we begin with a measurement along the z -axis, with observable σ_z , and select an ensemble with outcome $+1$. We denote the state-vector for this ensemble $|+Z\rangle$. Next, we measure along an axis at 45° to both the z - and x -axes: $\mathbf{r} = \frac{1}{\sqrt{2}}(1, 0, 1)$ with observable

$$\hat{B} = \frac{\sigma_x + \sigma_z}{\sqrt{2}} \quad (3.3)$$

The expectation value for this measurement is $\langle +Z | \hat{B} | +Z \rangle = \frac{1}{\sqrt{2}}$, which we obtain from averaging many runs of the experiment, each with outcome $+1$ or -1 . If we reduce the coupling between our measurement apparatus and the spin such that measurement of \hat{B} becomes a weak measurement, then we still obtain this expectation value with enough averaging, but individual measurement results are now in a continuous range, no longer mapping directly to the eigenvalues ± 1 . By reducing the coupling, we do not disturb the initial state of the spin significantly, and we may perform additional measurements on it.

After the weak measurement of \hat{B} , we perform a complete measurement along the x -axis. The expectation value for this measurement $\langle \sigma_x \rangle = 0$, corresponding to equal probability of measuring $+1$ or -1 . If we perform a post-selection on the results of the weak measurement, such that we only consider those runs where the result of the final measurement was $+1$, then we may use (3.2) to obtain the average value of the weak measurement results.

$$B_w = \frac{\langle +X | \frac{1}{\sqrt{2}} (\sigma_x + \sigma_z) | +Z \rangle}{\langle +X | +Z \rangle} \quad (3.4)$$

$$= \langle +X | \sigma_x | +Z \rangle + \langle +X | \sigma_z | +Z \rangle \quad (3.5)$$

$$= \sqrt{2}, \quad (3.6)$$

which is outside the eigenvalue spectrum of \hat{B} , and also corresponds to the result we might expect classically if we measured angular momentum along z then along x and got $+1$ for both.

Weak values have also been closely associated with interpretation of counterfactual statements and the apparent paradoxes they introduce. In particular, they have been proposed as a resolution to Hardy's Paradox [2].

The derivation of weak-values and experimental methods to access them assume a continuous-valued measurement system. I propose a method to perform a weak measurement where both the system to be measured and the "measuring device" will be a single qubit—there will be no need to introduce a continuous variable pointer system. This work is motivated by the observation that in quantum information, the emphasis is on obtaining exquisite quantum control over the discrete state space of qubits, whereas high precision measurement of continuous state spaces is often more difficult. As I shall demonstrate, using an appropriate "weakened" version of an entangling gate, we can reduce the problem of determining the weak value of an operator to measuring the spin state in an auxiliary qubit repeatedly along an appropriately chosen axis.

3.2.1 Review of continuous pointer measurements

We shall briefly review the measurement formalism developed by von Neumann [64], and its extension to weak measurements, as these will serve our

development of the qubit case by analogue.

We begin by considering a generic quantum system, S , with an operator \hat{M} corresponding to an observable we wish to measure. We also have a measurement system, D , called the “pointer” which consists of a one dimensional continuous state space. \hat{Q} and \hat{P} are conjugate observables ($[\hat{Q}, \hat{P}] = i$) whose spectra span the continuous state space. One possible pointer is a free particle, in which case, \hat{Q} and \hat{P} correspond to position and momentum (within a factor of \hbar), but many others are possible. We prepare the system in a Gaussian wavepacket centered around $Q = 0$, with width Δ .

$$\Psi_D(Q) = (\Delta^2 \pi)^{-1/4} e^{-Q^2/2\Delta^2} \quad (3.7)$$

The initial state of S will be arbitrary, but we shall express this as an eigenvalue decomposition in \hat{M}

$$\Psi_S = \sum_k \alpha_k |m_k\rangle \quad (3.8)$$

We allow the systems to briefly couple via an interaction Hamiltonian

$$\hat{H}_i = \lambda(t) \hat{M} \hat{P} \quad (3.9)$$

where we have written $\lambda(t)$ to indicate that we can control the coupling strength in time. In particular, we shall consider the case where λ is quickly switched on to a constant value and then switched off. The “on time” will be short compared to the timescales associated with the Hamiltonians of both the system and the pointer such that they can safely be neglected in our analysis of the measurement process. Since λ is constant during this process the time-dependent Schrödinger equation then gives rise to the propagator

$$\mathcal{U}(t) = e^{-i\lambda \hat{M} \hat{P} t} \quad (3.10)$$

Applying this to our initial state $\Psi_S \Psi_D$, we have

$$\Psi(t) = \mathcal{U}(t) \Psi_S \Psi_D(Q) \quad (3.11)$$

$$= e^{-i\lambda \hat{M} \hat{P} t} \sum_k \alpha_k |m_k\rangle \Psi_D(Q) \quad (3.12)$$

$$= \sum_k \alpha_k e^{-i\lambda m_k t \hat{P}} |m_k\rangle \Psi_D(Q) \quad (3.13)$$

Recalling that \hat{Q} and \hat{P} are conjugate operators, or alternatively, that momentum is the *generator* for translation, the operator $e^{-i\lambda m_k t \hat{P}}$ shifts the wavefunction by an amount $\lambda m_k t$ in the Q variable, i.e. $\Psi_D(Q) \rightarrow \Psi_D(Q - \lambda m_k t)$. Thus

$$\Psi(t) = \sum_k \alpha_k |m_k\rangle \Psi_D(Q - \lambda m_k t) \quad (3.14)$$

$$= (\Delta^2 \pi)^{-1/4} \sum_k \alpha_k |m_k\rangle e^{-(Q - \lambda m_k t)^2 / 2\Delta^2} \quad (3.15)$$

and it follows immediately that the probability distribution for the measurement of the pointer's position (by tracing over the system, S) is

$$\text{Prob}(Q) = (\Delta^2 \pi)^{-1/2} \sum_k |\alpha_k|^2 e^{-(Q - \lambda m_k t)^2 / \Delta^2} \quad (3.16)$$

If the $\lambda m_k t$ are all large compared with Δ , then this will be a distribution with sharp peaks centered on each eigenvalue and weighted by the probabilities in S . The probability of measuring a Q not near an eigenvalue of \hat{M} is essentially zero, and a measurement projects the system with near certainty into the corresponding eigenvalue.

In the opposite limit, where the $\lambda m_k t$ are all small compared with Δ (which we can accomplish either by making the coupling λ sufficiently weak or the interaction time sufficiently short), we no longer have well separated peaks and a single measurement cannot determine the state of S with certainty. Instead, we can approximate (3.16) as a single Gaussian shifted by the weighted average of the eigenvalues, i.e. the expectation value.

$$\text{Prob}(Q) \approx (\Delta^2 \pi)^{-1/2} e^{-(Q - \lambda M_w t)^2 / \Delta^2} \quad (3.17)$$

where $M_w = \sum_k |\alpha_k|^2 m_k$, is called the "weak value", which in this case is simply the expectation value. In general, as described in Aharonov & Vaidman [3], the weak value will depend on both pre- and post-selected measurement results, and so can take values other than the expectation value.

3.2.2 Qubit pointer measurements

We shall proceed by analogy to consider the case where both the system S and the pointer, D are single qubits.

Our prototype interaction is the *controlled-not*, or CNOT gate. This is an entangling operation between two qubits and, when used in the configuration where S is the control bit and D is the target bit, allows S to be measured by performing a measurement on D . Consider the case where D is prepared in the state $|0\rangle$, and S is in the canonical arbitrary state $a|0\rangle + b|1\rangle$. Then the CNOT gate maps this input state as follows:

$$(a|0\rangle + b|1\rangle) |0\rangle \longrightarrow a|00\rangle + b|11\rangle \quad (3.18)$$

We seek to describe a weakened version of this interaction analogous to the interaction Hamiltonian in (3.9). There is no difficulty in considering S to be a qubit, since we placed no constraints on it in the previous discussion other than that it have a finite spectrum. We take the observable in S that we wish to measure to be the ordinary computation basis, that is, \hat{Z}_S . And, recalling that the controlled-not gate is sometimes called the *controlled- X* gate, we take the pointer operator to be \hat{X}_D .

$$\hat{H}_i = \lambda(t) \hat{Z}_S \hat{X}_D \quad (3.19)$$

What is the effect of this Hamiltonian on the total system state? Recall that the effect of (3.9) was to entangle the system eigenstates of \hat{M} with a pointer translated by an amount λt times the corresponding eigenvalue. In this case, we have replaced the pointer conjugate momentum operator with the pointer's \hat{X} operator. If we imagine the pointer qubit as a spin, the \hat{X} operator is the angular momentum component along the x -axis. In the same way that our conjugate momentum generated a translation on D , the angular momentum in this Hamiltonian will generate a rotation about the x -axis.

Since we will be considering the effect of rotations in D about the x -axis, it is natural to parametrise the state of the pointer by the angle the state-vector makes with the $+z$ -axis in the yz -plane, which we shall denote θ . Note that we do not span the entire qubit state-space with a single parameter, but our initial state $|0\rangle$ and any subsequent state which can be created by rotations about the x -axis will be in the span of this parameter. Explicitly expressed in the computation basis

$$|\theta\rangle = \cos \frac{\theta}{2} |0\rangle - i \sin \frac{\theta}{2} |1\rangle \quad (3.20)$$

and conveniently,

$$|\theta = 0\rangle = |0\rangle \quad (3.21)$$

We now return to the effect of the propagator generated by the Hamiltonian (3.19)

$$\mathcal{U}(t) = e^{-i\lambda \hat{Z}_S \hat{X}_D t} \quad (3.22)$$

on the initial state

$$\Psi_S \Psi_D = a |0\rangle_S |0\rangle_D + b |1\rangle_S |0\rangle_D \quad (3.23)$$

The effect on S , since it is expressed in the computation basis, is to replace \hat{Z}_S with eigenvalues ± 1 , and the effect on D is to rotate:

$$\Psi(t) = \mathcal{U}(t) (a |0\rangle_S |0\rangle_D + b |1\rangle_S |0\rangle_D) \quad (3.24)$$

$$= ae^{-i\lambda t \hat{X}_D} |0\rangle_S |0\rangle_D + be^{+i\lambda t \hat{X}_D} |1\rangle_S |0\rangle_D \quad (3.25)$$

$$= a |0\rangle_S |\theta = \lambda t\rangle_D + b |1\rangle_S |\theta = -\lambda t\rangle_D \quad (3.26)$$

The pointer, in the Bloch representation, can be thought to “point,” like a hand on a dial, to the value of the \hat{Z} observable of S —the pointer is entangled with states of S . There is perfect correlation between measurements on D and the state of S when $\lambda t = \pi/2$, such that the two pointer states are diametrically opposite along the y -axis. Subsequent measurement of D in the \hat{Y} basis corresponds to a perfect “strong” measurement of S . We note however that the choice of initial state of D was arbitrary, and if we choose $|\theta = -\pi/2\rangle$ (the $+y$ eigenstate) as the initial pointer state, then we can make our measurement in the familiar computation basis, i.e.

$$\mathcal{U}\left(t = \frac{\pi}{2\lambda}\right) (a |0\rangle_S + b |1\rangle_S) \left|\theta = -\frac{\pi}{2}\right\rangle_D = a |0\rangle_S |0\rangle_D + b |1\rangle_S |1\rangle_D \quad (3.27)$$

For simplicity of notation in developing the ideas of measurement of the pointer, we shall proceed with this convention of the pointer initially in the $|\theta = -\pi/2\rangle$ state, and use $\Lambda \equiv \lambda t$ to reduce the “strength” of the measurement to a single parameter.

We perform a weak measurement by reducing the coupling such that $\Lambda \ll \pi/2$. In this limit, we do not disturb the pointer very far from its initial state, but, more importantly, we do not disturb S very far from its initial state. Thus, we do not have a significant effect on future evolution and measurements of S . In this limit we lose the ability to say anything of consequence about S from a single measurement of D , but with repeated measurements, we can determine the squared amplitudes $|a|^2$ and $|b|^2$ to any desired accuracy.

Let us calculate the expectation value of the measurements on D as a function of $|a|^2$, $|b|^2$, and of course Λ . After the interaction, the system is in the state

$$\Psi = a |0\rangle_S \left|\theta = -\frac{\pi}{2} + \Lambda\right\rangle_D + b |1\rangle_S \left|\theta = -\frac{\pi}{2} - \Lambda\right\rangle_D \quad (3.28)$$

which, after application of (3.20) and a few trigonometric manipulations can be written

$$\begin{aligned} \Psi = & \cos\left(\frac{\Lambda}{2} - \frac{\pi}{4}\right) (a |00\rangle + ib |11\rangle) + \\ & \cos\left(\frac{\Lambda}{2} + \frac{\pi}{4}\right) (ia |01\rangle + b |10\rangle) \end{aligned} \quad (3.29)$$

where we have condensed terms like $|0\rangle_S |1\rangle_D$ to $|01\rangle$. Since Λ is small we can perform a Taylor expansion

$$\cos\left(\frac{\Lambda}{2} \mp \frac{\pi}{4}\right) \approx \frac{1}{\sqrt{2}} \left(1 \pm \frac{\Lambda}{2}\right) \quad (3.30)$$

then we have

$$\Psi = \frac{1}{\sqrt{2}} \left[\left(1 + \frac{\Lambda}{2}\right) (a|00\rangle + ib|11\rangle) + \left(1 - \frac{\Lambda}{2}\right) (ia|01\rangle + ib|10\rangle) \right] \quad (3.31)$$

The expectation value of the \hat{Z}_D operator for this state is

$$\langle \hat{Z}_D \rangle = \frac{1}{2}(|a|^2 - |b|^2) \left[\left(1 + \frac{\Lambda}{2}\right)^2 - \left(1 - \frac{\Lambda}{2}\right)^2 \right] \quad (3.32)$$

$$= \frac{1}{2} \langle \hat{Z}_S \rangle \left[\left(1 + \frac{\Lambda}{2}\right)^2 - \left(1 - \frac{\Lambda}{2}\right)^2 \right] \quad (3.33)$$

$$= \langle \hat{Z}_S \rangle \Lambda \quad (3.34)$$

The expectation value for measurements of D is the expectation value for measurements of S scaled by the coupling. Furthermore, in the case of post-selected ensembles, a measurement interaction of this form will yield the weak value of \hat{Z}_S scaled by the coupling, as demonstrated by [8], who obtain a joint state similar to (3.28) via a different interaction scheme. [69] observe that spin-spin interactions like (3.19) are a cornerstone of trapped-ion experiments in recent years and may be implemented with relative ease—indeed our own group has implemented a suitable interaction in the demonstration of the 2-ion geometric phase gate [23].

3.3 Bloch vector representation of the density operator

A pure state of a single qubit is often represented by a point on the surface of the unit sphere in 3-dimensions, called the *Bloch sphere*. The point, a vector in \mathbb{R}^3 , given in spherical polar coordinates $\mathbf{r} = (1, \theta, \varphi)$ is called the *Bloch vector*, and describes a qubit in the state

$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\varphi}\sin\frac{\theta}{2}|1\rangle \quad (3.35)$$

which parametrises the entire single-qubit Hilbert space up to an unobservable overall phase factor.

With the convention that $|0\rangle$ and $|1\rangle$ have eigenvalues +1 and -1, we note that expressing the Bloch vector in the Cartesian coordinates gives the expectation values for the Pauli spin operators:

$$\mathbf{r} = (\langle \sigma_x \rangle, \langle \sigma_y \rangle, \langle \sigma_z \rangle). \quad (3.36)$$

From this perspective, we may generalise to represent mixed states as well. The density operator for a single qubit is 2×2 and Hermitian, and therefore can be expressed as superposition of 4 linearly independent operators. A convenient basis is that of the Pauli operators plus the identity, since coefficients in the superposition will be real. In fact, the coefficients are directly related to the expectation values of the corresponding operator. Again, with eigenvalue spectrum $\{+1, -1\}$ for each operator

$$\rho = \frac{1}{2} \hat{I} + \frac{\langle \sigma_x \rangle}{2} \sigma_x + \frac{\langle \sigma_y \rangle}{2} \sigma_y + \frac{\langle \sigma_z \rangle}{2} \sigma_z \quad (3.37)$$

$$= \frac{\hat{I} + \sigma \cdot \mathbf{r}}{2}. \quad (3.38)$$

where we have assumed ρ is appropriately normalised to give $\text{tr } \rho = 1$. This assumption fixes the first parameter, leaving just 3 real parameters which exactly correspond to the components of the Bloch vector in (3.36). Therefore, when the qubit is in a pure state, the Bloch vector coincides with the surface of the unit sphere, as before, and when it is mixed the Bloch vector may be in the interior of the sphere. Up to an unobservable overall phase, the set $\{|\mathbf{r}| \leq 1\}$ parametrises the set of normalised density operators of a single qubit.

Expressing the density operator as a real vector in real space puts it in a realm which we can easily imagine and analyse, but we caution that even though the space is mathematically real (as opposed to complex), in general, it does not correspond directly to any physical space like a laboratory*.

A quantum process is a map from the space of density matrices onto itself. Quantum processes which do not take into account the results of measurements will be linear maps on the space of density matrices.

On the other hand, if we take into account the results of measurements, the resulting density matrix will be of the form

$$\rho' = \frac{\hat{M}_m \rho \hat{M}_m^\dagger}{\text{tr}(\hat{M}_m \rho \hat{M}_m^\dagger)} \quad (3.39)$$

where a measurement result m is associated with measurement operator[†] \hat{M}_m . The numerator in this equation is a linear map, but the renormalisation associated with the denominator is non-linear. We have two options, then, for how to proceed with the analysis. If we ignore the denominator and work with

*There is one special case in which it does: an actual $J = \frac{1}{2}$ spin system. We return to a further generalisation of this idea in Section 7.1.2.

[†]we adopt the terminology of [42]

un-normalised density operators, then we can use all the tools and techniques which apply to linear maps (e.g. linear algebra). The disadvantage is that the trace of the density matrix is not preserved, and probabilities when considering future measurements will not sum to 1. For qubits, we also lose the useful geometric interpretation introduced above—an un-normalised density operator cannot be conveniently expressed as a Bloch vector.

Alternatively, we may use (3.39) as written. This choice is adopted in the following sections. The maps associated with our quantum processes will be non-linear, but the resulting states will be appropriately normalised. Having probabilities which sum to 1 is an important consideration to being able to discuss an ensemble of quantum systems as an entity which stands on its own, and in the case of qubits, makes available the Bloch vector interpretation.

We may visualise the effect of a single-qubit quantum process by considering its effect on the Bloch sphere. The maps we consider in this work are relatively smooth, such that we may readily grasp their essential character by considering the effect of a set of points on the surface of the sphere. This gives the mapping explicitly for pure states, but we may readily imagine a similar effect for mixed states.

A quantum process of unitary evolution, for example, is just a rotation of the Bloch sphere. Below, we see an example of a rotation about the y -axis by angle θ .




Figure 3.1: Rotation of the Bloch sphere about the y -axis by an angle θ .



Non-unitary processes deform the Bloch sphere. Consider the dephasing quantum process, which flips the phase of the qubit with probability $p/2$, where $0 \leq p \leq 1$ is the strength of the dephasing. This compresses the Bloch sphere along the x - and y -axes as phase information is gradually lost. Figure 3.2 displays this graphically.




Figure 3.2: Dephasing process as illustrated by its deformation of the Bloch sphere. p is the strength of dephasing (see text). In the case $p = 1$ the sphere has collapsed to a single line along the z -axis.


3.4 Partial collapse and uncollapse of the wavefunction

In Section 3.2 we considered the limiting case where a quantum measurement was so weak as to cause negligible disturbance to the measured system, with the projective measurement the opposite extreme. We now turn our attention to an intermediate case where the measurement process is strong enough to affect the measured system, but not strong enough to fully project it. We are motivated by the understanding, discussed in the first section, that the interaction of a quantum system with a measuring device is never ideal or instantaneous—any real interaction will take finite time to complete.

We regard the measurement process as a dynamical evolution of both systems with full projection as an endpoint—a special case of a much richer theory predicted by the measurement postulate. If the measurement interaction is halted before the systems are fully entangled, the resulting state of the system is not an eigenstate of the observable, and furthermore the final state is not fully determined by the result of the measurement alone; it depends on *both* the measurement outcome and the initial state. The system has undergone non-unitary evolution, but still contains some information about its original (in general, superposition-) state prior to measurement. This is the partial collapse of the wavefunction.

Since the quantum state following a measurement depends on the outcome of the measurement itself, we cannot assign a definite quantum state to an ensemble of all “shots” of the experiment where the measurement result is unspecified. We can only assign a definite quantum state to an ensemble selected on both the initial state and the result of the partial measurement.

Let us denote a partial measurement \diamond , and consider the measurement scenarios encountered so far:



Since we consider the state of the system after the partial collapse, we might be tempted to refer to this as a pre-selected ensemble in accordance with the two-state vector formalism of Section 3.1. I wish to call attention to an important distinction. In Section 3.1 the measurements in the pre-selection were complete, projective measurements (\bullet)—fully defining the forward-propagating state vector based on their outcome alone. In the partial collapse measurement-selected ensemble, we only partially define the state vector based on the result of partial measurement—it also depends on initialisation or other, complete measurements before the partial collapse. For this reason, I shall reserve use of the term pre-selected ensemble for $\bullet \rightarrow \diamond$, in which the selection is based only on complete measurements. Elsewhere in the literature, “weak” measurement is sometimes used to describe any measurement in which the entangling interaction has been weakened compared with a projective measurement [28]. In this thesis we shall reserve the term weak measurement to refer to a measurement of the weak value in the sense of (3.2).

3.4.1 Partial collapse

Let

$$\zeta \equiv \cos\left(\frac{\Lambda}{2} - \frac{\pi}{4}\right) \text{ and} \quad (3.40)$$

$$\eta \equiv \cos\left(\frac{\Lambda}{2} + \frac{\pi}{4}\right) \quad (3.41)$$

then (3.29) may be written

$$|\psi\rangle = \zeta(a|00\rangle + ib|11\rangle) + \eta(ia|01\rangle + b|10\rangle) \quad (3.42)$$

The state with aligned spins occurs with a probability $|\zeta|^2$, whereas the anti-aligned spins occurs with probability $|\eta|^2$. With $0 \leq \Lambda \leq \pi/2$, we have $\zeta > \eta \geq 0$. When we measure the pointer qubit, we extract some information about the system qubit—but our information is incomplete because if we make the “best guess” that the states agree then we will be wrong with probability η^2 . If we

have access to many copies of the system qubit, we may perform the partial measurement many times to determine $|a|$ and $|b|$ to any desired precision.

$$P(|0\rangle_D) = P(|00\rangle) + P(|10\rangle) \quad (3.43)$$

$$= \zeta^2 |a|^2 + \eta^2 |b|^2 \quad (3.44)$$

$$= (\zeta^2 - \eta^2) |a|^2 + \eta^2 \quad (3.45)$$


$s \equiv (\zeta^2 - \eta^2)$ is a measure of the strength of the measurement. For $s = 0$ we gain no information; $s = 1$ corresponds to projective measurement. For a fixed number of measurements, we will always retain larger uncertainty (by a factor of s^{-1}) using the partial measurement than we would if we had performed a projective measurement.

In the case of finding the pointer in $|0\rangle$, the system qubit is left in the state

$$|\psi\rangle_s = \frac{\zeta a |0\rangle + \eta b |1\rangle}{\zeta^2 |a|^2 + \eta^2 |b|^2} \quad (3.46)$$

(this is obtained by applying the appropriate measurement operator and tracing out the pointer, D . The measurement operator for the pointer in $|0\rangle$ is $\hat{M}_{D=0} = |00\rangle\langle 00| + |10\rangle\langle 10|$.) The amplitudes of the $|0\rangle$ and $|1\rangle$ states have been scaled by different amounts which depend on the initial state: a non-unitary transformation. However, the final state of the system qubit is pure. The partial collapse operation extracts incomplete information from the system qubit, altering its state, but does not “erase” information about the initial state, including, in particular, the coherence.

3.4.2 Uncollapse

Although the mapping of input state to output state of the partially-collapsed ensemble is non-unitary, the mapping is still one-to-one and remains so as the strength of the partial measurement is increased up until the point where it becomes a projective measurement. This leads us to consider the possibility of undoing the partial collapse by applying some process whose mapping is the inverse. Such a mapping will, of course, be non-unitary, so we will be unable to implement it by quantum gates (Hamiltonian evolution). However, it can be realised by a second measurement provided a particular result is obtained. In this way we can form a second ensemble, subset of the first, which contains systems in which the partial collapse was successfully reversed, a procedure we denote uncollapse, and draw .

The uncollapse procedure will, in general, succeed with probability $P \leq 1$, with $P = 1$ only if we either already have complete knowledge of the initial state, or if the coupling with the measurement device is identically 0, in which

case we gained no information from the measurement. However, we will be able to tell in each case whether or not the uncollapse succeeds, and the uncollapse procedure itself is independent of the initial state. The measurement-selected ensemble will be completely identical to the original ensemble before both measurements, albeit smaller in size (we shall shortly return to the question of how much smaller).

A quantum measurement is defined by a set of operators $\{\hat{M}_m\}$, with each m the associated measurement result. We may define POVM elements associated with each operator

$$\hat{E}_m = \hat{M}_m^\dagger \hat{M}_m \quad (3.47)$$

and note that each \hat{M}_m may be written

$$\hat{M}_m = \hat{U}_m \sqrt{\hat{E}_m} \quad (3.48)$$

where \hat{U}_m is a unitary operator.

The reversal of a partial collapse associated with a measurement result \hat{M}_m may be accomplished by a second measurement with associated operator [32]

$$\hat{L} = C \hat{E}_m^{-1/2}. \quad (3.49)$$

where C is a complex constant. More generally, the measurement operator could be \hat{L} multiplied on the left and right by unitary matrices \hat{U}_L and \hat{V}_L , which must be undone before and after the uncollapse measurement—this gives us more freedom to choose a convenient \hat{L} if we have complete coherent control over the system. The operator \hat{U}_m^\dagger is applied to undo the unitary part of \hat{M}_m , and then the measurement performed. If we get the result associated with \hat{L} then the final state is

$$|\psi_u\rangle = \frac{\hat{L} \hat{U}_m^\dagger \hat{M}_m |\psi\rangle}{\mathcal{N}} \quad (3.50)$$

$$= \frac{C |\psi\rangle}{\mathcal{N}} \quad (3.51)$$

$$= |\psi\rangle \quad (3.52)$$

with \mathcal{N} the usual normalisation constant, which we immediately identify as C . This constant is determined by the measurement protocol and there are some restrictions on its value related to the nature of the measurement we are attempting to uncollapse. For \hat{L} to be physically realisable, $\hat{L}^\dagger \hat{L}$ must have eigenvalues ≤ 1 . If $k_i^{(m)}$ are eigenvalues of \hat{E}_m , then the eigenvalues of $\hat{L}^\dagger \hat{L} = |C|^2 \hat{E}_m^{-1}$ are $|C|^2 / k_i^{(m)}$. This leads to the inequality [32]

$$|C|^2 \leq \min_i k_i^{(m)} = \min_{|\psi\rangle} P(m) \quad (3.53)$$

$P(m)$ is the probability of result m , and we take its minimum over all possible input states $|\psi\rangle$. Consider two extremes of this quantity: if $\min P(m) = 1$, then we always get result m regardless of the input state. The measurement extracts no information from the system and its associated operator is unitary, which we may trivially “uncollapse” by unitary evolution amounting to its inverse. If $\min P(m) = 0$, then the measurement operator is a projector, and it will not be possible to uncollapse an unknown state.

We noted previously that C was the normalisation constant for the uncollapse procedure. $|C|^2$ is therefore the probability of getting measurement result m and then undoing the measurement by successful uncollapse. A subtly different quantity, the probability of successful uncollapse given result m has already been obtained is

$$P_U = \frac{|C|^2}{P(m)} \quad (3.54)$$

Returning to our example, the measurement operators (which we now express in the space of operators on a *single* qubit—condensing the entanglement with the pointer and measurement of its state into a single operation on the system qubit, S) are

$$\hat{M}_0 = \zeta |0\rangle\langle 0| + \eta |1\rangle\langle 1| = \begin{pmatrix} \zeta & 0 \\ 0 & \eta \end{pmatrix} \quad (3.55)$$

$$\hat{M}_1 = \eta |0\rangle\langle 0| + \zeta |1\rangle\langle 1| = \begin{pmatrix} \eta & 0 \\ 0 & \zeta \end{pmatrix} \quad (3.56)$$

The POVM element associated with \hat{M}_0 is

$$\hat{E}_0 = \hat{M}_0^\dagger \hat{M}_0 \quad (3.57)$$

$$= \begin{pmatrix} \zeta^2 & 0 \\ 0 & \eta^2 \end{pmatrix} \quad (3.58)$$

from which we can see immediately that $\sqrt{\hat{E}_0} = \hat{M}_0$, since $\zeta > \eta \geq 0$. Thus the uncollapse measurement operator is

$$\hat{L} = C \hat{M}_0^{-1} \quad (3.59)$$

$$= \begin{pmatrix} \frac{C}{\zeta} & 0 \\ 0 & \frac{C}{\eta} \end{pmatrix} \quad (3.60)$$

The constraints on C discussed earlier give $C \leq \eta$. Equality in this statement is the ideal uncollapse procedure (which is to say, the one with greatest probability of success):

$$\hat{L}_{\text{ideal}} = \begin{pmatrix} \frac{\eta}{\zeta} & 0 \\ 0 & 1 \end{pmatrix} \quad (3.61)$$

After the first partial measurement, the two-qubit system, $S \otimes D$ is in a product state. Using (3.46):

$$|\psi'\rangle = \left(\frac{\zeta a |0\rangle + \eta b |1\rangle}{|\zeta|^2 |a|^2 + |\eta|^2 |b|^2} \right) \otimes |0\rangle \quad (3.62)$$

This has the form $a' |00\rangle + b' |10\rangle$. We may implement the ideal uncollapse procedure by applying the unitary transformation

$$|00\rangle \rightarrow \frac{\eta}{\zeta} |00\rangle + \frac{\sqrt{s}}{\zeta} |01\rangle \quad (3.63)$$

$$|01\rangle \rightarrow -\frac{\sqrt{s}}{\zeta} |00\rangle + \frac{\eta}{\zeta} |01\rangle \quad (3.64)$$

$$|10\rangle \rightarrow |10\rangle \quad (3.65)$$


$$|11\rangle \rightarrow |11\rangle \quad (3.66)$$

via another CNOT-type entangling interaction, and measuring D in the computation basis. If we detect $|0\rangle_D$, the uncollapse was successful.

3.4.3 Asymmetric case

The example we have used is symmetric in the sense that there exists an analogous uncollapse procedure for the case where we obtain $|1\rangle_D$ in the partial measurement. We have not shown this explicitly because it is readily derived in the same fashion as our example for $|0\rangle_D$. We now turn our attention to an asymmetric case: uncollapse shall only be possible for one of the outcomes of the partial measurement.

Consider a quantum network involving a qubit, $a |0\rangle + b |1\rangle$, which can couple to an auxiliary state $|2\rangle$.



The auxiliary state couples to the $|1\rangle$ state with strength $p = \sin^2 \varphi$:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \varphi & \sin \varphi \\ 0 & -\sin \varphi & \cos \varphi \end{pmatrix}. \quad (3.67)$$

The measurement of the auxiliary is some sort of PMT-style detector. We get “click” or “no click” to indicate whether or not the system is in $|2\rangle$. The projector associated with “no click” is

$$P = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (3.68)$$

And, finally the X operator is the quantum not:

$$X = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (3.69)$$

The effect of this network is a partial collapse after the first measurement, if we get “no click”, followed by uncollapse if we again get “no click” in the second measurement. Up to normalisation, overall we have $XPUPU|\psi\rangle$.

$$XPUPU = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \varphi & \sin \varphi \\ 0 & -\sin \varphi & \cos \varphi \end{pmatrix} \quad (3.70)$$

$$= \begin{pmatrix} 0 & \cos \varphi & \sin \varphi \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (3.71)$$

$$(XPUPU)^2 = \begin{pmatrix} 0 & \cos \varphi & \sin \varphi \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & \cos \varphi & \sin \varphi \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (3.72)$$

$$= \begin{pmatrix} \cos \varphi & 0 & 0 \\ 0 & \cos \varphi & \sin \varphi \\ 0 & 0 & 0 \end{pmatrix} \quad (3.73)$$

which, up to normalisation, is the identity for the qubit $\{|0\rangle, |1\rangle\}$ subspace.

As we shall see, this is an attractive place to start, experimentally, because the requirements for quantum control of the entangling operation and read-out are less demanding. This procedure has been previously demonstrated in a solid-state qubit system [28], and in Chapter 8 of this work we report the first demonstration in a single trapped ion qubit.

Let us examine this process in more detail. Instead of the 4-dimensional Hilbert space of 2 qubits, we shall require only 3: a qubit with the usual basis $|0\rangle, |1\rangle$, and an auxiliary state, $|2\rangle$. The auxiliary state is detectable with a boolean observable which measures whether or not the system is in $|2\rangle$. More precisely, the measurement operators are:

$$\hat{M}_{\text{dark}} = |0\rangle\langle 0| + |1\rangle\langle 1| \quad (3.74)$$

$$\hat{M}_{\text{bright}} = |2\rangle\langle 2| \quad (3.75)$$

to which we have suggestively applied labels “dark” and “bright,” whose meaning will become clear when we discuss the physical implementation with a trapped ion.

The system begins in the usual qubit state $|\psi\rangle = a|0\rangle + b|1\rangle$. A partial collapse measurement of the system shall consist of a transition between $|1\rangle$ and $|2\rangle$, with probability $p \leq 1$, followed by a measurement as indicated above. We obtain the “bright” result with probability $p|b|^2$. Repeated measurements allow us to determine $|b|^2$ to any desired precision, but for a fixed number of measurements, the uncertainty will be a factor of p^{-1} larger than if a projective measurement had been used.

When we obtain the “bright” result, the system is projected into $|2\rangle$, and thus uncollapse will not be possible. The “dark” result, however, does not project the system to a definite state—merely back into the $\{|0\rangle, |1\rangle\}$ subspace—and we may uncollapse if $p < 1$.

Unlike in the previous, symmetric example, the measurement interaction which transitions the system from $|1\rangle$ to $|2\rangle$ need not be coherent, since in any case we follow this operation with a measurement, destroying coherence between $|1\rangle$ and $|2\rangle$.

To analyse this operation, we proceed in the density operator formalism.

$$\rho_{\text{in}} = (a|0\rangle + b|1\rangle)(a^\dagger\langle 0| + b^\dagger\langle 1|) = \begin{pmatrix} |a|^2 & ab^\dagger & 0 \\ a^\dagger b & |b|^2 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (3.76)$$

Define a quantum process \mathcal{T}_p , which is the (incoherent) transfer of population from $|1\rangle$ to $|2\rangle$, with probability p .

$$\mathcal{T}_p : \begin{pmatrix} \rho_{00} & \rho_{01} & \rho_{02} \\ \rho_{10} & \rho_{11} & \rho_{12} \\ \rho_{20} & \rho_{21} & \rho_{22} \end{pmatrix} \rightarrow \begin{pmatrix} \rho_{00} & \rho_{01}\sqrt{1-p} & \rho_{02} \\ \rho_{10}\sqrt{1-p} & \rho_{11}(1-p) & \rho_{12}\sqrt{1-p} \\ \rho_{20} & \rho_{21}\sqrt{1-p} & \rho_{22} + \rho_{11}p \end{pmatrix} \quad (3.77)$$

The effect of this process is

$$\mathcal{T}_p(\rho_{\text{in}}) = \begin{pmatrix} |a|^2 & ab^\dagger\sqrt{1-p} & 0 \\ a^\dagger b\sqrt{1-p} & |b|^2(1-p) & 0 \\ 0 & 0 & |b|^2 p \end{pmatrix} \quad (3.78)$$

The state after measurement with the “dark” result is

$$\rho_{\text{pc}} = \frac{\hat{M}_{\text{dark}} \mathcal{T}_p(\rho_{\text{in}}) \hat{M}_{\text{dark}}^\dagger}{\text{tr}(\mathcal{T}_p(\rho_{\text{in}}) \hat{M}_{\text{dark}}^\dagger \hat{M}_{\text{dark}})} \quad (3.79)$$

$$= \begin{pmatrix} |a|^2 & ab^\dagger\sqrt{1-p} & 0 \\ a^\dagger b\sqrt{1-p} & |b|^2(1-p) & 0 \\ 0 & 0 & 0 \end{pmatrix} \cdot (1 - |b|^2 p)^{-1} \quad (3.80)$$

$P(\text{dark}) = 1 - |b|^2 p$ is the probability of obtaining the “dark” result. This is the pure state

$$|\psi_{\text{pc}}\rangle = \frac{a}{\sqrt{P(\text{dark})}} |0\rangle + \frac{b\sqrt{1-p}}{\sqrt{P(\text{dark})}} |1\rangle. \quad (3.81)$$

We may visualise the effect of this partial collapse by considering its effect on the Bloch sphere for different values of p (Figure 3.3.)




Figure 3.3: Partial collapse process as illustrated by its deformation of the Bloch sphere. p is the strength of collapse (see text). Values of p chosen in this example exponentially approach 1.

The partial collapse may use coherent or incoherent dynamics to implement the $|1\rangle \rightarrow |2\rangle$ transition, and yet, in either case, the measurement-selected ensemble for the “dark” result is in a pure state. The form of (3.81) is the clearest way to see that the measurement operator for the whole process acting on the qubit subspace is

$$\tilde{M}_{\text{dark}} = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-p} \end{pmatrix} \quad (3.82)$$

We now consider how to uncollapse this measurement. Using (3.49), we obtain the uncollapse measurement operator

$$\tilde{L} = C\tilde{E}_{\text{dark}}^{-1/2} \quad (3.83)$$

$$= C\tilde{M}_{\text{dark}}^{-1} \quad (3.84)$$

$$= \begin{pmatrix} C & 0 \\ 0 & \frac{C}{\sqrt{1-p}} \end{pmatrix}. \quad (3.85)$$

The ideal uncollapse has $C/\sqrt{1-p} = 1$, so

$$\tilde{L}_{\text{ideal}} = \begin{pmatrix} \sqrt{1-p} & 0 \\ 0 & 1 \end{pmatrix}. \quad (3.86)$$

Comparing this result with (3.82), we notice that it is essentially the same operation, but with $|0\rangle$ swapped with $|1\rangle$, i.e. a quantum-NOT operation followed by exactly the same partial measurement sequence. Let $\mathcal{P}_p(\rho)$ be the partial collapse process including transfer with probability p , measurement, and selection into an ensemble of “dark” results. $\mathcal{P}_p(\rho)$ is given by (3.79). The process of partial collapse followed by uncollapse, \mathcal{U}_p may be expressed

$$\mathcal{U}_p(\rho) = \hat{X}\mathcal{P}_p(\hat{X}\rho\hat{X}^\dagger)\hat{X}^\dagger \quad (3.87)$$

where $\hat{X} = |0\rangle\langle 1| + |1\rangle\langle 0|$ is the usual quantum-NOT operator. These process definitions are taken on the qubit subspace rather than the full $\{|0\rangle, |1\rangle, |2\rangle\}$ Hilbert space. For $p < 1$, $\mathcal{U}_p(\rho) = \rho$ [§]. The partial collapse followed by successful uncollapse occurs with probability $|C|^2 = 1 - p$, regardless of input state (including mixed states). This corresponds to observing the “dark” result in both partial measurements. With probability p , however, we observe the “bright” result in one or both of the partial measurements, and in that case, the system is projected to $|2\rangle$.

3.4.4 Practical value

We might wish to attempt to use the partial collapse followed by an uncollapse as a way to gain some information about an ensemble of unknown systems and still be able to leave them undisturbed. Of course, only in the limit $p = 0$ will uncollapse be assured, and in that case, our measurements tell us nothing of the system state. At best, we will be able to leave some fraction undisturbed. We therefore consider whether partial collapse followed by uncollapse is a better strategy than the straightforward one of projective measurements on some fraction equal to the number of systems we expect to be unable to uncollapse.

[§]here we only refer to the measurement-selected ensemble where both partial measurements resulted “dark”

Consider an ensemble of N systems prepared in an unknown state $a|0\rangle + b|1\rangle$, and we wish to measure $|b|^2$. By partially measuring them with strength p , and attempting to uncollapse we expect to disturb Np systems on average. However, we have the results of N partial measurements. Our uncertainty in $P(\text{dark})$ is

$$\Delta P(\text{dark}) \propto \frac{1}{\sqrt{N}} \quad (3.88)$$

so the uncertainty in $|b|^2$ is

$$\Delta|b|^2 \propto \frac{1}{p\sqrt{N}} \quad (3.89)$$

since $P(\text{dark}) = 1 - p|b|^2$. If we had simply measured Np systems projectively, our uncertainty would be

$$\Delta|b|^2 \propto \frac{1}{\sqrt{Np}}. \quad (3.90)$$

Since $p < 1$, this is clearly the better alternative. The partial collapse uncertainty is equivalent to measuring Np^2 systems projectively.

Despite our observation that uncollapse is not advantageous from the perspective of interrogating an unknown quantum state, there are situations in which it may have practical value in protecting coherence in physical systems. If a system has an incoherent or uncontrolled coupling to an unwanted auxiliary state or nearby system, and if the state of that auxiliary is measurable, then a successful uncollapse will completely undo the effects of that coupling. This is analogous to protection techniques like spin echo. Emphasising the difference, we quote Katz et al. [28]: “spin echo is the undoing of an unknown unitary transformation, while uncollapsing is the undoing of a known but non-unitary transformation.”

We claim one step further, which is that uncollapse may be useful in undoing an *unknown* non-unitary transformation, provided that the strength of the coupling we wish to undo does not change too quickly. If, as in our above example, one of our qubit levels decays with an (unknown) probability p in time τ to an auxiliary state, we may implement uncollapse by measuring, performing a quantum-NOT, waiting another time τ , and measuring again. If the coupling has not changed significantly between the two τ 's (in further analogy with spin echo), we still implement the correct uncollapse measurement operator, and uncollapse will be possible.

In this situation, uncollapse does not improve the fraction of systems which “survive” the unwanted coupling (that is to say, which are measured to have not collapsed to the auxiliary): $1 - p|b|^2$ without uncollapse, $1 - p$ with uncollapse. It undoes the non-unitary transformation of the state: a “purification” or “filtering” process.

An ideal uncollapse also acts as an integrity check which operates at the level of individual systems in an ensemble: a particular ensemble size may be *deterministically* generated in the presence of unwanted coupling by repetition until enough systems “pass” the uncollapse measurement.

3.4.5 Leaky transfer

We shall now consider an important case which results in a non-ideal partial measurement and a non-ideal uncollapse: that in which the transfer from $|1\rangle \rightarrow |2\rangle$ includes a small “leak” of population that transfers from $|1\rangle \rightarrow |0\rangle$. This case is of practical importance in general because finite population transfer errors occur in all real systems, and is of particular interest for our implementation in the $^{40}\text{Ca}^+$ ion which we further develop in Chapter 7.

We denote the “leaky” transfer \mathcal{D}_p , with strength p . It transfers a small portion, ϵp of the population ρ_{11} to ρ_{00} and $(1 - \epsilon)p$ to ρ_{22} :

$$\mathcal{D}_p : \begin{pmatrix} \rho_{00} & \rho_{01} & \rho_{02} \\ \rho_{10} & \rho_{11} & \rho_{12} \\ \rho_{20} & \rho_{21} & \rho_{22} \end{pmatrix} \rightarrow \begin{pmatrix} \rho_{00} + \rho_{11}\epsilon p & \rho_{01}\sqrt{1-p} & \rho_{02} \\ \rho_{10}\sqrt{1-p} & \rho_{11}(1-p) & \rho_{12}\sqrt{1-p} \\ \rho_{20} & \rho_{21}\sqrt{1-p} & \rho_{22} + \rho_{11}(1-\epsilon)p \end{pmatrix} \quad (3.91)$$

The leaky transfer \mathcal{D}_p takes the place of \mathcal{T}_p in the expressions for the partial and uncollapse procedures, which are (3.79) and (3.87), respectively. We denote the non-ideal versions \mathcal{P}'_p and \mathcal{U}'_p . The output states of these procedures will not be pure states for finite ϵ , and $\mathcal{U}'_p(\rho) \neq \rho$, although this will be approximately true in the limit $\epsilon \ll 1 - p$.

Explicit expressions for the output density matrices are obtained easily from the above substitution, but they are somewhat convoluted for fruitful analysis in that form. Instead, we turn to their effect on the Bloch sphere, illustrated by Figs. 3.4 and 3.5.

3.4.6 Implementation with a trapped ion qubit

The asymmetric case of partial and uncollapse of the wavefunction may be implemented in a single trapped ion qubit which has an auxiliary subspace which is detectable without disturbing the qubit states themselves in the case of a “dark” measurement.

In trapped $^{40}\text{Ca}^+$ ions, a high-fidelity detection may be implemented by driving fluorescence on the $4S_{1/2} \leftrightarrow 4P_{1/2} \leftrightarrow 3D_{3/2}$ manifold (Section 4.1). This detection event occurs when the ion is found to be in one of the aforementioned states, but if the ion is found to be in the metastable $3D_{5/2}$ subspace, it will not fluoresce and appear “dark.” Furthermore, the application of resonant lasers to

3.4. Partial collapse and uncollapse of the wavefunction




Figure 3.4: Partial collapse process, \mathcal{P}'_p , including a small population leak during transfer ϵ (see text) as illustrated by its deformation of the Bloch sphere. p is the strength of collapse. In this example, $\epsilon = 0.036$, and the values of p exponentially approach 1.




Figure 3.5: Uncollapse process, \mathcal{U}'_p , including a small population leak during transfer ϵ (see text) as illustrated by its deformation of the Bloch sphere. p is the strength of partial collapse. In this example, $\epsilon = 0.036$, and the values of p exponentially approach 1.

drive the fluorescence does not significantly perturb the states in the $3D_{5/2}$ subspace. Using states in the $3D_{5/2}$ subspace as a qubit is complicated by the fact that the magnetic dipole transitions between the Zeeman states all occur at the same frequency, a difficulty we overcome by the application of a quasi-resonant laser field to detune unwanted transitions. The partial collapse can be achieved by optical pumping out of the $3D_{5/2}$ level. However, this process is accompanied by a small transfer from $|1\rangle$ to $|0\rangle$, meaning that it is “leaky” in the sense of the previous subsection. This scheme is discussed in detail in Chapter 8.

4

Experimental Methods

We now turn to a description of the apparatus used to complete the experiments which make up the remainder of this work. We describe ion-trapping attempts in two different ion trap systems. The trap systems will be described separately in later chapters. We begin with the portions of the apparatus common to both. These include the systems to generate the laser beams required to photoionise ^{40}Ca and address transitions in $^{40}\text{Ca}^+$, the photon-counting systems, and experimental control systems.

4.1 Calcium atomic structure




Figure 4.1: Simplified ^{40}Ca level structure with transitions addressed by laser systems in this work.

Figure 4.1 is a simplified level diagram for neutral ^{40}Ca . We use a two-photon photoionisation (PI) process to create $^{40}\text{Ca}^+$ ions when we attempt to

load the trap. The frequency sensitivity of the $^1S_0 \leftrightarrow ^1P_1$ transition allows isotope selective loading [35] and the resonance enhancement increases the efficiency over a two-photon process using a single color of light.




Figure 4.2: $^{40}\text{Ca}^+$ low-lying level structure. Left: transitions addressed by laser systems in this work. Right: relative spontaneous decay rates indicated by the width of the connecting arrow. Decay from the D levels to the ground state are omitted. These decay rates are 6 orders of magnitude smaller than the smallest indicated transition and would have widths of roughly 3 \AA on this scale. Level separation is not to scale and greatly exaggerates the fine-structure splitting (a few THz). Color is qualitative only and differentiates the infra-red transitions from the ultra-violet.

Figure 4.2 shows low-lying energy levels for $^{40}\text{Ca}^+$, including transitions accessed with our laser systems. The $4S_{1/2} \leftrightarrow 4P_{1/2}$ transition is used for fluorescence detection and Doppler cooling. However, the natural decay of the $4P_{1/2}$ level to the $3D_{3/2}$ level requires the use of a repump laser to prevent accumulation of population in the latter. The state detection technique in which the ion is optically pumped to the metastable $3D_{5/2}$ level, removing it from the “bright” fluorescing manifold ($4S_{1/2} \leftrightarrow 4P_{1/2} \leftrightarrow 3D_{3/2}$) is often called “electron shelving,” so we denote the lasers which are used to optically pump in and out of the $3D_{5/2}$ level the “shelving” and “deshelving” lasers. Lastly, in our manipulations of the magnetic sub-states in the $3D_{5/2}$ level we require an intense laser field $\sim 100 \text{ GHz}$ detuned from resonance with the $4P_{3/2} \leftrightarrow 3D_{5/2}$ transition, which we denote the “light-shift laser.”

Table 4.1 summarises the laser systems in use. For brevity, here and henceforth, we omit electron configurations and refer to levels by their spectroscopic terms only. We shall also refer to laser systems by names which evoke their purpose (instead of wavelength). We hope this will improve readability for those not intimately familiar with $^{40}\text{Ca}^+$.

Name	Atom	Transition	f (THz)	λ (nm)	Polarization
Resonant PI	^{40}Ca	$^1\text{S}_0 \leftrightarrow ^1\text{P}_1$	709	423	uncontrolled
Nonresonant PI	^{40}Ca	$^1\text{P}_1 \leftrightarrow ^4\text{Ca}^+$	771	389	uncontrolled
Doppler	$^{40}\text{Ca}^+$	$\text{S}_{1/2} \leftrightarrow \text{P}_{1/2}$	755	397	$\sigma^+ + \sigma^-$
Repump	$^{40}\text{Ca}^+$	$\text{D}_{3/2} \leftrightarrow \text{P}_{1/2}$	346	866	$\sigma^+ + \sigma^-$
Shelving	$^{40}\text{Ca}^+$	$\text{S}_{1/2} \leftrightarrow \text{P}_{3/2}$	762	393	σ^-
Deshelving	$^{40}\text{Ca}^+$	$\text{D}_{5/2} \leftrightarrow \text{P}_{3/2}$	351	854	σ^-, π^\ddagger
Light-shift	$^{40}\text{Ca}^+$	$\text{D}_{5/2} \leftrightarrow \text{P}_{3/2}$	351 [†]	854	$\sigma^{-\ddagger}$

Table 4.1: Lasers used to manipulate ^{40}Ca and $^{40}\text{Ca}^+$. \ddagger indicates polarisations required with very high purity (see Section 7.2). † the light-shift laser is detuned ~ 100 GHz from resonance.

4.2 Stabilised diode lasers

All lasers in use are generated by diodes, and all except the nonresonant PI laser are stabilised to an external cavity consisting of the back face of the diode and a piezo-mounted grating in the Littrow configuration. This lowers the linewidth of the output laser beam, and improves tunability. Coarse tuning of grating-stabilised lasers is accomplished by changing the angle of the grating with respect to the diode, which alters the frequency of beams which are reflected back into the diode. For infra-red laser diodes this tuning can be 10 THz or more. Fine tuning is accomplished by changing the length of the external cavity via a piezo on which the grating is mounted. This alters the mode spectrum of the cavity. All of our grating-stabilised lasers are commercially available systems built by Toptica Photonics AG.

For the resonant PI laser, the stability provided by the external-grating cavity is sufficient for our needs. However, the lasers which address transitions in $^{40}\text{Ca}^+$ have more stringent requirements for stability and reproducibility of the tuning. For these systems we lock the laser outputs to ultra-low-drift etalons manufactured by the UK National Physical Laboratory (NPL). These etalons are manufactured out of materials with low thermal expansion, sealed inside high-vacuum chambers, and temperature controlled. They are also tunable via piezos used to mount one of the cavity mirrors.

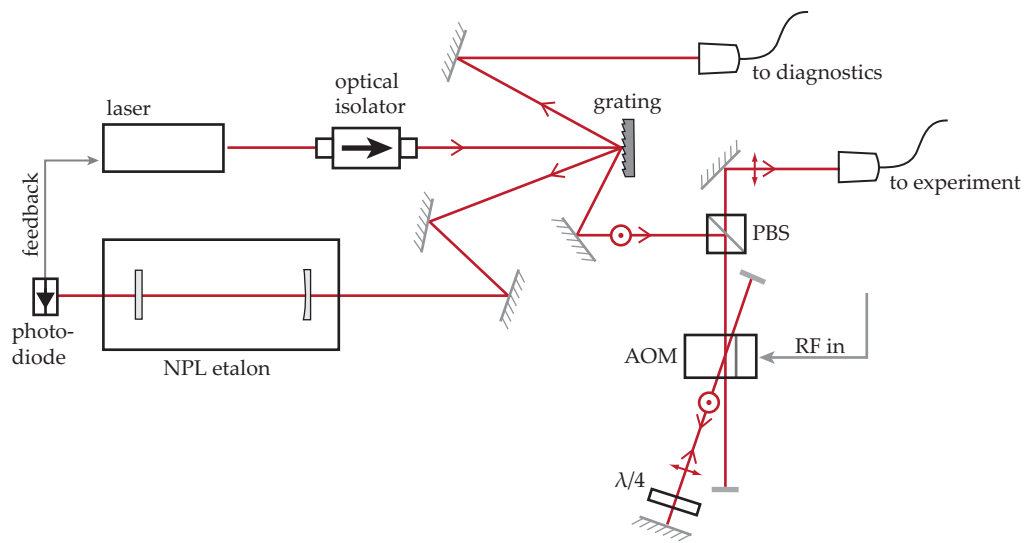


Figure 4.3: Typical laser system configuration depicted with simplified optics. See text for deviations.

In a typical configuration (Figure 4.3), the output beam from the stabilised diode laser is split into 3 beams with a diffraction grating: one beam used for locking to an etalon, one fed to a wavelength meter for diagnostics, and one being fed to the experiment. The grating also spectrally filters amplified spontaneous emission from the laser output. The beam used for locking is mode-matched to the etalon, and the transmission monitored via a photodiode. Locking is via the “side-of-fringe” configuration. The laser is tuned via the grating piezo to the side of one of the fringes in the etalon transmission spectrum, and a voltage set-point determined from the photodiode output at this location. The grating piezo is adjusted to keep the laser on the fringe via negative feedback (using a Toptica PID controller) on the difference between set-point and photodiode output. This locks the grating-stabilised laser to the etalon. Tuning is then accomplished via the etalon piezo, which shifts the etalon spectrum, relying on the servo control to scan the laser with it.

The beam used for diagnostics is coupled into a multimode fibre which terminates on a fibre-switcher. The fibre-switcher multiplexes the diagnostic beams from the different laser systems. Half the power is directed into a WS/7 Wavemeter built by HighFinesse GmbH, which gives a readout of the laser frequency. The wavemeter is typically accurate to within 1 GHz, and is used to coarsely tune the lasers when preparing for experiments, as well as various diagnostics on the systems. The other half of the diagnostic beam power is directed to an optical spectrum analyser coated for infra-red (repump, light-shift, and shelving beams). The spectra of the ultra-violet Doppler and shelving beams are obtained from another analyser, as detailed in their individual sections below.

The beams directed to the experiment are typically switched with acousto-

optic modulators (AOMs) in double-pass configurations: the first-order diffracted beam from the AOM is retro-reflected back through the AOM where it is isolated either via a polariser or picked-off by an edge mirror. The double-pass AOM used as a switch gives extinction ($P_{\text{off}}/P_{\text{on}}$) of typically $< 10^{-4}$, and can switch at frequencies up to several MHz. The output beam from this configuration is then coupled to a single-mode optical fibre which delivers it to the trap system.

The Repump laser operating at 346 THz (866 nm) is an example of a laser system which operates as described above. We shall now discuss relevant deviations of the individual laser systems from this “standard configuration.”

4.2.1 Photoionisation lasers

The lasers used to photoionise neutral ^{40}Ca have less stringent requirements for frequency stability. The nonresonant PI beam excites to a continuum of states, and thus does not need stabilisation via an external-grating cavity. It is driven by a ThorLabs LCS100 current/temperature controller, and is the only diode laser system not built by Toptica. It is also not coupled into the wavemeter. The resonant PI laser is grating-stabilised, but is not locked to an NPL etalon.

The two PI lasers are shared between experiments and a series of $\lambda/2$ waveplates and polarising beam splitter (PBS) cubes are used to steer the beams down different paths. The PI beams are overlapped onto a single mode fibre using a PBS cube. Instead of using an AOM as a switch, the PI beams are switched on and off together by a mechanical shutter situated in front of the fibre input coupler.

4.2.2 Doppler laser

The Doppler laser is arranged in the standard configuration except for a pick-off mirror which directs the (usually discarded) zeroth order mode from the first pass of the AOM down another beampath and into an optical spectrum analyser coated for ultra-violet light. When the beam is switched off at the experiment (AOM off), this path receives full power. When the beam is switched on at the experiment (AOM on) this path receives $\sim 30\%$ of the power.

4.2.3 Shelving laser

The shelving laser also has a pick-off to direct the zeroth order from the double-pass AOM to the spectrum analyser. In practice, we usually block one or other of these two beam paths with cards so that only one beam is incident on the analyser at a time.

In the shelving laser system, a series of beam splitters (instead of a grating) is used to create the experimental, locking, and diagnostic beams. A glass pick-off

(anti-reflection coated on one side) generates the locking beam, and a $\lambda/2$ waveplate followed by a PBS cube splits experimental and diagnostic beams. This allows easy alteration of the relative powers in these beams without misalignment of downstream components.

4.2.4 Deshelving lasers

For initial trapping attempts, the deshelving laser system was operated in the standard configuration, but for later $D_{5/2}$ spin experiments (Chapters 7 and 8), we required two independently switchable beams at this frequency. The main experimental beam was therefore further split by a $\lambda/2$ waveplate and PBS, and the two beams directed into two different AOMs in the double-pass configuration. These two beams were then coupled into appropriate fibres.

4.2.5 Light-shift laser

The light-shift laser system was primarily used to detune specific m -sublevel transitions out of resonance in the $D_{5/2}$ subspace, as detailed in Section 7.1.2. For a few experiments (e.g. Section 7.2.2), however, we needed a laser tuned to the $P_{3/2} \leftrightarrow D_{3/2}$ transition at 353 THz (851 nm). For these situations the light-shift laser system was retuned by ~ 2 THz to this resonance by slightly altering the angle the external cavity grating makes with the diode. This alters the beam path significantly enough to require realignment of all components in the system each time this frequency was changed.

4.3 Laser Control Unit

Laser switching is controlled by a PC-based Laser Control Unit (LCU), which allows precisely timed TTL-compatible* control pulses to be sent to the RF drivers for the AOMs [29]. It is timed by a 10 MHz clock, allowing pulses with a settable resolution of 100 ns. Specified TTL-transition times are < 50 ns. The LCU has 8 precision-timed outputs, which are hardware timed on the 10 MHz clock, and each pulse is followed by a $6 \mu\text{s}$ reprogramming time during which all precision-timed outputs are switched off. 8 additional outputs are available software timed: switched on and off during the $6 \mu\text{s}$ reprogramming period. The software timed outputs' switch-on and switch-off times bookend the precision-timed pulse, allowing devices which either need an extra μs or two to stabilise before use, or which need to be on before the application of the precision-timed output, to be used. An example of the latter is the application of the light-shift laser whose intensity must stabilise before the application of RF driving pulses (Section 7.4).

*transistor-transistor logic

Timed pulse sequences are entered into an experimental control program, written in Turbo PASCAL and running on an MS-DOS based PC. The experimental control program also interfaces with photon counting electronics, issues commands to RF synthesisers via RS-232, and controls fine-tuning of laser frequencies via analog connections to the piezo drivers which control the NPL etalons.

4.4 Photon counting

An imaging system allows the detection of ion fluorescence via a photo-multiplier tube (PMT) or charge-coupled device (CCD) camera. Figure 4.4 outlines the components of the imaging systems, which are nominally identical for the Lucent and Liverpool traps. A Nikon wide aperture compound lens system, situated near the main trap system view port, collects light from the trapping region and images it on an adjustable rectangular aperture, which allows us to select a region of interest. A beam splitter directs light to imaging paths to the CCD camera and PMT, in roughly the ratio of 7:1. The beam splitter is mounted on a flipper mount so it can be removed from the imaging path and all light directed to the PMT. A pair each of interference filters are typically inserted in front of the PMT and the CCD camera, transmitting light at 755 THz (397 nm)—the primary fluorescence signal. They are removed for experiments conducted prior to trapping which establish correct alignment and tuning of the resonant photoionisation laser beam at 709 THz (423 nm).

Figure 4.4 is an overview schematic of the system. Its main operating parameters are an $8\times$ optical magnification and an overall detection efficiency of $\eta = 1.6 \times 10^{-3}$ at 755 THz without the beam splitter and with interference filters in place. Design and characterisation are covered in detail in [41] and [27].

Photons are counted by an Electron Tubes Ltd. P25PC-12 PMT, which gives 10 ns “clicks” at TTL compatible voltage levels. A pulse stretcher is placed inline between the PMT and an Intel 8254 counter, stretching the pulses to 50 ns where they can be read by the counter without glitching. The counter is on board an expansion card in the MS-DOS based control computer.

The output of the pulse-stretcher is also connected to a time-to-amplitude converter (TAC), which we use to correlate photon arrival times to the phase of the trap-RF—a capability which is useful in micromotion compensation (Section 6.2.4). The PMT pulse triggers the “start” of a time accumulation on the converter, while a TTL-compatible square-wave synched with the trap RF triggers the “stop” on an edge. An analog-to-digital converter on the control computer reads the time between start and stop as a voltage level on the output of the time-to-amplitude converter [41]. This process can be repeated to build up histograms of photon arrival times.




Figure 4.4: Schematic drawing of imaging system showing approximate distances between elements. Figure is originally from [41] and we have made slight modifications to reflect changes since the system was designed.

5

Lucent Trap

The Lucent Technologies (formerly Bell Laboratories) Planar Ion Trap Rev. 3, hereafter the *Lucent trap*, was build in 2005 by Richard Slusher. We received the trap through a collaboration with what is now called the Intelligence Advanced Research Projects Activity (IARPA), formerly the Disruptive Technologies Office (DTO).

We designed and built a vacuum system to house the trap, and attempted to load Ca^+ ions. The vacuum system and chip mounting socket was a modified version of that which was successfully used to trap ions [27] in the Sandia National Laboratories built chip trap from the same (DTO) collaboration.

While we were unable to successfully load the Lucent trap, other collaborators at the University of Michigan [59] and the National Institute of Standards and Technology were. This chapter documents our attempts with the aim of highlighting the differences in our setup, and the areas of difficulty which might be addressed in ongoing [5] revisions to surface trap designs.

5.1 Apparatus

5.1.1 Design and fabrication

The Lucent trap is a micro-fabricated planar ion trap with aluminium-on-tungsten electrodes patterned on a silicon substrate [59]. The substrate is plated with $1\ \mu\text{m}$ of aluminium to act as a ground-plane, and the DC electrodes are atop a $0.4\ \mu\text{m}$ SiN insulating layer. It has 19 DC electrode segment pairs arranged around two long RF “rail” electrodes which are $20\ \mu\text{m}$ wide and separated by 150, 125, 100, and $75\ \mu\text{m}$ as they narrow in 4 sections along the axis of the trap (Figure 5.1). The RF rails are $\sim 10\ \mu\text{m}$ out of plane atop a layer of etched SiO_2 .

There are several versions of the trap, which differ in the geometry in the region between the RF rails. The first is the most basic, where the region between the rails is simply another DC electrode, called the centre control electrode, hereafter CCE. It is used to adjust the y -component of the DC electric field at the site

of the ion. This design has advantages in the creation of appropriate trapping potentials, as will be discussed later, but requires loading via an atomic vapour directed from the front of the trap chip. If vapour densities are not carefully controlled this can lead to plating of the trap electrodes. A second design includes a small loading slot, $80\ \mu\text{m}$ wide by $800\ \mu\text{m}$ long, cut straight through the CCE and through the back of the chip ($\sim 300\ \mu\text{m}$ thick). The slot is located in the region where the RF rails are widest and can be used to allow loading via an atomic vapour originating behind the trap chip.




Figure 5.1: Lucent trap micrograph [54].

The third design, used in the trap received by the Oxford group includes a “fully slotted” chip, where a laser access slot is cut through the CCE along the full length of the trapping region. This allows us to direct lasers to the trapping region from behind the chip in addition to skimming the surface. The disadvantage of this design is that it removes most of the CCE, leaving only narrow strips on either side of the access slot. In addition, the fabrication process by which the slot is cut suffers from low accuracy in the alignment of the slot with respect to the electrodes. Our slot is about $10\ \mu\text{m}$ off-axis in the x -direction. This broken symmetry means applying a voltage to the CCE to correct the y -component of the field introduces an unwanted x -component which must be compensated with a bias across x -opposite DC electrodes.

5.1.2 Electrical properties and faults

Prior to inserting the Lucent trap into the vacuum system, several electrical tests were conducted. An in-air electrical testing rig was built by wiring up a PCB-mount socket for the chip trap. The empty test socket showed $\sim 40\ \text{pF}$ of capacitance between any two pins in the socket. After transferring the Lucent trap to the test socket, the capacitance to ground was measured. We measured $53\ \text{pF}$ for the RF rails, $162\ \text{pF}$ for the CCE, and $\sim 100\ \text{pF}$ for each DC electrode. Resistance between all combinations of pins was measured with a handheld ohmmeter. In all cases the resistance was $> 1\ \text{G}\Omega$. We then proceeded to apply $\pm 1\ \text{V}$ in an alternating configuration to the DC electrodes, while the RF and CCE were left




Figure 5.2: Lucent trap schematic along the chip plane. To scale for the second trapping region (DC electrode pairs 6 and 7). p-doped Si substrate is $\sim 300 \mu\text{m}$ thick (i.e. only about 1/15th is shown). Redrawn from [59]

unconnected. The trap was left in this configuration for 8 days. The measured resistances were all still $> 1 \text{ G}\Omega$.

We then connected up a single power supply set to $+2 \text{ V}$ in series with a picoammeter and a $1 \text{ M}\Omega$ resistor. The power supply was connected to even numbered DC electrode pairs, with odd numbered pairs held at ground. In this configuration the picoammeter read 0.5 nA current to ground. After 7 days of no change in current reading, the supply was switched to odd numbered electrode pairs. This drew significantly higher current, so we disconnected the supply and measured the resistance. Electrode T5 showed $56 \text{ k}\Omega$ to ground.

As a final in-air test, the power supply/picoammeter setup was connected to electrode T19. We recorded the current drawn as the voltage was varied. The results are shown in Figure 5.3. Although the results clearly show deviation from Ohmic behavior, the differential resistance is several $\text{G}\Omega$ in the test range.

5.1.3 Vacuum system

The Lucent trap is mounted in a Kimball Physics 4.5 in. Magdeburg hemisphere. The chip socket is made of laser-machined Vespel (a high performance plastic which is UHV compatible). Two Vespel plates sandwich a set of pin-receptacles




Figure 5.3: Measured current from electrode T19 to ground.

arrayed to accept the Kyocera PGA-100 chip carrier (which holds the Lucent trap) by press fit.

The pumping system connects to the hemisphere via a 1.33 in. flange. A 20 L/s ion pump provides primary evacuation when the system is at UHV pressures and a chemical getter improves pumping of hydrogen and argon.

Electrical wiring

Although the Lucent trap has 41 independent control electrodes (2×19 DC segments, CCE, RF rails, and ground), we used a 25-pin DC feed through and only connected 12 pairs of DC electrodes. An error in wiring rendered only 24 of the 25 conductors on the DC feed through operable, which were used to connect the 12 pairs of DC electrodes. The final DC electrode, the CCE, was connected via a spare pin on the Ca oven feed through. The RF rails and ground were connected via the 2-conductor RF feed through.

The connected electrode pairs were 2–13 (see Figure 5.1 for numbering). Initially, we planned to leave all unused electrodes unconnected, but before inserting the Lucent trap into the vacuum system, we decided to explicitly ground electrode pair 1. An unconnected conductor in the vicinity of the trapping region behaves unpredictably since it can collect charge liberated elsewhere. Electrode pair 1 was grounded since it was the closest to the region we would attempt to load in, and because the pins on the reverse side of the chip carrier were adjacent to ground. I connected the pins by wrapping gold wire around the base and securing it with a dot of vacuum-compatible silver paint (Figure 5.4).

Calcium oven

One oven was used, filled with calcium metal in its natural isotope abundance. The Lucent trap has a very small receiving angle from the back of the chip, only




Figure 5.4: Photograph of reverse-side of Lucent trap chip carrier showing the wiring used to explicitly ground electrode pair 1.

a few degrees off the y -axis in the x - y plane. We therefore placed the oven directly behind the chip carrier on straight rods from the oven feed through. Since exact placement is difficult and the margin for error very tight, we made a test of the oven in situ by inserting a glass slide between the oven aperture and the chip-carrier, evacuating the vacuum system, and test-firing the oven until a visible calcium coating appeared on the slide. The location of the calcium spot confirmed correct oven placement, after which the slide was removed.

5.1.4 Optical setup

A large (4.5 in.) front viewport provides access for imaging and photon collection optics. Lasers were delivered via two 2.5 in viewports each at 45° to the trap plane. The trap was mounted such that the viewport axes were along the $y = \pm z, x = 0$ lines (given in trap coordinates).

The PI laser beams are overlapped onto a single-mode fibre. The fibre output is collimated, then steered into the vacuum system with dielectric mirrors, passing through a final focussing lens which gives a spot size ($2w_0$) roughly the width of the laser access slot.

The Doppler and repump cooling lasers are delivered to the trap table by

single-mode fibres. Fibre outputs are collimated and steered to co-propagate onto a broadband dielectric mirror which serves as a final adjustment for spot placement and allows the beams to be moved together in the trapping region. Since we collect fluorescence at the wavelength of the Doppler laser (755 THz, 397 nm), the final focussing lens was optimized to place the Doppler laser beam waist at the chip trap. This reduces Doppler laser beam scatter at the expense of increased scatter from the repump laser beam (346 THz, 866 nm) off the chip substrate and electrodes; this is readily filtered in the observation system.


5.1.5 Generation of DC trapping potentials

A model was developed of the electrode geometry which was then fed into *Charged Particle Optics* to obtain numerical approximations to the electric potential due to each electrode of interest. The Lucent trap is too complicated to simulate in its entirety due to hard-coded limits on the number of segments handled by the software. Only a limited set of electrodes around the chosen trapping region were modelled. We chose to attempt trapping in the second widest trapping region. We could not trap in the centre of the widest region because electrode pair 1 was grounded. We were also advised to stay away from angled regions where the RF rails go from one width to another, since these give a distorted RF field [60]. The second trapping region contains 4 electrode pairs, numbered 5–8 (Figure 5.1). The centre of the region is between electrode pairs 6 & 7.

The chip substrate is p-doped silicon, which has a resistivity of $0.018 \Omega\text{cm}$, and is connected to the aluminium ground plane. Via a simple model, we argue that for the purpose of this simulation, the substrate should be treated as grounded conductor. Consider a parallel plate capacitor where the plates sandwich a block of doped silicon, with thickness d , and plate area A . If this capacitor is charged up and then disconnected, it will discharge through the doped silicon (which has finite resistance) with a time constant RC , where R is the resistance of the slab of silicon and C is the capacitance. From the definition of resistivity ρ , we have $R = \rho d/A$, and $C = A\epsilon/d$ (where ϵ is the permittivity of silicon, $\sim 11\epsilon_0$). Thus $RC = \rho\epsilon \approx 2 \times 10^{-14}$ s. Since the period of one RF cycle at the drive frequency, 25.8MHz, is much longer than this, charge in the silicon has plenty of time to move and null the potential in response to the RF drive field. The walls of the laser-access slot are therefore modelled as conducting surfaces in the simulation.

Electrodes at the angled RF area were included in the simulation model. The dimensions and location of the CCE are measured from a micrograph of the trap [54]. The x -widths of relevant structures are summarised in Table 5.1.

With the laser access slot cut off centre, as mentioned in Section 5.1.1, there are no geometric symmetries (Section 2.4.3) which we can take advantage of to

Figure 5.5: Closeup micrograph of 2nd trapping region.

Trapping Region	Top CCE	Slot	Bottom CCE
1 st	5	90	30
2 nd	5	70	25
3 rd	5	55	20
4 th	0	35	20

Table 5.1: x -widths of laser-slot structures, as measured from micrograph. All dimensions in μm .

create our desired potential of the form

$$V(x, y, z) = Q\left(z^2 - \frac{1}{2}x^2 - \frac{1}{2}y^2\right) \quad (5.1)$$

where x, y are measured from the RF zero point, and z from the edge separating DC electrode pairs 6 and 7.

Using the method of linear constraints with quadratic polynomials, to achieve the potential exactly we must satisfy a total of 5 constraints*. We can simplify the generation of potentials considerably by making some compromises. We note that the symmetry along the z -axis is only broken in a location comparatively far from the centre of the trapping region. We guess, and can confirm via simulation, that if the potentials are applied symmetrically about the centre in the z -direction then the linear terms in the potential along z will approximately cancel. Any residual will have the effect of moving the trap centre along the z -

*linear and quadratic coefficients in x, y, z give 6, but satisfying Laplace's equation reduces this by one.

	V	Electrode #	Feedthrough Pins
CCE	-6.2237		
Centre Bottom	-2.6940	6, 7	9, 10
Centre Top	-7.0917	6, 7	4, 17
Endcaps	-0.2427	5, 8	13, 16, 21, 22

Table 5.2: Computed potentials and feedthrough pins for Trapping Solution 1

	V	Electrode #	Feedthrough Pins
CCE	0		
Centre	-1.4166	6, 7	4, 9, 10, 17
Endcaps	4.0211	5, 8	13, 16, 21, 22

Table 5.3: Computed potentials and feedthrough pins for Trapping Solution 2

axis, but only by a short distance, which well within the tolerance of laser-beam and collection optics placement.

This leaves 4 constraints, namely that the linear terms in the x - and y -directions cancel, and that the quadratic terms in the z - and x -directions sum to Q and $-Q/2$ respectively. One solution was computed based on these constraints, with electrodes organised as follows. The 4 electrodes in pairs 5 & 8 formed the “endcaps.” The 4 inner electrodes (pairs 6 & 7) were grouped into two sets, one on either side of the trapping axis. These formed the “centre bottom” and “centre top” sets. The CCE became the fourth degree of freedom, and all other electrodes were set to 0V. This configuration is known as Trapping Solution 1. The computed potentials for $\omega_z = 2\pi \times 1$ MHz axial trapping frequency are given in Table 5.2.

Another set of electrode potentials, called Trapping Solution 2, was calculated. This solution does not separate the centre electrode pairs (6 & 7) into “top” and “bottom” sets, which gives a simpler electrode wiring setup. It also reduces the linear constraints to 2: that the linear terms in y sum to zero (i.e. DC trap centre is at correct height above the plane) and that the quadratic terms in z sum to give a 2 MHz axial trapping frequency. This leaves one free degree of freedom, which was used to set the CCE to 0V. During some of the trapping attempts the CCE developed a finite resistance to ground, and this solution allowed us to continue trapping attempts without having to pass current through the CCE and risk exacerbating the fault. Since the split CCE is the primary asymmetry in the x -direction, setting its voltage to 0 also helps to limit excess field that causes micromotion. The disadvantage of this solution is that it does not constrain the quadratic terms in x or y . As it happens, the solution is strongly expelling in the x direction, so the RF must be driven at a higher voltage to maintain a stable trapping potential. The potentials are listed in Table 5.3.

5.1.6 Generation of RF trapping potential

In Chapter 2 we derived that the secular frequency of a trapped ion's motion is proportional to the RF voltage, and inversely proportional to the RF drive frequency. We may therefore conclude that, for a given ion species, the quantity $V_{\text{rf}}/\Omega_{\text{rf}}\omega_r$ is a constant which depends only on the trap geometry. Unlike other geometric factors, however, it includes the trap scale. In Hz frequency units this quantity for the Lucent trap is about $0.4 \frac{\text{V}}{\text{MHz}\cdot\text{MHz}}$; i.e. 0.4 V per MHz drive frequency per MHz secular frequency.

Our design goal was 2 MHz secular frequency, which means we could choose an RF drive frequency in the 20-30 MHz range without making the required voltages too difficult to achieve from a practical standpoint. Other groups had previously driven the trap at over 300 V, so we were also well within the design limits of the chip's breakdown characteristics.

A helical resonator [36] was used to impedance match and step-up the voltage from the RF drive to the trap system. Details of its construction can be found in [22]. When connected to the Lucent trap system, the resonance frequency was 26 MHz and the Q was 20. Since the voltage step-up in a resonant circuit is approximately Q, no amplifier was needed.

5.1.7 Magnetic field

Loading requires only coarse control over the magnetic field. To prevent optical pumping into the $D_{3/2}$ level a finite magnetic field needs to be applied. At zero magnetic field the polarisation of the repump laser field will define a set of bright states which the field couples to $P_{1/2}$ and a set of dark states which do not couple. Driving the usual cooling manifold results in optical pumping to these dark states. The application of a magnetic field with a different quantisation axis ensures that the dark states will not be stationary—the system will quickly precess to a state which can be repumped.

A small field coil of 104 turns was placed over the large front viewport. Approximate calibration was accomplished with a Hall probe, and current passed through the coil to give ~ 1 G field.

5.2 Neutral calcium fluorescence

We observed fluorescence in the trapping region with the calcium oven on and the resonant PI laser focussed through the laser access slot. Using the computer DAC fed into the laser's grating stabilisation, we could scan over ~ 2 GHz, using the wavemeter as a reference. We heated the calcium oven by passing 5.7 A of current through it, and detected fluorescence using the PMT. We then switched off the oven current and waited a few minutes before taking a nominally identi-

cal ‘background’ scan. Figure 5.6 shows the detected signal minus background versus laser frequency.




Figure 5.6: Neutral Ca fluorescence, background subtracted. f_0 is the nominal fluorescence peak.

The photoionisation laser interacts with the atoms at a 45° angle to their flight from the oven aperture to the trapping region. The spread of velocities in the atomic beam gives rise to a Doppler-broadened fluorescence profile with width of at least 1 GHz. The full extent of the profile was beyond the scan range of the laser.

After several failed loading attempts, we wanted more conclusive evidence that calcium atoms were reaching the trapping region. The observation of fluorescence, while strong evidence, does not directly demonstrate calcium atoms in the trapping region, since atoms behind the chip trap could conceivably scatter photons through the access slot even if they themselves were blocked from passing through the slot.

We therefore collected fluorescence on the CCD camera. These images of the trapping area (Figure 5.7), taken from the front side of the chip trap, show a visible ‘‘cloud’’ of atomic vapour which extends from the laser access slot outward. This is very convincing evidence that calcium atoms reach the trapping area.

5.3 Loading attempts

5.3.1 The Dust

After sealing the Lucent trap in the vacuum system and baking, we observed a small piece of foreign material, about $20\ \mu\text{m}$ across, stuck to one of the RF rails between electrode pairs 7 & 8.

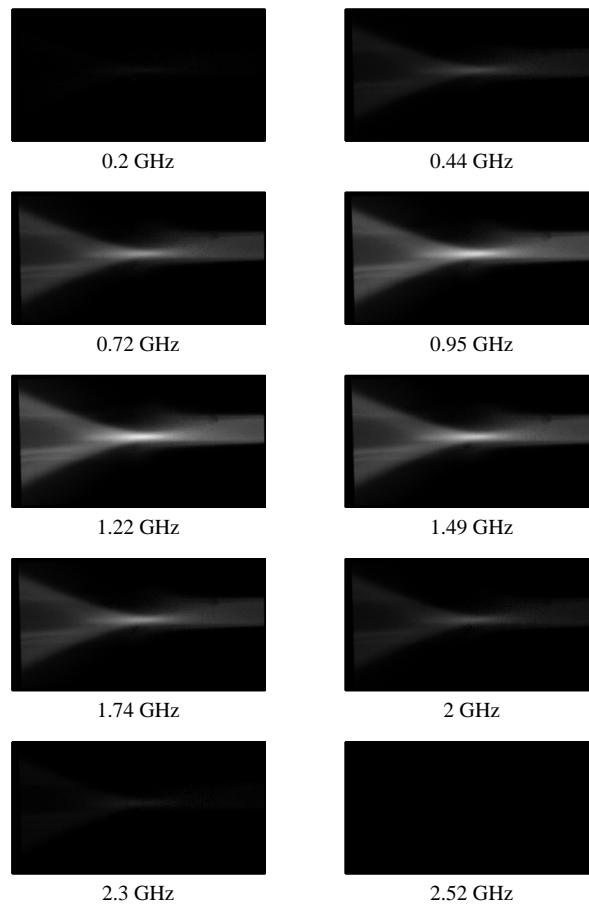


Figure 5.7: Neutral Ca fluorescence images. Detunings are relative to the nominal 'Doppler-free' resonance.

Fearing that it would distort the RF fields and prevent us from forming a stable trapping region, we attempted to knock it loose. We reasoned that if the material was insulating, we might be able to charge it up using a UV laser. Using the imaging camera, we aimed a small spot of $\sim 100 \mu\text{W}$ of 397nm light from the Doppler laser at the material. We then applied a DC potential of +35 V to the RF rail while keeping the DC electrodes and CCE grounded. There was no change in the appearance of the material as viewed on the camera. We proceeded to increase the voltage up to +110 V in 10 V increments. We also struck the vacuum can with the handle of a screwdriver, hoping that the vibrations would help shake the speck loose. There again was no change.

We reversed the polarity of the applied potential so the ground-plane of the

chip was at high voltage while the RF electrodes were held at ground. The applied voltage reached +8 V before starting to draw current. We disconnected the potentials and measured the resistance to ground from the pins on the DC feedthrough. Pin 16 (electrode T9) showed $240\ \Omega$ to ground. We removed the connection to Pin 16 and reapplied the DC potential, this time to 20 V before drawing current. Disconnecting, we again measured resistances to ground and found Pin 5 (electrode B3) showed $169\ \Omega$. At this point, after developing finite resistances on two electrodes, we abandoned our attempt to remove the foreign material and proceeded to attempt trapping in spite of its presence.

5.3.2 Background scatter

The narrow loading/laser access slot had the advantage of making optical alignment very simple with the help of the imaging CCD camera, since laser scatter on from the sides of the slot was readily imaged. However, it has proved a major obstacle in obtaining low background PMT counts. The ideal limit of a laser being cleanly threaded through the slot is only possible when the spot size at the beam waist, w_0 , is much smaller than the slot. Based on the beam width at the final lens, a spot size of $w_0 \approx 15\ \mu\text{m}$ would be expected at the trap if the beam were collimated and Gaussian at the lens. The slot is $70\ \mu\text{m}$ wide in the trapping region, so it should be possible to direct a Gaussian beam cleanly through.

With $30\ \mu\text{W}$ in the Doppler laser beam, no filters, no pinhole, and the imaging beam splitter in, there were 320 kHz counts of background scatter. With the beam splitter out this increased to 2270 kHz.

After replacing lenses, adjusting the vertical alignment, and closing the rectangular aperture we were able to reduce this to 150 kHz (for $36\ \mu\text{W}$ of Doppler laser beam power and beam splitter out). This suggests that either the laser spot size was not $15\ \mu\text{m}$, or the beam is not approximately Gaussian, or both. The trap view ports are glass, and we have observed both absorption and forward scatter from ultra-violet lasers passing through. Their flatness is specified as $\lambda/2$, so some distortion of the phase-front is to be expected as the Doppler laser is transmitted.

For comparison, with similar laser power there are ~ 1 kHz background counts in the Liverpool trap. With our collection efficiency, a cold ion gives ~ 30 kHz counts of fluorescence, but for a single hot ion the count rate is much lower. Therefore the background noise presented a major difficulty in detecting a successful load in the Lucent trap.

5.3.3 Procedure

We attempted trapping using both of the potential solutions described in Section 5.1.5, alternately scanning the Doppler cooling laser and the resonant PI laser in

frequency.

In order to better observe ion fluorescence above background counts, the counting software has a background-subtraction mode. In alternate counting bins, the repump laser beam is switched on and off, and the indicated count level is the difference between successive bins. Bins were typically 50–100 ms in duration. Fluorescence from a trapped ion drops to essentially zero in the bins when the repump is off, so this bin represents the background scatter from the Doppler laser. Both background subtraction and (normal) continuous fluorescence modes were used extensively in our loading attempts. We observed no signs of trapped ions.

Since the trap depth predicted by our models was relatively low, we attempted to maximise the number of ions in the trapping region with low velocities. By tuning the resonant PI laser near the nominal Doppler-free resonance, we preferentially ionise the slower-moving atoms, but at a cost of ionising lower total number of atoms.

Increasing the oven temperature will, of course, increase the number of atoms in the atomic beam, but it does not necessarily increase the number of low-speed atoms. The speed-distribution of the *volume* number density of an atomic beam emanating from an oven is

$$N(v) = \kappa N_0 \left(\frac{kT}{m} \right)^{-\frac{3}{2}} v^2 e^{-mv^2/2kT} dv \quad (5.2)$$

with N_0 the number density inside the oven, and κ a geometric factor which does not depend on temperature, T , or speed, v . Using the relation $p = N_0 kT$, we write

$$N(v) = \frac{p(T)}{kT} \times \kappa \left(\frac{kT}{m} \right)^{-\frac{3}{2}} v^2 e^{-mv^2/2kT} dv \quad (5.3)$$

Holding v constant at some small value such that $\frac{1}{2}mv^2 \ll kT$, we obtain the temperature dependence for low speeds:

$$N(v) \propto p(T) \times T^{-\frac{5}{2}} \quad (5.4)$$

$p(T)$ is experimentally known in the temperature range we normally operate the ovens (~ 700 K), and has a much stronger dependence on T than $T^{-5/2}$ (approximately T^{14} in that temperature range), so the overall number of slow atoms does increase.

When trapping failed at low oven currents, for our final attempts, we decided to increase number density in this manner by increasing the oven current. Neutral fluorescence was visible by naked eye between the chip trap and oven. After an afternoon of running the oven at this high temperature the backside of the chip and socket were visibly coated with calcium.

Electrical testing revealed multiple pin pairs with resistances $< 1000 \Omega$ between them, indicating conductance via the calcium coating. Attempts to load the trap were abandoned.

5.4 Comments

Failure to load in this configuration of the Lucent trap highlights some design considerations for future generations of segmented microtraps. The thinness of the $\sim 0.4 \mu\text{m}$ SiN insulating layer between the DC control electrodes and ground-plane, presumably combined with insufficiently cautious handling of the chip-trap in our laboratory, resulted in electrical breakdown on some electrodes. When testing a chip with many trapping regions, this is an inconvenience, but different fabrication and/or better handling will be required to successfully load a fully-functional trap for scalable quantum computing. While there has been some modest success with this trap design elsewhere [33], new trap designs and fabrication methods (e.g. [4]) are being explored.

The high background count rates also presented a major practical difficulty in our testing of the trap. The absolute size of the slot as well as the aspect ratio made it more difficult than other micro-scale trap systems [27] to achieve illumination of the trapping region without introducing high scattering rates. A redesign of the laser beam optics, including vacuum windows made from fused-silica (for better transparency to ultra-violet light) and better flatness might enable us to achieve a beam of sufficient quality to be passed cleanly through such an access slot. While it is attractive to be able to generate light fields with wave vectors out of the chip-plane, our subsequent experience with another surface trap [5] suggests that this is not essential. Delivering light fields and collecting ion fluorescence via fibres embedded into the chip substrate have also been proposed [31] [18].

6

Liverpool Trap




Figure 6.1: Photo of Liverpool trap front with schematic drawing of a side view. Shaded areas in the drawing are MACOR insulating layers.

6.1 Apparatus

6.1.1 Design and fabrication

The *Liverpool trap* was designed as a “mesoscopic”-scale segmented trap optimised for the fast separation of ions [25]. Splitting a trapping region (potential minimum) in two requires the application of an octupole potential [24] and the

Liverpool trap is designed to apply large octupole potentials without electrical breakdown in vacuo.

It was built at the University of Liverpool by S. Taylor and B. Brkic using miniature conventional machining techniques. The insulating spacers are MACOR; the conducting “pin” DC-electrodes and “plate” RF-electrodes are non-magnetic stainless steel.

Figure 6.1 shows the trapping region. The trap consists of 7 pairs of DC electrodes $560\ \mu\text{m}$ in diameter and spaced (centre-to-centre) along the trap z -axis by $730\ \mu\text{m}$, with a pair of RF electrodes in front and behind the DC pins. Unlike in a conventional four rods or blade trap, all 4 RF electrodes are driven in phase. The DC electrodes serve as an RF ground, creating the RF quadrupole potential required for radial confinement (Figure 6.2). The distance from the trapping region to the nearest DC electrode is $700\ \mu\text{m}$. 4 additional DC electrodes are mounted out of the plane of the DC “pin” electrodes and away from the trapping region. These are the compensation electrodes, used to null any stray fields that misalign DC potentials with the RF trapping pseudopotential (see Figure 6.1).




Figure 6.2: Liverpool trap view with RF-quadrupole potential

During final assembly or bakeout of the vacuum system, one of the RF electrodes was shifted to an angled position. This breaks some of the symmetries in geometry, requiring a more detailed computational model, but does not prevent trapping. Simulations of the new geometry show that it distorts the RF potential such that the null does not coincide with the axis defined by the DC electrodes. When trapping in a single region this presents very little additional difficulty, since the DC potential minimum can be aligned with the RF null via compen-

sation fields. This is demonstrated experimentally in later sections. However, when trapping in multiple regions it may not be possible to fully compensate two or more trapping regions simultaneously [27].

6.1.2 Vacuum system

The Liverpool trap is mounted in a Kimball Physics 4.5 in. Magdeburg hemisphere. Vacuum-compatible, conducting silver paint was used to connect DC electrodes to wires which terminated on a 32-conductor vacuum feed through. The RF electrodes were attached to strips of copper plating, secured with a machine screw, and then attached to a 2-conductor vacuum feed through with barrel connectors.




Figure 6.3: Liverpool trap system with indicated optical access and beam polarisations.

Two stainless steel tubes 2 mm in diameter were filled with calcium metal and crimped closed at each end. Each had a small hole ~ 0.5 mm in diameter drilled in the side. The tubes were spot-welded to stainless steel support rods to place them on either side of the trapping region, and the support rods connected to a third vacuum feed through to allow current to be driven through the tubes. Once under vacuum the tubes can be heated to ~ 600 K and calcium vapour effuses from the side hole as an atomic beam.

The hemisphere has a large 4.5 in. CF (ConFlat) flange and two smaller 2.75 in. CF flanges which were each fitted with view ports to allow laser access and collection of ion fluorescence. Figure 6.3 shows the viewports and beam configuration. Photoionisation beams were sent through a small view port fixed to a

1.33 in. CF flange on the back of the hemisphere at an elevation to the horizontal of about 33° .

A 1.33 in. CF flange connects the hemisphere to pumping system. A 20 L/s ion pump provides primary evacuation when the system is at ultra high vacuum (UHV) pressures. A chemical getter improves pumping of hydrogen and argon. The system is also fitted with an ion gauge to monitor the pressure.

The assembled vacuum components, not including the trap system, feed throughs or view ports, were evacuated with a turbomolecular pump and baked to 350°C ; then the system was fully assembled, evacuated and baked to 200°C for several weeks. The system was isolated from the pumping/baking rig with a UHV valve and then secured to the optics table for ion loading.

6.1.3 Generation of RF trapping potential

The required RF drive voltage for the Liverpool trap is about $50 \frac{\text{V}}{\text{MHz}\cdot\text{MHz}}$; i.e. 50 V per MHz drive frequency per MHz secular frequency.

Our design goal was 2 MHz secular frequency, so we chose 6 MHz as our RF drive frequency, since we hoped to be able to achieve 600 V amplitude on the RF electrodes. This is within the spec of our vacuum feedthrough, and seemed achievable with a few watts from an RF amplifier and a resonator circuit (e.g. 5 W at $50\ \Omega$ is 15 V, meaning the resonator would have to step up the voltage by $40\times$). Assuming an axial secular frequency $\omega_z = 2\pi \times 1\ \text{MHz}$ gives a Mathieu parameter $q = 0.78$, which is stable.

An air-core wire-wound inductor was constructed to resonate with the capacitance of the trap, cabling and feed through. The inductor was constructed with exposed wiring such that the inductance could be adjusted in situ by connecting the output wire to the coil at different locations. The input side of the coil was impedance matched to the $50\ \Omega$ drive system by connecting a length of coaxial cable, approximately one quarter the wavelength, to the input with a T junction. The cable was left unterminated, and its length adjusted by adding and removing sections until reflection (measured using one-way couplers) was minimised. The final, loaded quality factor was measured to be 32.




Figure 6.4: Resonator with impedance match

6.1.4 Generation of DC trapping potentials

The Liverpool trap was modelled in CPO to calculate the potentials given by the DC electrodes. This model included the displaced RF electrode, as measured by high-resolution photos of the trap before it was sealed in the vacuum system. The simulations indicated that the primary effect of the displaced electrode was to shift the RF potential null, but it did not affect the shape of the potentials in any significant way.




Figure 6.5: Liverpool trap DC electrode grouping.

The configuration space of the applied voltages was simplified by tying some electrodes together into control groups with the same voltage. This simplifies electrical connection and allows us to find a unique solution to the linear constraints problem. The DC control electrodes were organised into six groups, illustrated in Figure 6.5, with the DC compensation electrodes tied together in a seventh group. The basic trapping potential is created by raising the potential of groups 1, 3, 4, and 6 to the same (positive) voltage. An x -compensation field is generated by applying a differential voltage to groups on either side of the trap axis ($\{1, 2, 3\}$ and $\{4, 5, 6\}$) in linear combination with the trapping potential. This configuration compensates stray fields with very little additional distortion to the trapping potential. Finally, a voltage is applied to the compensation plates to create a field in the y -direction.

The voltages applied to the DC control electrodes are generated by a computer system operating a digital-to-analog converter (DAC) with multiple outputs. These outputs are limited to ± 10 V, so not all combinations in the configuration space will be available. We typically apply a 5 V trapping potential which gives $\omega_z \approx 2\pi \times 330$ kHz secular frequency along the trap z -axis. Voltages required to null the stray field in the y -direction are often outside the ± 10 V range, so the compensation electrodes are connected to a bench-top power supply with a range ± 30 V.

6.1.5 Magnetic field

Static magnetic fields are generated by three coils wound with magnet-wire and fed by low-noise current supplies. The coils are situated around the vacuum system in nominally orthogonal orientations such that the field may be set in any direction (up to, of course, the limit of overheating the coils by driving too much current through them).




Figure 6.6: Liverpool trap system with magnetic field coil placement.

The ratio for each coil of magnetic field at the ion site per unit current can be calculated approximately. When a thin coil of N turns and radius R is aligned such that its axis passes through a point a distance x away, the magnetic field is

$$B = \frac{\mu_0 N R^2 I}{2(R^2 + x^2)^{3/2}} \quad (6.1)$$

with I the applied current. The approximation that the coils are thin is not well satisfied in our setup, and the vacuum chamber can provide some shielding, so such calculations primarily drive design choices and are not used to calibrate the final setup. As we will discuss in Section 6.3, once we have loaded an ion, fluorescence spectra may be used to accurately deduce the actual field at the ion's location.

6.2 Loading

6.2.1 Beam alignment

Correctly overlapping the photoionisation and cooling laser beams in the trapping region is critical for successful loading. The four lasers necessary for loading must be aligned to all pass through a point in empty space, $700 \mu\text{m}$ away

from the nearest electrode, with tolerances of order the beam radii: perhaps 30-60 μm .

For initial observation of neutral Ca fluorescence, only the resonant PI laser is needed. Alignment along the trap z -axis was accomplished by imaging the centre DC electrode with the CCD camera, then illuminating the pin with the beam. The beam was placed half-way between the pins by adjusting the alignment screws on the final mirror mount until the beam clipped the top or bottom pin (observable by examining the laser spot after it exited the trap system), then setting the screw half-way between.

The imaging system lens is mounted on a 3-axis micrometre stage, allowing straightforward translation to the trapping region after imaging electrodes to calibrate the micrometre screw readings. With the imaging system in place, fine adjustments to the resonant PI laser position were made by imaging the beam while it interacted with neutral calcium vapour in the trapping region. The non-resonant PI laser beam is coupled to the same single mode optical fibre as the resonant PI laser beam, so the fibre output modes co-propagate so long as care is taken to ensure that the beams pass through the optical axes of components in the beam path (to minimise chromatic aberration).

The two cooling laser beams, Doppler and repump, were coarsely aligned to co-propagate through the trapping region by examining the beam spots before and after they passed through the vacuum system. A glass slide was secured in a “flipper” mount to allow it to be reproducibly flipped in and out of the beam path. It was flipped out for normal running, but flipped in for alignment: the slide reflected beams into an additional CCD placed such that the distance from slide to camera was approximately identical to the distance from slide to trapping region. Thus the CCD image was a direct analogue to position of the beam spots in the trapping region. By imaging the centre DC electrode with the imaging system camera, then shining the laser beams onto the electrodes, positions on the alignment CCD image could be calibrated, and the cooling beams aligned to the trapping region.

6.2.2 Neutral calcium fluorescence

The first observation of neutral calcium fluorescence involves some uncertainty in the precise frequency of the resonant PI laser and the current to pass through the calcium oven. Our method is to slowly scan the resonant PI over several GHz in the vicinity of the transition*, as measured by the wavemeter, and to slowly increase oven current until fluorescence is observed.

*the natural linewidth of the transition is 35 MHz, and we expect additional broadening due to transit-time of atoms across the laser beam, the uncollimated atomic beam, saturation effects, and misalignment of the crossed laser and atomic beams [35].

The interference filters were removed from the imaging system, and we monitored the photon count rate while scanning the resonant PI laser frequency. Neutral calcium fluorescence is distinguished from background scatter by its frequency sensitivity.

6.2.3 Ion trapping

Even with the beam alignments discussed in Section 6.2.1, some additional searching is required to successfully trap. Stray electric fields can distort the DC potential to move the potential minimum outside the laser spots, or induce severe micromotion making stable trapping difficult. Using the alignment CCD as a guide, we move the cooling laser beams to explore a small region around the calculated trapping centre, and adjust the compensation potentials over a wide range (± 60 V).

There is often some uncertainty in the frequency of the cooling lasers, due to finite accuracy in the wavemeter system, so these were also scanned when loading for the first time. The low-drift etalons are typically stable to < 10 MHz per day, so after a successful load, setting the appropriate frequency for subsequent attempts consists of locking to the correct etalon cavity mode (easily identified using the wavemeter as the modes are 1.6 GHz apart).

The presence of trapped ions can be confirmed either by sensitivity to the Doppler laser frequency, or by strobing the repump laser on and off to obtain fluorescence counts above the background scatter. Once cooled, single ions are readily imaged on the CCD camera:



Figure 6.7: Single ion of calcium in Liverpool trap.

6.2.4 Micromotion compensation

As discussed in Chapter 2, stable solutions to the Mathieu equations combine a slow secular harmonic oscillation with driven oscillation at the RF drive frequency, Ω_{RF} . This driven oscillation is called micromotion, and its amplitude is proportional to the amplitude of the RF electric field at the ion's location. A cold ion sits at the bottom of an overall potential given by the sum of the DC electric potential and the RF pseudo-potential (Section 2.2). If the minima of these two

potentials are aligned, the ion will be at the null of the oscillating RF potential, and micromotion will be at a minimum[†]. When they are not aligned, due to trap geometry errors or stray fields from charged surfaces in the vicinity of the trap, then driven micromotion occurs.

Micromotion can shorten ion confinement time in the absence of cooling and it adds an effective broadening to spectral features due to Doppler shifts arising from the ion motion [7]. It is to our advantage, then, to minimise micromotion, which we accomplish by applying compensating electric fields to bring the DC potential in alignment with the RF null. Coarse compensation can be guided by the ion's spectrum when scanning either the Doppler or repump laser across the resonance—poor compensation broadens and distorts the resonance peak.

The method of correlating photon arrival time with the RF drive gives increased sensitivity to micromotion [7]. This is accomplished with a time-to-amplitude converter gated on photon arrivals and RF drive sync pulses (Section 4.4) to build up histograms of number of arrivals versus RF phase. The stronger the correlation, the worse the micromotion.




Figure 6.8: Histogram of photon-arrival with RF phase. Width is one full RF period. The x -bias is applied via DC control electrodes (see text).

Compensation fields in the x -direction are applied with a bias to the DC control electrodes (Section 6.1.4). We calculate the size of the compensation field by trap simulations: 1150 V/m field per V of applied bias. Typical correlation histograms versus x -bias are in Figure 6.8. We easily discriminate photon histograms by eye for changes in the bias by 0.005 V corresponding to 6 V/m field.

Compensation fields in the y -direction are applied via the compensation electrodes (Figure 6.1, right). Trap simulations give 8.49 V/m per V applied to the electrodes. We discriminate photon histograms by eye for changes in the bias by 0.1 V. Better compensation in the y - than the x -direction stems from the beam geometry. The RF correlation method is sensitive only to the projection of the ion motion onto the laser wave vector. We detect y micromotion with a beam at $\sim 50^\circ$ to the y -axis, but the x micromotion with a beam at $\sim 80^\circ$ to the x -axis.

We do not attempt to compensate micromotion in the z -direction. The ideal RF potential is a 2-d quadrupole with no field components in the z -direction.

[†]some small micromotion always remains because the ion cannot actually be stationary due to finite temperature and quantum uncertainty effects.

The actual fields will, of course, have some components in the z -direction due to the finite length of the electrodes, but near the geometric centre of the trap these will be negligible.

6.3 Characterisation

Some rough bounds on the rate that the trap heats ions (in the absence of cooling) were established using the ‘‘Doppler recool’’ method [67]. Cooling is switched off for a time τ (in our case by turning off the repump laser beam) and the ion allowed to heat. We then monitor the ion fluorescence as a function of time when cooling is switched back on, by collecting photons in $50\ \mu\text{s}$ bins. This experiment is repeated for, typically, $N = 1024$ shots, to produce an average fluorescence versus time plot for the first 50 ms of the recool process. Significant changes in fluorescence occur when the Doppler shifts $\Delta_D = \mathbf{k} \cdot \mathbf{v}$ are of order the transition linewidth, Γ . For a trapped ion, maximal velocity occurs when all its energy is kinetic $v_{\text{max}} = \sqrt{2E/m}$. Thus we expect to see the effect on fluorescence when the ion has heated to $E = m\Gamma^2/2k_i^2 \approx k_B \times 0.4\ \text{K}$ for the Doppler transition[‡] in Ca^+ . We see these effects with τ as little as 100 ms, so this indicates a heating rate of several K/s.

This observed heating rate is quite high for a trap of this size. It is roughly three orders of magnitude higher than the heating rate observed in [45], where rates as low as 1 motional quantum in 190 ms were observed. The trap in that study was a ring-geometry Paul trap of similar size to the Liverpool trap, and the rate they observed is consistent with other traps when a scaling of d^{-4} with trap size is taken into account. Heating rates are often converted to an inferred electric field spectral noise density S_E to compare different ion species and secular frequencies. A heating rate of 4 K/s, assuming a trap frequency $\omega = 2\pi \times 300\ \text{kHz}$, indicates $S_E(\omega) = 6 \times 10^{-10}\ \text{V}^2\ \text{m}^{-2}\ \text{Hz}^{-1}$. See Figure 7.12 of [27] for a survey of reported heating rates.

We have not yet made a sustained effort to eliminate the possibility that our equipment is injecting noise which is responsible for the anomalous heating, so the rate we report should be considered preliminary. The partial collapse and uncollapse experiments detailed in Chapters 7 and 8 do not require coherent control of the motional state of the ion, but rather depend on reliable trapping with long ion lifetimes, and a low rate of drift of static magnetic and electric fields. In these respects the Liverpool trap has been very successful.

[‡]with k_i the projection of the wavevector onto one of the trapping axes, $|\mathbf{k}|/\sqrt{2}$ in our system, and k_B the Boltzman constant.

We obtain high accuracy fluorescence spectra of trapped ions by scanning either the Doppler or repump laser across the resonance. Relative laser frequencies can be controlled very accurately by adjusting the piezo voltages on the NPL etalons used to stabilise them, and absolute frequency is determined by the position of the resonance peak. From these spectra we may deduce the laser intensity and linewidth, as well as the applied magnetic field by curve fitting. These fitted parameters allow accurate calibration of laser intensity from the power measured with a photodiode, and of magnetic field from the measured coil current.

The theory curves are generated by numerical solution of the optical Bloch equations governing the ion dynamics. Working in a density matrix formalism, we include incoherent processes of spontaneous decay and dephasing from finite laser linewidths, along with stimulated transitions from the lasers. See [56] or [37] for detail on how these equations are set up. The scan speeds are all very slow compared with the timescales associated with ion dynamics, so fitting a scan is reduced to finding the steady state solution for each detuning—which amounts to inverting a $n^2 \times n^2$ matrix for an n -level system. The dynamics of the $S_{1/2} \leftrightarrow P_{1/2} \leftrightarrow D_{3/2}$ cooling manifold involve $n = 8$ states in $^{40}\text{Ca}^+$ which has no hyperfine structure. Since the fluorescence rate in steady state is proportional to the population in the $P_{1/2}$ states, we scale the detected photon counts by a constant to compare experiment and theory.



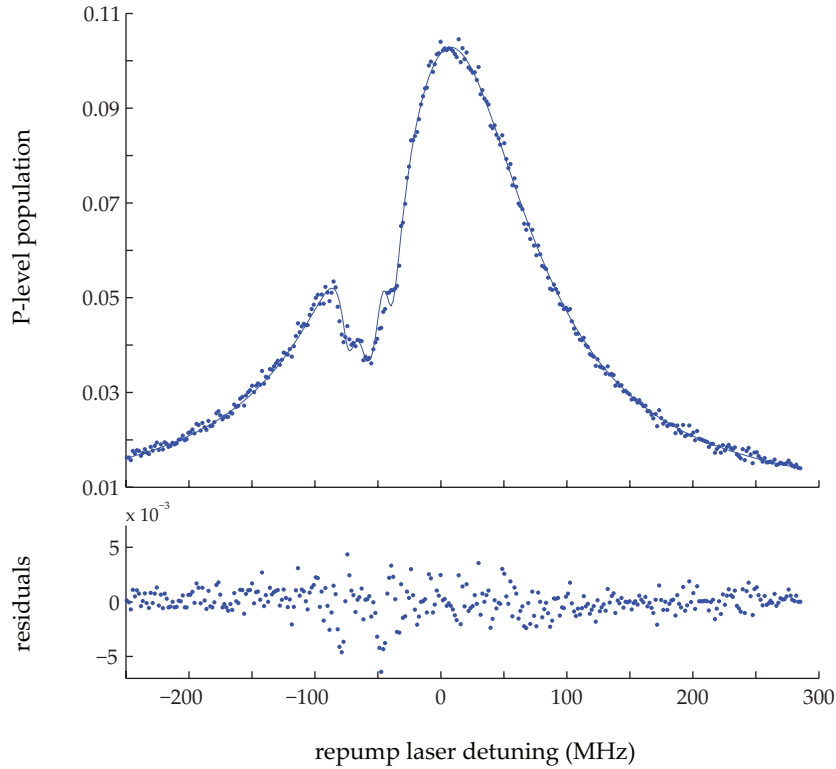


Figure 6.9: Fluorescence spectrum from scanning repump laser. Fluorescence has been scaled to represent population in the P levels, with theory curve calculated using an optical Bloch equation solver. The parameters fitted were Doppler laser intensity, linewidth, and detuning; repump laser intensity and linewidth; magnetic field; and background scattering rate.

Cooling the ion requires that the Doppler laser detuning be kept red of the transition resonance, so we typically choose to scan the repump laser to obtain a spectrum, since we may then explore both positive and negative detunings. Figure 6.9 shows an example fluorescence spectrum, with counts scaled to represent P-level population along with a theory curve fit to the data. For the scan depicted, the background count rate was 0.9 kHz, which is typical when the Doppler laser power is $30 \mu\text{W}$. The peak is 11.4 kHz above the background. The Doppler laser was detuned ~ 60 MHz to the red of the transition. The dips in fluorescence on the red side of the peak are due to a *dark resonance*: when the Doppler and repump lasers have the same detuning they excite a coherent two-photon transition from $S_{1/2} \leftrightarrow D_{3/2}$ which does not populate the $P_{1/2}$ level. Since the $D_{3/2}$ level is not appreciably broadened by spontaneous decay, the dark resonance is narrow enough to resolve the Zeeman components which are responsible for the structure in the dip. From this structure we may find the magnetic field strength. The residuals indicate that the dark resonance structure is not completely accounted for. This is primarily a consequence of the fact that at some detunings within the dark resonance, the lasers can heat the ion enough

to disrupt the fluorescence.

Figure 6.9 is part of a series of spectra used to calibrate the primary (Figure 6.6) magnetic field coil, which generates the static field to split the $D_{5/2}$ Zeeman states for the experiments in Chapters 7 and 8. At zero magnetic field, driving the ion with the Doppler and repump lasers results in optical pumping to $D_{3/2}$. This happens no matter the polarisation of the repump laser, since with no magnetic field, the repump laser itself determines the quantisation axis. This gives a clear experimental indication of when the field is nulled—ion fluorescence drops to essentially zero, and remains zero even when the polarisation of the repump laser is rotated using a waveplate. Using this technique, we determined the coil currents required to null the field.

Coil	Current (A) to null	Calibration (G/A)
Vert. trim	2.3	3.4*
Horz. trim	0.5	2.1*
Primary field	0.0	1.9

Table 6.1: Coil currents to null the magnetic field and calibrations. * trim coil calibrations are approximate, calculated using (6.1)

We then recorded a series of spectra, scanning the repump laser, increasing the current in the primary field coil in 1 A steps up to 4 A. Fits to the spectra give the magnetic fields, and a linear fit to the extracted data gives 1.9 G/A (which is 0.9 times the approximate value calculated using (6.1)). Detailed calibrations of the trim coils were not conducted.

Interestingly, the coil currents required to null the field indicate that the background field with no currents is ~ 7 G and nearly vertical—a fact that we have confirmed with additional fluorescence spectra. This is surprising since the trap itself and the vacuum system are all made from nominally non-magnetic materials.

7

$D_{5/2}$ Spin Experiments

In this chapter we describe a qubit stored in two Zeeman states of the $D_{5/2}$ subspace. Our implementation is a true “qubit” in the sense that it can be initialised to a known state, manipulated via single-qubit gates, and read out with high fidelity. There are, of course, several other schemes for implementing a qubit in $^{40}\text{Ca}^+$ [34] [40] [49], each with associated advantages. The motivation for developing this particular qubit is to allow partial measurements of the type discussed in section 3.4, and thus the partial collapse and uncollapse of a qubit wavefunction. We begin with a theoretical discussion of the scheme—deriving the main operational parameters—and conclude with experimental results demonstrating manipulation and readout of the qubit. We present a demonstration of the partial collapse and restoration of the qubit wavefunction in the next chapter.

7.1 Theory




Figure 7.1: σ^- -coupled transitions

The qubit levels are the $m = -\frac{3}{2}$ and $m = -\frac{5}{2}$ Zeeman states in $D_{5/2}$, which we manipulate by driving magnetic transitions with an RF coil. For the duration of this chapter, we take m to mean m_J , the z -projection of the total angular momentum. At low magnetic field, all the magnetic dipole transitions within the

unperturbed $D_{5/2}$ subspace occur at the same frequency. We isolate the qubit transition by the application of an intense laser field in the vicinity of the $P_{3/2} \leftrightarrow D_{5/2}$ transition, but polarised σ^- , such that it only couples to the $m \geq -\frac{1}{2}$ states in $D_{5/2}$. This is the well-known *AC Stark effect* or *light-shift*, and has the effect of altering the energies of these levels, shifting them out of resonance with the qubit transition.

The initial state preparation is via optical pumping (Figure 7.2). Measurement of the qubit's state is accomplished by using the polarised deshelling lasers to deshelve the $m = -\frac{3}{2}$ level, followed by fluorescence detection using the Doppler and repump lasers. The "strength" of this measurement can be continuously adjusted by altering the duration or intensity of the deshelve pulse.

It is worth noting that the choice of $m = -\frac{3}{2}$ and $-\frac{5}{2}$ for qubit states is arbitrary. We could just as easily choose $m = +\frac{3}{2}$ and $+\frac{5}{2}$, if we choose opposite senses of the σ polarisation components, and detune to the opposite side of the transition with the light-shift laser (see section 7.1.2).

7.1.1 State preparation

We prepare the initial state via optical pumping. Figure 7.2 shows the state preparation procedure schematically, with all illustrated lasers on for the entire pulse. The shelving laser drives $S_{1/2} \leftrightarrow P_{3/2}$, from which population can spontaneously decay back to the ground state, or to one of the two D levels: 5.3% branching to $D_{5/2}$ and 0.63% to $D_{3/2}$. The repump laser drives $D_{3/2} \leftrightarrow P_{1/2}$ and thus clears any population in the $D_{3/2}$, and two polarised beams clear all the $D_{5/2}$ states, save $m = -\frac{5}{2}$. These beams are derived from the deshelling laser driving $D_{5/2} \leftrightarrow P_{3/2}$; one is polarised π , the other σ^- .

This scheme relies on spontaneous (and therefore incoherent) dynamics of the ion. This allows us to consider only the state populations and not the coherences in our analysis. Coherences affect the dynamics when stimulated emission and absorption dominate the transitions involving a given level—when spontaneous transitions dominate (either into or out of a level) the coherences can safely be neglected in an approximate analysis. For $^{40}\text{Ca}^+$ the system is very simple, requiring only 18 real parameters, one for each Zeeman state. Under this approximation, the dynamics of the atom are reduced to a set of rate equations of the form [65]

$$\dot{\mathbf{n}} = \mathbf{M}\mathbf{n} \quad (7.1)$$

where n_i are the state populations, and \mathbf{M} is a matrix detailing processes which cause population to enter or leave each level. \mathbf{M} can be decomposed into a spontaneous component and a stimulated component, $\mathbf{M} = \mathbf{M}_{\text{spont}} + \mathbf{M}_{\text{stim}}$, with the former a constant matrix depending only on the choice of atom, and the latter




Figure 7.2: Schematic rendering of preparation of the $D_{5/2}, m = -\frac{5}{2}$ by optical pumping. Fine and Zeeman structure exaggerated.

determined by the intensity and detuning of applied laser fields. The spontaneous emission matrix is generated by considering the observed *level* decay rates (the Einstein A coefficients), and scaling them by Clebsch-Gordan coefficients to obtain the transition rates between the particular magnetic states. The stimulated transition rate is proportional to the intensity of the laser accessing the transition and has the same dependence on magnetic states as the spontaneous emission rate. The stimulated transition rate is also scaled by a lineshape function which depends on the detuning. A complete treatment of the construction of these matrices can be found in [61].

Given the matrix M , we may solve for steady-state populations by setting $\dot{\mathbf{n}} = 0$ and imposing one additional linear constraint, that the populations sum to 1, and numerically solving in the usual way. If we are interested in the system dynamics, then we may numerically integrate (7.1). In particular, for constant M

$$\mathbf{n}(t) = e^{Mt} \mathbf{n}_0 \quad (7.2)$$

or, in the case where laser pulses characterised by matrices $\{M_i\}$ are turned on for times $\{\tau_i\}$, the final state will be

$$\mathbf{n} = \left(\prod_i e^{M_i \tau_i} \right) \mathbf{n}_0 = \left(\prod_i U_i \right) \mathbf{n}_0 \quad (7.3)$$

The extremely long lifetime of the $D_{5/2}$ level ($\tau = 1.168(7)$ s [6]) against spontaneous decay means that ideally, so long as the overall rate at which population is transferred to this level, R_{shelve} , is fast compared with the spontaneous decay rate (i.e. $R_{\text{shelve}}\tau \gg 1$), we should attain nearly 100% state preparation fidelity. In practice, imperfections in the polarisation of the deshelving lasers, not spontaneous emission, will be the dominant process which depopulates the $D_{5/2}$ $m = -\frac{5}{2}$ sub-level. Also, population decaying from $P_{3/2}$ is distributed across all m sub-levels, which must be effectively cleared via the deshelving lasers. Given impure polarisation, there is a trade-off for the deshelving laser power: too weak and we do not clear the $m \neq -\frac{5}{2}$ states fast enough, too strong and we start to deshelve $m = -\frac{5}{2}$.

This type of optimization problem is non-linear, but fortunately (7.3) can be evaluated quickly enough to allow numerical optimization to complete in a reasonable amount of time. We use *MATLAB*'s `fmincon` command, which computes the minimum of a multi-variable scalar function subject to constraints. The scalar function, f , is a figure-of-merit for how good the state preparation is with the given parameters. We take as the parameters to be optimised the intensities of the shelving, π -deshelve and σ^- -deshelve lasers and the length of the state-preparation pulse: $f(I_s, I_\pi, I_\sigma, \tau)$. The state preparation is not sensitive to the intensity of the repump laser in the regime in which we typically operate, which is moderate saturation. We compute the matrix M , given these parameters and experimentally determined polarisation purities, solve (7.3), and take f to be

$$f = \gamma B - G \quad (7.4)$$

where B is the sum of populations in the $D_{5/2}$ $m \neq -\frac{5}{2}$ states (the ‘‘bad’’ population) and G is the $m = -\frac{5}{2}$ population (the ‘‘good’’ population). These are computed after applying the state preparation pulse, assuming initial populations equally distributed in the ground-state ($S_{1/2}$) states. γ is a biasing factor which we can use to adjust the relative importance of minimizing bad population to maximising good population. Why might one of these things be more important? Because we are going to be doing experiments in which we select ensembles based on measuring the ion to be in the $D_{5/2}$ subspace, any population in the ground state is eliminated by this selection (assuming high-fidelity readout). Whereas, any stray population in the $m \neq -\frac{5}{2}$ states of the $D_{5/2}$ subspace is not eliminated by our measurement selection. Cast in this light, the good population represents the probability of successfully preparing the state, and the bad population represents the impurity of the state we finally prepare. A low G simply makes our data collection take longer, since fewer runs will be selected into our final ensemble, a high B lowers the fidelity.

`fmincon` also allows linear constraints on the parameters to be specified, which

we use to keep the optimization routine from choosing laser powers too small to be conveniently measured in the lab. Final operating parameters were determined by choosing a convenient minimum laser power, then a maximum bad population B , and increasing the bias γ until `fmincon` produced a solution which had B below this threshold.

7.1.2 Coherent manipulation

The Zeeman effect, characterised by the $-\boldsymbol{\mu} \cdot \mathbf{B}$ term in the Hamiltonian, splits levels defined by their total angular momentum, J , into $2J + 1$ components, each an eigenstate of J_z with eigenvalue m . For weak fields, the splittings are all equal and given by the Larmor frequency, $\omega_0 = g_J \frac{\mu_B}{\hbar} B$, with g_J the Landé g -factor.

Consider the case where, in addition to a constant magnetic field, we also apply an oscillating field $\mathbf{B} = \mathbf{B}_0 + \mathbf{B}' \cos \omega t$, with $B_0 \gg B'$ such that perturbation theory is appropriate. The constant field establishes the $|J, m\rangle$ basis states and our perturbing Hamiltonian is given by $\boldsymbol{\mu} \cdot \mathbf{B}'$. If \mathbf{B}' is not parallel to \mathbf{B}_0 , then the Hamiltonian will have off-diagonal matrix elements within the subspace defined by J , and we may excite transitions between the m states. The resonant frequency for each of these transitions is none other than the Larmor frequency, and because we have a *dipole* matrix element, the selection rules are $\Delta m = \pm 1$.

The degeneracy in transition frequency can be lifted with the application of a light-shift, allowing us to coherently manipulate a superposition of just two $|m\rangle$ states. We want to be able to characterise our experimental apparatus for magnetic driving without the added complexity of the light-shift laser, so it will be helpful to first consider a theoretical description of magnetic resonance in the $D_{5/2}$ subspace without a light-shift. We will then tackle the light-shift itself, and conclude with the description of magnetic driving in the light-shifted subspace.

Unshifted $D_{5/2}$ subspace

Taking \mathbf{B}' to be oriented along $\hat{\mathbf{x}}$ with the constant field oriented along $\hat{\mathbf{z}}$, the interaction Hamiltonian can be written [68]

$$\hat{H}_I = \boldsymbol{\mu} \cdot \mathbf{B}' \cos \omega t = g_J \mu_B B' \hat{J}_x \cos \omega t \quad (7.5)$$

We take $\Omega_B = g_J \mu_B B' / 2\hbar$ as the *reduced Rabi frequency*, analogous to the Larmor frequency, but proportional to the magnitude of the oscillating field (divided by two to make later math easier). We do not consider transitions out of the $D_{5/2}$ subspace as these involve terahertz or optical frequencies and will be far off-resonant from our RF driving field. We remove the time dependence from (7.5) by transforming to a reference frame rotating at the driving frequency, ω , and making the usual rotating wave approximation. This calculation is given in

detail in Appendix A.1. We shall describe the physical interpretation and simply quote the results.

The Hamiltonian associated with the constant magnetic field rotates the system about the z -axis at the Larmor frequency. Our reference frame, however, rotates at the RF driving frequency. When we absorb this evolution into the interaction picture, we introduce a rotation at the difference frequency $\delta = \omega_0 - \omega$ about the z -axis. This results in an extra term proportional to \hat{J}_z . The now time-independent Hamiltonian of the system in the rotating frame is

$$\hat{H}_R = \delta \hat{J}_z + \Omega_B \hat{J}_x \quad (7.6)$$

This gives rise to a time evolution operator, via the Schrödinger Equation, of

$$\mathcal{U}(\tau) = \exp\left(\frac{-i(\delta \hat{J}_z + \Omega_B \hat{J}_x)\tau}{\hbar}\right) \quad (7.7)$$

which, if we further define $W = \sqrt{\delta^2 + \Omega_B^2}$ and a direction unit vector $\hat{\mathbf{r}} = W^{-1} \langle \Omega_B, 0, \delta \rangle$, we can rewrite in the form

$$\exp\left(\frac{-i(\delta \hat{J}_z + \Omega_B \hat{J}_x)\tau}{\hbar}\right) = \exp\left(\frac{-i\mathbf{J} \cdot \hat{\mathbf{r}} W \tau}{\hbar}\right). \quad (7.8)$$

This is a rotation through angle $W\tau$ about the axis defined by $\hat{\mathbf{r}}$.




Figure 7.3: Axis of rotation, \mathbf{r} .

A spin- $\frac{1}{2}$ system (or indeed, any two level system) interacting with monochromatic radiation is commonly represented as a Bloch vector being rotated about different axes. Our calculation illustrates the fact that we need not restrict this powerful physical picture to the spin- $\frac{1}{2}$ case. Indeed *any* system whose interaction Hamiltonian can be expressed as a linear combination of J operators has a time evolution which consists simply of rotations. However, for $J > \frac{1}{2}$, we can no longer regard the quantum state as a simple vector in the Bloch sphere.

Real space no longer affords sufficient dimensions to express the complexity of a multilevel quantum superposition! Nonetheless, the time evolution operators of the system are rotation operators which can be computed analytically, as we now demonstrate.

Any rigid-body rotation may be decomposed into three steps known as *Euler rotations*, each parameterized by a single real angle. In quantum mechanics, the standard textbook [46] procedure is to make these rotations about fixed-space axes:

$$R(\alpha, \beta, \gamma) = R_z(\alpha)R_y(\beta)R_z(\gamma). \quad (7.9)$$

We shall find it more convenient, however, to decompose them into y - z - y rotations, because the angles then take a convenient form. Define $\theta = \tan^{-1}\left(\frac{\Omega}{\delta}\right)$, the inclination angle of \mathbf{r} (Figure 7.3). We then express the rotation in (7.8) as

$$R(\hat{\mathbf{r}}, W\tau) = R_y(\theta)R_z(W\tau)R_y(-\theta). \quad (7.10)$$

We may calculate the operators in the $|J, m\rangle$ basis corresponding to each rotation above. The matrix elements of rotations about z take a particularly simple form, owing to the fact that we are in the eigenbasis of \hat{J}_z :

$$\langle J, m' | \exp\left(\frac{-i\hat{J}_z\phi}{\hbar}\right) | J, m \rangle = e^{-im\phi}\delta_{m,m'} \quad (7.11)$$

with $\delta_{m,m'}$ the Kronecker Delta. To obtain matrix elements of y -axis rotations we use *Wigner's formula*:

$$\begin{aligned} \langle J, m' | \exp\left(\frac{-i\hat{J}_y\beta}{\hbar}\right) | J, m \rangle = \\ \sum_k (-1)^{k-m+m'} \frac{\sqrt{(J+m)!(J-m)!(J+m')!(J-m')!}}{(J+m-k)!k!(J-k-m')!(k-m+m')!} \\ \times \left(\cos\frac{\beta}{2}\right)^{2J-2k+m-m'} \left(\sin\frac{\beta}{2}\right)^{2k-m+m'} \end{aligned} \quad (7.12)$$

where the sum over k is taken only over terms for which the factorials in the denominator have non-negative arguments [46].

◇

For the $D_{5/2}$ subspace, each of the above operators will be a 6×6 matrix. An explicit functional form for the y -axis rotation would take an entire page to print, so we instead we present, in Appendix A.2, *Mathematica* code to evaluate the operator. This code can be used to generate plots like Figure 7.4, which shows a frequency scan of a magnetic pulse depleting population initially prepared in $|m = \frac{5}{2}\rangle$




Figure 7.4: $m = \frac{5}{2}$ population remaining after magnetic driving pulse, $\tau = \frac{\pi}{\Omega}$.

Light shifts

We now present the calculation of the light-shifts in the $D_{5/2}$ subspace obtained by applying an intense laser beam off-resonant with the $D_{5/2} \leftrightarrow P_{3/2}$ transition. The laser field couples states in the upper and lower levels such that the stationary states are a mixture of the two. Using perturbation theory, we may calculate the energy shifts of the new stationary states relative to the unperturbed ones. This amounts to calculating the eigenvalues of the interaction Hamiltonian. After making the electric dipole and rotating wave approximations, and transforming to an interaction picture (Appendix A.3 for details), the time-independent perturbing Hamiltonian is

$$\hat{H}' = -\hbar\delta\hat{I}^{(P_{3/2})} + \hbar\omega_B\hat{Z} - \frac{E_0}{2}\hat{\mu}. \quad (7.13)$$

δ is the detuning of the laser (with $\delta < 0$ red-detuning) with respect to the centre-of-mass of the Zeeman-split resonance, and the superscript on the \hat{I} means identity for the $P_{3/2}$ states, and 0 for all others. \hat{Z} is a diagonal matrix containing the Zeeman shifts (including different Landé g-factors) for the levels, with

$\omega_B = \frac{\mu_B}{\hbar} |\mathbf{B}|$. E_0 is the amplitude of the electric field and $\hat{\mu}$ is the dipole interaction matrix

$$\hat{\mu}_{jk} = e \langle j | \hat{\epsilon} \cdot \mathbf{r} | k \rangle. \quad (7.14)$$

We evaluate the angular dependence of $\hat{\mu}$ by decomposing it into a scalar reduced dipole matrix element, d , multiplied by an angular matrix, \hat{M} . \hat{M} may be evaluated by Clebsch-Gordan coefficients, or alternately, as Wigner-3j symbols, as follows:

$$\langle \alpha', J', m' | \hat{\mu} | \alpha, J, m \rangle = (-1)^{J'-m'} \begin{pmatrix} J' & 1 & J \\ -m' & q & m \end{pmatrix} d \quad (7.15)$$

$$= \hat{M}_{q,m',m} d \quad (7.16)$$

where $q = \{-1, 0, +1\}$ correspond to $\{\sigma^-, \pi, \sigma^+\}$ polarisation components, and we choose the convention that this matrix be real-valued.

We substitute this into (7.13), and for ease of calculation, drop the \hbar 's and express everything in frequency units. We also absorb E_0 and the dipole matrix element into a light-shift Rabi frequency $\hbar\Omega_{LS} = E_0 d$,

$$\hat{H}_{LS} = -\delta \hat{I}^{(P_{3/2})} + \omega_B \hat{Z} - \frac{\Omega_{LS}}{2} \hat{M}_q. \quad (7.17)$$

For σ^- light this is

$$\begin{array}{c} \text{D}_{5/2} \qquad \qquad \qquad \text{P}_{3/2} \\ \hline \begin{matrix} m = \frac{5}{2} & & & & & & m = -\frac{5}{2} & m = \frac{3}{2} & & & m = -\frac{3}{2} \end{matrix} \\ \left(\begin{array}{ccccccccccc} 3\omega_B & 0 & 0 & 0 & 0 & 0 & -\frac{\Omega_{LS}}{2\sqrt{6}} & 0 & 0 & 0 & 0 \\ 0 & \frac{9\omega_B}{5} & 0 & 0 & 0 & 0 & 0 & -\frac{\Omega_{LS}}{2\sqrt{10}} & 0 & 0 & 0 \\ 0 & 0 & \frac{3\omega_B}{5} & 0 & 0 & 0 & 0 & 0 & -\frac{\Omega_{LS}}{4\sqrt{5}} & 0 & 0 \\ 0 & 0 & 0 & -\frac{3\omega_B}{5} & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{\Omega_{LS}}{4\sqrt{15}} \\ 0 & 0 & 0 & 0 & -\frac{9\omega_B}{5} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -3\omega_B & 0 & 0 & 0 & 0 & 0 \\ -\frac{\Omega_{LS}}{2\sqrt{6}} & 0 & 0 & 0 & 0 & 0 & 2\omega_B - \delta & 0 & 0 & 0 & 0 \\ 0 & -\frac{\Omega_{LS}}{2\sqrt{10}} & 0 & 0 & 0 & 0 & 0 & \frac{2\omega_B}{3} - \delta & 0 & 0 & 0 \\ 0 & 0 & -\frac{\Omega_{LS}}{4\sqrt{5}} & 0 & 0 & 0 & 0 & 0 & -\delta - \frac{2\omega_B}{3} & 0 & 0 \\ 0 & 0 & 0 & -\frac{\Omega_{LS}}{4\sqrt{15}} & 0 & 0 & 0 & 0 & 0 & -\delta - 2\omega_B & 0 \end{array} \right) \end{array} \quad (7.18)$$

with $g_j = \frac{4}{3}$ for $P_{3/2}$ and $\frac{6}{5}$ for $D_{5/2}$. The eigenvalues are plotted in Figure 7.5 for $\delta > 0$. We see that the $m = -\frac{5}{2}$ and $-\frac{3}{2}$ states are unperturbed, as desired. It is important to choose the correct sign of the detuning, otherwise the crossings seen in the $P_{3/2}$ states will be in the $D_{5/2}$ states. If we had instead chosen the σ^+




Figure 7.5: Light shifted energy levels for σ^- light, and blue detuning. Energies (ω) and light intensity (Ω_{LS}^2) are scaled by δ , which is the separation between the $D_{5/2}$ and $P_{3/2}$ centre-of-mass energies (horizontal black lines). We do *not* have avoided crossings because there is no coupling between different m states of the same J .

polarisation for our light-shift, the $m = \frac{5}{2}$ and $\frac{3}{2}$ states would be unperturbed, and $\delta < 0$ would keep the $D_{5/2}$ states from crossing.

To minimise scatter out of the $D_{5/2}$ subspace during coherent manipulation, we work at large detuning, and it is appropriate to expand the eigenvalues in powers of Ω_{LS}^2 . We also take ω_B to be zero, since $\omega_B \ll \delta$. We are concerned with the shift in the transition frequency, which will be the difference in the light shift of adjacent levels. The eigenvalues of (7.18) are listed in full in Appendix A.4,

and in Table 7.1 we list the truncated power series.

m -transition	shift
$\frac{5}{2} \rightarrow \frac{3}{2}$	$\frac{4}{240} \frac{\Omega_{LS}^2}{\delta} + \dots$
$\frac{3}{2} \rightarrow \frac{1}{2}$	$\frac{3}{240} \frac{\Omega_{LS}^2}{\delta} + \dots$
$\frac{1}{2} \rightarrow -\frac{1}{2}$	$\frac{2}{240} \frac{\Omega_{LS}^2}{\delta} + \dots$
$-\frac{1}{2} \rightarrow -\frac{3}{2}$	$\frac{1}{240} \frac{\Omega_{LS}^2}{\delta} + \dots$
$-\frac{3}{2} \rightarrow -\frac{5}{2}$	0

Table 7.1: Light-shifted magnetic transition frequencies

Now, the square of the Rabi frequency is proportional to the intensity of the driving laser. The constant of proportionality depends on the reduced dipole matrix element for the transition, and it is worked out in detail in Appendix A.5. For the $P_{3/2} \leftrightarrow D_{5/2}$ transition in $^{40}\text{Ca}^+$,

$$\Omega_{LS}^2 = 4.61 \times 10^{13} \left[\frac{\text{m}^2}{\text{Js}} \right] \times I_L. \quad (7.19)$$

We now consider the question of how to choose the detuning for the light-shift laser. We examine two major error processes: off-resonant optical excitation to the $P_{3/2}$ level due to imperfections in the polarisation of the light-shift laser, and off-resonant magnetic scattering on the $m = -\frac{3}{2} \rightarrow -\frac{1}{2}$ transition. For a given laser intensity, increasing δ decreases the former, but increases the latter by reducing the light-shift.

The optical excitation rates are given by *Fermi's Golden Rule*:

$$R_{ij} = \frac{\pi}{2} \Omega_{ij}^2 g(\delta) \quad (7.20)$$

with $g(\delta)$ the lineshape given by

$$g(\delta) = \frac{1}{2\pi} \frac{\Gamma}{\delta^2 + \Gamma^2/4}. \quad (7.21)$$

Now, the Ω_{ij} are the Rabi frequencies for the transition $|i\rangle \rightarrow |j\rangle$, which in terms of the light-shift Rabi frequency ($\hbar\Omega_{LS} = dE_0$) are

$$\Omega_{ij} = \langle j | \hat{M}_q | i \rangle \Omega_{LS}, \quad (7.22)$$

\hat{M}_{-1} , corresponding to σ^- light, has no matrix elements connecting to $|D_{5/2}, -\frac{5}{2}\rangle$ or $|D_{5/2}, -\frac{3}{2}\rangle$, so the scattering rate depends on the polarisation impurity introducing π or σ^+ light. For simplicity, and since polarisation purity varies over the course of an experimental run, we shall assume that both I_{σ^+} and I_{π} are given by εI_L . If the polarisation components were known to sufficient precision, we could obtain a more accurate theory by explicitly calculating the matrix elements in (7.22), but instead we simply take a representative value ($10^{-1/2}$) and write

$$R_{\text{opt}} \approx \frac{\varepsilon\pi}{20} \Omega_{LS}^2 g(\delta) \quad (7.23)$$

$$\approx \frac{\varepsilon}{40} \Omega_{LS}^2 \frac{\Gamma}{\delta^2} \quad (7.24)$$

where we have used a series approximation for the lineshape since $\delta \gg \Gamma$.

We shall now consider the magnetic scattering rate. For simplicity of notation, we label the $D_{5/2}$ m -states $\{0 \dots 5\}$ for $m = \{-\frac{5}{2} \dots \frac{5}{2}\}$. We assume that the populations of all states save $|0\rangle$ and $|1\rangle$ remain small, so that the only transition we are concerned with is $|1\rangle \rightarrow |2\rangle$. The probability of exciting a transition is

$$P_{\text{mag}} = \left(\Omega_{12} \frac{\sin(-\delta_B t/2)}{\delta_B} \right)^2, \quad (7.25)$$

where δ_B is the detuning of the RF magnetic field. At large detunings, relatively small changes in δ_B result in large changes in the probability, due to the $\sin(-\delta_B t/2)$ term in the numerator. We set this to 1, to give the worst case. We also note that the magnetic driving Rabi frequencies occur in a fixed ratio, which can be derived directly from the matrix elements of \hat{J}_x . For $J = \frac{5}{2}$, $\Omega_{12} = 2\sqrt{2/5}\Omega_{01}$. We rewrite (7.25)

$$P_{\text{mag}} = \left(\frac{2\sqrt{2/5}\Omega_{01}}{\Delta_{LS}} \right)^2 \quad (7.26)$$

where we have substituted $\Delta_{LS} = \delta_B$ to indicate that the detuning of this transition comes entirely from the applied light-shift.

To compare with the optical scattering, we rewrite the rate (7.24) as a probability by multiplying by the time the light-shift laser needs to be on, which is approximately Ω_{01}^{-1} . We may then use (7.19) and Table 7.1 to write (7.26) in terms of the light-shift laser detuning, set it equal to the optical scattering probability, and solve for δ . We note that the probability of magnetic scattering is $\propto (\Omega_{01}^2 \delta^2 / I_L^2)$, while for the optical case it is $\propto (I_L / \delta^2 \Omega_{01})$. The optimal value for δ is $\propto (I_L / \Omega_{01})^{3/4}$, and the final probability of scattering at optimal δ is

$\propto (\Omega_{01}/I_L)^{1/2}$. It is to our advantage to maximise the intensity of the light-shift laser. We could also lower the magnetic Rabi frequency Ω_{01} , but this choice will be dependent on the decoherence characteristics of the qubit (Section 7.4).

Shifted $D_{5/2}$ subspace

The application of the light-shift alters the interaction Hamiltonian we developed in (7.6). The shifting of the $D_{5/2}$ sublevels' energy spectrum means that the transformation which leaves the Hamiltonian time-independent is no longer a rotation about the z -axis. As a result, we are unable to write the diagonal entries in the (time-independent) interaction Hamiltonian as a multiple of \hat{J}_z and our analytic interpretation of the time-evolution as a rotation in real-space breaks down. Although perhaps not as mathematically elegant, we may obtain the time-evolution of the system by numerical integration—solving the Optical Bloch Equations.

We begin by noting that we do not change the time-*dependent* form of the Hamiltonian; it is still the oscillating magnetic dipole. The application of our light-shift is a perturbation on the original $|J, m\rangle$ states, and acts to mix the $D_{5/2}$ states with those of $P_{3/2}$. In the limit where $\delta \gg \Omega$, the mixing is small and we neglect the $P_{3/2}$ contribution to the magnetic dipole matrix elements in the shifted $D_{5/2}$ subspace. This means the off-diagonal matrix elements in the interaction Hamiltonian are identical to the unshifted case. Since the reference frame transformation does not affect the off-diagonal matrix elements, these will still be given by $\Omega_B \hat{J}_x$.



Examining (7.6), the time-independent version of the interaction for the unshifted subspace, we note that the diagonal matrix elements, given by

$$\langle J, m | \delta \hat{J}_z | J, m \rangle = \delta m \quad (7.27)$$

are the “accumulated” detunings of the levels, which is to say the detuning of the level from resonance summed with the detunings of each level below it. Figure 7.6 depicts this schematically.




Figure 7.6: Transformation to remove the time dependence due to a single driving field oscillating at angular frequency ω . Detunings from resonance become the splittings in the transformed energy spectrum.

With this in mind, we need not repeat the whole calculation to transform the Hamiltonian to time-independent form, but simply insert the light shifts calculated in the previous subsection to obtain the new detunings, remembering to accumulate these detunings as we move up the energy spectrum. Combining Table 7.1 with (7.27), we obtain

$$\hat{H} = \begin{pmatrix} m = -\frac{5}{2} & & & & & & m = \frac{5}{2} \\ -\frac{5}{2}\delta & \sqrt{5} & 0 & 0 & 0 & 0 & \\ \sqrt{5} & -\frac{3}{2}\delta & 2\sqrt{2} & 0 & 0 & 0 & \\ 0 & 2\sqrt{2} & -\frac{1}{2}\delta - \Delta_{LS} & 3 & 0 & 0 & \\ 0 & 0 & 3 & \frac{1}{2}\delta - 3\Delta_{LS} & 2\sqrt{2} & 0 & \\ 0 & 0 & 0 & 2\sqrt{2} & \frac{3}{2}\delta - 6\Delta_{LS} & \sqrt{5} & \\ 0 & 0 & 0 & 0 & \sqrt{5} & \frac{5}{2}\delta - 10\Delta_{LS} & \end{pmatrix} \quad (7.28)$$

with $\Delta_{LS} = \Omega_{LS}^2/240\delta_{LS}$ the *reduced light-shift*, defined this way since all light-shifts occur as integer multiples of this frequency. We obtain the time-evolution by evaluating the matrix exponential.

We contrast this method of solving for the dynamics, which implicitly takes into account the coherences between states, with the method employed in Section 7.1.1, which does not. A fully coherent analysis of state preparation involves 18 levels, and thus $18^2 = 324$ parameters, but we obtain accurate results with just 18 (the populations). Here we consider coherent dynamics of just 6 levels, and thus 36 parameters.

7.1.3 Readout

The novel readout of this qubit was our motivation to study it in detail. Here we consider its advantages and limitations.

We may discriminate between a $^{40}\text{Ca}^+$ ion in its ground state and one in the $D_{5/2}$ subspace with high fidelity by illuminating it with lasers on the $S_{1/2} \leftrightarrow P_{1/2}$ and $P_{1/2} \leftrightarrow D_{3/2}$ transitions. Scattered photons at 755 THz (397 nm) signal an ion that started in the ground state (the “bright” state), and the absence of photons (the “dark” state) indicates the ion is in the $D_{5/2}$ subspace. Because of background scatter off electrode surfaces and PMT dark counts, the detection of several photons is required to indicate a bright ion with high confidence. The most straightforward method is to count photons for a fixed time t_b , and define a threshold count n_c . If $\geq n_c$ photons are collected, the ion is measured bright, if $< n_c$, it is measured dark.

Statistical fluctuations in the number of detected photons result in incorrect detection of the ion’s state. The discrimination error decreases approximately exponentially as t_b increases by reducing the overlap of the bright and dark count probability distributions, but only until the error due to statistical fluctuations is reduced to of order the probability of a dark ion spontaneously decaying to the bright subspace, which for $t_b \ll \tau$ increases linearly as t_b/τ [10]. More sophisticated discrimination methods exist [40], but for this experiment other errors (cf. branching ratios below) are expected to dominate and we implement a simple threshold.

Scattering photons necessarily randomises and destroys any coherence among the ground state levels, but a coherent state in the $D_{5/2}$ subspace remains undisturbed.

In this way, the $^{40}\text{Ca}^+$ ground state is used as an auxiliary “detectable” subspace which we may selectively couple to the $D_{5/2}$ qubit described above. Recalling our description of the asymmetric partial- and uncollapse in Section 3.4.3, we consider the Hilbert space spanned by the qubit states $|0\rangle$, $|1\rangle$, and an auxiliary $|2\rangle$: the detectable state. The quantum state is taken as ρ , the density operator on that space. The quantum process $\mathcal{D}_p(\rho)$ incoherently transfers population from $|1\rangle$ to $|2\rangle$ with probability p . We also have the set of two measurement

operators suggestively named

$$\hat{M}_{\text{dark}} = |0\rangle\langle 0| + |1\rangle\langle 1|, \quad \text{and} \quad (7.29)$$

$$\hat{M}_{\text{bright}} = |2\rangle\langle 2|. \quad (7.30)$$

A full von Neumann-style measurement on the qubit is obtained by applying \mathcal{D}_p to the qubit with $p = 1$ and then measuring in the {dark, bright} basis. With $p < 1$, we obtain a partial measurement as described in Section 3.4.

We implement \mathcal{D}_p by optical pumping. The π -polarized deshelving laser is pulsed on, exciting transitions from $|1\rangle$, the $D_{5/2}$, $m = -\frac{3}{2}$ state to the $P_{3/2}$, $m = -\frac{3}{2}$ state where it predominantly decays to the ground state. The length and intensity of this pulse determine the strength, p . The $P_{3/2}$, $m = -\frac{3}{2}$ state also decays to $D_{3/2}$ and back to $D_{5/2}$. Since $D_{3/2}$ population is repumped during readout, decays here are functionally equivalent to decays to the ground state in this scheme. However, decays back to the $D_{5/2}$ corrupt the ideal process. The branching ratios are given in Table 7.2

Level	m		ratio
$S_{1/2}$	(all)		0.9405
$D_{3/2}$	(all)		0.0063
$D_{5/2}$	$-\frac{5}{2}$	$ 0\rangle$	0.0355
$D_{5/2}$	$-\frac{3}{2}$	$ 1\rangle$	0.0142
$D_{5/2}$	$-\frac{1}{2}$		0.0035

Table 7.2: Branching ratios from $P_{3/2}$, $m = -\frac{3}{2}$

These branching ratios result in several effects which deviate from the ideal. Firstly, the readout process becomes biased such that the probability of getting a bright ion, $P(\text{bright}) \neq p\rho_{11}$. Secondly, the $|0\rangle$ state is disturbed by an incoherent population transfer from $|1\rangle$ via the $P_{3/2}$ level. There are additional effects, like a slight decrease in the ρ_{01} coherence due to decay back to $|1\rangle$, and transfer of population elsewhere in $D_{5/2}$, but these are of comparatively less importance. The primary features of the readout bias and disturbing of $|0\rangle$ are captured by a simple model where \mathcal{D}_p is modified to include only the decay to $|0\rangle$, with branching ratio $\epsilon = 0.0355$. This is the “leaky transfer” described in Section 3.4.5. It will transfer ϵp of the population ρ_{11} to ρ_{00} , and $(1 - \epsilon)p$ to ρ_{22} :

$$\begin{pmatrix} \rho_{00} & \rho_{01} & \rho_{02} \\ \rho_{10} & \rho_{11} & \rho_{12} \\ \rho_{20} & \rho_{21} & \rho_{22} \end{pmatrix} \rightarrow \begin{pmatrix} \rho_{00} + \rho_{11}\epsilon p & \rho_{01}\sqrt{1-p} & \rho_{02} \\ \rho_{10}\sqrt{1-p} & \rho_{11}(1-p) & \rho_{12}\sqrt{1-p} \\ \rho_{20} & \rho_{21}\sqrt{1-p} & \rho_{22} + \rho_{11}(1-\epsilon)p \end{pmatrix} \quad (7.31)$$

Consider the case where $p = 1$, the full measurement of the qubit, initially in some superposition $\rho = |\psi\rangle\langle\psi|$ with $|\psi\rangle = a|0\rangle + b|1\rangle$.

$$P(\text{bright}) = \rho_{22} + \rho_{11}(1 - \epsilon)p \quad (7.32)$$

$$= |b|^2(1 - \epsilon) \quad (7.33)$$

and

$$P(\text{dark}) = \rho_{00} + \rho_{11}\epsilon p \quad (7.34)$$

$$= |a|^2 + |b|^2\epsilon \quad (7.35)$$

$$= |a|^2(1 - \epsilon) + \epsilon. \quad (7.36)$$

Inverting (7.33) and (7.36) give the qubit populations $|a|^2$ and $|b|^2$ in terms of experimentally measured parameters.

7.2 Polarization

The fidelity of state preparation and deshelling pulses depend critically on the application of laser light which is highly pure in one of the $\{\sigma^-, \pi, \sigma^+\}$ polarisation components. Here I present the setup of the apparatus to generate these polarised beams and experimental data which characterized them.

3 lasers have very stringent polarisation requirements: 2 deshelling lasers polarised σ^- and π , and one light-shift laser polarised σ^- . The others are not so critical; refer to Table 4.1 for nominal polarisations.

7.2.1 Optical setup




Figure 7.7: Optical setup to overlap the σ laser beams onto the same optical fibre input.

The two σ^- lasers were overlapped onto the same single-mode fibre using a 90-10 non-polarising beam splitter, and both adjusted to the polarisation-maintaining axis of the highly birefringent fibre (Figure 7.7). Since they are only separated by about 100 GHz (or $\sim 10^{-3}\lambda$), the output modes for the two lasers should be nearly identical. So, when one laser's polarisation is highly pure, the

other should be as well. The π -polarised deshelling laser was overlapped onto the same optical fibre as the repump laser using a polarising beam splitter (PBS) cube such that the two lasers travel down the fibre in nominally orthogonal polarisations.




Figure 7.8: Optical setup around the Liverpool trap for $D_{5/2}$ spin experiments. All beam splitter cubes are polarising. Scale approximate.

Figure 7.8 illustrates the optical setup around the trap system. The σ^- beams pass through a $\lambda/2$ waveplate and a PBS cube to clean up the output polarisation from the fibre. They reflect from a steering mirror, and through a further PBS, and finally a $\lambda/4$ waveplate aligned at 45° with respect to the PBS, which converts the plane-polarised light to circular. When the magnetic field at the ion's location is aligned with the Poynting vector of the circular light, this is σ^- or σ^+ polarisation. The handedness with respect to the magnetic field determines whether it is σ^+ or σ^- . We may operate the experiment with either polarisation; σ^- (σ^+) requires that the light-shift laser be detuned blue (red) from resonance to ensure that the shifted Zeeman states' energies do not cross as the light-shift is increased.

To create π -polarisation, we need plane-polarised light with the magnetic field at the ion in the plane of polarisation and perpendicular to the Poynting vector. This means our beam-line must be at a right angle to the σ^- beam-line, with the polarisation rotated into alignment with the magnetic field. The π -deshelve laser beam emerges from the fibre coupler and passes through a $\lambda/2$

waveplate and a PBS to separate it from the repump laser. It then reflects off a steering mirror, then passes through a further $\lambda/2$ waveplate and a final PBS cube in a 360° rotation-mount. The waveplate is used to rotate the plane of polarisation into alignment with the magnetic field, and the PBS cube is aligned as a final polarisation filter.

7.2.2 Alignment

Coarse alignment was done by setting mirror and lens positions so that the two beams were at right angles, as nearly as practical. The σ^- beam path was also set to travel horizontally with respect to the optical table. Care was taken to ensure beams passed through the centres of the final lenses, to avoid deviations which change the Poynting vector with respect to the rest of the beam path.

Optical pumping

Fine adjustments were made to the setup using optical pumping diagnosed via the ion fluorescence as a measure of polarisation purity. With an ion trapped and fluorescing with the Doppler and repump lasers near resonance, we switch on the shelving and one of the deshelving (σ^- or π) lasers. The fluorescence rate of the ion is approximately proportional to population in the $P_{1/2}$ states, which is diminished by optical pumping to the $D_{5/2}$ states due to the action of the shelving laser. This is balanced by the action of the deshelving laser, which returns population to the fluorescing subspace. When the deshelving laser is perfectly polarised, some of the $D_{5/2}$ m states are uncoupled and population piles up in them, making the ion dark. Realistically, there is always some unwanted polarisation that couples to these states, allowing the ion to fluoresce at a reduced rate. The absolute sensitivity, in terms of intensity in the unwanted polarisation components, can be adjusted over several orders of magnitude by adjusting the intensity of the shelving laser, tipping the balance of shelving rate versus deshelving rate. In practice, the shelving laser intensity is adjusted until the fluorescence drops noticeably from its maximum level. Adjustments are then made to the apparatus to minimise the ion fluorescence. If the fluorescence drops near 0, sensitivity is increased by increasing the deshelving laser intensity, or decreasing the shelving laser intensity, or both.

Alignment procedure

After coarse alignment, the apparatus needs numerous fine adjustments. The magnetic field direction is trimmed by a set of coils, the beam pointing in the two different paths is adjusted via steering mirrors, and the polarisation of the beams is adjusted via rotations of waveplates. The coarse-adjusted beam pointing for the σ^- beam was taken as fixed. The polarisation in that beam and mag-

netic field trimming were then adjusted, via a slightly modified version of the optical pumping procedure described above. This fixes everything except the adjustments to the π -beam, adjusted last because it is easier to move on the optical table. The beam pointing was adjusted by “walking” the beam (turning one steering mirror to move the spot off the ion, then the other to restore it, but at a different angle) until it minimised the fluorescence in the scheme described above. The final adjustment was to rotate the final polariser and $\lambda/2$ waveplate to minimise fluorescence while still maximising transmission through the polariser.

Modified optical pumping

A modified version of optical pumping used in setting the σ^- polariser and magnetic field trimming is needed because when both settings are unknown, these two adjustments are not independent! When the $\lambda/4$ waveplate is misaligned, we have elliptically polarised light, not circular*. Consider the polarisation impurity resulting from a misalignment of the magnetic field with respect to an elliptically polarised beam. Let A_+ (A_-) be the major (minor) amplitudes of the elliptic field, and consider its projection on axes defined by the magnetic field along \hat{z} . Let the ellipse’s axes be along \hat{x} and \hat{y} when projected into the xy -plane, and let the Poynting vector of the light have an inclination (azimuth) of φ (0).




Figure 7.9: Ellipse with tilted normal

We then have,

$$x : A_+ \cos \varphi \cos \omega t \quad (7.37)$$

$$y : A_- \sin \omega t \quad (7.38)$$

$$z : A_+ \sin \varphi \cos \omega t. \quad (7.39)$$

*it is not sufficient to make the polarisation circular outside the vacuum chamber because the windows are birefringent.

By inspection, we note that if $A_- = A_+ \cos \varphi$, then the xy components are circular, so one of the σ components drops out. If we imagine looking at an ellipse at an inclination from its major axis, there is a particular angle at which it appears circular. At this angle, the beam will have π and σ^- , but no σ^+ component in its polarisation. This is a problem when optically pumping from levels with $J = \frac{5}{2}$ to $J = \frac{3}{2}$ because this polarisation will optically pump to the $J = \frac{5}{2}, m = -\frac{5}{2}$ state. Therefore, a misalignment in the $\lambda/4$ waveplate will lead us to misaligning the magnetic field to compensate the ellipticity. Without some independent means, we have no way of finding an optimum for both these adjustments.

A slight modification of the optical pumping setup allows us to overcome this by pumping instead to the $D_{3/2}$ level. In this configuration, only when the magnetic field *and* the waveplate are properly aligned will we get good optical pumping. This is because when driving $J = \frac{3}{2}$ to $J = \frac{3}{2}$, the $\pi + \sigma^-$ combination does not optically pump. Only when we have good σ^- polarisation do we expect the fluorescence to drop.

To observe this type of signal, we switch off the repump laser, and instead drive population out of $D_{3/2}$ by tuning the light-shift laser the 7 THz to be resonant with the $D_{3/2} \leftrightarrow P_{3/2}$ transition. The π -deshelve laser is polarised $\sigma^+ + \sigma^-$ by rotating the last polariser and waveplate to clear out all population in the $D_{5/2}$ level. In this configuration, we set the magnetic field coils and the σ^- beam path waveplate. We then return the light-shift laser to its original tuning, and proceed to optimise the π beam path.

7.2.3 Results

Fluorescence rate measurements are fast and continuous, making them very useful for realtime diagnostics while adjusting the apparatus. To be more quantitative, we turn to timed experiments. We transfer the ion from the ground state to the $D_{5/2}$ level with a pulse from the shelving laser, and then follow with a timed pulse from one of our deshelving lasers. The shelving pulse puts the ion in an incoherent mixture of m states, due to spontaneous decay from the $P_{3/2}$ level. The deshelving lasers deshelve different m states at different rates. When the polarisation is highly pure, some states will be quickly deshelled (with rates all of the same order, and determined by Clebsch-Gordan coefficients)—these are the states which couple to the intended polarisation. The remaining states only couple to the impurity components of the polarisation, and deshelve much more slowly. By measuring the shelved populations as a function of pulse time, we have observed these two timescales. Their ratio gives a rough figure of merit for the polarisation purity, and more precise results can be obtained by fitting the data with curves generated by rate equations.




Figure 7.10: π deshelve experiment, showing two datasets which contribute to a single fit. Around 40% of the population is deshelled in $20 \mu\text{s}$ by the π component of the deshelve laser. The remaining population decay over 10 ms is due to the polarisation impurity. Error bars are shot noise with 500 repetitions per point.

Figure 7.10 shows typical results for the π -deshelve beam. Two sets of data were taken, one for $0-20 \mu\text{s}$, and one for $0-5000 \mu\text{s}$ deshelve pulses. This allows clear resolution of the two different deshelling timescales. The fits to the data give a polarisation impurity of $\varepsilon = 3.6 \times 10^{-4}$, where the σ^- and σ^+ component intensities are both set to εI_π . The fitted intensity is $I_\pi = 0.27 I_s$, where

$$I_s = \frac{\hbar\omega^3\Gamma}{6\pi c^2} \quad (7.40)$$

is the saturation intensity of the transition. Γ is the linewidth of the upper state and ω the transition frequency. Defined this way, when $I = I_s$ the stimulated transition rate is equal to the total spontaneous decay rate of the upper level.






Figure 7.11: σ^- deshelve, short timescale, with shelving laser σ^+ . Fits to this data determine I_{σ^-} , the intensity of the primary polarisation component.

Figures 7.11 and 7.12 show typical results for the σ^- -deshelve beam. Two data sets were taken, over the timescales 0-20 μs and 0-5000 μs . The normal operating configuration for the shelving laser is nominally σ^- polarisation, as this increases efficiency of the $m = -\frac{5}{2}$ state preparation. In this configuration, a shelving pulse from the ground state leaves roughly 90% of the population in the m states which are not deshelled by the σ^- -deshelve laser. This makes it much harder to observe the short timescale dynamics, since they only occur in the upper 10% of the shelved population. Changing the shelving laser to σ^+ polarisation results in nearly all the population in m states which are quickly deshelled by the σ^- -deshelve laser, making these dynamics easier to observe (Figure 7.11). The shelving laser is returned to its normal configuration for the observation of the slow dynamics (Figure 7.12). Fits give $I_{\sigma^-} = 0.10I_s$ and $\varepsilon = 5.3 \times 10^{-4}$.

The observed polarisation impurities of $\lesssim 10^{-3}$ are consistent with the extinction ratio (T_s/T_p) for the polarising beam cubes used in our setup, quoted at 10^{-3} . Our group has also previously observed some distortion of laser po-




Figure 7.12: σ^- deshelve, long timescale, with shelving laser σ^- . Fixing the overall laser intensity to the value from fits to Figure 7.11 allow determination of the impurity, ε . Fast deshelving (i.e. the elbow between first and second data point) is due to the intensity of the primary polarisation component.

larisation as beams pass through vacuum system view ports. The polarisation purity achieved is sufficient to observe partial-collapse and uncollapse quantum processes in our system.

7.3 Magnetic driving

Coherent control over the state of our system is achieved by driving magnetic dipole transitions between adjacent m levels.

7.3.1 Magnetic coil and control electronics

The oscillating dipole is generated by a modest field coil ~ 70 mm in diameter with 10 turns. The coil is mounted against the front view port, putting the ion approximately 50 mm from its centre along its axis. The coil is connected in

series with an air-gap variable capacitor, adjustable over ~ 10 – 225 pF. This resonant circuit tunes via the capacitor between 2.5 and 6 MHz, and has an in situ quality factor of 10.3 measured with the resonance at 4.9 MHz.

A pair of Stanford Research Systems DS345 signal generators generate the waveforms applied to the coils. The outputs of the generators are fed to an RF switch (Mini Circuits ZASWA-2-50DR), where a digital output from the Laser Control Unit (LCU) selects the source. This is then cascaded to another RF switch where a second LCU output is used to pulse the signal on and off during a timed sequence. A Frankonia FLL-25 amplifier steps up the source signal +47 dB before it is applied to the resonant coil circuit.

The signal generators are locked to the same clock source. This allows them to maintain phase coherence when set to the same frequency. However, the relative phase between the generators changes whenever the generators' frequencies are different. The sync outputs of the generators are connected to inputs of a digital oscilloscope, which measures the relative phase to display on screen. The relative phase can be set to any desired value by adjusting the absolute phase offset on one of the generators. One of the generators is connected to the experimental control computer via an RS-232 interface, which allows control over the frequency, amplitude and phase during an experimental run.



7.3.2 Results

Figure 7.13 shows typical data for magnetic driving in the unshifted $D_{5/2}$ subspace (that is, with Zeeman splitting but no light-shift). State preparation is to the $m = -\frac{5}{2}$ sublevel, as described in Section 7.1.1, and readout is accomplished by using long pulses of the π - and σ^- -deshelve lasers to map all $m \neq -\frac{5}{2}$ states to the bright subspace. Fitting is via the analytic theory developed in the first part of Section 7.1.2, with the Rabi frequency Ω_B , and the detuning δ determined by the fits. The fitting program multiplies by a prefactor A to account for state preparation errors.




Figure 7.13: Magnetic driving in the unshifted $D_{5/2}$ subspace, time scan of RF pulse length versus shelved population, corresponding to population in the $m = -5/2$ state. Fitted parameters: $\Omega_B = 2\pi \times 36.8$ kHz, $\delta = 2\pi \times 24.8$ kHz, and $A = 0.89$.

Figure 7.14 shows typical data for scanning the frequency of the driving pulse (nominally a π rotation on resonance). Preparation and readout are as before, with fitted parameters A , Ω_B , and the Larmor frequency, ω_0 .




Figure 7.14: Frequency scan in the unshifted $D_{5/2}$ subspace, versus population in the $m = -5/2$ state. The pulse length $t = 12.6 \mu\text{s}$. Fitted parameters: $\Omega_B = 2\pi \times 37.4 \text{ kHz}$, $\omega_0 = 2\pi \times 2746.6 \text{ kHz}$, and $A = 0.937$.

Both figures show agreement with the analytic theory, but there appears to be some structure left in the residuals, particularly near 0 shelved population. This is likely due to the “leaky” transfer (Section 7.1.3) to the ground state. Figure 7.13 also shows evidence of decoherence, since the peaks in the shelved population appear to decay over the scan. We discuss this type of decoherence for the light-shifted driving in Section 7.4.3. The lineshapes in the $J = \frac{5}{2}$ subspace are significantly different from the familiar ones for the two level system, but some of the general features, such as periodicity of the time scan with frequency Ω_B , still hold true.

7.4 The $D_{5/2}$ qubit

We now discuss the remaining functional pieces to this experiment: the application of the light-shift laser to break the degeneracy in transition frequency, and implementation of “dynamic decoupling” pulse-sequences to protect the coherence of the qubit.

7.4.1 Application of light shifts

The light-shift laser beam travels down the same fibre as the σ^- -deshelving laser beam, so once the optical path for that beam is optimised, we require no further adjustments to set the light-shift laser beam. We wish to characterise the light-shift so that we may optimise its detuning. In the laboratory, we directly measure laser power, but may infer the intensity at the ion by measuring the size of the light shift and using the theory set out Section 7.1.2. The conversion factor between the power and intensity should remain constant so long as the beam optics are not altered and the beam is kept well aligned on the trapped ion. This conversion factor is an effective area of the laser spot, $I = P/A_{\text{eff}}$.

For brevity, in the remainder of this section, we shall denote the $D_{5/2}$ sub-levels by $|m\rangle$. Recalling that when the light-shift laser is well polarised, the primary qubit transition ($|-\frac{5}{2}\rangle \rightarrow |-\frac{3}{2}\rangle$) is unshifted, we characterise the light-shift by measuring the shift in the $|-\frac{3}{2}\rangle \rightarrow |-\frac{1}{2}\rangle$ transition. After preparing $|-\frac{5}{2}\rangle$, we apply a π -pulse on the qubit transition frequency to transfer to $|-\frac{3}{2}\rangle$. We then apply a second pulse whose frequency we scan. Using the matrix elements of \hat{J}_x from (7.6), the Rabi frequencies for the qubit and $|-\frac{3}{2}\rangle \rightarrow |-\frac{1}{2}\rangle$ transitions are $\sqrt{5}\Omega_B$ and $2\sqrt{2}\Omega_B$ respectively.



To read out the final state, we apply only the σ^- -deselve laser, which couples to the $m \geq -\frac{1}{2}$ states. Thus any population transferred on the $|-\frac{3}{2}\rangle \rightarrow |-\frac{1}{2}\rangle$ transition will be transferred to the bright subspace. Data from a typical run of this experiment are presented in Figure 7.15. The theory curve is generated by solving the optical Bloch equations of (7.28), and fitting the applied light-shift Δ_{LS} , the Rabi frequency Ω_B , and a normalisation to take into account state preparation errors. The residuals indicate that our theory only crudely accounts for the observed lineshape. This could be because of the presence of decoherence processes, such as fluctuating power in the light-shift laser. However, the fits still determine the line centre and therefore Δ_{LS} with high precision, so these discrepancies have not been pursued.




Figure 7.15: $|-\frac{3}{2}\rangle \rightarrow |-\frac{1}{2}\rangle$ transition frequency probe, with fit generated by optical Bloch equation solver for the light-shifted $D_{5/2}$ subspace. The fit approximates the lineshape, but the centre frequency can be determined with high confidence.

Table 7.3 summarises the results of varying the light-shift laser power and detuning, including the primary fitted parameter A_{eff} . These results were used

to obtaining a representative value of $A_{\text{eff}} = 10 \times 10^{-10} \text{ m}^2$ for a well aligned laser. Together with the maximum available power of 15 mW, this allows calculation of the optimal detuning for the light shift laser using the theory laid out in Section 7.1.2. Assuming $\Omega_{01} = 2\pi \times 50 \text{ kHz}$, and taking into account the full experimental pulse sequence described in Section 8.1, we obtain $\delta_{LS} = 2\pi \times 110 \text{ GHz}$.

P (mW)	δ_{LS} ($2\pi \cdot \text{GHz}$)	Δ_{LS} ($2\pi \cdot \text{kHz}$)	A_{eff} (10^{-10} m^2)
1.0	11.2	442	9.8
1.0	14.9	312	10.4
1.6	14.9	572	9.2
3.3	62.7	258	10.0
3.3	41.8	429	9.0
3.3	20.9	885	8.6

Table 7.3: Measured light-shifts with varying power and detuning. A_{eff} is fitted from the theory in Section 7.1.2






Figure 7.16: $D_{5/2}$ qubit timescan with fit to $Ae^{-t/\tau} \left(1 - \frac{\Omega^2}{W} \sin^2 \frac{Wt}{2}\right) + B$. Fitted parameters: $A = 0.75$, $B = 0.12$, $\tau = 1490 \mu s$, $\Omega = 2\pi \times 51.2 \text{ kHz}$, and $W = \sqrt{\Omega^2 + \delta^2}$ with $\delta = 0$.

7.4.2 Qubit manipulations

Figures 7.16 and 7.17 show typical results of coherent manipulations of the $\{|-\frac{5}{2}\rangle, |-\frac{3}{2}\rangle\}$ subspace with RF driving pulses, scanning the time and frequency respectively. They are fitted to a simple two-level system analytic model (e.g. [16]) multiplied by a factor to account for state preparation and readout errors. Agreement to the theory indicates that the qubit levels are well isolated from the rest of the $D_{5/2}$ subspace.




Figure 7.17: $D_{5/2}$ qubit frequency scan with fit $A \left(1 - \frac{\Omega^2}{W^2} \sin^2 \frac{Wt}{2} \right) + B$, with $W = \sqrt{\Omega^2 + (2\pi f - \omega_B)^2}$ and $t = 12.6 \mu\text{s}$. Fitted parameters: $A = 0.77$, $B = 0.10$, $\Omega = 2\pi \times 50.1 \text{ kHz}$, $\omega_B = 2\pi \times 2701 \text{ kHz}$.

7.4.3 Decoherence and protection sequences

Since the $D_{5/2}$ qubit splitting is caused by the Zeeman effect, fluctuations in the \mathbf{B} field couple to the qubit and can cause decoherence. In the RF domain spontaneous radiative decay is negligible on the timescale of our experiments, so we are primarily concerned with dephasing due to fluctuations with $\omega \ll \omega_B$.

One important source of noise is the 50 Hz mains lines powering the equipment in our lab. Magnetic shielding can reduce this noise, but a simpler solution is to take advantage of the fact that the noise is relatively stable and periodic. A triggering device monitors the mains and outputs a TTL-compatible pulse on the positive zero-crossing of the mains voltage; this is fed into the laser control unit to start subsequent experimental runs from the same time in the mains cycle of 20 ms. By triggering two shots of a frequency scan experiment from a single mains-cycle trigger with a variable delay—that is to say: trigger, prepare, manipulate, readout, delay, re-prepare, manipulate, readout—and comparing the difference in the centre frequency, we extract the frequency shift as a function of location in the mains cycle.




Figure 7.18: Magnetic field shifts measured by shift in $f_0 = \omega_B/2\pi$ versus delay between shots of the experiment (see text). All runs are mains-triggered. The variation in magnetic field at 50 Hz is clear (field sensitivity is $\mu_B g_J / \hbar = 1.68$ kHz/mG).

Figure 7.18 shows the variation of the magnetic field (scaled by the gyromagnetic ratio $\mu_B g_J$) over the 50 Hz mains cycle. For experimental sequences whose overall length is $\ll 20$ ms, we may approximate a stationary field by triggering our experiments from the mains cycle. Then, these variations, if stable, will be reduced to a constant phase shift. Experimental sequences which have overall lengths comparable to 20 ms may be sensitive to the variation in the field. For the experiments described in Section 8.1, a judicious choice of trigger point was made, in an attempt to find a place in the cycle when the field was at a stationary point. However, the consequent improvement in the data was marginal.

We now examine more elaborate qubit manipulations designed to be robust against such fluctuations.

The partial- and uncollapse experiments described in Chapter 8 require maintaining qubit coherence during one or two bright/dark readouts (Section 7.1.3), which are the primary limitation on speed of the procedure. The time required to read out with sufficient fidelity was determined experimentally by alternately observing a bright ion then a dark ion by pulsing the repumping laser on and off. Photon counts were recorded for each bright and each dark observation, and a threshold chosen which minimised the error. The error is then the fraction of observations misclassified by the threshold chosen. The observation time was varied in order to choose the smallest time which gave less than 1% error, which for typical bright and dark count rates of $R_B = 31$ kHz and $R_D = 0.4$ kHz (cf. [10]) was $t_b = 300$ μ s.

Even with triggering on the mains cycle, significant decoherence of the qubit occurred in 300 μ s. Figure 7.16, for example, shows a noticeable decay in contrast near the end of the 200 μ s scan. By applying a sequence of RF driving pulses (a set of techniques called *dynamic decoupling* generally), the coherence of a qubit can be extended to many times its unprotected coherence time. We investigated two such sequences: firstly, the *Carr-Purcell-Meiboom-Gill* (CPMG) sequence [17] and what has become known as the *Uhrig Dynamic Decoupling* (UDD) sequence [63].

CPMG uses n equal-spaced π -pulses, and is a straightforward extension of spin-echo. The pulse times are given by [61]

$$\alpha_i = \frac{i - \frac{1}{2}}{n} \quad (7.41)$$

where the times are $t_i = \alpha_i \tau$ and τ the time to protect the qubit.

The UDD pulse sequence uses n unequal pulses given by

$$\alpha_i = \sin^2 \left(\frac{\pi}{2} \frac{i}{n+1} \right) \quad (7.42)$$

and can be shown [61] to cancel extraneous phase caused by any perturbation of the magnetic field which is a polynomial in time of order $< n - 1$.

The pulse sequences were tested using the $D_{5/2}$ qubit in a Ramsey experiment where the phase of the second $\pi/2$ -pulse was scanned to reveal a set of fringes. The contrast of the fringes versus τ , the time between the $\pi/2$ -pulses, gives the coherence properties of the qubit. Here we are concerned with the contrast at $\tau = 300$ μ s: the duration of a readout operation. Experiments revealed no significant difference between the performance of the UDD and CPMG pulse sequences on such timescales. A sequence of 3 π -pulses greatly improved the contrast over both 0- and 1- π -pulse sequences. The results of a study [61] in a similar trap system were consistent with the observation that UDD and CPMG

perform similarly for $\tau = 300 \mu\text{s}$, and indicated that moving to sequences with greater than 3 π -pulses would be unlikely to yield additional improvements in contrast on the timescale of a readout operation.

Our laser control unit cannot issue additional timed switching commands during another longer pulse, so it is not possible to have a continuous readout while simultaneously issuing the CPMG or UDD protection pulses. The readout was therefore broken into several bins such that the total photon counting time equalled the desired $t_b = 300 \mu\text{s}$, with protection π -pulses interleaved. This results in a somewhat longer sequence due to the mandatory gaps between pulses during which the laser control unit TTL-outputs are reprogrammed, but it is still well within the observed coherence time for 3-pulse protection sequences. In the main experiments, the CPMG pulse sequence was chosen since the equal-spaced pulses make insertion of readout photon counting bins more simple to implement.

8

Uncollapse

We present the results of an experimental demonstration of qubit uncollapse of the type introduced in Section 3.4. The uncollapse procedure was characterised by quantum process tomography. An ideal partial collapse followed by uncollapse procedure is equivalent to the identity operation. We also examine the initial claim about the uncollapse procedure: that we are able to undo a non-unitary operation on a qubit, with heralded but probabilistic success. This was evaluated by observing the partial collapse of the qubit wavefunction after the first partial measurement.

8.1 Pulse sequence

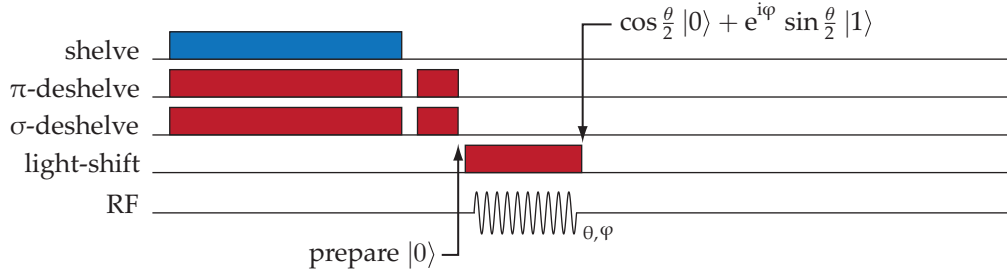
A partial collapse experiment can be broken down into three parts: state preparation, partial measurement, and a tomography step to measure the resultant state. Figure 8.1 shows sequence diagrams for these parts. State preparation proceeds as described previously, with one additional step: a “clear-out” pulse following the optical pumping to $|0\rangle$ ($D_{5/2}$, $m = -\frac{5}{2}$). In this pulse, the polarised deshelve lasers are turned on for a short time to help clear out any population in the $m \neq -\frac{5}{2}$ levels. This pulse also deshelves some of the population in $|0\rangle$, but the measurement-selected ensemble we construct from the results will filter out any experiments where this happens. Following the preparation of $|0\rangle$ we apply an RF pulse to create the desired superposition

$$|\psi\rangle = \cos\frac{\theta}{2}|0\rangle + e^{i\varphi}\sin\frac{\theta}{2}|1\rangle \quad (8.1)$$


with $\theta = \Omega_{01}\tau$ and φ the RF phase.

Partial measurement is accomplished with pulses from the deshelve lasers. This pulse is timed to have a “strength” or probability p of deshelve $|1\rangle$ ($D_{5/2}$, $m = -\frac{3}{2}$), with pulse times determined experimentally or via solving rate equations for the fitted deshelve laser intensities. The π -deshelve laser primarily accomplishes this, with the σ^- -deshelve simply helping clear out any population

State Preparation



Detection & Protection



Tomography




Figure 8.1: Partial- and uncollapse sequence diagrams. See text for explanation of full sequence. The tomography sequence is implemented with the RF pulse having y -phase, x -phase, or being omitted to measure \hat{X} , \hat{Y} , and \hat{Z} respectively. The repump laser is left on continuously and is omitted from these sequence diagrams. Pulse widths are shown for readability and do not accurately represent the timing.

off-resonantly driven into other m states. The Doppler laser and PMT counter gate are then switched on twice, for a total of $t_b = 300 \mu\text{s}$, while interspersed RF pulses ($\theta = \pi$, $\varphi = 0$) protect the coherence of the $D_{5/2}$ qubit. This part of the sequence is denoted “detection and protection.” Note that since we have an odd number of π -pulses, the net result of this sequence is a single π -rotation about the y -axis.

Lastly, the state following the partial collapse is measured by tomography in the Bloch basis. An RF pulse rotates the desired operator basis to the computation ($|0\rangle, |1\rangle$) basis. To reduce the number of required RF phases, we actually

measure $-\hat{X}, +\hat{Y}, +\hat{Z}$ by applying $(\theta = \pi/2, \varphi = 0)$, $(\theta = \pi/2, \varphi = -\pi/2)$, and no RF pulse, respectively. A projective ($p = 1$) measurement is then performed.

To perform an uncollapse experiment, a second detection and protection sequence is inserted immediately following the first one. Another advantage of the CPMG protection sequence is that stringing two n -pulse sequences in succession has exactly the same form as a $2n$ -pulse sequence of length 2τ (something which is not true of UDD). This is followed by exactly the same tomography sequence.

Partial Collapse	Uncollapse
State Prep	State Prep
Detection & Protection	Detection & Protection
Tomography	Detection & Protection Tomography

Figure 8.1 does not show accurate relative timing, because the readout pulses are so much longer than other pulses they would make the diagram unreadable. An example pulse sequence with timing information is compiled in Table B.1.

The experiments were conducted with typical Zeeman splitting $\omega_0 = 2\pi \times 2700$ kHz (1.9 G), 15 mW light-shift laser power, and light-shift laser detuning $\delta_{LS} = 2\pi \times 110$ GHz. This results in a light shift of $\Delta_{LS} = 670$ kHz. Magnetic transitions were driven at a Rabi frequency $\Omega_{01} = 2\pi \times 52$ kHz.

8.2 Process tomography

A quantum process is a completely positive map from input states (which, in this case are given by density operators) to output states in the same Hilbert space. In Section 3.4, we gave theoretical descriptions of two different versions of the partial collapse and uncollapse processes. We denote the ideal partial collapse process \mathcal{P}_p and the ideal uncollapse process \mathcal{U}_p . We also introduced the “leaky” transfer of population during a partial measurement. When these leaks are included, we denote the partial collapse \mathcal{P}'_p and the uncollapse \mathcal{U}'_p . Finally, we denote the actual processes realised in our experiments $\mathcal{P}_p^{\text{expt}}$ and $\mathcal{U}_p^{\text{expt}}$.

The ideal partial collapse process, \mathcal{P}_p , is non-linear* because of the normalisation associated with the measurement-selected ensemble (Section 3.3), but the ideal uncollapse, \mathcal{U}_p , by virtue of the fact that it reduces to the identity, is linear. Experimental errors and the “leaky transfer” will result in non-linearity for the \mathcal{U}'_p and $\mathcal{U}_p^{\text{expt}}$ processes as well, but we proceed with analysis under the assumption of linearity, which will be approximately true when errors are small.

*i.e. it is not characterised by a linear transformation from the input Hilbert space to the output Hilbert space

We do this for two reasons. Firstly, in order to fully quantify a linear process, it is sufficient to determine the action of the quantum process on a set of input states which span the space. The Hilbert space for operators on a single qubit has 4 dimensions (unlike the *ket*-space, which has only 2), therefore we need to perform our experiment on 4 suitably chosen input states. Secondly, we wish to define some quantitative measure of how well we are able to uncollapse. There exists a straightforward measure of the fidelity of a linear process which quantifies how close a noisy process comes to implementing the ideal. We shall return to the definition of fidelity later in the chapter.

To quantify the effect of $\mathcal{P}_p^{\text{expt}}$ and $\mathcal{U}_p^{\text{expt}}$ on each of the 4 input states we need to perform quantum *state* tomography on each of the 4 output states. As discussed in Section 3.3, for a single qubit, we may obtain this information by measuring the expectation values $\langle \hat{X} \rangle$, $\langle \hat{Y} \rangle$, $\langle \hat{Z} \rangle$, for a total of $4 \times 3 = 12$ measurements.

It is noteworthy that we need 4 linearly independent states to span the space of density operators when we may parametrise it with just 3 real numbers. This is because the complete description of a quantum state is a ray in the Hilbert space, not a particular vector. The convention of working with normalised state vectors is a mathematical convenience, but the set of normalised state vectors is *not* a linear subspace of the full Hilbert space—in the same sense that a unit sphere is not a subspace of 3-dimensional Euclidean space. So even though we can parametrise the set of density matrices with just three components, our parametrisation will not be linear. We require 4 linearly independent vectors to span the space, but can obtain the output states by measuring 3 expectation values.

I chose the pure states $|0\rangle\langle 0|$, $|+x\rangle\langle +x|$, $|-y\rangle\langle -y|$, and $|1\rangle\langle 1|$ as inputs, where

$$|\pm x\rangle = \frac{|0\rangle \pm |1\rangle}{\sqrt{2}} \quad \text{and} \quad (8.2)$$

$$|\pm y\rangle = \frac{|0\rangle \pm i|1\rangle}{\sqrt{2}}. \quad (8.3)$$

8.3 Results

Each of the 12 measurements for quantum process tomography were measured by repeating the pulse sequence detailed above several thousand times (Appendix B.2 tabulates the ensemble sizes in detail) and recording the number of counted photons in each experimental readout period (2 readouts in the case of partial collapse, 3 in the case of uncollapse). At the end of each experimental sequence, a reset sequence consisting of alternating π -pulses and strong deshelve pulses was applied to reset the ion to the ground state for the next experiment. At the end of the reset sequence, a final readout period verifies that the ion is

successfully re-prepared in its ground state, and that it remains trapped and cold. Momentary heating of the ion due to background gas collisions or temporary laser frequency fluctuations cause the ion to go dark during these check bins. These shots of the experiment are removed from the final ensemble in post-processing. This also allows for monitoring of ion loss during data taking.

For partial collapse experiments, there are 3 photon count bins: C, B, and A which are the partial measurement result, tomography result, and check result respectively. All counts are converted to bright or dark results by the threshold method, and a measurement-selected ensemble is created from the shots where both the C bin was dark and the A bin was bright. For uncollapse experiments, there are 4 photon count bins: D, C, B, and A, with the D bin the additional partial measurement result. The ensemble in this case consists of shots where both D and C bins were dark and A was bright. The statistics of the B bins in the measurement-selected ensemble are the results of the experiment.

Let us consider the main experimental limitations which contribute systematic errors to our results.

The deshelving process includes small population leaks back to the $D_{5/2}$ subspace, most notably an $\epsilon = 0.036$ transfer to $|0\rangle$ (Section 7.1.3). These effects are not significant for $p \approx 0$, but become very important when p is near 1.

The main qubit transition frequency has linear dependence on the static magnetic field, so fluctuations contribute strongly to phase decoherence. This decoherence is greatly suppressed by triggering experiments on the mains cycle and the introduction of dynamic decoupling sequences (Section 7.4.3), but is still significant at the level of a few percent.

Magnetic driving on the qubit resonance will also couple to transitions which are detuned by the light shift. The most important of these is the $m = -\frac{3}{2} \rightarrow -\frac{1}{2}$ transition. In choosing the detuning of the light-shift laser there is a trade-off between the off-resonant magnetic process and off-resonant optical scattering by the light-shift laser. The latter is suppressed by the polarisation purity of laser beam, and the laser detuning is chosen to minimise the sum of these errors (Section 7.1.2). Under typical experimental conditions we estimate the sum of these probabilities is $\sim 1\%$.

Finally, long term drifts in the intensity of the deshelving laser beam alters the strength of the partial collapse, p . Experimental runs for quantum process tomography take of order 1 hour for each nominal p value, so drifts can corrupt the process tomography results. These drifts are primarily the result of polarisation drift in the single-mode fibres used to deliver the beams to the experiment. The polarising optics convert polarisation changes into intensity changes.

The most straightforward way to examine the data is to plot the output states as points on the Bloch sphere, with data taken directly from the $\langle \hat{X} \rangle$, $\langle \hat{Y} \rangle$, $\langle \hat{Z} \rangle$ measurements. Figure 8.2 shows the results of the partial collapse experiments.

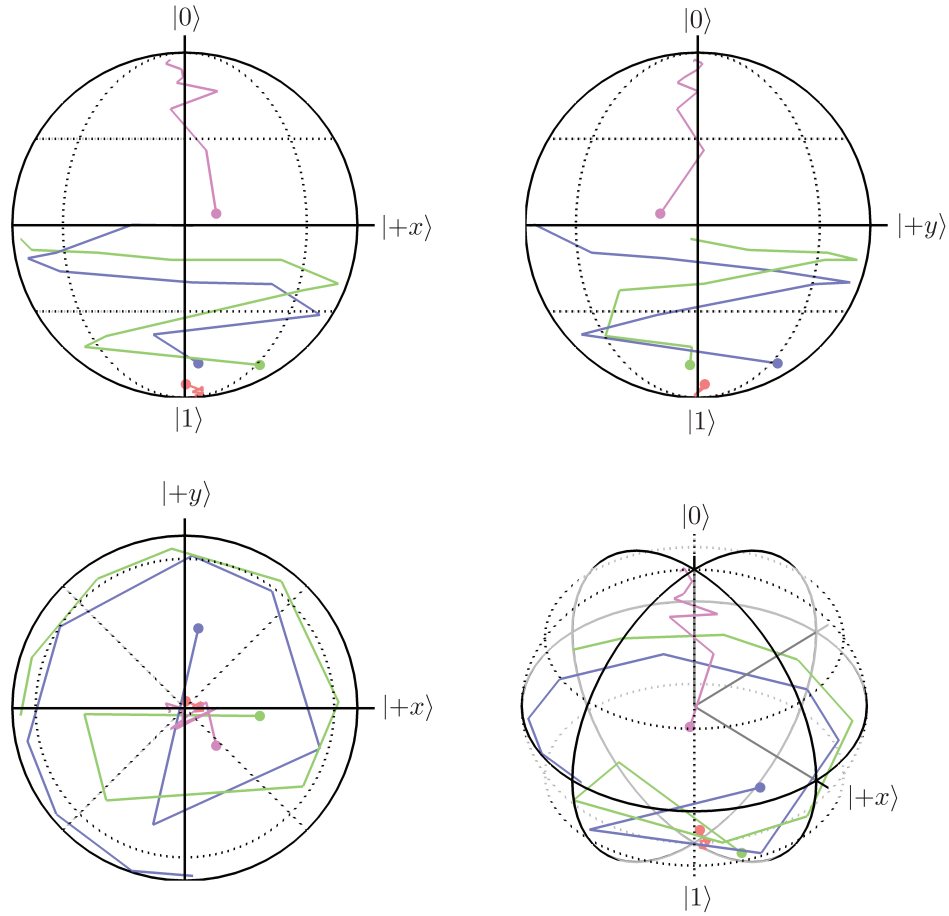


Figure 8.2: Partial-collapse results as Bloch vectors. Clockwise from lower left we have xy -, xz -, yz -projection and isometric views of the Bloch sphere. Lines show increasing values of p , with highest value at dotted head. Experimentally measured p values are [0.02, 0.10, 0.22, 0.31, 0.42, 0.55, 0.62, 0.75, 0.81, 0.93]. Prepared states are $|0\rangle$, $|1\rangle$, $|+x\rangle$, and $|+y\rangle$. Note this is *after* the effect of the π -rotation in the CPMG protection sequence.

These plots are intended to be reminiscent of the deformed Bloch spheres we encountered in Section 3.4.5. As p is increased, we see the $|+x\rangle$ and $|+y\rangle$ preparations spiralling down toward the $|1\rangle$ pole. These results are taken *after* the CPMG protection sequence, which results in an overall π -rotation, so we interpret this data as the expected collapse of the wavefunction toward the $|0\rangle$ state. The spiralling about the z -axis is due to perturbations of the qubit energy caused by the deshelling lasers (near resonant light-shifts). This perturbation causes the qubit to precess at a different frequency than our reference frame set by the RF oscillator used to drive transitions.

Figure 8.3 shows the results of the uncollapse experiments. The most striking difference in the results is the suppression of the rotations about the z -axis.

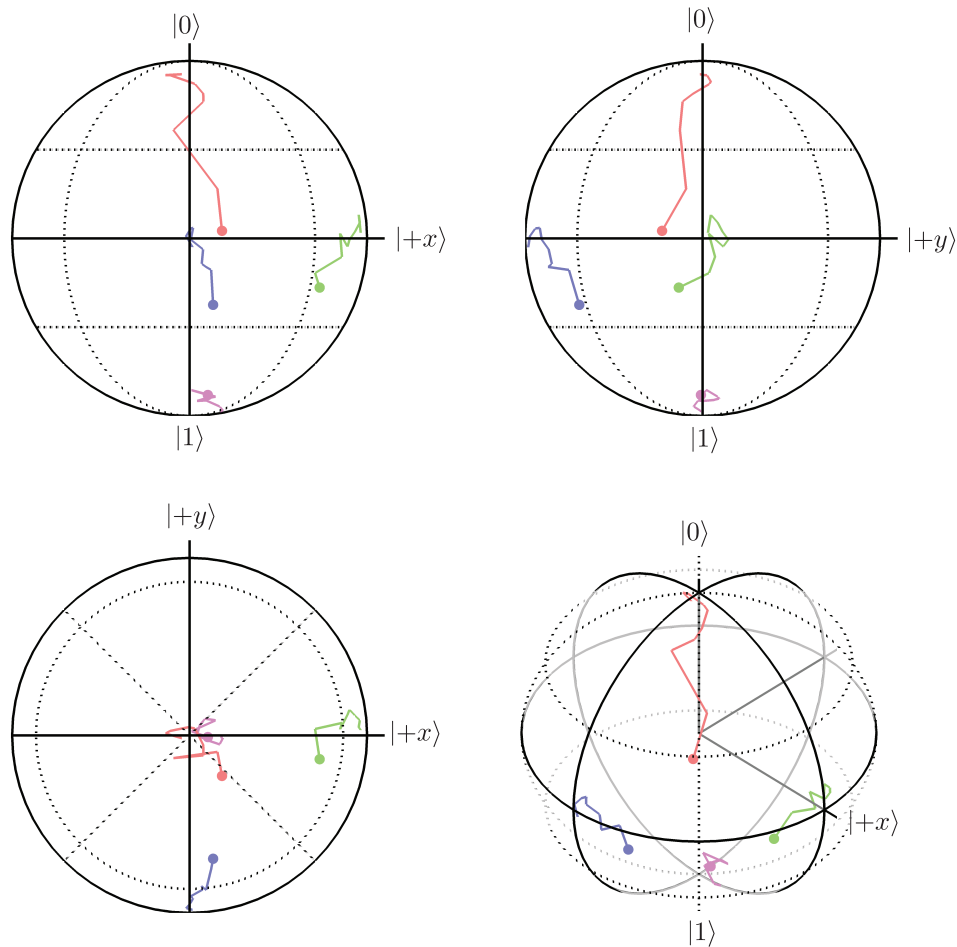


Figure 8.3: Uncollapse results as Bloch vectors. Clockwise from lower left we have xy -, xz -, yz -projection and isometric views of the Bloch sphere. Lines show increasing values of p , with highest value at dotted head. Experimentally measured p values are [0.01, 0.09, 0.22, 0.31, 0.43, 0.56, 0.64, 0.75, 0.83, 0.94] Prepared states are $|0\rangle$, $|1\rangle$, $|+x\rangle$, and $|+y\rangle$.

However, this is not due to any effect specific to the uncollapse procedure—the rotations we saw in the partial-collapse were unitary, and performing the partial collapse again with the same strength simply unwinds the extra precession in a spin-echo effect. The collapse of the $|+x\rangle$ and $|+y\rangle$ states toward to the lower pole is greatly suppressed. The behavior of the state which starts near $|0\rangle$ ($|0\rangle$ for the uncollapse, but $|1\rangle$ for the partial collapse) is very similar. Both these effects are almost entirely due to the finite branching (leaky transfer) back to the qubit from the $P_{3/2}$ level discussed in Section 7.1.3. Neither partial collapse nor uncollapse will have any effect on $|0\rangle$ or $|1\rangle$ states in the ideal case. When the branching ratio $\epsilon \ll (1 - p)$, the leakage effect is small, but when $\epsilon \approx (1 - p)$, it dominates the dynamics in the measurement-selected ensemble.

We might wonder why the finite branching ratio effect is not symmetric with respect to the $|0\rangle$ and $|1\rangle$ states, since we apply the partial measurement twice with a π -rotation in between. Consider the effect of a partial measurement when p is close to 1. When in $|0\rangle$ this has no effect, but for $|1\rangle$, we transfer most of the population to the bright state, but we also transfer $p\epsilon$ to $|0\rangle$, which will be comparable to the $(1-p)$ that we leave in $|1\rangle$. Now we create a selected ensemble in which the ion remained dark. For the $|0\rangle$ preparation, this again has no effect, but for the $|1\rangle$, this renormalises the populations in the qubit levels. After a π -pulse, we see that our state going into the second partial-measurement is not simply the reverse of what we had to start with—the state that began in $|1\rangle$ is either nearly depolarised or else starting to be repolarised the opposite direction!

The Bloch representation is a good tool for getting an intuitive, geometric understanding of the effect of our processes, but we would like a quantitative measure of the fidelity of our process compared to ideal processes, or to the theoretical limits set by the leaky transfer.

The fidelity of two quantum states is (using the convention of [42])

$$F(\rho_o, \rho_i) \equiv \text{tr} \sqrt{\rho_i^{1/2} \rho_o \rho_i^{1/2}}. \quad (8.4)$$

Fidelity is a measure of the “distance” between two quantum states. In the case of pure states, $\rho_i = |\psi_i\rangle\langle\psi_i|$ and $\rho_o = |\psi_o\rangle\langle\psi_o|$, the fidelity is the absolute value of the inner product, $|\langle\psi_o|\psi_i\rangle|$.

We may examine the uncollapse experimental data by considering the fidelity of the states measured after the uncollapse with respect to the inputs, $F(\mathcal{U}_p^{\text{expt}}(\rho), \rho)$. The fidelity is plotted versus partial collapse strength p in Figure 8.4 for each of the 4 input state preparations.

The fidelity is affected by all the errors in our experiment: state-preparation, readout, qubit decoherence, as well as errors in the uncollapse procedure itself. It will be helpful to consider only the errors introduced by the application of the uncollapse procedure. For this we compute the fidelity using the measured $p = 0$ states as the inputs. These are the experimentally determined quantum states when the deshelving lasers are not switched on during the experiment. They represent the best estimate of the states actually adopted by the system in the presence of systematic errors in the state preparation, decoherence, the CPMG protection sequence, and the final readout. The fidelity computed from these experimentally determined inputs should be more representative of errors in the uncollapse procedure itself. Note that a consequence of using the $p = 0$ states as inputs is that, by definition, the fidelity equals 1 for these states.

We also plot the fidelity of the 4 theoretical processes. The lower black line represents the state fidelity of the ideal partial collapse, $F(\mathcal{P}_p(\rho), \rho)$. This quantifies effect of the first partial transfer if we perform the fluorescence detection and select only the ions which were dark. An ideal uncollapse restores the fi-

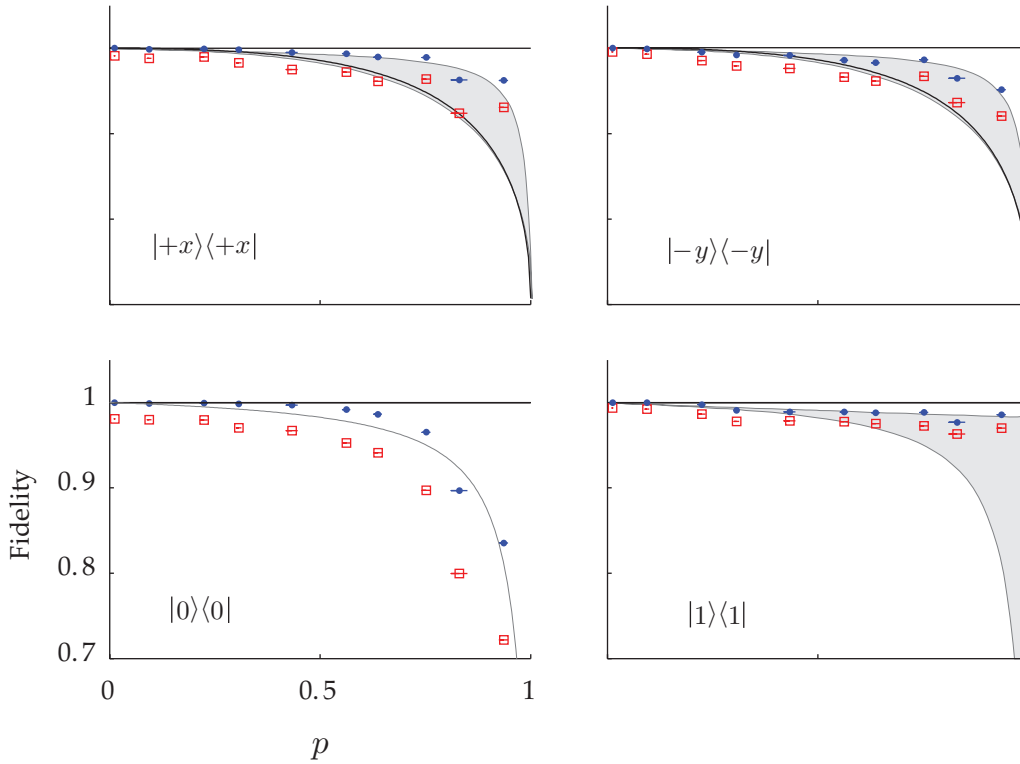


Figure 8.4: Uncollapse experimental state fidelity. (\square) assume ideal input states; (\bullet) use experimentally determined input states (see text). Values of p are experimentally determined, and error bars indicate the standard deviation of experiments taken at the same nominal p . The upper (lower) black line indicates expected fidelity from an ideal uncollapse (partial collapse) procedure. The upper (lower) gray line indicates expected fidelity from a leaky uncollapse (partial collapse) procedure. The area between the leaky procedure curves is shaded, and in the lower left, the shading is omitted because the leaky partial collapse has a fidelity of 1.

delity of the state, $F(\mathcal{U}_p(\rho), \rho) = 1$ for $p < 1$, indicated by the upper black lines. The gray lines represent the state fidelity of the leaky processes: \mathcal{P}'_p and \mathcal{U}'_p . The gray shaded areas indicate the scope of improvement achievable by using the leaky uncollapse.

Note that for the $|0\rangle\langle 0|$ and $|1\rangle\langle 1|$ input states, neither the ideal partial collapse, \mathcal{P}_p , nor the ideal uncollapse, \mathcal{U}_p , affect the state. However, the leaky partial collapse, \mathcal{P}'_p , reduces the fidelity of the $|1\rangle\langle 1|$ state, and the leaky uncollapse, \mathcal{U}'_p is able to increase the fidelity above this level. For the $|0\rangle\langle 0|$ state, \mathcal{P}'_p does not affect the state, but \mathcal{U}'_p actually introduces errors (the gray line printed without shading)! This is because the leaky transfer occurs in the $|1\rangle \rightarrow |0\rangle$ sense, so the $|0\rangle\langle 0|$ is unaffected.

For the $|+x\rangle\langle +x|$ and $|-y\rangle\langle -y|$ states, the relative improvement in fidelity increases with increasing p for the ideal case. However, in the leaky case, the improvement in fidelity is only marginal for p near 1.

When considering the fidelity with respect to ideal inputs, our uncollapse

procedure, $\mathcal{U}_p^{\text{expt}}$, improves the state (i.e. results in fidelity above that expected for the leaky partial collapse) only for the higher values of p , with the exception of the $|0\rangle\langle 0|$ state (which, as mentioned, can only be damaged by the uncollapse). Considering the fidelity with respect to experimentally determined input states, our uncollapse procedure improves the state, with more dramatic improvement at high p , again with the exception of the $|0\rangle\langle 0|$ state. Qualitatively, the fidelity with respect to experimentally determined input states tracks the leaky uncollapse fidelity very well.

Figure 8.4 gives the fidelity of 4 specific states, but we now consider a description of the process itself—including a measure of the overall *process fidelity*. As mentioned above, we do this by measuring the effect of the process on 4 linearly independent states. We may represent the action of a linear quantum process by expanding it in an operator sum representation

$$\mathcal{E}(\rho) = \sum_k \hat{E}_k \rho \hat{E}_k^\dagger \quad (8.5)$$

where $\{\hat{E}_k\}$ are a set of operation elements and satisfy a completeness relation

$$\sum_k \hat{E}_k \hat{E}_k^\dagger = I. \quad (8.6)$$

Instead of using a set of operation elements unique to each quantum process, it is often convenient to represent the quantum process with respect to a fixed set of operators (following the convention of [42]), \tilde{E}_k which span the space of operators on ρ . We then expand each operation element

$$\hat{E}_k = \sum_j e_{kj} \tilde{E}_j \quad (8.7)$$

and rewrite the quantum process as

$$\mathcal{E}(\rho) = \sum_{mn} \tilde{E}_m \rho \tilde{E}_n^\dagger \chi_{mn}. \quad (8.8)$$

χ is a matrix which represents the quantum process with respect to the fixed basis set of operators, and its matrix elements are given by

$$\chi_{mn} = \sum_k e_{km} e_{kn}^*. \quad (8.9)$$

The χ -matrix serves as a compact representation of the action of the quantum process.

To compute χ from the experimental data, we first express each of the 4 output density matrices as a linear combination of the 4 input matrices—this is a 4×4 matrix, λ , which expresses the action of the process as a linear transformation using the input states as the basis. We also compute, for each combination

of operators $\{\tilde{E}_m, \tilde{E}_n\}$, a matrix β^{mn} , which expresses the operators' transformation in the same way. Thus, we have a matrix λ , which we wish to express as a linear combination of β^{mn} matrices. The coefficients in this linear combination are the matrix elements of χ . Some shape manipulations allow us to re-express λ as a column vector, and the β 's as a matrix such that $\beta\chi = \lambda$. Solution to this equation yields χ . This procedure is described in detail in [42].

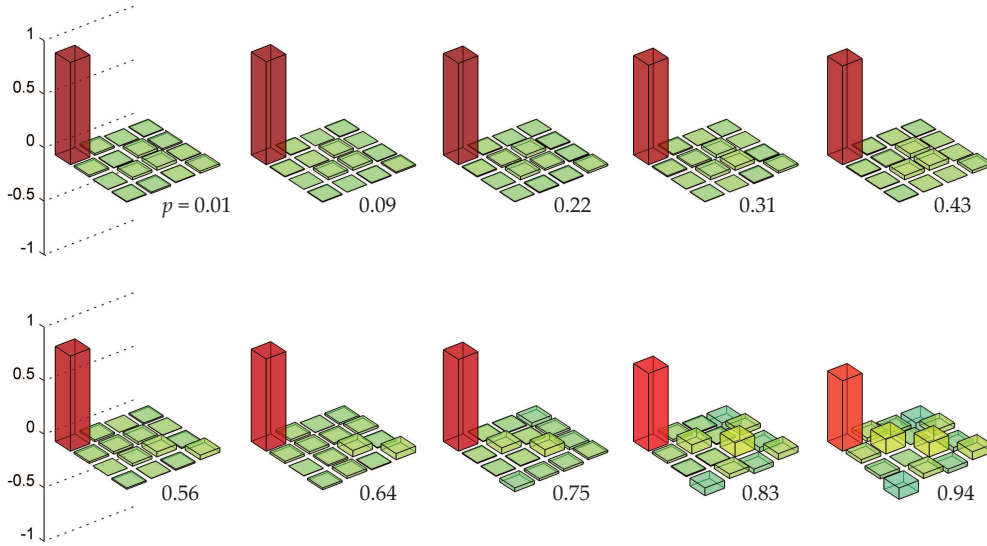


Figure 8.5: $\text{Real}(\chi)$ for uncollapse process, expressed in the Pauli $(I, \hat{X}, \hat{Y}, \hat{Z})$ basis. Increasing p left to right, top to bottom. The large peak at (I, I) indicates the uncollapse process is very nearly the identity.

Figure 8.5 shows the χ -matrix representation of the uncollapse procedure in the Pauli operator basis $(I, \hat{X}, \hat{Y}, \hat{Z})$. The values of p were determined experimentally by considering the statistics of the D and C photon counting bins and comparing to the theory set out in Section 7.1.3. The listed value is the mean of the 12 experimental runs used to compute the process tomography for each value of p .

The ideal uncollapse procedure is simply the identity, which would have a χ -matrix consisting of a 1 in the (I, I) position, and zeros elsewhere. The experimental data match this rather well, even at fairly high values of p . Only the real components of χ are plotted. The imaginary components, in the ideal case, are all 0, and the imaginary components of the experimental data are all small (typically a few %) so are omitted to reduce clutter. Near $p = 1$ we begin to see (\hat{X}, \hat{X}) and (\hat{Y}, \hat{Y}) components appear. These correspond to spin-flips about y or x and their presence indicates that the qubit is undergoing depolarisation. In the last few frames there also appear (\hat{Z}, \hat{Z}) components, which are phase flips, indicating qubit dephasing. Compare with Figure 8.3: depolarisation is partic-

ularly clear with the $|0\rangle$ input state, and the effects of dephasing move the $|+x\rangle$ and $| -y\rangle$ states inward toward the z -axis.

The χ -matrix representation allows us to define the quantum process fidelity as the overlap: $\text{Tr}(\chi_{\text{expt}}\chi_{\text{ideal}}^\dagger)$. For the uncollapse, the overlap with the ideal is just the real part of the (I, I) component. In computing χ for Figure 8.5, we took the input states to be the ideal density operators $|0\rangle\langle 0|$, $|+x\rangle\langle +x|$, $| -y\rangle\langle -y|$, and $|1\rangle\langle 1|$.

The fidelity of the uncollapse procedure is plotted in Figure 8.6, including series computed with respect to ideal inputs, and from experimentally determined inputs. Again, when using experimental inputs states, the fidelity from the $p = 0$ data is 1 by definition. We have also included the theoretical maximum fidelity achievable when the finite branching ratio from $P_{3/2}$ back to the qubit states of $\epsilon = 0.036$ is included.




Figure 8.6: Uncollapse fidelity versus p . Solid line represents the theoretical maximum including the $\epsilon = 0.036$ branching ratio. (*) assume ideal input states; (•) use experimentally determined input states (see text). Vertical error bars include expected statistical fluctuations only. Horizontal error bars are the standard deviation of experimentally determined p values for the data used to compute the point. Errors are identical for the two series plotted.

The values of p were determined experimentally, as before, and the error bar is the standard deviation of the p values in the experimental runs which make up the quantum process tomography. The alternative would be to use the standard error in the mean, using expected deviations from a binomial distribution. However, the scatter on the computed p values over the 12 runs for each nominal strength is larger than would be expected from such statistical fluctuations

alone. This is likely due to experimental parameter (e.g. deshelving laser intensity, detuning) drift, so the standard deviation of these values is probably a better estimate for the error.

The fidelity, however, depends on all 12 experimental runs, and thus we cannot take a deviation to estimate the error as we did for p . We instead compute statistical errors on the 12 values, binomially distributed, and propagate them in quadrature using the usual error formula. The formula for the fidelity (in the case of ideal input states) is linear in extracted experimental data, so this is a straightforward procedure. We use the same error bars for both fidelity series. This has the effect, for the purposes of error analysis, of ignoring any errors that may be present in the input states for the infidelities calculated from experimental input states, perhaps slightly under-representing their size for that series. In any case, it is clear that the statistically derived errors do not account for the deviation of those points from the theory curve for ideal uncollapsing in the presence of leaky transfer back to the qubit.

Attention should also be drawn to the second point in the uncollapse fidelity with respect to experimental input states series which has fidelity > 1 . This is simply the result of experimental errors causing the outputs from the second run ($p = 0.09$) to be more pure than the first ($p = 0$). No real “black box” style quantum process can do this, and in this case it is simply experimental error. However, we should note that in going from an ensemble which contains all the shots of the experiment to one in which we select shots based on measurement, we increase the purity of the resultant state. This increase in purity is akin to filtering where we increase fidelity by throwing away some of the data.

A completely random quantum process on a qubit will average a fidelity of 0.25 with respect to another random process due to the requirement that χ have trace 1, and the dimension of the operator space being 4. For $p = 1$ the theory curve drops to 0.25 since the uncollapse procedure, leaky transfer included, maps every input state to $|1\rangle$. Our experimental data explores the range 0.01–0.94, and fidelity remains above 0.65 over this entire range. For $p < 0.5$, it is nearly ideal with respect to the procedure to prepare, protect, and readout the qubit with no partial measurement at all.

9

Conclusions

In this thesis we presented experimental results under two themes in quantum information processing with trapped ions: the move to segmented ion traps to scale up the size of the quantum system under coherent control, and some initial investigations into the physics of measurement-selected ensembles. We have successfully loaded $^{40}\text{Ca}^+$ ions into a small-scale segmented trap, and used that trap system to demonstrate the partial collapse and uncollapse of a qubit wavefunction. This is the first demonstration of uncollapsing a quantum measurement in trapped ion qubits, and we report significantly higher fidelity than the only previous demonstration, in a solid-state qubit [28].

9.1 Comments on ion trapping

The method given in Chapter 2 for parametrising electrode voltages has proved to be of great practical value. It was used extensively in the experimental work with ion traps presented in this thesis. Other treatments of this topic (e.g. [26]) also develop the idea of computing a numerical approximation to the linear map from electrode configuration space to the function space of the electric potential. Our work develops this further by decomposing the electric potential space into functions which capture intuitive characteristics of the desired trapping potential. This allows us to choose electrode voltages by manipulating only a few parameters with clear purpose. This general strategy is widely applicable to many segmented trap designs, and will have greater utility as the complexity of these designs increase. The work in this thesis centres around creating a single-well trapping potential, but other situations may be analysed using a different basis set of scalar functions to decompose the potential. For example, the creation of a double-well potential for separation of ions involves an octupole [24]. There are already traps being tested with more than 60 independent electrode segments [50], and in this regime an efficient parametrisation of the configuration space is a matter of necessity.

The difficulties we encountered in our failed attempt to load ions in the Lucent trap highlight issues to be addressed in both the construction of ion traps and in the optical systems which deliver the required laser fields. The Oxford group continues to collaborate with organisations which develop chip-top ion traps [4] and has also begun to fabricate traps in-house [5].

The electrical faults developed during assembly and testing probably stem from the relatively thin ($\sim 0.4 \mu\text{m}$) oxide layer insulating the DC electrodes from a conducting ground plane. A thicker oxide layer or removal of the ground plane could address this issue. The addition of capacitors connected near the chip in vacuo can offset the reduction in capacitance* due to a redesign or removal of the ground plane.

The high background scatter observed in the Lucent trap highlights the difficulty in delivering laser beams to micro-scale ion traps, and surface traps in particular. The laser access slot in the Lucent trap is $70 \mu\text{m}$ wide, and we would expect our laser beam, if it were Gaussian, to cleanly pass through. Upgrading our vacuum view ports to fused silica, which has better transmission at 755 THz (397 nm, the Doppler laser frequency), and upgrading to better flatness are probably necessary. Our research group has recently loaded a surface trap by directing a 755 THz laser beam $150 \mu\text{m}$ above the surface of the trap plane. This is roughly $4\times$ further than in the Lucent trap from the nominal beam centre to the nearest surface, but it indicates that having optical access out of the trap plane may not be essential. Trapping has been demonstrated as close as $40 \mu\text{m}$ from the surface of a planar ion trap [51].

Using the Liverpool trap, we have developed a reliable trap system which performs well in terms of stability against drifting electric and magnetic fields, and with efficient fluorescence detection sufficient for high fidelity readout of the state of a single ion. The heretofore observed heating rate in the trap is significantly higher than would be expected for a trap of its size, but no sustained effort has yet been made to rule out the possibility of noise on the electrical supplies being responsible for some of the heating. The trap has also not been evaluated in terms of its performance in the task for which it was designed: fast separation of ions. Its geometry allows a large octupole potential to be generated along the trapping axis [24]. This may still be of interest because, even though the state of the art in trap fabrication has moved far beyond the Liverpool trap, we know of no other traps built with similar geometry.

9.2 Comments on measurement-selected ensembles

In Chapter 3 we extended two measurement-selected ensemble scenarios, weak measurements and partial measurements, specifically to the case of qubits. Quan-

*needed to ensure the DC electrodes form a good RF ground

tum information processing techniques are well suited for experimental demonstrations of these scenarios.

The Liverpool Trap system was used to demonstrate partially-collapsed and uncollapsed measurement-selected ensembles with a single trapped $^{40}\text{Ca}^+$ ion. These experiments required the development of a new qubit in $^{40}\text{Ca}^+$. The uncollapse we present is non-ideal both because of finite precision in the control of experimental parameters and because of the leaky transfer out of the $|1\rangle$ qubit state. Nonetheless, we achieve a clear demonstration of the essential physics of uncollapse, and substantially repair the “damage” done to a qubit state by the partial collapse: we achieve a process fidelity >0.65 over the entire range of collapse strengths explored: $p = 0.01 - 0.94$.

The data show that the uncollapse procedure gives a final state fidelity well above that which would be achieved if the partial transfer process were merely detected by a perfect measurement, without the subsequent uncollapse. This point has not to our knowledge been explored in previous work in this area.

9.2.1 Higher fidelity

The qubit transition is linearly dependent on the magnetic field, and thus sensitive to fluctuations, which introduce phase decoherence. These fluctuations could be reduced by better shielding or active stabilisation. Errors of this type accumulate over time, so shortening the experimental sequence would be beneficial. The sequence length is currently dominated by the $t_b = 300 \mu\text{s}$ photon collection time bins. Increasing photon detection efficiency, or reducing the background scatter would allow shorter collection bins without sacrificing readout fidelity.

The leaky transfer limits the overall measurement fidelity. Errors introduced by these population leaks dominate the dynamics of partial collapse and uncollapse when $\epsilon \approx 1 - p$. There are several reconfigurations of the deshelling process which limit these population leaks. The most straightforward solution is to directly drive a transition from the $|1\rangle$ state to another state which is in the bright manifold: a laser at 411 THz (729 nm) can drive the electric quadrupole transition $S_{1/2}$ [48], or a Raman transition could drive a transfer to $D_{3/2}$ [62] [55].

The transfer from the $D_{5/2}$ could also be initiated from a different $|m\rangle$ state such that population decaying from the $P_{3/2}$ does not end up in the $|-\frac{5}{2}\rangle$ state. For example, first a π -pulse on the magnetic $|-\frac{3}{2}\rangle \leftrightarrow |-\frac{1}{2}\rangle$ resonance transfers population to the $|-\frac{1}{2}\rangle$ state. A π -polarised deshelling laser then drives a transition to the $P_{3/2}$, $m = -\frac{1}{2}$ state, which is dipole-forbidden to decay to $|-\frac{5}{2}\rangle$. Finally, another π -pulse on $|-\frac{3}{2}\rangle \leftrightarrow |-\frac{1}{2}\rangle$ returns the population to $|-\frac{3}{2}\rangle$. In a similar scheme, the two extra π -pulses could be avoided by choosing $|-\frac{5}{2}\rangle$ and $|-\frac{1}{2}\rangle$ as the qubit levels and driving the two-photon transition between them to manipulate the qubit. However, population in the $|-\frac{1}{2}\rangle$ state is more sensitive

off-resonant scattering by the light-shift laser by a factor ε^{-1} (i.e. several orders of magnitude), since this state is coupled by the primary polarisation of the laser instead of the impure polarisation.

Off-resonant magnetic coupling on the $|-\frac{3}{2}\rangle \leftrightarrow |-\frac{1}{2}\rangle$ transition also introduces errors. These can be reduced by either increasing the light shift or by driving the magnetic transitions more slowly. The former is in a trade-off with reducing the off-resonant optical scattering by the light-shift laser. Increasing the light-shift laser intensity or its polarisation purity would allow a decrease in both of the error processes (Section 7.1.2). Driving the magnetic transitions more slowly increases the overall length of the experimental sequence, but as long as the experiment duration is dominated by readouts, this is a minor issue. The real problem is that the reduced bandwidth of the π -pulses makes them more sensitive to errors in the qubit transition frequency (caused primarily by magnetic field noise). Errors in the π -pulses reduce the effectiveness of dynamic decoupling to protect against dephasing and can accumulate over the course of the sequence.

Fluctuations in laser intensity also contribute to experimental errors. The deshelving pulse in the first partial measurement must be identical in strength to the second partial measurement to implement uncollapse correctly. A major source of such fluctuations are the single mode optical fibres used to deliver laser beams to the trap system. These fibres suppress polarisation fluctuations due to their high birefringence, but those that remain are converted to intensity fluctuations by the polarisers in the beam paths. Active stabilisation may be employed to suppress these fluctuations.

9.2.2 Symmetric uncollapse

For the partial collapse we have implemented, it is only possible to uncollapse if a particular result is obtained from the partial measurement (i.e. the dark result). This is the asymmetric case discussed in Section 3.4.3. Using quantum information techniques in trapped ions it should be possible to observe a symmetric partial collapse (e.g. Section 3.4.2) in which it will be possible to attempt uncollapse for either partial measurement result. The correct uncollapse measurement will depend on the result of the first partial measurement. This requires two qubits which can be manipulated independently, an entangling interaction, and independent readout of the qubits.

9.2.3 Weak measurements

Essentially the same experimental setup as described above could be used to demonstrate the weak measurement of a qubit. An appropriate preparation of the qubits before a CNOT gate [8], or a weakened version of the entangling

interaction (Section 3.2.2) can be used to measure the weak value of a qubit observable. Thus it is expected that under pre- and post-selection on the results of suitably chosen strong measurements on the system qubit, the weak value could be observed to lie outside the range of the eigenvalues of the qubit observable being measured. The experimental requirements for this demonstration are less stringent than the symmetric uncollapse in one respect: we may simultaneously measure both qubits. For the symmetric uncollapse we require the ability to measure the pointer qubit without measuring the system qubit during the first partial measurement. For the weak measurement it is sufficient to always measure both at the same time. This could be accomplished by fluorescence detection with a CCD camera [9].

Bibliography

- [1] Y. Aharonov, P. G. Bergmann, and J. L. Lebowitz. Time symmetry in the quantum process of measurement. *Phys. Rev.*, 134:B1410–B1416, 1964. [cited at p. 22, 23]
- [2] Yakir Aharonov, Alonso Botero, Sandu Popescu, Benni Reznik, and Jeff Tollaksen. Revisiting Hardy’s paradox: Counterfactual statements, real measurements, entanglement and weak values. arXiv:quant-ph/0104062v1, 2001. [cited at p. 25]
- [3] Yakir Aharonov and Lev Vaidman. The two-state vector formalism: an updated review. In Gonzalo Muga, R. Sala Mayato, and Inigo Egusquiza, editors, *Time in Quantum Mechanics*. Springer, 2007. [cited at p. 22, 24, 27]
- [4] D. T. C. Allcock. The Oxford planar ion trap project. Conference poster, International Conference on Quantum Optics, Obergurgl, Tirol, Austria. <http://www.uibk.ac.at/th-physik/obergurgl2010/> accessed 24 Feb., 2010. [cited at p. 68, 134]
- [5] D. T. C. Allcock, J. A. Sherman, D. N. Stacey, A. H. Burrell, M. J. Curtis, G. Imreh, N. M. Linke, D. J. Szwer, S. C. Webster, A. M. Steane, and D. M. Lucas. Implementation of a symmetric surface electrode ion trap with field compensation using a modulated Raman effect. *New J. Phys.*, 12:053026, 2010. [cited at p. 55, 68, 134]
- [6] P. A. Barton, C. J. S. Donald, D. M. Lucas, D. A. Stevens, A. M. Steane, and D. N. Stacey. Measurement of the lifetime of the $3d^2D_{5/2}$ state in $^{40}\text{Ca}^+$. *Phys. Rev. A*, 62(3):032503, Aug 2000. [cited at p. 86]
- [7] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *J. Appl. Phys.*, 83(10):5025–5033, February 1998. [cited at p. 77]
- [8] Todd A. Brun, Lajos Diósi, and Walter T. Strunz. Test of weak measurement on a two- or three-qubit computer. *Phys. Rev. A*, 77:032101, 2008. [cited at p. 30, 136]
- [9] A. H. Burrell, D. J. Szwer, S. C. Webster, and D. M. Lucas. Scalable simultaneous multi-qubit readout with 99.99% single-shot fidelity. Preprint, arXiv:0906.3304v1, *Phys. Rev. A* forthcoming, 2010. [cited at p. 137]
- [10] Alice H. Burrell. *High Fidelity Readout of Trapped Ion Qubits*. PhD thesis, University of Oxford, 2010. [cited at p. 97, 117]
- [11] C.-W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband. Frequency comparison of two high-accuracy Al⁺ optical clocks. Preprint, arXiv:0911.4527v2, *Phys. Rev. Lett.* forthcoming, 2010. [cited at p. 5]
- [12] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74(20):4091–4094, May 1995. [cited at p. 5]
- [13] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe. Scaling and suppression of anomalous heating in ion traps. *Phys. Rev. Lett.*, 97(10):103007, Sep 2006. [cited at p. 6]

- [14] Charles J. Donald. *Development of an Ion Trap Quantum Information Processor*. PhD thesis, University of Oxford, 2000. [cited at p. 16]
- [15] A. R. Edmonds. *Angular Momentum in Quantum Mechanics*. Princeton, 1957. [cited at p. 148]
- [16] Christopher J. Foot. *Atomic Physics*, pages 127–8. Oxford University Press, 2005. [cited at p. 114]
- [17] Ray Freeman. *Spin Choreography: Basic Steps in High Resolution NMR*. Oxford University Press, 1999. [cited at p. 117]
- [18] Elizabeth Marie George. Fiber optic integration in planar ion traps. Bachelor's Thesis, Massachusetts Institute of Technology, 2008. [cited at p. 68]
- [19] Pradip K. Ghosh. *Ion Traps*. Oxford University Press, 1995. [cited at p. 16]
- [20] Pradip K. Ghosh. *Ion Traps*, page 12. Oxford University Press, 1995. [cited at p. 16]
- [21] D. Hanneke, J. P. Home, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland. Realization of a programmable two-qubit quantum processor. *Nature Physics*, 6:13–16, 2009. [cited at p. 5]
- [22] J. P. Home. Trapping and cooling of ions for quantum information processing. First year report, University of Oxford, 2005. [cited at p. 63]
- [23] J. P. Home, M. J. McDonnell, D. M. Lucas, G. Imreh, B. C. Keitch, D. J. Szwer, N. R. Thomas, S. C. Webster, D. N. Stacey, and A. M. Steane. Deterministic entanglement and tomography of ion-spin qubits. *New Journal of Physics*, 8(9):188, 2006. [cited at p. 30]
- [24] J. P. Home and A. M. Steane. Electrode configurations for fast separation of trapped ions. *Quantum Information and Computation*, 6(4 & 5):289–325, 2006. [cited at p. 69, 133, 134]
- [25] Jonathan P. Home. *Entanglement of Two Trapped-Ion Spin Qubits*. PhD thesis, University of Oxford, 2006. [cited at p. 69]
- [26] D. Hucul, M. Yeo, W. K. Hensinger, J. Rabchuk, S. Olmschenk, and C. Monroe. On the transport of atomic ions in linear and multidimensional ion trap arrays. Preprint, arXiv:quant-ph/0702175v3, 2008. [cited at p. 133]
- [27] Gergely Imreh. *Implementing Segmented Ion Trap Designs for Quantum Computing*. PhD thesis, University of Oxford, 2008. [cited at p. 53, 55, 68, 71, 78]
- [28] Nadav Katz, Matthew Neeley, M. Ansmann, Radoslaw C. Bialczak, M. Hofheinz, Erik Lucero, A. O'Connell, H. Wang, A. N. Cleland, John M. Martinis, and Alexander N. Korotkov. Reversal of the weak measurement of a quantum state in a superconducting phase qubit. *Phys. Rev. Lett.*, 101(20):200401, Nov 2008. [cited at p. 34, 39, 43, 133]
- [29] Benjamin Keitch. *A Quantum Memory Qubit in Calcium-43*. PhD thesis, University of Oxford, 2007. [cited at p. 52]
- [30] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, June 2002. [cited at p. 6]
- [31] Jungsang Kim and Changsoon Kim. Integrated optical approach to trapped ion quantum computation. *Quant. Inf. Comput.*, 9(2), 2009. [cited at p. 68]

-
- [32] Alexander N. Korotkov and Andrew N. Jordan. Undoing a weak quantum measurement of a solid-state qubit. *Phys. Rev. Lett.*, 97(16):166805, Oct 2006. [cited at p. 36]
- [33] D. R. Leibbrandt, J. Labaziewicz, R. J. Clark, I. L. Chuang, R. J. Epstein, C. Ospelkaus, J. H. Wesenberg, J. J. Bollinger, D. Leibfried, D. J. Wineland, D. Stick, J. Sterk, C. Monroe, C.-S. Pai, Y. Low, R. Frahm, and R. E. Slusher. Demonstration of a scalable, multiplexed ion trap for quantum information processing. arXiv:0904.2599v2, 2009. [cited at p. 68]
- [34] D. M. Lucas, C. J. S. Donald, J. P. Home, M. J. McDonnell, A. Ramos, D. N. Stacey, J.-P. Stacey, A. M. Steane, and Simon Webster. Oxford ion-trap quantum computing project. *Philos. Transact. A Math. Phys. Eng. Sci.*, 361:1401–1408, 2003. [cited at p. 83]
- [35] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Nakayama, J.-P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. Isotope-selective photoionization for calcium ion trapping. *Phys. Rev. A*, 69:012711, January 2004. [cited at p. 48, 75]
- [36] W. W. Macalpine and R. O. Schildknecht. Coaxial resonators with helical inner conductor. *Proc. IRE*, 47(12):2099–2105, December 1959. [cited at p. 63]
- [37] Matthew McDonnell. *Two-Photon Readout Methods for an Ion Trap Quantum Information Processor*. PhD thesis, University of Oxford, 2003. [cited at p. 79]
- [38] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer, 1999. [cited at p. 147]
- [39] T. Monz, K. Kim, W. Hänsel, M. Riebe, A. S. Villar, P. Schindler, M. Chwalla, M. Hennrich, and R. Blatt. Realization of the quantum Toffoli gate with trapped ions. *Phys. Rev. Lett.*, 102(4):040501, Jan 2009. [cited at p. 5]
- [40] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Phys. Rev. Lett.*, 100(20):200502, May 2008. [cited at p. 5, 83, 97]
- [41] Alice Myerson. High-precision qubit readout and microfabricated ion traps for quantum information processing. First year report, University of Oxford, Aug 2007. [cited at p. 53, 54]
- [42] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge, 2000. [cited at p. 31, 126, 128, 129]
- [43] W. Paul and M. Raether. Das elektrische Massenfiter. *Z. Phys.*, 140(3):262–273, 1955. [cited at p. 8]
- [44] John Preskill. Fault-tolerant quantum computation. Preprint, arXiv:quant-ph/9712048v1, 1997. [cited at p. 9]
- [45] Ch. Roos, Th. Zeiger, H. Rohde, H. C. Nägerl, J. Eschner, D. Leibfried, F. Schmidt-Kaler, and R. Blatt. Quantum state engineering on an optical transition and decoherence in a Paul trap. *Phys. Rev. Lett.*, 83(23):4713–4716, Dec 1999. [cited at p. 78]
- [46] J. J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, revised edition, 1994. [cited at p. 89]
- [47] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland. Spectroscopy using quantum logic. *Science*, 309(5735):749–752, 2005. [cited at p. 5]

- [48] F. Schmidt-Kaler, H. Häffner, S. Gulde, M. Riebe, G.P.T. Lancaster, T. Deuschle, C. Becher, W. Hänsel, J. Eschner, C.F. Roos, and R. Blatt. How to realize a universal quantum gate with trapped ions. *Appl. Phys. B*, 77:789–796, 2003. [cited at p. 135]
- [49] Ferdinand Schmidt-Kaler, Hartmut Häffner, Mark Riebe, Stephan Gulde, Gavin P. T. Lancaster, Thomas Deuschle, Christoph Becher, Christian F. Roos, Jürgen Eschner, and Rainer Blatt. Realization of the Cirac-Zoller controlled-not quantum gate. *Nature*, 422:408–411, March 2003. [cited at p. 83]
- [50] Stephan A. Schulz, Ulrich Poschinger, Frank Ziesel, and Ferdinand Schmidt-Kaler. Sideband cooling and coherent dynamics in a microchip multi-segmented ion trap. *New J. Phys.*, 10:045007, 2008. [cited at p. 133]
- [51] S. Seidelin, J. Chiaverini, R. Reichle, J. J. Bollinger, D. Leibfried, J. Britton, J. H. Wesenberg, R. B. Blakestad, R. J. Epstein, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland. Microfabricated surface-electrode ion trap for scalable quantum information processing. *Phys. Rev. Lett.*, 96(25):253003, Jun 2006. [cited at p. 134]
- [52] P. Shor. Scheme for reducing decoherence in a quantum computer memory. *Phys. Rev. A*, 52:2493, 1995. [cited at p. 9]
- [53] P. W. Shor. Fault-tolerant quantum computation. In *Proc., 37th Annual Symposium on Fundamentals of Computer Science*, pages 56–65. IEEE, IEEE Press, 1996. [cited at p. 9]
- [54] Richard Slusher. private communication, 2008. [cited at p. 56, 60]
- [55] Jens L Sørensen, Ditte Møller, Theis Iversen, Jakob B Thomsen, Frank Jensen, Peter Staunum, Dirk Voigt, and Michael Drewsen. Efficient coherent internal state transfer in trapped ions using stimulated Raman adiabatic passage. *New Journal of Physics*, 8(11):261, 2006. [cited at p. 135]
- [56] D. N. Stacey, D. M. Lucas, D. T. C. Allcock, D. J. Szwer, and S. C. Webster. Optical Bloch equations with multiply connected states. *J. Phys. B: At. Mol. Opt. Phys.*, 41(8):085502, 2008. [cited at p. 79]
- [57] A. M. Steane. Error correcting codes in quantum theory. *Phys. Rev. Lett.*, 77:793, 1996. [cited at p. 9]
- [58] A. M. Steane. The ion trap quantum information processor. *Appl. Phys. B*, 64:623–642, April 1997. [cited at p. 7]
- [59] Daniel L. Stick. *Fabrication and Characterization of Semiconductor Ion Traps for Quantum Information Processing*. PhD thesis, University of Michigan, 2007. [cited at p. 55, 57]
- [60] Daniel L. Stick. private communication, 2008. [cited at p. 60]
- [61] David Szwer. *High Fidelity Readout and Protection of a $^{43}\text{Ca}^+$ Trapped Ion Qubit*. PhD thesis, University of Oxford, 2009. [cited at p. 85, 117]
- [62] K. Toyoda, H. Shiibara, S. Haze, R. Yamazaki, and S. Urabe. Experimental study of the coherence of a terahertz-separated metastable-state qubit in $^{40}\text{Ca}^+$. *Phys. Rev. A*, 79(2):023419, Feb 2009. [cited at p. 135]
- [63] Götz S. Uhrig. Keeping a quantum bit alive by optimized π -pulse sequences. *Phys. Rev. Lett.*, 98(10):100504, Mar 2007. [cited at p. 117]

- [64] John von Neumann. *Mathematical foundations of quantum mechanics*. Princeton landmarks in mathematics and physics. Princeton University Press, 1996. [cited at p. 25]
- [65] Simon C. Webster. *Raman Sideband Cooling and Coherent Manipulation of Trapped Ions*. PhD thesis, University of Oxford, 2005. [cited at p. 84]
- [66] J. H. Wesenberg. Electrostatics of surface-electrode ion traps. *Phys. Rev. A*, 78(6):063410, Dec 2008. [cited at p. 17]
- [67] J. H. Wesenberg, R. J. Epstein, D. Leibfried, R. B. Blakestad, J. Britton, J. P. Home, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Fluorescence during Doppler cooling of a single trapped atom. *Phys. Rev. A*, 76(5):053416, Nov 2007. [cited at p. 78]
- [68] G. K. Woodgate. *Elementary Atomic Structure*, pages 146–148. Oxford University Press, 2nd edition, 1980. [cited at p. 87]
- [69] Shengjun Wu and Klaus Mølmer. Weak measurements with a qubit meter. *Physics Letters A*, 374(1):34–39, 2009. [cited at p. 30]

A

Details of $^{40}\text{Ca}^+$ Calculations

A.1 Transforming the magnetic dipole Hamiltonian

Consider the Hamiltonian due to the interaction of the atomic magnetic dipole with an oscillating field

$$\hat{H} = \omega_0 \hat{J}_z + 2\Omega_B \hat{J}_x \cos \omega t \quad (\text{A.1})$$

with $\omega_0 = g_j \mu_B B_0$ the Larmour frequency, $\Omega_B = g_j \mu_B B'/2$ the reduced Rabi frequency, and ω the driving frequency. As we shall demonstrate, we can remove the time dependence of the Hamiltonian by transforming to a reference frame rotating about the z -axis at angular speed ω and making the *rotating wave approximation*.

We shall transform the state kets with a unitary operator $\hat{U} |\psi\rangle = |\psi'\rangle$. Applying this transformation to the Schrödinger Equation, we have

$$\frac{d}{dt} (\hat{U}^\dagger |\psi'\rangle) = \frac{i}{\hbar} \hat{H} \hat{U}^\dagger |\psi'\rangle \quad (\text{A.2})$$

$$\hat{U}^\dagger \frac{d}{dt} |\psi'\rangle = \frac{i}{\hbar} \hat{H} \hat{U}^\dagger |\psi'\rangle - \frac{d}{dt} \hat{U}^\dagger |\psi'\rangle \quad (\text{A.3})$$

$$\frac{d}{dt} |\psi'\rangle = \frac{i}{\hbar} \left[\hat{U} \hat{H} \hat{U}^\dagger - \frac{\hbar}{i} \hat{U} \frac{d}{dt} \hat{U}^\dagger \right] |\psi'\rangle. \quad (\text{A.4})$$

this allows us to write

$$\hat{H}' = \hat{U} \hat{H} \hat{U}^\dagger - \frac{\hbar}{i} \hat{U} \frac{d}{dt} \hat{U}^\dagger. \quad (\text{A.5})$$

The unitary transformation for a rotation about the z -axis is

$$\hat{U} = \exp \left(\frac{-i \hat{J}_z \omega t}{\hbar} \right). \quad (\text{A.6})$$

The first term in (A.5) is

$$\hat{U}\hat{H}\hat{U}^\dagger = \exp\left(\frac{-i\hat{J}_z\omega t}{\hbar}\right) \left[\omega_0\hat{J}_z + 2\Omega_B\hat{J}_x \cos\omega t\right] \exp\left(\frac{-i\hat{J}_z\omega t}{\hbar}\right) \quad (\text{A.7})$$

$$= \omega_0\hat{J}_z + 2\Omega_B \cos\omega t \left[\hat{J}_x \cos\omega t + \hat{J}_y \sin\omega t\right] \quad (\text{A.8})$$

$$= \omega_0\hat{J}_z + \Omega_B \left[\hat{J}_x + \hat{J}_x \cos 2\omega t + \hat{J}_y \sin 2\omega t\right]. \quad (\text{A.9})$$

In the rotating wave approximation, we drop the terms oscillating at $2\omega t$. The second term in (A.5) is

$$\frac{\hbar}{i}\hat{U}\frac{d}{dt}\hat{U}^\dagger = \frac{\hbar}{i}\exp\left(\frac{-i\hat{J}_z\omega t}{\hbar}\right)\frac{d}{dt}\exp\left(\frac{i\hat{J}_z\omega t}{\hbar}\right) \quad (\text{A.10})$$

$$= \frac{\hbar}{i}\exp\left(\frac{-i\hat{J}_z\omega t}{\hbar}\right)\frac{i\omega}{\hbar}\hat{J}_z\exp\left(\frac{i\hat{J}_z\omega t}{\hbar}\right) \quad (\text{A.11})$$

$$= \omega\hat{J}_z. \quad (\text{A.12})$$

Combining these terms

$$\hat{H}' = \omega_0\hat{J}_z + \Omega_B\hat{J}_x - \omega\hat{J}_z \quad (\text{A.13})$$

$$= \delta\hat{J}_z + \Omega_B\hat{J}_x \quad (\text{A.14})$$

A.2 *Mathematica* code for rotation matrices

The following *Mathematica* package computes the unitary operator representing the time evolution of the unshifted $D_{5/2}$ subspace driven by an oscillating magnetic field (Section 7.1.2). This is accomplished by evaluating rotation matrices given in (7.10). This package computes the rotation operators for $J = \frac{5}{2}$, but it is straightforwardly generalised to other J .

```
BeginPackage["JSystemEvolution`"]
```

```
U::usage="U[Ω,δ,t] is the j=5/2 propagator for a state vector
{m=-5/2,...,m=5/2} with Rabi frequency Ω, detuning δ, and time t."
```

```
Begin["`Private`"]
```

```
RyKFunc[j_,k_,m_,n_] := Sqrt[(j+m)!(j-m)!(j+n)!(j-n)! (-1)^(k-m
+n)/((j+m-k)! k! (k-m+n)! (j-k-n)!) (Cos[β/2])^(2j-2k+m-
n) (Sin[β/2])^(2k-m+n)
```

```
RxKFunc[j_,k_,m_,n_] := Sqrt[(j+m)!(j-m)!(j+n)!(j-n)! (-1)^(2k-
m+n)/((j+m-k)! k! (k-m+n)! (j-k-n)!) (Cos[λ/2])^(2j-2k+m-
```

```

n) (Sin[λ / 2])^(2k-m+n)
DRx[m-, n-] := Sum[RxKFunc[5/2, k, m, n], {k, m-n, 5/2+Min[m, -n]}]
DRy[m-, n-] := Sum[RyKFunc[5/2, k, m, n], {k, m-n, 5/2+Min[m, -n]}]
Dx = Array[DRx, {6, 6}, {-5/2, -5/2}]
Dy = Array[DRy, {6, 6}, {-5/2, -5/2}]
U[Ω-, δ-, t-] := (Dx/.λ -> ArcTan[δ/Ω]) . (Dy/.β -> t Sqrt[Ω^2+δ^2])
. (Dx/.λ -> -ArcTan[δ/Ω])

End[]
EndPackage[]

```

A.3 Time-independent light shift Hamiltonian

Here we present a more detailed derivation of the transformation of the light-shift Hamiltonian to a time-independent form presented in Section 7.1.2. We begin with the Schrödinger Equation

$$\hat{H}\Psi(\mathbf{r}, t) = i\hbar \frac{\partial \Psi(\mathbf{r}, t)}{\partial t} \quad (\text{A.15})$$

with \hat{H} the total Hamiltonian, which we can expand into atomic and interaction terms as:

$$\hat{H} = \hat{H}_{\text{atom}} + \hat{H}'(t) \quad (\text{A.16})$$

We can then expand $\Psi(\mathbf{r}, t)$ in terms of the eigenfunctions of \hat{H}_{atom} , $\phi_j(\mathbf{r})$

$$\Psi(\mathbf{r}, t) = \sum_j c_j \phi_j(\mathbf{r}) e^{-i\omega_j t}. \quad (\text{A.17})$$

We can then write the Schrödinger Equation as a set of equations on the c_j 's [38]

$$i\hbar \dot{c}_j = \sum_k c_k \hat{H}'_{jk}(t) e^{i\omega_{jk} t} \quad (\text{A.18})$$

The interaction Hamiltonian is

$$\hat{H}'(t) = -e\mathcal{E}(\mathbf{r}, t) \cdot \mathbf{r} \quad (\text{A.19})$$

$$= -eE_0 \hat{\mathbf{e}} \cdot \mathbf{r} \cos(\omega_L t) \quad (\text{A.20})$$

where we have made the *electric dipole approximation* with E_0 the electric field strength, $\hat{\mathbf{e}}$ the polarization unit vector, and ω_L the laser frequency. We then absorb the $(e\hat{\mathbf{e}} \cdot \mathbf{r})$ into a matrix independent of field strength

$$\hat{\mu}_{jk} = e \langle \phi_j | \hat{\mathbf{e}} \cdot \mathbf{r} | \phi_k \rangle \quad (\text{A.21})$$

such that

$$\hat{H}'_{jk}(t) = -\hat{\mu}_{jk}E_0 \cos(\omega_L t) \quad (\text{A.22})$$

$$= \frac{-\hat{\mu}_{jk}E_0}{2} [e^{i\omega_L t} + e^{-i\omega_L t}] \quad (\text{A.23})$$

Inserting the above into (A.18) and selecting only co-rotating terms gives

$$i\hbar\dot{c}_j = \sum_k c_k \frac{-\hat{\mu}_{jk}E_0}{2} e^{i(\pm\delta + g_j m_j \omega_B - g_k m_k \omega_B)t} \quad (\text{A.24})$$

where we have $-\delta$ for the (upper) $\text{P}_{3/2}$ states, and $+\delta$ for the (lower) $\text{D}_{5/2}$ states, with $\delta = \omega_L - \omega_a$ (i.e. red-detuning gives a negative δ). ω_a is taken as the centre-of-mass of the Zeeman-split resonance. The $g_j m_j \omega_B$ are the Zeeman shifts of the levels in frequency units. We can absorb this time dependence into the Hamiltonian by transforming to

$$c'_j(t) = \begin{cases} c_j(t) e^{i(\delta - g_j m_j \omega_B)t} & j \in \text{P}_{3/2} \\ c_j(t) e^{i(-g_j m_j \omega_B)t} & j \in \text{D}_{5/2}. \end{cases} \quad (\text{A.25})$$

After the transformation we have

$$\forall j \in \text{P}_{3/2} \quad i\hbar\dot{c}'_j = -\hbar(\delta - g_j m_j \omega_B) c'_j - \sum_{k \in \text{D}_{5/2}} c'_k \frac{-\hat{\mu}_{jk}E_0}{2} \quad (\text{A.26})$$

$$\forall j \in \text{D}_{5/2} \quad i\hbar\dot{c}'_j = \hbar(g_j m_j \omega_B) c'_j + \sum_{k \in \text{P}_{5/2}} c'_k \frac{-\hat{\mu}_{jk}E_0}{2} \quad (\text{A.27})$$

which we can write in more compact matrix form, for the primed states,

$$\hat{H}'' = -\hbar\delta\hat{I}^{(\text{P}_{3/2})} + \hbar\omega_B\hat{Z} - \frac{E_0}{2}\hat{\mu} \quad (\text{A.28})$$

where the superscript on the \hat{I} means identity for the $\text{P}_{3/2}$ states, and 0 for all others. \hat{Z} is a diagonal matrix containing the Zeeman shifts (including different Landé g-factors) for the levels. The eigenvalues of \hat{H}'' are the eigenvalues of the perturbed system, and their deviations from the unperturbed system are the light shifts.

Next, we decompose $\hat{\mu}$ into a scalar reduced dipole matrix element, d , and an angular matrix, $\hat{\mu} = d\hat{M}$. Using the *Wigner-Eckhart Theorem* (in the convention of [15]), we have

$$\langle \alpha', J', m' | \hat{\mu}_q | \alpha, J, m \rangle = (-1)^{J'-m'} \begin{pmatrix} J' & 1 & J \\ -m' & q & m \end{pmatrix} d \quad (\text{A.29})$$

$$= \hat{M}_{q,m',m} d \quad (\text{A.30})$$

where $d = \langle \alpha', J' || \hat{\mu} || \alpha J \rangle$ and the $(:::)$ constructions are $3j$ -symbols. The double-bar matrix element has no angular dependence, and has a simple relationship with the observed line-strength, Einstein A -coefficient, or other observed atomic data (see Section A.5). All angular dependence is encoded in the \hat{M}_q matrix, where $q = \{-1, 0, 1\}$ correspond to $\{\sigma^-, \pi, \sigma^+\}$ polarisation components, and we choose the convention that this matrix be real-valued.

A.4 Eigenvalues of light shift Hamiltonian

These are the eigenvalues of (7.18).

$$\begin{array}{ll} |L_J, m\rangle : & \text{eigenvalue} \\ |D_{5/2}, -5/2\rangle : & -3\omega_B \end{array} \quad (\text{A.31})$$

$$|D_{5/2}, -3/2\rangle : \quad -\frac{9}{5}\omega_B \quad (\text{A.32})$$

$$|P_{3/2}, +3/2\rangle : \quad -\frac{1}{2}\delta + \frac{5}{2}\omega_B - \sqrt{\left(\frac{1}{2}\delta + \frac{1}{2}\omega_B\right)^2 + \frac{1}{24}\Omega^2} \quad (\text{A.33})$$

$$|D_{5/2}, +5/2\rangle : \quad -\frac{1}{2}\delta + \frac{5}{2}\omega_B + \sqrt{\left(\frac{1}{2}\delta + \frac{1}{2}\omega_B\right)^2 + \frac{1}{24}\Omega^2} \quad (\text{A.34})$$

$$|P_{3/2}, +1/2\rangle : \quad -\frac{1}{2}\delta + \frac{37}{30}\omega_B - \sqrt{\left(\frac{1}{2}\delta + \frac{17}{30}\omega_B\right)^2 + \frac{1}{40}\Omega^2} \quad (\text{A.35})$$

$$|D_{5/2}, +3/2\rangle : \quad -\frac{1}{2}\delta + \frac{37}{30}\omega_B + \sqrt{\left(\frac{1}{2}\delta + \frac{17}{30}\omega_B\right)^2 + \frac{1}{40}\Omega^2} \quad (\text{A.36})$$

$$|P_{3/2}, -1/2\rangle : \quad -\frac{1}{2}\delta - \frac{1}{30}\omega_B - \sqrt{\left(\frac{1}{2}\delta + \frac{19}{30}\omega_B\right)^2 + \frac{1}{80}\Omega^2} \quad (\text{A.37})$$

$$|D_{5/2}, +1/2\rangle : \quad -\frac{1}{2}\delta - \frac{1}{30}\omega_B + \sqrt{\left(\frac{1}{2}\delta + \frac{19}{30}\omega_B\right)^2 + \frac{1}{80}\Omega^2} \quad (\text{A.38})$$

$$|P_{3/2}, -3/2\rangle : \quad -\frac{1}{2}\delta + \frac{13}{10}\omega_B - \sqrt{\left(\frac{1}{2}\delta + \frac{7}{5}\omega_B\right)^2 + \frac{1}{240}\Omega^2} \quad (\text{A.39})$$

$$|D_{5/2}, -1/2\rangle : \quad -\frac{1}{2}\delta + \frac{13}{10}\omega_B + \sqrt{\left(\frac{1}{2}\delta + \frac{7}{5}\omega_B\right)^2 + \frac{1}{240}\Omega^2} \quad (\text{A.40})$$

A.5 Light shift Rabi frequency

Here we calculate the square of the Rabi frequency for the light shift laser. From the definition of the light shift Rabi frequency (7.17),

$$\Omega_{LS}^2 = \frac{d^2 E_0^2}{\hbar^2}. \quad (\text{A.41})$$

Using $d^2 = S$, the line strength, and $\frac{1}{2}\epsilon_0 c E_0^2 = I_L$, the laser intensity we have

$$\Omega_{LS}^2 = \frac{2SI_L}{\epsilon_0 c \hbar^2}. \quad (\text{A.42})$$

In terms of the Einstein A-coefficient, S obtains

$$S = g_f \frac{3\epsilon_0 \hbar c^3}{2\omega^3} A \quad (\text{A.43})$$

where $g_f = 2J_f + 1$ is the statistical weight of the upper level. For $\text{P}_{3/2}$, $g_f = 4$. For $\text{P}_{3/2} \leftrightarrow \text{D}_{5/2}$, we also have $A = 7.7 \times 10^6 \text{ s}^{-1}$ and $\omega = 2\pi \times 351 \text{ THz}$, and thus

$$\Omega_{LS}^2 = 4.61 \times 10^{13} \left[\frac{\text{m}^2}{\text{J s}} \right] \times I_L. \quad (\text{A.44})$$

B

Partial and Uncollapse Experiment Detailed Data Tables



B.1 Pulse Timings

Description	τ (μs)	Active Pulses			
shelve	100.0	S	π	σ^-	
clear out	10.0		π	σ^-	
$\pi/2$ -pulse	4.8				LS RF
partial deshelve	2.4		π	σ^-	
wait	71.3				
π -pulse	9.6				LS RF
readout "D"	150.2				D PMT
PMT gap	26.0				
π -pulse	9.6				LS RF
readout "D"	150.2				D PMT
PMT gap	26.0				
π -pulse	9.6				LS RF
wait	90.5				
partial deshelve	2.4		π	σ^-	
wait	71.3				
π -pulse	9.6				LS RF
readout "C"	150.2				D PMT
PMT gap	26.0				
π -pulse	9.6				LS RF
readout "C"	150.2				D PMT
PMT gap	26.0				
π -pulse	9.6				LS RF
wait	79.7				
$\pi/2$ -pulse	4.8				LS RF
full deshelve	15.0		π	σ^-	
readout "B"	150.2				D PMT
PMT gap	26.0				
readout "B"	150.2				D PMT
PMT gap	26.0				

Table B.1: Uncollapse pulse timing. Pulses are **S**helving laser, π -deshelving laser, σ^- -deshelving laser, **L**ight-**S**hift laser, **R**F magnetic driving, **D**oppler laser, and **P**M**T** counter gate. Every pulse is followed by a $6.0 \mu\text{s}$ gap while timers and outputs are reprogrammed; listed times do *not* include this gap.

B.2 Partial and uncollapse experiment ensemble sizes

We present the ensemble sizes, before (N_T) and after (N_E) selection into the measurement selected ensemble for the partial collapse and uncollapse experimental runs presented in Chapter 8.

Partial Collapse									
Nom. p	Exp. p	$ 0\rangle\langle 0 $		$ +x\rangle\langle +x $		$ -y\rangle\langle -y $		$ 1\rangle\langle 1 $	
		N_T	N_E	N_T	N_E	N_T	N_E	N_T	N_E
0.00	0.02	4982	4940	4982	4918	4984	4923	4983	4928
0.10	0.10	4982	4933	4976	4732	4985	4732	4980	4553
0.20	0.22	6231	6171	6207	5549	6223	5538	6211	4919
0.30	0.31	7064	6978	7037	6002	7007	5954	7010	4990
0.40	0.42	8213	8133	8200	6579	8171	6430	8152	4871
0.50	0.55	9880	9785	9865	7296	9858	7183	9851	4654
0.60	0.62	12224	12097	12221	8668	12206	8558	12377	4954
0.70	0.75	12395	12203	12395	7898	12373	7842	12366	3729
0.80	0.81	12363	12167	12335	7568	12199	7272	12105	2846
0.90	0.93	12105	11838	12134	6761	12244	6565	12255	1531

Uncollapse									
Nom. p	Exp. p	$ 0\rangle\langle 0 $		$ +x\rangle\langle +x $		$ -y\rangle\langle -y $		$ 1\rangle\langle 1 $	
		N_T	N_E	N_T	N_E	N_T	N_E	N_T	N_E
0.00	0.01	4980	4926	4982	4928	4985	4922	4988	4918
0.10	0.09	4979	4537	4981	4522	4979	4523	4979	4523
0.20	0.22	6215	4935	6216	4889	6207	4843	6216	4812
0.30	0.31	7040	5002	7028	4967	7034	4928	7040	4912
0.40	0.43	8189	4843	8153	4888	8208	4690	8235	4673
0.50	0.56	9828	4614	9847	4585	9823	4470	9765	4295
0.60	0.64	12370	4930	12365	4792	12291	4818	12378	4508
0.70	0.75	12365	3646	12383	3376	12392	3462	12340	3116
0.80	0.83	12116	2836	12123	2511	12115	2324	12129	1989
0.90	0.94	12130	1499	12159	1199	12134	1217	12144	820

Table B.2: Partial and uncollapse experiments ensemble sizes, by nominal and experimentally determined p . N_T is the total number of experiments, excluding those that failed the A “check” bin, averaged over the 3 tomography measurements. N_E is the number in the measurement selected ensemble with the C bin (C and D bins for uncollapse) showing the “dark” result, again averaged over the 3 tomography measurements.

List of Symbols and Abbreviations

Symbol	Description	Page
◇	page break for figure placement	
●	projective measurement	23
◐	partial measurement	34
○	weak measurement	24
●	“selection” measurement	23
◆	measurement for which quantum state is defined	23
AOM	acousto-optic modulator	51
BEM	boundary element method	17
CCD	charge coupled device	53
CCE	centre control electrode	55
CF	ConFlat	71
CNOT	controlled-not gate	27
CPMG	Carr-Purcell-Meiboom-Gill pulse sequence	117
CPO	<i>Charged Particle Optics</i>	17
$D_{3/2}$	$3d \ ^2D_{3/2}$ level	48
$D_{5/2}$	$3d \ ^2D_{5/2}$ level	48
DAC	digital-to-analog converter	73
Δ_{LS}	reduced light-shift	96
\mathcal{D}_p	“leaky” transfer quantum process	44
K	pseudo-matrix transform from electrode configuration space to electric potential	16
k_n	electric potential with 1 V on electrode n	16
LCU	Laser Control Unit	52
MACOR	a machinable glass ceramic manufactured by Corning Inc.	11
m	magnetic quantum number; eigenvalue of \hat{J}_z	83
\hat{M}_{bright}	“bright” state measurement operator	98
\hat{M}_{dark}	“dark” state measurement operator	98
NPL	UK National Physical Laboratory	49
Ω_B	reduced magnetic Rabi frequency	87
Ω_{LS}	light-shift Rabi frequency	91
Ω_{rf}	trap RF drive frequency	14
Ω_{01}	qubit transition Rabi frequency	94
Ω_{12}	off-resonant magnetic scattering Rabi frequency	94
$\omega_x, \omega_y, \omega_z$	trapped ion secular motion frequencies	15
ω_0	Larmor frequency	87
ω_B	reduced Zeeman splitting	91
$P_{1/2}$	$4p \ ^2P_{1/2}$ level	48
$P_{3/2}$	$4p \ ^2P_{3/2}$ level	48

LIST OF SYMBOLS AND ABBREVIATIONS

Symbol	Description	Page
PBS	polarising beam splitter	51
PI	photoionisation	47
PMT	photo-multiplier tube	53
\mathcal{P}_p	partial collapse quantum process	121
Q_{rf}	quadratic coefficient of RF trapping potential	15
Q_{dc}	quadratic coefficient of DC trapping potential	15
RF	radio-frequency	7
$S_{1/2}$	$4s \ ^2S_{1/2}$ level	48
\mathcal{T}_p	incoherent transfer quantum process	40
TSVF	two-state vector formalism	22
UDD	Uhrig Dynamic Decoupling pulse sequence	117
UHV	ultra-high vacuum	72
\mathcal{U}_p	uncollapse quantum process	121
V	electric potential	13
V_{pseu}	radio-frequency pseudo-potential	14

List of Figures

3.1	Rotation of the Bloch sphere about the y -axis	32
3.2	Dephasing process deformation of the Bloch sphere.	33
3.3	Partial collapse process deformation of the Bloch sphere.	41
3.4	“Leaky” partial collapse process deformation of the Bloch sphere.	45
3.5	“Leaky” uncollapse process deformation of the Bloch sphere.	45
4.1	Simplified ^{40}Ca level structure	47
4.2	$^{40}\text{Ca}^+$ level structure with stimulated and spontaneous transitions	48
4.3	Typical laser system configuration.	50
4.4	Schematic drawing of imaging system	54
5.1	Lucent trap micrograph	56
5.2	Lucent trap schematic along the chip plane	57
5.3	Measured current from electrode T19 to ground.	58
5.4	Photograph of reverse-side of Lucent trap chip carrier	59
5.5	Closeup micrograph of trapping region	61
5.6	Neutral Ca fluorescence in Lucent trap vs. frequency	64
5.7	Neutral Ca fluorescence images in Lucent trap	65
6.1	Photo and drawing of Liverpool trap	69
6.2	Liverpool trap view with RF-quadrupole potential	70
6.3	Liverpool trap system optical access	71
6.4	Resonator with impedance match	72
6.5	Liverpool trap DC electrode grouping.	73
6.6	Liverpool magnetic field coil placement	74
6.7	Single ion of calcium in Liverpool trap.	76
6.8	RF correlation histograms	77
6.9	Fluorescence spectrum from scanning repump laser	80
7.1	σ^- -coupled transitions	83
7.2	Schematic rendering of state preparation by optical pumping	85
7.3	Axis of rotation, \mathbf{r}	88
7.4	$m = \frac{5}{2}$ population remaining after magnetic driving pulse, $\tau = \frac{\pi}{\Omega}$	90
7.5	Lightshifted energy levels for σ^- light.	92
7.6	Transformation of the energy spectrum	96
7.7	Laser table optical setup for σ beams	99
7.8	Liverpool trap beam optical setup	100
7.9	Ellipse with tilted normal	102
7.10	π -deshelve polarisation experiment	104

LIST OF FIGURES

7.11	σ^- deshelve, short timescale	106
7.12	σ^- deshelve, long timescale	107
7.13	Magnetic driving in the unshifted $D_{5/2}$ subspace	109
7.14	Frequency scan in the unshifted $D_{5/2}$ subspace	110
7.15	$ -\frac{3}{2}\rangle \rightarrow -\frac{1}{2}\rangle$ transition frequency probe	112
7.16	$D_{5/2}$ qubit timescan with fit.	114
7.17	$D_{5/2}$ qubit frequency scan with fit.	115
7.18	Magnetic field shifts at mains oscillation	116
8.1	Partial- and uncollapse pulse sequence diagram	120
8.2	Partial-collapse results as Bloch vectors	124
8.3	Uncollapse results as Bloch vectors	125
8.4	Uncollapse experimental state fidelity	127
8.5	Uncollapse process χ -matrix	129
8.6	Uncollapse fidelity versus p	130

List of Tables

4.1	Lasers used to manipulate ^{40}Ca and $^{40}\text{Ca}^+$	49
5.1	x -widths of laser-slot structures	61
5.2	Computed potentials and feedthrough pins for Trapping Solution 1 .	62
5.3	Computed potentials and feedthrough pins for Trapping Solution 2 .	62
6.1	Coil currents to null the magnetic field and calibrations	81
7.1	Light-shifted magnetic transition frequencies	93
7.2	Branching ratios from $P_{3/2}, m = -\frac{3}{2}$	98
7.3	Measured light-shifts with varying power and detuning	113
B.1	Uncollapse Pulse Timing	152
B.2	Partial and uncollapse experiment ensemble sizes	153