


High-Fidelity Microwave-Driven Quantum Logic in Intermediate-Field $^{43}\text{Ca}^+$



Thomas P. Harty

A thesis submitted for the degree of
Doctor of Philosophy

Balliol College
Michaelmas Term, 2013

Department of Physics, University of Oxford

Abstract

High-Fidelity Microwave-Driven Quantum Logic in Intermediate-Field $^{43}\text{Ca}^+$

Thomas P. Harty

A thesis submitted for the degree of Doctor of Philosophy

Michaelmas Term 2013

Balliol College, Oxford

This thesis is concerned with the development of an intermediate magnetic field “clock-qubit” in $^{43}\text{Ca}^+$ at 146 G and techniques to manipulate this qubit using microwaves and lasers. While $^{43}\text{Ca}^+$ has previously been used as a qubit, its relatively complicated level structure – with a nuclear spin of $7/2$ and low-lying D-states – makes cooling it in the intermediate field an intimidating prospect. As a result, previous experiments have used small magnetic fields of a few gauss where coherence times are limited and off-resonant excitation is a significant source of experimental error.

We demonstrate a simple scheme that allows $^{43}\text{Ca}^+$ to be cooled in the intermediate field without any additional experimental complexity compared with low fields. Using the clock-qubit, we achieve a coherence time of $T_2^* = 50(10)$ s – the longest demonstrated in any single qubit. We also demonstrate a combined state preparation and measurement error of $6.8(6) \times 10^{-4}$ – the lowest achieved for a hyperfine trapped ion qubit [NVG⁺13] – and single-qubit logic gates with average errors of $1.0(3) \times 10^{-6}$ – more than an order of magnitude better than the previous record [BWC⁺11]. These results represent the state-of-the-art in the field of single-qubit control. Moreover, we achieve them all in a single scalable room-temperature ion trap using experimentally robust techniques and without relying on the use of narrow-linewidth lasers, magnetic field screening or dynamical decoupling techniques.

We also present work on a recent scheme [OWC⁺11] to drive two-qubit gates using microwaves. We have constructed an ion trap with integrated microwave circuitry to perform these gates. Using this trap, we have driven motional sideband transitions, demonstrating the spin-motion coupling that underlies the two-qubit gate. We present an analysis of likely sources of experimental error during a future two-qubit gate and the design and preliminary characterisation of apparatus to minimise the main error contributions. Using this apparatus, we hope to perform a two-qubit gate in the near future.

Acknowledgements

Perhaps the pleasantest aspect of writing a DPhil thesis is being given the opportunity to properly thank the people who have made my time at Oxford so happy. My gratitude must go first to my supervisor Dr. David Lucas for taking a lost philosophy of physics student and putting him to proper work in the lab. It is hard to imagine my DPhil without David's kindness and his willingness to teach me the intricacies (and basics) of atomic physics and ion trapping.

Prof. Derek Stacey figured out how to cool $^{43}\text{Ca}^+$ at 146 G. The fact that this beast has never intimidated me is testimony to his ability to understand and simplify complicated physics and to his mastery of the optical Bloch equations. I have enjoyed many happy conversations with Derek, without which my understanding of the physical world would be much the poorer.

The work I've spent these past years pursuing all stems from foundations laid by Dr. David "D3" Allcock. D3 taught me that covering things in glue can work so much better than "gold plating" them and made sure that the lab's playlist was always interesting. I remain, however, unimpressed by his photographic skills.

Chris Ballance taught me (almost) everything I know about two-qubit gates and has endured far too many conversations about phase noise and return currents. Many of the best results in this thesis are only here because Chris asked "Well, why *can't* you do that? Surely it's trivial?".

Luca Guidoni spent a year away from his own research group to come and work with us in the Clarendon. Always a voice of calm, the experiment never worked so well as when he was here.

Thank you also to Martin Sepiol, who is carrying the microwave experiments on to bigger and better things, and to the rest of the Clarendon basement rabble: Norbert Linke, Hugo Janacek (who worked on the cooling with Derek), Diana Prado Lopes Aude Craik, Sarah Woodrow, Vera Schäfer, Simon Webster, Jeff Sherman – who first showed me Horowitz and Hill and whose enthusiasm for physics and electronics made me feel that I was where I belonged from my first days in the group – , David Szwer, Alice Myerson, Michael Curtis & Ben Keitch, as well as to Prof. Andrew Steane, the other head of our group.

Graham Quelch, our research technician, built some integral parts of the apparatus and taught me many interesting things (most of which he probably shouldn't have). My thanks also go to the Clarendon Mechanical Workshop team, including Simon Moulder, Mathew Newport, Jonathan Denton and Rob

Harris. Rob, in particular, taught me how to build things and always new how to turn half-baked ideas into reality.

The ion trappers at NIST, including Christian Ospelkaus (now at Leibniz Universität Hannover) and Dietrich Leibfried, gave us much helpful advice on the microwave gates.

Outside the world of physics, the Maxwell clan, Jesse and Wren have given me much kindness and food over my many years in Oxford and have often offered me an oasis away from life's many stresses.

Through these past years Emily has been there for me in one way or another with so much support and kindness. Her encouragement and patience kept me going, for which I will be forever grateful.

Finally, I want to thank my parents and my two siblings, Patrick and Julia. You have all shaped how I look at the world and without your love and support none of this would have been possible.

Contents

Contents	vii
1 Introduction	1
1.1 Quantum simulation and computation	1
1.2 More fundamental interests	3
1.3 The trapped-ion quantum information processor	4
1.4 Structure of the thesis	5
2 Motivation and Theory	7
2.1 What to look for in an ion qubit	8
2.2 Calcium qubits	11
2.3 Level structure in the intermediate field	13
2.3.1 The ground level	14
2.3.2 Transitions within the ground level	19
2.3.3 Clock transitions	20
2.4 Photoionization	23
2.5 Cooling at 146 G	24
2.5.1 The two-state system	25
2.5.2 Cooling calcium	26
2.5.3 The intermediate field	27
2.5.4 Dynamic effects	30
2.5.5 Low-temperature cooling on a dark resonance	31
2.6 Single-qubit manipulations	34
2.6.1 Optical pumping to the stretched state	34
2.6.2 Coherent manipulation with microwaves	36
2.6.3 Readout	39
2.7 Two-qubit gates using trapped ions	44
2.7.1 Microwaves or lasers?	46
2.7.2 Microwave mediated spin-motion coupling	49
2.7.3 The Mølmer-Sørensen force	54
3 Noise Master Equations	59
3.1 The Stochastic Liouville-von Neumann Equation	61
3.2 White noise	62
3.3 Noise with arbitrary correlations	64

3.4	Three faces of white noise	66
3.4.1	White amplitude noise	66
3.4.2	Frequency and phase noise	68
3.4.3	White phase noise	69
3.4.4	White frequency noise	70
3.5	The Noisy Quantum Harmonic Oscillator	72
3.5.1	Frequency fluctuations	72
3.5.2	Frequency fluctuations due to RF noise	74
3.5.3	Coupling to a noisy electric field	76
4	Apparatus	79
4.1	The “microwave trap”	79
4.1.1	Design	79
4.1.2	Fabrication	81
4.1.3	Characterisation	82
4.2	Vacuum system	84
4.3	Electrode Connections	88
4.3.1	Diplexer and transformer	88
4.3.2	DC supply	90
4.3.3	Trapping RF	90
4.4	Coils	91
4.4.1	146 G coils	91
4.4.2	Trim coils	92
4.5	Lasers and optics	93
4.5.1	Lasers	93
4.5.2	Trap optics	95
4.6	Imaging system	100
4.7	Microwave sources	103
4.7.1	Rubidium and GPS frequency standards	103
4.7.2	Microwave synths	104
4.7.3	Quadrupled HP8656B	106
4.7.4	Octupled AD9910 DDS	107
4.7.5	IQ Mixers modulated at D.C.	112
4.7.6	Calibration	114
5	Single-Qubit Results and Analysis	117
5.1	Doppler cooling	117
5.1.1	Setting up and calibrating experimental parameters	118
5.1.2	Results	121
5.2	State preparation	124
5.2.1	Optical pumping to the stretched state	124
5.2.2	Microwave transfer pulses	127

5.3	Readout	132
5.3.1	Shelving	132
5.3.2	Results and analysis	141
5.4	Magnetic field servo	145
5.5	Ramsey spectroscopy on the clock transition	146
5.5.1	Qubit frequency as a function of B-field	148
5.5.2	Systematic shifts	149
5.5.3	Rubidium reference calibration	149
5.5.4	Quadratic Zeeman effect	150
5.5.5	A.C. Stark Shift	151
5.5.6	AC Zeeman shift due to the trap RF	152
5.5.7	Conclusions	156
5.6	Coherence time	158
5.6.1	Experimental considerations	160
5.6.2	Results and discussion	161
5.7	Randomized benchmarking of a single qubit	164
5.7.1	Randomized benchmarking protocol	166
5.7.2	Experimental implementation	167
5.7.3	Results	169
5.7.4	Error analysis	170
5.7.5	Effects of sample size	179
6	Towards a Microwave-Driven Two-Qubit Gate	181
6.1	Introduction	181
6.2	Characterising and controlling the microwave field	182
6.2.1	The field	182
6.2.2	Rabi spectroscopy	183
6.2.3	Ramsey spectroscopy of A.C. Zeeman shifts	184
6.2.4	Displacing the ion	185
6.2.5	Polarization angles	186
6.2.6	Field gradient measurement	187
6.2.7	Field nulling	190
6.2.8	Outlook	191
6.3	Low temperature cooling on a dark resonance	192
6.3.1	Sideband thermometry	193
6.3.2	The Doppler Limit	194
6.3.3	Experimental results	195
6.4	Heating rate	199
6.5	Trapping RF fluctuations	202
6.5.1	Measuring AM noise	205
6.5.2	AM noise of various sources	209
6.5.3	7400 squarer	210

6.5.4	Diode limiter	212
6.5.5	Close-in amplifier noise	212
6.6	Motional mode cross-coupling	214
6.7	Secular frequency drift	215
6.7.1	RF power drifts	217
6.7.2	Resonator drifts	217
6.7.3	Measured drift rate	221
6.8	Concluding remarks	222
7	Conclusion	223
7.1	Comparison with other results	224
7.2	Outlook	226
A	Waveplates	233
A.1	Tilting quarter-wave plates	234
A.2	Using quarter-wave plates at multiple wavelengths	236
A.3	Characterising waveplates	236
A.4	397 nm Doppler Polarization	237
B	Definitions and useful results	241
B.1	Rotating frames and the interaction picture	241
B.2	Two-state systems and the Pauli matrices	242
B.3	The spherical basis	243
B.4	The Magnus Expansion	246
B.5	The A.C. Zeeman Shift	247
B.5.1	Transverse polarization	247
B.5.2	Longitudinal polarization	249
B.6	Rabi flopping and Ramsey Experiments	249
B.6.1	π -qubits	251
B.6.2	Ramsey experiments	252
B.6.3	Ramsey experiments as a phase noise probe	252
C	Modulation theory	255
C.1	Amplitude modulation	255
C.1.1	Modulation at multiple frequencies	256
C.1.2	Modulation by a continuous spectrum	257
C.1.3	The Wiener-Khinchin theorem	258
C.1.4	Coherent demodulation	260
C.2	Phase modulation	261
C.2.1	PM at multiple frequencies	263
C.2.2	Modulation by a continuous spectrum	263
C.2.3	PM Wiener-Khinchin relations	264

C.2.4	Frequency multiplication	265
C.3	The relationship between AM and PM	266
C.4	Modulation due to noise	267
D	Electronic Circuits	269
D.1	7400 limiting amplifier	269
D.2	RF levelling loop	271
E	Atomic structure of $^{43}\text{Ca}^+$	279
E.1	Atomic constants	279
E.2	Properties of the ground level	279
E.2.1	Expansion coefficients	279
E.3	Matrix elements	282
	Bibliography	285

1

Introduction

1.1 Quantum simulation and computation

The story of quantum computation arguably began with a keynote address given by Feynman to the MIT Physics of Computation conference in 1981 [Fey82]. In this address, Feynman considered the question of what kinds of physical system may realistically be simulated using a computer, suggesting that

The rule of simulation that I would like to have is that the number of computer elements required to simulate a large physical system is only to be proportional to the space-time volume of the physical system. I don't want to have an explosion. That is, if you say I want to explain this much physics, I can do it exactly and I need a certain-sized computer. If doubling the volume of space and time means I'll need an *exponentially* larger computer, I consider that against the rules (I make up the rules, I'm allowed to do that).

One may ask then, "what types of physical system are simulatable?" Consider first a classical system composed of n binary elements. At any time, the overall system inhabits one of 2^n different states, and hence may be represented by an n -digit binary number. Each time we add an extra element to the system,

we add an extra bit to our number, so the situation appears to be fundamentally simulatable.

Turning to the corresponding system composed of binary quantum elements – so-called “qubits” – we encounter the strange world of superposition and entanglement. We find that one cannot simply describe the system by pointing to a single one of the 2^n states which it inhabits. Instead, one must describe how much and with what phase the system inhabits each of these possibilities, requiring a total of $2^n - 1$ complex numbers to describe its state. Now each time we wish to extend our system to incorporate another element, we double the amount of information we need to store; quantum systems are not simulatable according to Feynman’s rule. To put this scaling into perspective, we note that to describe a system of 400 qubits would require $2^{400} \simeq 10^{120}$ complex numbers, which is more than the number of fundamental particles in the universe [MK13]. Clearly, this is a battle that a classical computer is not going to win for long.

Feynman’s solution to this problem was to use a new kind of device built on fundamentally quantum principles, noting

And I’m not happy with all the analyses that go with just the classical theory, because nature isn’t classical, dammit, and if you want to make a simulation of nature, you’d better make it quantum mechanical, and by golly it’s a wonderful problem, because it doesn’t look so easy.

The idea is thus to find a quantum system which we can control and to use it to simulate other quantum systems that are beyond the reach of any classical computer.

Inspired by this one might wonder if a device based on quantum principles could solve other interesting computational tasks that are classically intractable. This thought led Deutsch to the notion of a universal quantum computer [Deu85] and later to discoveries such as Shor’s factorisation algorithm

[Sho97] and Grover’s algorithm for searching unstructured lists [Gro97]. Shor’s algorithm is particularly significant since the difficulty of factorising large numbers underpins the widely used RSA cryptography system. Constructing a quantum computer thus represents a goal of immense scientific and technological importance.

1.2 More fundamental interests

At its heart, constructing a quantum computer is about developing the techniques to control simple quantum systems with exquisite precision. One side effect of the efforts put into achieving this goal is that first year graduate students now routinely perform the kinds of experiment of which Schrödinger once remarked [Sch52]

We never experiment with just one electron or atom or (small) molecule. In thought-experiments we sometimes assume that we do; this invariably entails ridiculous consequences... it is fair to state that we are not *experimenting* with single particles, any more than we can raise Ichthyosauria in the zoo.

The recent progress in this field has allowed us to probe the heart of quantum mechanics in a way that would not have seemed feasible mere decades ago. This has opened up the ability to begin to ask fundamental questions, such as what happens when we begin to create macroscopic quantum systems – can one build Schrödinger’s cat one qubit at a time? Moreover, the techniques developed in pursuit of the dream of a quantum computer have enabled some of the most precise measurements ever made and offer new avenues to test our fundamental theories [IBR⁺09].

1.3 The trapped-ion quantum information processor

Almost any system that demonstrates quantum mechanical behaviour can in principle be used as a qubit, with research currently being conducted in fields as diverse as photons, trapped neutral atoms, quantum dots, superconducting circuits, solid state nuclear spins, liquid molecule NMR and trapped atomic ions. Out of all of these system, trapped ions have emerged as perhaps the most promising candidate.

Trapped atomic ions possess several features that make them near ideal vessels for quantum information. Firstly, all atoms of the same species are identical and unchanging. Unlike a computer built from man-made elements, such as a superconducting circuit, one does not have to calibrate the characteristics of each ion or worry about them changing over time – they are simply handed to us by the universe and we take them as they come.

The other great benefit of ions may perhaps be seen in the resolution of a paradox. To construct a quantum computer, we need a system which is exquisitely isolated from its environment, allowing fragile quantum information to be stored for the duration of many logic operations. Yet, we must be able to interact with and manipulate the system's state. This takes the form of both coherent logic operations and the non-unitary, dissipative couplings required for measurements. The tension between these two requirements may be seen in the difficulties faced by many implementations of quantum computing, in which the couplings one relies on to interact with the qubits also lead to uncontrolled interactions between the qubits and their environment – including neighbouring qubits – resulting in decoherence.

By confining individual atomic ions in an ultra-high vacuum system, tens of microns away from the closest other particle, we create a near-perfectly isolated environment. Information stored in the ions' long-lived internal (hyperfine electronic) states may then be preserved for minutes [BHI⁺91], even without “re-

freshing” of the memory [SSS⁺13]. Yet, despite this isolation, we may engineer the couplings we need to operate our computer. Using laser or RF/microwave radiation we may coherently manipulate the ions’ states to perform logic gates on individual ions, while multiple ions may be entangled by engineering a coupling using their shared motion as an intermediary. Using lasers we may selectively couple our qubit states to further short-lived atomic states, scattering millions of photons and allowing us to “read out” the qubits’ states.

As Devoret remarked in a recent review of superconducting circuits [DS13],

It is no surprise that the leading quantum information technology has been and today remains the trapped ions, which are the best clocks ever built.

1.4 Structure of the thesis

This thesis is concerned with the development of an intermediate magnetic field “clock-qubit” in $^{43}\text{Ca}^+$ at 146 G, as well as techniques to manipulate this qubit using microwaves and lasers. We begin in chapters 2 and 3 with a discussion of the conceptual and theoretical background of the work. In chapter 2 we describe the physics underlying microwave-driven quantum information processing using the $^{43}\text{Ca}^+$ clock-qubit as well as the potential advantages of this system. In chapter 3 we derive master equations describing the impact of classical fluctuations in experimental parameters on coherent logic operations, providing the basis for later discussion of experimental errors.

Chapter 4 then provides a description of the apparatus we have constructed to perform these experiments. In chapter 5 we describe experimental work involving a single qubit. We discuss cooling of a single $^{43}\text{Ca}^+$ ion at 146 G with the aim of producing a high rate of fluorescence in order to allow efficient readout. We discuss spectroscopy of the qubit transition, on which we measure a coherence time of $T_2^* = 50$ s – the longest coherence time so far measured on any single

physical qubit – as well as characterisation of the fundamental single-qubit operations of preparation, coherent rotations and readout all with state-of-the-art fidelities.

In chapter 6 we move on to preliminary work aimed at performing a microwave-driven two-qubit Mølmer-Sørensen gate on the clock-qubit. We provide results of a simple technique based on a two-photon dark-resonance that allows $^{43}\text{Ca}^+$ ions to be cooled to near the Doppler limit. We present measurements of the microwave field gradients that will drive a future gate as well as an analysis of likely sources of error in such work. We discuss design and initial characterisation of experimental apparatus to minimise these errors to the point where a gate should be feasible in the near future.

Chapter 7 concludes: we discuss the work that has been conducted and how it relates to other results in the field. We finish by presenting the author’s perspective on the work to be done in the near future.

The main body of the thesis is supplemented by five appendices. Appendix A contains a discussion of the waveplate settings used to produce the laser beam polarizations required for Doppler cooling, state-preparation and shelving, as well as measurements of the polarization purities that may be produced. Appendix B defines various conventions and elements of notation used throughout the work, as well as providing derivations of a few “standard” results of quantum optics and atomic physics that may prove useful to the reader. In appendix C we review some basic aspects of modulation and noise theory which underpin our various discussions of noise. Appendix D provides schematics and descriptions of a few electronic circuits that have proved useful in tackling experimental problems we have encountered. Finally, appendix E provides a reference on various atomic properties of calcium, including ground-level transition frequencies and magnetic dipole matrix elements at 146 G and 288 G.

2

Motivation and Theory

In this chapter we explain the physics behind microwave-driven quantum information processing in $^{43}\text{Ca}^+$ at 146 G, and why this makes such a promising qubit system. We begin in section 2.1 with a discussion of desirable features in a trapped ion qubit. In section 2.2 we describe the use of calcium for quantum information processing and give an overview of the different qubits available. Section 2.3 describes the details of calcium's level structure in the intermediate field. After discussing photoionization in section 2.4, we turn in section 2.5 to the crucial question of how one should go about cooling an ion whose cooling cycle consists of 64 spectroscopically resolved states spread across three levels.

In section 2.6 we describe implementation of the basic single-qubit operations of state-preparation, logic gates and readout. We turn in section 2.7 to two-qubit gates in ion traps. We review the basics of two-qubit logic gates using trapped ions and explain why one might (or might not) want to replace traditional laser-driven entanglement schemes with one based on microwave frequency magnetic field gradients. We derive the basic Hamiltonian describing the microwave mediated spin-motion coupling and show that it can be used to implement a Mølmer-Sørensen gate.

2.1 What to look for in an ion qubit

When selecting a trapped ion qubit one has to make two choices: which atom?; and, which states? Suitable atoms should have a strong optical transition suitable for laser cooling, optical pumping and readout, as well as a sufficiently simple level structure to allow precise control over the ions' state. The most commonly used atoms are the lower mass group II alkaline earth atoms beryllium, magnesium, calcium and strontium owing to the relatively simple, alkali-like level structure of their remaining valence electron. A few non-alkaline earth atoms have also been used, most notably ytterbium and cadmium.

One qubit state is almost invariably a state in the ground level. The other may be either a second state in the ground level or a state in a metastable excited level, separated from the ground level by a forbidden optical frequency transition. While some extremely important results have been achieved using such optical qubits (see, e.g. [BKRB08b] and [SNM⁺]) they suffer from two serious limitations. Firstly, the finite lifetime of the excited level – typically of order 1 s – places a fundamental limit on how long quantum information may be stored.

Secondly, the laser used to drive the qubit transition must be extremely spectrally pure and the entire optical path between the laser and ion must be interferometrically stable. Otherwise, even a relatively small line width in the beam reaching the ion will cause the laser to lose phase coherence with the qubit over the time-scale of a computation. Moreover, because the laser drives the qubit transition directly, any spectral impurities megahertz away from the carrier frequency are liable to drive spectator transitions or to produce difficult to control light shifts.

Ground level qubits, on the other hand, have a qubit lifetime limited only by ion loss or collisions with background gas (typically once every ~ 30 minutes at UHV pressures). Moreover, transitions are driven by RF/microwave sources either directly or indirectly using stimulated Raman transitions. RF/microwave

sources with the required noise and stability levels may be purchased commercially “off the shelf”. Because of these advantages, we will only consider ground-level qubits for the rest of this thesis – despite the impressive results that have been achieved using optical qubits.

How good a particular qubit is depends on the requirements of the experiment to be performed and the equipment available at the time (for example, which laser diodes are currently being manufactured). One may, nonetheless, identify several features which are likely to be important for ion trap QIP:

1. **Mass** Heavier ions are harder to move and require more intense radiation fields to achieve a given two-qubit gate speed. This increase in intensity leads to larger errors due to off-resonant excitations and, in the case of laser-driven gates, worse errors due to photon scattering. Since the total power is often limited, a larger mass may mean longer gate times and greater susceptibility to noise. Heavier ions also require larger trapping and control voltages, which are more complicated to produce and can lead to difficulties in trap fabrication [SRW⁺].
2. **Hyperfine structure** Isotopes with hyperfine structure offer two significant advantages. Firstly, they permit so-called clock-transitions with no first-order magnetic field sensitivity. Qubits based on such transitions are robust to dephasing due to noisy external magnetic fields, which is a main noise source in many experiments¹ [MSB⁺11]. Secondly, the hyperfine splitting separates the qubit states by typically ~ 1 GHz. This splitting is large compared with the natural linewidths of the transitions used for cooling and optical pumping, allowing the qubit states to be addressed spectroscopically – an essential component of the readout scheme used in

¹Magnetic field perturbations are much more significant than those due to noisy electric fields for two reasons. Firstly, low-frequency electric fields are effectively screened by metal trap structures and the vacuum system. Secondly, unlike Zeeman shifts, Stark-shifts are a second-order perturbation so small residual electric fields do not cause significant energy shifts.

this work. The main disadvantage of hyperfine structure is that it complicates the ion's level structure and introduces spectator states in the ground level. It is also difficult to produce a differential light shift on a clock transition, making clock-qubits unsuitable for some two-qubit gate schemes (see section 2.3.3).

3. **Closeness to a two-level system** While an ideal qubit should be a perfect two-level system, practical qubits are often part of a larger manifold of states. Off-resonant excitation of spectator transitions from the qubit states limits gate speeds and introduces errors. The closer the ground-level is to a two-level system – either because it does not contain spectator states or because the spectator states are separated from the qubit by a large energy splitting – the better.
4. **Low-lying D-states** provide a meta-stable shelf-level which can be used for robust high-fidelity readout by optical pumping (see section 2.6.3) or as an optical qubit. If a quadrupole laser is used then the rich atomic structure provided by the shelf allows many useful tricks to be played (see [SNM⁺]). The disadvantage of the D-states is that they in general open the otherwise closed cooling cycle up to include the entire ground level. This complicates cooling and readout – we note that, prior to this work, no ion with low-lying D-states had been used as a qubit in an intermediate magnetic field.
5. **Visible cooling transitions** are important since lasers in the UV are more likely to charge trap dielectrics, causing drifts in compensation voltages and ion secular frequencies [HBHB10] [WHL⁺11]. In addition, many optical components are unavailable or perform poorly in the UV. This may make ions with UV cooling transitions unsuitable for schemes for scalable quantum computing involving integrated optics systems.
6. **Convenient lasers** It is advantageous if the lasers required to cool and ma-




Figure 2.1: Basic level structure of Ca^+ , consisting of low-lying S, P and D levels, with dipole allowed transitions and state lifetimes shown. The metastable $D_{5/2}$ level provides a “shelf” for readout or use as an optical qubit.

nipulate the ion may be produced from simple, readily available sources such as diodes, without the need for frequency doubling or mixing.

2.2 Calcium qubits

Calcium is a medium mass ion, with isotopes used for quantum information having mass numbers of 40 and 43 (compared with $^9\text{Be}^+$ [LOJ⁺05] or $^{171}\text{Yb}^+$). Like the alkali neutral atoms it has a single valence electron with a basic level structure consisting of low-lying S, P and D levels as shown in figure 2.1. It has a visible cooling transition at 397 nm, with other electric dipole transitions at 393 nm, 866 nm, 850 nm and 854 nm. These wavelengths are convenient for three reasons. Firstly, they may all be produced directly by commercially available diodes without frequency doubling. Secondly, they fall into two narrow

wavelength ranges (393 nm-397 nm and 850 nm-866 nm), which I will refer to as “blue” and “red”, allowing standardisation of most optics. Thirdly, the blue transitions have sufficiently long wavelengths to be compatible with high quality fibres and other optical components and do not cause significant problems with charging dielectrics in our experiments (compare with the $^9\text{Be}^+$ cooling transition at 313 nm [LOJ⁺05] or $^{25}\text{Mg}^+$ at 280 nm).

Of the two isotopes referred to, $^{40}\text{Ca}^+$ has normally been used for QIP since its lack of hyperfine structure simplifies its level structure, with a spin-1/2 ground-level presenting an exact two-level qubit system. $^{40}\text{Ca}^+$ is typically used with static magnetic fields of order 1 G providing a few megahertz of qubit splitting. Since this splitting is small compared with the 23 MHz natural line width of the dipole transitions from the ground state, readout must proceed either using a coherent population trapping technique [MSW⁺04] (limited to 95 % fidelity) or using a narrow-linewidth laser on the $S_{1/2} \leftrightarrow D_{5/2}$ quadrupole transition laser to shelve the ion.

$^{43}\text{Ca}^+$, the only stable odd isotope, has a nuclear-spin of $I = 7/2$ giving it hyperfine structure. Its ground level consists of 16 states separated into two hyperfine manifolds with $F = 3$ and $F = 4$, split by 3.2 GHz at low magnetic fields. Near zero magnetic field, the frequency of the $|F = 4, M_F = 0\rangle$ to $|F = 3, M_F = 0\rangle$ transition is independent of magnetic field to first order and has been used as a clock-qubit in previous work [LKH⁺11]. The drawback of this qubit is that the splitting between states is relatively small (usually < 5 MHz) making off-resonant excitation a problem. By using a larger magnetic field one can increase this splitting, but only at the expense of introducing a non-negligible first-order field sensitivity.

In this work, we use the clock-transition between states $|F = 4, M_F = 0\rangle$ and $|F = 3, M_F = +1\rangle$ at 146.0942 G. At this field the Zeeman splittings between states are sufficiently large (~ 50 MHz) that gates may be driven at microsecond

level speeds without introducing significant errors due to off-resonant excitation. Moreover, by operating the qubit exactly at the field-independent point, we obtain a coherence time of almost a minute without spin-echo – the longest reported for a single qubit.

These benefits make the 146 G clock-qubit a very attractive option for QIP. Its principal disadvantage is its complicated level structure. A key result of our work has been to demonstrate that – despite a cooling cycle consisting of 64 spectroscopically resolved states spread across three levels – one may still cool the ion using a simple scheme. Indeed, our scheme (described in section 2.5) is no more complicated than that used to cool $^{43}\text{Ca}^+$ at low fields and involves only the addition of a single EOM to the scheme used to cool $^{40}\text{Ca}^+$.

2.3 Level structure in the intermediate field

At 146 G the Zeeman splittings are still small compared with the ion’s fine structure, so the overall level diagram remains little changed from the low field case. There is, however, a significant change to the structure of the states within each level (see figures 2.2 to 2.6). In all levels considered apart from the $3\text{D}_{5/2}$ level, the field is neither strong nor weak with respect to the ion’s hyperfine structure. Thus, whilst the ion’s nuclear spin and orbital angular momentum, I and J , remain good quantum numbers, its total angular momentum, F , and the projections of its angular momenta onto the quantization axis, M_I and M_J do not.

Despite F not being a good quantum number, we will often label states in the low-lying $4\text{S}_{\frac{1}{2}}$ and $4\text{P}_{\frac{1}{2}}$ levels by values of $|F, M_F\rangle$. This is justified by the fact that at 146 G the Zeeman shifts, whilst not negligible, are still relatively small compared with the hyperfine splitting in these levels (~ 50 MHz compared with 3.2 GHz in the ground level). Thus the states still divide into clearly separated manifolds corresponding to different values of F . To be rigorous, “ $|F, M_F\rangle$ ” should be taken to be a convenient shorthand for “the state that one would ar-

rive at beginning in the low-field state $|F, M_F\rangle$ and ramping the field up adiabatically.”

We will not use this notation for states in the $4P_{3/2}$ and $3D_{3/2}$ levels, where the hyperfine structure is sufficiently small to be scrambled by the Zeeman splittings. In these levels, states must be identified by the projection of their total angular momentum, M , and their energy. In the $3D_{5/2}$ level, the hyperfine structure splitting is sufficiently small that we are entering the high-field region by 146 G and states are best labelled by values of $|M_I, M_J\rangle$.

2.3.1 The ground level

Much of the work conducted in this thesis centres on transitions between states in the ground level. It is thus useful to derive some basic properties of these states. In the next two sections, I derive the expansion coefficients used to express the intermediate-field states in terms of high-field angular momentum eigenstates and the matrix elements for transitions between them.




Figure 2.2: Energies of states in the $4S_{1/2}$ ground-level of $^{43}\text{Ca}^+$ as a function of magnetic field. At 146 G, despite the large Zeeman splittings, the levels divide into two manifolds, corresponding to low-field states with $F = 3$ and $F = 4$. The $|4, 0\rangle \Leftrightarrow |3, +1\rangle$ clock-qubit transition at 146 G is indicated by a vertical line.




Figure 2.3: Energies of states in the $4P_{\frac{1}{2}}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field. At 146 G, despite the large Zeeman splittings, the levels divide into two manifolds, corresponding to low-field states with $F = 3$ and $F = 4$.




Figure 2.4: Energies of states in the $4P_{\frac{3}{2}}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field. At 146 G, the Zeeman splittings scramble the hyperfine-manifold structure. By 300 G we are approaching the high-field regime and states may be labelled in terms of $|M_I, M_J\rangle$.




Figure 2.5: Energies of states in the $3D_{3/2}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field. The hyperfine manifold structure is scrambled by the Zeeman splittings.




Figure 2.6: Energies of states in the $3D_{5/2}$ level of $^{43}\text{Ca}^+$ as a function of magnetic field. The hyperfine splitting is small enough that by 146 G, we are in the high-field regime so states may be distinguished by labelling in terms of $|M_I, M_J\rangle$.




Figure 2.7: The $4S_{1/2}$ ground level of $^{43}\text{Ca}^+$ at 146 G, showing important transitions. Despite being in the intermediate field, we label states as $|F, M_F\rangle$ (see text). When discussing microwave transitions between the manifolds, I use $|M_n\rangle$ to refer to the state in the higher-energy $F = 3$ level and $|M_m\rangle$ to refer to the state in the $F = 4$ manifold. The corresponding magnetic dipole matrix element is R_{nm} . The qubit states (blue) are $|4, 0\rangle$ and $|3, +1\rangle$. The green transitions are used to transfer the ion from the $|4, +4\rangle$ “stretched” state to the qubit after optical pumping. The π transitions shown in brown are used to improve the fidelity of optical pumping.

As noted previously, the intermediate field states are not angular momentum eigenstates. One may nonetheless represent them as superpositions of the angular momentum eigenstates. One could in principle choose either the low-field $|F, M_F\rangle$ states or the high-field $|M_I, M_J\rangle$ states as a basis; however, the high-field states will prove the more convenient since the weakness of the nuclear magnetic moment compared with the electronic magnetic moment ensures that M_I is, to a good approximation, conserved during transitions.

Since we are in the intermediate-field regime, we treat the hyperfine and Ze-

man interactions together as a single perturbation with Hamiltonian [Woo80]

$$\begin{aligned} H_I &= \mathbf{A}\mathbf{I} \cdot \mathbf{J} - \boldsymbol{\mu}_J \cdot \mathbf{B} - \boldsymbol{\mu}_I \cdot \mathbf{B} \\ &= A \left\{ I_z J_z + \frac{1}{2} (I_+ J_- + I_- J_+) \right\} + g_J \mu_B B J_z - g_I \mu_N B I_z \end{aligned} \quad (2.1)$$

where we have chosen our quantization axis to lie along the direction of the B-field so that $\mathbf{B} = B\hat{\mathbf{z}}$ and we have used

$$\begin{aligned} I_{\pm} &:= I_x \pm iI_y \\ J_{\pm} &:= J_x \pm iJ_y \\ \boldsymbol{\mu}_J &= -g_J \mu_B \mathbf{J} \\ \boldsymbol{\mu}_I &= g_I \mu_N \mathbf{I} \end{aligned} \quad (2.2)$$

The Hamiltonian is thus composed of products of angular momentum projection operators and products of one raising and one lowering operator. In both cases, the matrix elements of the Hamiltonian can only be non-vanishing between states with the same value of $M := M_I + M_J$. Since $J = \frac{1}{2}$ there are at most two such states for any given M . We may thus express the energy eigenstates in the form

$$|F, M\rangle_{n,m} = \alpha_{n,m} |M_I = M + \frac{1}{2}, M_J = -\frac{1}{2}\rangle + \beta_{n,m} |M_I = M - \frac{1}{2}, M_J = +\frac{1}{2}\rangle \quad (2.3)$$

where the $M = \pm 4$ “stretched” states are the only states that are angular momentum eigenstates at all fields.

We can find the α and β expansion coefficients as well as the energy of each state by diagonalizing the Hamiltonian. To do this, we recall that the matrix elements of the operators involved in the Hamiltonian are given by

$$\begin{aligned} J_{\pm} |M_J\rangle &= \sqrt{(J \mp M)(J \pm M + 1)} |M_J \pm 1\rangle \\ J_z |M_J\rangle &= M_J |M_J\rangle \end{aligned} \quad (2.4)$$

Since $J = \frac{1}{2}$ in the ground-state, the energies may be also be found directly using the Breit-Rabi formula [Woo80]

$$E_{\pm}(B, M) = -\frac{E_{\text{hfs}}}{2(2I + 1)} - g_I \mu_N B M \pm \frac{E_{\text{hfs}}}{2} \sqrt{1 + \frac{4\chi B M}{2I + 1} + \chi^2 B^2} \quad (2.5)$$

where $\chi := (g_I\mu_N + g_J\mu_B)/E_{\text{hfs}}$ and E_{hfs} is the zero-field hyperfine structure splitting.

2.3.2 Transitions within the ground level

Since all states in the ground level arise from the same configuration, transitions between them conserve parity and are hence forbidden by the electric dipole selection rules. Transitions thus proceed via the magnetic dipole interaction with Hamiltonian $H_I = -\boldsymbol{\mu} \cdot \mathbf{B}$, where $\boldsymbol{\mu} := \boldsymbol{\mu}_I + \boldsymbol{\mu}_J$ and $\mathbf{B}$ is the driving magnetic field.

We consider a transition between the states $|F = 3, M3\rangle$ and $|F = 4, M4\rangle$, characterised by the polarization $Q := M3 - M4$ and driven by the magnetic field²

$$\mathbf{B}(t) = (-1)^Q B \hat{\mathbf{e}}_Q \cos \omega t \quad (2.6)$$

The interaction Hamiltonian is thus

$$H = (-1)^{Q+1} B \mu_Q \cos \omega t \quad (2.7)$$

We define the transition Rabi frequency to be

$$\hbar\Omega := B R_{M3,M4} \quad (2.8)$$

where R_{nm} is the matrix element of the magnetic dipole moment, defined by

$$R_{nm} := (-1)^{Q+1} \langle n | \mu_Q | m \rangle \quad (2.9)$$

These matrix elements can readily be calculated using the expansion coefficients derived in the previous section and are tabulated in appendix E. Note that according to this definition, a π -pulse takes a time of $t_\pi = \pi/\Omega$.

For a field with arbitrary polarization, we make the obvious extension

$$\hbar\Omega := B_{-Q} R_{M3,M4} \quad (2.10)$$

²See appendix B.3 for a discussion of the spherical basis notation used in this section.

2.3.3 Clock transitions

In both the high- and low-field regimes, the energies of electronic states vary linearly with applied magnetic field. The slopes, however, are different in the two cases. At intermediate fields the energies follow curved trajectories between these two asymptotes. This curvature permits the existence of points where the energies of particular transitions become independent of applied magnetic field to first order. These are the so-called “clock-transitions”, named after their use in atomic time standards. All ions with hyperfine structure and a nuclear spin greater than 1/2 exhibit clock states. The ions previously used as clock-qubits are listed in table 2.1.

Ion	States	Field (G)	$\frac{d^2 f}{dB^2}$ (kHz/G ²)	T ₂ (s)	Reference
⁹ Be ⁺	$ 1, +1\rangle \Leftrightarrow 2, 0\rangle$	119.4	6.10	14.7(1.6)	[LOJ ⁺ 05]
²⁵ Mg ⁺	$ 2, 1\rangle \Leftrightarrow 3, 1\rangle$	213	4.6		[OWC ⁺ 11]
⁴³ Ca ⁺	$ 3, 0\rangle \Leftrightarrow 4, 0\rangle$	0	2.4	1.2(2)	[LKH ⁺ 11]
¹¹¹ Cd ⁺	$ 0, 0\rangle \Leftrightarrow 1, 0\rangle$	0	0.6		[HBD ⁺ 05]
¹⁷¹ Yb ⁺	$ 0, 0\rangle \Leftrightarrow 1, 0\rangle$	0	0.6	2.5(3)	[OYM ⁺ 07]

Table 2.1: Hyperfine transitions in trapped ions previously used as clock-qubits along with the fields they become insensitive at, the second-order field sensitivities and measured coherence times (where specified in the original work). The ⁴³Ca⁺ qubit was actually used at a field of 1.78 G to break the degeneracy between states, inducing a first-order field sensitivity of 4.3kHz/G (see [BKR08a] for a characterisation of this effect). Note that a clock-qubit based on an optical transition has also been demonstrated in calcium [BKR⁺07].

From (2.5), we see that the energy of the transition between the states $|4, M_4\rangle$ and $|3, M_3\rangle$ is given as a function of magnetic field by

$$\Delta E = -g_I \mu_N B Q + \frac{E_{\text{hfs}}}{2} \sqrt{1 + \frac{4\chi B M_3}{2I + 1} + \chi^2 B^2} + \frac{E_{\text{hfs}}}{2} \sqrt{1 + \frac{4\chi B M_4}{2I + 1} + \chi^2 B^2} \quad (2.11)$$

where, as usual, $Q := M_3 - M_4$. The clock transitions are found by determining the fields at which the derivative of this expression vanishes.

The lowest magnetic field clock transition to occur in ⁴³Ca⁺ is between the states $|3, 0\rangle$ and $|4, 0\rangle$ at zero magnetic field. However, as discussed previously,

this qubit suffers from the limitation that the field required to break the degeneracy of states in each hyperfine manifold induces a first-order field sensitivity. One must thus either live with a nearly degenerate manifold or a non-negligible field sensitivity.

The first few clock transitions to occur at non-zero magnetic fields are listed in table 2.2. The qubit used in this work is the first of these transitions, occurring at a field of 146.094 G between the states $|4, 0\rangle$ and $|3, +1\rangle$. This qubit is more convenient than the $|4, +1\rangle \Leftrightarrow |3, 0\rangle$ qubit because it may be prepared from the stretched state using one fewer π -pulse. We chose to use the 146 G qubit over the 288 G qubit because the lower magnetic field is simpler to produce in the lab and because the larger Zeeman splittings occurring at 288 G were expected to further complicate the Doppler cooling.

We note, however, that the 288 G qubit has two potential advantages over the 146 G qubit: it is a π -polarized transition; and, its matrix element is larger by a factor of 1.4. Having a π -polarized qubit is useful because our trap can only produce linearly polarized microwaves and can thus couple more efficiently to a π -polarized qubit than a σ qubit. Each of these factors makes a difference of approximately $\sqrt{2}$ in how fast we can drive the qubit transition with a given microwave power. Since the achievable Rabi frequency is likely to be a limiting factor in microwave-driven two qubit gates (section 2.7), we may move to this qubit in future work.

To demonstrate the advantage of using a clock qubit, we consider qubit dephasing due to a noisy magnetic field. If a qubit's transition frequency fluctuates by an amount δf , then we expect it to dephase (to 1 radian) after a time of order

$$T_2 \sim 1 / (2\pi\delta f) \quad (2.12)$$

Typical lab field fluctuations are of order 1 mG. For the $|4, +4\rangle$ to $|3, +3\rangle$ stretched transition, this corresponds to a frequency fluctuation of 2.4 kHz leading to coherence times of order 70 μ s. In contrast, for the 146 G clock qubit

Transition	Field independent point (G)	f_0 (Hz)	$\frac{d^2 f}{dB^2}$ (kHz/G ²)
$ 4, 0\rangle \Leftrightarrow 3, +1\rangle$	146.094 2	3 199 941 076.87 (5)	2.416
$ 4, +1\rangle \Leftrightarrow 3, 0\rangle$	146.331 4	3 199 941 144.8 (3)	2.416
$ 4, +1\rangle \Leftrightarrow 3, +1\rangle$	287.782 7	3 123 181 793.7 (3)	2.514

Table 2.2: Field independent points, transition frequencies and second-order sensitivities for the first three clock-states occurring at non-zero magnetic fields in $^{43}\text{Ca}^+$ (N.B. we consider only transitions for which $Q = \{\pm 1, 0\}$). For a discussion of uncertainties in the transition frequencies, see section 5.5. The matrix elements for transitions within the ground-level as well as the α and β coefficients are tabulated in appendix E for the qubits at 146.094 G and 287.782 G.

the frequency fluctuations are only 1.2 mHz, giving a coherence time of order 2 minutes – an improvement of 6 orders of magnitude.

Finally, we note two other features of clock qubits. Firstly, from (2.11), we see that transitions with polarizations Q and $-Q$ are nearly degenerate, split only by the small factor $2g_I\mu_N B$. This means that $\sigma_{\pm}$ -polarized transitions occur in pairs (e.g. in table 2.2). This has implications for readout of π -polarized qubits (see section 2.6.2.1). The second, related, point is that only a vanishingly small differential A.C. Stark shift is produced on a clock transition by a laser that is detuned from an excited level by an amount that is large compared with the clock-transition frequency³. This renders clock-qubits unsuitable for most – although not all (see, e.g. [AKB⁺07]) – two-qubit gates based on A.C. Stark shifts (see, e.g. [LDM⁺03]).

³For a derivation, see [LBD⁺05]. In brief, however, the point is this: clock transitions exist because the first-order field sensitivities of two states are the same. For this to be the case, ignoring a correction of order $g_I/g_J \sim 10^{-3}$ due to the $\mathbf{I} \cdot \mathbf{B}$ interaction, the two states must contain the same ratio of electron spin up to spin down components. We find, $|\alpha_n| \simeq |\alpha_m|$ and $|\beta_n| \simeq |\beta_m|$. The two states thus differ only in the phases of their expansion coefficients and, if the transition is σ -polarized, in the values of M_I mixed in. However, the Stark shift is a second-order electric dipole perturbation, so it is insensitive to both the phase of the expansion coefficients and to the nuclear spin. This means that the matrix elements appearing in any Stark shifts are the same for both states, so a differential shift can only arise due to a laser beam with a different detuning from the two states. In the limit where the laser detuning is large compared with the clock-transition frequency, the detuning from each state is approximately the same so differential shifts are strongly suppressed.




Figure 2.8: Photoionization of neutral calcium. Atoms are excited from the ground level to the $4s4p\ ^1P_1$ excited level using radiation at 423 nm. From here, they are excited to the continuum using radiation of wavelength less than 390 nm (we use a laser diode at ~ 389 nm).

2.4 Photoionization

We ionize neutral calcium atoms using the two-stage, isotope selective photoionization scheme originally described in [GRB⁺01] and [LRH⁺04], illustrated in figure 2.8. A thermal beam of atoms initially in their ground level are excited by a laser tuned to the 423 nm $4s^2\ ^1S_0 \leftrightarrow 4s4p\ ^1P_1$ transition. Photoionization from this excited level is accomplished using radiation of wavelength less than 390 nm. Isotope selectivity is achieved spectrally, using the 612 MHz isotope shift between $^{40}\text{Ca}^+$ and $^{43}\text{Ca}^+$. This requires the 423 nm laser to be roughly perpendicular to the atomic beam, ensuring the Doppler broadening is less than the isotope shift.

Photoionization is essentially the same at 146 G as in the low field. States in the ground level are all nearly degenerate due to the lack of electronic angular momentum and the weakness of the nuclear magnetic moment. The structure of the $4s4p\ ^1P_1$ level is shown in figure 2.9. At 146 G we are approaching the high-field regime with respect to the atom's hyperfine structure and the states split by into three manifolds according to $M_J = \{\pm 1, 0\}$, separated by ~ 200 MHz. The ~ 20 MHz splitting across the $M_J = \pm 1$ manifolds (corresponding to the relative orientations of the nuclear and electronic spins) is unresolved in our




Figure 2.9: States in the $4s4p\ ^1P_1$ level of neutral ^{43}Ca as a function of magnetic field. At 146 G we are approaching the high-field regime with respect to the atom's hyperfine structure and states may be labelled in terms of M_I and M_J . States split into three manifolds according to $M_J = \{\pm 1, 0\}$, with $M_J = +1$ having a positive detuning from the level's centre of gravity.

experiments due to Doppler broadening and transition natural linewidths of 35 MHz.

With appropriate 423 nm polarization, ionization can be achieved by driving transitions to any of these Zeeman manifolds. In our experiments the 423 nm beam passes through the same polarization optics as the 397 nm but in polarization quadrature (see section 4.5) and is hence expected to be mainly composed of σ^- and π polarization in the power ratio of roughly 2 : 1. Correspondingly, we find loading most efficient using the $\Delta M_J = -1$ transition at ~ -200 MHz relative to the centre of gravity (+400 MHz relative to low-field ^{40}Ca).

2.5 Cooling at 146 G

Ion trap QIP experiments place two particular requirements on laser cooling⁴: a large fluorescence signal for efficient readout (section 2.6.3) and cooling to a low motional state for high-fidelity two-qubit manipulations (see chapter 6). Ini-

⁴For a discussion of cooling in the more general context of metrology and spectroscopy using trapped ions, see e.g. [IBBW95].

tial cooling is typically achieved using the Doppler cooling method, consisting of coupling the ground level to a short-lived excited level using a red-detuned laser. The Doppler shift in combination with the red detuning ensures that ions absorb more photons when in motion towards the laser than away from it, leading to cooling. Equilibrium obtains at the Doppler limit when this cooling is balanced by recoil heating due to spontaneous emission from the excited level. If further cooling below the Doppler limit is required then resolved sideband cooling techniques are typically employed [IBBW95].

2.5.1 The two-state system

Doppler cooling of a two-state system provides a useful benchmark for evaluating the performance of cooling techniques applied to more complicated systems. For a two state system, in the absence of saturation effects, the lowest temperature is reached when the cooling laser is detuned from resonance by half of the transition's natural linewidth. At this point the ion's temperature is of order

$$k_B T_D = \frac{\hbar \Gamma}{2} \quad (2.13)$$

where Γ is the rate of decay from the excited state. For the $4P_{1/2}$ level of $^{40}\text{Ca}^+$ $\Gamma = 140 \times 10^6 \text{ s}^{-1}$, giving a Doppler limit of 0.5 mK.

The highest fluorescence, on the other hand, is obtained when the cooling transition is saturated, distributing population evenly between the upper and lower states. Photons are then scattered at a rate of $\frac{1}{2}A$, where A is the cooling transition's Einstein A coefficient. For a two-state system $A = \Gamma$, however for calcium A is slightly smaller than Γ due to infrequent ($8.4 \times 10^6 \text{ s}^{-1}$) decays to the $3D_{3/2}$ level, giving a maximum fluorescence rate of 66×10^6 photons per second.

2.5.2 Cooling calcium

Cooling calcium is complicated by two factors compared with the two-state system. Firstly the $4P_{1/2}$ level may decay to the metastable $3D_{3/2}$ level (6 % branching ratio), requiring repumping via either the $4P_{1/2}$ level using an 866 nm laser or via the $4P_{3/2}$ and $3D_{5/2}$ levels using lasers at 850 nm and 854 nm (see figure 2.1). The latter approach has the advantage that it results in a relatively simple quasi-two level cooling system; however, since it involves the $3D_{5/2}$ shelf level it is not suitable for readout.

Cooling is thus usually performed using an 866 nm repumper. Whilst this avoids inclusion of the shelf in the cooling cycle, it comes at the cost of introducing the potential for two-photon dark states since both the 397 nm and 866 nm lasers drive transitions to the same excited states. Moreover, since the 866 nm laser couples the $3D_{3/2}$ level into the cooling cycle, it reduces the peak $4P_{1/2}$ population attainable, decreasing fluorescence. If both the 397 nm and 866 nm transitions are saturated then the peak attainable $4P_{1/2}$ population is reduced to 25 %⁵.

If one uses an odd isotope, such as $^{43}\text{Ca}^+$, then hyperfine structure presents further complications. Because the ground-level hyperfine splitting is very large compared with the natural linewidths of the 397 nm transitions, we require one laser to excite transitions from each hyperfine manifold⁶. This is typically provided by a single laser source with 3.2 GHz EOM sidebands. Furthermore, the hyperfine structure increases the number of states accessible to the atom during the cooling process and, correspondingly, the number of potential dark resonances. Fields of a few gauss have typically been used to destabilise these reso-

⁵See [Szw09] for a quantitative discussion of repumping techniques in Ca^+ based on the rate equations.

⁶Notice that this is not necessary for the 866 nm repumper for two reasons. Firstly, the hyperfine structure in the $3D_{3/2}$ level is significantly smaller than in the ground level. Secondly, in order to cool efficiently, one must not saturate the 397 nm transition, limiting the amount of power one may use. In contrast, we are free to use a large 866 nm intensity to saturation-broaden transitions out of the D level.




Figure 2.10: Cooling scheme in the intermediate field. A pair of predominantly σ^+ -polarized 397 nm lasers are used to localise the ion in the $4S_{1/2}|4, +4\rangle$ stretched state and the few other states of highest M . The π -polarized $4S_{1/2}|4, +4\rangle$ to $4P_{1/2}|4, +4\rangle$ transition then acts as a quasi-closed cooling transition. Decays to the metastable $3D_{3/2}$ level are repumped using an 866 nm laser.

nances [BB02].

2.5.3 The intermediate field

While both $^{40}\text{Ca}^+$ and $^{43}\text{Ca}^+$ have previously been cooled at fields of a few gauss, it was unclear how effectively cooling would work in the intermediate field. Here the Zeeman splittings are larger than the natural linewidths of the 397 nm transitions, so even with a pair of lasers one can no longer efficiently excite transitions from all states in the ground level. Moreover, there are now dozens of spectrally resolved dark resonances which could render cooling inefficient (see figure 2.12b).

Despite these complications, we find that one may cool $^{43}\text{Ca}^+$ in the intermediate field using essentially the same scheme as at low field. The important insight is that one can dramatically reduce the number of states involved in the cooling cycle by appropriate choice of the polarization of the 397 nm laser beam. We consider the case where one 397 nm frequency component is tuned half a natural line width to the red of the $4S_{1/2}|4, +4\rangle$ to $4P_{1/2}|4, +4\rangle$ transition and the other is resonant with the $4S_{1/2}|3, +3\rangle$ to $4P_{1/2}|4, +4\rangle$ transition, with the polarizations of both components chosen to be predominantly σ^+ with a small π component (figure 2.10).

In this configuration the $4S_{1/2}|4, +4\rangle$ to $4P_{1/2}|4, +4\rangle$ π -transition acts as our cooling transition. Whenever the ion scatters from $4P_{1/2}|4, +4\rangle$ to states with $M < 4$ in the ground level, the relatively intense σ^+ component acts to drive it back into the stretched state. The result is a quasi-cycling transition with the ion spending almost all of its time in the few states with largest M .

The devil is, however, very much in the detail; with laser polarizations, intensities and detunings having a large influence on the efficiency of the cooling and the rate of fluorescence produced. To study these effects, we modelled the complete S-P-D system using the optical Bloch equations [SLA⁺08], allowing us to investigate choices of parameters numerically [Jan].

Our main focus in this computational work was to maximise the rate of fluorescence that could be obtained, with only secondary concern given to the ion's temperature. This reflects the fact that for single-qubit manipulations, the ion's temperature is not important so long as it is not sufficiently large to produce significant Doppler broadening of the shelving transition. The fluorescence, on the other hand, is crucial to achieving a high readout fidelity (see section 5.3 for a discussion of both of these points).

In section 5.1, we give details of our cooling parameters alongside experimental results. Here, we note that with well chosen parameters we are able

to achieve a P-state population of 13 %, comparable to that achievable using $^{43}\text{Ca}^+$ in the low-field. With a collection efficiency of 0.3 % we are able to collect $\sim 50,000$ photons per second, which is sufficient to allow readout fidelities of $\sim 99.99\%$ ⁷.

The main drawback of this cooling scheme is that the intensity of the 397 nm beam's σ^+ component needs to be 10 – 100 saturation intensities to allow re-pumping from the few Zeeman states with significant population. This is large compared with the intensities used to cool in the low field, leading to an increased amount of light scattered from the trap's surface and, correspondingly, to worse readout errors. Moreover, given our beam geometry, we are forced to use a 397 nm polarization ratio of $2/3 \sigma^+$ to $1/3 \pi$ (see section A.4), meaning that the intensity of the π -polarized beam is similarly high. This leads to saturation broadening of the cooling transition, which prevents us from reaching the Doppler limit with this technique.

It would also be possible to cool using σ^- radiation to confine the ion around the $M = -4$ end of the ground-level. The two schemes are essentially equivalent apart from one point: if we work in the $M = +4$ end of the ground level then the remaining 397 nm transitions are blue-detuned from the dominant $4S_{1/2}|4, +4\rangle$ to $4P_{1/2}|4, +4\rangle$ transition. These transitions add a large amount of structure to 397 nm detuning scans once the laser is blue of resonance. However, since cooling is always achieved with the laser tuned to the red we do not see this structure and – apart from the 397 nm-866 nm dark resonances discussed in the next section – we find that fluorescence changes monotonically with the 397 nm detuning, simplifying the cooling process.

One might worry that if the ion became sufficiently hot that its Doppler broadening became comparable to the detunings of these transitions (for example due to a collision with a background gas particle) then they will lead to

⁷This includes error due to a background signal of $\sim 10,000 \text{ s}^{-1}$ due to laser scatter from the trap's surface.

further heating and potentially to ion loss. We do not, however, find this to be a problem in our trap, with standard cooled single-ion lifetimes being ~ 3 hours.

2.5.4 Dynamic effects

We now consider an additional, more subtle effect on the cooling process arising from the equilibration time of the atomic population in the D states. Since 866 nm photons are scattered rarely and have a low momentum compared with 397 nm photons, one may neglect the momentum they impart to the ion. Cooling is thus, to a good approximation, solely due to the absorption and emission of 397 nm photons, with the 866 nm laser only having an effect by the way it modulates the probabilities of these events occurring.

Now, a naïve picture of cooling is given by the assumption that atomic populations respond quickly to changing Doppler shifts due to the ion's secular harmonic motion in the trap. To the extent that this is the case, one may model cooling by calculating the change in fluorescence as a function of Doppler shift for a stationary ion at equilibrium with the radiation fields. If the fluorescence is greater for Doppler shifts corresponding to motion towards the laser than away the ion is cooled and, if not, it is heated.

Since the 866 nm detuning modulates the ion's fluorescence, one expects it to play a significant role in determining whether or not the ion is cooled. For example, consider the case where the 397 nm and 866 nm lasers co-propagate and the 866 nm laser is blue-detuned from its resonance. In this case, according to our naïve time-independent model, because of the Doppler shift on the 866 nm laser, the fluorescence may decrease as the ion moves towards the lasers, leading to heating – even if the 397 nm laser is red-detuned.

Now, the assumption that the ion's motion is essentially adiabatic is generally justified for the S-P system, whose equilibration time-scale is set by the P level's A coefficient of $132 \times 10^6 \text{ s}^{-1}$ which is fast when compared with typical

secular frequencies (~ 1 MHz). However, using typical laser parameters this is not so for the population in the D level, whose equilibration time constant can be of the order of a few microseconds. In this case, the D-level population cannot respond to the sinusoidally varying Doppler shift on the 866 nm laser. Thus the Doppler shift on the 866 nm laser does not lead to a change in fluorescence with the phase of the ion's motion, so our previous argument cannot be correct.

Indeed, in many cases, it is likely that the ion's motion will act as an effective line broadening mechanism on the 866 nm transition, rather than modulating the fluorescence at the frequency of secular motion. Thus, we see that some conditions which, under the naïve picture one would expect to lead to heating, still cool the ion. The lesson to be learnt from this is that the effect of D-state lifetime must be treated using time-dependent methods. For a discussion of how this can be done, see [Jan]. However, with an intense 866 nm laser close to resonance, as used by our standard cooling parameters (section 5.1), this effect is not expected to be significant.

2.5.5 Low-temperature cooling on a dark resonance

While the Doppler cooling scheme described above provided a good basis for our single-qubit work, it was not deemed sufficient for investigations into the microwave-driven spin-motion coupling (section 2.7). For this, we wanted our ions close to their motional ground states (see chapter 6). In principle, this could have been achieved by following initial Doppler cooling with a stage of resolved sideband cooling. However, this would have added extra complexity into the experimental set-up. Moreover, many schemes for resolved sideband cooling would not be expected to work effectively on the high thermal state left by our initial cooling.

The final temperature reached by the ion during Doppler cooling depends on both the absolute rate of fluorescence – which leads directly to recoil heating

– and on the difference between fluorescence during forwards and backwards motions, which can cool. To achieve the lowest temperature, we wish to have a large change in fluorescence with Doppler shift, while maintaining a low absolute rate of fluorescence.

Now, for a two-state system the absolute fluorescence and its gradient with respect to laser frequency are tied together by the cooling transition's Lorentzian response function. However, for the more complicated $^{43}\text{Ca}^+$ cooling cycle, the fluorescence does not change in such a simple manner with the laser detunings, so there is scope for achieving lower temperatures. Indeed, as can be seen from experimental scans of the 866 nm and 397 nm laser detunings (figure 2.12b), ideal conditions are found near the bottom of a dark resonance.




Figure 2.11: A dark resonance in the lambda system composed of states A & B together with an excited state, C. When the detuning of the fields coupling A & B to C are equal there exist coherent superpositions of A and B that do not couple to C. However, if either A or B is coupled to a further state, D, then the dark resonance is destabilised.

Two-photon dark resonances occur when two states are both coupled to a common excited state (figure 2.11). When the radiation fields have equal detunings, there exist superpositions of the two states which do not couple to the excited state, suppressing fluorescence. Due to the resonant nature of this phenomenon, fluorescence is rapidly restored as the radiation's detuning is changed, providing the steep gradient required for cooling.




Figure 2.12: The 397 nm-866 nm dark resonance used for low-temperature Doppler cooling. (a) Level diagram showing the dark resonance. (b) Experimental scan and numerical simulation of the $4P_{1/2}$ population as a function of the 866 nm laser's detuning using our standard low-temperature cooling parameters (section 6.3). The 866 nm frequency scale is relative to the transition's centre of gravity ($-3\,464.3$ MHz relative to zero-field $^{40}\text{Ca}^+$). We achieve the lowest temperatures by tuning the laser near the bottom of the main dark resonance at -170 MHz.

While there are many dark resonances in the S-P-D system, most are impractical for cooling since one or both of the dark states is coupled to further states. This allows population to escape from the dark superposition state, restoring fluorescence (figure 2.11). The result is a definite decrease in fluorescence when the resonance condition is met, but not enough to make the feature worthwhile for cooling. Given our choice of 397 nm polarization (σ^+ and π , with negligible σ^-) there is only one dark resonance that forms an ideal lambda system: that between $4S_{1/2}|4, +4\rangle$ and $3D_{3/2}|+5\rangle$ ⁸ (figure 2.12a). The only transition out of $4S_{1/2}|4, +4\rangle$ is the π -transition to the state $3P_{1/2}|4, +4\rangle$, while the only induced transition from $3D_{3/2}|+5\rangle$ is the σ^- transition to the same state. Resonance occurs when the detunings of the 397 nm and 866 nm lasers from these transitions are equal. This resonance is clearly visible as the deepest feature on experimental scans (figure 2.12b). For a stationary ion and monochromatic laser beams, the fluorescence on resonance would drop to zero.

To utilise this feature for cooling, we tune our lasers so that the 397 nm is near to resonance with the $4S_{1/2}|4, +4\rangle$ to $4P_{1/2}|4, +4\rangle$ transition, while the 866 nm is set near to the bottom of the dark resonance. We set-up our beam paths so that the 397 nm and 866 nm lasers counter-propagate. This ensures that the sign of the Doppler shifts on the two lasers are different, maximising the change in fluorescence as a function of Doppler shift. Using this technique, we achieved temperatures close to the $^{40}\text{Ca}^+$ Doppler limit (see section 6.3).

2.6 Single-qubit manipulations

2.6.1 Optical pumping to the stretched state

The first preliminary to any computation is initialisation of the qubit to a known state. Like most ion trap QIP experiments, we achieve this using optical pump-

⁸Notice that in the $3D_{3/2}$ level, F is not a good quantum number. However, since there is only one state with $M = 5$ identification is unambiguous in this instance.




Figure 2.13: Optical pumping to the $4S_{1/2}|4, +4\rangle$ stretched state, using a pair of σ^+ polarized 397 nm beams, one addressing each hyperfine manifold, and the standard 866 nm repumper.

ing to prepare the ion in the stretched state, $4S_{1/2}|4, +4\rangle$ in our case. We do this using a pair of σ^+ -polarized 397 nm beams, one addressing each hyperfine manifold, alongside our usual 866 nm repumper (figure 2.13). The fidelity of this process is limited by the polarization purity of the 397 nm beams. We estimate that in this work we were able to reduce the fractional intensity of the impurities to $\lesssim 10^{-3}$. Achieving a higher purity is possible, but experimentally somewhat demanding.

Polarization impurities causes some population to be “lost” from the desired $4S_{1/2}|4, +4\rangle$ state; this is distributed predominantly in states of highest M in the $4S_{1/2}$ level, specifically in the $M = +3$ and $M = +2$ states. We are able to reclaim most of this lost population using a simple technique based on microwave π -pulses and optical pumping. We begin by applying the 397 nm beam addressing the $F = 3$ manifold in the ground-level without the $F = 4$ beam (figure 2.14a). This clears out population left in the $F = 3$ manifold, some of which is transferred to the stretched state. This is experimentally easy since our two 397 nm beams derive from the same laser using an EOM (see section 4.5). So long as we choose the carrier to address the transitions from the $F = 3$ manifold

and the first blue sideband to address the $F = 4$ manifold, we are able to extinguish the $F = 4$ beam by turning off the EOM. The remaining beam is 3 GHz off resonance from transitions out of the stretched-state, so the effect of polarization impurities is greatly reduced.

We are now left with lost population in the $|4, M < 4\rangle$ states, but negligible population in those with $|3, M \leq 3\rangle$. We now use microwave π -pulses to transfer population from the $|4, M < 4\rangle$ states into the $F = 3$ manifold, before re-applying our 397 nm carrier beam to move this population back into the $F = 4$ level. After a few cycles of this process, the majority of the lost population is transferred into the stretched state (figure 2.14).

We only apply microwaves on the transitions from the $|4, +3\rangle$ and $|4, +2\rangle$ states (t & q in figure E.1) as the population in states with $M < 2$ is sufficiently small that further microwave pulses offer little improvement. We estimate (section 5.2) that this technique reduces the infidelity of our optical pumping by an order of magnitude. Since applying extra microwave frequencies may be easily accomplished in our set-up, we get this improvement at very little cost in terms of experimental complexity. This improvement may be particularly important in future ion traps using integrated optics, where the achievable polarization purities may be limited by the poorer quality of the miniaturised optics.

2.6.2 Coherent manipulation with microwaves

Our next task is to transfer the ion from the stretched-state to the clock-qubit. We achieve this by applying microwave π -pulses resonant with transitions A , B & C (figure 2.15) to map the state $|4, +4\rangle$ to $|3, +1\rangle$. Once the ion is in one of the qubit states, we can use further microwave pulses with appropriate durations and phases to perform arbitrary single-qubit logic operations (see section 5.7).

Before we can readout the final state of our qubit (section 2.6.3), we must map the clock-qubit onto a qubit composed of the state $|4, +4\rangle$ and a state in the




Figure 2.14: Enhanced optical pumping procedure. (a) After initial optical pumping with two 397 nm σ^+ beams (figure 2.13), we use a single 397 nm beam to re-pump any population left in the $F = 3$ manifold back into the $F = 4$ manifold, some of which makes its way into the stretched state. (b) Next, we apply microwave π -pulses to move population in $|4, M < 4\rangle$ states back into the $F = 3$ manifold. After a few cycles of these two processes, the infidelity of optical pumping may be reduced by roughly an order of magnitude.




Figure 2.15: Microwave transitions used for single-qubit manipulations. Transitions A, B & C are used to transfer between the qubit and the stretched state, whilst the qubit transition, Q, is used for single-qubit manipulations.

$F = 3$ manifold. We do this using the series of π -pulses $\{CBAQ\}$ (figure 2.16) to implement the mapping

$$\begin{aligned} |3, +1\rangle &\Rightarrow |4, +4\rangle \\ |4, 0\rangle &\Rightarrow |3, +1\rangle \end{aligned} \quad (2.14)$$




Figure 2.16: We use the pulse sequence $\{ABC\}$ to transfer the ion from the stretched state to the $|3, +1\rangle$ qubit state. After performing our single-qubit logic operations, we use the pulse sequence $\{CBAQ\}$ to map the qubit onto the temporary $\{|3, +1\rangle, |4, +4\rangle\}$ qubit for readout.

2.6.2.1 π -qubits

We can use pulse sequences analogous to that of figure 2.16 to prepare and read out arbitrary qubits in the ground level. We use many of these qubits for various diagnostic and calibration experiments, such as AC Zeeman shift measurements used to calibrate the microwave B-field (section 6.2). If the qubit transition is σ -polarized then the required changes to the pulse sequence are straightforward. However, for π -qubits – pairs of states separated by a π -polarized transition – there is a slight complication due to the fact that $\sigma^\pm$ transitions occur in near-degenerate pairs.




Figure 2.17: Transitions around the $|4, 1\rangle \Leftrightarrow |3, +1\rangle$ qubit. The $|4, +1\rangle \Leftrightarrow |3, +2\rangle$ transition (red) is less than 100 kHz detuned from C, leading to errors when reading the qubit out.

To illustrate this point, consider attempting to readout the qubit composed of the states $|4, +1\rangle$ and $|3, +1\rangle$. To do this, we need to map the state $|3, +1\rangle$ onto the stretched state. We might try using our usual $CBA(Q')$ pulse sequence

(figure 2.17). However, the C pulse is only detuned from the $|4, +1\rangle \Leftrightarrow |3, +2\rangle$ transition by 84 kHz. Since we are only capable of producing linearly polarized radiation in our set-up, the C pulse will cause significant off-resonant excitation of this transition, leading to readout errors.




Figure 2.18: Simple composite pulse sequence used to readout π -qubits. We replace the C π -pulse with the composite $\frac{\pi}{2}$ -delay- $\frac{\pi}{2}$ sequence on the same transition.

We use the simple composite pulse sequence illustrated in figure 2.18 to get around this problem⁹. Here, we replace our π -pulse with two $\frac{\pi}{2}$ -pulses, separated by a delay of $\tau \simeq 1/(2\delta)$, where δ is the frequency splitting between the $|4, +2\rangle \Leftrightarrow |3, +1\rangle$ and $|4, +1\rangle \Leftrightarrow |3, +2\rangle$ transitions¹⁰.

Since the microwaves are resonant with the $|4, +2\rangle \Leftrightarrow |3, +1\rangle$ transition, the composite sequence is simply equivalent to a single π -pulse if the ion begins in the state $|3, +1\rangle$. If the ion begins in the $|4, +1\rangle$ state then it is transferred to a superposition of $|4, +1\rangle$ and $|3, +2\rangle$. During the delay, the microwaves accumulate a phase difference of π relative to this superposition. Thus, the second $\frac{\pi}{2}$ -pulse acts to undo the effect of the first, moving the ion back to the $|4, +1\rangle$ state. Subsequent readout pulses then proceed as usual.

2.6.3 Readout

Ion qubits are typically read out using state-dependent fluorescence detection. This technique, originally proposed by Dehmelt [Deh75], involves coupling one qubit state to a short-lived ancillary level via a cycling transition which does not involve the other qubit state. By watching for photons scattered on this transition one can tell which state the ion is in.

The accuracy of this technique is determined by the rate at which fluorescence photons can be detected from the bright qubit state, compared with the

⁹The use of this pulse sequence was originally suggested to us by Christian Ospelkaus.

¹⁰An exact expression for the appropriate delay is given in section B.6.

rates at which the dark state gets pumped into the cycling transition manifold or the bright state gets pumped out of it. Since decay rates of ancillary levels are typically $\sim 10^6 \text{ s}^{-1}$, many photons may be collected before the undesired pumping processes occur, despite typical collection efficiencies of $\sim 0.1 \%$.




Figure 2.19: Fluorescence detection in $^9\text{Be}^+$ at 119.4 G. The ion is read out from a qubit consisting of the bright $|2, +2\rangle$ state and the dark $|1, +1\rangle$ state. σ^+ -polarized light at 313 nm creates a cycling transition between the bright state and the excited $2P_{3/2}|+3\rangle$ state, scattering millions of photons per second. This radiation is far detuned from any transition out of the dark state, heavily suppressing fluorescence.

Application of this technique is straightforward for ion qubits which permit a suitable cycling transition. As an example we consider $^9\text{Be}^+$ at a field of 119.4 G, as described in [Lan06]. There exists a cycling transition at 313 nm between the bright state $2S_{1/2}|2, +2\rangle$ and the ancillary state $2P_{1/2}|3, +3\rangle$. The lifetime of the P-state is 8.2 ns, so when a resonant σ^+ polarized laser at 313 nm is applied with of order one saturation intensity the ion scatters approximately 20 million photons per second. However, if the ion begins in the dark $2S_{1/2}|1, +1\rangle$ state then the laser is far detuned so photon scattering is greatly suppressed.

In calcium, matters are complicated by presence of low-lying D-levels which the $4P_{3/2}|5, +5\rangle$ state may decay to. The lasers required to repump from these levels provide pathways for the ion to escape to states with $M < 4$, opening

up the fluorescence cycle to involve the entire ground-level. Thus, ground-level qubits in calcium cannot be read out directly by fluorescence detection¹¹.




Figure 2.20: Shelving of the $|4, +4\rangle$ state by optical pumping: a $393\text{ nm } \sigma^+$ laser connects the $|4, +4\rangle$ state to the $M = +5$ state in the $3P_{3/2}$ level. From here, it can decay back to the $|4, +4\rangle$ state (93.5 % branching ratio) or into either the $3D_{3/2}$ or $3D_{5/2}$ levels. Decays to the $3D_{3/2}$ level (0.66 % branching ratio) must be to states with $M = 4$ or $M = 5$ and are repumped to the $4P_{3/2}|5, +5\rangle$ state using 850 nm beams with σ^+ - and π -polarizations respectively. Once the ion decays to the $3D_{5/2}$ level (5.87 % branching ratio), it is removed from the 397 nm - 866 nm cooling cycle and is said to be “shelved”. States in the $F = 3$ manifold are not shelved due to the 3.2 GHz ground-level hyperfine splitting.

Since the normal 397 nm - 866 nm cooling cycle does not involve the metastable $3D_{5/2}$ level one may, however, directly read out optical qubits composed of a state in the ground level and a state in the $3D_{5/2}$ level. We utilise this fact by incoherently mapping the $4S_{1/2}|4, +4\rangle$ state to the $3D_{5/2}$ “shelf” level using the optical pumping technique originally described in [MSW⁺08]. Read-

¹¹In fact, as the field becomes larger, a certain degree of state-selectivity is obtained in the fluorescence. However, this is unlikely to be enough to be useful for readout.

out of this optical qubit has already been demonstrated in $^{40}\text{Ca}^+$ with infidelity of $0.87(11) \times 10^{-4}$ by detecting photons scattered on the usual 397 nm-866 nm cooling transition [MSW⁺08].

The shelving process begins with the application of a 393 nm σ^+ -polarized laser tuned to resonance with the $4\text{S}_{1/2}|4, +4\rangle \leftrightarrow 4\text{P}_{1/2}|5, +5\rangle$ transition (figure 2.20). This populates the $4\text{P}_{3/2}|5, +5\rangle$ state, from which the ion can decay either back to the $|4, +4\rangle$ state (93.5 % branching ratio) or to one of the $3\text{D}_{3/2}$ or $3\text{D}_{5/2}$ levels. Decays to the $3\text{D}_{3/2}$ level (0.66 % branching ratio) must be to states with $M = 4$ or $M = 5$ and are repumped using an 850 nm laser with σ^+ - and π -polarizations respectively. Once the ion decays to the $3\text{D}_{5/2}$ level it is removed from the 397 nm-866 nm cooling cycle and is said to be shelved.

Our dark state is thus the shelved $|4, +4\rangle$ state. Since the 393 nm laser is 3 GHz from resonance with transitions out of the $F = 3$ manifold, states in this manifold remain unshelved. The bright state may thus be chosen to be any state in the $F = 3$ manifold.

The simplest 850 nm repumping technique uses a single beam at a slight angle to the static magnetic field, thus producing predominantly σ^+ light with a small π component. This beam can be applied at the same time as the 393 nm beam. In principle, fidelities of 99.96 % should be achievable using this method at a magnetic field of a few gauss [Szw09].

This simple scheme has two major drawbacks. Firstly it uses a single 850 nm frequency to repump several states. This requires saturation broadening the transitions, increasing errors due to polarization impurities. To overcome this limitation we use three 850 nm frequencies – one for each transition. These are derived from a single laser using AOMs.

Secondly, applying the lasers at the same time reduces the effectiveness of each of them. If the 850 nm is applied at the same time as the 393 nm then it will not only repump population that decays to the $3\text{D}_{3/2}$ level, it will also pump

population into it. Furthermore, by applying the 850 nm π laser at the same time as the 850 nm σ^+ laser, we allow population that decays into $M = 4$ states in the $3D_{3/2}$ level to be repumped into $M = 4$ states in the P-level, from which it can decay into $M < 4$ states in the ground-level.

The most effective method of shelving – and the method used in this thesis – is thus to pulse the lasers on in the sequence

$$\{393 \text{ nm}, 850 \text{ nm } \sigma^+, 850 \text{ nm } \pi\}^n 393 \text{ nm} \quad (2.15)$$

where n is the number of repeats. The infidelity of this technique is limited by both off-resonant shelving of the dark state and decay of the shelf-level, to $\simeq 1 \times 10^{-4}$.

It is interesting to contrast the readout schemes used for $^9\text{Be}^+$ and $^{43}\text{Ca}^+$. Both have the advantage that they utilise incoherent optical pumping techniques on dipole allowed transitions. As such they only require low ($\sim 1 \mu\text{W}$) laser intensities and can tolerate relatively large ($\sim 1 \text{ MHz}$) laser linewidths. Moreover, the Zeeman shifts induced by the intermediate field push unwanted transitions several linewidths from resonance, reducing the effects of polarization impurities. This makes them robust and experimentally easy to implement.

The more direct $^9\text{Be}^+$ scheme has the advantage that it only requires a single laser, whilst the $^{43}\text{Ca}^+$ scheme requires an additional repumping laser. However, in practice it is likely that the calcium lasers will be significantly easier to work with, as discussed previously.

The accuracy of the calcium scheme is fundamentally limited to $\simeq 1 \times 10^{-4}$, whilst the beryllium scheme may produce much smaller errors *if* the scattered photons can be detected with sufficient efficiency. However, in practice the detection efficiencies of typical imaging systems are an order of magnitude smaller than the branching ratio of the $4P_{3/2}$ decays to the shelf ($\sim 0.3\%$ detection efficiency compared with 5.3% branching ratio). This is significant because it means that shelving the ion is likely to involve scattering fewer photons than direct

fluorescence detection, limiting the amount of time for unwanted transitions between the dark and bright states. Since detection efficiencies are likely to be limited in implementations of scalable quantum computing, this may prove to be a significant benefit.

2.7 Two-qubit gates using trapped ions

The strength of the Coulomb force leads to ion-ion separations of order $1\text{ }\mu\text{m}$ – $10\text{ }\mu\text{m}$ in typical traps, which is orders of magnitude larger than the characteristic length scale of any direct coupling between electronic states. This makes directly entangling the states of ion qubits a near impossibility. The key to overcoming this limitation, originally understood by Cirac and Zoller [CZ95], is that one may use the internal-state-independent Coulomb interaction to couple ions' internal states indirectly, using their shared motional modes as a bus for quantum information.

We note that, in some senses, the lack of a direct internal state coupling is one of the chief benefits of trapped ion qubits; in many other systems the qubits interact constantly in ways that are very hard to control, leading to decoherence. In contrast, with trapped ions there are only the couplings we create.

The original Cirac and Zoller proposal involved first mapping the state of an ion onto the motional mode bus, before entangling the state of a second ion with the same bus mode. In this way the two ions are entangled using the motion as an intermediary. In practice, this is difficult to implement for several reasons: the ions must be addressed individually; the bus mode must begin in its ground state¹²; and, while the ion's state is mapped onto the bus mode, it is susceptible to loss of information due to interactions between the bus and the environment.

¹²Notice that this is a particularly strict requirement since any error in initialising the motional state directly produces a gate error. Thus, to produce a gate with an error of 1×10^{-4} one would have to ensure that the ions had a 99.99 % probability of being in their motional ground state. This would be *extremely* challenging experimentally.

Because of these complications, the Cirac and Zoller proposal was not experimentally demonstrated until 2003 [SKHR⁺03].

In the mean time, several two-qubit gates were proposed which used the spin-motion coupling in a more subtle way [SM99], [SM00], [MSJ00] & [SFZ99]. These proposals use a radiation field to exert a spin-dependent force on the ions. By controlling the phase of this force with respect to the ions' motion, we drive the multi-ion system around a closed circuit in the x - p phase space of one of its vibrational modes. During this process the ions' motion is transiently excited. When the circuit is completed the motion has returned to its original state; however the ions have accrued a phase that is proportional to the area enclosed by the circuit. Since the original force was dependent on the ions' collective spin state, so is the accumulated phase, leading to entanglement. This process may be seen as an application of the Berry geometric phase that accumulates when a system's Hamiltonian is adiabatically transported around a closed circuit in parameter space [Ber84].

The great advantage of these proposals is that by using an appropriate driving field, the force can be made to resemble a classical force with little dependence on the ions' motional state. This makes these gates insensitive to the ions' initial motional state, which greatly simplifies the experimental requirements.

There are two main ways one can produce the required force. Firstly, one can use spatially dependent A.C. Stark shifts to modulate the qubits' energy splitting at a frequency that is nearly resonant with a vibrational mode. I will refer to this as a $\sigma_z \otimes \sigma_z$ or "phase" gate. Secondly, one can drive motional sidebands of the qubit transition, which I will refer to as a $\sigma_\phi \otimes \sigma_\phi$ or Mølmer-Sørensen gate.

Phase gates are often preferred because they do not couple to the qubit transition directly, making them insensitive to changes in the qubit frequency, and because they may often be implemented with a simpler optical set-up when driven by lasers. However, as discussed previously, they are in general not ap-

plicable to clock states and so will not be considered further in this work. For a discussion of the relative merits of the two kinds of gate, see e.g. [Roo08].

2.7.1 Microwaves or lasers?

If a radiation field is to exert a force on an ion then it must change appreciably across the extent of the ion's wavepacket. For free-space radiation this change may be quantified by the Lamb-Dicke parameter, η , defined as

$$\eta := 2\pi q^0 / \lambda \quad (2.16)$$

where q^0 is the RMS extent of the ion's ground-state wavepacket and λ is the radiation's wavelength. To provide an effective spin-motion coupling, we typically want $\eta \sim 0.1$.

Since our qubit frequency is 3.2 GHz, one might suggest using microwave radiation to exert our force. However, for a $^{43}\text{Ca}^+$ ion confined in a trap with a secular frequency of 3 MHz, $q^0 = 6$ nm. This is orders of magnitude smaller than the 10 cm wavelength of 3.2 GHz radiation, leading to a negligible Lamb-Dicke parameter of $\eta = 4 \times 10^{-7}$. To get a Lamb-Dicke parameter of 0.1, we require radiation with a wavelength of 400 nm. Because of this vast discrepancy in Lamb-Dicke parameters, two-qubit gates have traditionally always been driven using lasers.

However, the advent of surface traps opened the possibility of trapping ions in the near-field regime tens of microns above a current-carrying waveguide. In this regime the field gradient depends on the waveguide geometry and not on the radiation's free-space wavelength. This allows a non-negligible coupling to the ion's motion, even at microwave frequencies. A two-qubit gate based on this effect was originally proposed in [OLA⁺08] and demonstrated in [OWC⁺11] (although we note that the idea can be found in earlier works such as [WMI⁺98]).

Using microwaves instead of lasers has several advantages: stable, low-noise sources are commercially available at these frequencies; control of the phase, frequency and amplitude is typically easier with electronic than optical sources; producing an interferometrically stable set-up is fundamentally easier at microwave frequencies than optical frequencies; waveguides may be integrated into trap structures in a fundamentally scalable way [AHB⁺13] [ALH⁺13]¹³; by incorporating microwave structures into the trap we eliminate pointing noise; and, since the qubit is not coupled to an excited state, errors due to photon scattering ([OIB⁺07]) are eliminated.

A further advantage of the near-field force is its high degree of insensitivity to the ions' motional state. A free-space optical radiation field exhibits non-linear structure over a length scale of $\lambda/2\pi$, which is typically only an order of magnitude larger than the spatial extent of the ions' wavepacket ($\eta \sim 0.1$). This non-linearity results in ions in different thermal states experiencing different forces, leading to gate infidelities. Because of this, ions must typically be cooled to near their motional ground state before two-qubit laser gates.

The first non-linear contribution to the radiation force is due to the field's third spatial derivative. For near-field radiation produced by typical electrode geometries, this contribution is suppressed by a factor of order $(q_0/d)^2$ relative to the linear term, where d is the ion-electrode separation¹⁴. This will be negligible in most cases, allowing high-fidelity gates to be achieved with ions which are only Doppler cooled. We note that q_0/d may be thought of as an effective Lamb-Dicke parameter for near-field radiation, so this thermal insensitivity may be seen as the up side of having a Lamb-Dicke parameter of order 10^{-4} .

Despite these potential advantages, microwave driving presents two major

¹³We note that significant progress has also been made towards integrating optical components into surface traps (see, e.g. [BEM⁺11]).

¹⁴Scaling factors used in this section are based on the field produced by a long, thin, straight current carrying wire, $B = \mu_0 I / (2\pi d)$ [OLA⁺08]. While the exact scaling will depend on the details of the trap geometry employed, this is expected to give an order of magnitude estimate in most cases.

challenges. The first is to produce a sufficient field gradient to allow two-qubit gates to be completed before heating and other noise mechanisms set in. The gradient is limited by the ion-electrode separation and tolerable power dissipation in trap structures. In general, we expect the gradient to increase with decreasing ion-trap separation as roughly d^{-2} . Thus, it is important to bring the ion as close to the trap's surface as possible. The extent to which one can do this is likely to be limited by anomalous heating, which typically scales as $\sim d^{-4}$. The expected error-per-gate due to trap heating is thus expected to scale as roughly d^{-2} . We note, however, that it may be possible to suppress this effect by cooling the trap to cryogenic temperatures [LGA⁺08] [CS13] or by employing surface cleaning techniques [AGH⁺11], [HCW⁺12], [DGB⁺13]. A further restriction on how close the ion can be brought to the trap surface arises from the decreasing width of the required trap electrodes. How thin one can make these electrodes will be limited by fabrication tolerances and power dissipation requirements.

The second challenge arises from the fact that strong field gradients are typically accompanied by strong fields. The presence of a strong field only a few megahertz detuned from the qubit transition would lead to off-resonant excitation of the qubit, producing significant light-shifts and potentially causing unwanted spin-flips. To reduce these effects to a tolerable level, a bias microwave field must be used to cancel the field at the ion's equilibrium position. In practice, suppression of the field by a factor of at least thousand is likely to be necessary, placing stringent requirements on the stability of the microwave system. We note, however, that the overall field at the ion scales roughly as d^{-1} , compared with the scaling d^{-2} of the gradient. Thus, the ratio of the field gradient to the field decreases with d^{-1} , providing further incentive to bring the ion as close to the trap's surface as possible. As discussed in the next section, this is essentially a further manifestation of the fact that the effective Lamb-Dicke parameter is $\tilde{\eta} \sim q^0/d$.

An alternative approach to RF-driven two-qubit gates was proposed in [MW01] and demonstrated in [JBT⁺09] and [KPS⁺12]. This proposal relies on the use of a large static magnetic field gradient to couple an RF field to the ions' motion. We note two potential disadvantages of this scheme. Firstly, the required magnetic field gradients produce large Zeeman splittings, which cause all qubits to operate at different frequencies. Whilst this allows spectral addressing of qubits, it introduces additional experimental complexity since the qubits' phases precess at an RF frequency relative to each other. This problem is likely to be particularly acute in architectures for trapped ion QIP which rely on shuttling ions around a two-dimensional trap. In such a system, the qubits would acquire phases that depend sensitively on how they were shuttled.

Secondly, this proposal cannot be implemented using magnetic field insensitive qubits, making it susceptible to field noise. A certain degree of robustness is typically gained by the use of qubits based on RF dressed states. However, this technique requires additional experimental overhead and coherence times are typically significantly shorter than achieved using clock qubits.

2.7.2 Microwave mediated spin-motion coupling

In this section we derive the Hamiltonian describing the coupling between a string of trapped ion qubits and an oscillating magnetic field. We keep the derivation general so that it applies to any qubit composed of two states separated by a magnetic dipole transition.

We begin by considering the displacements from equilibrium of a string of n ions along an axis $\hat{\mathbf{q}}$. Due the restrictions on trap geometry [All11], we are typically interested in the trap's radial directions. For a classical system the displacement of ion n from equilibrium is then [Jam98]

$$\mathbf{q}_n = \hat{\mathbf{q}} \sum_{j=1}^n b_{jn} q_j \quad (2.17)$$

where q_j is the amplitude of the ion chain's j^{th} normal mode and b_{jn} is its eigenvector, giving the phase and amplitude of motion for ion n .

We quantize the ions' motion by introducing the normal mode annihilation and creation operators, a_j and $a_j^\dagger$, such that $q_j = q_j^0 (a_j + a_j^\dagger)$. Here $q_j^0 = \sqrt{\hbar/(2m\omega_j)}$ and ω_j are the spatial extent of the ion's ground-state wavefunction and mode frequency for the j^{th} normal mode and m is the mass of each ion. We thus find that the quantized motion of the ions about their equilibrium positions may be expressed as

$$\mathbf{q}_n = \hat{\mathbf{q}} \sum_{j=1}^n b_{jn} q_j^0 (a_j + a_j^\dagger) \quad (2.18)$$

We next consider applying an oscillating magnetic field. We choose our axes such that the trap's symmetry axis lies along $\hat{\mathbf{x}}$ and the static B-field, which produces our quantization axis, lies along $\hat{\mathbf{z}}$. We assume the field to have no component along $\hat{\mathbf{x}}$ and further to have translational symmetry along $\hat{\mathbf{x}}$. We assume that the field may be expressed as a linear function of position in the region explored by the ions. Thus the field seen by the n^{th} ion is

$$\begin{aligned} \mathbf{B}(\mathbf{q}_n, t) &= \mathbf{B}(\mathbf{q}_n) \cos(\omega t - \phi_0) \\ &= (\mathbf{B}_n^0 + q_n \mathbf{B}') \cos(\omega t - \phi_0) \\ &= \{(B_{n,z}^0 + q_n B_z') \hat{\mathbf{z}} + (B_{n,y}^0 + q_n B_y') \hat{\mathbf{y}}\} \cos(\omega t - \phi_0) \end{aligned} \quad (2.19)$$

where $\mathbf{B}_n^0$ is the field amplitude at the ion's equilibrium position, $\mathbf{B}' := \frac{\partial \mathbf{B}}{\partial \mathbf{q}_n}$ is the field gradient along the normal mode axis and ω is the frequency of the applied radiation.

The Hamiltonian for the interaction between the ion and the oscillating field is¹⁵

$$\begin{aligned} H_n &= -\boldsymbol{\mu} \cdot \mathbf{B} \cos(\omega t - \phi_0) \\ &= -\{B_z \mu_z - B_{+1} \mu_{-1} - B_{-1} \mu_{+1}\} \cos(\omega t - \phi_0) \end{aligned} \quad (2.20)$$

We ignore all atomic states apart from the qubit states, which we label $|0\rangle$ and $|1\rangle$, taking $|1\rangle$ to represent the state with higher energy. In the case of the

¹⁵See B.3 for a discussion of the spherical basis used in the second line of equation 2.20.

clock qubit used in this thesis $|0\rangle := |4, 0\rangle$ and $|1\rangle := |3, +1\rangle$. We represent our state by the vector¹⁶

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \alpha|1\rangle + \beta|0\rangle \quad (2.21)$$

Since the Hamiltonian's matrix elements are real with appropriate choice of energy datum it may be expressed as

$$H_n = \frac{1}{2}(H_{11} - H_{00})\sigma_{z,n} + H_{10}\sigma_{x,n} \quad (2.22)$$

where H_{ij} are the Hamiltonian's matrix elements and $\sigma_{i,n}$ ($i = \{x, y, z, +, -\}$) is a Pauli operator acting on the state of the n^{th} ion.

The σ_z term, representing the differential Zeeman shift produced by B_z , is given by

$$\frac{1}{2}(H_{11} - H_{00}) = -\frac{1}{2}B_z (\langle 1|\mu_z|1\rangle - \langle 0|\mu_z|0\rangle) \cos(\omega t - \phi_0) \quad (2.23)$$

Note that this vanishes for any clock transition.

The σ_x term represents the coupling between qubit states due to the B-field. This arises from the polarization component of the B-field which matches the polarization of the qubit transition. Thus,

$$\begin{aligned} H_{10} &= (-1)^{Q+1} B_{-Q} \langle 1|\mu_Q|0\rangle \cos(\omega t - \phi_0) \\ &:= RB_{-Q} \cos(\omega t - \phi_0) \end{aligned} \quad (2.24)$$

where $Q := M_1 - M_0$ is the polarization of the qubit transition and R is the magnetic dipole matrix element for the qubit transition as defined in section 2.3.

We can simplify this notation somewhat by imagining our qubit as a pure spin- $\frac{1}{2}$ system with magnetic dipole moment $\tilde{\mu}$ interacting with a (fictitious) B-

¹⁶See B.2 for a discussion of the conventions I use to represent the ion's state in a two-dimensional sub-space as well as some useful results.

field $\tilde{\mathbf{B}}$. The elements of $\tilde{\boldsymbol{\mu}}$ and $\tilde{\mathbf{B}}$ are

$$\begin{aligned}\tilde{\mu}_x &:= -R \\ \tilde{\mu}_z &:= \frac{1}{2} (\langle 1|\mu_z|1\rangle - \langle 0|\mu_z|0\rangle) \\ \tilde{B}_x &:= B_{-Q} \\ \tilde{B}_y &:= 0 \\ \tilde{B}_z &:= B_z\end{aligned}\tag{2.25}$$

Our Hamiltonian then becomes

$$H_n = - \left(\tilde{B}_z \tilde{\mu}_z \sigma_{z,n} + \tilde{B}_x \tilde{\mu}_x \sigma_{x,n} \right) \cos(\omega t - \phi_0)\tag{2.26}$$

Using (2.19) to expand the B-field to first order in the ion's position, we find

$$\begin{aligned}H_n = & -\frac{1}{2} (e^{i(\omega t - \phi_0)} + e^{-i(\omega t - \phi_0)}) \left\{ \tilde{\mu}_x \tilde{B}_x^0 \sigma_{x,n} + \tilde{\mu}_z \tilde{B}_z^0 \sigma_{z,n} + \right. \\ & \left. \sum_{j=1}^n b_{jn} q_j^0 (a_j + a_j^\dagger) (\tilde{\mu}_x \tilde{B}_x' \sigma_{x,n} + \tilde{\mu}_z \tilde{B}_z' \sigma_{z,n}) \right\}\end{aligned}\tag{2.27}$$

We next move into the interaction picture with respect to the free energy of the ions' internal and motional states

$$H_0 = \frac{1}{2} \hbar \omega_0 \sigma_z + \sum_j \hbar \omega_j \left(\frac{1}{2} + a_j^\dagger a_j \right)\tag{2.28}$$

where ω_0 is the qubit frequency. We define the detuning of the B-field, δ , such that $\omega := \omega_0 + \delta$ and make a rotating wave approximation, dropping all terms rotating at $2\omega_0$. Our Hamiltonian becomes

$$\begin{aligned}H_n = & \frac{1}{2} \hbar \Omega_x e^{-i(\delta t - \phi_0)} \sigma_{+,n} + \frac{1}{2} \hbar \Omega_z e^{-i(\omega t - \phi_0)} \sigma_{z,n} \\ & + \frac{1}{2} \hbar \sum_{j=1}^n (\Omega_{x,jn} e^{-i(\delta t - \phi_0)} \sigma_{+,n} + \Omega_{z,jn} e^{-i(\omega t - \phi_0)} \sigma_{z,n}) (a_j e^{-i\omega_j t} + a_j^\dagger e^{i\omega_j t}) \\ & + \text{H.C.}\end{aligned}\tag{2.29}$$

where we have defined the Rabi frequencies

$$\begin{aligned}\hbar \Omega_x &:= -\tilde{\mu}_x \tilde{B}_x^0 = R B_{-Q}^0 \\ \hbar \Omega_z &:= -\tilde{\mu}_z \tilde{B}_z^0 = -\frac{1}{2} (\langle 1|\mu_z|1\rangle - \langle 0|\mu_z|0\rangle) B_z^0 \\ \hbar \Omega_{x,jn} &:= -\tilde{\mu}_x \tilde{B}_x' b_{jn} q_j^0 = R B_{-Q}' b_{jn} q_j^0 \\ \hbar \Omega_{z,jn} &:= -\tilde{\mu}_z \tilde{B}_z' b_{jn} q_j^0 = -\frac{1}{2} (\langle 1|\mu_z|1\rangle - \langle 0|\mu_z|0\rangle) B_z' b_{jn} q_j^0\end{aligned}\tag{2.30}$$

The microwave coupling Hamiltonian of eq (2.29) contains both σ_x terms, representing a coupling between the qubit states, and σ_z terms representing a differential Zeeman shift on the qubit transition. The first pair of terms act only on the ions' internal states and so drive global single-qubit rotations. The second pair, however, couple the spins to the ions' motion and can be used to create multi-ion entanglement.

Either of the two of the spin-motion terms could, in principle, be used to entangle ions, producing $\sigma_\phi \otimes \sigma_\phi$ or $\sigma_z \otimes \sigma_z$ couplings. However, driving the σ_z terms in (2.29) requires radiation nearly resonant with the ions' vibrational modes. This is problematic since the large currents required to produce the magnetic field gradient are likely to be accompanied by non-negligible voltages. These voltages will couple to the ions' charge, creating a strong, state-independent force, leading to significant heating. Moreover, as discussed previously, phase gates can only be used on magnetic-field sensitive transitions. Since these gates are likely to be relatively slow for currently realisable field gradients, this poses a potentially serious source of infidelity.

We will thus be concerned with the $\sigma_\phi \otimes \sigma_\phi$ coupling, driven by radiation close to resonance with the qubit's motional sideband transitions. In this case the σ_z terms are far from resonance and can be ignored, so our Hamiltonian becomes

$$\begin{aligned}
H_n = & \frac{1}{2}\hbar\Omega_x e^{-i(\delta t - \phi_0)}\sigma_{+,n} + \frac{1}{2}\hbar \sum_{j=1}^n \Omega_{x,jn} e^{-i(\delta t - \phi_0)}\sigma_{+,n} \left(a_j e^{-i\omega_j t} + a_j^\dagger e^{i\omega_j t} \right) \\
& + \text{H.C.}
\end{aligned}
\tag{2.31}$$

As discussed previously, since we are not using free-space radiation there is no well-defined Lamb-Dicke parameter for this interaction. It is nonetheless useful to define an effective Lamb-Dicke parameter, $\tilde{\eta}$, as the ratio of the strengths of the motional and carrier coupling terms in our Hamiltonian. This definition reflects the corresponding result for laser gates, which holds to first order in η .

Thus,

$$\begin{aligned}
\tilde{\eta} &:= \frac{\Omega_{x,jn}}{\Omega_x} \\
&= \frac{B'_{-Q}}{B^0_{-Q}} b_{jn} q_j^0 \\
&\sim q_j^0/d
\end{aligned} \tag{2.32}$$

where the last line represents an order of magnitude scaling estimate based on the field around a long, thin current-carrying wire as discussed in the previous section. For $d = 100 \text{ um}$, we expect $\tilde{\eta} \sim 1 \times 10^{-4}$. Thus, to achieve an effective Lamb-Dicke parameter of 0.1, comparable to laser-driven gates, one must use separate bias fields to reduce the carrier term by a factor of roughly 1000.

2.7.3 The Mølmer-Sørensen force

The Mølmer-Sørensen gate uses a bichromatic radiation field with components near the qubit's first red and blue motional sideband transitions to create a spin-dependent force. In this section we derive the Hamiltonian for this interaction, as well as an exact expression for the propagator describing the two-qubit Mølmer-Sørensen gate in the absence of errors, such as heating. We assume that there is no oscillating B-field at the ions' position (it has been "nulled") so that $\Omega_x = 0$. We further assume that the interaction only couples to a single normal mode with frequency ω_j and that the two sidebands have the same Rabi frequency. Finally, we make the rotating wave approximation, discarding all terms that are off-resonant by twice the ions' vibrational frequency.

We begin with a pair of oscillating fields with frequencies $\omega = \omega_0 \pm \omega_j \pm \delta$ and phases ϕ_r and ϕ_b . We will refer to these as "red" (−) and "blue" (+). The Hamiltonians for the two sideband interactions are

$$\begin{aligned}
H_{\text{BSB},n} &= \frac{1}{2} \hbar \Omega_n \sigma_{+,n} \left(a^\dagger e^{-i(\delta t - \phi_b)} + a e^{-i(2\omega_j t + \delta t - \phi_b)} \right) + \text{H.C.} \\
H_{\text{RSB},n} &= \frac{1}{2} \hbar \Omega_n \sigma_{+,n} \left(a e^{i(\delta t + \phi_r)} + a^\dagger e^{i(2\omega_j t + \delta t + \phi_r)} \right) + \text{H.C.}
\end{aligned} \tag{2.33}$$

where Ω_n is the Rabi frequency for the red and blue sideband couplings to ion

n , defined by

$$\begin{aligned}\hbar\Omega_n &:= \hbar\Omega_{x,jn} \\ &:= RB'_{-Q}b_{jn}q_j^0\end{aligned}\quad (2.34)$$

Discarding the far from resonant terms at $2\omega_j$, the sideband Hamiltonians become

$$\begin{aligned}H_{\text{BSB},n} &= \frac{1}{2}\hbar\Omega_n\sigma_{+,n}a^\dagger e^{-i(\delta t - \phi_b)} + \frac{1}{2}\hbar\Omega_n\sigma_{-,n}ae^{i(\delta t - \phi_b)} \\ H_{\text{RSB},n} &= \frac{1}{2}\hbar\Omega_n\sigma_{+,n}ae^{i(\delta t + \phi_r)} + \frac{1}{2}\hbar\Omega_n\sigma_{-,n}a^\dagger e^{-i(\delta t + \phi_r)}\end{aligned}\quad (2.35)$$

The Hamiltonian for the bichromatic field is thus

$$\begin{aligned}H_{\text{MS},n} &= H_{\text{BSB},n} + H_{\text{RSB},n} \\ &= \frac{1}{2}\hbar\Omega_n\left(\sigma_{+,n}e^{i\bar{\phi}} + \sigma_{-,n}e^{-i\bar{\phi}}\right)\left(ae^{i\delta t}e^{-\frac{1}{2}i\delta\phi} + a^\dagger e^{-i\delta t}e^{\frac{1}{2}i\delta\phi}\right) \\ &= \frac{1}{2}\hbar\Omega_n\sigma_{\bar{\phi},n}\left(ae^{i\delta t}e^{-\frac{1}{2}i\delta\phi} + a^\dagger e^{-i\delta t}e^{\frac{1}{2}i\delta\phi}\right)\end{aligned}\quad (2.36)$$

where we have defined the radiation field's mean phase and phase difference as

$$\begin{aligned}\bar{\phi} &:= \frac{1}{2}(\phi_b + \phi_r) \\ \delta\phi &:= (\phi_b - \phi_r)\end{aligned}\quad (2.37)$$

and $\sigma_{\bar{\phi},n} := \sigma_{x,n} \cos \bar{\phi} - \sigma_{y,n} \sin \bar{\phi}$.

The combined Hamiltonian for the N -ion system is a sum of such single-ion Hamiltonians, given by

$$H_{\text{MS}} = \frac{1}{2} \sum_{n=1}^N \hbar\Omega_n \sigma_{\bar{\phi},n} \left(ae^{i\delta t} e^{-\frac{1}{2}i\delta\phi} + a^\dagger e^{-i\delta t} e^{\frac{1}{2}i\delta\phi} \right) \quad (2.38)$$

For a system of two ions this becomes

$$H_{\text{MS}} = \frac{1}{2} \hbar\Omega S_{\bar{\phi}} \left(ae^{i\delta t} e^{-\frac{1}{2}i\delta\phi} + a^\dagger e^{-i\delta t} e^{\frac{1}{2}i\delta\phi} \right) \quad (2.39)$$

where we have defined the two-ion Rabi frequency and collective spin operators, Ω and $S_{\bar{\phi}}$, by

$$\begin{aligned}\Omega &:= \hbar\Omega_{x,jn} \text{sgn}(b_{jn}) = RB'_{-Q}q_j^0|b_{jn}| \\ S_{\bar{\phi}} &:= \text{sgn}(b_{j1})\sigma_{\bar{\phi},1} + \text{sgn}(b_{j2})\sigma_{\bar{\phi},2}\end{aligned}\quad (2.40)$$

Equation (2.39) describes the 2-qubit Mølmer-Sørensen gate in the absence of errors. When this Hamiltonian is applied to an eigenstate of the collective spin operator it corresponds to a harmonic oscillator driven by a classical force, whose phase and amplitude depend on the spin state¹⁷. We can integrate the Schrödinger equation for this Hamiltonian exactly using the Magnus expansion¹⁸. The commutators for the Hamiltonian at different times are given by

$$\begin{aligned} [H_{MS}(t), H_{MS}(t')] &= \frac{1}{2}i \left(\hbar \Omega S_{\bar{\phi}} \right)^2 \sin[\delta(t - t')] \\ [H_{MS}(t), [H_{MS}(t'), H_{MS}(t'')]] &= 0 \end{aligned} \quad (2.41)$$

so only the first two terms in the Magnus expansion are non-vanishing. The exact propagator is thus given by

$$U_{MS} = D(S_{\bar{\phi}}\alpha) \exp(-i\Phi(t)) \quad (2.42)$$

where

$$\begin{aligned} \alpha &:= -i\frac{\Omega}{\delta}e^{-\frac{1}{2}i\delta t}e^{\frac{1}{2}i\delta\phi}\sin\left(\frac{1}{2}\delta t\right) \\ \Phi(t) &:= \left(\frac{\Omega S_{\bar{\phi}}}{2\delta}\right)^2(\delta t - \sin\delta t) \end{aligned} \quad (2.43)$$

and D is the phase space displacement operator, defined by

$$D(\alpha) := \exp\left(\alpha a^\dagger - \alpha^* a\right) \quad (2.44)$$

The action of this propagator is most obvious in the eigenbasis of $S_{\bar{\phi}}$ for which

$$U_{MS} = \sum_{\lambda} D(\lambda\alpha) \exp(-i\lambda^2\Phi_0(t)) |\lambda\rangle\langle\lambda| \quad (2.45)$$

where

$$\Phi_0(t) := \left(\frac{\Omega}{2\delta}\right)^2(\delta t - \sin\delta t) \quad (2.46)$$

We see that different spin states are driven around different paths in phase space, depending on the phase and magnitude of their eigenvalue. As they traverse these paths, they accumulate a geometric phase of $\lambda^2\Phi_0(t)$. In general,

¹⁷See, e.g. [CN65], for a review of driven quantum harmonic oscillators.

¹⁸See Appendix B.4 for a review of the Magnus expansion.

this propagator entangles the ions' spins with their motion; however, at times $\tau = \frac{2n\pi}{\delta}$ the phase-space paths form closed loops, leaving the ions' spins and motion in a separable state. The propagator then reduces to

$$U_{\text{MS}} = \sum_{\lambda} \exp[-i\lambda^2\Phi_0(2n\pi/\delta)] |\lambda\rangle\langle\lambda| \quad (2.47)$$

The ions have thus accumulated a geometric phase that is dependent on their initial spin state. Importantly, however, this phase is independent of their initial motional state.

As an example of how this interaction can be used to implement a two-qubit gate, we consider the simplest case where

$$S_{\bar{\phi}} = \sigma_{x,1} + \sigma_{x,2} \quad (2.48)$$

Here, the closed loop propagator becomes

$$U_{\text{MS}} = \exp[-4i\Phi_0(2n\pi/\delta)] \left\{ |++\rangle\langle++| + |--\rangle\langle--| \right\} + \left\{ |+-\rangle\langle+-| + |-+\rangle\langle-+| \right\} \quad (2.49)$$

where

$$|\pm\rangle = 1/\sqrt{2} (|1\rangle \pm |0\rangle) \quad (2.50)$$

are the eigenstates of σ_x with corresponding eigenvalues of $\lambda_{\pm} = \pm 1$.

Choosing our detuning so that $\delta^2/\Omega^2 = 4n$ ensures that $4\Phi_0(2n\pi/\delta) = \pi/2 (S_{\bar{\phi}})^2$ so our propagator becomes

$$U_{\text{MS}} = \left\{ |+-\rangle\langle+-| + |-+\rangle\langle-+| \right\} - i \left\{ |++\rangle\langle++| + |--\rangle\langle--| \right\} \quad (2.51)$$

which represents an entangling gate closely related to the controlled-phase gate. By application of suitable single-qubit operations, this can be used to implement a controlled-NOT gate in the computational basis (see, e.g. [SM99]).

3

Noise Master Equations

Random fluctuations of classical parameters, such as Rabi frequencies and detunings, are one of the most pervasive difficulties in experimental quantum computing. In this chapter we develop a simple master equation-based approach, which will allow us to model and understand some of the important noise sources present in the experiments described in this thesis.

In general, understanding the dynamics of even relatively simple systems subject to noise is a difficult task. There are, however, three particular limiting cases in which the problem is greatly simplified. Firstly, when the noise is small enough that it does not significantly alter the system's trajectory and the Born approximation may be applied.

Secondly, when the correlation time of the noise process is short compared with the time-scale of the system's dynamics. In this case, we may approximate the noise as white noise, with a flat power spectrum and infinite bandwidth. Finally, when the noise correlation time is long compared with the time we wish to follow the system for. In this case, the system's evolution may be understood by assuming a constant, random offset in the fluctuating parameter. If we observe an ensemble of identical systems with independent fluctuations then the noise will be evident as shot-to-shot variations over the ensemble.

If, on the other hand, the noise's correlation time is comparable to the time-

scale of the system's evolution and, moreover, if the noise produces a significant change to the system's trajectory, then Monte-Carlo simulations must generally be performed.

Our focus in this chapter will be the first two of these cases: white noise and the small noise approximation. We begin in section 3.2 with a treatment of idealised white noise based on stochastic calculus. Next, in section 3.3, we provide a perturbative treatment of fluctuations with arbitrary power spectra. The final two sections of the chapter, sections 3.4 and 3.5, then apply the previously derived white noise master equation to two cases of particular importance: the two-state system and the harmonic oscillator. Our treatment of Rabi flopping on a two-state system in section 3.4 will cover white noise in the driving field's amplitude, phase and frequency. We derive master equations for these cases, which we integrate analytically to obtain expressions for the fidelity of a π -pulse in the presence of noise.

Our treatment of the noisy quantum harmonic oscillator in section 3.4 covers both dephasing due to oscillator frequency fluctuations and heating due a stochastic driving force. These two problems are of practical importance to us because of the use of the ions' motional modes during two-qubit gates. Gates based on trap radial modes, such as the microwave-driven gates considered in this thesis, are likely to be particularly prone to dephasing due to small amounts of noise on the trap RF. To understand this dephasing process, we derive expressions for the relation between the RF's AM and PM noise spectra and the dephasing terms in our master equation.

We use the results derived here in chapters 5 and 6 to develop models of sources of experimental infidelity during single-qubit randomized benchmarking experiments and two-qubit gates. The material in this chapter is supplemented by appendix C, which presents an overview of some basic aspects of modulation theory. This provides the link between the autocorrelation functions

used in the master equations and the noise spectra of the fields they represent.

The noise terms considered in this chapter will be classical, reflecting the fact that the experiments we perform do not involve coupling ions to fields containing a small number of quanta. This allows some simplification of our treatment. Nonetheless, the equations we derive are consistent with the relevant limits of quantum optics treatments in which fluctuations arise from couplings to quantised fields (see, e.g. [TTM⁺00] & [WM08]).

The stochastic treatments of fluctuations presented here are, of course, far from original and draw on a long history of work. Early treatments of quantum master equations describing Hamiltonians with multiplicative classical fluctuations can be found, for example, in [Aga78], [Aga76] and [Ebe76]. The treatment of multiplicative white noise using stochastic differential equations (SDEs) presented here is essentially that of Dyrting and Milburn [DM96] and Schneider and Milburn [SM98]. Readers unfamiliar with the stochastic methods used in this section will find an excellent introduction in the classic text [Gar90].

3.1 The Stochastic Liouville-von Neumann Equation

Our focus in this chapter is on quantum systems subject to multiplicative noise in a classical parameter, $\xi(t)$, with Hamiltonians of the form

$$H_T(t) = H_1(t) + \xi(t)H_2(t) \quad (3.1)$$

where $\xi(t)$ represents a real valued, zero mean stationary process.

The system's density matrix evolves according to the stochastic Liouville-von Neumann equation

$$i\hbar d\rho = [H_1, \rho] dt + \xi[H_2, \rho] dt \quad (3.2)$$

This equation describes the evolution of the density matrix under a particular realisation of the noise process $\xi(t)$. Experimentally, however, one can only

measure components of the density matrix by repeating an experiment many times and averaging the results. Our aim through this section will thus be to derive master equations describing the evolution of the density matrix averaged over the fluctuations.

3.2 White noise

We begin then by considering the limit where the fluctuations described by ξ have a flat frequency spectrum stretching out to a bandwidth that is large compared with the frequency at which the system can respond. In this limit, one may consider ξ to be a white noise process characterised by

$$\begin{aligned} \langle \xi(\tau)\xi(0) \rangle &= 0 & (\tau \neq 0) \\ \int_{-\infty}^{\infty} \langle \xi(\tau)\xi(0) \rangle d\tau &= \Gamma \end{aligned} \quad (3.3)$$

We may thus identify $\langle \xi(\tau)\xi(0) \rangle = \Gamma \delta(t)$ and interpret ξdt as the increment of the Wiener process, so that

$$\xi(\tau) dt \mapsto \sqrt{\Gamma} dW(t) \quad (3.4)$$

Equation (3.2) then becomes the Stratonovich stochastic differential equation (SDE)

$$i\hbar d\rho = [H_1, \rho] dt + \sqrt{\Gamma} [H_2, \rho] \circ dW \quad (3.5)$$

Averaging over realisations of ξ in (3.5) is complicated by the fact that for Stratonovich SDEs, $dW(t)$ and $\rho(t)$ are not statistically independent. We thus convert (3.5) to the Ito form before averaging. To do this, we assume that $\rho(t)$ is the solution to an Ito SDE of the form

$$d\rho(t) = A(\rho(t), t) dt + B(\rho(t), t) dW \quad (3.6)$$

where A and B are operator valued functions of $\rho(t)$ and t .

We recall that the Stratonovich integral is defined by

$$\int_{t_0}^t G(\rho(t'), t') \circ dW(t') \stackrel{\text{ms-lim}}{n \rightarrow \infty} \sum_{i=1}^n G\left\{\frac{1}{2}[\rho(t_i) + \rho(t_{i-1})], t_{i-1}\right\} [W(t_i) - W(t_{i-1})] \quad (3.7)$$

where the equality holds in the mean square limit [Gar90]. Now,

$$d\rho(t) := \rho(t + dt) - \rho(t) \quad (3.8)$$

so

$$G\left\{\frac{1}{2}[\rho(t_i) + \rho(t_{i-1})], t_{i-1}\right\} = G\left\{\rho(t_{i-1}) + \frac{1}{2}d\rho(t_{i-1}), t_{i-1}\right\} \quad (3.9)$$

Since all functions are now evaluated at the same time, we will stop labelling times from here on for the sake of clarity. In our case,

$$i\hbar G(\rho(t), t) = \sqrt{\Gamma}[H_2, \rho] \quad (3.10)$$

so

$$\begin{aligned} i\hbar G\left(\rho + \frac{1}{2}d\rho, t\right) &= \sqrt{\Gamma}[H_2, \rho + \frac{1}{2}d\rho] \\ &= \sqrt{\Gamma}[H_2, \rho] + \frac{1}{2}\sqrt{\Gamma}[H_2, A] dt + \frac{1}{2}\sqrt{\Gamma}[H_2, B] dW \end{aligned} \quad (3.11)$$

Making the usual identifications that

$$\begin{aligned} dW^2(t) &= dt \\ dt^2(t) &= 0 \\ dt(t) dW(t) &= 0 \end{aligned} \quad (3.12)$$

we find

$$\sqrt{\Gamma}[H_2, \rho] \circ dW = \sqrt{\Gamma}[H_2, \rho] dW + \frac{1}{2}\sqrt{\Gamma}[H_2, B] dt \quad (3.13)$$

Comparing (3.6) with (3.5) and (3.13), we arrive at the Ito SDE

$$i\hbar d\rho(t) = [H_1, \rho] dt - \frac{1}{2}\frac{i}{\hbar}\Gamma[H_2, [H_2, \rho]] dt + \sqrt{\Gamma}[H_2, \rho] dW \quad (3.14)$$

Now $dW(t)$ is statistically independent from $\rho(t)$ ¹, so we may average over realisations of (3.14) to arrive at our master equation

$$i\hbar \frac{d\langle \rho \rangle}{dt} = [H_1, \langle \rho \rangle] - \frac{1}{2}\frac{i}{\hbar}\Gamma[H_2, [H_2, \langle \rho \rangle]] \quad (3.15)$$

¹Here, we assume that H_2 is non-anticipating.

To simplify our notation, from here on we will drop the expectation brackets around ρ . It should be understood that we will always be referring to the density matrix averaged over the fluctuations and *not* to its evolution over any particular realisation of the noise process.

It is often convenient to express our master equation in the Lindblad form [Lin76], which may be done by expanding the commutator in (3.15). Our master equation then becomes

$$\begin{aligned}\frac{d\rho}{dt} &= -\frac{i}{\hbar}[H_1, \rho] - \frac{1}{2\hbar^2}\Gamma \{H_2 H_2 \rho - 2H_2 \rho H_2 + \rho H_2 H_2\} \\ &= -\frac{i}{\hbar}[H_1, \rho] + L\rho L^\dagger - \frac{1}{2}(L^\dagger L\rho + \rho L^\dagger L)\end{aligned}\tag{3.16}$$

where $L := \frac{\sqrt{\Gamma}}{\hbar} H_2$ and we have used the fact that $H_2 = H_2^\dagger$.

3.3 Noise with arbitrary correlations

If the noise spectrum cannot be approximated as white then one must take account of the details of its correlation function. We begin with the stochastic Liouville-von Neumann equation

$$i\hbar\dot{\rho}(t) = [H_1(t), \rho(t)] + \xi(t)[H_2(t), \rho(t)]\tag{3.17}$$

We now move into an interaction picture with respect to the deterministic component of the dynamics, H_1 , by making the transformation

$$\rho \mapsto \tilde{\rho} := U^\dagger \rho U\tag{3.18}$$

where $U(t)$ is the propagator corresponding to H_1 , defined as the solution to the differential equation

$$\frac{dU}{dt} = -\frac{i}{\hbar}H_1 U\tag{3.19}$$

that satisfies $U(t) = 1$ (see section B.4). The system then evolves according to

$$\begin{aligned}i\hbar\dot{\tilde{\rho}} &= i\hbar U^\dagger \dot{\rho} U - U^\dagger [H_1, \rho] U \\ &= \xi(t) U^\dagger [H_2, \rho] U \\ &= \xi(t) [\tilde{H}_2, \tilde{\rho}]\end{aligned}\tag{3.20}$$

where we have defined the interaction picture Hamiltonian, $\tilde{H}_2$, as

$$\tilde{H}_2 := U^\dagger H_2 U \quad (3.21)$$

From here on, we will drop the tildes for convenience.

Next, we formally integrate (3.20) so that

$$\rho(t) = \rho(0) - \frac{i}{\hbar} \int_0^t \xi(t') [H_2(t'), \rho(t')] dt' \quad (3.22)$$

By repeatedly substituting the solution (3.22) into itself, we arrive at the series solution

$$\begin{aligned} \rho(t) = & \rho(0) + \sum_{n=1}^{\infty} \left(-\frac{i}{\hbar}\right)^n \int_0^t dt_1 \int_0^{t_1} dt_2 \cdots \int_0^{t_{n-1}} dt_n \xi(t_1) \cdots \xi(t_n) \\ & \times [H_2(t_1), [H_2(t_2), [\dots, \rho(0)]]] \end{aligned} \quad (3.23)$$

Averaging over the fluctuations, this becomes

$$\begin{aligned} \langle \rho(t) \rangle = & \rho_0(t) + \sum_{n=2}^{\infty} \left(-\frac{i}{\hbar}\right)^n \int_0^t dt_1 \int_0^{t_1} dt_2 \cdots \int_0^{t_{n-1}} dt_n \langle \xi(t_1) \cdots \xi(t_n) \rangle \\ & \times [H_2(t_1), [H_2(t_2), [\dots, \rho(0)]]] \end{aligned} \quad (3.24)$$

where the first-order term has vanished because of the assumption that $\xi(t)$ is a zero-mean process. Once again, we drop the expectation brackets around ρ .

We now make the Born approximation that the effect of the noise over the duration of the integration is sufficiently small that we may neglect terms in (3.24) of higher than second order in ξ . This is equivalent to making the approximation that we may evaluate the impact of the noise term using the density matrix that would be obtained in the absence of noise. We thus arrive at the interaction picture master equation

$$\rho(t) = \rho(0) - \frac{1}{\hbar^2} \int_0^t dt_1 \int_0^{t_1} dt_2 G(t_1 - t_2) [H_2(t_1), [H_2(t_2), \rho(0)]] \quad (3.25)$$

where, as usual, $G(\tau)$ is the noise autocorrelation function, defined by

$$G(\tau) := \langle \xi(t + \tau) \xi(t) \rangle \quad (3.26)$$

3.4 Three faces of white noise

3.4.1 White amplitude noise

We begin by considering Rabi flopping due to a monochromatic field with fluctuating amplitude. In the frame co-rotating with the driving field and after making the rotating wave approximation, we have

$$\begin{aligned} H &= -\frac{1}{2}\hbar\delta\sigma_z + \frac{1}{2}\hbar\Omega(t)\sigma_x \\ &= -\frac{1}{2}\hbar\delta\sigma_z + \frac{1}{2}\hbar\Omega_0\sigma_x + \frac{1}{2}\hbar\Delta\Omega(t)\sigma_x \end{aligned} \quad (3.27)$$

where, as usual, Ω_0 and δ are the mean Rabi frequency and detuning in angular units and $\Delta\Omega(t)$ represents the fluctuating component of the Rabi frequency.

We assume that the Rabi flopping is driven by a field with amplitude V , such that

$$\Omega(t) \propto V(t) \quad (3.28)$$

thus

$$\frac{\Delta\Omega(t)}{\Omega_0} = \frac{\Delta V(t)}{V_0} \quad (3.29)$$

From (C.24), we have

$$\begin{aligned} G(\tau) &:= \langle \Delta\Omega(t+\tau)\Delta\Omega(t) \rangle \\ &= \Omega_0^2 \left\langle \frac{\Delta V(t+\tau)\Delta V(t)}{V_0^2} \right\rangle \\ &= \frac{\Omega_0^2}{P_0} \int_{-\infty}^{\infty} e^{-2\pi i f \tau} N_0(f) df \end{aligned} \quad (3.30)$$

where P_0 is the power in the carrier and $N_0(f)$ is the single-sideband AM noise power spectral density. In the limit where the noise is white, we have

$$G(\tau) = \Omega_0^2 \frac{N_0}{P_0} \delta(\tau) \quad (3.31)$$

where δ is the Dirac delta function.

Substituting (3.27) and (3.31) into (3.15), we arrive at the master equation

$$\frac{d\rho}{dt} = \frac{1}{2}i\delta[\sigma_z, \rho] - \frac{1}{2}i\Omega_0[\sigma_x, \rho] - \frac{N_0}{8P_0}\Omega_0^2[\sigma_x, [\sigma_x, \rho]] \quad (3.32)$$

To illustrate the effect of the depolarizing term in this master equation, we consider the simple case of resonant Rabi flopping, beginning in the state $|0\rangle$. Integrating (3.32) we find that

$$\rho(t) = \begin{pmatrix} \frac{1}{2} \left(1 - e^{-\frac{N_0 t \Omega_0^2}{2P_0}} \cos(t\Omega_0) \right) & -\frac{1}{2} i e^{-\frac{N_0 t \Omega_0^2}{2P_0}} \sin(t\Omega_0) \\ \frac{1}{2} i e^{-\frac{N_0 t \Omega_0^2}{2P_0}} \sin(t\Omega_0) & \frac{1}{2} \left(1 + e^{-\frac{N_0 t \Omega_0^2}{2P_0}} \cos(t\Omega_0) \right) \end{pmatrix} \quad (3.33)$$

To first order in N_0/P_0 , the error after n π -pulses ($\Omega_0 t = n\pi$) is given by

$$1 - P_1 = n\pi \frac{N_0}{4P_0} \Omega_0 \quad (3.34)$$

This result reflects the fact that in the case of resonant Rabi flopping, the fidelity of a π -pulse depends only on the total pulse area, $A := \int \Omega(t) dt$. Given a pulse area of $A = n\pi + \Delta A$ ($\Delta A \ll 1$), we expect an error of

$$1 - P_1 = \frac{1}{4} (\Delta A)^2 \quad (3.35)$$

Since our fluctuations in the Rabi frequency are of the form

$$\Delta\Omega dt = \sqrt{\Omega_0^2 \frac{N_0}{P_0}} dW \quad (3.36)$$

we expect Gaussian fluctuations in the pulse area with variance

$$\begin{aligned} \langle \Delta A(t)^2 \rangle &= t \Omega_0^2 \frac{N_0}{P_0} \\ &= n\pi \Omega_0 \frac{N_0}{P_0} \end{aligned} \quad (3.37)$$

We see that the RMS pulse area error increases as the square root of n , as expected for a random walk process. For a fixed mean pulse area, the error also increases linearly with Ω_0 , reflecting the fact that the faster we perform our π -pulses, the larger the noise bandwidth we are sensitive to.

More generally, one can show that if the system begins in the pure state

$$|\psi(t=0)\rangle = \cos\left(\frac{\theta_i}{2}\right)|0\rangle + e^{i\phi_i} \sin\left(\frac{\theta_i}{2}\right)|1\rangle \quad (3.38)$$

then the infidelity in the density matrix after a time $t = n\pi$ is ²

$$1 - F = \frac{1}{2} (1 - \cos^2 \phi_i \sin^2 \theta_i) \left(1 - e^{-n\pi \frac{N_0}{2P_0} \Omega_0} \right) \quad (3.39)$$

Integrating over all possible starting states, we find that the average infidelity of a series of π -rotations from random initial states is

$$\langle 1 - F \rangle = \frac{1}{3} \left(1 - e^{-n\pi \frac{N_0}{2P_0} \Omega_0} \right) \quad (3.40)$$

3.4.2 Frequency and phase noise

Fluctuations in the frequency/phase of the driving field or system can often be approximated using one of two simple models. Firstly, one can imagine a broadened carrier with a Lorentzian lineshape. Secondly, one can imagine a pure carrier (represented by a delta function in the frequency domain), superposed on a white noise pedestal.

In the former case, the phase noise power spectral density (PSD) decreases with $1/f^2$ for frequencies larger than the Lorentzian linewidth (3 dB frequency). This may be thought of as white noise in the driving field's frequency, which produces a random walk in its phase (often referred to as Wiener-Lévy phase diffusion).

In general, the phase noise spectrum in any experiment will contain both a white noise region and a random-walk region. If the source is a microwave/RF signal generator then the linewidth of the Lorentzian region of the spectrum will typically be < 1 mHz. In this case, the random-walk fluctuations in phase over any experiment will typically be negligible, so it is reasonable to model the fluctuations as entirely arising from the white phase noise region of the spectrum.

If, on the other hand, the source is a laser³ then the Lorentzian linewidth will

²We use the standard definition of fidelity whereby $F := \text{Tr}(\rho(t) \cdot \rho_{\text{ideal}}(t))$, where ρ_{ideal} is the density matrix one would obtain in the absence of noise.

³Notice that if the transition in question is a stimulated Raman transition, then the relevant phase is the phase offset between the two beams. In general, this is produced by modulating a single laser source at microwave/RF frequencies. For noise purposes, we thus think of Raman transitions as a special case of a microwave/RF source.

typically not be negligible – although lasers with linewidths of $\lesssim 40$ mHz have been constructed [KHG⁺12], linewidths of lasers used to drive optical qubits are typically ~ 20 Hz [BKRB08b]. In this case it will often suffice to model the noise as entirely a phase diffusion process, neglecting smaller fast fluctuations.

3.4.3 White phase noise

We now consider Rabi flopping due to a field with constant amplitude and white Gaussian phase fluctuations. In the interaction picture, and making the rotating wave approximation, our Hamiltonian is

$$H = \frac{1}{2}\hbar\Omega e^{-i(\delta t + \phi(t))}\sigma_+ + \frac{1}{2}\hbar\Omega e^{i(\delta t + \phi(t))}\sigma_- \quad (3.41)$$

where $\phi(t)$ represents the phase fluctuations.

This Hamiltonian is a non-linear function of the phase fluctuations. To linearise it, we assume that $\phi(t) \ll 1$ and expand to first order in the noise

$$H \simeq \frac{1}{2}\hbar\Omega \left(e^{-i\delta t}\sigma_+ + e^{i\delta t}\sigma_- \right) + \frac{1}{2}i\hbar\Omega\phi(t) \left(e^{i\delta t}\sigma_- - e^{-i\delta t}\sigma_+ \right) \quad (3.42)$$

Taking the white noise limit of (C.56), we have

$$\begin{aligned} G(\tau) &:= \langle \phi(t + \tau)\phi(t) \rangle \\ &= \frac{N_0}{P_c} \int_{-\infty}^{\infty} e^{-2\pi i f \tau} df \\ &= \frac{N_0}{P_c} \delta(\tau) \end{aligned} \quad (3.43)$$

To illustrate the effect of white phase noise, we consider the simplified case of resonant Rabi-flopping ($\delta = 0$) where

$$H = \frac{1}{2}\hbar\Omega\sigma_x + \frac{1}{2}\hbar\Omega\phi(t)\sigma_y \quad (3.44)$$

Our white-noise master equation is thus

$$\frac{d\rho}{dt} = -\frac{1}{2}i\Omega[\sigma_x, \rho] - \frac{1}{8}\Omega^2\frac{N_0}{P_0}[\sigma_y, [\sigma_y, \rho]] \quad (3.45)$$

Integrating (3.45), we find that if the ion begins in state $|0\rangle$ then its probability of being in state $|1\rangle$ after a time $t = \frac{n\pi}{\Omega}$ is given by

$$P_1 = \frac{1}{2} \left\{ 1 - e^{-\frac{1}{4}n\pi \frac{N_0\Omega}{P_0}} \left[\cos(n\pi\kappa) - \frac{N_0\Omega}{4P_0\kappa} \sin(n\pi\kappa) \right] \right\} \quad (3.46)$$

where

$$\kappa := \sqrt{1 - \left(\frac{N_0\Omega}{4P_0} \right)^2} \quad (3.47)$$

Dropping terms of greater than first order in $\frac{N_0\Omega}{4P_0}$, we find that the error after n π -pulses is given by

$$1 - P_1 = n\pi \frac{N_0\Omega}{8P_0} \quad (3.48)$$

More generally, one can show that if the system begins in a pure state then, to first order in the noise, the infidelity in the density matrix after a time $t = n\pi$ is

$$1 - F = n\pi \frac{N_0\Omega}{8P_c} (1 + \cos^2 \phi_i \sin^2 \theta_i) \quad (3.49)$$

Integrating over all possible starting states, we find that the average infidelity of a series of π -rotations from random initial states is

$$\langle 1 - F \rangle = n\pi \frac{N_0\Omega}{6P_c} \quad (3.50)$$

Equations (3.49) and (3.50) were derived previously by Chen *et al.* [CBWT12] by considering the rotation of the Bloch vector due to a resonant field with phase noise. We note that the present derivation has the potential advantage that we work within the general framework of noise master equations. This allows extension to more complicated systems which are not readily solvable using the Bloch vector framework, such as non two-level systems.

3.4.4 White frequency noise

We turn now to the case of Rabi flopping due to a field with white frequency fluctuations. Our Hamiltonian is

$$H = \frac{1}{2}\hbar\Omega e^{-i(\delta t + \phi(t))}\sigma_+ + \frac{1}{2}\hbar\Omega e^{i(\delta t + \phi(t))}\sigma_- \quad (3.51)$$

where $\phi(t)$ is the time-dependent phase resulting from the frequency fluctuations around the mean detuning δ .

To linearise this Hamiltonian with respect to the noise term, we move into a frame that co-rotates with the frequency fluctuations by making the stochastic transformation⁴

$$\begin{aligned}\rho &\mapsto \tilde{\rho} := U^\dagger \rho U \\ H &\mapsto \tilde{H} := U^\dagger H U = \frac{1}{2} \hbar \Omega \sigma_+ e^{-i\delta t} + \text{H.c.}\end{aligned}\tag{3.52}$$

where

$$U = e^{-i\phi(t)\sigma_+\sigma_-}\tag{3.53}$$

We then have

$$\begin{aligned}d\tilde{\rho} &= (dU^\dagger)\rho U + U^\dagger \rho (dU) + U^\dagger (d\rho) U \\ &= i\sigma_+\sigma_- U^\dagger \rho U d\phi - iU^\dagger \rho U \sigma_+\sigma_- d\phi - \frac{i}{\hbar} U^\dagger [H, \rho] U dt \\ &= -\frac{i}{\hbar} [\tilde{H}, \tilde{\rho}] dt + i[\sigma_+\sigma_-, \tilde{\rho}] d\phi \\ &= -\frac{i}{\hbar} [\tilde{H}, \tilde{\rho}] dt + i\delta f[\sigma_+\sigma_-, \tilde{\rho}] dt\end{aligned}\tag{3.54}$$

where $\delta f(t) := \dot{\phi}(t)$ represents the frequency fluctuations.

Next, we take the limit where

$$\langle \delta f(t + \tau) \delta f(t) \rangle = \Gamma \delta(\tau)\tag{3.55}$$

where Γ is related to the phase noise PSD by (see C.2.3)

$$\Gamma = \frac{N_0(f_0)}{P_0} f_0^2\tag{3.56}$$

and we have assumed that the phase noise PSD has a $1/f^2$ spectrum, so that

$$N_0(f) = N_0(f_0) \left(\frac{f_0}{f} \right)^2\tag{3.57}$$

for some frequency f_0 .

⁴Notice that this transformation adds random phases to the coherences. Since we will be concerned with populations in this section, we ignore these phases; however, they will be significant in some cases.

Using (3.15) and dropping the tildes, we arrive at the master equation

$$\dot{\rho} = -\frac{1}{2}i\Omega[e^{-i\delta t}\sigma_+ + e^{i\delta t}\sigma_-, \rho] - \frac{1}{2}\Gamma[\sigma_+\sigma_-, [\sigma_+\sigma_-, \rho]] \quad (3.58)$$

Once again, we demonstrate the effect of frequency noise by considering the resonant case, $\delta = 0$. After applying the driving field for a time t to a system beginning in either qubit state, the infidelity is

$$1-F = \frac{1}{2} \left\{ 1 - e^{-\frac{1}{4}\Gamma t} \left[\frac{\Gamma}{4\Omega_e} \cos(\Omega t) \sin(\Omega_e t) + \cos(\Omega t) \cos(\Omega_e t) + \frac{\Omega}{\Omega_e} \sin(\Omega t) \sin(\Omega_e t) \right] \right\} \quad (3.59)$$

where the effective Rabi frequency is defined by

$$\Omega_e := \sqrt{\Omega^2 - \left(\frac{\Gamma}{4}\right)^2} \quad (3.60)$$

Dropping terms of greater than first order in Γ ,

$$1 - F \simeq \frac{1}{8}\Gamma t - \frac{1}{16\Omega}\Gamma \sin(2\Omega t) \quad (3.61)$$

Starting from a random initial state, the infidelity after a time t is

$$\begin{aligned} \langle 1 - F \rangle = \frac{1}{2} \left\{ 1 - \frac{1}{3}e^{-\frac{1}{2}\Gamma t} + \frac{1}{3\Omega_e}e^{-\frac{1}{4}\Gamma t} [(\Omega - \Omega_e) \cos(t(\Omega + \Omega_e)) \right. \\ \left. - (\Omega + \Omega_e) \cos(t(\Omega - \Omega_e))] \right\} \end{aligned} \quad (3.62)$$

To first order in the noise, we find

$$\langle 1 - F \rangle \simeq \frac{1}{6}t\Gamma \quad (3.63)$$

3.5 The Noisy Quantum Harmonic Oscillator

3.5.1 Frequency fluctuations

We begin by considering the problem of a quantum harmonic oscillator of frequency ω which undergoes white fluctuations, such that $\omega(t) = \omega_0 + \delta\omega(t)$. Here, ω_0 represents the oscillator's average frequency and $\delta\omega$ represents the fluctuations. The Hamiltonian for the oscillator's free evolution is

$$H_0 = \hbar\omega_0 \left(a^\dagger a + \frac{1}{2} \right) + \hbar\delta\omega \left(a^\dagger a + \frac{1}{2} \right) \quad (3.64)$$

In the interaction picture with respect to the oscillator's mean energy, the oscillator's dynamics are completely determined by the fluctuating component, so that

$$H_{0,I} = \hbar\delta\omega \left(a^\dagger a + \frac{1}{2} \right) \quad (3.65)$$

We assume $\delta\omega$ to be white with

$$\langle \delta\omega(t + \tau) \delta\omega(t) \rangle = \Gamma \delta(\tau) \quad (3.66)$$

where $\delta(\tau)$ is the Dirac delta function.

The resulting master equation is

$$\begin{aligned} \frac{d\rho}{dt} &= -\frac{1}{2}\Gamma [a^\dagger a, [a^\dagger a, \rho]] \\ &= \frac{1}{2}\Gamma \left[2a^\dagger a \rho a^\dagger a - (a^\dagger a)^2 \rho - \rho (a^\dagger a)^2 \right] \end{aligned} \quad (3.67)$$

We note that this equation is often derived by considering a harmonic oscillator coupled to a quantum heat bath (see e.g. [TTM⁺00]).

To understand the effect of the dephasing term here, we consider the evolution of the density matrix in the basis of motional Fock states

$$\rho(t) = \sum_{n,m} \rho_{nm}(t) |n\rangle \langle m| \quad (3.68)$$

where $|n\rangle$ denotes the oscillator's n^{th} Fock state. The density matrix evolves according to

$$\begin{aligned} \dot{\rho}_{ij}(t) &= \frac{1}{2}\Gamma \sum_{n,m} \rho_{nm} \{ 2\langle i|a^\dagger a|n\rangle \langle m|a^\dagger a|j\rangle \\ &\quad - \langle i|(a^\dagger a)^2|n\rangle \langle m|j\rangle - \langle i|n\rangle \langle m|(a^\dagger a)^2|j\rangle \} \\ &= \frac{1}{2}\Gamma \rho_{ij}(t) \{ 2ij - i^2 - j^2 \} \\ &= -\frac{1}{2}\Gamma (i - j)^2 \rho_{ij}(t) \end{aligned} \quad (3.69)$$

We see that the populations (ρ_{ii}) remain constant, however the coherences decay as

$$\begin{aligned} \rho_{ij}(t) &= e^{-\frac{1}{2}\Gamma (i-j)^2 t} \rho_{ij}(0) \\ &= e^{-(i-j)^2 t/\tau_c} \rho_{ij}(0) \end{aligned} \quad (3.70)$$

where we have defined the motional coherence, $\tau_c := 2/\Gamma$, as the time constant for the coherence decay between Fock states separated by a single vibrational quantum.

If the harmonic oscillator interacts with another system so that

$$H = H_0 + H_1 \quad (3.71)$$

then our resulting master equation is

$$i\hbar \frac{d\rho}{dt} = [\tilde{H}_1, \rho] - \frac{1}{2}i\hbar\Gamma[a^\dagger a, [a^\dagger a, \rho]] \quad (3.72)$$

where $\tilde{H}_1$ represents H_1 after transformation to the interaction picture with respect to the harmonic oscillator's mean free energy. In particular, we note that motional frequency fluctuations during a Mølmer-Sørensen gate may be described using the master equation

$$i\hbar \frac{d\rho}{dt} = [H_{\text{MS}}, \rho] - \frac{1}{2}i\hbar\Gamma[a^\dagger a, [a^\dagger a, \rho]] \quad (3.73)$$

where H_{MS} is the interaction picture Mølmer-Sørensen Hamiltonian of equation (2.39).

If the noise cannot be accurately described as white then it may be convenient to use the perturbative master equation (3.25).

3.5.2 Frequency fluctuations due to RF noise

Due to restrictions on trap geometry we are only able to produce microwave fields which couple to the ions' motion along the trap's radial directions. The frequencies of these modes are determined by the frequency and amplitude of the trapping RF. In this section, we consider how one can relate the RF field's noise properties to the master equations derived in the previous section. We will assume that the noise is slow compared with the RF frequency so that the pseudopotential approximation can be applied. This is a good approximation in most systems, since the RF is coupled onto the trap's electrodes using a high-Q

resonant circuit. Quality factors in the range of 50-500 are typical, setting a large separation of time-scales.

We consider a radial mode with fluctuating secular frequency such that $\omega(t) = \omega_0 + \delta\omega(t)$, where ω_0 represents the mean frequency and $\delta\omega$ represents the fluctuations. We may relate $\omega(t)$ to the trapping RF amplitude and frequency, $V(t) = V_0 + \delta V(t)$ and $\Omega(t) = \Omega_0 + \delta\Omega(t)$ as follows⁵ [WMI⁺98]

$$\omega(t) = \omega_0 \times \frac{V(t)}{V_0} \times \frac{\Omega_0}{\Omega(t)} \quad (3.74)$$

For pure amplitude noise equation (3.74) becomes

$$\delta\omega(t) = \omega_0 \delta V(t) / V_0 \quad (3.75)$$

Now, for white AM noise with a power spectral density (PSD) of N_0^{AM} , we have

$$\begin{aligned} \langle \delta\omega(t + \tau) \delta\omega(t) \rangle &= \omega_0^2 \left\langle \frac{\delta V(t + \tau)}{V_0} \frac{\delta V(t)}{V_0} \right\rangle \\ &= \omega_0^2 \frac{N_0^{\text{AM}}}{P_c} \delta(f) \end{aligned} \quad (3.76)$$

where P_c is the power in the carrier (see appendix C for a derivation of this result). Thus, the dephasing factor Γ is related to the RF AM noise spectrum by

$$\Gamma = \omega_0^2 \frac{N_0^{\text{AM}}}{P_c} \quad (3.77)$$

Similarly, for frequency noise we have

$$\begin{aligned} \delta\omega(t) &= \omega_0 \left(\frac{1}{1 + \delta\Omega(t)/\Omega_0} - 1 \right) \\ &\simeq -\omega_0 \delta\Omega(t) / \Omega_0 \end{aligned} \quad (3.78)$$

So, for a $1/f^2$ phase noise spectrum (white FM noise), given by

$$N_0(f) = N_0(f_n) \left(\frac{f_n}{f} \right)^2 \quad (3.79)$$

for some frequency, f_n , we have

$$\begin{aligned} \langle \delta\omega(t + \tau) \delta\omega(t) \rangle &= \left(\frac{\omega_0}{\Omega_0} \right)^2 \langle \delta\Omega(t + \tau) \delta\Omega(t) \rangle \\ &= \omega_0^2 \left(\frac{2\pi f_n}{\Omega_0} \right)^2 \frac{N_0(f_n)}{P_c} \delta(f) \end{aligned} \quad (3.80)$$

⁵Notice that here we ignore anti-trapping due to DC potentials.

So

$$\Gamma = \omega_0^2 \left(\frac{2\pi f_n}{\Omega_0} \right)^2 \frac{N_0(f_n)}{P_c} \quad (3.81)$$

In most frequency sources, the AM and PM noise PSDs will be similar. However, comparing (3.81) to (3.76), we see that the contribution from PM noise at a frequency f is suppressed by a factor of $\left(\frac{2\pi f}{\Omega_0} \right)^2$ compared with AM noise. For a typical experiment, $\Omega_0/(2\pi)$ is in the range 10 MHz - 100 MHz, whilst the system response time is limited to < 100 kHz. Thus, $\left(\frac{f}{\Omega_0} \right)^2$ represents a large suppression factor. We thus conclude that frequency noise is unlikely to be a problem in any experiment.

In practice, then, we find that

$$\Gamma = \omega_0^2 \frac{N_0^{\text{AM}}}{P_c} \quad (3.82)$$

giving a motional coherence time of

$$\tau_c = \left(\frac{1}{2} \omega_0^2 \frac{N_0^{\text{AM}}}{P_c} \right)^{-1} \quad (3.83)$$

3.5.3 Coupling to a noisy electric field

Finally we turn to the case of a quantum harmonic oscillator driven by a fluctuating classical force. Our Hamiltonian is [CN65]

$$H = \hbar\omega \left(a^\dagger a + \frac{1}{2} \right) - q_0 F(t) (a + a^\dagger) \quad (3.84)$$

where ω is the oscillator's frequency, $q_0 := \sqrt{\hbar/2m\omega}$ is the RMS extent of the oscillator's ground-state wavefunction and $F(t)$ represents a real, zero-mean fluctuating classical force, characterised by

$$\langle F(t + \tau) F(t) \rangle = \Gamma \delta(\tau) \quad (3.85)$$

In the interaction picture with respect to the oscillator's free energy, the Hamiltonian is

$$H = -q_0 F(t) (a e^{-i\omega t} + a^\dagger e^{i\omega t}) \quad (3.86)$$

The corresponding master equation is

$$\begin{aligned}\dot{\rho} &= -\frac{1}{4\hbar m\omega}\Gamma[ae^{-i\omega t} + a^\dagger e^{i\omega t}, [ae^{-i\omega t} + a^\dagger e^{i\omega t}, \rho]] \\ &= -\frac{1}{4\hbar m\omega}\Gamma\left[e^{-2i\omega t}(a^2\rho - 2a\rho a + \rho a^2) + e^{2i\omega t}\left((a^\dagger)^2\rho - 2a^\dagger\rho a^\dagger + \rho(a^\dagger)^2\right) + \right. \\ &\quad \left. aa^\dagger\rho + a^\dagger a\rho - 2a^\dagger\rho a - 2a\rho a^\dagger + \rho a^\dagger a + \rho aa^\dagger\right]\end{aligned}\quad (3.87)$$

Dropping counter-rotating terms we find

$$\dot{\rho} = \frac{1}{4\hbar m\omega}\Gamma[2a^\dagger\rho a + 2a\rho a^\dagger - aa^\dagger\rho - a^\dagger a\rho - \rho a^\dagger a - \rho aa^\dagger]\quad (3.88)$$

We will generally be concerned with heating due to fluctuating electric fields coupling to the ion's charge. In this case

$$F(t) = qE(t)\quad (3.89)$$

where $E(t)$ represents a fluctuating spatially uniform field. We define the PSD of field fluctuations by⁶

$$\langle E(t+\tau)E(t) \rangle = \frac{1}{2}S_E\delta(\tau)\quad (3.90)$$

in which case, our master equation becomes

$$\dot{\rho} = \frac{q^2 S_E}{8\hbar m\omega} [2a^\dagger\rho a + 2a\rho a^\dagger - aa^\dagger\rho - a^\dagger a\rho - \rho a^\dagger a - \rho aa^\dagger]\quad (3.91)$$

It is easy to show (see e.g. [WMI⁺98]) that if the oscillator begins in a thermal state characterised by a mean occupancy, $\bar{n}$, then the stochastic force results in heating of the oscillator's motion such that

$$\dot{\bar{n}} = \frac{q^2 S_E}{4\hbar m\omega}\quad (3.92)$$

In practice, the PSD for electric field fluctuations is often non-white. However, since the ions' motion is only lightly damped, we expect it to have a narrow

⁶Notice the factor of $\frac{1}{2}$ in this definition, arising from the fact that S_E represents base-band noise and hence is defined using a one-sided spectral power density (i.e. we only integrate over positive frequencies). This definition is standard in the literature (see, e.g. [TKK⁺00])

lineshape. So long as variations in the noise spectrum are small over the motional linewidth, one may still use the above analysis even with non-white noise. In this case, we replace $S_E \mapsto S_E(\omega)$ (where ω is the oscillator's frequency). If the ions experience micromotion then electric fields resonant with sidebands of the RF may also couple to the ions' motion. See [TKK⁺00] for a discussion of this.

4

Apparatus

The work presented in this thesis builds on the previous efforts of ion trappers in the “New Lab” at Oxford. As such, much of the apparatus it relies on has already been documented elsewhere. While we will provide a brief overview of the pre-existing equipment to assist the reader’s understanding of the present work, we will avoid detailed repetition of previous discussions. Further details can be found in [All11] and the references therein.

4.1 The “microwave trap”

The experiments described in this thesis were conducted using the ion trap shown in figure 4.1. It is a symmetric, 5-wire surface-electrode design featuring integrated microwave circuitry and a $75\text{ }\mu\text{m}$ ion-surface separation, whose design and initial characterisation were discussed in [All11] and [AHB⁺13].

4.1.1 Design

The trap design was optimised for use with the microwave-driven two-qubit gate scheme described in section 2.7. To create the necessary microwave field gradients while maintaining a field null at the trap’s centre, we require three microwave electrodes. Since both ions must lie in the field null at the same




Figure 4.1: The ion trap used during this work. (a) Photograph showing the ion trap mounted on its aluminium holder. (b) Model of the trap, showing integrated microwave elements (CPW stands for coplanar waveguide). (c) Schematic of the trap centre, showing voltages and currents applied to each electrode. The direction of the 146 G static B-field is indicated by the arrow $\mathbf{B}_0$. Reprinted with permission from [AHB⁺13]. Copyright 2013, AIP Publishing LLC.

time, these electrodes must run parallel to the trap's symmetry axis, alongside electrodes carrying the trapping RF and DC control voltages.

We conducted numerical simulations using Ansys HFSS software to study several potential trap geometries. These simulations demonstrated that the largest magnetic field gradients are achieved when the trapping RF, DC control voltages and microwaves are combined onto the same three central electrodes. This combination of different signals onto the same electrodes places several constraints on the final design: the electrodes must be DC coupled and not short-circuited at any point; they should act as a $50\ \Omega$ load around the 3.2 GHz qubit frequency; and, since the qubit is a magnetic dipole transition, the ions should

be trapped above a current anti-node if possible to maximise the field gradient they experience.

We opted for the simplest design satisfying these constraints, utilising half-wave cavities resonant at 3.2 GHz and matched to $50\ \Omega$ using quarter-wave coupling elements. The cavities and matching elements were designed to produce a modest Q of ~ 5 to provide a reasonably broad-band $50\ \Omega$ match. As well as allowing the circuit to be matched to $50\ \Omega$, the cavities provide a current build-up, increasing the magnetic field that may be produced by a given input power. The coupling elements are connected to in-vacuum SMA connectors using $50\ \Omega$ co-planar waveguide.

Outside the three central electrodes are six further DC control electrodes used to create the static trap potential. They are set back from the central electrodes using ground plane to ensure that they do not interfere with the microwave return currents which would lead to impedance discontinuities.

Confinement in the trap’s radial directions is achieved using 72 V of RF at 38.7 MHz to create a 4 MHz potential with a stability parameter $q = 0.3$ and an estimated trap depth of 59 meV. DC voltages of $\lesssim 10$ V provide axial confinement of up to 500 kHz; create a static radial quadrupole used to rotate the trap’s radial normal modes, allowing efficient cooling [ASS⁺10]; and compensate stray electric fields to minimise micromotion.

4.1.2 Fabrication

The trap was fabricated in-house by David Allcock and the author by electroplating gold electrodes onto a sapphire substrate. The fabrication process has been described in detail in [All11]. Rather than repeat the process in full detail here, we simply note a few key features.

We chose to use sapphire instead of the fused silica substrates used for previous traps because of its superior thermal conductivity and electrical proper-

ties. The trap substrate was bonded to an aluminium holder using a thin layer of Epoxy Technology 353ND epoxy. Aluminium was chosen because its high thermal conductivity allows it to act as an efficient heat-sink, minimising temperature changes in the trap due to dissipated microwave power. DC connections were made by wire-bonding the trap to copper wires¹ glued to the trap holder. To ensure that the DC electrodes are connected to RF ground by a low-impedance connection, each DC connection was also wire-bonded to an 820 pF single layer capacitor² attached to the aluminium trap mount using conductive epoxy³. Microwave connections were made by modifying 2.92 mm connectors⁴, connected to the trap using launch pins and conductive epoxy.

4.1.3 Characterisation

We characterised the trap's microwave properties using a vector network analyser (VNA). The data, shown in figure 4.2, demonstrates that the cavities' resonance frequencies were $\sim 5\%$ higher than their design frequency. This is thought to have been due to insufficient refining of the mesh used during simulations and not due to fabrication imperfections. However, because of the cavities' modest Q , more than 65 % of the input power is nonetheless coupled in at 3.22 GHz. After the measured resonance frequencies were taken into account, simulations were found to agree well with preliminary measurements of the microwave field made using a single trapped ion [AHB⁺13].

By applying microwaves to the three central electrodes with appropriate amplitudes and phases, one may produce a strong magnetic field gradient with only a small residual field at the trap's centre, as required for the microwave-

¹Kurt Lesker FTAK06010, Kapton dipped copper wire with 0.69 mm outer diameter and 0.63 mm conductor diameter.

²American Technical Ceramics 116UL821M100TT.

³Epoxy Technology H20E silver loaded epoxy.

⁴Original connectors were Southwest Microwave 1012-16SF with 290-13P launch pins. Connectors were modified by removing the capture bead and dust cover. The capture bead was replaced with a new bead made from PEEK. See [All11] for details.




Figure 4.2: Comparison of VNA data (solid line) and simulations (dashed lines) of the total fraction of power incident on a given microwave port that is coupled into the trap. Figure taken from [All11].

driven gates. Using 1.71 mW, 0.48 mW and 1.83 mW on ports 1,2 and 3 respectively (see figure 4.1), we were able to produce a quadrupole field, characterised by gradients of $\frac{\partial B_z}{\partial y} = 0.304$ T/m and $\frac{\partial B_y}{\partial z} = 0.293$ T/m. This agrees well with simulations, which predict $\frac{\partial B_z}{\partial y} = \frac{\partial B_y}{\partial z} = 0.289$ T/m when experimentally measured values for the power coupled into each port are taken into account. Based on thermal measurements in similar neutral atom chip traps, we expect to be able to use input powers exceeding 1 W without damaging the trap [GKW⁺04]. This would be sufficient to produce gradients approaching 10 T/m, allowing two-qubit gate speeds of order 1 ms.

Using a modified Doppler recool method [ASS⁺10] [WEL⁺07], we measured the single ion heating rate to be $\dot{n} = 1.4(3)$ quanta/ms at an axial secular frequency of $\omega = 2\pi \times 500$ kHz, corresponding to an electric field spectral noise density of $\omega S_E = 1.6(3) \times 10^{-5}$ V²/m². This is extremely low for a room-temperature surface trap and comparable to many cryogenic traps (see, e.g. [HCW⁺12] for a recent review of heating rates in ion traps). Since the mechanisms leading to anomalous heating are still not properly understood, it is hard to know why the heating rate is so low. We note, however, that attention was

paid to screening and filtering of DC control electronics to minimise excess technical noise.

We characterise and compensate the ion's micromotion using a single $^{40}\text{Ca}^+$ ion with the technique described in [ASS⁺10]. Using 38.7 MHz trap RF and a secular frequency of 3 MHz, we find that we are able to compensate micromotion in the in-plane ($\hat{z}$) direction to $\lesssim 1 \text{ V/m}$, corresponding to a micromotion amplitude of $\lesssim 1 \text{ nm}$. In the out-of-plane ($\hat{y}$) direction, however, we are unable to compensate the micromotion to better than $\lesssim 30 \text{ V/m}$, corresponding to a micromotion amplitude of $\lesssim 20 \text{ nm}$. The most likely cause of this uncompensated micromotion is RF pick-up on the trap's centre electrode [AHJ⁺11]. This electrode is only connected to RF ground outside the vacuum can, after the microwave block filter in the Diplexer (see section 4.3.1), presenting a significant inductance. The inability to ground this electrode at RF frequencies inside the vacuum chamber is a side-effect of multiplexing the microwaves onto the DC electrode. However, the micromotion amplitude is sufficiently small that it is not anticipated to cause problems.

4.2 Vacuum system

The trap is housed in a stainless steel UHV system constructed around a 6'' spherical octagon⁵ (see figure 4.3). The system is pumped using an ion pump⁶ and a non-evaporable getter⁷. An ion gauge⁸ is used to monitor the pressure, which is $\lesssim 1 \times 10^{-11} \text{ Torr}$, even with the calcium oven running at our typical loading current of 5 A (see below).

Optical access for laser beams is provided by six 2.75'' viewports⁹ attached

⁵Kimball Physics MCF600-SphOct-F2C8, 316L stainless steel with 2x 6'' flanges and 8x 2.75'' CF flanges.

⁶Varian VacIon Plus 20 Diode, 271/s.

⁷Saes Getters GP100 Mk5, ST707 alloy, 240 l/s H_2 pumping speed.

⁸Varian UHV-24P with x-ray limit of $5 \times 10^{-12} \text{ Torr}$.

⁹Torr Scientific VPZ38QBBAR-LN, 38CF 3 mm thick fused silica viewport using 316LN stainless steel with Kovar weld rings. Broad-band anti-reflection coating (BBAR) on both sides, speci-




Figure 4.3: Schematic of the vacuum system, showing the static magnetic field direction, B_0 , and laser beams. The red beam path shown dashed enters and exits the vacuum system at 45° to the plane of the diagram through the imaging window.

to the sides of the octagon. A further 6" viewport¹⁰ attached to the top of the octagon is used for imaging as well as 45° IR beam paths (see figure 4.3).

On the bottom of the octagon is a 6" flange custom made by LewVac, containing a 25-pin D-type feedthrough used for the DC connections, a pair of 2-pin 15 A power feedthroughs intended for ovens¹¹ and tapped holes to mount the trap base.

In-vacuum SMA cables¹² are used to connect the trap mount to an SMA

fied reflectance per surface is 0.5 % at 397 nm and 9.5 % at 866 nm.

¹⁰Torr Scientific VPZ100-ITO/AR-LN 4 mm thick Kodial glass using 316LN stainless steel with Kovar weld rings. 1QWOT AR coating on air-side chosen for transmission at 397 nm, thin layer of ITO on vacuum side to prevent charging. We measured the transmission of this window to be 76 % at 397 nm at normal incidence.

¹¹While we only use a single calcium oven in this experiment, we wished to leave room on the feedthrough for a second oven. This could be used, for example, to allow a second ion species to be loaded for sympathetic cooling.

¹²LewVac ZCC-50-SMA/SMA-105, 105 mm long SMA cables with gold-plated brass plugs,

feedthrough¹³. The feedthrough is offset from the main vacuum chamber using a nipple¹⁴, allowing the SMA cables to be slightly longer than would otherwise be possible, providing some room for cables to flex, reducing strains.

The calcium oven consists of a ~ 10 mm length of metal tube¹⁵, containing $\simeq 5$ 1 mm x 1 mm x 0.65 mm granules of enriched $^{43}\text{Ca}^+$, composed of 12 % $^{43}\text{Ca}^+$ and 88 % $^{40}\text{Ca}^+$ supplied by Oakridge National Laboratory [All11]. The tube has a single 0.5 mm hole in its centre to allow calcium vapour to effuse. It is crimped at both ends and spot-welded to sections of 2 mm diameter stainless steel rod, which are connected to the pins of the feedthrough using barrel connectors. During loading we resistively heat the oven using 5 A, which requires 1 V including the voltage drop across the cabling between the power supply and the oven. Once an ion has been trapped we turn the oven current down to ~ 4.3 A for increased ion lifetime. The oven feedthroughs are positioned so that the flux of neutral calcium atoms makes an angle of $\simeq 83^\circ$ to the photoionization beams, reducing Doppler shifts on the 423 nm transition.

The oven heating current generates a magnetic field at the trap's centre of $\simeq 5$ mG/A measured using the frequency shift of the $|4, +4\rangle \Leftrightarrow |3 + 3\rangle$ stretched transition.




Figure 4.4: Schematic (top) and photograph (bottom) of diplexer and RF transformer with metal lids removed. The third diplexer channel for the centre trap electrode (port 2) is mounted in a separate box below the two channels shown. It is identical except that the RF connection is grounded. Figure taken from [All11].

4.3 Electrode Connections

4.3.1 Diplexer and transformer

Combination of the DC control voltages, RF trapping voltages and microwave currents is achieved outside the vacuum system using the diplexer shown in figure 4.4. The diplexer is designed to isolate the various sources from each other, while providing low insertion losses and minimizing the capacitance on the RF part of the circuit. It consists of a coupled-line bandpass filter (A) to pass microwaves, but block RF and DC, and a stub filter (B) to pass RF and DC but block microwaves. Further inductors¹⁶ and capacitors¹⁷ are used to form a bias-Tee to allow the combination of RF and DC.

We designed the diplexer using Microwave Office software¹⁸. It is fabricated on ceramic-PTFE composite board with low loss and high temperature stability¹⁹. The back side of the diplexer acts as a ground plane and is bonded to its brass housing using conductive epoxy. Network analyser measurements on the diplexer show reasonable agreement with simulations (figure 4.5), however the return loss is never better than 10 dB. We believe that this is due to the SMA connectors we used and hope to improve it in a future revision.

At 40 MHz the diplexer and trap act as a lumped capacitive element. We use a resonant transformer to match this to our 50 Ω RF sources. The transformer

gold-plated copper pins, PEEK insulator, assembled using UHV-compatible solder. Cables initially have a Teflon insulator, which was replaced with PEEK by the manufacturer. This changes the impedance to roughly 45 Ω . We measured the insertion loss to be <0.41 dB and the return loss to be >22 dB at 3.2 GHz. Cables have a capacitance of 95 pF/m.

¹³Ceramtech 21363-01-CF, 2.75" CF flange containing 3 double-ended SMA feedthroughs with grounded shields. Air dielectric and glass-ceramic seals allow baking to 350 °C. We measured the insertion loss to be <0.12 dB and the return loss to be >26 dB at 3.2 GHz.

¹⁴LewVac ZFL-TFA-40LN 316LN steel tapped flange adaptor, 58.5 mm long.

¹⁵316L stainless steel tube supplied by Unimed S.A., 2 mm outer diameter, 1.8 mm inner diameter.

¹⁶Epcos B82111EC25 ferrite core inductor, 100 μ H, 500 V DC/AC, 55 MHz self-resonance frequency.

¹⁷Kemet 2225 10 nF COG 500 V, ± 30 ppm/°C.

¹⁸This was generously donated by the AWR corporation.

¹⁹Rodgers Corporation RO3003, 0.5 mm thick, 17 μ m copper clad, $\epsilon_r = 3.00$ (temperature coefficient +13 ppm/°C), dissipation factor 0.0013.




Figure 4.5: Comparison between microwave office (MWO) simulations and measured performance of the diplexer. VNA data is shown for one of the diplexer channels which is connected to the RF transformer as well as for the centre channel, which is not connected to the transformer. S_{ij} denote the 2-port scattering parameters (see, e.g. [Poz05] for the definition of this quantity). Figure taken from [All11].

is wound on an iron-powder toroidal core²⁰, providing a circuit Q of 65 and a voltage step-up²¹ of 9.6. This is sufficient to allow 4 MHz radial frequencies to be achieved using $\lesssim 1$ W of RF power.

The coupled-line bandpass filter only provides ~ 25 dB attenuation at 40 MHz, which allows ~ 1 mW of RF to be coupled out of the microwave ports. To increase the isolation, we use further high-pass filters²² on the microwave ports.

²⁰Micrometals T80-10

²¹The voltage step-up was calculated by comparing measurements of the ion's secular frequencies with an electrostatic model of the trap using a known input RF power.

²²MiniCircuits VHF-1320+ 1700 MHz to 5000 MHz high-pass filter, 63.2 dB insertion loss at 50 MHz, 0.25 dB insertion loss at 3.25 GHz.

4.3.2 DC supply

We create our DC voltages using a commercial 16-bit 32-channel digital to analogue converter (DAC) evaluation board²³ connected to the experimental control PC via USB. The DAC is wired to the 6 DC electrodes using screened cable with the screen grounded at the DAC end. Connections to the central electrodes are made via SMA cables connected to the diplexer. All channels are filtered using a low-pass filter with a C-R-C-L-C topology described in [AHJ⁺11]. For the 6 DC electrodes, this filter board is mounted in a screened box bolted directly onto the D-type feedthrough on the vacuum system.

4.3.3 Trapping RF

Confinement in the trap's radial directions is achieved by applying an RF voltage to the central electrodes. The secular frequency associated with the resulting pseudo-potential is proportional to the RF voltage's amplitude and inversely proportional to its frequency.

Manipulations of a single ion are insensitive to small fluctuations in the ion's secular frequencies and so place no particular requirements on the trapping RF. The microwave-driven two-qubit gates, on the other hand, rely on driving motional sideband transitions using the radial modes²⁴ and so are directly sensitive to such fluctuations. This makes ensuring a high degree of stability in the radial secular frequency a primary experimental concern for these gates. In practice, it is easy to ensure a low level of drift and noise in the RF's frequency (c. f. section 3.5), so our main concern will be fluctuations in its amplitude.

For our single-qubit work, we generated our RF using an HP 8656B frequency synthesiser and a 1 W amplifier²⁵. Using this set-up, we found that the ion's radial secular frequency fluctuated by ~ 1 kHz with the temperature

²³Analog devices AD5372 DAC chip on an EVAL-ADUM4160 evaluation board.

²⁴As discussed in section 4.1, the use of radial rather than axial modes is forced on us by the requirement that both ions should lie in the microwave field null.

²⁵Motorola CA2810C.

of our lab due to changes in the RF power (see [AHJ⁺11]). This is comparable to the maximum sideband Rabi frequency that this trap was designed to produce and so clearly is not acceptable for two-qubit work. To improve this, we have considered several alternate RF sources, whose design and characterisation are described in section 6.7.

4.4 Coils

4.4.1 146 G coils

The 146 G magnetic field is provided by a pair of rectangular water-cooled coils, each consisting of 16 turns of hollow copper wire²⁶. The coils are set onto aluminium formers using epoxy resin²⁷. They have inner dimensions²⁸ of 137 mm x 226 mm ($\hat{y} \times \hat{x}$) and the perpendicular distance from the closest face of each coil to the ion is 52 mm. From the coil dimensions and the Biot-Savart law, we predict a magnetic field of 1.2 G/A. Experimentally, we find that we require 135 A to produce the 146 G field in reasonable agreement with predictions.

The coils are operated in series electrically and in parallel for water flow. They are powered using a pair of high-current power supplies²⁹ in series, one operating in constant current mode and one in constant voltage. To produce 146 G we require 135 A and 10 V, including voltage drop across welding cables used to connect the coils to the supplies.

We cool the coils using ambient-temperature water at mains pressure with a combined flow of 600 ml/minute. We designed the coils so that the length

²⁶Wolverine Tube Europe OFHC copper 3.2 mm x 3.2 mm square cross section with a 2 mm diameter bore, Kapton wrapped by S & W Wire Company.

²⁷Araldite 2011 epoxy, glass transition at 45 °C, max operating temperature 80 °C.

²⁸All dimensions given in this section refer to the size of the coil itself, and not to the former it is wound on.

²⁹Agilent 6671A 2 kW, 8 V, 220 A power supply. Specified RMS ripple and noise current is 200 mA between 20 Hz to 20 MHz, however the actual noise current depends on the load and is typically significantly smaller for inductive loads such as coils.

of wire in each coil is not much more than 10 m to avoid the flow becoming turbulent. When running at 135 A current, they reach a temperature of $\sim 32^\circ\text{C}$.

Since the coils dissipate over a kilowatt of power, we felt it important to interlock them, which we do in two ways. Firstly, we use the active-low disable logic input on the high-current power supplies. We connect this to a flow sensor³⁰, which monitors the water exiting the coils, and to 50°C thermal switches glued to each coil. Secondly, we use a microprocessor to monitor the state of the flow sensor, thermal switches and a pair of additional digital temperature sensors³¹ attached to the coils. The microprocessor reports to the experimental control computer at regular intervals over USB. If abnormal conditions are detected then the computer shuts the power supplies down using their GPIB interface.

We actively stabilise the current flowing through the coils using a Hall probe current sensor³². The sensor is connected to a channel on the FPGA used for our laser intensity stabilisation [All11], which feeds back to a Darlington transistor used to shunt current away from the coils. We have not characterised this system in detail, however preliminary measurements on the ion suggest that it reduces the RMS noise to $\sim 1\text{ mA}$, which was deemed sufficient for the present experiments.

4.4.2 Trim coils

Since the high-current supplies do not provide mA resolution, we use an additional 92 turn cylindrical trim coil to provide fine adjustment of the field. This coil is placed next to the 397 nm σ^+ viewport and produces a field of 0.46 mG/mA.

For high-quality optical pumping, it is important that the σ^+ laser beams are

³⁰Gems Sensors 155425BSPP

³¹DS18B20

³²LEM LA 205-S 200 A current sensor.

accurately aligned onto the magnetic field axis. Whilst this could be achieved by steering the beam onto the axis provided by the combination of the main coils and background lab field, in practice it is much more convenient to be able to align the field onto the optical pumping beams. To this end, we use further coils to provide fields of a few Gauss along the $\hat{x}$ and $\hat{y}$ axes.

The $\hat{y}$ field is provided by a pair of 30 turn cylindrical coils bolted onto the top and bottom of the octagon. These coils were designed to fit over the 6" flanges and have inner and outer diameters of 156.5 mm and 176 mm respectively. They are 18 mm thick and their closest face is 38.3 mm from the ion. They produce a combined field of approximately 3 G/A.

The vacuum pumps make it difficult to fit a coil over the far end of the vacuum system (see figure 4.3), so the $\hat{x}$ field is produced using a single 168 turn cylindrical coil, which fits over the microwave feedthrough. It has inner and outer diameters of 80 mm and 136 mm respectively and a thickness of 18 mm. Its closest surface is 133 mm from the ion and it produces approximately 0.9 G/A.

All trim coils are powered by Thurlby Thandar Instruments (TTi) QL series power supplies connected to the experimental control PC via USB.

4.5 Lasers and optics

4.5.1 Lasers

We use lasers to photoionize neutral Ca and to cool, prepare and readout Ca^+ ions (sections 2.4-2.6). Photoionization is accomplished using lasers at 423 nm and 389 nm. Ca^+ ions are manipulated using lasers at 393 nm, 397 nm, 866 nm, 850 nm & 854 nm to drive all five electric dipole transitions between the 4S, 4P & 3D levels (figure 2.1).

All of our lasers are commercially available Toptica DL 100 model extended cavity diode lasers in the Littrow configuration apart from the 389 nm which

is a free-running UV laser diode in a home-made mount. They are located on a separate optics table to the main experiment and the beams are transferred using single-mode fibres.

The Ca^+ lasers are frequency stabilised by locking to low-drift etalons³³. The 850 nm and 854 nm lasers are locked using the side-of-fringe technique with Toptica PID controllers, while the other lasers are locked using PDH locks and home-made electronics [Kei07]. Piezo elements incorporated in the etalons allow the lasers' frequencies to be adjusted by the experimental control computer.

To ensure that lasers operate on a single longitudinal mode at the correct frequency, we monitor them using a multi-channel diagnostic system containing a wavemeter and optical spectrum analysers (OSAs) [All11]. The 423 nm laser is not locked to a cavity, but is monitored on the wavemeter and manually tuned to resonance.

The Ca^+ beams used in experiments are switched on/off and amplitude modulated using double passed acousto-optic modulators (AOMs) located before the fibres, while the PI beams are switched using a mechanical shutter. For long experiments requiring an exceptionally high level of laser extinction, we also switch the 397 nm and 393 nm beams using mechanical shutters before the fibres.

We use blazed diffraction gratings after the 397 nm, 850 nm & 866 nm lasers to prevent amplified spontaneous emission (ASE) at other wavelengths from reaching the ion, which could lead to unwanted shelving/deshelving.

To cool $^{43}\text{Ca}^+$ we add 2.93 GHz sidebands to the 397 nm laser using an electro-optic modulator (EOM)³⁴.




Figure 4.6: Doppler and photoionization beam path.

4.5.2 Trap optics

Figure 4.3 provides an overview of the laser beam paths used in these experiments. We use a beam path at 45° to the magnetic field containing both UV & IR beams for cooling, photoionization and in-plane micromotion compensation. For our standard cooling and readout experiments, both the 397 nm and 866 nm beams propagate along this path. However, for low-temperature cooling using a dark resonance (section 2.5) we used a counter-propagating 866 nm beam entering from the opposite end of the vacuum system. We perform optical pumping and shelving using a pair of counter-propagating σ^+ -polarized UV/IR

³³Etalons supplied by National Physical Laboratories with a specified drift rate of less than 0.5 MHz/hour in laboratory conditions.

³⁴New Focus 4431M resonant EOM, tuneable from ~ 2.4 GHz to ~ 3.7 GHz.

beams parallel to the magnetic field. An IR beam path at 45° to the trap plane provides an 850 nm π -polarized shelving beam and an 866 nm beam for vertical micromotion compensation.

The optical layouts of these beam paths are shown in figures 4.6-4.9. The beam paths share several common features. Polarizing beam splitters (PBSs) after the fibres are used to ensure that the beams have well defined polarization states. Pick-offs and photodiodes are used as part of an FPGA-based “noise eater”, which feeds back to the AOM RF level to stabilise the beam power [All11]. To minimise scatter from the UV beams, which can result in charging of dielectric trap structures and degrades the background signal during read-out, we avoid the use of glass components in the UV beam paths, instead using fused silica components where possible.

All beams are focused onto the trap centre using 200 mm focal length achromatic doublet lenses. UV fibre collimators are chosen to produce a roughly $50\text{ }\mu\text{m}$ beam waist ($1/e^2$ beam diameter) to ensure the beam does not clip any part of the trap. Clipping is less important for the IR beams, for which we typically use waists of $\sim 80\text{ }\mu\text{m}$.

We use a beam profiler³⁵ to help align focusing optics and beam positions. The beam profiler is carefully positioned to be the same distance from the final lens as the ion (allowing for the extra optical thickness of the vacuum viewport). During alignment, beams may be diverted onto the beam profiler using aluminium mirrors in flip mounts. See [All11] for more details about trap alignment techniques.

4.5.2.1 Doppler paths

The layout of the Doppler beam path used for cooling, photoionization (PI) and in-plane micromotion compensation is shown in figure 4.6. It is aligned at 45°

³⁵Thorlabs BP104-UV, $0.5\text{ }\mu\text{m}$ resolution.

to the trap's symmetry axis to ensure that all trap normal modes are cooled.

The 397 nm fibre is a Schäfter+ Kirchoff GmbH PMC-S405Si³⁶. To minimise scatter from the 397 nm beam, we use a fibre collimator consisting of a multi-element $f = 14.5$ mm lens³⁷. Unlike other collimators we tried, this produces an output beam with no visible defects when viewed on a card. We are less sensitive to scatter from the PI beams and we collimate them using a standard asphere collimator³⁸. To match the PI spot size to the 397 nm beam, we set the PI collimating lens to produce a diverging beam, which we collimate using a 300 mm achromat approximately 250 mm from the fibre. The IR beam is collimated using a 30 mm achromatic doublet. This is chosen to be approximately twice the focal length of the 397 nm collimator to ensure a suitable spot size at the trap's centre.

The 397 nm and PI beams are combined with orthogonal polarizations using a polarizing beam splitter (PBS). The IR beams are then superposed onto the UV beams by passing them through the back of a UV dielectric mirror. The final mirror before the trap is a Thorlabs UFM10R dual band 400 nm and 800 nm high reflecting mirror suitable for both UV and IR beams.

The 397 nm polarization ratio is set to 2/3 of σ^+ to 1/3 of π using $\lambda/2$ & $\lambda/4$ waveplates (section A.4). The polarization ratio of the IR beams is set to contain equal σ^+ and σ^- with no π -component using IR $\lambda/2$ & $\lambda/4$ waveplates placed in the IR beam before it is combined with the UV. These waveplates are set to compensate for the effect of the UV waveplates which the IR beam passes through.

The counter-propagating 866 nm beam path used for dark-resonance cooling is illustrated in figure 4.7. This beam path uses similar optics to the co-

³⁶PMC-S405Si-2.9-NA011-3-APC-1000-P 405 nm single-mode polarization-maintaining fibre with a pure silica core and a NA of 0.11.

³⁷CVI Melles Griot GLC-14.5-8.0-405, composed of a diffraction-limited meniscus and achromatic doublet combination with a NA of 0.28.

³⁸Thorlabs PAF-2-X-A, $f = 2$ mm




Figure 4.7: Counter-propagating 866 nm path used for dark-resonance cooling.

propagating 866 nm beam path.

4.5.2.2 Sigma paths

We use a pair of counter-propagating σ^+ -polarized beams parallel to the magnetic field for optical pumping. The blue sigma beam path is shown in figure 4.8, with the red beam path set up in the same way. Both σ^+ paths use Glan Taylor polarisers³⁹ to ensure a high-level of polarization purity. The $\lambda/4$ waveplates are mounted on kinematic mirror mounts, allowing them to be tilted and rotated to fine-tune the degree of retardation they produce (section A.1). This is important since they were not designed for the exact wavelengths we are using and thus do not act as perfect $\lambda/4$ waveplates unless tilted appropriately. Furthermore, strains in the viewports induce a birefringence, which may be partially compensated for by tilting the waveplates a small amount (see section 5.1 for further discussion of optimization of waveplate settings).

³⁹Thorlabs TG-10A/GT-10B, 100,000:1 nominal extinction.




Figure 4.8: Blue sigma beam path. “GT” denotes a Glan-Taylor prism.

4.5.2.3 850 nm- π and 866 nm vertical micromotion-motion

The vertical IR beam path is shown in figure 4.9. Apart from the fact that this beam path enters the trap through a periscope it follows a similar pattern to other beam paths.

The vertical micromotion beam is never used at the same time as the main IR Doppler beam. This allows us to split it off from the Doppler beam using a half-wave plate and PBS (figure 4.6). When it is not in use, we extinguish it using a beam block.




Figure 4.9: Vertical infra-red beam path for 850 nm π -polarized readout beam and 866 nm in-plane micromotion compensation.

4.6 Imaging system




Figure 4.10: Schematic of the imaging system used in this work. It is a two-stage design, with the first stage composed of the Innsbruck (IBK) lens and a further 250 mm achromatic doublet. An adjustable aperture in the plane of the image formed by this stage allows background scatter to be reduced.

The imaging system provides us with a means of observing ions by collecting light they scatter and focusing it onto a camera or photomultiplier tube (PMT). As well as serving many important diagnostic purposes, it allows us to make

projective measurements on our qubits by taking advantage of state-dependent fluorescence techniques (section 2.6.3).

The main objective is a custom 5-element lens designed for the Innsbruck ion trapping group (IBK) by Sill Optics. Its design was optimised for focusing 397 nm and 729 nm light to a near diffraction-limited spot through 6 mm of fused silica. It has a specified transmission of 95 % at 397 nm and an effective focal length of 66.83 mm. It was specified to produce a magnification of 20 when placed so that the lens' first surface is 59.2 mm from the ion along the optical axis (including 6 mm of fused silica). In this configuration, the image is formed 1.4 m from the lens and the numerical aperture is approximately 0.3. The lens is contained in an aluminium housing with an M47 x 0.75 thread at one end.

The imaging system uses a two-stage design, shown in figure 4.10. The first stage is composed of the IBK lens and a second $f = 250$ mm achromatic doublet⁴⁰, which serves to decrease the effective focal length of the IBK lens, allowing a more compact design. An adjustable aperture⁴¹ in the first stage's image plane is used to block stray light. This aperture may be opened during beam alignment to allow a $500 \mu\text{m} \times 500 \mu\text{m}$ area of the trap's surface to be viewed on the camera. A removable beam-splitter allows one to direct light to either the camera or the PMT.

The second stage optics re-image the ions onto the camera/PMT. The combination of $f = 40$ mm and $f = 250$ mm lenses in the camera are chosen to produce a total magnification of 16. The two lenses in the PMT arm are non-catalogue UV achromats from Comar with transmission guaranteed to be $> 90\%$ in the range 360 nm-700 nm and typical transmission of 97.5 % in the range 390 nm-425 nm. These are not guaranteed diffraction limited, but are designed for focusing UV beams to a tight spot. The 200 mm lens is a standard

⁴⁰Comar 50DQ250 2" diffraction limited achromatic doublet, $\sim 90\%$ transmission at 397 nm.

⁴¹Ealing Catalog Inc. MiniMech rectangular aperture. Four independent knife edges adjustable from fully closed to 12 mm x 12 mm.

Comar achromat.

We use an Electron tubes P25PC-12 PMT (serial number 9584) with a specified QE of $\simeq 27\%$ at 400 nm and a dark count rate of 100 s^{-1} at 20°C . The camera was an Andor Technology Luca DL-604M-OEM electron-multiplying CCD, with 1004×1002 $8\text{ }\mu\text{m} \times 8\text{ }\mu\text{m}$ pixels, 40 % QE at 400 nm and 13.5 MHz maximum pixel readout rate.

During normal operation we use filters in front of the PMT and camera to minimise background light. For the PMT we use Semrock FF01-377/50-25 and FF01-406/15-25 (combined transmission of 86 % at 397 nm) to block out room lights as well as all UV lasers apart from the 397 nm (the PMT is insensitive to IR). When observing fluorescence from neutral Ca during initial loading, we used a Semrock FF01-435/40-25 with a transmission of 93 % at 423 nm to reduce the background due to room lights. We use a Schott BG38 glass filter to block IR light on the camera. This is removed when aligning IR beams using the camera. For high-fidelity readout experiments (section 5.3) we removed all filters to maximise the detection efficiency. However, this necessitated turning off most of the room lights when running the experiment.

Using the technique described in [SKDB10], we measure the typical overall detection efficiency of the imaging system to be $\sim 0.23\%$ depending on the width of the aperture and which filters are in use.

We measured the imaging system's magnification by taking advantage of the fact that the trap structures are photolithographically defined and hence have well-known dimensions. We focused the imaging system on the trap using a 400 nm LED⁴² for illumination to minimise chromatic effects. Comparing the image on the camera to the known trap dimensions gives a magnification of 15.7, so that each pixel on the camera represents $0.51\text{ }\mu\text{m}$ in the focal plane.

⁴²Bivar LED5-UV-400-30

4.7 Microwave sources

We used a variety of microwave frequency sources during the course of this work, each chosen to suit a different need. In this section we describe these sources, explaining significant specifications and measurements made.

4.7.1 Rubidium and GPS frequency standards

All frequency sources used in this work were phase-locked to the 10 MHz output of an SRS FS725 rubidium frequency standard [SRS13]. The standard was purchased in 2005 with a specified initial accuracy of $\pm 5 \times 10^{-11}$ (± 0.16 Hz)⁴³. Its specified ageing rate after the first year is $< 5 \times 10^{-10}$ /year (1.6 Hz/year), corresponding to an error of less than 4×10^{-9} (12 Hz) at the time of writing. Its specified temperature and magnetic field sensitivities are $\frac{\Delta f}{f} < \pm 1 \times 10^{-10}$ over the range $+10^\circ\text{C}$ to $+40^\circ\text{C}$ and $\frac{\Delta f}{f} < \pm 2 \times 10^{-10}$ for a 1 G field reversal respectively. Given that our lab is typically stable to $\pm 0.5^\circ\text{C}$ and ± 10 mG, we expect temperature and field effects to be less than 10^{-11} (32 mHz).

For spectroscopy on the clock qubit, we wished to set the absolute frequency of our microwaves to better than 1 Hz. To do this we needed to calibrate our rubidium reference relative to TAI (international atomic time), which we did using a Trimble Thunderbolt GPS-disciplined oscillator. We logged the phase difference between the 10 MHz outputs from the two standards on an Agilent 53131A universal counter for a period of 24 hours. Over this averaging period the Thunderbolt has a specified accuracy of 1.16×10^{-12} (3.7 mHz). We performed this calibration twice, once in August 2012 and once in September 2013. In August 2012, we measured the rubidium reference to have a fractional error of $+3.22 \times 10^{-10}$ (+1.03 Hz), and in September 2013 we measured an error of $+4.03 \times 10^{-10}$ (+1.29 Hz)⁴⁴. As a consistency check, we took the September cali-

⁴³In this section, bracketed quantities denote errors referred to a 3.2 GHz reference frequency.

⁴⁴N.B. on sign conventions: in both cases the rubidium standard ran slightly faster than 10 MHz.




Figure 4.11: Single-sideband phase noise PSD for two Agilent synths. These graphs are taken from the manufacturer’s datasheets.

bration twice on adjacent days and found that the results agreed to 1.2×10^{-13} . For a period of one week after calibration, we expect the rubidium reference to accurately represent TAI to about 1×10^{-11} (32 mHz).

4.7.2 Microwave synths

We used three models of microwave frequency synthesiser during these experiments: Agilent E8247C PSG; Agilent E4422B ESG; and, Agilent E4426B ESG-AP⁴⁵. Details of these units may be found in the manufacturer’s data sheets and operator’s manuals. Here we note a few relevant features of the synthesisers.

The specified single-sideband phase noise PSD for the E8247C and E4426B are shown in figure 4.11. The E4426B-AP phase-noise is shown for a carrier frequency of 1 GHz, corresponding to the frequency used to clock the octupled DDS system discussed below. Phase noise for the E4422B model is significantly worse than for the E4426B: < -116 dBc/Hz at a 20 kHz offset from a 1 GHz carrier, compared with < -130 dBc/Hz for the E4426B.

For some experiments, such as the Ramsey experiments used to measure qubit coherence times, it is important to use two microwave sources with a well-defined phase difference. Since our synths are all phase locked to the rubidium

⁴⁵The E4426B is a high high spectral purity version of the E4422B.




Figure 4.12: Phase drift between the 3.2 GHz outputs of two synths (E4422B and E4426B) which have been phase-locked together at 10 MHz. The drift is likely to be mainly thermal, with the ~ 20 minute periodicity corresponding to our lab air conditioning cycles.

standard, one expects phase differences between them to be relatively stable over time. However, since the phase lock to the rubidium standard is at 10 MHz, it must be multiplied by a factor of 320 to produce a 3.2 GHz frequency. Any drifts in the phase lock point will be multiplied by this same factor, making us sensitive to small drifts.

To characterise such drifts, we used a MiniCircuits ZX05-73L-S+ mixer as a phase detector. We used an E4426B set to 4 dBm at 3.2 GHz as the mixer's local oscillator (LO) and an E4422B set to 0 dBm also at 3.2 GHz to drive the RF port. We 50 Ω terminated the mixer's IF port, logging the voltage across the terminator with a digital volt-meter (Fluke 8842A). We locked the two synths' 10 MHz references together. Under these conditions, the IF port is a DC voltage which varies with the sine of the phase difference between the LO and RF inputs, allowing the mixer to be used as a phase detector.

To calibrate the mixer, we recorded the IF voltage as a function of the phase set on the E4426B. The amplitude of the resulting sinusoid (150 mV in this case) determines the phase to voltage conversion. Having established this calibration




Figure 4.13: Apparatus used to multiplex and quadruple the outputs of the five HP8656B frequency synthesisers used for state preparation.

constant, we set the initial phase difference between the synths so that the IF voltage was approximately zero, maximising the measurement sensitivity. We then logged the IF voltage over a period of 24 hours. The results are shown in figure 4.12. The bulk of this drift is likely to be thermal, corresponding to changes in lab air temperature – in particular, the ~ 20 minute periodicity corresponds to our lab air conditioning cycles.

4.7.3 Quadrupled HP8656B

To prepare the qubit state, we use five microwave frequencies in the range 2.873 GHz to 3.094 GHz (see sections 2.6.1 and 2.6.2). We achieve this by quadrupling five HP8656B frequency synthesisers using the arrangement shown in figure 4.13.

The four switches (ZASWA-2-50DR+⁴⁶ absorptive high-isolation) are used to select between the five input frequencies. The first frequency doubler is a ZX90-2-11-S+ with typical conversion loss of 12.3 dB at a 2 GHz output frequency with 5 dBm input power (see datasheet for harmonic content specifications). The amplifier is a ZX60-33LN-S+ with gain of 16.5 dB at 1.6 GHz. The 7 dB attenuator (BW-S7WS+) ensures the power level is appropriate for the second doubler (ZX90-2-19-S+, conversion loss 9.7 dB for a 3.2 GHz output frequency with 5 dBm input power). The two filters (VLF-3000+, 3 GHz low pass, and VHF-2700A+, 2.9 GHz high-pass) remove harmonic content from the output signal before the final amplifier (another ZX60-33LN-S+).

⁴⁶All components used in this quadrupler were manufactured by Mini-Circuits.

We note two useful features arising from the non-linearity of the quadrupling process. Firstly, it provides very fine control of the microwave amplitude. This is important since our experimental control hardware has a timing resolution of $0.1\ \mu\text{s}$, which is extremely coarse for $1\ \mu\text{s}$ pulses. The HP8656B synths only have an amplitude resolution of $0.1\ \text{dB}$ corresponding to a fractional change in Rabi frequency of $1\ \%$, which is better but still insufficient for high-fidelity work. However, the non-linearity of the multiplication process greatly reduces the change in output power as a function of synth power.

The second advantage is that, since the quadrupling process is highly non-linear, any signals that are not directly connected to the quadrupler by the final switch are not multiplied. This eliminates leakage from all synths apart from the one which is connected to the multiplier at any point. This is important for long experiments such as coherence time measurements, which are sensitive to leakage. So long as the switches are set so that the frequency that is connected to the quadrupler does not correspond to a transition out of either qubit state, effects of leakage are virtually eliminated.

4.7.4 Octupled AD9910 DDS

For our single-qubit manipulations we required a phase-agile $3.2\ \text{GHz}$ source, which we produced using an octupled AD9910 DDS running at $400\ \text{MHz}$, clocked at $1\ \text{GHz}$ from an E4426B synth. The DDS is based on a $1\ \text{GSPS}$, 14-bit DAC and a 32-bit accumulator, allowing 6×10^{-5} fractional amplitude control and $\sim 0.23\ \text{Hz}$ frequency tuning resolution. We use the DDS in single tone mode, which provides 8 profiles with independent frequency, amplitude and 16-bit (0.005°) phase control.

Our DDS was an Analog Devices evaluation board (AD9910/PCBZ) with two modifications. Firstly, we replaced the reference current resistor, R6, (originally a Panasonic ERJ-2RKF1002X $10\ \text{k} \pm 100\ \text{ppm}/^\circ\text{C}$) with a $\pm 10\ \text{ppm}/^\circ\text{C}$ temp-

co. resistor for improved power stability.

The second modification was necessary because the evaluation boards are designed to be controlled by either an on-board programmable logic device (PLD) via USB or by externally connected logic inputs, with a jumper to select between the two choices. However, we wished to program the device via USB, whilst controlling the profile select and IO UPDATE pins from external inputs. To avoid a bus conflict, we cut the profile select and IO UPDATE lines from the PLD and routed them through a SN74AC32N quadruple 2-input positive-OR gate running on a 3.3 V supply. We then connected our external logic signals to the second input of these OR gates to allow a limited degree of bus sharing. Since our experimental control hardware produces a logic high of 5 V we use potential dividers to reduce the logic level to 3.3 V.

The DDS has a minimum profile toggle period of 2 SYNC_CLK cycles. This means that when changing profile, one must either change all profile select pins on the same clock cycle or more than two clock cycles apart. Because our experimental control hardware is not synchronised with the DDS clock, we cannot ensure that all profile pins are switched on the same clock cycle. Thus, we drive the profile select pins from experimental control computer outputs which are chosen to switch $\sim 1 \mu\text{s}$ apart during the experimental dead-time. In a future system without this dead time, we will implement some latching logic on an FPGA to synchronise the control hardware with the DDS clock.

To maintain the DDS's low phase noise specification, we powered the analogue portion of the DDS from LT1963A low noise regulators⁴⁷. We have not yet measured whether these regulators degrade the DDS's phase noise.

The active octupler was built from standard MiniCircuits parts and is de-

⁴⁷For a thorough discussion of AD9910 DDS phase and amplitude noise, see http://martein.home.xs4all.nl/pa3ake/hmode/dds_intro.html and, in particular, http://martein.home.xs4all.nl/pa3ake/hmode/dds_ad9910_pmnoise.html. Note that using LM317s for the analog 1.8 V clock circuitry increases the phase noise by more than 10 dB at some frequencies compared to a home-brew ultra-low noise regulator.




Figure 4.14: (a) AD9910 residual phase (taken from manufacturer’s data sheet). (b) Black line : Agilent E4426B phase noise at 1 GHz (from data sheet). Orange line: AD9910 phase at 400 MHz increased by 8 dB to account for frequency multiplication to 1 GHz.

scribed in table 4.1. The design was originally chosen and optimised for the frequencies used during state preparation and operates with output frequencies between 2.9 GHz and 3.1 GHz. It produces an output power of $\sim 10\text{ mW}$ at 3.2 GHz. The amplitude of the output signal varies approximately linearly with small changes in the DDS amplitude scale factor (ASF) as

$$\Delta V_{\text{out}} = 1.39 \Delta \text{ASF} \quad (4.1)$$

Over a period of a few hours, we see typical amplitude fluctuations of $\lesssim 10^{-3}$. We believe that this was one of the dominant sources of error during our randomized benchmarking experiments (see section 5.7.3).

4.7.4.1 Phase noise

We have not directly measured the phase noise characteristics of the octupled system, however one can estimate it based on the specifications of the components used. The residual phase noise PSD (i.e. the phase noise added by the DDS to the clock’s phase noise) is shown in figure 4.14a. We can approximate the phase noise PSD, $L(f)$, for the 400 MHz curve as

$$L(f) = 10^{-12.5} \times \frac{1\text{ kHz}}{f} \text{ Hz}^{-1} + 10^{-15.2} \text{ Hz}^{-1} \quad (4.2)$$

Part number	Notes
AD9910/PCBZ	DDS evaluation board with minor modifications noted in the text. Max output power at 400 MHz $\sim$ 1 mW (inv. sinc filter off).
ZFDC-10-1-S+	10.5 dB directional coupler, allows the base-band RF to be monitored on a fast 'scope for diagnostic purposes.
BW-S12W2+	12 dB fixed precision attenuator, temperature coefficient 0.0004 dB/dB/°C (46 ppm/°C fractional amplitude stability). Sets input power for first doubler.
SLP-450+	DC-400 MHz low-pass filter. Removes spurs from DDS before doubling.
SHP-300+	290 MHz high-pass filter.
ZX60-33LN-S+	Low noise amplifier (1.1 dB NF), 21 dB gain @ 400 MHz.
ZX90-2-13-S+	Frequency doubler. 11 dB conversion loss.
BW-S6W2+	6 dB attenuator.
SHP-600+	600 MHz high-pass filter
SLP-850+	DC-780 MHz low-pass filter.
ZX60-33LN-S+	Amplifier, 19 dB gain @ 800 MHz.
ZX90-2-11-S+	Frequency doubler 10 dB conversion loss @ 1.6 GHz output.
BW-S3W2+	3 dB attenuator.
VHF-1200+	1.22 GHz high-pass filter.
VLF-1575+	DC-1.575 GHz low-pass filter.
ZX60-33LN-s+	17 dB gain @ 1.6 GHz.
ZX90-2-19-S+	Frequency doubler 9.7 dB conversion loss @ 3.2 GHz output.
VHF-2275+	2.45 GHz high-pass filter.
VLF-3000+	DC-3 GHz low-pass filter.
ZX60-33LN-S+	Amplifier, 12 dB gain @ 3 GHz.

Table 4.1: List of parts used in the octupled DDS system. Components are listed in order from DDS to the final amplifier. Components were connected using semi-rigid SMA or SMA 'T' pieces and were mounted in a mechanically rigid fashion using hot glue.




Figure 4.15: AM noise on AD9910 DDS. The blue curve shows measurements made using a diode detector with the DDS running at 37.1 MHz (see section 6.5). The black points represent measurements taken using a notch filter method with the DDS running at 14.4 MHz (see http://martein.home.xs4all.nl/pa3ake/hmode/dds_ad9910_amnoise.html).

The black curve in figure 4.14b shows the phase noise of the E4426B synth used to provide the 1 GHz DDS clock. The orange curve in this figure shows the phase noise of the AD9910 running at 400 MHz, increased by 8 dB to account for the factor of 2.5 frequency difference. We see that the two sources are comparable.

Since the power at the input to the multipliers is always $\gtrsim 5$ dBm, it is likely that the multiplication will not add any noise on top of the 18 dB associated with multiplication to 3.2 GHz.

To a reasonable approximation, we thus assume that neither the clock nor the multiplication adds any excess noise, leading to an overall system phase noise of roughly

$$L(f) = 10^{-10.7} \times \frac{1 \text{ kHz}}{f} \text{ Hz}^{-1} - 10^{-13.4} \text{ Hz}^{-1} \quad (4.3)$$

4.7.4.2 Amplitude noise

During our investigation into amplitude noise on RF sources (section 6.5), we measured the AM noise of an AD9910 evaluation board at 37.1 MHz. The

results are plotted in figure 4.15. Also plotted in this figure are the results of similar measurements taken from http://martein.home.xs4all.nl/pa3ake/hmode/dds_ad9910_amnoise.html using a notch filter method and running the DDS at 14.3 MHz. The two data sets show reasonable agreement to 3 dB.

It is likely that this noise is simply additive noise on the DAC and should thus be independent of the output RF frequency. Further measurements would be needed to determine the effect of the octupler on the overall system AM noise. This noise is sufficiently large that it is unlikely to be increased by the octupler – indeed, it is possible that the saturated nature of the octupler may act as a limiter and reduce this noise.

4.7.5 IQ Mixers modulated at D.C.

To produce a microwave field gradient while maintaining a field null at the ions' position, we will need to drive each of the trap's three electrodes with microwaves with independently controllable amplitudes and phases. The sources used for this must be stable to of order 0.1 % and 0.1° over a period of a day so that frequent re-calibration is not required to maintain the field null. We achieve this using IQ mixers (Eclipse microwave IQ2040M 2 GHz-4 GHz, 10 dBm LO input) modulated at DC.

The basic form of an IQ mixer is shown in figure 4.16. The input signal is split into in-phase and quadrature (sine and cosine) components, which are independently amplitude modulated before being re-combined. The output signal is thus of the form

$$V_{\text{out}} = A \sin \omega t + B \cos \omega t = \sqrt{A^2 + B^2} \sin(\omega t + \phi) \quad (4.4)$$

where

$$\phi = \begin{cases} \arcsin \frac{B}{\sqrt{A^2+B^2}} & A \geq 0 \\ \pi - \arcsin \frac{B}{\sqrt{A^2+B^2}} & A < 0 \end{cases} \quad (4.5)$$




Figure 4.16: Schematic of an IQ mixer. The LO is split into in-phase and quadrature components by a 90° hybrid. The quadratures are then modulated independently by double-balanced mixers before being re-combined by a 0° hybrid.

With appropriate modulation of the two ports one can produce an output signal with an arbitrary amplitude and phase.

One can modulate the I and Q components either at DC, producing output RF at the LO frequency, or at an intermediate frequency (IF), producing output RF with components at $LO \pm IF$ ⁴⁸. An important feature of IQ mixers stems from the fact that the RF components at $LO \pm IF$ have an opposite phase relationship with respect to the carrier (see equations C.66). Thus, by setting the phase difference between the I and Q ports appropriately, one can create RF with only a single frequency component at either $LO + IF$ or $LO - IF$. This process is termed single-sideband mixing for obvious reasons⁴⁹.

In this work, we chose to bias our mixers at DC. The reason for this choice over a single-sideband approach was simplicity since DC currents were easier for us to generate than the required RF. In the future, we plan to switch to a single-sideband system in which the I and Q inputs are driven at ~ 400 MHz by a pair of DDSs. The advantage of the DDS-driven single-sideband approach is that the amplitude and phase of the output RF can be controlled directly from

⁴⁸Here, we assume that the same IF frequency is used for the I and Q ports.

⁴⁹Note that in practice some of the carrier leakage will be present at the RF output as well as the unwanted sideband. This leakage can be reduced by adding a DC bias to the I and Q IF.




Figure 4.17: Test set-up used to calibrate IQ mixers. See text for details.

the DDS with extremely high resolution and without the need for a complicated calibration procedure (see below). This makes dynamic control of the RF phase and amplitude straightforward. We expect the phase drift between DDS chips clocked from the same 1 GHz clock to be extremely low.

We generate the I and Q bias currents using an AD5372 32-channel, 16-bit DAC. Each DAC channel is buffered using an OpAmp-based⁵⁰ buffer with a voltage gain of 0.15 and a 3 dB-frequency of 30 Hz. 50 Ω resistors on the buffers' outputs act to limit the current through the mixers.

4.7.6 Calibration

We calibrated the RF power and phase produced by the IQ mixers as a function of I and Q DAC voltages using the set-up shown in 4.17. The calibration procedure was automated using a *MatLab* script and test and measurement equipment connected to a PC via USB and GPIB. All components were bolted to an optics table and connected using semi-rigid SMA cabling for thermal and mechanical stability.

The mixer's LO is driven at 3.2 GHz with 10 dBm by an E8247C synth. A directional coupler (pulsar CS20-04-436/13) and a power detector (Agilent V3500A) allow us to measure the RF output power. Our phase detector is based on a second mixer (MiniCircuits ZX05-73L-S+), driven at 3.2 GHz with 4 dBm by an E4422B phase locked to the E8247C. We monitor the IF voltage from this mixer using a digital multimeter (Fluke 8842A).

⁵⁰OP27

We define zero phase to be the phase, ϕ , set on the E8247C which produces 0 V on the phase detector with positive $\frac{dV}{d\phi}$ when the I and Q DACs are both set to 10 V. To find the phase produced by a given combination of I and Q DAC voltages, we use an iterative procedure to find the E8247C phase which produces zero volts on the phase detector. By checking the zero-phase point at regular intervals, we are able to remove effects of long-term drift between the two synths.

We found that the IQ mixer's VSWR changed with the I and Q voltages. Since we were attempting to make a relatively precise measurement and all components of the test setup are sensitive to VSWR – for example, the output power from the synths changes with the load VSWR – we found it necessary to insert isolators (Sierra Microwave Technology 1CY63) liberally through the test setup. We also found that the IQ mixer's RF output contained some harmonic content, which interfered with our measurement. To remove this, we added a pair of low-pass filters (MiniCircuits VLF-3800+).

We measured the RF power and phase for I and Q voltages between ± 10 V at 200 mV intervals. The result is a surface like that of figure 4.18. To use the data in an experiment, we need a lookup table of I and Q voltages, corresponding to a given power and phase. To this end, we fitted contours of constant power to the measured data, which we interpolated to find points at regular phase intervals.

We find that the IQ mixers have a useful dynamic range of about 10 dB and a maximum output power of ~ 0 dBm under these bias conditions. To estimate errors in the calibration process, we used our lookup table to attempt to set the mixers to particular powers and phases, while using the calibration setup to see what was actually produced. We find the calibration is accurate to $\sim \pm 0.1$ dB ($\sim \pm 1\%$ signal amplitude) and $\sim \pm 0.2^\circ$, which is expected to be sufficient for the experiments we intend to carry out. We also used the calibration setup to test the mixer's temperature stability. Over 24 hours, we found fluctuations of




Figure 4.18: Measured power and phase of IQ mixer RF as a function of I and Q voltage. Black dots show fitted contours of constant power and phase.

order ± 0.01 dB ($\sim \pm 0.1$ % amplitude) and $\pm 0.1^\circ$.

5

Single-Qubit Results and Analysis

In this chapter we present experimental work involving single trapped $^{43}\text{Ca}^+$ ions. This work may be broadly divided into four experiments: cooling $^{43}\text{Ca}^+$ at 146 G with the aim of obtaining a high rate of fluorescence (section 5.1)¹; high-fidelity preparation and read out the clock-qubit (sections 5.2 and 5.3); spectroscopy of the clock transition and measurement of qubit coherence time (sections 5.5 and 5.6); and, finally, randomized benchmarking of single qubit gates (section 5.7).

For each of these experiments we describe the set up and calibration of experimental parameters, give experimental results and discuss sources of error.

5.1 Doppler cooling

In this section we present an experimental realisation of the cooling scheme described in section 2.5. Our primary goal is the production a large rate of fluorescence in order to allow high-fidelity state detection (see section 5.3). In contrast, we are less concerned with the ion's temperature, which does not affect the fi-

¹We discuss cooling optimised instead for lower temperatures by taking advantage of a dark resonance in section 6.3.

delity of any single-qubit operation until the ion becomes sufficiently hot for the 393 nm shelving transition to be Doppler broadened (see sections 2.6.3 and 5.6).

5.1.1 Setting up and calibrating experimental parameters

5.1.1.1 397 nm and 866 nm polarizations

To efficiently cool $^{43}\text{Ca}^+$ at 146 G the 397 nm cooling beam should contain no σ^- component. Given the requirement that the 397 nm beam should couple to all of the ion's normal modes, this can only be achieved using a beam with polarization ratio of 2/3 of σ^+ to 1/3 of π . We produce this polarization using a half-wave and a quarter-wave plate set to appropriate angles (see section A.4).

Numerical modelling [Jan] suggested that the cooling process is relatively insensitive to the exact polarization state of the 866 nm laser due to the low branching ratio ($\simeq 5\%$) for decays from the $4\text{P}_{1/2}$ to the $3\text{D}_{3/2}$ states. It is, however, convenient to use a beam whose polarization state is well known, since this simplifies comparison between theory and experiment. We chose to use equal σ^- and σ^+ with no π light since this is easy to implement and should produce as good a rate of fluorescence as other choices.

To set the 866 nm beam's polarization state, we temporarily placed a PBS cube after the last polarization optic before the trap (figure 4.6) to act as a polarization diagnostic. We then adjusted the red half- and quarter-wave plates until no light was transmitted by this polarizer, corresponding to a beam with vertical polarization. This empirical method was necessary to compensate for the effect of the 397 nm wave plates on the 866 nm polarization.

5.1.1.2 Laser intensities

For most experiments we calculate our laser intensities at the trap's centre from measured beam powers and waist sizes. For the data presented in this section we performed an independent calibration by loading a $^{40}\text{Ca}^+$ ion at a field of

$B_0 \simeq 10$ G and recording the ion fluorescence as a function of 866 nm detuning at different laser powers. The 397 nm and 866 nm beam intensities were then deduced from fits to the experimental data.

We used $^{40}\text{Ca}^+$ rather than $^{43}\text{Ca}^+$ for this since its simpler level structure makes fitting the resulting scans simpler and more reliable. The complication with $^{43}\text{Ca}^+$ largely stems from the fact that when the 866 nm detuning is set to place the laser on the blue side of a two-photon dark resonance (see section 6.3), the cooling lasers heat the ion. The observed fluorescence is strongly affected by the motion of the ion because the Doppler shifts then vary over a large range. As there are so many dark resonances, we find that there are only relatively small regions of the 866 nm scan where the ion is sufficiently cold that Doppler shifts do not distort the fluorescence profile. Since the fits are performed on the basis of a stationary ion, this results in a poor match between theory and experiment.²

5.1.1.3 397 nm EOM

As discussed in section 2.5, cooling $^{43}\text{Ca}^+$ requires a pair of 397 nm beams to drive transitions from both ground-level hyperfine manifolds. In our experiment both of these beams are derived from the same laser source using an EOM. We drive this EOM at 2.930 GHz using a VCO³ amplified to ~ 0.5 W. At this power level the carrier and sidebands have comparable intensities when viewed on an optical spectrum analyser⁴. This frequency and power level were chosen on the basis of simulations, optimising fluorescence.

We use the modulated laser's carrier to drive transitions from the $F = 3$ manifold and its first blue sideband to drive transitions from the $F = 4$ manifold.

²The 146 G $^{43}\text{Ca}^+$ dark resonances do, however, come in useful when interpreting the $^{40}\text{Ca}^+$ scans by allowing us to calibrate our 866 nm frequency scale: we scan our 866 nm laser's detuning by applying a voltage to a piezo in the reference etalon the laser is locked to (section 4.5). Since the dark resonances have well known spacings, one may use them to provide an accurate calibration of the change in 866 nm frequency as a function of piezo voltage.

³MiniCircuits ZX95-3050.

⁴In the experiments reported in this section each sideband was $\sim 85\%$ as intense as the carrier, although the exact value does not make a significant difference to the experimental results.

This choice is not significant for cooling, but allows improved state preparation (section 2.6.1).

5.1.1.4 397 nm & 866 nm frequencies

During initial loading we set our 397 nm & 866 nm frequencies using a wavemeter⁵. To ensure the wavemeter's absolute accuracy we first calibrate it by loading a $^{40}\text{Ca}^+$ ion at $B_0 \sim 1$ G and finding the lasers' resonances. Because $^{40}\text{Ca}^+$ at low field makes such a convenient frequency reference, we will tend to specify laser detunings relative to the respective transitions in $^{40}\text{Ca}^+$ in zero field.

For optimal fluorescence the 397 nm laser should be set so that carrier (blue EOM sideband) detuning relative to zero-field $^{40}\text{Ca}^+$ is approximately -1200 MHz ($+1725$ MHz). At this detuning the blue EOM sideband is close to resonance with the $4S_{1/2}|4, +4\rangle \Leftrightarrow 4P_{1/2}|4, +4\rangle$ cooling transition, while the carrier is slightly to the red of the transitions from the $4S_{1/2}|3, +3\rangle$ state (see below).

When loading ions using the wavemeter for frequency calibration we typically set our laser ~ 100 MHz to the red of this. Once an ion has been loaded, we tune the 397 nm closer to resonance to improve fluorescence. We typically scan the 397 nm laser to the blue until the cooling lasers' net effect is to heat the ion, resulting in a disappearance of fluorescence and often in ion loss⁶. We find that in our experiment the frequency that this point occurs at is relatively insensitive to both the 866 nm frequency and the intensities of either beam. We typically operate with the 397 nm ~ 5 MHz to the red of this point.

The 866 nm detuning is then scanned to find the detuning that gives the highest fluorescence (see figure 5.1). Care must be taken, however, since it is possible to pick 866 nm detunings where there is significant heating, resulting in a de-

⁵High-Finesse WS-7 with 10 MHz resolution and 60 MHz absolute accuracy.

⁶Note, however, that there are 397 nm detunings ~ 400 MHz to be blue of this point where it is possible to load and cool an ion. Under these conditions the ion is typically significantly hotter than normal and fluorescence is decreased by a factor of roughly two.

crease in readout fidelity. This can typically be avoided by red detuning a few megahertz from the point of highest fluorescence.

5.1.1.5 Detection efficiency

The fluorescence one measures from an ion is the product of three quantities: the population of the $3P_{1/2}$ level, the Einstein A coefficient for the $4P_{1/2}$ – $4S_{1/2}$ transition and the imaging system’s net detection efficiency. When comparing experimentally measured fluorescence rates with theoretical predictions, whose currency is atomic populations, it is necessary to know the collection efficiency to sufficient accuracy for the purpose. We measure this using a single trapped $^{40}\text{Ca}^+$ ion with the technique described in [SKDB10].

5.1.2 Results

Our typical cooling diagnostic is provided by scanning the 866 nm laser while recording the ion’s fluorescence. We scan the 866 nm rather than the 397 nm laser because it may be taken over a much larger frequency range without disrupting the cooling too severely and hence provides a better diagnostic – if the 397 nm is taken too far to the blue then heating results and the ion is often lost, while if it is taken too far to the red then cooling becomes extremely inefficient and the fluorescence drops off rapidly.

Figure 5.1 shows an 866 nm scan using our standard high-fluorescence cooling parameters. Here, the 397 nm and 866 nm intensities are $116(20) I_0$ and $750(150) I_0$ respectively. The 397 nm power is divided between the carrier and sidebands, with $43(8) I_0$ in the carrier and $37(7) I_0$ in each sideband. The 397 nm carrier (sideband) detuning is $-1200(15) \text{ MHz}$ ($+1725(15) \text{ MHz}$) relative to zero-field $^{40}\text{Ca}^+$, which is -1893 MHz ($+1037(15) \text{ MHz}$) relative to the 397 nm transition’s centre of gravity. We achieve a peak $4P_{1/2}$ -state population of $13(1) \%$ at an 866 nm detuning of $-33643(15) \text{ MHz}$ relative to zero-field $^{40}\text{Ca}^+$




Figure 5.1: Detuning scan of the 866 nm laser using our standard high-fluorescence cooling parameters (see text). The 866 nm frequency scale is relative to the transition’s centre of gravity ($-3\,464.3$ MHz relative to zero-field $^{40}\text{Ca}^+$). We achieve a peak $4P_{1/2}$ -state population of 13 (1) % using 866 nm detunings of approximately $-3\,364.3$ MHz relative to zero-field $^{40}\text{Ca}^+$ ($+100$ MHz relative to the 866 nm transition’s centre of gravity).

($+100$ (15) MHz relative to the 866 nm transition’s centre of gravity).

The dashed line in figure 5.1 shows a simulation using our estimated parameters, assuming a laser line-width of 100 kHz and a stationary ion. There is a pleasing agreement between theory and experiment across much of the measured profile. The cooling lasers tend to heat the ion in regions around the blue side of dark resonances, causing Doppler broadening, which distorts the fluorescence profile.

It is instructive to consider where these parameters place the 397 nm laser’s detuning relative to transitions from states with $M \geq 3$ in the ground level, as shown in figure 5.2. We see that the 397 nm carrier is close to resonance with transitions from $4S_{1/2}|3, +3\rangle$ to the $F = 4$ states in the $4P_{1/2}$ level. Similarly, the blue sideband is close to resonance with transitions from the $4S_{1/2}|4, +4\rangle$ and $4S_{1/2}|4, +3\rangle$ states to the same $F = 4$ states in the $4P_{1/2}$ level. We do not show an equivalent diagram for the 866 nm transition since the presence of both σ^+ and σ^- polarization components in this beam means that there are enough transitions that the diagram is difficult to interpret.




Figure 5.2: Frequencies and branching ratios of 397 nm σ^+ and π -polarized transitions from states in the $4S_{1/2}$ level with $M \geq 3$. The transition between states $4S_{1/2}|F, M_F\rangle$ and $4P_{1/2}|F', M'_F\rangle$ is labelled as $(F, M_F) \leftrightarrow (F', M'_F)$. The frequency scale is relative to the 397 nm transition's centre of gravity. The laser detunings are based on fits to the data and have an uncertainty of approximately 20 MHz.




Figure 5.3: Detuning scan of the 866 nm laser using different 866 nm intensities. All parameters apart from the 866 nm intensity are kept fixed at our standard high-fluorescence cooling parameters (as illustrated in figure 5.1).

Figure 5.3 illustrates the effects of changing intensity of the 866 nm laser. Here, all parameters apart from the 866 nm intensity are kept fixed at our standard high-fluorescence cooling parameters (as illustrated in figure 5.1). In general, the fluorescence decreases with decreasing 866 nm intensity. Above $600 I_0$, the fluorescence continues to increase, however the ion becomes sufficiently hot that Doppler broadening begins to reduce our readout fidelity (section 5.3).

5.2 State preparation

As described in section 2.6.1, we prepare the ion in the $|3, +1\rangle$ qubit state using a two-stage process consisting of optical pumping to the stretched state, followed by a series of three microwave transfer pulses. In this section we discuss how the experimental parameters affecting both of these processes may be chosen and controlled to maximise the state-preparation fidelity. We describe set up of the 397 nm σ^+ optical pumping beam, characterisation of the microwave-enhanced optical pumping technique and the choice of π -time on microwave transfer pulses to minimise errors due to magnetic field noise and off-resonant excitation.

Since it is difficult to measure the state preparation and readout fidelities independently, we defer discussion of the measurement of this fidelity until section 5.3, where we present measurements of the combined state preparation and measurement (SPAM) fidelity.

5.2.1 Optical pumping to the stretched state

High-fidelity optical pumping to the stretched state relies on accurate control of the 397 nm σ^+ beam polarization. We achieve this using a high-quality polarizer⁷ followed by a quarter-wave plate (section 4.5).

⁷Thorlabs GT-10A Glan Taylor polarizer with 100,000 : 1 extinction ratio.

5.2.1.1 397 nm σ^+ set-up

We set up the 397 nm σ^+ beam path using a single $^{40}\text{Ca}^+$ ion at a field of approximately 10 G. Alignment consists of cooling the ion using the 397 nm σ^+ beam with intense ($\sim 500 I_0$) 850 nm and 854 nm repumpers. If the waveplate angle and tilt are set correctly and the beam is aligned along the magnetic field axis then the ion is optically pumped to $|M_J = +1/2\rangle$ and fluorescence stops. Thus, by iteratively adjusting these degrees of freedom, the fluorescence may be minimised, maximising the polarization purity.

We use $^{40}\text{Ca}^+$ with 850 nm and 854 nm cooling for this to maximise fluorescence and facilitate alignment. To ensure that the ion is not lost during alignment, we alternate counting fluorescence photons for 10 ms with 10 ms of cooling with the 397 nm Doppler laser.

Once the waveplate has been set properly and the magnetic field aligned along the beam axis at a low field, we turn the 146 G coils on and load a $^{43}\text{Ca}^+$ ion. Moving to the intermediate field does not affect the waveplate's angle or tilt, but will change the magnetic field axis somewhat. To re-align the field with the σ^+ beam, we turn the 397 nm σ^+ beam power up to maximum power ($\sim 100 I_0$) and measure the state-preparation and readout fidelity. At this high power, effects of polarization impurity are increased, causing state preparation errors. This allows small errors in the magnetic field angle to be detected.

We measure the state preparation and measurement fidelity (SPAM) as a function of the current in each of the coils in turn, fitting a quadratic to the result, from which the optimum field may be found.

5.2.1.2 Microwave-enhanced optical pumping

In section 2.6.1 we described a scheme for improving optical pumping by using microwave π -pulses on the $|4, +3\rangle \Leftrightarrow |3, +3\rangle$ and $|4, +2\rangle \Leftrightarrow |3, +2\rangle$ “D” and “E” transitions. Crude simulations using the rate equations suggest that this




Figure 5.4: Pulse sequence used to test the microwave-enhanced optical pumping technique. The ion is first optically pumped with the 397 nm σ^+ beam for 100 μ s using the carrier and EOM sideband to address both hyperfine manifolds. Next, the EOM sideband addressing the $F = 4$ manifold is turned off for a further 50 μ s of optical pumping. Microwave π -pulses are then applied to the D and E transitions before further EOM-off optical pumping. After applying n 397 nm σ^+ pulses without the EOM and $(n - 1)$ microwave pulses the ion's state is measured.

technique should improve the fidelity of state preparation by roughly an order of magnitude, depending on the exact laser parameters used.

When the 397 nm σ^+ optical pumping beam is set optimally, the stretched state is prepared with fidelity of $\gtrsim 99.9\%$, making it difficult to assess the impact that the microwave pulses have on the preparation fidelity. To make this easier to measure, we deliberately mis-aligned our magnetic field with respect to the blue σ^+ beam by changing the currents in our coils. This induced a state preparation error of $\sim 30\%$. We then performed the experiment shown in figure 5.4 to determine how this error was affected by the microwave technique. We chose our coil currents to tilt the field by an angle of approximately 8° to the horizontal while maintaining a field magnitude of 146 G.

To understand the effect of this tilt, we recall (see, e.g. [Szw09]) that a σ^+ -polarized beam of intensity I_0 at an angle θ to the magnetic field contains polarization components

$$\begin{aligned} I_\pi &= I_0 \frac{(\sin \theta)^2}{2} = 0.01 I_0 \\ I_{\sigma^+} &= I_0 \frac{(1 + \cos \theta)^2}{4} = 0.99 I_0 \\ I_{\sigma^-} &= I_0 \frac{(1 - \cos \theta)^2}{4} = 2 \times 10^{-5} I_0 \end{aligned} \tag{5.1}$$

Thus, the dominant effect of this angle is to introduce a 1% π -polarized component into the optical pumping beam. The polarization impurity will decrease the fidelity of both the readout and the state-preparation, however the state




Figure 5.5: Microwave enhanced optical pumping with a deliberately mis-aligned magnetic field. The abscissa labels the number of 397 nm σ^+ pulses with the EOM off. Thus, the first data point ($n = 0$) represents normal optical pumping with the carrier and sideband. The second data point has one EOM-off pulse, with no microwaves, while subsequent points have $n - 1$ microwave pulses. We see that the error in state preparation is reduced by more than a factor of two by a single EOM-off pulse and is improved by more than an order of magnitude after $\simeq 8$ microwave pulses.

preparation is expected to be much more sensitive to such errors (see section 5.3).

The results of these experiments are shown in figure 5.5. During these experiments our readout fidelity was only 99 %, even with the magnetic field properly aligned. This was not carefully measured and normalised in the plotted data and so places a limit on the sensitivity of the measurement. We see that this limit is reached by the time $n = 8$, implying that we can correct a 30 % error in state preparation to $\lesssim 1$ %. This is consistent with our previous claim that the microwave pulses allow an increase in state preparation fidelity of roughly an order of magnitude.

5.2.2 Microwave transfer pulses

We transfer the ion from the stretched state to the $|3, +1\rangle$ qubit state using a series of three microwave π -pulses on the transitions "A", "B" & "C" (section 2.6.1). These pulses are affected by three main sources of infidelity: pulse area

(Rabi frequency / pulse length) errors, detuning errors due to fluctuations in the ambient magnetic field and off-resonant excitation of spectator transitions. Other sources of error, including fast phase and amplitude fluctuations in the microwave source and fast fluctuations of the transition frequency (e.g. due to magnetic field noise), are expected to be small and will not be discussed here.

5.2.2.1 Pulse area errors

To understand the effects of pulse area errors, we consider driving a two-level system using a resonant field. We assume that the system starts in one state and is driven for a period $\Omega t = \pi + \theta$, where θ represents the pulse area error. At the end of the pulse the population of the state we are driving the system into is given by

$$P = \cos^2 \frac{1}{2} \Omega t \simeq 1 - \frac{1}{4} \theta^2 \quad (5.2)$$

Thus, to perform a π -pulse with an error of 1×10^{-4} one requires the Rabi frequency/pulse duration to be correct to within 0.7 %. In practice, setting pulse lengths with this accuracy is simple. Moreover, typical drifts in the microwave amplitude are closer to 0.1 %. We thus do not expect pulse area errors to be a significant source of error in these experiments.

5.2.2.2 Magnetic field noise

The transitions used during state preparation have frequencies which depend on magnetic field to first order, making our preparation pulses relatively sensitive to field noise.

Our lab's magnetic field noise spectrum is a product of many different sources. We see a contribution of approximately 2 mG at 50 Hz and harmonics associated with the mains' electricity. To avoid errors due to this field, we typically synchronise our experiment with the mains cycle. For long experiments where we care about small preparation errors, we take care to ensure that all

pulses on field-sensitive transitions occur at the same phase of the mains cycle (i.e. separated by integral multiples of 20 ms). There are several ways of reducing the effect of noise that is not synchronous with the mains' cycle. The simplest is to ensure that the pulse is driven with a Rabi frequency that is large compared with the noise-induced frequency shift; however, composite pulse techniques may also be used [Lev86].

5.2.2.3 Off-resonant excitation

In our experiment we are unable to produce a microwave field with pure circular polarization. Thus, the field used to drive transitions between a pair of states in the ground-level will couple to all other "spectator" transitions that involve those states. Off-resonant excitation of these spectator transitions leads to two effects. Firstly, the frequency of the transition we are trying to drive is altered – the light or AC Zeeman shift. Secondly, small amounts of population will be transferred ("leak") into the spectator states.

The higher the Rabi frequency we drive transitions with the more severe are the effects off-resonant excitation. This is an important factor in our experiment since we wish to drive transitions as fast as possible to minimise their sensitivity to magnetic field noise.

To understand the effects of off-resonant excitation, we consider a set of states $\{|n\rangle\}$ in the $F = 3$ manifold coupled to states $\{|m\rangle\}$ in the $F = 4$ manifold via the Hamiltonian

$$H_1 = -\boldsymbol{\mu} \cdot \mathbf{B} \cos(\omega t - \phi_0) \quad (5.3)$$

We will work in the basis of interaction-free energy eigenstates, so that an arbitrary state may be expressed as

$$|\psi\rangle := \sum_n c_n(t) |n\rangle + \sum_m c_m(t) |m\rangle \quad (5.4)$$

We move into an interaction picture with respect to the free-atom Hamiltonian, defined by the transformation

$$\begin{aligned} |n\rangle &\mapsto e^{iE_n t/\hbar} |n\rangle \\ |m\rangle &\mapsto e^{iE_m t/\hbar} |m\rangle \end{aligned} \quad (5.5)$$

where E_i is the energy of the state $|i\rangle$ in the absence of the microwave field.

Our interaction picture Hamiltonian then has matrix elements

$$\begin{aligned} (H_I)_{nm} &= e^{i\omega_{nm}t} \langle n | -\boldsymbol{\mu} \cdot \mathbf{B} | m \rangle \cos(\omega t - \phi_0) \\ &:= \hbar\Omega_{nm} e^{i\omega_{nm}t} \cos(\omega t - \phi_0) \end{aligned} \quad (5.6)$$

where we have defined the transition frequencies $\hbar\omega_{ij} := E_i - E_j$.

Defining the detunings $\omega := \omega_{nm} + \delta_{nm}$ and making the rotating wave approximation, our Hamiltonian becomes

$$\begin{aligned} (H_I)_{nm} &= \frac{1}{2} \hbar\Omega_{nm} (e^{i(\omega_{nm}+\omega)t} e^{-i\phi_0} + e^{i(\omega_{nm}-\omega)t} e^{i\phi_0}) \\ &= \frac{1}{2} \hbar\Omega_{nm} e^{-i\delta_{nm}t} e^{i\phi_0} \end{aligned} \quad (5.7)$$

When we come to look at randomized benchmarking experiments in section 5.7 it will be convenient to be able to calculate a single propagator that is independent of the microwave phase. To do this, we hide the microwave phase in our definition of the atomic states by making the transformation

$$|\psi\rangle \mapsto |\tilde{\psi}\rangle := R^\dagger |\psi\rangle \quad (5.8)$$

where R is defined by the diagonal matrix elements

$$\begin{aligned} R_{mm} &:= e^{-i\phi_0/2} \\ R_{nn} &:= e^{i\phi_0/2} \end{aligned} \quad (5.9)$$

Our Hamiltonian then becomes

$$(\tilde{H}_I)_{nm} = \frac{1}{2} \hbar\Omega_{nm} e^{-i\delta_{nm}t} \quad (5.10)$$

We define the propagator, $U(t)$, by

$$|\tilde{\psi}(t)\rangle := U(t) |\tilde{\psi}(t=0)\rangle \quad (5.11)$$

which can be found as the solution to the differential equation

$$i\hbar \frac{dU(t)}{dt} = \tilde{H}_I(t)U(t) \quad (5.12)$$

subject to the initial condition $U = 1$. We solve this by numerical integration using *MatLab*'s *ode45* solver. Finally, one may re-introduce the microwave phase by inverting the transformation of (5.8) so that

$$U(t) \mapsto RU(t)R^\dagger \quad (5.13)$$

5.2.2.4 Predicted errors

To understand the trade-off between off-resonant excitation and sensitivity to magnetic field noise, we simulate the error of a π -pulse on each of the three transitions A, B & C for different π -times. We consider both resonant microwaves and microwaves with detunings corresponding to varying levels of magnetic field offsets.

The results of these simulations are shown in figures 5.6a to 5.6c. In each case we calculated the microwave field required to achieve the desired π -time given the measured polarization angle at the transition frequency (section 6.2). We see that errors of 1×10^{-4} should be possible on all transitions given field fluctuations of $\lesssim 1$ mG, with optimum π -times of around $1 \mu\text{s}$ – $2 \mu\text{s}$.

Due to the presence of high-frequency Rabi flopping on the spectator transitions, pulse fidelities are not in general maximised for pulse durations equal to the π -time. The frequency of this dynamics is set by the ~ 100 MHz detunings of the spectator transitions, giving a periodicity of roughly 10 ns. More detailed simulations suggest that one could, in principle, take advantage of this fact to reduce errors due to off-resonant excitation by a factor of roughly two by careful choice of π -time. This behaviour is visible (although heavily aliased) in the simulations of transitions B & C, where we see that the error at the π -time may be minimised by careful choice of the π -time. We do not, however, attempt to take advantage of this fact here.

In these simulations we have set the microwaves to be resonant with the nominal transition frequency, without attempting to correct for AC Zeeman shifts. In figure 5.6d, we simulate the error of a π -pulse on transition B, with and without compensation for AC Zeeman shifts. We see that the improvement due to AC Zeeman shift correction is relatively small, suggesting that the dominant source of error is leakage into spectator states rather than the detuning error due to the shift itself.

5.3 Readout

The readout scheme, described in 2.6.3, consists of first selectively shelving one qubit state in the $3D_{\frac{5}{2}}$ level by optical pumping before turning on the 397 nm & 866 nm cooling lasers and detecting fluorescence photons. In this section, we discuss optimal experimental parameters for both of these processes. We present experimental results for the combined fidelity of state preparation and measurement (SPAM) and analyse sources of error present in these experiments.

5.3.1 Shelving

5.3.1.1 Optimising parameters

We shelve the ion using the pulse sequence (illustrated in figure 5.7)

$$\{393 \text{ nm}, 850 \text{ nm } \sigma^+, 850 \text{ nm } \pi\}^n \text{ 393 nm} \quad (5.14)$$

We model this process in *MatLab* using rate equations [Szw09]. The experimental parameters are: the frequencies and intensities of the 393 nm and the three 850 nm lasers; the length of each pulse; and, the number of repetitions of the 393 nm-850 nm-850 nm cycle (n). Additionally, one has the choice of which state in the $F = 3$ manifold to use as the bright (unshelved) state.

We performed experiments using two different choices of bright state. We used the $|3, +3\rangle$ state in conjunction with the $|4, +4\rangle$ stretched-state to form a




Figure 5.6: (a)-(c): errors during π -pulses on transitions "A" "B" and "C" due to constant magnetic field offsets and off-resonant excitation of spectator states. We assume that the transition is driven by a microwave field that is detuned by an amount corresponding to a constant magnetic field error. Solid lines represent simulations on the full 16 state ground-level and so represent errors due to both detuning and off-resonant excitation. Dashed lines represent simulations of a two-level system and so do not include off-resonant effects. Simulations assume the measured polarization angles given in section 6.2. (d) Simulation of errors due to off-resonant excitation on transition B. The "uncorrected" curve shows errors for microwaves that are resonant with the nominal transition frequency, while the "corrected" curve shows errors for a microwave field whose frequency is set to attempt to compensate for AC Zeeman shifts.




Figure 5.7: Shelving of the $|4, +4\rangle$ state by optical pumping (repeat of figure 2.20): a $393\text{ nm } \sigma^+$ laser connects the $|4, +4\rangle$ state to the $M = +5$ state in the $3P_{3/2}$ level. From here, it can decay back to the $|4, +4\rangle$ state (93.5 % branching ratio) or into either the $3D_{3/2}$ or $3D_{5/2}$ levels. Decays to the $3D_{3/2}$ level (0.66 % branching ratio) must be to states with $M = 4$ or $M = 5$ and are repumped to the $|P_{3/2}, M = 5\rangle$ state using 850 nm beams with σ^+ - and π -polarizations respectively. Once the ion decays to the $3D_{5/2}$ level (5.87 % branching ratio), it is removed from the 397 nm - 866 nm cooling cycle and is said to be “shelved”. States in the $|S_{1/2}, F = 3\rangle$ manifold are not shelved due to the 3.2 GHz ground-level hyperfine splitting.

“stretched qubit”. This is particularly convenient for use in various calibration and diagnostic experiments. Secondly, we used $|3, +1\rangle$ when reading out the clock-qubit because it may be reached from the $|4, 0\rangle$ qubit state using a single Q pulse.

Since we use one laser frequency for each transition we wish to address, the lasers are all tuned to their respective resonances, with the (slight) exception of the 850 nm - σ^+ beams. These lasers are on simultaneously, allowing the possibility of a dark state, which would reduce the repumping efficiency. To avoid this

possibility, we tune the beam repumping the $|D_{\frac{3}{2}}, 5, +4\rangle$ state +15 MHz from resonance. This detuning should be sufficient to avoid dark states without significantly reducing its effectiveness as a repumper.

For the sake of experimental simplicity, we use the same pulse length for the two 850 nm σ^+ lasers and we take the final 393 nm pulse to be the same length as the other 393 nm pulses. All other pulse lengths are free variables of our optimisation.

The probability of shelving the ion after a single 393 nm pulse is limited to 89.9 % by the ratio of decays to the $3D_{\frac{3}{2}}$ level compared with decays to the shelf. Because of this, each 393 nm-850 nm-850 nm cycle may only decrease the shelving error by a factor of roughly 10. To reach a 10^{-4} infidelity we thus expect to require a minimum of four 393 nm pulses ($n = 3$). This increase in shelving fidelity with the number of repeats of the 393 nm-850 nm-850 nm cycle is illustrated in figure 5.8. We see that the fidelity increases rapidly with the number of repeats until $n \sim 6$, after which there is little further increase. In practice, however, the sequence will become difficult to implement when the number of repeats becomes too large since the lengths of each pulse become extremely small. We chose to use 8 repeats as a balance between these factors.

The optimisation thus consists of choosing an intensity for each laser beam as well as lengths for the 393 nm, 850 nm- σ^+ and 850 nm- π pulses. The best choice of these parameters depends slightly on how pure the laser polarizations are. The experimental parameters used during the work reported in this thesis were based on optimisations assuming the following polarization impurities: for the 393 nm and 850 nm σ^+ beams, we assumed 0.5% of the intensity is in each of the σ^- and π components. For the 850 nm π beam path, which was the least carefully aligned, we assumed a 3% σ^- impurity. However, the low error rates we measured indicate that these estimates were unduly pessimistic (see below).

I tabulate the optimised readout parameters and simulated shelving fidelity-

Laser	Intensity (I_0)	Pulse length (μs)
393 nm σ^+	0.0141	11.3
850 nm π	2.89	1.9
850 nm σ^+ ($ 5, +4\rangle$)	6.15	4.4
850 nm σ^+ ($ 4, +4\rangle$)	4.72	4.4
e_{bright}	4.24×10^{-4}	
e_{dark}	1.22×10^{-4}	
$e_T := \frac{1}{2}(e_{\text{bright}} + e_{\text{dark}})$	2.73×10^{-4}	

Table 5.1: Optimised shelving parameters and simulated errors using the “stretch-qubit” bright state $|3, +3\rangle$. Optimisation was performed assuming polarization impurities, however quoted errors assume perfect polarization purity. The 850 nm σ^+ lasers are labelled by the $3D_{\frac{3}{2}}$ state they repump.

Laser	Intensity (I_0)	Pulse length (μs)
393 nm σ^+	0.0146	12.5
850 nm π	2.08	2.25
850 nm σ^+ ($ 5, +4\rangle$)	4.49	5.5
850 nm σ^+ ($ 4, +4\rangle$)	3.53	5.5
e_{bright}	2.07×10^{-4}	
e_{dark}	9.26×10^{-5}	
$e_T := \frac{1}{2}(e_{\text{bright}} + e_{\text{dark}})$	1.50×10^{-4}	

Table 5.2: Optimised shelving parameters and simulated errors using the bright state $|3, +1\rangle$ used to read out the clock-qubit. Optimisation was performed assuming polarization impurities, however quoted errors assume perfect polarization purity. The 850 nm σ^+ lasers are labelled by the $3D_{\frac{3}{2}}$ state they repump.

ties for these choices of bright state in tables 5.1 and 5.2. Here, we define e_{dark} to be the probability of the dark state remaining unshelved and e_{bright} to be the probability that the bright state is shelved. We define the average shelving error to be $e_T := \frac{1}{2}(e_{\text{bright}} + e_{\text{dark}})$.

On analysing the results of our SPAM measurements (section 5.3.2), it became clear that our polarization estimates were unduly pessimistic of what was achieved in the lab. In future experiments we will thus use parameters chosen on the basis of optimisations assuming pure polarizations. The results of such optimisation are given in tables 5.3 and 5.4. We see that the predicted error is reduced by 20 % for the clock qubit. We note, however, that these parameters

Laser	Intensity (I_0)	Pulse length (μs)
393 nm σ^+	0.0138	11.1
850 nm π	4.06	2.3
850 nm σ^+ ($ 5, +4\rangle$)	12.44	7.7
850 nm σ^+ ($ 4, +4\rangle$)	11.29	7.7
e_{bright}	4.04×10^{-4}	
e_{dark}	8.69×10^{-5}	
$e_T := \frac{1}{2} (e_{\text{bright}} + e_{\text{dark}})$	2.46×10^{-4}	

Table 5.3: Optimised shelving parameters and simulated errors using the “stretch-qubit” bright state $|3, +3\rangle$. Optimisation was performed assuming pure polarizations. The 850 nm σ^+ lasers are labelled by the $3D_{\frac{3}{2}}$ state they repump.

Laser	Intensity (I_0)	Pulse length (μs)
393 nm σ^+	0.0138	12.3
850 nm π	3.20	2.4
850 nm σ^+ ($ 5, +4\rangle$)	6.56	9.6
850 nm σ^+ ($ 4, +4\rangle$)	7.36	9.6
e_{bright}	1.91×10^{-4}	
e_{dark}	5.50×10^{-5}	
$e_T := \frac{1}{2} (e_{\text{bright}} + e_{\text{dark}})$	1.23×10^{-4}	

Table 5.4: Optimised shelving parameters and simulated errors using the bright state $|3, +1\rangle$ used to read out the clock-qubit. Optimisation was performed assuming pure polarizations. The 850 nm σ^+ lasers are labelled by the $3D_{\frac{3}{2}}$ state they repump.

were not used during any experiments reported in this thesis.

It is interesting to consider how likely the ion is to be shelved from the various states in the ground level. This is shown in table 5.5, assuming the parameters of table 5.2. We see that the shelving process is essentially manifold selective, although there are a few states in the $F = 4$ manifold which remain unshelved with $> 95\%$ probability. We also note that the likelihood of being shelved does not decrease monotonically with decreasing M . This is due to the large number of transitions between the S and P levels. For example, $|4, -4\rangle$ state is more than twice as likely to be shelved than its neighbouring states. This is because the 393 nm laser is only ~ 200 kHz detuned from the $|S_{\frac{1}{2}}, 4, -4\rangle \Leftrightarrow |P_{\frac{3}{2}}, 4, -3\rangle$ transition. Thus, this state is extremely likely to be

$F = 3$		$F = 4$	
State	P_{shelve}	State	P_{shelve}
$ 4, +4\rangle$	0.99991	$ 3, +3\rangle$	5.02×10^{-4}
$ 4, +3\rangle$	2.74×10^{-1}	$ 3, +2\rangle$	2.79×10^{-4}
$ 4, +2\rangle$	2.31×10^{-1}	$ 3, +1\rangle$	2.07×10^{-4}
$ 4, +1\rangle$	1.26×10^{-1}	$ 3, 0\rangle$	1.73×10^{-4}
$ 4, 0\rangle$	4.23×10^{-2}	$ 3, -1\rangle$	1.76×10^{-4}
$ 4, -1\rangle$	2.20×10^{-2}	$ 3, -2\rangle$	1.89×10^{-4}
$ 4, -2\rangle$	1.98×10^{-2}	$ 3, -3\rangle$	2.01×10^{-4}
$ 4, -3\rangle$	3.20×10^{-2}		
$ 4, -4\rangle$	7.17×10^{-2}		

Table 5.5: Probability of shelving different states in the ground-level of $^{43}\text{Ca}^+$ at 146 G using the optimised parameters from table 5.2.

excited to the P-level, from which it has a 5 % probability of being shelved, otherwise it decays a state with $|4, M > -4\rangle$.




Figure 5.8: Shelving error as a function of the number of repeats of the 393 nm-850 nm-850 nm cycle (n). The readout parameters are optimised at each point, assuming the $|3, +1\rangle$ bright state used for clock-qubit readout (so the point at $n = 8$ corresponds to table 5.2). The dotted line shows the asymptote at $e_T = 1.5 \times 10^{-4}$. We see that the fidelity increases with n until $n \gtrsim 6$, after which there is little improvement. For $n \gtrsim 10$ the lengths of the individual pulses start to become small, making the sequence difficult to implement.

For large-scale quantum computation, it is beneficial to perform all operations as quickly as possible. It is thus interesting to investigate how the shelving fidelity changes with the total time to shelve the ion. This is plotted in figure 5.9.

We see that the shelving fidelity decreases rapidly for times less than $30\ \mu\text{s}$, after which there is little gain. Compared with typical fluorescence detection times of $100\ \mu\text{s}$, this is not expected to be a limitation. For times longer than $\sim 200\ \mu\text{s}$ the shelving fidelity begins to decrease due to decay of the shelf level.




Figure 5.9: Shelving error as a function of total time taken to shelve. We optimise the shelving parameters at each point, assuming the $|3, +1\rangle$ bright state used for qubit readout. During the optimisation, we constrain the total shelving time but vary how that time is divided between the different laser pulses. The red dot corresponds to the conditions of table 5.2. We see that the fidelity decreases rapidly for times less than $\sim 30\ \mu\text{s}$. This is because the rapid shelving requires increased 393 nm intensity, leading to off-resonant shelving of the bright state. For times longer than $\sim 200\ \mu\text{s}$ the fidelity begins to decrease due to decay of the shelf level.

5.3.1.2 Beam alignment

The 393 nm σ^+ laser beam propagates along the same path as the 397 nm σ^+ , the set up of which was described in section 5.2. The 850 nm σ^+ was crudely aligned by overlapping it with the this beam path, along which it counter-propagates. More careful alignment was then performed using 850 nm and 854 nm repumpers to cool a $^{43}\text{Ca}^+$ ion at 146 G. When repumping with an 850 nm laser that is purely σ^+ the ion will be pumped into the $3D_{3/2}|M = +5\rangle$ state and fluorescence will cease. There are three transitions out of this state that will be driven by any polarization impurities present in the 850 nm beam:

the π transition to $3P_{3/2}|5, +5\rangle$ and the two σ^- transitions to $3P_{3/2}|5, +4\rangle$ and $3P_{3/2}|4, +4\rangle$.

To eliminate the beam's σ^- component, we tuned the 850 nm laser to resonance with one of the two σ^- transitions and adjust the wave plate's angle and tilt to minimise fluorescence. We then minimise the beam's π -component by using a pair of mirrors to change the 850 nm σ^+ beam angle, while monitoring the readout fidelity to ensure that the beam is always overlapped with the ion. We alternate walking the beam a small distance with checking for fluorescence on the π -transition.

The 850 nm π beam path is the least critical beam used during readout. The beam path was constructed so that the beam is approximately perpendicular to the magnetic field. We set the waveplate angle by temporarily placing a PBS in the beam path before the trap. The waveplate angle was then adjusted until the beam was polarized along the magnetic field axis.

5.3.1.3 Calibrating laser intensities

We calibrate our laser intensities by measuring optical pumping rates from different states and fitting the results to rate equations-based models. In the case of the 393 nm laser, we begin by preparing the ion in either of the $|4, +4\rangle$ or $|3, +3\rangle$ states in the ground level. We next apply the 393 nm laser for a varied amount of time to pump the ion into the $3D_{5/2}$ level before measuring the shelved population. The 393 nm laser is tuned to resonance with the transition from the $|4, +4\rangle$ state for both choices of initial state.

For a cold ion driven by a monochromatic source we expect to obtain the same intensity from both experiments. However, if the ion is hot then Doppler broadening reduces the pumping rate from $|4, +4\rangle$ and we correspondingly calculate a lower intensity. Since the transition from the $|3, +3\rangle$ state is detuned by 3 GHz, Doppler broadening does not have such a significant effect so the mea-

sured pumping rate remains largely unchanged. A similar effect is expected if the laser contains a significant amount of spectral impurity due, for example, to ASE.

We observe that if the cooling lasers' frequencies are set to produce significant heating then the pumping rate from $|4, +4\rangle$ can be reduced by a factor of two or more. Even with the cooling lasers set up properly, we typically find that the intensity measured using the $|3, +3\rangle$ state is about 15 % higher than predicted from the pumping rate from $|4, +4\rangle$. This is expected to lead to a small degradation in shelving fidelity.

To calibrate the 850 nm π beam intensity, we measure the pumping rate from the $3D_{3/2}|5, +5\rangle$ state. We prepare this state from $4S_{1/2}|4, +4\rangle$ by applying the 393 nm, 397 nm σ^+ , 850 nm σ^+ and 854 nm lasers together with the 397 nm EOM off. After preparing this state, we apply the 850 nm π laser for a varied amount of time. This transfers population from this state to the ground level $|4, +4\rangle$ state, from where it is shelved to the $3D_{5/2}$ level by applying a 393 nm pulse. Shelved population is then detected as usual.

For the two 850 nm σ^+ lasers, we prepare a mixture of states in the $3D_{3/2}$ level starting, as usual, from the $4S_{1/2}|4, +4\rangle$ state and applying the 393 nm, 397 nm σ^+ and 854 nm lasers with the 397 nm EOM off. This experiment relies on the spectroscopic isolation of the transitions from each state in this level. The contrast of the data is limited by the populations in the two $|M = 4\rangle$ states. However, due to our choice of lasers, these states have significant populations, allowing a contrast of order 20 %. This is sufficient for determination of the laser intensities with an error of $\lesssim 10\%$, which is satisfactory for these experiments.

5.3.2 Results and analysis

We performed experiments to measure the state preparation and measurement (SPAM) fidelity using both the stretched- and clock-qubits. These experiments,




Figure 5.10: Experiments used to measure the state-preparation and readout fidelities for the stretched- and clock-qubits. Pulses shown in grey with dashed borders are used for one qubit state only, determining which state is characterised. All other pulses are common to both states. Each experiment begins with preparation of $|4, +4\rangle$ using microwave enhanced optical pumping, followed by microwave π -pulses to transfer the ion to one of the qubit states. In the case of the clock-qubit, we then use further π -pulses to map the clock-qubit onto the $\{|4, +4\rangle, |3, +1\rangle\}$ temporary readout qubit. Finally, we apply the shelving sequence before using the state-dependent fluorescence technique to measure the population in the $3D_{5/2}$ level.

illustrated in figure 5.10, begin with the preparation of $|4, +4\rangle$ by microwave-enhanced optical pumping, followed by microwave π -pulses to transfer the ion to one of the qubit states. In the case of the clock-qubit, we then use further π -pulses to map the qubit onto the $\{|4, +4\rangle, |3, +1\rangle\}$ temporary readout qubit. Finally, we apply the shelving sequence before using the time-resolved fluorescence technique [Bur10] to measure the population in the $3D_{5/2}$ level.

We characterised one qubit at a time, taking data in runs of 25,000 shots on one of the two qubit states. We alternated which state was used in each run to ensure that both states were characterised under the same conditions, taking a total of $\sim 150,000$ shots on each state. The experiment was synchronised with the mains cycle to minimise shot to shot magnetic field fluctuations and the magnetic field servo (section 5.4) was run every 5,000 shots to remove slow drifts in the magnetic field.

The results of these experiments are shown in tables 5.6 and 5.7 along with a list of the estimated contributions due to different sources of error.

The quoted shelving error was calculated using rate equations-based simulations (section 5.3.1.1), taking into account the discrepancy between the 393 nm

	$ 4, +4\rangle$	$ 3, +3\rangle$	Average
Measured Error	$3.6(5) \times 10^{-4}$	$9.4(9) \times 10^{-4}$	$6.5(7) \times 10^{-4}$
Pumping to $ 4, +4\rangle$	$\lesssim 1 \times 10^{-4}$	$\lesssim 1 \times 10^{-4}$	$\lesssim 1 \times 10^{-4}$
Transfer pulses		3.8×10^{-4}	1.9×10^{-4}
Shelving	1.0×10^{-4}	5.4×10^{-4}	3.2×10^{-4}
Fluorescence detection	2.9×10^{-4}	0.2×10^{-4}	1.5×10^{-4}

Table 5.6: Measured SPAM errors and estimated error contributions for the stretched-qubit.

	$ 4, 0\rangle$	$ 3, +1\rangle$	Average
Measured Error	$4.1(5) \times 10^{-4}$	$9.5(8) \times 10^{-4}$	$6.8(6) \times 10^{-4}$
Pumping to $ 4, +4\rangle$	$\lesssim 1 \times 10^{-4}$	$\lesssim 1 \times 10^{-4}$	$\lesssim 1 \times 10^{-4}$
Transfer pulses	0.4×10^{-4}	6.7×10^{-4}	3.6×10^{-4}
Shelving	0.8×10^{-4}	2.6×10^{-4}	1.7×10^{-4}
Fluorescence detection	2.9×10^{-4}	0.2×10^{-4}	1.5×10^{-4}

Table 5.7: Measured SPAM errors and estimated error contributions for the clock-qubit.

intensity we measured when pumping out of $|4, +4\rangle$ and when pumping out of $|3, +3\rangle$ or $|3, +1\rangle$. The error quoted for each state assumes the 393 nm intensity calculated from the pumping rate out of the relevant state. All other parameters are assumed to be set to their optimised values and all polarizations are assumed pure. The simulation thus represents a lower limit on the shelving error.

We calculate the expected error due to the state-dependent fluorescence measurement using Monte-Carlo simulations based on our achieved fluorescence and background signals [Bur10]. In these experiments, we used the “standard” high-fluorescence cooling parameters discussed in section 5.1 with the one modification: the 866 nm intensity was $800 I_0$ rather than $600 I_0$. These parameters were chosen to provide an optimal balance between the ion’s fluorescence and the background due to scattered laser light according to the simulations in [Bur10]. Using these parameters, we achieved a fluorescence signal of 60,000 per second including a background of 10,000 per second due to 397 nm scat-

ter. Given our collection efficiency of 0.3 %, this corresponds to a $4P_{1/2}$ population of approximately 13 %. The Monte-Carlo simulations predict that our time-resolved photon counting should contribute an error of 2.87×10^{-4} to the dark (shelved) state and 0.21×10^{-4} to the bright state (1.54×10^{-4} average error).

We estimate that our state-preparation contributes an error of $\lesssim 1 \times 10^{-4}$ on the basis of two pieces of evidence. Firstly, since the shelving process is essentially *state-selective*, one would expect any state preparation error to be detected by our experiment on the stretched state. However, the error we measure on this state is, within the statistical uncertainty, equal to the sum of the minimum predicted shelving and fluorescence detection errors. This suggests that preparation error of the $|4, +4\rangle$ state was $\lesssim 0.5 \times 10^{-4}$. Secondly, as a control experiment, we also took the stretched qubit measurement using standard optical pumping with the 397 nm EOM on. We expect this to increase the state preparation error by roughly an order of magnitude. However, no increase in the SPAM error was detected to within the measurement's 1×10^{-3} sensitivity.

Between each shot of the experiment, we apply an 854 nm laser to ensure that the ion is pumped back into the normal 866 nm/397 nm cooling cycle (“deshelving”). If this laser is not properly extinguished by its AOM then some 854 nm light will reach the ion during fluorescence detection, effectively reducing the $3D_{5/2}$ shelf lifetime. To check that this was not the case, we measured the $3D_{5/2}$ lifetime in a separate experiment by looking at quantum jumps [BDL⁺00]. This experiment showed that the D-state lifetime was unchanged at the 10 % level and thus that 854 nm leakage was not a significant issue. This conclusion is further supported by the fact that the $|4, +4\rangle$ measured SPAM error agrees well with theoretical prediction based on expected fluorescence and shelving errors.

After accounting for state-preparation, shelving and fluorescence detection errors, the only remaining source of error is expected to be the microwave trans-

Transition	π -time (μs)
A	1.4
B	3.0
C	1.1
Q	4.2

Table 5.8: Microwave π -times used during SPAM measurements.

fer pulses. We did not attempt to characterise these pulses separately to the main experiment and so do not have an independent estimate of their errors. However, since other sources of error have been accounted for with a reasonable degree of certainty, we attribute the remaining error to the transfer pulses. Comparing this error to that expected for the microwave π -times used (table 5.8), shown in figure 5.6, the net transfer pulse error of $\sim 3 \times 10^{-4}$ is comparable to that expected for field fluctuations of ~ 1 mG.

5.4 Magnetic field servo

We compensate long-term drifts in the magnetic field – usually caused when someone uses the building’s lift (field change of ~ 1 mG per floor that the lift moves) or turns on a 14 T magnet in a nearby lab (~ 20 mG in our lab) – by servoing the field to the ion. The servo routine uses Rabi spectroscopy on the stretched transition to determine the magnetic field, before adjusting the current in a trim coil to return it to 146.094 G. The routine is automated, allowing slow magnetic field drifts over the course of a long experiment to be eliminated.

The Rabi experiment runs as follows. We begin by optically pumping the ion to the state $|4, +4\rangle$ before applying a π -pulse on the stretched transition and measuring the ion’s probability of remaining in $|4, +4\rangle$. The pulse’s π -time is chosen to be $22.2 \mu\text{s}$, giving a 36 kHz FWHM. We perform this experiment twice, with the microwaves detuned by ± 18 kHz from the stretched transition’s frequency at 146.0942 G.

If the field is 146.094 G then pulses at both detunings result in half of the population being transferred to $|3, +3\rangle$. Any imbalance in the measured populations thus indicates a field error. The routine calculates this error and adjusts the coil current appropriately. With our parameters, we find that the servo sets the field to an accuracy of $\lesssim 1$ mG, which is sufficient to allow preparation of the clock qubit with an error below 10^{-4} (see section 5.2).

5.5 Ramsey spectroscopy on the clock transition

From the Breit-Rabi formula, one expects the frequency of the clock transition to become independent of magnetic field to first order at 146.094 2 G (section 2.3.3). To confirm this prediction and experimentally determine the qubit frequency at the field independent point, we performed Ramsey spectroscopy on the clock transition.

We prepare the ion in the state $|3, +1\rangle$ before applying a pair of $\pi/2$ -pulses, detuned from the clock transition by a frequency δ and separated by a Ramsey delay τ_R (figure 5.11), and reading out the ion's resulting state. As the detuning of the two $\pi/2$ -pulses is changed the populations of the qubit states oscillate, producing a characteristic fringe pattern from which the transition frequency may be determined (figure 5.12).




Figure 5.11: Ramsey experiment on the clock-qubit. The ion is first prepared in the state $|3, +1\rangle$ before applying a pair of $\pi/2$ -pulses, detuned from the clock transition by a frequency δ and separated by a Ramsey delay τ_R , and reading out the ion's resulting state.

When the detuning is large compared with the $\pi/2$ -pulses' bandwidth⁸ the amplitude of the fringes decreases. This allows the centre of the fringe pattern to be located, providing an absolute measurement of the transition frequency

⁸The reciprocal of their duration.




Figure 5.12: Ramsey fringes on the clock transition. In both cases the duration of the $\pi/2$ -pulses was $478.62 \mu\text{s}$. (a) A wide detuning scan with $\tau_R = 1.051$ ms. Here, the detuning span is greater than the $\pi/2$ -pulse bandwidth, allowing the central fringe to be readily identified. (b) A narrow scan around the central fringe using $\tau_R = 24.93$ ms, showing a fringe offset of $8.4(1)$ Hz. Here, the detuning span is narrow in comparison with the $\pi/2$ -pulse bandwidth so it is not obvious from the pattern which is the central fringe.

(figure 5.12a). However, to determine the frequency of the centre of the fringe pattern with a high degree of accuracy, small detuning spans and long Ramsey delays must be used. In this case all fringes have approximately the same amplitude, making identification of the centre of the fringe pattern slightly more complicated (figure 5.12b).

In practice, one typically begins with a scan with a short delay and large detuning span so that the central fringe may be located unambiguously. One then takes a series of scans with decreasing detunings and longer Ramsey delays. So long as the detuning is decreased by sufficiently small increments between scans, the centre of the fringe pattern is always known. One may also confirm that a particular fringe is the central fringe by altering the Ramsey delay, which has the effect of moving all fringes except the central one. The fractional change in the Ramsey delay should be close to an irrational number to minimise the potential for aliasing effects.

5.5.1 Qubit frequency as a function of B-field

To demonstrate the qubit's field independence, we performed a series of Ramsey experiments at different static magnetic fields near 146.0942 G. We set the magnetic field for each scan using the automated B-field servo routine (section 5.4), which alters the current in the $\hat{z}$ trim coil until the stretched transition has a specified frequency. This technique uses the stretched transition's well known field dependence (-2.36 kHz/mG at 146 G) as a measure of the magnetic field. We estimate that this allowed us to set the average field with an absolute accuracy of $\lesssim 1$ mG, although we see fluctuations $\lesssim 5$ mG of at 50 Hz and other frequencies.




Figure 5.13: Clock transition frequency offset from 3 199 941 070.566 Hz against magnetic field offset from 146.0942 G. The black curve is a fit to the data based on the Breit-Rabi formula, using a frequency offset as the sole fitted parameter. The uncertainty in each frequency measurement is estimated to be ~ 0.1 Hz based on the fits, while the B-field is known to $\lesssim 5$ mG due to noise at 50 Hz and other frequencies.

The results of these experiments are shown in figure 5.13. As noted in the caption, the black curve is a fit to the data based on the Breit-Rabi formula, using a frequency offset as the sole fitted parameter. This offset is used to account for systematic shifts in the qubit frequency as well as uncertainty in the nuclear g -factor for $^{43}\text{Ca}^+$, as discussed in the following sections. Apart from

this frequency offset, the data are consistent with the theory. Based on the fit, and not correcting for systematic errors, we measure a qubit frequency of 3.199 941 070 566 (3) GHz.

5.5.2 Systematic shifts

While it is not our present purpose to obtain the most accurate possible measurement of the clock-transition's frequency, it is interesting to consider briefly the systematic errors present in the previous measurement. In the following sections we will discuss those effects which we believed to be most significant (or interesting). A more thorough discussion of sources of error in this kind of experiment may be found in the extensive literature on trapped ion metrology (see, for example, [BMB⁺98a]).

In the following discussions we will use a sign convention whereby frequency shifts represent changes in the qubit frequency due to a particular effect, with a positive shift indicating that the qubit frequency has increased. The shifts must thus be subtracted from the measured qubit frequency to arrive at the unperturbed value.

5.5.3 Rubidium reference calibration

We generated our $\pi/2$ -pulses using a commercial microwave synth phase locked to a rubidium reference (section 4.7.1). To correct errors in this time standard we calibrated it to TAI by comparing the phase of its 10 MHz output to a GPS disciplined oscillator for a period of 24 hours immediately after taking the data. This calibration showed that the rubidium reference was fast compared with TAI with a fractional error⁹ of $+4.13(1) \times 10^{-10}$. This corresponds to a shift of -1.322 Hz on the 3.199 941 GHz qubit transition.

⁹The error quoted here represents the GPS reference's specified Allan Variance of 10^{-12} over 24 hours.

5.5.4 Quadratic Zeeman effect

Fluctuations in the magnetic field lead to shifts in the qubit's frequency given (to second order) by

$$f(B_0 + \delta B) = f(B_0) + \delta B \frac{df}{dB}(B_0) + \frac{1}{2}(\delta B)^2 \frac{d^2f}{dB^2}(B_0) \quad (5.15)$$

where B_0 is the average magnetic field and δB represents a zero-mean fluctuation. Here, we have assumed that the fluctuations are far from resonance with any transition from either qubit state and so may be treated as a quasi-static modulation of the magnetic field. This is not so for the magnetic field associated with the trapping RF, whose frequency is close to the ground-level Zeeman splitting. We will consider this case in section 5.5.6.

When we perform a Ramsey experiment, we measure the average transition frequency during the delay interval, given by

$$\langle f \rangle = f(B_0) + \frac{1}{2} \langle \delta B^2 \rangle \frac{d^2f}{dB^2}(B_0) \quad (5.16)$$

For the 146 G clock transition, this corresponds to a shift of $1.2 \text{ mHz/mG}_{\text{RMS}}^2$.

The dominant source of magnetic field noise in our experiment is the 50 Hz mains electricity, which produces a field of approximately 3 mG at the ion's position¹⁰¹¹ Taking 5 mG as a somewhat pessimistic upper bound on the total field noise, we expect the quadratic Zeeman effect to lead to a frequency shift of $\lesssim 30 \text{ mHz}$, giving a fractional error of $\lesssim 1 \times 10^{-11}$.

5.5.4.1 Second-order Doppler shift

As the ion moves the frequency of the clock-transition is altered due to the Doppler effect. Since it is confined within the trap the ion has no mean velocity, so this effect enters only at second order. The resulting fractional shift of the

¹⁰Measured using Ramsey experiments on the stretched transition.

¹¹In this section we will ignore the magnetic field due the trap RF, which we consider separately in section 5.5.6.

transition's frequency is given by

$$\frac{\Delta\nu_{D2}}{\nu} = -\frac{1}{2} \frac{\langle V^2 \rangle}{c^2} \quad (5.17)$$

The ion's motion consists of two components: secular motion and micromotion. The former is a random motion that may be thought of as the ion being at a finite temperature, while the latter is the result of driven motion due to the trapping RF. Rather than attempt to provide a detailed calculation based on estimates of the ion's temperature, we note that the trap depth is only $\lesssim 100$ meV, leading to an RMS escape velocity of 670 m/s. This corresponds to a fractional frequency shift of -2.5×10^{-12} . In practice, the ion will move much more slowly than escape velocity. Since the ion is laser cooled and in a well-compensated trap with only a small degree of residual motion micromotion, we expect the second-order Doppler shift to be extremely small we conclude that Doppler shifts will be negligible in comparison with the other systematics in this experiment.

5.5.5 A.C. Stark Shift

The A.C. Stark shift is a second-order perturbation arising from mixing between the qubit states and states of opposite parity, to an excellent approximation those in the P-levels, due to an external electric field. The energy shift of a state $|s\rangle$ in the ground level due to an electric field $\mathbf{E}$ is given by

$$\delta E_s = \sum_i \frac{|\langle s | \mathbf{d} \cdot \mathbf{E} | p_i \rangle|^2}{E_s - E_{p_i}} \quad (5.18)$$

where $\{|p_i\rangle\}$ are the states in the 3P levels and E_s and E_{p_i} are the state's (unshifted) energies.

We detect electric fields that are either static or coherent with the trapping RF by measuring the ion's micromotion [BMB⁺98b]. We find that the measured field is dominated by the ~ 30 V/m of trapping RF present at the trap's centre (see section 4.1).

This technique is not sensitive to fluctuating electric fields that are incoherent with the trap RF. One can place an upper bound on the magnitude of such fields by looking for micromotion sidebands on optical transitions, which would be present if an electric field displaced the ion from the trap's centre. The fact that we do not see sidebands when we scan our lasers implies that such fields have a magnitude of $\lesssim 100 \text{ V/m}$.

To estimate the magnitude the Stark shift, we approximate the ion's electric dipole moment by

$$d \sim a_0 e \quad (5.19)$$

where a_0 is the Bohr radius and e is the electronic charge. The resulting shift is of order $10^{-7} \text{ Hz}/(\text{V/m})^2$. For a 100 V/m field, this produces a qubit frequency shift of approximately 1 mHz , corresponding to a fractional error of $\sim 1 \times 10^{-12}$. This is negligible in comparison with other considered effects.

5.5.6 AC Zeeman shift due to the trap RF

The RF used to produce the radial trapping potential creates both electric and magnetic fields. Since the ion is confined to a small region around the RF electric field null, the residual electric field it "sees" is small, leading to a correspondingly small Stark shift (see previous section). The magnetic field, on the other hand, may be non-negligible even at the trap's centre, leading to a significant A.C. Zeeman shift.

At 146 G the Zeeman splittings between states in the ground level are roughly 50 MHz , which is comparable to the 38.2 MHz RF frequency used during these experiments. Because of this, there is a significant AC Zeeman shift due to off-resonant excitation of σ -polarized RF transitions within each ground-level hyperfine manifold (figure 5.14). Thus it is not sufficient to consider the effect of the trap RF as a simple modulation of the magnetic field strength as we did in section 5.5.4. Instead, we calculate the A.C. Zeeman shifts resulting from




Figure 5.14: Ground-level magnetic dipole transitions from the qubit states (blue). There are 5 microwave transitions at ~ 3 GHz as well as four σ -polarized RF transitions at ~ 50 MHz. The 38.2 MHz trap RF is red-detuned from all of these transitions, so the A.C. Zeeman shifts tend to increase the qubit splitting. The only exception to this are the $|4, 0\rangle \leftrightarrow |4, -1\rangle$ and $|3, +1\rangle \leftrightarrow |3, 0\rangle$ transitions, which lead to a decrease in qubit splitting.

all possible magnetic dipole transitions between the qubit states and other states in the ground level (table 5.9)¹².

We note that the total shift due to the π -transitions is exactly equivalent to the quadratic Zeeman shift discussed previously¹³. This may be seen by noting that when one considers the quadratic Zeeman shift as a second order perturbation of the qubit's energy, it is the matrix elements in table 5.9 that determine the transition's second-order field sensitivity.

To determine the A.C. Zeeman shift on our qubit frequency measurement we performed a series of Ramsey experiments at 146.094 G using different RF powers. The results of these experiments are shown in figure 5.15. We measured the RF power for each experiment using a directional coupler placed between

¹²See section B.5 for a discussion of this calculation.

¹³Note that the factor of two difference between the number quoted here and that of section 5.5.4 arises from the fact that in section 5.5.4 we considered the field's RMS, while here we consider its amplitude.

Transition	Matrix element (μ_B)	RF Detuning (MHz)	A.C. Zeeman shift (Hz/ μT^2)
$ 4, 0\rangle \Leftrightarrow 3, -1\rangle$	+0.528	-3262.7	+0.0083
$ 4, 0\rangle \Leftrightarrow 3, 0\rangle$	-0.993	-3212.8	+0.0297
$ 4, 0\rangle \Leftrightarrow 3, +1\rangle$	+0.702	-3161.2	+0.0302
$ 4, +1\rangle \Leftrightarrow 3, +1\rangle$	-0.993	-3109.6	+0.0307
$ 4, +2\rangle \Leftrightarrow 3, +1\rangle$	+0.887	-3056.2	+0.0249
$ 4, -1\rangle \Leftrightarrow 4, 0\rangle$	+0.773	-11.3	-2.8050
$ 4, 0\rangle \Leftrightarrow 4, +1\rangle$	-0.798	-12.9	+2.6681
$ 3, +1\rangle \Leftrightarrow 3, 0\rangle$	-0.618	-12.8	-1.6112
$ 3, +2\rangle \Leftrightarrow 3, +1\rangle$	+0.584	-14.6	+1.2872
Total π -shift			+0.0604
Total σ -shift			-0.3976

Table 5.9: A.C. Zeeman shifts on the clock transition due to an R.F. magnetic field at 38.2 MHz.




Figure 5.15: Qubit A.C. Zeeman shift as a function of trap RF power. A fit linear fit to the data gives a total shift of -0.0127 Hz/mW corresponding to -4.98 Hz at the 25.93 dBm power level used in the data of figure 5.13 (shown in red).




Figure 5.16: (a) Simplified magnetostatic model of the trap. We model the three trap electrodes and ground-plane as thin wires separated by 90.5 μm (the distance between the centres of the RF electrodes and the trap’s symmetry axis). We assume that a current I flows through each RF electrode and returns evenly distributed between the ground plane and centre electrodes, which each carry a current of $-\frac{2}{3}I$. (b) Simulated angle between the RF magnetic field and the $\hat{z}$ axis as a function of displacement from the trap’s centre along the $\hat{z}$ axis at a height of 75 μm above the trap’s surface. We see a rapid change in field angle as we approach the point 16 μm from the trap’s centre where the field along the $\hat{z}$ direction changes sign. By 5.5 μm away from the trap’s centre, the field has reached an angle of 50°.

the RF amplifier and the resonator. Since the shift depends on the square of the RF current, one expects it to scale linearly with RF power, in good agreement with the measurements. The measured shift is -0.0127 Hz/mW , corresponding to a shift of -4.98 Hz at the 400 mW power level used in the data of figure 5.13 (shown in red).

It is interesting to note that the sign of this shift suggests that the RF magnetic field contains a large σ -polarized component. This was initially surprising since we expected the σ component to be extremely small due the trap’s symmetry. To investigate this effect further, we repeated this measurement on the low-field $|4, 0\rangle \leftrightarrow |3, 0\rangle$ clock-transition at a field of 1 G. For this transition, both σ^- and π -polarized fields increase the transition frequency. By comparing the results of the two measurements, we estimate that the trap’s magnetic field has a magnitude of order 10 μT at an angle of around 50° to the static magnetic field.

To understand the origin of this field angle, we consider a simple magneto-static model of the trap, as shown in figure 5.16. From this model, we can see that the field angle reaches 50° to the $\hat{z}$ axis at a distance of roughly $5.5\ \mu\text{m}$ from the trap's centre. To investigate this effect experimentally, we used the trap's DC electrodes to create electric fields which displace the ion known distances from the electric field null, measuring the qubit frequency at each point. The resulting data are shown in figure 5.17. We see that the electric field null is displaced by $3.5\ \mu\text{m}$ along the $\hat{z}$ axis from the point of symmetry of the magnetic field. It thus seems plausible that the large field angle we measure is the result of a small asymmetry – for example, slightly different RF voltages being coupled onto the two RF electrodes – moving the electric field null away from the geometric centre of the trap.

We note that, due to the large frequency dependence of the σ AC Zeeman shifts, once the RF magnetic field angle has been measured it is possible to choose an RF frequency for which the σ and π shifts cancel exactly. For our trap, this would be roughly 28 MHz, which would be experimentally feasible. This “magic frequency” might be useful for some trapped ion metrology experiments where AC Zeeman shifts can be a leading cause of error.

5.5.7 Conclusions

Without compensating for systematic errors, the measured transition frequency was 3.199 941 070 566 (3) GHz (section 5.5.1). We believe that the dominant systematic errors were the rubidium reference calibration and the RF A.C. Zeeman shift, producing a combined offset of -6.30 Hz. The main uncertainty in the measurement arises from the fact that the RF power was only controlled to $\lesssim 1\%$, producing an uncertainty of 50 mHz. Taking these corrections into account, we arrive at a measured transition frequency of 3.199 941 076 866 (50) GHz.

It is interesting to consider how this measurement relates to more fundamen-




Figure 5.17: A.C. Zeeman shifts as a function of ion position in the $\hat{y}$ and $\hat{z}$ directions (no shift was found in the $\hat{x}$ direction at the 0.1 Hz level). Points along the $\hat{y}$ axis were taken with an RF power of 400 mW, while points along the $\hat{x}$ axis were taken at three different RF powers. The light blue curve represents a quadratic fit to the data, while the black and red curves are the result of rescaling the blue curve in proportion to the measured RF power. The point of symmetry in the $\hat{z}$ direction lies $3.5 \mu\text{m}$ away from the electric field null (as determined by separate micromotion compensation experiments).

tal atomic constants. We recall that the energies of the two qubit states are given by the Breit-Rabi formula

$$E_{\pm}(B, M) = -\frac{E_{\text{hfs}}}{2(2I+1)} - g_I \mu_N B M \pm \frac{E_{\text{hfs}}}{2} \sqrt{1 + \frac{4\chi B M}{2I+1} + \chi^2 B^2} \quad (5.20)$$

where $\chi := (g_I \mu_N + g_J \mu_B) / E_{\text{hfs}}$ and E_{hfs} is the zero-field hyperfine structure splitting.

The transition frequency thus depends on three properties of the ground-level: its zero-field hyperfine structure splitting and its Landé and nuclear g -factors. The hyperfine splitting has been measured for $^{43}\text{Ca}^+$ ions in a Paul trap using optical-microwave double resonance to be $3.225\,608\,286\,4(3)$ GHz [ABG⁺94]. The Landé g -factor has also measured by double resonance, this time on a cloud of $^{40}\text{Ca}^+$ ions in a Penning trap at 1.43 T [TPR⁺03], where it was found to be $g_J = 2.002\,256\,64(9)$. The uncertainty in the Landé g -factor produces a 2 mHz uncertainty in the qubit frequency, which is dwarfed by the

300 mHz uncertainty of the hyperfine splitting.

The author is aware of several measurements of the nuclear g -factor for calcium 43, which are shown in table 5.10. [Jef53] and [LSU73] used NMR on liquid CaBr_2 at 1 T and aqueous CaCl_2 , CaBr_2 and $\text{Ca}(\text{NO}_3)_2$ at 1.8 T respectively, while [Ols72] used double-resonance on an atomic vapour. All measurements give the g -factor for neutral ^{43}Ca (after correction for diamagnetic shifts in the case of the NMR measurements). There will be a slight correction for the $^{43}\text{Ca}^+$ ion due to different diamagnetic screening of the nucleus, however we expect this to be small. Ignoring this correction, we find that our current measurement is consistent with [Jef53] & [Ols72] but not with [LSU73]. The reason for this discrepancy is unclear.

Measured g_I	Predicted qubit frequency (GHz)	Measurement
$-1.315\,2\,(2)$	$3.199\,941\,072\,(6)$	[Jef53], liquid NMR
$-1.315\,37\,(60)$	$3.199\,941\,078\,(19)$	[Ols72], atomic vapour
$-1.315\,645\,(7)$	$3.199\,941\,086\,3\,(2)$	[LSU73], liquid NMR
$-1.315\,348\,(9)$	$3.199\,941\,076\,866\,(300)$	This work

Table 5.10: Measurements of the nuclear g -factor of ^{43}Ca and the resulting qubit frequency. Predictions of the qubit frequency assume the values of the hyperfine structure splitting and Landé g -factors given in [ABG⁺94] and [TPR⁺03], with the stated uncertainty only representing the contribution due the the nuclear g -factor. The uncertainty on the nuclear g -factor calculated from the measurements made in this work is dominated by the uncertainty in the hyperfine structure splitting.

5.6 Coherence time

An important benchmark for any qubit is its coherence time – the duration for which quantum information may be stored for before being lost to noise. This is typically measured using Ramsey experiments similar to those discussed in the previous section. In this case, however, the two $\pi/2$ -pulses are both resonant with the qubit transition and it is their relative phase that is scanned to give fringes. Ideally the fringes should have unit amplitude, indicating that the

second $\pi/2$ pulse may add to or subtract from the effect of the first with perfect coherence.

In practice, however, noise always intervenes during the Ramsey delay, reducing the phase coherence between the qubit and microwaves and, correspondingly, decreasing the fringe amplitude. In the regime where the Ramsey delay is long compared with the noise's correlation time, the decay will be exponential¹⁴ so that

$$A(\tau_R) = e^{-\tau_R/T_2} \quad (5.21)$$

where A is the fringe amplitude, τ_R the Ramsey delay and T_2 the coherence time.

As an additional probe of the qubit's coherence, we also performed a CPMG-type experiment [CP54] [Lev86]. This differed from a normal Ramsey experiment by containing a series of π -pulses in the Ramsey delay. These pulses act as an extended spin-echo, reducing the effect of noise. By choosing the pulses' phase to differ by 90° from that of the initial $\pi/2$ -pulse, the sequence is made extremely insensitive to fluctuations in the π -pulses' Rabi frequency or duration. So long as the delay between π -pulses can be made short compared with the noise correlation time, one expects this type of sequence to be extremely robust to external perturbations.

We note that the CPMG sequence's robustness to Rabi frequency/pulse length errors is only obtained because one knows the qubit's state – one can picture the first $\pi/2$ -pulse as preparing the qubit on the $\hat{y}$ axis of the Bloch sphere. The π -pulses then correspond to 180° rotations about this axis. In the absence of noise, the qubit will stay on the $\hat{y}$ axis and the π -pulses do nothing, so one is entirely insensitive to errors in their duration/Rabi frequency. Clearly, this technique does not work if the qubit is in an unknown state on the Bloch sphere and so cannot be applied to general purpose quantum memories. There are, of course, other techniques which are useful for a general purpose memory, how-

¹⁴See section B.6 for a discussion of this point.

ever they sacrifice some aspects of the CPMG sequence's robustness [Lev86] [UBV⁺09].

5.6.1 Experimental considerations

These experiments proceeded much along the lines of the qubit frequency measurement discussed previously apart from a few special considerations due to the long (up to 16 s) Ramsey delays used.

Since our AOMs provide only a finite degree of extinction, a small amount of laser light may reach the ion, even with the AOMs nominally switched off. Light at 397 nm or 393 nm is capable of driving transitions out of the qubit states, which would decrease the measured fringe contrast. To eliminate this effect, we used mechanical shutters to block the beams during the longer delays.

A similar problem occurs due to small levels of microwave leakage through the RF switches. To eliminate leakage from the state-preparation synths, we ensured that the synth responsible for driving the $|4, +4\rangle$ to $|3, +3\rangle$ transition was connected to the quadrupler during the Ramsey delay (see section 4.7.3). This synth was chosen because it does not drive transitions out of the qubit state. Due to the non-linearity of the frequency multiplication process, extinction from all other state-preparation synths is then extremely high.

The Ramsey experiments used a pair of microwave synths to generate the $\pi/2$ -pulses. We used an RF switch to control which synth was connected to the trap and a further pair of switches in series to allow both synths to be connected or disconnected from the trap with high isolation. Despite the use of two switches¹⁵, we still observed ~ 1 Hz carrier Rabi-flopping during the delay interval. This was eliminated by using the synth's pulse modulation input, which is specified to provide an additional >60 dB of isolation with a 150 ns settling time.

¹⁵A third switch was not found to improve extinction, suggesting that the dominant path for microwaves to get to the ion had become a parasitic path such as pick-up on amplifier power cables.




Figure 5.18: 16 s second Ramsey (a) and CPMG (b) experiments, showing oscillations in the population of the state $|3, +1\rangle$ as the relative phase of the $\pi/2$ pulses is scanned.

We used a frequency octupled DDS (section 4.7.4) to drive both π - and $\pi/2$ -pulses during the CPMG experiment. The DDS power was set to produce a π -time of $24.4 \mu\text{s}$. Its amplitude was set to zero during inter-pulse delays to eliminate leakage effects.

For both experiments, we used the magnetic field servo to remove long-term drifts in the magnetic field. The servo was run once for each different phase point used in a run.

We interleaved the Ramsey experiments with short check experiments similar to those used in section 5.3 to measure the readout levels (typically 99 % for these experiments). The measured readout levels were used to normalise the fitted fringe contrasts.

5.6.2 Results and discussion

We performed Ramsey experiments with delays of up to 16 s and a single CPMG experiment with a 16 s delay containing 800 π -pulses (20 ms between pulses). The fringe contrast for each experiment was calculated by fitting a sinusoid to

the measured data, so that

$$P_1(\phi) = \frac{1}{2} \{1 + A \cos [\kappa \times (\phi - \phi_0)]\} + P_1^{(0)} \quad (5.22)$$

where $P_1(\phi)$ gives the probability of measuring the ion to be in the state $|1\rangle$ when the phase difference between the two $\pi/2$ -pulses is ϕ and A is the fringe contrast.

The skew factor κ is used to account for a phase offset that changes linearly over the course of the experiment [LOJ⁺05]. In [LOJ⁺05] this occurred because of magnetic field drifts. We servo the magnetic field between phase points, so we do not expect this to be an issue. However, in the Ramsey experiments the two $\pi/2$ -pulses were generated by different synths whose relative phase is expected to drift (see section 4.7.2). In the CPMG experiment this was not a problem, so κ is manually set to 1 in the fit.

The offset $P_1^{(0)}$ can account for differences in the fidelity of readout for the two qubit states. It was small in all the fits and did make a significant difference to the fitted fringe contrast.

Figure 5.18 shows the data and fits for the 16 s Ramsey and CPMG experiments. We fit an exponential decay to the fringe contrast as a function of Ramsey delay (5.19), which gives a coherence time¹⁶ of $T_2^* = 50(10)$ s.

As discussed in section 2.3.3, we expect the coherence time to be of order

$$T_2 \sim 1/(2\pi\delta f) \quad (5.23)$$

where δf characterises the qubit frequency fluctuations. A coherence time of 50 s corresponds to frequency fluctuations of order 3 mHz. This is roughly what one would expect from the analysis of section 5.5 given fluctuations in the trap RF power and magnetic field of order 0.1 % and 2 mG respectively. These estimates are consistent with typical values for our experiment.

¹⁶Notice that T_2^* indicates the coherence time averaged over an ensemble. In this case, the ensemble consists of measurements made on the same ion at different times.




Figure 5.19: Fringe contrast measured during Ramsey (blue) and CPMG (red) experiments. The black curve is an exponential fit to the measured Ramsey fringe contrast (as equation (5.21)), giving a coherence time of $T_2^* = 50(10)$ s. The fringe contrast for the CPMG experiment was $90(4)\%$ at a delay time of 16 s.

We are aware of two effects which could lead to the qubit coherence time being underestimated by these experiments. Firstly, the readout levels were found using short experiments without long delays between state preparation and measurement. However, during our Ramsey experiments the ion is left uncooled for up to 16 s. Over the course of this interval we expect it to undergo significant heating, leading to Doppler broadening of the 393 nm shelving transition and to a corresponding reduction in readout fidelity. Such heating-induced measurement errors have been found to have a significant impact on measured fringe contrasts in related experiments [Lan06].

This effect could have been characterised by adding a delay equal to the Ramsey delay to the readout level measurements, allowing the fitted fringe contrasts to be normalised appropriately. Alternatively, the intensities of the readout lasers could have been increased to accommodate the Doppler broadening. Neither of these techniques were applied to this data, however. This problem

highlights the fact that a large-scale ion trap quantum computer is likely to require sympathetic cooling using a different ion species [BDS⁺03].

The second effect stems from the use of separate frequency synthesisers to generate each Ramsey $\pi/2$ -pulse. The longer experiments took a period of several hours to complete, over which time we typically observe significant phase fluctuations between phase-locked microwave synths (section 4.7.2). Given the $\approx 45^\circ$ skew that was needed to fit the 16 s Ramsey experiment it is likely that this drift had a significant impact on the measured fringe contrast.

It is interesting to note that the CPMG experiment did not show as large an increase in contrast compared with the Ramsey experiment as one might expect. One interpretation of this fact would be that the bulk of the noise is fast compared with the 50 Hz π -pulse rate used during the CPMG sequence. A further possibility is that the measured loss of contrast is predominantly due to decreasing readout fidelity as the ion becomes hot.

5.7 Randomized benchmarking of a single qubit

To assess the usefulness of an implementation of quantum computing, one needs a measure of the accuracy with which quantum logic gates may be performed. The most complete characterisation of a quantum gate is provided by quantum process tomography (QPT) [CN97] [PCZ97], which affords one a full description of the gate's effect on an arbitrary input state.

Whilst QPT has been used successfully to gather useful information about quantum gates [RKS⁺06], it suffers from several limitations. Firstly, to fully characterise a quantum process one must take a large number of measurements: $16^n - 4^n$ for an n -qubit system. This exponential scaling limits the usefulness of QPT to systems composed of a small number of qubits.

Secondly, the fidelity of the gate we wish to characterise is often as good as or better than the fidelity with which we can perform the state preparation and

measurement operations required by QPT. This makes accurate estimation of the errors due to the gate alone extremely challenging.

Finally, for a gate to be useful for quantum computation, it is important that it should have a high fidelity whenever it is performed. In practice, however, the fidelity of a gate will often be sensitive to factors such as experimental duty cycle or the timing of the gate relative to the 50 Hz mains cycle. Thus, whilst QPT gives a full characterisation of a gate performed under a particular set of circumstances, it does not provide one with a guarantee of how the gate will perform when used in a realistic computation.

Randomized benchmarking offers an alternative to QPT which overcomes these limitations [EA05] [KLR⁺08] and has become the standard technique for characterising errors in experimental quantum computing (see [ECMG13] and references therein). Benchmarking consists of applying sequences of gates chosen randomly from a specified group and measuring the average decay in fidelity of the resulting state as a function of the sequence length.

The usefulness of randomized benchmarking rests on two claims. Firstly that the fidelity measured by benchmarking is a good approximation to the fidelity of the gates used averaged over possible input states. Secondly that the measured fidelity converges to the average over all random sequences, even when only a small number of different random sequences are used. Recent theoretical work has shown that both of these claims hold over a wide range of realistic error models [MGE11] [ECMG13].

In this section we describe randomized benchmarking experiments used to measure the fidelity of single-qubit manipulations in our set-up. We provide a detailed analysis of likely sources of experimental error, which appears to provide good agreement with the results obtained.

5.7.1 Randomized benchmarking protocol

We use a randomized benchmarking protocol consisting of

1. Choose N_G sequences of “computational gates”, each consisting of l $\pi/2$ -pulses drawn randomly from the set $\{e^{\pm i\sigma_c\pi/4} : c = \{x, y\}\}$, where $\sigma_{x,y}$ are the Pauli matrices.
2. Randomize each sequence by inserting a “randomization pulse”, drawn randomly from the set $\{e^{\pm i\sigma_r\pi/2} : r = \{0, x, y, z\}\}$, before each computational gate¹⁷. We will refer to the combination of a randomization pulse followed by a computational gate as a “randomized computational gate”. Repeat this randomization process N_R times for each sequence of computational gates to arrive at $N_G N_R$ sequences of randomized computational gates.
3. Calculate the outcome expected when each sequence is applied to a chosen qubit state (we used $|3, +1\rangle$ in our experiment). In the absence of experimental errors, the ion will be in an eigenstate of a Pauli operator. There are thus two choices of $\pi/2$ -pulse from the set $\{\pm x, \pm y\}$, which will bring the ion back into an eigenstate of σ_z and hence yield a deterministic result on measurement. Append a randomization pulse, followed by one of these two choices of $\pi/2$ -pulse (selected randomly, of course!) to each sequence of randomized computational gates.
4. Append a final randomization π -pulse to each sequence.
5. Repeat the above steps for N_l different values of l to obtain a total of $N_G N_R N_l$ different sequences.
6. Implement the sequences experimentally, repeating each N_e times and measuring the ion’s resulting state. The fidelity of a sequence is given

¹⁷Here, σ_0 denotes the identity transformation.

by $F = 1 - \frac{e_T}{N_e}$ where e_T is the total number of times that the measured outcome differed from the expected outcome.

7. Average the fidelities of the $N_G N_R$ sequences of length l to calculate the average fidelity for each length, F_l . Fit these average fidelities to the model $F_l = \frac{1}{2} + \frac{1}{2} (1 - 2e_{\text{SPAM}}) (1 - 2e_{\text{gate}})^l$ where e_{SPAM} is the error associated with state preparation and measurement and e_{gate} is the average error per *randomized computational gate*.

The above protocol is essentially that of [KLR⁺08] with one minor change: we chose N_G independent sequences of computational gates for each sequence length. In [KLR⁺08] they chose N_G sequences which were as long as the longest sequence they wished to implement. They then truncated each sequence to each of the N_l lengths used to form $N_G N_l$ sequences. This difference was purely practical and is not expected to impact on the results of the experiment.

5.7.2 Experimental implementation




Figure 5.20: Phase-agile single qubit microwave drive system. See text for details.

We performed single-qubit gates using the phase-agile microwave system shown in figure 5.20. Microwaves were generated by an octupled AD9910 DDS running at a base-band frequency of 400 MHz (section 4.7.4). This was clocked at 1 GHz from an Agilent E4426B ESG-AP series frequency synthesiser locked to the rubidium frequency standard. A 9 dB attenuator between the DDS and first

switch provided coarse power (Rabi frequency) tuning and helped to minimise VSWR.

We ran the DDS and octupler continuously and at a constant power level to attempt to minimise effects of experimental duty cycle. A series connected pair of high isolation switches¹⁸ switch its output between the trap and external 50 Ω loads. Two switches were used since the 60 dB of isolation provided by a single switch was not sufficient for these experiments. The switches were connected to the trap's right electrode via a circulator and the microwave diplexer (section 4.3.1).

The DDS's amplitude was set to produce a $\pi/2$ -time of 12.1 μs (roughly -10 dBm at the input to the diplexer). Calibration was performed using a sequence of 751 $\pi/2$ -pulses on the qubit transition, with the DDS amplitude adjusted to produce an equal probability of finding the ion in either qubit state. There was a dead-time of 14 μs between each $\pi/2$ -pulse due to limitations of our current experimental control hardware.

We followed [KLR⁺08] and [BWC⁺11] in several details of implementation: π -pulses were implemented as a pair of $\pi/2$ -pulses; identity pulses were implemented as a delay of the same length as other pulses, but with the microwaves turned off; z -rotations were implemented as an identity pulse accompanied by a change in the phase of the microwave drive field.

Unlike the clock-qubit, the microwave transitions used during SPAM are magnetic field sensitive to first order. To minimise errors due to field fluctuations associated with the 50 Hz mains electricity, we synchronised the start of our experiment with the mains cycle. However, while this ensures that preparation pulses always occur at the same phase of the mains cycle, the timing of the pulses used during readout depend on the length of the random sequence used.

¹⁸Mini-Circuits ZASWA-2-50DR+ high isolation absorptive switch. Specified performance at 3.5 GHz: 2.25 dB insertion loss; 59.87 dB off-state isolation; 1.25 : 1 input VSWR (equivalent to a return loss of 19.08 dB).

To ensure that all sequences had the same SPAM errors, we chose our sequence lengths to occupy integral numbers of mains cycles, keeping all SPAM pulses at the same phase of the mains cycle. We used sequences of length $l = \{10, 250, 500, 750, 1000, 1500, 2000\}$, with $N_G = 4$ different sequences at each length, each randomized $N_R = 8$ times. The 10-gate sequence occupied one cycle (20 ms), while the 2000 gate sequence occupied 9. For all sequences apart from the 10-gate sequence, we used mechanical shutters to block the 397 nm and 393 nm beams.

We grouped the data collection into blocks of approximately three minutes; we repeated each $l = 10$ experiment 250 times and the $l = 2000$ experiments 50 times. We servoed the magnetic field and re-calibrated the microwave power between each block.

5.7.3 Results




Figure 5.21: Single-qubit randomized benchmarking results, showing the average number of errors per sequence for each block of data. The black line represents a least-square fit to the results giving a mean error per gate of $1.0(3) \times 10^{-6}$ and an average SPAM error of $7(2) \times 10^{-4}$. Overlapping points are offset horizontally for clarity.

The results of these experiments are shown in figure 5.21. A fit to the data gives a mean error per gate of $1.0(3) \times 10^{-6}$ and an average SPAM error of $7(2) \times 10^{-4}$.

5.7.4 Error analysis

It is interesting to attempt to understand the various contributions to the measured error rate. In order to keep our investigation concrete we will focus our attention on the $l = 2000$ sequence used in these experiments, using the experimentally realized values of pulse durations, experimental dead times, etc.

5.7.4.1 Frequency error




Figure 5.22: Simulated and experimentally measured errors due to a constant microwave detuning from the qubit transition. We simulate errors both with and without a $14 \mu\text{s}$ experimental dead time between $\pi/2$ -pulses, finding that the dead time increases the experiment's sensitivity to detuning by a factor of roughly four. The measured data is taken by deliberately detuning the microwaves from the qubit transition and is not normalised to account for SPAM and other sources of error.

We begin then by considering errors due to a constant microwave detuning from the qubit transition. We model this error using the two-state system prop-

agator derived in section B.6. The propagators for each (imperfect) $\pi/2$ -pulse, along with experimental dead times, are multiplied together to calculate a total sequence propagator, from which error rates may be predicted. The results of this simulation are shown in figure 5.22. Also shown in this figure are measured error rates using the $l = 2000$ sequence with different microwave detunings.

We were not as careful to control experimental parameters which can shift the qubit frequency during the randomised benchmarking as we were during the spectroscopy of section 5.5. As a result, the microwaves were not set to be exactly resonant with the qubit frequency during the benchmarking experiments. Based on measurements made after the data was taken, we estimate that the microwave's detuning from the qubit transition was approximately +4.5 Hz, which would produce an error per gate of 6.7×10^{-7} (1.6×10^{-7} excluding the experimental dead time).

5.7.4.2 Pulse area error

Next, we consider errors due to a constant deviation of the microwave pulse area – the product of pulse duration and Rabi frequency – from $\pi/2$. This describes both systematic mis-settings of the microwave power/pulse length and fluctuations which are slow compared with the sequence length (our sequences ranged from 20 ms to 160 ms).

We apply the techniques discussed in the previous section to model gate errors due to a constant Rabi frequency error both for the case where the microwaves are resonant with the qubit transition and for the case where they are detuned by +4.5 Hz. The results of these simulations are shown in figure 5.23. Also shown are measured error rates using the $l = 2000$ sequence at a detuning of +4.5 Hz, with the Rabi frequency deliberately mis-set by a known amount.

Between each block of data in section 5.7.3 we re-calibrated the microwave power required to produce a $\pi/2$ -time of 12.1 μs , recording the result. We ob-




Figure 5.23: Simulated and experimentally measured errors due to a constant Rabi frequency error. We consider two cases: resonant microwaves and microwaves detuned by +4.5 Hz from the qubit transition. The SPAM error is not subtracted from the measured data, providing a baseline equivalent to an error per gate of $\sim 3 \times 10^{-6}$.

served fluctuations in the Rabi frequency of $\lesssim 5 \times 10^{-4}$ between blocks, which we believe to have been mainly due to thermal effects in the octupled DDS system. From our simulations, we expect this to produce an error of $\sim 3 \times 10^{-7}$ neglecting the 4.5 Hz detuning error, or 1.1×10^{-6} including it.

5.7.4.3 AC Zeeman shifts and off-resonant excitation

Off-resonant excitation of spectator transitions leads to two effects: AC Zeeman shifts and excitation of the ion from the qubit manifold. While these two phenomena both result from the same interaction, they have very different effects on the qubit. AC Zeeman shifts may be entirely compensated for by changing the microwave frequency during $\pi/2$ -pulses compared with during dead times and identity pulses. Alternatively, the microwave phase may be adjusted to provide a first-order compensation of the AC Zeeman shift's integrated effect (this was




Figure 5.24: Errors due to AC Zeeman shifts. We calculate the microwave field required to produce a given pulse π -time and the resulting AC Zeeman shifts using our measured microwave field angle of 63° . We then calculate the corresponding error by modelling the ion as a two-level system, driven by detuned radiation. Identity pulses and dead times are assumed error-free.

done, for example, by [KLR⁺08]). Off-resonant excitation out of the qubit manifold, on the other hand, may not be remedied so easily. It is thus useful to model the two effects separately to allow one to understand their relative importance.

We begin, then, by considering the qubit as a two-level system, driven by radiation that is detuned by an amount equal to the AC Zeeman shift. The results of this simulation are shown figure 5.24. Here, we have taken the microwave field to lie at 63° to the static magnetic field (section 6.2). The microwave field required to produce a $\pi/2$ -time of $12.1 \mu\text{s}$ leads to a AC Zeeman shift of -0.95 Hz (equivalent to a microwave detuning of $+0.95 \text{ Hz}$), producing a negligible gate error of 2×10^{-9} . We note that the AC Zeeman shift only affects the $\pi/2$ -pulses and not the identity pulses or the dead times, and so produces a much smaller error than the constant detunings considered in section 5.7.4.1.

Next, we consider the effect of off-resonant excitation out of the qubit man-




Figure 5.25: Errors due to off-resonant excitation of spectator transitions in the ground level. The blue curve represents a full simulation of benchmarking, including all states in the ground level. As in figure 5.24, the black curve represents a simulation of a two-level system driven by radiation detuned by the AC Zeeman shift. Comparing with 5.24, one sees that for the parameters considered in this thesis, excitation out of the qubit manifold dominates effects of AC Zeeman shifts.

ifold. We model the radiation field’s effect on the full ground-level of calcium using the propagator derived in section 5.2.2.3. As usual, we simulate the randomized benchmarking sequence by multiplying the propagators for each pulse to calculate a total sequence propagator (now for the full 16 state system). The results of this simulation are shown in figure 5.25. We see that errors due to off-resonant excitation of the ion from the qubit manifold dominate effects of AC Zeeman shifts. For the $12.1 \mu\text{s}$ $\pi/2$ time used in this work, we expect an error per gate of approximately 1×10^{-7} .

We note that the “noise” on the 16-state system simulation results from aliasing of high-frequency dynamics: for off-resonant excitation of a transition by a radiation field results in Rabi flopping at a rate approximately equal to the radiation’s detuning. In our case, the detunings are ~ 100 MHz, leading to dynamics with a periodicity of order 10 ns.

5.7.4.4 Qubit dephasing

Taking the ratio of our qubit's coherence time of 50 s to the gate time of 12.1 μ s, one might expect an error per gate of order 2×10^{-7} . Based on this crude estimate, one might expect qubit dephasing to have contributed a non-negligible amount to our measured error rate (particularly since each gate actually consists of 3 separate $\pi/2$ pulses separated by 14 μ s of dead time).

In fact, this model is extremely crude since it does not take the qubit's frequency noise spectrum into account. If, as is likely to be the case, the majority of the noise is at a low frequency compared with the 38 kHz pulse repetition rate then the noise will be strongly correlated from one pulse to the next. In this regime, the analysis of section 5.7.4.1 would be expected to apply and, since qubit frequency deviations are expected to be $\lesssim 10$ mHz, one would expect a negligible error.

Without taking more detailed measurements of the qubit's coherence over a shorter timescale, it is hard to estimate what effects qubit dephasing really had on this experiment.

5.7.4.5 Microwave phase and amplitude noise

The octupled DDS's single-sideband phase noise power spectral density is expected to be approximately (section 4.7.4)

$$L(f) = 10^{-10.7} \times \frac{1 \text{ kHz}}{f} \text{ Hz}^{-1} + 10^{-13.4} \text{ Hz}^{-1} \quad (5.24)$$

The $1/f$ component of the PM spectrum corresponds to a random walk in phase, producing an exponential loss of phase coherence. The resulting phase coherence time is given by (section C.2.3)

$$\tau_c = 1 / (2\pi f_{3\text{dB}}) \quad (5.25)$$

where $f_{3\text{dB}}$ is the noise's 3 dB frequency, which is approximately $10^{-7.4}$ Hz in this case. The resulting coherence time is many orders of magnitude larger than

the qubit's own coherence time and so clearly will not lead to significant errors.

One could provide a complete treatment of the effect of the DDS's white phase noise on a randomized benchmarking sequence by solving the master equation of section 3.4.3 to calculate a process map describing a $\pi/2$ -pulse in the presence of phase noise. This could be solved for a single value of the microwave phase, with other phases introduced by suitable rotations as in section 5.2.2.3.

We will (perhaps uncharacteristically) opt for a simpler approach, however. We use the analytic expression for the fidelity of a series of n π -pulses in the presence of phase noise given in (3.49), repeated below

$$\begin{aligned}\langle 1 - F \rangle &= n\pi \frac{N_0 \Omega}{6P_c} \\ &= \frac{n\pi^2}{6t_\pi} L\end{aligned}\tag{5.26}$$

where $t_\pi := \pi/\Omega$ is the pulse π -time. In our case $L = 10^{-13.4} \text{ Hz}^{-1}$, giving an error per gate of 3×10^{-9} at a π -time of $24 \mu\text{s}$.

Like its phase noise, the DDS's amplitude noise spectrum consists of both low-frequency and white components. The low-frequency AM noise is strongly correlated over the period of an experiment and so is covered by the analysis of section 5.7.4.2. Its broad-band single-sideband AM noise is expected to be around -140 dBc/Hz , dropping off above $\sim 100 \text{ kHz}$ (section 4.7.4).

As with the PM, we model this using the analytic expression for a series of π -pulses, given in section 3.4.1, repeated below

$$\begin{aligned}\langle 1 - F \rangle &= \frac{1}{3} \left(1 - e^{-n\pi \frac{N_0}{2P_0} \Omega_0} \right) \\ &\simeq \frac{n\pi^2}{6t_\pi} L\end{aligned}\tag{5.27}$$

For a π -time of $24 \mu\text{s}$, this corresponds to an error per gate of 7×10^{-10} .

On the basis of the previous analysis, we conclude that microwave phase and amplitude noise are unlikely to be a serious concern in these experiments.

5.7.4.6 Pulse length jitter

We now consider errors due to fluctuations in the $\pi/2$ -pulses' durations. We performed two kinds of simulation both shown in figure 5.26. Firstly, we considered the case in which the pulse length was mis-set by a constant amount that is the same for each pulse. Secondly, we considered the case where the lengths of each pulse undergo independent Gaussian fluctuations about their mean. We note that, under the conditions considered here, the two cases produce very similar errors.

We estimate timing jitter in our experiment to be $\lesssim 1$ ns, which should produce a negligible error.




Figure 5.26: Simulated effect of pulse length errors on the randomized benchmarking experiment. We consider both the case in which the $\pi/2$ -pulse duration is mis-set by the same amount for all pulses and in which the length of each pulse is randomly varied. In the latter case, the fluctuations are assumed to be Gaussian and independent for each pulse, with the abscissa denoting the RMS timing jitter. In both cases, the microwaves are assumed to be resonant with the qubit transition.

5.7.4.7 Microwave-ion coupling fluctuations

Since the microwave field is produced by the same trap electrodes responsible for creating the RF pseudopotential, we expect coupling fluctuations to be negligible in this experiment. Moreover, since the microwave field is extremely linear in the region explored by the ion, we do not expect effects of ion motion to produce a significant change in the ion-microwave coupling. We thus expect the Rabi frequency produced by a constant input power to the trap to be extremely stable over time.

As a brief experimental check of this, we used the trap's DC electrodes to move the ion $1.25\text{ }\mu\text{m}$ away from the trap's centre to induce a micromotion of 100 nm . We then recalibrated the DDS amplitude required to achieve a $\pi/2$ -time of $12.1\text{ }\mu\text{s}$ before performing the $l = 2000$ sequence. We detected no change in the total sequence error to within the measurement's sensitivity (limited by statistics) of 1×10^{-3} (equivalent to an error per gate of 5×10^{-7}).

5.7.4.8 Microwave and laser leakage and ion heating

To check for effects of laser/microwave leakage and ion heating-induced measurement errors, we performed a pair of check experiments. First, we prepared and read out both qubit states to measure the readout levels. Then, we performed a sequence, consisting of preparing one qubit state before waiting 40 ms and reading out the ion's resulting state. We measured no degradation of fidelity to within 0.1% with the 40 ms delay, suggesting that leakage and other effects were not significant.

Notice, however, that if the microwaves had been set to resonance (rather than detuned by 4.5 Hz) then the microwave leakage might have become a significant issue.

5.7.4.9 Effects of microwave duty-cycle

To estimate the impact of experimental duty cycle on the error rate, we calibrated the microwave amplitude with sequences of both 1501 and 751 $\pi/2$ -pulses. We did not measure a change in the DDS amplitude required, suggesting that duty-cycle effects were not a significant issue at the power-levels we are working at.

5.7.4.10 Discussion

We see that the bulk of our error rate comes from the 7×10^{-7} contribution due to combination of the 4.5 Hz microwave detuning with the experimental dead-time. The other main source of error stems from slow power fluctuations in the octupled DDS, contributing an error of roughly 3×10^{-7} . Off-resonant excitation is also expected to play a minor role at about the 1×10^{-7} level.

5.7.5 Effects of sample size

In the introduction to this section, we suggested that the usefulness of randomized benchmarking as a technique relies on the claim that the errors we measure are representative of a randomly chosen sequence. It is interesting to see how true this claim is in the case of our experiments. To test this, we simulated the errors on 500 different randomized sequences, all with $l = 2000$, $N_G = 4$ and $N_R = 8$, like the sequence we used. We assumed 12.1 μs $\pi/2$ -pulses with a constant +4.5 Hz detuning error, a constant 5×10^{-4} Rabi frequency error and a 14 μs dead time (i.e. approximately our experimental conditions).

The results of this simulation are shown in figure 5.27. We see that the data is reasonably well approximated by a normal distribution with $\mu = 0.81(1) \times 10^{-6}$ and $\sigma = 1.4(1) \times 10^{-7}$, with the particular $l = 2000$ sequence used in our experiments about 2 standard deviations above the mean error. The standard deviation of these points is an order of magnitude smaller than their mean. We thus conclude that randomization used here is sufficient to ensure that, with a

high degree of probability, a randomly chosen sequence will be representative of the mean error to better than a factor of two.





Figure 5.27: Simulation of errors in 500 different $l = 2000$ sequences (all with $N_G = 4$ and $N_R = 8$), under our experimental conditions ($12.1 \mu\text{s}$ $\pi/2$ -pulses with a $+4.5 \text{ Hz}$ detuning error, a 5×10^{-4} Rabi frequency error and a $14 \mu\text{s}$ dead time). The black line is a fit to a normal distribution with $\mu = 0.81(1) \times 10^{-6}$ and $\sigma = 1.4(1) \times 10^{-7}$. The blue point represents the $l = 2000$ sequence used in our randomized benchmarking experiments.

6

Towards a Microwave-Driven Two-Qubit Gate



Reproduced by kind permission of PRIVATE EYE magazine/Jon & Mick/Modern Toss.

6.1 Introduction

In this chapter we describe work related to the microwave-driven ion-ion entanglement scheme discussed in section 2.7. We begin in section 6.2 by describing experiments used to characterise the fields generated by our trap's three microwave electrodes. In section 6.3 we present results of an experimental implementation of the dark-resonance cooling technique discussed in section 2.5. We

show that even at 146 G one may cool $^{43}\text{Ca}^+$ to a temperature close to the $^{40}\text{Ca}^+$ Doppler limit using only the standard 397 nm and 866 nm lasers.

The remainder of the chapter is devoted to a discussion of the effects we expect to present the dominant sources of error in any future microwave-driven two-qubit gate: heating of the motional mode (section 6.4); dephasing of the motional mode due to fast RF amplitude fluctuations (section 6.5); coupling between the ions' stretched and rocking modes (section 6.6); and, drifts in the ions' radial secular frequency (section 6.7).

Most of the work presented in this chapter is part of an on-going investigation and experimental results should be treated as preliminary data to be backed up and expanded upon by future work.

6.2 Characterising and controlling the microwave field

If we are to produce a controlled ion-microwave interaction then we must have a method of characterising the field created by a given configuration of currents. Moreover, to prevent unwanted single-qubit rotations during two-qubit gates, we need a technique for adjusting the amplitudes and phases of the microwave current in each electrode to produce a strong field gradient with no net field.

In this section, I discuss techniques for achieving this based on Rabi and Ramsey spectroscopy. Results for our experiment using Rabi spectroscopy were originally presented in [AHB⁺13], while the Ramsey technique is based on the work of [WOC⁺12].

6.2.1 The field

The microwave magnetic field at a point $\mathbf{r}$ near the trap's centre may be expressed as

$$\mathbf{B}(\mathbf{r}, t) = B_z(z, y)\hat{\mathbf{z}} \cos \omega t + B_y(z, y)\hat{\mathbf{y}} \cos(\omega t + \phi) \quad (6.1)$$

where we have assumed that due to the trap's axial symmetry the field has no component along $\hat{x}$ (the trap's symmetry axis) and, moreover, that it has translational symmetry along $\hat{x}$.

The z -component of the field lies along the quantization axis and thus represents π -polarized radiation, whilst the y -component represents equal quantities of $\sigma^\pm$ polarization. In the spherical basis the field has components

$$\begin{aligned} B_{\pm 1}(z, y) &= \mp \frac{1}{\sqrt{2}} B_y(z, y) \\ B_\pi(z, y) &= B_z(z, y) \end{aligned} \tag{6.2}$$

Our goal is to find a way of measuring these field components as a function of position in the region of space around the trap's centre.

6.2.2 Rabi spectroscopy

A simple means of achieving this goal would be to measure the rates of Rabi flopping on a π - and a σ -polarized transition induced by a given configuration of microwave currents. The field components may then be calculated from the measured Rabi frequencies. The ion may then be displaced a known distance by applying bias voltages to the trap electrodes (section 6.2.4) before repeating the measurements. In this way one may build up a map of $B_\pi(z, y)$ and $B_{\pm 1}(z, y)$.

This technique provides a simple method of characterising the microwave field. It suffers, however, from the drawback that the frequency of the microwaves must be tuned to resonance with the transitions used. In the intermediate field, this involves changing the microwave frequency by tens of megahertz. Since the amplitudes and, in particular, the phases of the microwaves will vary with frequency, this introduces an error into the measurement. We would thus like a way of measuring the microwave field as close to the frequency we intend to use as possible.

Despite its limitations, this technique was used during initial characterisation of this experiment [AHB⁺13]. During this work we measured the Rabi rate




Figure 6.1: Experiment used to measure the A.C. Zeeman shift on the transition between a pair of ground-level states $|i\rangle$ and $|j\rangle$. We begin by preparing $|i\rangle$ using optical pumping followed by microwave transfer pulses before applying a $\pi/2$ -pulse on the $|i\rangle \leftrightarrow |j\rangle$ transition. We next wait for a Ramsey delay τ before applying a spin-echo π -pulse. The microwave field to be characterised is then applied for part of the second arm of the Ramsey interferometer, which is closed by a final $\pi/2$ -pulse. The ion's resulting state is then mapped onto a suitable readout qubit using further microwave transfer pulses. From here, it is read out as usual.

on the transitions $|4, 0\rangle \leftrightarrow |3, 0\rangle$ and $|4, 0\rangle \leftrightarrow |3, +1\rangle$ at a static magnetic field of 23 G. At this field, the two transitions used are only split by 8.1 MHz, reducing the size of the required frequency shifts.

6.2.3 Ramsey spectroscopy of A.C. Zeeman shifts

A second approach [WOC⁺12] uses Ramsey spectroscopy to measure the A.C. Zeeman shifts the microwaves induce on transitions within the ground level. The total shift on each transition arises from off-resonant excitation of spectator transitions, in general, with both π - and σ -polarizations. The shifts thus contain information about both field components and so, by measurement of two such shifts, one obtains sufficient information to fully determine the field at a point.

The principal advantage of this method is that the microwaves do not have to be tuned to resonance with the shifted transitions. One may thus characterise the exact field one wishes to use.

The Ramsey experiment used to measure the AC Zeeman shifts is illustrated in figure 6.1. Since most of the transitions of interest are magnetic field sensitive, the Ramsey experiment is conducted using a spin-echo-type sequence, which eliminates effects of slow field fluctuations. During these experiments, we keep the total Ramsey delay constant and only vary the fraction of the sequence for which the microwaves are applied. This helps to minimise measurement errors

due to systematic effects such as 50 Hz magnetic field noise by keeping their effect constant over the course of the experiment.

6.2.4 Displacing the ion

We move the ion by applying DC voltages to the trap's 9 electrodes. By application of appropriate voltages, we are able to produce a well-known static E-field, $\mathbf{E}_{\text{dc}}$, which displaces the ion along $\hat{\mathbf{x}}_i$ by a distance equal to [BMB⁺98a]

$$\delta x_i = \frac{Q\mathbf{E}_{\text{dc}} \cdot \hat{\mathbf{x}}_i}{m\omega_i^2} \quad (6.3)$$

where Q and m are the ion's charge and mass respectively and ω_i is the frequency of secular motion along $\hat{\mathbf{x}}_i$.

We calculate the electrode voltages required to produce a given E-field using an electrostatic model of the trap created with commercial *Charged Particle Optics* (CPO) numeric boundary element software [AHJ⁺11].

To allow laser beams propagating parallel to the trap's surface to cool all of the ions' normal modes, we apply a static DC "tilt" quadrupole [AHJ⁺11]. This field breaks the degeneracy of the normal modes, producing secular frequencies $\omega_{z'}$ and $\omega_{y'}$ along the rotated axes $\hat{\mathbf{z}}' := \frac{1}{\sqrt{2}}(\hat{\mathbf{z}} + \hat{\mathbf{y}})$ and $\hat{\mathbf{y}}' := \frac{1}{\sqrt{2}}(\hat{\mathbf{z}} - \hat{\mathbf{y}})$. The tilt field is chosen so that $\omega_{y'} < \omega_{z'}$. This couples the z- and y-motion so that

$$\left. \begin{aligned} \delta z &= \frac{QE_{\text{dc}}}{2m} \left(\frac{1}{\omega_{z'}^2} + \frac{1}{\omega_{y'}^2} \right) \\ \delta y &= \frac{QE_{\text{dc}}}{2m} \left(\frac{1}{\omega_{z'}^2} - \frac{1}{\omega_{y'}^2} \right) \end{aligned} \right\} \mathbf{E}_{\text{dc}} = E_{\text{dc}}\hat{\mathbf{z}} \quad (6.4)$$

$$\left. \begin{aligned} \delta z &= \frac{QE_{\text{dc}}}{2m} \left(\frac{1}{\omega_{z'}^2} - \frac{1}{\omega_{y'}^2} \right) \\ \delta y &= \frac{QE_{\text{dc}}}{2m} \left(\frac{1}{\omega_{z'}^2} + \frac{1}{\omega_{y'}^2} \right) \end{aligned} \right\} \mathbf{E}_{\text{dc}} = E_{\text{dc}}\hat{\mathbf{y}}$$

To confirm our model, we made two kinds of measurement. Firstly we applied E-fields along $\hat{\mathbf{z}}$ and $\hat{\mathbf{y}}$ while monitoring the ion's position in the x - z plane on the imaging system's camera. Since the imaging system's magnification and

pixel size are well known, this provides accurate information about the ion's location as a function of electrode voltages. Secondly, we applied varying fields along $\hat{y}$ and found the $\hat{z}$ field that produced the smallest micromotion in the x - z plane. This provides a measure of the z - y cross-coupling induced by the tilt quadrupole.

We observed that applying a field nominally in the $\hat{z}$ direction moves the ion in the x - z plane at $\sim 17^\circ$ to the z -axis. The cause of this angle is not known. Other than that, the measurements were in agreement with the model to within 5 %. These results suggest that the ion's displacement can be inferred from the applied voltages to within $\lesssim 10\%$.

We note that the displacement from the RF null also produces micromotion with amplitude [BMB⁺98a]

$$a_i = \frac{1}{2}q_i\delta x_i \simeq \sqrt{2}\frac{\omega_i}{\Omega_{\text{rf}}}\delta x_i \quad (6.5)$$

where q_i is the Mathieu q (stability) parameter and Ω_{rf} is the trapping RF frequency. For this trap, $\Omega_{\text{rf}} = 38.7$ MHz and $\omega_i \simeq 3$ MHz, giving $q = 0.2$. We thus expect the displacement to produce micromotion with amplitude equal to 10% of the displacement. Provided the microwave field is linear over this region, this micromotion will not affect the measurement.

6.2.5 Polarization angles

To assess the impact of off-resonant excitation on single-qubit operations one needs to know the polarization angles¹ of the microwave fields used for state-preparation and single-qubit rotations. Since these operations span a frequency range of 300 MHz, which is comparable to the bandwidth of our trap's microwave cavities, one must measure the polarization angle at the frequency of each transition.

¹The angle the field's polarization vector makes to the $\hat{z}$ -axis.

Transition	Frequency (GHz)	AC Zeeman			
A	2.874	$ 4, +2\rangle \Leftrightarrow$	$ 3, +1\rangle$	$ 4, 0\rangle \Leftrightarrow$	$ 3, +1\rangle$
B	2.986	$ 4, +1\rangle \Leftrightarrow$	$ 3, +1\rangle$	$ 4, 0\rangle \Leftrightarrow$	$ 3, +1\rangle$
C	3.095	$ 4, +4\rangle \Leftrightarrow$	$ 3, +3\rangle$	$ 4, +3\rangle \Leftrightarrow$	$ 3, +3\rangle$
Q	3.200	$ 4, +4\rangle \Leftrightarrow$	$ 3, +3\rangle$	$ 4, +2\rangle \Leftrightarrow$	$ 3, +3\rangle$

Table 6.1: We determine the microwave field’s polarization angle at frequencies corresponding to transitions of interest by measuring the rate of Rabi flopping on the transition, as well as the AC Zeeman shift the field induces on two other transitions. This table list the transitions we characterised in this manner, as well as their frequencies and the extra transitions whose AC Zeeman shift we measured.

Electrode	Polarization angle at transition (°)			
	A	B	C	Q
Left	42.2	73.5	59.1	58.4
Centre	2.1	1.1	3.5	3.1
Right	38.3	73.2	63.0	62.9

Table 6.2: Measured polarization angles of the fields produced when microwaves with frequencies corresponding to the four transitions used during single-qubit operations are applied to each of the trap’s three electrodes.

We achieved this using a combination of Rabi and Ramsey spectroscopy. We begin by measuring the Rabi frequency on the transition of interest to determine the field’s σ component. Next, we measure the AC Zeeman shifts on two further transitions. By combining these AC Zeeman shift measurements with the known field’s σ component, one may determine the field’s π component and hence its polarization angle.

The transitions used for these measurements are listed in table 6.1, and the measured polarization angles are given in table 6.2.

6.2.6 Field gradient measurement

To characterise the field gradient that will be responsible for driving our two-qubit gate, we need to measure the microwave magnetic field as a function of position around the trap’s centre. We perform this measurement using microwaves tuned near to the red and blue motional sidebands of the qubit transition since

these are the frequencies that will be used during the gate.

We measure the field at a point using Rabi spectroscopy to determine the AC Zeeman shifts on the three transitions: $|4, 0\rangle \Leftrightarrow |3, +1\rangle$, $|4, 0\rangle \Leftrightarrow |3, 0\rangle$ & $|4, +1\rangle \Leftrightarrow |3, +1\rangle$. We chose these transitions because their different dependences on the field's π and σ components allow an accurate determination of the field (see tables 6.3 - 6.5): the qubit transition's AC Zeeman shift is extremely insensitive to the field's π -component because the high degree of symmetry means that the contributing shifts cancel. The other two transitions have a greater sensitivity to the π field. Moreover, on one transition the π -shift adds to the σ , while on the other it subtracts.

Transition	Matrix element (μ_B)	Detuning (MHz)	Red shift (Hz/ μT^2)	Blue shift (Hz/ μT^2)
$ 4, +0\rangle \Leftrightarrow 3, -1\rangle$	+0.528	+50.0	+0.133	+0.141
$ 4, +0\rangle \Leftrightarrow 3, +0\rangle$	-0.993	+0.0	+1.786	+2.005
$ 4, +0\rangle \Leftrightarrow 3, +1\rangle$	+0.702	-51.5	+8.051	-8.044
$ 4, -1\rangle \Leftrightarrow 3, +0\rangle$	-0.877	+50.0	+0.366	+0.388
$ 4, +1\rangle \Leftrightarrow 3, +0\rangle$	+0.702	-51.6	+8.285	-7.828
Total π -shift			+1.786	+2.005
Total σ -shift			+16.836	-15.344

Table 6.3: AC Zeeman shifts on the $|4, 0\rangle \Leftrightarrow |3, 0\rangle$ transition due to radiation detuned by ± 3 MHz from the qubit transition.

Transition	Matrix element (μ_B)	Detuning (MHz)	Red shift (Hz/ μT^2)	Blue shift (Hz/ μT^2)
$ 4, 0\rangle \Leftrightarrow 3, -1\rangle$	+0.528	+101.5	+0.133	+0.141
$ 4, 0\rangle \Leftrightarrow 3, 0\rangle$	-0.993	+51.5	+0.893	+1.002
$ 4, 0\rangle \Leftrightarrow 3, +1\rangle$	+0.702		+16.103	-16.088
$ 4, +1\rangle \Leftrightarrow 3, +1\rangle$	-0.993	-51.6	-0.986	-0.876
$ 4, +2\rangle \Leftrightarrow 3, +1\rangle$	+0.887	-105.0	-0.372	-0.351
Total π -shift			-0.0926	+0.126
Total σ -shift			+15.864	-16.298

Table 6.4: AC Zeeman shifts on the $|4, 0\rangle \Leftrightarrow |3, +1\rangle$ qubit transition due to radiation detuned by ± 3 MHz from the qubit transition.

We have measured the fields produced by the trap's centre and right electrodes using microwaves tuned +4 MHz from the qubit transition (i.e. the blue




Figure 6.2: Measured field components produced when microwaves detuned by +4 MHz from the qubit transition (i.e. the blue motional sideband) are applied to the trap's centre or right electrodes. Contours show the field's σ and π components in μT as a function of position around the trap's centre after normalisation to an input power of 1 mW measured before the diplexer. The π component of the right electrode's field is not shown because it is not sufficiently well determined by the fits to give useful information (this is not the case for the centre electrode, since the π field is significantly larger).

Transition		Matrix element (μ_B)	Detuning (MHz)	Red shift (Hz/ μT^2)	Blue shift
$ 4, +1\rangle \Leftrightarrow$	$ 3, +0\rangle$	+0.702	+51.5	+8.285	-7.828
$ 4, +1\rangle \Leftrightarrow$	$ 3, +1\rangle$	-0.993	+0.0	-1.971	-1.753
$ 4, +1\rangle \Leftrightarrow$	$ 3, +2\rangle$	-0.514	-53.3	-0.125	-0.118
$ 4, +0\rangle \Leftrightarrow$	$ 3, +1\rangle$	+0.702	+51.6	+8.051	-8.044
$ 4, +2\rangle \Leftrightarrow$	$ 3, +1\rangle$	+0.887	-53.4	-0.372	-0.351
Total π -shift				-1.971	-1.753
Total σ -shift				+15.840	-16.341

Table 6.5: AC Zeeman shifts on the $|4, +1\rangle \Leftrightarrow |3, +1\rangle$ transition due to radiation detuned by ± 3 MHz from the qubit transition.

Electrode	$\frac{\partial B_\sigma}{\partial z'} \text{ (T/m)}$	$\frac{\partial B_\sigma}{\partial y'} \text{ (T/m)}$
Centre	0.128	0.136
Right	0.046	0.087

Table 6.6: Measured field gradients along the $\hat{z}'$ and $\hat{y}'$ axes normalised to an input power of 1 mW measured before the diplexer.

motional sideband). The results of these measurements are shown in figure 6.2. To calculate the corresponding field gradients, we fitted the measured data to the model

$$B_\sigma(z, y) = B_\sigma^0 + \frac{\partial B_\sigma}{\partial z'} z' + \frac{\partial B_\sigma}{\partial y'} y' \quad (6.6)$$

The fitted gradients are given in table 6.6.

6.2.7 Field nulling

To null the field at the trap's centre, we first apply a microwave current to one of the trap's electrodes. Next, we apply currents to the other two electrodes while monitoring the field using Rabi or Ramsey techniques. We then adjust the amplitudes and phases of these currents until the field is minimised.

We have previously nulled the field using Rabi spectroscopy as described in [AHB⁺13]. Because of the drawbacks of this technique discussed above, we intend to use AC Zeeman shift measurements in future work. For this, we will pick a transition where the σ and π AC Zeeman shifts combine with the same

sign so that, by minimising the AC Zeeman shift, one also minimises the field. From tables 6.3 and 6.5, we see that the $|4, 0\rangle \leftrightarrow |3, 0\rangle$ transition is suitable for the red sideband, while the $|4, +1\rangle \leftrightarrow |3, +1\rangle$ is suitable for the blue sideband.

6.2.8 Outlook

In future experiments we intend to apply peak microwave powers of ~ 10 W to all three trap electrodes to produce a gradient of ~ 10 T/m whilst maintaining only a small residual field at the trap's centre. During the preliminary work conducted so far, we have limited the peak power to ~ 100 mW. This decision has been motivated by uncertainty over the trap's damage threshold.

We currently only apply microwave power to the centre electrode since, as one can see from table 6.6, this produces the strongest field gradient for a given input power. We recall that the Rabi frequency for sideband transitions is given by (section 2.7)

$$\hbar\Omega_{j,n} = R_{10}B'_\sigma b_{j,n}q_j^0 \quad (6.7)$$

For a single ion mode ($b_{j,n} = 1$) in a 3 MHz trap ($q_j^0 = 6.2$ nm), given the gradients of table 6.6 scaled to an input power of 100 mW, we expect sideband Rabi frequencies of

$$\begin{aligned} \Omega_{z'} &= 2\pi \times 78 \text{ Hz} \\ \Omega_{y'} &= 2\pi \times 83 \text{ Hz} \end{aligned} \quad (6.8)$$

Without nulling from the left and right electrodes, this gradient is accompanied by a field of

$$\begin{aligned} B_\sigma &= 6.6 \mu\text{T} \\ B_\pi &= 173.1 \mu\text{T} \end{aligned} \quad (6.9)$$

This field leads to a AC Zeeman shift of 3 kHz (-2 kHz) when the microwaves are tuned to resonance with a 3 MHz blue (red) sideband. When both sidebands are on at the same time (as during a gate) the total shift is 981 Hz. Since the microwave field is stable to better than 1 %, this shift may be measured and

compensated for by adjusting the microwave frequency and so is not expected to lead to errors.

The field also leads to low-amplitude Rabi flopping on the various spectator transitions, causing oscillations of the atomic populations with amplitude

$$P = \frac{\Omega^2}{\Omega^2 + \delta^2} \quad (6.10)$$

where, as usual, δ is the transition's detuning in radians/second. The resulting oscillations have amplitudes of 5×10^{-4} on the carrier transition and 2×10^{-3} on the two π transitions from the qubit states. This is negligible for our current experiments.

6.3 Low temperature cooling on a dark resonance

The cooling parameters we considered in section 5.1 allow one to trap ions and to attain a high rate of fluorescence; however, as discussed in section 2.5, they leave the cooled ions far from their motional ground states. While one of the benefits of the microwave-driven entanglement scheme is its insensitivity to the ions' initial motional states, it is still advantageous for them to be as cold as possible, as this simplifies the gate dynamics and allows for better diagnostics (this is discussed in the following sections, see also [KBZ⁺09]).

One could follow initial Doppler cooling with resolved sideband cooling methods using either lasers or microwaves; however, these techniques require additional laser/microwave overhead and are typically inefficient when beginning from high thermal states.

In section 2.5 we described a simple scheme to achieve lower temperatures from Doppler cooling by tuning the 397 nm and 866 nm lasers to meet a two-photon resonance condition. In this section, we discuss experimental implementation of this scheme, demonstrating that one may cool $^{43}\text{Ca}^+$ ions at 146 G to near the $^{40}\text{Ca}^+$ Doppler limit.




Figure 6.3: Pulse sequences used to drive (a) blue and (b) red sideband transitions. The ion is prepared in $|3, +1\rangle$ before two pulses are applied: a π -pulse on the qubit transition and a microwave pulse whose frequency is tuned one motional frequency below the qubit transition. By choosing the order of these pulses, we determine whether the microwave pulse adds or subtracts momentum from the ion and hence whether it is a “blue” or “red” sideband pulse. This point is demonstrated in (c), which shows the possible transitions out of each qubit state.

6.3.1 Sideband thermometry

We measure the temperature of a $^{43}\text{Ca}^+$ ion by performing sideband spectroscopy using the microwave gradient. This technique takes advantage of the fact that when an ion is cooled sufficiently for its motional quantization to become significant the Rabi frequencies for blue and red motional sideband transitions are different. We recall that the Rabi frequency for the transition between the motional Fock states $|n\rangle$ and $|n+1\rangle$ is [WMI⁺98]

$$\Omega_{n,n+1} = \sqrt{n+1} \Omega \quad (6.11)$$

where Ω is the sideband Rabi frequency, as defined in section 2.7. We see that the red sideband transition is always slower than the blue², vanishing completely for an ion in the ground motional state.

After Doppler cooling the ion is left in a random Fock state, distributed according to the Boltzmann distribution. We may describe this thermal state by

²Here, we ignore effects of curvature in the driving field that can become significant for large n .

the density operator

$$\rho_{\text{th}} = \sum_{n=0}^{\infty} \frac{\bar{n}^n}{(\bar{n} + 1)^{n+1}} |n\rangle\langle n| \quad (6.12)$$

where $\bar{n}$ is the mean motional quantum number (occupancy).

When we drive sideband transitions on an ion in such a thermal state, we observe a weighted sum of independent Rabi oscillations for the different motional Fock states. In particular, if the ion begins in the qubit state $|0\rangle$ then its probability of being in $|1\rangle$ after applying a red/blue sideband pulse for a period t is given by

$$\begin{aligned} P_1^b(t) &= \sum_{n=0}^{\infty} \frac{\bar{n}^n}{(\bar{n}+1)^{n+1}} \sin^2\left(\frac{1}{2}\Omega_{n,n+1}t\right) \\ P_1^r(t) &= \sum_{n=1}^{\infty} \frac{\bar{n}^n}{(\bar{n}+1)^{n+1}} \sin^2\left(\frac{1}{2}\Omega_{n,n-1}t\right) \end{aligned} \quad (6.13)$$

Combining these equations, we find

$$\frac{P_1^r(t)}{P_1^b(t)} = \frac{\bar{n}}{\bar{n} + 1} \quad (6.14)$$

and hence that the ion's temperature may be deduced from the asymmetry in the measured populations. An important feature of equation (6.14) is that it does not involve the sideband Rabi frequency. The temperature may thus be deduced even in the face of uncertainty in Ω .

We measure the temperature using the experiment illustrated in figure 6.3. We drive the sideband transitions by applying $\simeq 100$ mW of microwave power to the trap's centre electrode. To minimise systematic errors due for example to changes in microwave power with frequency, we drive both the red and the blue sideband using the same microwave frequency, using a π -pulse on the qubit transition to determine which sideband is driven (figure 6.6).

6.3.2 The Doppler Limit

The Doppler limit of a two-level system (section 2.5) provides a useful benchmark for evaluating the performance of a cooling scheme. We recall that the

Doppler limit is given by

$$k_B T_D = \frac{\hbar \Gamma}{2} \quad (6.15)$$

where Γ is the rate of decay from the upper level. Our sideband measurements determine the ion's mean vibrational state, $\bar{n}$, which is related to its temperature by

$$k_B T_{\bar{n}} = \hbar \omega \bar{n} \quad (6.16)$$

Comparing (6.15) and (6.16) we find

$$\omega \bar{n} = \frac{1}{2} \Gamma \quad (6.17)$$

For $^{40}\text{Ca}^+$ cooled using the 397 nm transition, $\Gamma = 140 \times 10^6 \text{ s}^{-1}$, giving a Doppler limit of 0.5 mK. For a 3 MHz trap frequency, this corresponds to a mean thermal occupancy of $\bar{n} = 3.5$.

6.3.3 Experimental results

We have performed preliminary experiments to optimise the cooling parameters to minimise the ion's temperature. We use the same EOM settings used for the high-fluorescence cooling, that is: a 2.930 GHz EOM frequency, with the carrier addressing the $F = 3$ manifold and the blue sideband addressing the $F = 4$ manifold; and with the sideband intensities 85 % of the carrier.

Like the high fluorescence cooling discussed in section 5.1.2, we find it beneficial for the 397 nm laser to be tuned as far to the blue as conveniently possible. In practice, we set the 397 nm laser frequency by detuning $\lesssim 10$ MHz to the red of the frequency at which the lasers' net effect is to heat the ion. This produces a typical 397 nm carrier (sideband) detuning of $-1893(15)$ MHz ($+1037(15)$ MHz) relative to the transition's centre of gravity, which is $-1205(15)$ MHz ($+1725(15)$ MHz) relative to $^{40}\text{Ca}^+$ at zero field. At this detuning, the sideband is approximately resonant with the $4S_{1/2}|4, +4\rangle \Leftrightarrow 4P_{1/2}|4, +4\rangle$ cooling transition, whose frequency is $+1\,020$ MHz relative to the




Figure 6.4: Scan of the 866 nm laser detuning taken using our standard low-temperature cooling parameters (see text). The 866 nm detuning is specified relative to the transition’s centre of gravity. This may be related to the $4P_{1/2} \leftrightarrow 3D_{3/2}$ transition in $^{40}\text{Ca}^+$ at zero field by adding the isotope shift of $-3\,464.3(3)$ MHz. The coldest temperatures are achieved when the 866 nm laser is tuned near to the bottom of the dark resonance feature at -170 MHz ($-3\,6434$ MHz relative to $^{40}\text{Ca}^+$). The dashed line shows results of a simulation using our estimated parameters and assuming a 100 kHz laser linewidth. The match between simulation and measurements is generally good apart from regions around the blue sides of dark resonances, where the ion is expected to heat.

397 nm transition’s centre of gravity. The carrier is $\lesssim 10$ MHz red detuned from the $4S_{1/2}|3, +3\rangle \leftrightarrow 4P_{1/2}4, +3$ transition at -1.874 MHz.

Our standard 397 nm and 866 nm intensities are $30(6) I_0$ and $240(48) I_0$ respectively. Here, as throughout this section, the 397 nm intensity refers to the total intensity in the carrier and two sidebands. The power in the carrier (sidebands) is thus $11(2) I_0$ ($9.5(20) I_0$).

A scan of the 866 nm laser’s detuning under these conditions is shown in figure 6.4. For best cooling, we tune the 866 nm near to the bottom of the main dark resonance feature at approximately -170 MHz relative to the transition’s centre of gravity ($-3\,6434$ MHz relative to $^{40}\text{Ca}^+$).

Figure 6.5 shows a pair of red and blue sideband scans taken under these conditions. One can see a visible asymmetry between the red and blue sidebands, indicating that the ion is cooled into the quantum regime. A fit to the




Figure 6.5: Scans of the microwave detuning during a red and blue sideband pulse. We see an asymmetry between the heights of the two sidebands, from which the temperature may be deduced. Fitting the two scans while floating the Rabi frequency and temperature gives $\Omega = 2\pi \times 103(4)$ Hz, $\bar{n} = 2.5(3)$. Here the error bars represent statistical uncertainties, which are expected to be significantly smaller than systematic errors.

data allowing both the Rabi frequency and temperature to vary gives $\Omega = 2\pi \times 103(4)$ Hz and $\bar{n} = 2.5(3)$. However, care must be taken in interpreting this result: since the asymmetry is only small, the fit is able to trade off Rabi frequency against temperature to a large degree. From the microwave field gradient measurements, we expect our Rabi frequency to be $\Omega = 2\pi \times 75$ Hz. Fitting the scans with this Rabi frequency gives a temperature of $\bar{n} = 5.1(1)$ with a similar quality of fit. Based on fits to the data, one can conclude that the ion is within a factor of two of the Doppler limit with some confidence, however the main point of these experiments is to demonstrate that we reach a temperature at which microwaves can be used to cool.

Figure 6.6 shows the result of varying the 866 nm detuning while keeping all other parameters at their optimised values. We began and ended this experiment by taking detuning scans of the 866 nm laser (figure 6.6 (a)). In between these scans, we took red and blue sideband scans at different 866 nm detunings, beginning at the hot (blue) side of the resonance and moving monotonically to the red. One can see a drift in the dark resonance position of roughly 7 MHz




Figure 6.6: Measured ion temperature ($\bar{n}$) as a function of 866 nm detuning using our standard cooling parameters (see text). The 866 nm detuning is specified relative to the transition's centre of gravity. This may be related to the $4P_{1/2} \leftrightarrow 3D_{3/2}$ transition in $^{40}\text{Ca}^+$ at zero field by adding the isotope shift of $-3\,464.3\,(3)$ MHz. Error bars represent statistical uncertainty and do not account for systematic effects as discussed in the text.

over the course of the data set. This is far in excess of the specified drift rate of the reference etalons the lasers are locked to³ and prevents any quantitative correlation between figures 6.6a and 6.6b. However, it is clear that the decreasing slope and increased absolute fluorescence as one leaves the resonance in 6.6a are reflected in the higher temperature.

We see that the lowest temperature is achieved with the 866 nm tuned near to the bottom of the dark resonance feature, however the detuning may be changed by several MHz without significantly impacting the temperature.

The effects of varying the 397 nm and 866 nm intensities are illustrated in figure 6.7. Increasing the 397 nm intensity increases the fluorescence rate and, correspondingly, the heating due to photon recoil. We find that doubling the 397 nm intensity from our optimised value roughly doubles the temperature. We do not find a significant change in the ion's temperature with decreasing 397 nm intensity. The temperature is not a strong function of the 866 nm intensity, however we find that changing the 866 nm intensity by a factor of 2 makes a roughly 30 % difference to the final temperature.

The temperature increases monotonically with increasing 397 nm detuning to the red. Increasing the detuning by 10 MHz makes a roughly 50 % difference in the final temperature.

6.4 Heating rate

The two-qubit gate couples ions indirectly, relying on one of their shared motional modes to act as an intermediary. While the ions' spins are entangled with their motional states they become vulnerable to any noise process that affects the motion. A significant such noise source arises from stray fluctuating electric fields coupling to the ions' charges, leading to heating of the motional modes.

³Specified drift rate is $\lesssim 0.5$ MHz/hour, compared with a data acquisition time of roughly 30 minutes.




Figure 6.7: Detuning scans of the 866 nm laser with different 397 nm and 866 nm intensities. In (a) the 866 nm intensity is fixed at $240 I_0$ (our standard value) and we vary the 397 nm intensity. As usual, the 397 nm power indicates the total power in the carrier and both sidebands (see text), with $30 I_0$ being our standard value. (b) the 397 nm intensity is fixed at $30 I_0$ and we vary the 866 nm intensity. In both figures the 866 nm detuning is specified relative to the transition's centre of gravity. This may be related to the $4P_{1/2} \leftrightarrow 3D_{3/2}$ transition in $^{40}\text{Ca}^+$ at zero field by adding the isotope shift of $-3\,464.3(3)$ MHz.




Figure 6.8: (a) Simulated gate error as a function of heating rate for a single-loop Mølmer-Sørensen gate on two ions. (b) Measured temperature of a single-ion 2.8 MHz radial mode as a function of heating delay. The fit gives a heating rate of 77 quanta per second.

To understand how heating affects a gate, we consider the master equation (see sections 2.7 & 3.5.3)

$$\dot{\rho} = -\frac{i}{\hbar}[H_{\text{MS}}, \rho] + \frac{1}{2}\dot{n} \left[2a^\dagger \rho a + 2a \rho a^\dagger - aa^\dagger \rho - a^\dagger a \rho - \rho a^\dagger a - \rho a a^\dagger \right] \quad (6.18)$$

where H_{MS} is the Mølmer-Sørensen Hamiltonian, given by

$$H_{\text{MS}} = \frac{1}{2}\hbar\Omega S_{\bar{\phi}} \left(ae^{i\delta t} e^{-\frac{1}{2}i\delta\phi} + a^\dagger e^{-i\delta t} e^{\frac{1}{2}i\delta\phi} \right) \quad (6.19)$$

and $\dot{n}$ is the motional heating rate in units of quanta per second.

To calculate the gate fidelity, we integrate this master equation numerically. We consider a single-loop gate, with gate time and detuning $\tau_g = \pi\sqrt{2}/\Omega_1$ & $\delta = \sqrt{2}\Omega_1$ respectively (where Ω_1 is the single-ion sideband Rabi frequency)⁴. The results of this simulation are shown in figure 6.8a.

We measure the heating rate using sideband thermometry with an experiment similar to that of figure 6.3. Here, we insert a delay between state preparation and the microwave pulses, during which the motional modes are allowed

⁴Notice that since we do not include a detuning error into the model the simulation does not depend on the ion's initial motional states.

to heat. We then measure the ion's temperature as a function of the length of this delay. Using this technique, we have measured the heating rate on a single-ion 2.8 MHz radial mode to be 77 quanta per second (figure 6.8b). This measurement is subject to the same uncertainties as the temperature measurements described in section 6.3, and so is only expected to be accurate to a factor of two.

A two-qubit gate would use the two-ion rocking mode since this is usually found to have a heating rate an order of magnitude or more lower than the single-ion mode [KWM⁺98], [WOC⁺12]. However, we have not yet measured this rate for our experiment.

At our current power levels, we are limited to a sideband Rabi frequency of $\Omega \simeq 2\pi \times 100$ Hz, giving a gate time of order 5.6 ms. With our measured heating rate, this should be sufficient to perform a two-qubit gate with $\sim 80\%$ fidelity. By using the the rocking mode and increasing our Rabi frequency, we hope to achieve an order of magnitude decrease in $\dot{n}t_g$, which would allow a gate fidelity of 97 %. Based on these simulations, we tentatively conclude that heating is unlikely to be a significant limitation in these experiments.

6.5 Trapping RF fluctuations

A second source of motional dephasing arises from amplitude fluctuations in the trapping RF voltage (see section 3.5). We model this using the master equation

$$\dot{\rho} = -\frac{i}{\hbar}[H_{\text{MS}}, \rho] - \frac{1}{\tau_c}[a^\dagger a, [a^\dagger a, \rho]] \quad (6.20)$$

where τ_c is the motional coherence time which, assuming the noise is white, is related to the single-sideband AM noise power spectral density, N_0^{AM} , by

$$\tau_c^{-1} = \frac{1}{2}\omega_r^2 \frac{N_0^{\text{AM}}}{P_c} \quad (6.21)$$

Here P_c is the power in the RF carrier and ω_r is the motional secular frequency in angular units. Once again, we integrate this master equation for a single-loop gate, the results of which are shown in figure 6.9.




Figure 6.9: Gate errors due to motional dephasing for a single-loop Mølmer-Sørensen gate on two ions beginning in different thermal states.

We see that the gate becomes much more sensitive to motional dephasing as the ions' initial temperature increases. To understand this effect, we recall that the gate operates by driving the ions' wave packets around different paths in motional phase space. At the gate's end these paths ideally form closed loops, leaving the ions' spins and motion in a separable state. However, if the motional mode's phase does not remain coherent with the driving field then the loops will not close and the spins remain entangled with the motion. When we measure the resulting spin state we trace over the motion, leading to a gate error that depends on the degree of residual spin-motion entanglement.

To see how this effect scales with temperature, we consider an ion initially in a motional Fock state undergoing a displacement α . The resulting state is

$$|\alpha, n\rangle := \mathcal{D}(\alpha) |n\rangle \quad (6.22)$$

where $\mathcal{D}$ is the displacement operator, defined by

$$\mathcal{D}(\alpha) := \exp(\alpha a^\dagger - \alpha^* a) \quad (6.23)$$

The overlap between two such states is given by [Wün91]

$$\langle \alpha, n | \beta, m \rangle = \langle \alpha | \beta \rangle \sqrt{\frac{n!}{m!}} (\alpha^* - \beta^*)^{m-n} L_n^{m-n}[(\alpha - \beta)(\alpha^* - \beta^*)] \quad (6.24)$$




Figure 6.10: Simulated gate dynamics for a single-loop Mølmer-Sørensen gate on a pair of ions starting from thermal states with $\bar{n} = 5$. (a) shows a perfect gate without dephasing, while (b) shows a gate with $t_c = 10t_g$. We plot the probabilities that one or both ions undergo a spin flip, P_{1e} and P_{ee} , as well as the gate fidelity, F .

where $L_n^\nu(x)$ denotes the Laguerre polynomials and

$$\langle \alpha | \beta \rangle = \exp \left(-\frac{1}{2} |\alpha|^2 - \frac{1}{2} |\beta|^2 + \alpha^* \beta \right) \quad (6.25)$$

The overlap between a pair of displaced thermal states may thus be shown to be

$$\langle \alpha, \bar{n} | \beta, \bar{n} \rangle = \langle \alpha | \beta \rangle \exp \{ -\bar{n} [(\alpha - \beta)(\alpha^* - \beta^*)] \} \quad (6.26)$$

Thus, if at the end of the gate the ions' wave packets are displaced by $\pm \alpha$ we expect errors that scale as

$$\langle \alpha, \bar{n} | -\alpha, \bar{n} \rangle = \exp \left[-2 |\alpha|^2 (2\bar{n} + 1) \right] \quad (6.27)$$

so we expect the error to increase exponentially as the ions' temperature increases.

Figure 6.10 shows the gate dynamics for a pair of ions initially in a thermal state with $\bar{n} = 5$. We consider both an ideal gate without dephasing and a gate with a motional coherence time of $t_c = 10t_g$. Ideally, at the end of the gate, the two ions should be in a maximally entangled state so that the probability of measuring a single ion to have undergone a spin flip vanishes. Loss of entangle-

ment due to dephasing is visible in the non-zero single-ion spin-flip probability at $t = t_g$.

In practice, one wants the motional coherence time due to RF fluctuations to be longer than the heating time ($1/\dot{n}$). In our experiment, we would thus ideally like a motional coherence time of at least 100 ms. For a 3 MHz radial secular frequency, this corresponds to an AM noise power density of $\lesssim -133$ dBc/Hz. Given the previous discussion, if the ions begin in an initial thermal state with $\bar{n} = 5$ then our motional coherence time needs to be roughly an order of magnitude longer than this, corresponding to a noise density of $\lesssim -143$ dBc/Hz.

6.5.1 Measuring AM noise

Oscillator AM noise is typically measured using either a diode or a mixer to demodulate the noise sidebands from the carrier, producing a baseband signal which may be viewed on a low-frequency analyser [HP174] [LFP98]. We chose to use a diode as the demodulator since we had a suitable detector available⁵, however we note that mixers often provide a better measurement noise floor (see, e.g. [Wal97]).

Measurement proceeds using the setup illustrated in figure 6.11. We first calibrate the apparatus by using a directional coupler to superpose a second low-level signal (typically -50 dBc to -100 dBc) onto the RF signal we wish to characterise (the “carrier” hereafter). The frequency of this signal is chosen to be close to the carrier frequency. At this low power level, the second signal may be thought of as a pair of equal power upper and lower AM and PM sidebands of the carrier⁶, with the power in each of the upper and lower AM sidebands

⁵Agilent 423B low-barrier Schottky detector. Negative polarity response, with a maximum input power of 200 mW.

⁶Of course, only one sideband – the upper or lower, depending on whether the calibration source is set to a higher or lower frequency than the carrier – is seen when the signal is viewed on a spectrum analyser because the AM and PM sidebands interfere destructively for the other sideband. Note that while the remaining AM and PM sidebands are both 6 dB lower than the total signal power, they are in phase with each other and so their voltages and not powers add to give the overall signal. See appendix C for a discussion of these and related points.




Figure 6.11: Test setup used to measure the AM noise of a signal source (DUT). A directional coupler is used to superpose a second, heavily attenuated signal onto the DUT's output for calibration purposes. The combined signal is then connected either to a spectrum analyser (again, for calibration purposes) or to a diode detector (we use an Agilent 423B low-barrier Schottky detector). The baseband signal from the diode detector is then filtered to remove any residual carrier signal before being passed to a low-noise pre-amp (LNA). Finally, the amplified base-band signal is viewed on a suitable signal analyser.

being 6 dB lower than the total power in the superposed signal.

We now view the carrier and superposed signal on a spectrum analyser and measure both the carrier power and the total (AM and PM) power in the superposed signal as a function of the power set on the calibration source. Subtracting 6 dB from this gives us a calibration factor, C_1 , relating the power set on the calibration source to the equivalent single-sideband AM power. Thus, for example, if we measure the power in the carrier to be 10 dBm and if setting 0 dBm on the calibration source produces a sideband with a total power of -60 dBm then $C_1 = 76$ dB. We note that it is important to ensure that the comparatively large carrier power does not saturate the spectrum analyser, leading to errors when measuring the sideband power.

Now that we are able to inject a pair of AM sidebands with known powers onto the carrier, we replace the spectrum analyser with the diode detector. The detector's output is filtered to remove any residual carrier signal before being passed through a high-gain low-noise amplifier and viewed on an low-

frequency signal analyser, such as an FFT analyser. To complete the calibration, we again look at injected sidebands using the calibration source. We measure a second coupling factor, C_2 , relating the power set on the calibration source to the measured baseband power. We define this coupling factor so that if setting 0 dBm on the calibration source produces a baseband power of +10 dBm then $C_2 = 10$ dB.

By combining these coupling factors, one may simply relate the measured baseband power spectrum, $P(f)$, to the single-sideband AM noise power at an offset f from the carrier, $N_0^{AM}(f)$, as follows⁷

$$10 \log_{10} \left(\frac{N_0^{AM}(f)}{P_c} \right) = 10 \log_{10} \left(\frac{P(f)}{P_c} \right) - C_1 - C_2 \quad (6.28)$$

During the calibration, it is important to establish that the detector has been biased into its linear regime (this typically takes around 10 dBm of RF power). This may be checked by confirming that the measured baseband sideband power changes linearly with the power set on the calibration source. One must also confirm that the circuit has sufficient bandwidth by checking that the measured baseband power is not a function of frequency in the bandwidth of interest. It is also important to establish that neither the LNA or analyser are in saturation, for example due to residual carrier signal or other spurious signals.

To ensure that the LNA does not degrade the system's noise floor, it is important that it be well matched to the diode's output impedance and signal level. To determine these quantities, we applied varying levels of 40 MHz RF power to the detector and measured its DC output voltage under loads of both 1 M Ω and 1 k Ω . The results of these measurements are shown in figure 6.12. We see that the diode response is relatively linear over the power range considered, reaching -1.5 V at an input power of 17 dBm. The diode's output impedance is

⁷Notice, that we have assumed that the upper and lower AM noise sidebands add in phase at the detector's output (as they do for the calibration source). This will often, but not always be the case. If they do not add coherently (so that the base-band powers, rather than voltages add) then the measured noise power is 3 dB higher than given by this formula.




Figure 6.12: Properties of Agilent 423 low-barrier Schottky diode detector as a function of input power. (a) shows the diode's output voltage when driving a $1\text{ M}\Omega$ load. At the $+17\text{ dBm}$ input power used during these measurements, the diode's output voltage squared is linear in the input power and the D.C. signal level is -1.5 V . (b) shows results of a rough measurement of the diode's output impedance made by comparing the detector's output voltages when driving $1\text{ k}\Omega$ and $1\text{ M}\Omega$ loads.

around $600\text{ }\Omega$ at this power level. We note, however, that without knowing how the diode is biased one cannot provide a noise model for the detector.

We used a pre-amp based on an AD8099 op-amp in the non-inverting configuration. The amplifier has input and output impedances of $10\text{ k}\Omega$ and $50\text{ }\Omega$ respectively (both AC coupled), 3 dB bandwidths of 0.5 Hz and 270 kHz and a gain of 500 into a $50\text{ }\Omega$ load. It has equivalent input voltage and current noise densities of $1.3\text{ nV}/\sqrt{\text{Hz}}$ and $2.5\text{ pA}/\sqrt{\text{Hz}}$ respectively.

The overall noise floor of this system has not been determined yet, although the measurements presented below provide an upper limit. We note that a lower noise floor may be achievable by either replacing the diode with a mixer or using a two-channel cross-correlation technique if needed (see e.g. [WCF⁺88], [Wal97], [Wal92], and [SFF99]).




Figure 6.13: Single-sideband AM noise power densities for various signal sources. We show the measured AM noise spectra for three signal generators: an HP8656B frequency synthesiser; an AD9910 DDS; and, a homebrew limiter built from 7400 series logic (see text for description). All three of these measurements were made using a Motorola CA2810C amplifier after to boost the signal levels to the measurement's 17 dBm power level. The right-hand ordinate gives the motional coherence time for a 3 MHz mode that would be obtained with *white* amplitude noise at the level given by the left-hand ordinate. Spurs in the measurements are thought to be due to ground-loops and base-band pick-up in the measurement circuitry. We also plot the noise for a Wenzel Associates Inc. 40 MHz ultra-low noise crystal oscillator. This data is an estimate provided by the manufacturer.

6.5.2 AM noise of various sources

Using the above techniques, we have measured the AM noise spectra of various signal sources, which are shown in figure 6.13. Since the signal sources we considered were not capable of producing the 17 dBm output power required for the measurement, we amplified them using a Motorola CA2810C RF amplifier. This amplifies the signals to approximately 26 dBm (roughly the RF power used in our experiments), from where they are reduced to 17 dBm for the measurement using a fixed attenuator. The graph also shows the manufacture's estimate of the AM noise spectrum for a Wenzel Associates Inc. 40 MHz ultra-low noise

crystal oscillator.

We expect to be most sensitive to noise on a similar timescale to the gate duration which, using our current parameters, will be of order 5 ms. This gives a characteristic frequency of 200 Hz. At this frequency the noise spectra for all of these sources are far from white so one must treat our previous analysis with caution. Nonetheless, we expect modelling the noise as white noise with a power density equal to the measured density at 200 Hz to provide an order-of-magnitude estimate of the dephasing error.

With this caveat in mind, one sees that the coherence time for the HP8656B synth and DDS will only be of order 10 ms, which will not be sufficient to perform a gate – particularly given the ion temperatures we can currently achieve.

The Wenzel oscillator, on the other hand, has an extremely low noise density, even on the long timescales we currently intend to work at. Assuming the oscillator meets the manufacturer’s expectations, it seems sufficient to eliminate motional dephasing due to RF noise.

The disadvantage of this oscillator is that it is not frequency agile. This is inconvenient since we require the frequency of our RF source to be kept within ~ 1 kHz of the frequency of our resonator (see section 6.7). To achieve this using a fixed frequency oscillator, one would need to tune the resonator’s frequency to match the oscillator. This could be done, for example, using a varactor to alter the circuit capacitance. This should be relatively easy to accomplish in practice.

6.5.3 7400 squarer

An alternative to using a fixed-frequency oscillator is to attempt to lower the noise spectrum of a frequency agile source. One could go about this in several ways. One possibility would be to use a detector to coherently measure the signal’s AM noise, which could then be removed using a feed-forwards amplifier [HNN06]. This might be particularly effective for the close-in AM noise that we




Figure 6.14: Measured single-sideband AM noise power density of the 7400 squarer driven by an HP8656B synth at 37.1 MHz. The blue curve shows a direct measurement of the 1.8 dBm RF output from the squarer, while the black curve shows a measurement made using the CA2810C amplifier followed by a fixed attenuator to boost the signal to 17 dBm. Spurs are thought to be due to pick-up and ground loops in the baseband circuitry and are hence significantly worse for the lower power measurement.

are sensitive to.

A second possibility is to use a non-linear element to “clip” the input signal, producing a constant amplitude square-wave. Once the harmonics have been filtered from this square-wave, a signal at the original frequency is recovered with the AM noise (ideally) removed. Such limiters can be passive – typically one or more diodes with appropriate biasing – or active, such as an amplifier designed to run in heavy saturation.

In principle, good results should be obtainable with any of these methods. We chose to investigate the use of a 7400-series inverter as an extremely simple – simple enough for a physicist to design – limiting amplifier, using the circuit described in section D.1. At 40 MHz, the amplifier puts out approximately 3 dBm

of power when driven with an input power of $\gtrsim -10$ dBm.

We measured the device's AM noise when driven from an HP8656B synth at 37.1 MHz both with and without the CA2810C amplifier. The results of these measurements are shown in figure 6.14. Care must be taken in interpreting these results since the RF power levels present at the detector differ significantly between the two measurements. Since the noise properties of the measurement setup have not been characterised yet, a direct comparison is not possible.

As a preliminary assessment of the device's performance, however, we note that the squarer appears to reduce the synth's noise by more than 10 dB across the frequency range. We also note that the noise density measured using the amplifier is larger than the noise measured without it. A possible interpretation of this fact is that the amplifier's close-in AM noise has become the dominant limitation in this measurement. This tentative conclusion is supported by the fact that the noise appears to be more or less white without the amplifier, but gains a $1/f$ -type characteristic with it.

6.5.4 Diode limiter

We had a MiniCircuits VLM-52+ limiter conveniently to hand in the lab and so decided to investigate its AM suppression. This limiter clips input signals to ~ 10 dBm, with an output change of 0.2 dB per 1 dB change in input power. The use of this limiter to reduce AM noise is illustrated in figure 6.15, where we see the limiter reducing the noise at the output of an amplifier by roughly 10 dB. A diode limiter of this kind may provide a convenient means of reducing the AM noise to an acceptable level.

6.5.5 Close-in amplifier noise

The previous investigation suggests that close-in ($1/f$) amplifier AM noise may be a significant noise source. To investigate this, we measured the noise of




Figure 6.15: Single-sideband AM noise power densities measured using the 7400 squarer as well as various amplifiers and limiters. See text for details.

the 7400 squarer when amplified by both the Motorola CA2810C and a MiniCircuits ZX60-33-LN-S+ amplifier. The results of these measurements are shown in figure 6.15. We see that the MiniCircuits amplifier significantly degrades the squarer's performance.

To ensure that the amplifier does not degrade the system's noise, it may be important to use an RF amplifier with good close-in noise properties⁸. In practice, this is likely to require the use of a low phase-noise amplifier⁹. Given

⁸RF amplifiers are often specified in terms of a noise figure, which indicates the amplifier's broadband (white) noise far from the carrier. Given the long time scales we require RF amplitude coherence on, we are more interested in the noise around the carrier, which typically increases as $1/f$ or faster. The noise figure does not tell one anything about the amplifier's behaviour in this regime. Further to this, the noise figure is typically measured in the absence of a carrier; while the noise is usually a function of the signal level due to non-linear intermodulation processes. The noise figure is thus not even a reliable specification of the broadband noise. For situations where one cares about the noise, the amplifier noise spectrum must be measured under the same conditions it will be operated in. See e.g. [HHWW06] for further discussion of these points.

⁹Technically, we want a low AM noise amplifier, but amplifiers are very rarely specified in terms of AM noise. In practice, amplifier noise tends to be uncorrelated with the carrier and so is composed of equal PM and AM components. The PM specifications are thus usually a good indication of AM performance.

our frequency range and power level, one possible amplifier is a Spectrum Microwave BX3110¹⁰. The manufacturers of this amplifier have measured its noise properties using a 160 MHz input signal with the amplifier running at its 1 dB compression point. The results of this measurement are shown in figure 6.15. The noise from this amplifier should be sufficient to ensure that it does not present an experimental limitation in any work we intend to conduct.

6.6 Motional mode cross-coupling

A further source of motional dephasing arises from the cross-coupling between the ions' "stretch" motion along $\hat{x}$ and the radial rocking mode that we will use for our gate. This coupling stems from the non-linearity of the Coulomb interaction [NRJ09]. For a pair of ions in a symmetric linear trap characterised by axial and radial frequencies ω_x & $\omega_\perp$ ($\omega_x < \omega_\perp$), the coupling takes the form

$$H = \hbar\chi a^\dagger a b^\dagger b \quad (6.29)$$

where a and b are the annihilation operators for the rocking and stretched modes respectively, and

$$\chi := -\omega_s \left(\frac{2\hbar\omega_x}{\alpha^2 m c^2} \right)^{1/3} \left(\frac{\omega_x}{\omega_r} \right) \left[\frac{1}{2} + \frac{\omega_s^2/2}{4\omega_r^2 - \omega_s^2} \right] \quad (6.30)$$

Here, ω_r and ω_s are the frequencies of rocking and stretched modes, which are given by

$$\begin{aligned} \omega_r &= \sqrt{\omega_\perp^2 - \omega_x^2} \\ \omega_s &= \sqrt{3}\omega_x \end{aligned} \quad (6.31)$$

while m is the mass of each ion and α the fine structure constant.

This Hamiltonian will apply to our system with only a small correction due to the DC "tilt" quadrupole. For a trap with radial frequencies of 3 MHz and axial frequencies of 475 kHz, $\hbar\chi = -0.8$ rad/s. To estimate the size of the resulting

¹⁰Low phase noise 10 MHz - 200 MHz RF amplifier. 15.5 dB gain, 30 dBm @ 1 dB compression.




Figure 6.16: Simulated gate errors due to constant detunings during (a) one- and (b) two-loop gates. In (b) dashed lines indicate the use of Walsh-modulation of the Hamiltonian (in this case, equivalent to changing the Hamiltonian’s sign half-way through the gate). The abscissa gives the detuning offset scaled by the single-ion sideband Rabi frequency.

coupling, we assume that the stretch mode is in a thermal state with

$$\langle b^\dagger b \rangle \simeq \frac{\omega_r}{\omega_x} \langle a^\dagger a \rangle \simeq 20 \quad (6.32)$$

We thus expect a coupling of order 15 Hz. This is unlikely to be the dominant limitation in our current work, but could be significant in future studies, requiring sub-Doppler cooling of the stretch motion.

6.7 Secular frequency drift

If the microwave fields are not resonant with the blue and red motional sideband transitions then the ions’ spins and motion will remain entangled at the end of the gate, leading to errors as discussed in section 6.5. The magnitude of this error may be reduced by modulating the gate Hamiltonian’s sign according to a Walsh function [HCD⁺12]. This may be accomplished, for example, by changing either of the gate phases, $\bar{\phi}$ or $\delta\phi$, (section 2.7) by π during the gate or by performing single-qubit π -pulses on the two ions.

Simulations of gate errors due to constant detuning offsets are shown in figure 6.16. We simulate both a single-loop gate without Walsh modulation and a two-loop gate with and without modulation. Errors for the single-loop gate and two-loop gate without modulation are similar, while the modulation provides a significant decrease in the errors for small detuning offsets.

Based on these simulations we expect to be able to achieve a gate fidelity of $\sim 90\%$ even with detuning errors as large as $\lesssim 0.2 \times \Omega_1$. Given our 100 Hz sideband Rabi frequency, this is equivalent to an absolute detuning error of $\lesssim 20$ Hz. Over long timescales, the microwave frequency may be servoed to track drifts in the secular frequency. However, this is only likely to be practical on timescales of a few minutes, setting a limit on the tolerable drift rate to $\lesssim 10$ Hz/minute.

Further robustness could be obtained by using more loops and a higher-order Walsh function; however, this would increase the gate time, making it more susceptible to dephasing and heating. Moreover, if one imposes a limit on the microwave duty cycle to avoid damaging the trap then increasing the gate detuning will slow down the experimental repetition rate, potentially making drifts worse.

The secular frequency is proportional to the trapping RF voltage amplitude so, if we are to achieve a mode stability of $\lesssim 10$ Hz using a 3 MHz secular frequency, this amplitude must have a fractional stability of $\lesssim 3$ ppm. There are two main sources of RF amplitude drift. Firstly, the power entering the resonator may drift due, for example, to changes in the power produced by the signal source or the gain of the amplifier. Secondly, the impedances in the RF resonant circuit may change, altering the voltage produced on the trap's electrodes. We consider these in turn in the next sub-sections.

6.7.1 RF power drifts

When we first setup the experiment, we produced the trapping RF using a commercial frequency synthesiser and a Motorola CA2810C amplifier. We observed drifts in the radial modes of a single $^{43}\text{Ca}^+$ ion in a 4 MHz trap of order 1 kHz, which were strongly correlated with our laboratory air conditioning cycle [All11]. To investigate the source of these drifts, we logged the RF power output from the synth over a period of 24 hours. We measured power fluctuations of order 5×10^{-4} ; which is of an appropriate order of magnitude to explain the observed drifts.

One solution to this problem might be to use a temperature stabilised signal source and amplifier to attempt to reduce these drifts to an acceptable level. A DDS might be a particularly convenient source for this since its compact nature makes it easier to temperature stabilise. The disadvantage of this approach is that the entire RF chain must be maintained at a highly stable temperature, which is inconvenient. Instead, we opted to build a levelling loop, based around a commercially available ADL5511 RF power detector.

The levelling loop, described in detail in appendix D, uses a directional coupler to pick-off a sample of the RF power, which is then measured by the detector. A PI loop then creates an error signal, which is fed back to a voltage variable attenuator after the synth. The directional coupler, detector and PI loop are contained within a small box that is maintained at a constant temperature to within 1 mK. The temperature coefficient of this detector was measured to be 170 ppm/ $^{\circ}\text{C}$, which we hope will be sufficient to allow a radial mode stability of $\lesssim 10$ Hz.

6.7.2 Resonator drifts

We consider the simplified lumped-element model of the RF resonant circuit shown in figure 6.17. The capacitor in this model represents the trap as well




Figure 6.17: Simplified model of the trap RF circuit as a series RLC circuit inductively coupled to a $50\ \Omega$ signal source.

as other capacitances due, for example, to the SMA cables and diplexer PCB. The voltage across this capacitor is equal to the trapping voltage. The secular frequency, $f_{\perp}$, is proportional to this voltage so that

$$\frac{1}{f_{\perp}} \frac{df_{\perp}}{dt} = \frac{1}{V} \frac{dV}{dt} \quad (6.33)$$

6.7.2.1 Resistive losses

We begin by considering changes to the circuit's loss, R . For a given input RF power, P , coupled into the circuit, the trapping voltage is given by $V = \sqrt{PR}$, so that

$$\frac{df_{\perp}}{dt} = \frac{1}{2} f_{\perp} \times \frac{1}{R} \frac{dR}{dt} \simeq (1.5 \times 10^6 \text{ Hz}) \times \frac{1}{R} \frac{dR}{dt} \quad (6.34)$$

where we have assumed the secular frequency to be 3 MHz, which is typical for these experiments.

The losses in the circuit arise from loss in the copper and gold conductors as well as core losses in the matching transformer, with comparable losses in the core and metals (dielectric losses are expected to be small). Both copper and gold have temperature coefficients of electrical resistivity of $\sim 4000 \text{ ppm}/^{\circ}\text{C}$ at 400 K (see, e.g. [HL11]), corresponding to a frequency change of $5 \text{ kHz}/^{\circ}\text{C}$ ¹¹. The temperature dependence of the core loss is harder to calculate.

¹¹We note that a similar effect will also arise due to losses in the cabling connecting the RF amplifier to the matching transformer. However, since the voltage step up of the matching network is $\gtrsim 10$, we are more than two orders of magnitude more sensitive to loss inside the resonant circuit than outside it.

Changes in the loss will also alter the amount of power coupled into the circuit. The power coupled from a transmission line of (real) characteristic impedance R_0 terminated into a load R_L is given by (assuming the generator is matched to the line) [Poz05]

$$P = P_0 (1 - |\Gamma|^2) \quad (6.35)$$

where Γ is the reflection coefficient, given by

$$\Gamma = \frac{R_L - R_0}{R_L + R_0} \quad (6.36)$$

so that¹² for $R_L \simeq Z_0$

$$\frac{1}{P} \frac{dP}{dT} \simeq -\Gamma \frac{dR_L}{dT} \quad (6.37)$$

combining these equations, we find that for a well matched line

$$\frac{df_{\perp}}{dt} \simeq -\frac{1}{2} f_{\perp} \Gamma \times \frac{1}{R_L} \frac{dR_L}{dt} \quad (6.38)$$

This is essentially the same as equation (6.34) except for the factor of Γ . If the resonator is well matched to 50Ω , we expect $\Gamma > 10$ so this expression does not contribute to the circuit's temperature dependence significantly.

6.7.2.2 Complex impedances

In general, we expect the complex impedances to have relatively low temperature coefficients compared with the resistive losses discussed above. For example, the specified temperature coefficient of the Micrometals number 10 material used in the transformer is 150 ppm/°C. However, since the circuit has a relatively high quality factor, the power coupled into the circuit may change significantly even for small changes to L or C .

The resonance frequency of an LC circuit is given by

$$f_0 = \frac{1}{2\pi\sqrt{LC}} \quad (6.39)$$

¹²Here, we assume that $R_L > R_0$. If the opposite is true then the sign of this expression changes.

so that

$$\frac{1}{f_0} \frac{df_0}{dT} = -\frac{1}{2} \frac{1}{L} \frac{dL}{dT} \quad (6.40)$$

and similarly for changes to C . For the 150 ppm/°C temperature coefficient of L with $f_0 = 40$ MHz, this corresponds to a change in resonance frequency of 3 kHz/°C.

Assuming the circuit has a Lorentzian lineshape, we expect

$$V(f) = \frac{V_0}{\sqrt{1 + (f - f_0)^2 / \gamma^2}} \quad (6.41)$$

where γ is the circuits 3 dB frequency which, in our case, is 380 kHz. The temperature coefficient of the coupled power is thus

$$\frac{1}{V_0} \frac{dV_0}{dT} = -\frac{\delta f / \gamma^2}{1 + (\delta f / \gamma)^2} \times \frac{df_0}{dT} \simeq -\delta f / \gamma^2 \times \frac{df_0}{dT} \quad (6.42)$$

where $\delta f := f - f_0$. We thus find that

$$\frac{df_{\perp}}{dT} = 0.06 \delta f \text{ } ^\circ\text{C}^{-1} \quad (6.43)$$

So long as δf is of order 1 kHz, this effect should be negligible compared with changes in the circuit's losses. To ensure that this is the case, we servo the RF frequency to match the circuit's resonance. We do this by measuring the power coupled into the resonant circuit using a power detector¹³ connected to one of the DC input ports on the diplexer. A small amount ($\sim 100 \mu\text{W}$) of RF power is coupled out of this port due to imperfect filtering by the diplexer. By monitoring the amount of coupled power as the RF frequency is changed we are able to find the circuit's resonance frequency.

Since the pickoff is itself frequency dependent, the frequency at which the measured power is greatest does not correspond exactly to the circuit's resonance. To measure this offset, we measured the ion's secular frequency as a function of RF frequency. We find that the maximum secular frequency occurs at an RF frequency 9 kHz higher than the maximum measured power.

¹³MiniCircuits ZX47-40-LN-S+




Figure 6.18

6.7.3 Measured drift rate

To attempt to minimise the secular frequency drift, we installed the RF levelling loop and temperature stabilised the parts of the RF resonant circuit outside the vacuum system (the diplexer and matching transformer). We then measured the drift of the radial secular frequency of a single $^{43}\text{Ca}^+$ ion using sideband spectroscopy while recording both the lab air temperature and the temperature of the vacuum system. The results of this are shown in figure 6.18.

The drift rate over a period of 5 minute appears to be of order 20 Hz which is around the limit of what we can tolerate. It is unclear from the present data whether the drifts are correlated with the air/vacuum system temperature or

not.

6.8 Concluding remarks

We have considered several error mechanisms that could limit the two-qubit gate fidelity we are ultimately able to achieve. The most fundamental of these was the heating rate. Based on the heating rate alone, we would expect to be able to achieve two-qubit gate fidelities of $\gtrsim 90\%$ even with our current low sideband Rabi frequencies; with the fidelity depending on to what extent the heating rate is suppressed on the stretch mode.

The other sources of error were all to do with the stability of the ions' motional mode frequency. Given our current high thermal states, we are extremely sensitive to these errors. It thus seems that sub-Doppler cooling using either microwaves or Raman lasers is likely to be essential before the fidelity limit set by the heating rate can be reached. With the ions cooled close to their motional ground state, errors due to stretch-rocking mode cross-coupling could be made insignificant.

It thus seems that factors that are most likely to prevent us from reaching a heating-rate-limited fidelity are drift and jitter of the radial mode frequency. However, we have constructed preliminary apparatus that appears to significantly reduce these errors.

7

Conclusion

In this thesis we have presented the experimental realisation of an intermediate-field clock qubit in $^{43}\text{Ca}^+$ at 146 G, as well as progress towards a microwave-driven two-qubit gate on this qubit. We have demonstrated that, despite its relatively complex level structure, $^{43}\text{Ca}^+$ may be efficiently laser-cooled in the intermediate field, producing both a high rate of fluorescence and temperatures close to the Doppler limit. We showed that the qubit may be used as a quantum memory with minute-long coherence time, as well as demonstrating the fundamental single-qubit operations of state-preparation, readout and coherent rotations with high fidelities.

A significant feature of this work is that all of these operations may be accomplished in a single, room-temperature experiment using a scalable trap design, while still producing state-of-the-art fidelities. Moreover, the techniques we use are fundamentally robust and do not rely on narrow-linewidth lasers, extensive magnetic field screening or the use of dynamical decoupling techniques.

Using microwave magnetic-field gradients, we have driven motional sideband transitions on a single ion, demonstrating the spin-motion coupling that underlies the two-qubit gate. We have provided an analysis of likely sources of experimental error during a future two-qubit gate and have designed apparatus to minimise these errors. Using our currently attainable Rabi frequencies, we hope to demonstrate a gate with a fidelity of $\gtrsim 80\%$ in the near future. By oper-

ating the trap at its full design power, we hope to achieve Rabi frequencies close to 1 kHz, which should allow a significant improvement in fidelity.

7.1 Comparison with other results

The field of experimental quantum computing has grown to encompass a large number of physical systems each with its own benefits and drawbacks. Given the extreme diversity of these systems and the different challenges they face, it would not be possible to compare our current work with all other proposed implementations of QIP here. Instead, we will focus on the two implementations which the author believes to offer the most potential for the realisation of a large-scale quantum information processor: trapped atomic ions and superconducting circuits. Recent reviews of these fields may be found, for example, in [MK13], [WL11] and [DS13].

One must take care in comparing even these two relatively similar fields since the challenges they face – in particular in terms of what must be done to turn currently realisable devices into a large scale information processor (see [DS13] for a good discussion of this point) – and the relevant measures of performance are somewhat different. With these difficulties in mind, however, we have summarised what we hope is a representative selection of results for these two systems in table 7.1.

A few caveats should be mentioned. We begin by noting that gate times differ significantly between trapped ion and superconducting circuit QIP. Typical single-qubit gate times are 1 μ s-10 μ s for ions and 5 ns-50 ns for circuits. When comparing coherence times the most useful figure of merit is perhaps the ratio of the coherence time to the gate time. For circuits, this is of order one thousand, while for our system it is 5 million.

A higher fidelity for preparation and measurement of an optical qubit in calcium has previously been reported by our group [MSW⁺08]. However, for

Operation	This work	Other hyperfine ion qubit	Superconducting circuit
SPAM	$6.8(6) \times 10^{-4}$	$8.5(7) \times 10^{-4}$ [NVG ⁺ 13]	2×10^{-2} [HSM ⁺ 13]
Coherence time	50 (10) s	10 s [LOJ ⁺ 05]	100 μ s [SPC ⁺ 12] [RGP ⁺ 12]
Single-qubit gate	1×10^{-6} in 10 μ s	$2.0(2) \times 10^{-5}$ in 20 μ s [BWC ⁺ 11]	8×10^{-4} in 10 ns [BKM ⁺ 14]

Table 7.1: State of the art errors achieved for fundamental single-qubit operations in both ion-trap and superconducting circuit QIP. See text for caveats to these numbers. We also list the approximate durations of single-qubit logic gates to allow comparison with coherence time.

the reasons discussed in section 2.2, we agree with Steane [Ste97] that “[t]he choice of RF rather than optical transitions for information processing appears so advantageous as to be forced upon us.”. Thus, while we acknowledge the exquisite results that have been achieved using optical qubits (see, for example [BKRB08b]), we exclude this result from our present comparison.

We note that coherence times exceeding 10 minutes have been achieved using trapped ions [FSL⁺95] [BHI⁺91], however these experiments have generally not been suitable for quantum information processing. The work of [FSL⁺95], used the $|F = 0, m_F = 0\rangle \Leftrightarrow |F = 1, m_F = 0\rangle$ clock transition in $^{171}\text{Yb}^+$ at zero field. As discussed in section 2.2, for this ion to be used as a qubit, a non-zero magnetic field must be applied to introduce a splitting between hyperfine states. This field would introduce a first-order field dependence to the system, significantly reducing the coherence time.

The work of [BHI⁺91] used the 303 MHz $|m_I = -1/2, m_J = 1/2\rangle \Leftrightarrow |m_I = -3/2, m_J = 1/2\rangle$ nuclear transition in $^9\text{Be}^+$ at 0.8194 T. This experiment used a cloud of 5 000 – 10, 000 $^9\text{Be}^+$ ions confined in a Penning trap with a further cloud of 50 000 – 150 000 $^{26}\text{Mg}^+$ ions, which were used for sympathetic cooling. We exclude this result since it involved a cloud of ions and not individual qubits. We further note that at this field the nuclear and electronic spins are effectively decoupled so that the clock-transition may not be driven via optical Raman transitions [Lan06]. Thus, this transition will not be suitable for use as a qubit until two-qubit gates and single-ion addressing can be performed with

high-fidelity using RF fields (see below). We note that even if microwave/RF techniques do achieve comparable fidelities to their laser-based counterparts, the superconducting magnets required to reach this field would still present an experimental inconvenience.

Finally, we re-iterate the fact that we are comparing our one experiment with an entire field. These results of table 7.1 rely on somewhat contradictory properties of a system. For example, in superconducting qubits state-preparation and readout require circuit quality factors that are several orders of magnitude lower than those used for long-term memories. Combining these two elements remains one of the large challenges facing this field [DS13].

7.2 Outlook

Large-scale quantum information processing will require the techniques of quantum error correction to prevent accumulation of errors and ensure that calculations have a high probability of success. For error correction to be applied, the fidelity of each operation must exceed a threshold known as the “fault-tolerant limit”. Exactly where this limit lies depends heavily on the details of the error correction algorithm used and the amount of overhead one is prepared to tolerate. A commonly used bound is 10^{-3} infidelity for single-qubit operations and 10^{-4} for two-qubit gates [Ste03].

The single-qubit operations demonstrated in this work exceed this limit and may thus be considered suitable for fault-tolerant quantum computation. Since the overheads required by error correction algorithms typically decrease rapidly as the fidelity of each operation is increased, it is likely to be important for fidelities to exceed this limit by an order of magnitude or more. It is thus significant that the single-qubit gate demonstrated in this work exceed this estimate of the fault-tolerant limit by three orders of magnitude.

If the techniques developed in this thesis are to be applied to the construction

of a large scale information processor then two limitations must be overcome. Firstly, we must successfully demonstrate a high-fidelity two-qubit gate to go alongside our single-qubit operations. Two-qubit gates using microwaves have already been demonstrated [OWC⁺11], however only with a fidelity of 76 % – well below the 99.99 % fidelity likely to be required for large-scale quantum computation [Ste03].

We hope that the work done in this thesis will contribute to the closing of this daunting gap and it seems that there is much scope for improvement on this front. Indeed, there is no *in principle* reason why errors exceeding this threshold should not be achieved in future work using achievable experimental parameters. A cynical experimentalist, however, may find that of little comfort and it remains an open question as to what will ultimately limit future work.

Looking to the future, our group has already begun work on the next iteration of this experiment which, we hope, will produce significantly higher fidelities. In our current experiment, the fidelity of the gate is expected to be limited by the amount of microwave power we can apply to the trap without damage or causing excessive change in the motional mode frequencies due to thermal effects. In the next experiment, we intend to tackle this problem using several means. Firstly, we will move the ion closer to the trap's surface ($\sim 50 \mu\text{m}$), increasing the field gradients produced by a given microwave power.

Secondly, we will cool the trap close to 77 K using liquid nitrogen. At this low temperature, the electrical and thermal resistivities of the trap electrodes will be greatly decreased, allowing both more power to be applied to the trap and a greater field gradient to be achieved for a given power. In addition, we expect the heating rate to be decreased by an order of magnitude [CS13], which should allow a further reduction in errors. Finally, we are considering moving to the 288 G clock qubit. This qubit is based on a π -polarized transition, which would couple to the gradients produced by our trap more effectively (see section

2.7).

If technical limitations prevent microwave-driven gates from achieving the required fidelities then laser-based schemes offer an excellent alternative. The clock-qubit we have developed is ideally suited to Mølmer-Sørensen gates driven using stimulated Raman transitions. In other work by our group, a geometric phase (“wobble”) gate [LDM⁺03] using the stretched qubit transition has recently been demonstrated on $^{43}\text{Ca}^+$ at a field of ~ 1 G in a macroscopic trap with a fidelity of $\gtrsim 99\%$. One of the limitations in this gate was dephasing due to magnetic field noise and thus better fidelities may be achievable in future experiments on the clock-qubit.

The second limitation to be overcome is microwave cross-talk. In a large scale device, many ions will be confined in close quarters and it is essential that operations may be performed on one ion without affecting its neighbours. Single-qubit addressing using microwaves with an error of 6×10^{-4} has already been demonstrated [WOC⁺13]. This work was conducted in a trap similar to the one developed in this thesis and took advantage of the strong microwave gradient to allow spatially separated ions to be addressed. This technique may end up being of limited use, however, since it is only suitable to linear arrays of ions and requires precise and stable positioning of the ions.

To overcome these limitations, our group has designed and constructed a prototype trap specifically designed to allow localised microwave fields to be produced [ALH⁺13]. The trap’s geometry is designed so that a field addressing one ion creates only a small ($\lesssim 1\%$) residual field at neighbouring ions. Further electrodes then allow “compensation currents” to be applied to reduce this field to a low level ($\sim 0.01\%$ for realistic microwave amplitude and phase fluctuations).

It remains an open question as to what technologies will ultimately prove most fruitful in the quest to build larger scale, more accurate trapped ion in-

formation processors – although, the author admits to having a soft spot for hyperfine qubits and microwaves. Many more lab hours will be required to understand how much we can really achieve which I for one look forward to immensely.

Appendices

A

Waveplates

The cooling (sections 2.5, 5.1 and 6.3), state-preparation (sections 2.6.1 and 5.2) and shelving (sections 2.6.3 and 5.3) techniques used in this thesis require laser beams with specific polarizations. These beams all initially pass through PBS cubes or Glan-Taylor prisms (section 4.5), producing light with a well-defined linear polarization. This is then transformed to the desired polarization state using half- and quarter-wave plates. In this appendix, we discuss how these waveplates should be set and what polarizations may be achieved.

We begin in section A.1 with a discussion of tilting quarter-wave plates to alter the degree of retardance they produce. This allows one to ensure that a particular waveplate operates as a quarter-wave plate at a chosen wavelength or to adjust the waveplate to compensate for small, unwanted birefringences in the beam path (for example due to vacuum viewports). This is particularly important for the state-preparation and readout beams, for which the fractional intensities of unwanted polarization components must be reduced to $\lesssim 0.1\%$. Next, in section A.2, we discuss how the performance of quarter-wave plates varies as a function of the wavelength of light used. This is significant in our experiments since both the 397 nm state-preparation and 393 nm shelving beams pass through the same polarization optics, which must consequently work well at both wavelengths.

In section A.3, we discuss the characterisation of quarter-wave plates using a polarimeter. We demonstrate that with appropriate tilting a $\gtrsim 100,000 : 1$ ratio of wanted to unwanted polarization components may be achieved. In the presence of a viewport, this ratio is reduced to $5,000 : 1$, suggesting that not all of the birefringence introduced by the viewport can be corrected by angle-tuning the quarter-wave plate.

Finally, in section A.4, we consider the polarization state of the 397 nm Doppler beam. To cool $^{43}\text{Ca}^+$ efficiently at 146 G this beam should contain as little σ^- light as possible (section 2.5). However, if it is to couple to all of the trap's normal modes then it must propagate at approximately 45° to the trap's symmetry axis and hence also at 45° to the static magnetic field, which defines our quantization axis. This limits our control over the beam's polarization. We show that, despite this, with appropriately set half- and quarter-wave plates, a polarization ratio of $2/3$ of σ^+ to $1/3$ of π may be produced.




Figure A.1: Typical use of a quarter-wave plate. A linearly polarized input beam enters normal to the plate's surface with its polarization vector at 45° to the plate's axis.

A.1 Tilting quarter-wave plates

We consider the quarter-wave plate illustrated in figure A.1. The input beam is polarized at 45° to the plate axis so that the plate introduces an optical path difference of

$$s = \Delta nt = (n_e - n_o) t \quad (\text{A.1})$$

between the e- and o-rays, where t is the plate's thickness and n_e and n_o are the refractive indices seen by the e- and o-rays. For the plate to operate as a quarter-wave plate at a wavelength λ_0 , one requires that $s = (p + 1/4)\lambda_0$ for integer p .

Suppose that one tilts the plate by an angle ϕ about the crystal axis, so that the beam enters at a small angle to the plate's surface, but still at right angles to the crystal axis. In this case, the refractive indices of both the e- and o-rays remain unchanged; however, the plate's thickness is effectively increased, so that

$$s_\phi = \frac{\Delta nt}{\cos \phi} = \frac{(p + 1/4) \lambda_0}{\cos \phi} \simeq (p + 1/4) \lambda_0 \times (1 + \phi^2/2) \quad (\text{A.2})$$

The waveplate now acts as a quarter-wave plate for a wavelength λ_ϕ such that

$$\lambda_\phi = \frac{\lambda_0}{\cos \phi} \simeq \lambda_0 (1 + \phi^2/2) > \lambda_0 \quad (\text{A.3})$$

If, on the other hand, one tilts the plate so that the light enters non-normal to the crystal axis then the o-ray no longer sees a refractive index of n_o , but rather

$$n_e(\phi) = \left(\frac{\cos^2 \phi}{n_o^2} + \frac{\sin^2 \phi}{n_e^2} \right)^{-2} \simeq n_o + \Delta n (1 - \phi^2) \quad (\text{A.4})$$

The optical path difference introduced by the plate is then

$$\begin{aligned} s_\phi &= \frac{n_e(\phi) - n_o}{\cos(\phi)} t \\ &\simeq \frac{\Delta n t (1 - \phi^2)}{1 - \phi^2/2} \\ &= (p + 1/4) \lambda_0 \times (1 - \phi^2/2) \end{aligned} \quad (\text{A.5})$$

The waveplate now acts as a quarter-wave plate for a wavelength

$$\lambda(\phi) = \lambda_0 (1 - \phi^2/2) < \lambda_0 \quad (\text{A.6})$$

Thus, one may either increase or decrease the wavelength for which the plate acts as a quarter-wave plate by tilting about an appropriate axis. We note that the previous results were independent of the plate's order, p .

In fact, the waveplates we use are “zero-order” waveplates, consisting of a pair of plates oriented at right angles, so that the extraordinary ray for a beam passing through the first becomes the ordinary in the second and *vice-versa*. The optical path difference in each plate is

$$\begin{aligned} s_1 &= \Delta n t_1 \\ s_2 &= -\Delta n t_2 \end{aligned} \quad (\text{A.7})$$

Using our previous expressions, we see that when the zero-order plate is tilted through a small angle about the axis of the first plate, the path differences become

$$\begin{aligned} s_1 &= \Delta n t_1 (1 + \phi^2/2) \\ s_2 &= -\Delta n t_2 (1 - \phi^2/2) \end{aligned} \quad (\text{A.8})$$

so that the total path difference is

$$\begin{aligned} S_\phi &:= s_1 + s_2 \\ &\simeq \Delta n \{ (t_1 - t_2) + \phi^2/2 (t_1 + t_2) \} \end{aligned} \quad (\text{A.9})$$

For the zero-order plate, we have

$$(p + 1/4) \lambda_0 = \Delta n (t_1 - t_2) \quad (\text{A.10})$$

for integer p . Thus, the tilted zero-order plate now operates as a quarter wave-plate for a wavelength

$$\lambda_\phi = \lambda_0 \times (1 + \alpha \phi^2) \quad (\text{A.11})$$

where

$$\alpha := \frac{t_1 + t_2}{2(t_1 - t_2)} \quad (\text{A.12})$$

This factor of α makes zero-order plates much more sensitive to angle tuning.

In these experiments, we used quartz zero-order plates of thickness $t_1 + t_2 \sim 1$ mm designed for 400 nm light so that $t_1 - t_2 = \lambda_0 / (4\Delta n) = 10 \mu\text{m}$ and $\alpha = 50$. Thus, tuning the waveplate by 3 nm to operate at 397 nm requires a tilt of roughly 0.7° .

A.2 Using quarter-wave plates at multiple wavelengths

We wish to know how well a waveplate that has been tuned to operate as a quarter-wave plate at one wavelength will perform at a different wavelength. In the following, we will assume that the waveplate is a two-slab zero-order plate as described previously. We assume that the waveplate introduces a phase shift of $\pi/2$ between E_x and E_y at a wavelength of λ_0 , so that

$$\frac{2\pi}{\lambda_0} \Delta t \Delta n = \frac{\pi}{2} \quad (\text{A.13})$$

where Δt is the difference in thickness between the two slabs forming the plate and Δn the difference in refractive indices for the plate's two principal directions.

At a wavelength $\lambda := \lambda_0 + \delta\lambda$, assuming that Δn does not change significantly between the two wavelengths, the plate introduces a phase shift of

$$\frac{2\pi}{\lambda_0 + \delta\lambda} \Delta t \Delta n = \frac{\pi}{2} + \theta \quad (\text{A.14})$$

so, assuming $\delta\lambda/\lambda \ll 1$

$$|\theta| \simeq \frac{\delta\lambda}{\lambda} \times \frac{\pi}{2} \quad (\text{A.15})$$

For $\lambda_0 = 397 \text{ nm}$ and $\lambda = 393 \text{ nm}$, this corresponds to a phase of 16 milli-radians.

The field at λ thus exits the waveplate with components

$$\begin{aligned} E_x &= a \cos(\omega t) \\ E_y &= a \sin(\omega t + \theta) \end{aligned} \quad (\text{A.16})$$

which we wish to express in terms of two circular components so that

$$\begin{aligned} E_x &= A \cos(\omega t + \phi_1) + B \cos(\omega t + \phi_2) \\ E_y &= A \sin(\omega t + \phi_1) - B \sin(\omega t + \phi_2) \end{aligned} \quad (\text{A.17})$$

After some algebra, one finds that

$$\begin{aligned} A^2 &= \frac{1}{2} (1 + \cos \theta) \simeq 1 - \frac{1}{4} \theta^2 \\ B^2 &= \frac{1}{2} (1 - \cos \theta) \simeq \frac{1}{4} \theta^2 \end{aligned} \quad (\text{A.18})$$

The intensity ratio of the two components is thus equal to $\theta^2/4$, which is 6×10^{-5} for the case where $\lambda_0 = 397 \text{ nm}$ and $\lambda = 393 \text{ nm}$. This error is insignificant for our experiments and we expect it to be swamped by polarization errors due to spatially-varying strain-induced birefringence in the vacuum viewports (see below).

A.3 Characterising waveplates

We investigated the polarization purity that may be achieved using our waveplates with the polarimeter shown in figure A.2. Here, a pair of nominally identical quarter-wave plates are placed in between a pair of polarizers. They are




Figure A.2: Polarimeter used to characterise quarter-wave plates. We begin by using a half-wave plate and PBS to align the polarization vector of an input laser beam with the horizontal ($\hat{x}$). The beam then passes through a pair of nominally identical quarter-wave plates, followed by a second polarizer. The quarter-wave plates are adjusted to minimise the intensity of light transmitted by the second polarizer.

then rotated and tilted to minimise the intensity of light transmitted by the second polarizer. At this point the two quarter-wave plates are aligned to form a half-wave plate. We began with 397 nm light, for which we were able to achieve an extinction ratio of 100,000 : 1, limited by the extinction of the polarizers used.

To investigate how well these waveplates work at 393 nm, we replaced the 397 nm light with 393 nm and measured the extinction without adjusting the waveplates. The extinction was found to be 10,000 : 1, which is consistent with our previous analysis to within a factor of two.

In our experiment, the optical pumping beams must pass through a quartz viewport, which is expected to have non-negligible birefringence – particularly when placed under strain by the internal vacuum. To investigate the effect of the viewport on our achievable polarization purity, we placed a viewport in the polarimeter with our two quarter-wave plates. Once again, we rotated and tilted the plates to maximise extinction. This time, we were only able to achieve an extinction of order 5,000 : 1 at 397 nm.

A.4 397 nm Doppler Polarization

We consider the situation shown in figure A.3, with a laser beam propagating along the $\hat{z}$ direction at 45° to a second axis $\hat{Z}$. The light's polarization may be described as

$$\begin{aligned} E_x &= ra \cos(\omega t - \phi) \\ E_y &= a \cos \omega t \\ E_z &= 0 \end{aligned} \quad (\text{A.19})$$

which, with respect to the $\hat{X}$, $\hat{Y}$ & $\hat{Z}$ axes, is

$$\begin{aligned} E_X &= \frac{ra}{\sqrt{2}} \cos(\omega t - \phi) \\ E_Y &= a \cos \omega t \\ E_Z &= \frac{ra}{\sqrt{2}} \cos(\omega t - \phi) \end{aligned} \quad (\text{A.20})$$

Here, E_Z corresponds to the field's π component, while E_X and E_Y give rise to the two σ components. We see from (A.20) that with appropriate choice of r




Figure A.3: We consider a laser beam – our 397 nm Doppler in this case – propagating along the $\hat{z}$ axis with a polarization defined relative to a set of Cartesian axes $\{x, y, z\}$. We wish to find its polarization with respect to a second set of axes $\{X, Y, Z\}$, where the static B-field, B_0 , defining our quantization axis lies along $\hat{Z}$.




Figure A.4: Polarimeter used to find the axes of half- and quarter-wave plates. We begin by using a half-wave plate and PBS to align the polarization vector of a 397 nm laser beam with the horizontal ($\hat{x}$). The light then passes through the waveplate to be characterised (DUT), followed by a second polarizer. The waveplate is adjusted to minimise the intensity of light reflected by the second polarizer. At this point the waveplate's axes are aligned along $\hat{x}$ and $\hat{y}$.

and ϕ , one may create a field with arbitrary ratio of σ^+ to σ^- . For a given σ^+ to σ^- ratio, however, the amplitude of the π component is then fixed to be $ra/\sqrt{2}$.

If one wishes one σ component to be absent then E_X and E_Y must have equal amplitudes and be $\pi/2$ out of phase. For this to be the case, $r = \sqrt{2}$ and $\phi = \pm\pi/2$. The amplitudes of all three field components are then equal. The intensity in the π component is thus 1/3 of the total, leaving 2/3 in the single σ component.

We wish our 397 nm cooling beam to contain no σ^- component. Given our beam geometry we are thus forced to use a 397 nm polarization ratio of 2/3 of σ^+ to 1/3 of π . Numerical modelling of the cooling process (see [Jan]) indicates that this polarization combination should allow near optimal fluorescence and thus should be well suited to these experiments. To produce light with this polarization, we need an input beam with an amplitude ratio $E_x/E_y = \sqrt{2}$, corresponding to linearly polarized light at an angle $\theta = \tan^{-1}(\sqrt{2}) = 54.7^\circ$ to the vertical ($\hat{y}$). In our beam path (figure 4.6), the 397 nm beam passes through a

PBS before the half- and quarter-wave plates and so is initially polarized along $\hat{x}$. We thus need our half-wave plate to be set to an angle of $(\pi/2 + \theta)/2 = 72.4^\circ$ to the vertical. To produce σ^+ rather than σ^- light, we need E_y to lag E_x by $\pi/2$ (section B.3). The quarter-wave plate should thus be aligned with its slow axis along $\hat{y}$.

To set the wave plates to the correct angles, one needs know the orientations of their axes relative to the mounts they are contained in. This was not known for our waveplates and so was measured using the polarimeter shown in figure A.4.

This experiment does not identify which is the fast and which the slow axis of the quarter-wave plate (although, this could be done, for example, using a Babinet-Soleil compensator). In practice, however, there are only two possibilities, leading to a beam with predominantly σ^+ or σ^- light and, from experience, it is possible to load and cool an ion with either. However, if the polarization used is not that appropriate to the choice of detunings used then the the fluorescence will be greatly suppressed and the ion will be relatively hot (evident by a blurring out of expected features and a large degree of hysteresis when the 866 nm laser's detuning is scanned). If this is the case then the waveplate may be rotated to correct the polarization.

B

Definitions and useful results

B.1 Rotating frames and the interaction picture

In this section we review the use of the interaction picture to separate the free evolution of a system from its response to a time-dependent interaction. We begin by considering a system with time-independent Hamiltonian H_0 , subject to a time-dependent interaction H_1 . The system evolves according to the time-dependent Schrödinger equation

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H_T |\psi(t)\rangle \quad (\text{B.1})$$

where $H_T = H_0 + H_1$ is the system's total Hamiltonian.

We make the unitary transformation

$$|\tilde{\psi}\rangle := U^\dagger |\psi\rangle \quad (\text{B.2})$$

where

$$U := e^{-iH_0 t/\hbar} \quad (\text{B.3})$$

After this transformation our Schrödinger equation becomes

$$\begin{aligned} i\hbar \frac{\partial}{\partial t} |\tilde{\psi}\rangle &:= i\hbar \frac{\partial}{\partial t} (U^\dagger |\psi(t)\rangle) \\ &= i\hbar \frac{\partial U^\dagger}{\partial t} |\psi\rangle + U^\dagger H_T |\psi\rangle \\ &= i\hbar \frac{\partial U^\dagger}{\partial t} U |\tilde{\psi}\rangle + U^\dagger H_T U |\tilde{\psi}\rangle \\ &= -H_0 |\tilde{\psi}\rangle + U^\dagger H_T U |\tilde{\psi}\rangle \\ &= U^\dagger H_1 U |\tilde{\psi}\rangle \\ &=: H_I |\tilde{\psi}\rangle \end{aligned} \quad (\text{B.4})$$

In the interaction picture the system's dynamics are entirely determined by H_I . Since the transformation between the Schrödinger and interaction pictures is unitary it can be thought of as a rotation. The transformation is thus pictured as moving into a frame which is co-rotating with H_0 , removing all dynamics due to H_0 , leaving only the interaction we wish to study.

Note that no approximation is involved in this transformation – in particular, we haven't assumed that H_1 represents a small perturbation. However, we will

often follow the transformation to the rotating frame with the rotating-wave approximation, in which we drop all terms in our interaction picture Hamiltonian which oscillate quickly compared with the time-scale of the system's dynamics.

B.2 Two-state systems and the Pauli matrices

When dealing with two-state systems, we will label our states $|0\rangle$ and $|1\rangle$, where $|0\rangle$ has the lower energy. On the whole, they will represent the clock-qubit states so that $|0\rangle := |4, +0\rangle$ and $|1\rangle := |3, +1\rangle$. We represent the ion's state as the vector

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} := \alpha|1\rangle + \beta|0\rangle \quad (\text{B.5})$$

We describe operations on our state in terms of the Pauli matrices

$$\begin{aligned} \sigma_x &:= |0\rangle\langle 1| + |1\rangle\langle 0| = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \\ \sigma_y &:= i(|0\rangle\langle 1| - |1\rangle\langle 0|) = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \\ \sigma_z &:= |1\rangle\langle 1| - |0\rangle\langle 0| = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \end{aligned} \quad (\text{B.6})$$

where the σ_x and σ_y operators represent couplings between the two-qubit states and the σ_z operator represents their energy splitting. Their commutation relations are

$$[\sigma_n, \sigma_m] = 2 \sum_c \epsilon_{nmc} \sigma_c \quad (\text{B.7})$$

where $n, m = \{x, y, z\}$ and ϵ_{nmc} is the Levi-Civita symbol.

These operators, in conjunction with the identity operator, form a basis for the set of operators on our two-dimensional space. Terms proportional to the identity operator represent shifts in the centre of gravity the system's energy and so can always be removed by appropriate choice of energy datum.

An important Pauli operator identity is the analogue of Euler's formula

$$e^{ia(\hat{\mathbf{n}} \cdot \boldsymbol{\sigma})} = \mathbb{1} \cos a + i(\hat{\mathbf{n}} \cdot \boldsymbol{\sigma}) \sin a \quad (\text{B.8})$$

where $\boldsymbol{\sigma} := \sum_{i=\{x,y,z\}} \hat{\mathbf{e}}_i \sigma_i$, with $\hat{\mathbf{e}}_i$ the Cartesian unit vectors and $\hat{\mathbf{n}}$ is a unit vector.

The raising and lowering operators are defined by

$$\begin{aligned} \sigma_+ &:= \frac{1}{2}(\sigma_x + i\sigma_y) = |1\rangle\langle 0| = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \\ \sigma_- &:= \frac{1}{2}(\sigma_x - i\sigma_y) = |0\rangle\langle 1| = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \end{aligned} \quad (\text{B.9})$$

Some important commutation relations for these operators are

$$\begin{aligned} [\sigma_+, \sigma_-] &= \sigma_z \\ [\sigma_+, \sigma_z] &= -2\sigma_+ \end{aligned} \quad (\text{B.10})$$

We will often work in the interaction picture with respect to a Hamiltonian of the form

$$H_0 = \begin{pmatrix} +\frac{1}{2}\hbar\omega_0 & 0 \\ 0 & -\frac{1}{2}\hbar\omega_0 \end{pmatrix} \quad (\text{B.11})$$

where $\omega_0 := \frac{1}{\hbar}(E_1 - E_0)$ is the frequency splitting between the two states. In this case the transformation U is given by

$$U = \begin{pmatrix} e^{-\frac{1}{2}i\omega_0 t} & 0 \\ 0 & e^{+\frac{1}{2}i\omega_0 t} \end{pmatrix} \quad (\text{B.12})$$

The Pauli matrices then transform as

$$\begin{aligned} \tilde{\sigma}_x &:= U^\dagger \sigma_x U = \begin{pmatrix} 0 & e^{i\omega_0 t} \\ e^{-i\omega_0 t} & 0 \end{pmatrix} \\ \tilde{\sigma}_y &:= U^\dagger \sigma_y U = \begin{pmatrix} 0 & -ie^{i\omega_0 t} \\ ie^{-i\omega_0 t} & 0 \end{pmatrix} \\ \tilde{\sigma}_z &:= U^\dagger \sigma_z U = \sigma_z \end{aligned} \quad (\text{B.13})$$

and the raising and lowering operators transform as

$$\begin{aligned} \tilde{\sigma}_+ &:= U^\dagger \sigma_+ U = \begin{pmatrix} 0 & e^{i\omega_0 t} \\ 0 & 0 \end{pmatrix} \\ \tilde{\sigma}_- &:= U^\dagger \sigma_- U = \begin{pmatrix} 0 & 0 \\ e^{-i\omega_0 t} & 0 \end{pmatrix} \end{aligned} \quad (\text{B.14})$$

B.3 The spherical basis

We can characterise atomic transitions by their polarization $q := M_n - M_m$, where M_n and M_m are the projections of angular momentum in the higher and lower energy states respectively. When transitions are driven by electromagnetic fields, the $\hat{z}$ component¹ of the field can only drive transitions in which $q = 0$, whilst its $\hat{x}$ and $\hat{y}$ components both drive those in which $q = \pm 1$.

It will often prove useful to work in a basis in which each field component is only capable of driving transitions with a single value of q . A convenient choice is the spherical basis, whose elements are

$$\begin{aligned} \hat{\mathbf{e}}_{\pm 1} &:= \mp \frac{1}{\sqrt{2}} (\hat{\mathbf{x}} \pm i\hat{\mathbf{y}}) = -(\hat{\mathbf{e}}_{\mp 1})^* \\ \hat{\mathbf{e}}_0 &:= \hat{\mathbf{z}} = (\hat{\mathbf{e}}_0)^* \end{aligned} \quad (\text{B.15})$$

¹Here, we take $\hat{z}$ to be our quantization axis.

These vectors satisfy the orthonormality relations

$$\begin{aligned}\hat{\mathbf{e}}_{\pm 1} \cdot \hat{\mathbf{e}}_{\pm 1} &= 0 \\ \hat{\mathbf{e}}_{\pm 1} \cdot \hat{\mathbf{e}}_{\mp 1} &= -1 \\ \hat{\mathbf{e}}_z \cdot \hat{\mathbf{e}}_q &= \delta_{q,0}\end{aligned}\tag{B.16}$$

For a vector $\mathbf{A}$, whose Cartesian components are $\mathbf{A} = A_x \hat{\mathbf{x}} + A_y \hat{\mathbf{y}} + A_z \hat{\mathbf{z}}$, we define the spherical components $A_q := \hat{\mathbf{e}}_q \cdot \mathbf{A}$. From which it follows that

$$\begin{aligned}A_{\pm 1} &= \mp \frac{1}{\sqrt{2}} (A_x \pm i A_y) \\ A_0 &= A_z\end{aligned}\tag{B.17}$$

$$\mathbf{A} = \sum_q (-1)^q A_q \hat{\mathbf{e}}_{-q} = \sum_q A_q \hat{\mathbf{e}}_q^* \tag{B.18}$$

We note that in this basis the product of two vectors is given by

$$\mathbf{A} \cdot \mathbf{B} = \sum_q (-1)^q A_q B_{-q} \tag{B.19}$$

Let us now consider, by way of example, the magnetic dipole interaction, whose Hamiltonian is

$$H = -\boldsymbol{\mu} \cdot \mathbf{B} \tag{B.20}$$

where $\boldsymbol{\mu}$ is the atom's magnetic dipole moment. We can use the above identities to re-write this in terms of spherical basis components as

$$\begin{aligned}H &= \sum_q (-1)^{q+1} \{\mu_q B_{-q}\} \\ &= -\{B_z \mu_z - B_{-1} \mu_{+1} - B_{+1} \mu_{-1}\}\end{aligned}\tag{B.21}$$

For a transition with polarization Q , the matrix elements of this Hamiltonian are given by

$$\begin{aligned}H_{nm} &:= \langle M_n | H | M_m \rangle \\ &:= \langle M_m + Q | H | M_m \rangle \\ &= \sum_q (-1)^{q+1} \{ \langle M_m + Q | \mu_q B_{-q} | M_m \rangle \}\end{aligned}\tag{B.22}$$

The atom's magnetic dipole moment is related to its angular momentum by an equation of the form $\boldsymbol{\mu} = -\mu_B g_J \mathbf{J}$, where $\mathbf{J}$ is the atom's angular momentum². To evaluate these matrix elements, we thus need to know the spherical components of the angular momentum operators. These are related to the raising and lowering "ladder" operators, $J_{\pm}$, by $J_{\pm 1} = \mp \frac{1}{\sqrt{2}} J_{\pm}$. When acting on an angular momentum eigenstate, their eigenvalues are (in units of $\hbar$)

$$\begin{aligned}J_{\pm 1} |J, M_J\rangle &= \mp \frac{1}{\sqrt{2}} \sqrt{(J \mp M_J)(J \pm M_J + 1)} |J, M_J \pm 1\rangle \\ J_z |J, M_J\rangle &= M_J |J, M_J\rangle\end{aligned}\tag{B.23}$$

²It the atom has hyperfine structure then this becomes $\boldsymbol{\mu} = -\mu_B g_J \mathbf{J} + \mu_N g_I \mathbf{I}$.

Thus, the only non-vanishing component of the sum in equation B.22 has $q = Q$. In other words, the spherical component B_q will only drive transitions in which $q = -Q$.

It is useful to demonstrate the connection between radiation with π or circular ($\sigma^\pm$) polarization and the spherical basis vectors. In both the Cartesian and the spherical basis, π -polarized radiation is defined by

$$\mathbf{B}_\pi = B_0 \hat{\mathbf{z}} \cos \omega t \quad (\text{B.24})$$

In Cartesian coordinates circularly polarized radiation is defined by

$$\begin{aligned} \mathbf{B}_{\sigma^\pm} &= B_0 \begin{pmatrix} \cos \omega t \\ \mp \cos(\omega t + \frac{\pi}{2}) \\ 0 \end{pmatrix} \\ &= \frac{1}{2} B_0 e^{i\omega t} \begin{pmatrix} 1 \\ \mp i \\ 0 \end{pmatrix} + \frac{1}{2} B_0 e^{-i\omega t} \begin{pmatrix} 1 \\ \pm i \\ 0 \end{pmatrix} \\ &= \frac{1}{\sqrt{2}} B_0 (-\hat{\mathbf{e}}_{+1} e^{\mp i\omega t} + \hat{\mathbf{e}}_{-1} e^{\pm i\omega t}) \end{aligned} \quad (\text{B.25})$$

In the spherical basis, circularly polarized radiation thus has components

$$\begin{aligned} \mathbf{B}_{\sigma^\pm}^{(q=+1)} &:= \hat{\mathbf{e}}_{+1} \cdot \mathbf{B}_{\sigma^\pm} = -\frac{1}{\sqrt{2}} B_0 e^{\pm i\omega t} \\ \mathbf{B}_{\sigma^\pm}^{(q=-1)} &:= \hat{\mathbf{e}}_{-1} \cdot \mathbf{B}_{\sigma^\pm} = \frac{1}{\sqrt{2}} B_0 e^{\mp i\omega t} \end{aligned} \quad (\text{B.26})$$

Note that a circularly polarized field is not given purely by a $B_{\pm 1}$ component, but contains both $B_{\pm 1}$ and a $B_{\mp 1}$ components. Indeed, we can see from equation B.17 that the $B_{\pm 1}$ fields are complex and hence can never exist without their conjugates. The difference between σ^+ and σ^- radiation thus lies purely in the sign of the phases in the $e^{i\omega t}$ components.

To make this clearer, let's re-consider the magnetic dipole interaction in the case where a transition with polarization $Q = \pm 1$ is driven by a circularly polarized magnetic field. Our matrix elements are

$$\begin{aligned} H_{nm}^\pm &:= \langle M_m + Q | -\boldsymbol{\mu} \cdot \mathbf{B}_{\sigma^\pm} | M_m \rangle \\ &\propto \langle M_m + Q | e^{\mp i(\omega_0 + \delta)t} J_{+1} - e^{\pm i(\omega_0 + \delta)t} J_{-1} | M_m \rangle \\ &\propto e^{\mp iQ(\omega_0 + \delta)t} \end{aligned} \quad (\text{B.27})$$

Where we have defined the detuning of the driving field, δ , such that $\omega = \omega_0 + \delta$, where $\omega_0 = \frac{1}{\hbar} (E_n - E_m)$ is the transition frequency.

The Hamiltonian thus has the form

$$H = \begin{pmatrix} 0 & e^{\mp iQ(\omega_0 + \delta)t} \\ e^{\pm iQ(\omega_0 + \delta)t} & 0 \end{pmatrix} \quad (\text{B.28})$$

If we now move into a rotating frame with respect to the free atom's energy, we end up with an interaction picture Hamiltonian of the form

$$\tilde{H}_I = \begin{pmatrix} 0 & e^{\mp iQ\delta t + i\omega_0 t(1 \mp Q)} \\ e^{\pm iQ\delta t - i\omega_0 t(1 \mp Q)} & 0 \end{pmatrix} \quad (\text{B.29})$$

If $Q = \pm 1$ then the ω_0 terms cancel and we end up with a near-resonant interaction with detuning δ . However, if $Q = \mp 1$ then the ω_0 terms counter-rotate and can be neglected using the rotating wave approximation.

In other words, so long as the rotating wave approximation is valid, the field $\mathbf{B}_{\sigma\pm}$ can only drive $Q = \pm 1$ transitions, despite the presence of both J_{+1} and J_{-1} terms in our Hamiltonian.

B.4 The Magnus Expansion

The Magnus expansion [Mag54][BCOR09] is a series solution to the differential equation

$$\frac{dY}{dt} = AY \quad (\text{B.30})$$

where $A(t)$ and $Y(t)$ are linear operators dependent on the real variable t and $Y(0) = \mathbb{1}$ ($\mathbb{1}$ is the identity operator). If $[A(t), A(t')] \neq 0$ then this has a non-trivial solution, which we can express using the exponential the expansion

$$Y(t) = \exp\left(\sum_{k=1}^{\infty} \Omega_k(t)\right) Y(0) \quad (\text{B.31})$$

where

$$\begin{aligned} \Omega_1(t) &= \int_0^t A(t_1) dt_1 \\ \Omega_2(t) &= \frac{1}{2} \int_0^t dt_1 \int_0^{t_1} dt_2 [A(t_1), A(t_2)] \\ \Omega_3(t) &= \frac{1}{6} \int_0^t dt_1 \int_0^{t_1} dt_2 \int_0^{t_2} dt_3 \{ [A(t_1), [A(t_2), A(t_3)]] + [A(t_3), [A(t_2), A(t_1)]] \} \\ &\vdots \end{aligned} \quad (\text{B.32})$$

A sufficient (although not necessary) condition for the series $\Omega = \Omega_1 + \Omega_2 + \dots + \Omega_k$ to terminate at a finite k is that

$$[A(t_1), [A(t_2), [A(t_3), \dots, [A(t_k), A(t_{k+1})] \dots]]] = 0 \quad (\text{B.33})$$

for all $t_1, \dots, t_{k+1}$.

Our interest in the Magnus expansion will lie in the construction of solutions to the time-dependent Schrödinger equation. Defining the propagator $U(t)$, such that

$$|\psi(t)\rangle := U(t) |\psi(0)\rangle \quad (\text{B.34})$$

we have

$$\begin{aligned} H(t)|\psi(t)\rangle &= H U(t) |\psi(0)\rangle \\ &= i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle \\ &= i\hbar \frac{dU(t)}{dt} |\psi(0)\rangle \end{aligned} \quad (\text{B.35})$$

Our propagator is thus given by the series solution to the equation

$$\frac{dU(t)}{dt} = \frac{1}{i\hbar} H U(t) \quad (\text{B.36})$$

Perhaps the most striking feature of the Magnus expansion is that the series expansion of Ω , when truncated at any order, belongs to the same Lie algebra as A and hence its exponential belongs to the associated Lie group. This ensures that the truncated series solution retains the symmetries of the original problem. In particular, it ensures that our approximate propagators will be unitary and hence will conserve probability.

B.5 The A.C. Zeeman Shift

In this section we will derive an expression for the A.C. Zeeman shift – the energy shift induced when an oscillating magnetic field off-resonantly excites an atomic transition – including the effects of counter-rotating terms.

We will assume throughout that a large static magnetic field, $\mathbf{B}_0$, along $\hat{z}$ defines our quantization axis. We assume further that the perturbation induced by the oscillating field is sufficiently small that shifts on different pairs of states simply add. We may thus consider, without loss of generality, the shift on a two-level system with polarization $Q := M_1 - M_0$.

B.5.1 Transverse polarization

We begin with the case of radiation polarized along $\hat{x}$ interacting with a σ -polarized transition ($Q = \pm 1$). Our field is thus

$$\mathbf{B}(t) = B\hat{x} \cos \omega t \quad (\text{B.37})$$

so, in the spherical basis

$$B_{\pm 1}(t) = \mp \frac{1}{\sqrt{2}} B \cos \omega t \quad (\text{B.38})$$

The Hamiltonian for the interaction is thus

$$\begin{aligned} H_1 &= -\boldsymbol{\mu} \cdot \mathbf{B}(t) \\ &= \frac{1}{\sqrt{2}} B \cos \omega t (\mu_{+1} - \mu_{-1}) \\ &= \hbar \Omega \sigma_x \cos \omega t \end{aligned} \quad (\text{B.39})$$

where we have used our usual definition of the Rabi frequency, $\hbar \Omega := B_{-Q} \langle 1 | \mu_Q | 0 \rangle$, and we have expressed our Hamiltonian in the basis of unperturbed (no oscillating field) energy eigenstates.

With suitable choice of energy datum, our full system Hamiltonian is thus

$$H = \hbar \begin{pmatrix} \omega_0 & \Omega \cos \omega t \\ \Omega \cos \omega t & 0 \end{pmatrix} \quad (\text{B.40})$$

where ω_0 is the transition frequency in the absence of the oscillating field.

We may express the evolution of an arbitrary perturbed state, $|A(t)\rangle$, in terms of the unperturbed states as

$$|A(t)\rangle = a_1(t) e^{-i\omega_0 t} |1\rangle + a_0(t) |0\rangle \quad (\text{B.41})$$

where

$$\begin{aligned} |1(t)\rangle &= e^{-i\omega_0 t} |1\rangle \\ |0(t)\rangle &= |0\rangle \end{aligned} \quad (\text{B.42})$$

are the unperturbed energy eigenstates.

The time-dependent Schrödinger equation then becomes

$$i\hbar \frac{d}{dt} \begin{pmatrix} a_1(t)e^{-i\omega_0 t} \\ a_0(t) \end{pmatrix} = \hbar \begin{pmatrix} \omega_0 a_1(t)e^{-i\omega_0 t} + \Omega a_0(t) \cos \omega t \\ \Omega a_1(t)e^{-i\omega_0 t} \cos \omega t \end{pmatrix} \quad (\text{B.43})$$

so

$$\begin{aligned} \dot{a}_1(t) &= -i\Omega a_0(t)e^{i\omega_0 t} \cos \omega t \\ \dot{a}_0(t) &= -i\Omega a_1(t)e^{-i\omega_0 t} \cos \omega t \end{aligned} \quad (\text{B.44})$$

To calculate the energy shift on the higher-energy state, we make the assumption that $a_0(0) = 0$ and $a_1(t) \simeq e^{-i\Delta_{\text{LS}} t}$, which implies

$$\begin{aligned} a_1(t)\Delta_{\text{LS}} &= \Omega a_0(t)e^{i\omega_0 t} \cos \omega t \\ \dot{a}_0(t) &= -i\Omega e^{-i(\omega_0 + \Delta_{\text{LS}})t} \cos \omega t \\ &= -\frac{1}{2}i\Omega e^{i(\omega - \omega_0 - \Delta_{\text{LS}})t} - \frac{1}{2}i\Omega e^{-i(\omega + \omega_0 + \Delta_{\text{LS}})t} \end{aligned} \quad (\text{B.45})$$

Integrating the above equation for $a_0(t)$, we find

$$a_0(t) = -\frac{1}{2}\Omega \frac{e^{i(\omega - \omega_0 - \Delta_{\text{LS}})t}}{\omega - \omega_0 - \Delta_{\text{LS}}} + \frac{1}{2}\Omega \frac{e^{-i(\omega + \omega_0 + \Delta_{\text{LS}})t}}{\omega + \omega_0 + \Delta_{\text{LS}}} - \frac{1}{2}\Omega \gamma \quad (\text{B.46})$$

where $\gamma := -2\frac{\omega_0 + \Delta_{\text{LS}}}{\omega^2 - (\omega_0 + \Delta_{\text{LS}})^2}$.

Inserting this into (B.45)

$$e^{-i\Delta_{\text{LS}} t} \Delta_{\text{LS}} = -\frac{1}{2}\Omega^2 e^{i\omega_0 t} \cos \omega t \left(\frac{e^{i(\omega - \omega_0 - \Delta_{\text{LS}})t}}{\omega - \omega_0 - \Delta_{\text{LS}}} - \frac{e^{-i(\omega + \omega_0 + \Delta_{\text{LS}})t}}{\omega + \omega_0 + \Delta_{\text{LS}}} + \gamma \right) \quad (\text{B.47})$$

We thus find the shift is given by

$$\begin{aligned} \Delta_{\text{LS}} &= -\frac{1}{4}\Omega^2 (e^{i\omega t} + e^{-i\omega t}) \left(\frac{e^{i\omega t}}{\omega - \omega_0 - \Delta_{\text{LS}}} - \frac{e^{-i\omega t}}{\omega + \omega_0 + \Delta_{\text{LS}}} + \gamma e^{i\omega_0 t} \right) \\ \langle \Delta_{\text{LS}} \rangle &= \frac{1}{4}\Omega^2 \left(\frac{1}{\omega_0 + \Delta_{\text{LS}} + \omega} + \frac{1}{\omega_0 + \Delta_{\text{LS}} - \omega} \right) \\ &= \frac{1}{2}\Omega^2 \left(\frac{\omega_0}{\omega_0^2 - \omega^2} \right) \end{aligned} \quad (\text{B.48})$$

where, in the last line, we have used the approximation $\omega_0^2 \simeq (\omega_0 + \Delta_{\text{LS}})^2$.

By assuming $a_1(0) = 0$, $a_0(t) \simeq e^{-i\Delta_{\text{LS}} t}$, one can show that the energy shift of the lower-energy state is equal to (B.48), but with opposite sign.

The sign of (B.48) is such that blue-detuned fields ($\omega > \omega_0$) create a negative energy shift, pulling the states together, whereas red-detuned radiation ($\omega < \omega_0$) gives a positive shift, pushing the states apart. This is illustrated by figure B.1.




Figure B.1: AC Zeeman shift, Δ_{LS} , on a transition due to detuned radiation with frequency $\omega = \omega_0 + \Delta$. The left hand side shows red detuned radiation increasing the energy splitting between a pair of states, whilst the right hand side shows the effect blue-detuned radiation.

B.5.2 Longitudinal polarization

Next, we consider a field polarized along $\hat{z}$, so that

$$\mathbf{B}(t) = B\hat{z} \cos \omega t \quad (\text{B.49})$$

Our Hamiltonian is thus

$$H_1 = -B\mu_z \cos \omega t \quad (\text{B.50})$$

Unlike the case of a transverse polarized field, the diagonal matrix elements of this Hamiltonian are non-zero. With suitable choice of energy datum, our total Hamiltonian is thus

$$H = \begin{pmatrix} \hbar\omega_0 + BR_{11} \cos \omega t & \hbar\Omega \cos \omega t \\ \hbar\Omega \cos \omega t & BR_{00} \cos \omega t \end{pmatrix} \quad (\text{B.51})$$

where $R_{ii} := -\langle i|\mu_z|i\rangle$ and we have defined our Rabi frequency as $\hbar\Omega := -B\langle 1|\mu_z|0\rangle$.

This Hamiltonian is identical to that of (B.40) except for the additional diagonal terms, representing a small modulation of the transition frequency. So long as this is small compared with the Rabi frequency, it has no effect on the perturbation and we find that the result of (B.48) is valid, regardless of the polarization of the field.

B.6 Rabi flopping and Ramsey Experiments

In this section we consider the interaction between a two-state atomic system and an oscillating magnetic field within rotating wave approximation. We begin by deriving the propagator for the interaction before applying it to some situations of interest.

We consider a transition with polarization $Q := M_1 - M_0$, driven by a magnetic field

$$\mathbf{B}(t) = \mathbf{B} \cos(\omega t - \phi_0) \quad (\text{B.52})$$

The magnetic dipole Hamiltonian is

$$H_1 = -\boldsymbol{\mu} \cdot \mathbf{B}(t) \quad (\text{B.53})$$

We ignore the on-diagonal elements of this Hamiltonian, which drop out when we make the rotating wave approximation. Working in the basis of free-atom energy eigenstates, we are thus left with the coupling Hamiltonian

$$H_1 = \hbar\Omega\sigma_x \cos(\omega t - \phi_0) \quad (\text{B.54})$$

where, as usual, we have defined our Rabi frequency by $\hbar\Omega := B_{-Q}R_{M_1, M_0}$ and the matrix elements by $R_{nm} := (-1)^{Q+1}\langle n|\mu_Q|m\rangle$.

Our complete Hamiltonian is

$$\begin{aligned} H_T &= H_0 + H_1 \\ &= \frac{1}{2}\hbar\omega_0\sigma_z + \hbar\Omega\sigma_x \cos(\omega t - \phi_0) \end{aligned} \quad (\text{B.55})$$

where H_0 is the free-atom Hamiltonian.

We can remove the time-dependence from this equation by making the unitary transformation

$$|\psi\rangle \mapsto |\tilde{\psi}\rangle := U^\dagger |\psi\rangle \quad (\text{B.56})$$

where

$$U := \begin{pmatrix} e^{-\frac{1}{2}i\omega t} & 0 \\ 0 & e^{\frac{1}{2}i\omega t} \end{pmatrix} \quad (\text{B.57})$$

In the new basis the Schrödinger equation becomes

$$i\hbar \frac{\partial}{\partial t} |\tilde{\psi}\rangle = \tilde{H} |\tilde{\psi}\rangle \quad (\text{B.58})$$

where

$$\begin{aligned} \tilde{H} &:= i\hbar \frac{\partial U^\dagger}{\partial t} U + U^\dagger H U \\ &= -\frac{1}{2}\hbar\omega\sigma_z + \frac{1}{2}\hbar\omega_0\sigma_z + \frac{1}{2}\hbar\Omega (e^{i\phi_0}\sigma_+ + e^{-i\phi_0}\sigma_-) \\ &= -\frac{1}{2}\hbar\delta\sigma_z + \frac{1}{2}\hbar\Omega (e^{i\phi_0}\sigma_+ + e^{-i\phi_0}\sigma_-) \end{aligned} \quad (\text{B.59})$$

and we have made the rotating wave approximation and used our usual definition of the radiation's detuning from the transition $\omega := \omega_0 + \delta$.

We define our propagator by

$$|\tilde{\psi}(t_1)\rangle := U(t_1, t_0) |\tilde{\psi}(t_0)\rangle \quad (\text{B.60})$$

Notice that from here on U will refer to the propagator and *not* to the rotation matrix we used to change our basis. A small notational inconvenience, I admit, but I've seen – and used – far worse...

Since our Hamiltonian is time-independent, we can immediately write down the corresponding propagator

$$U(t_1 - t_0) = e^{-\frac{i}{\hbar}\tilde{H}(t_1 - t_0)} \quad (\text{B.61})$$

To put this into a more useful form, we use the Pauli matrix identity

$$e^{ia(\hat{\mathbf{n}} \cdot \boldsymbol{\sigma})} = \mathbb{1} \cos a + i(\hat{\mathbf{n}} \cdot \boldsymbol{\sigma}) \sin a \quad (\text{B.62})$$

where $\boldsymbol{\sigma} := \sum_{i=\{x,y,z\}} \hat{\mathbf{e}}_i \sigma_i$, with $\hat{\mathbf{e}}_i$ the Cartesian unit vectors. Now,

$$\begin{aligned} -\frac{1}{\hbar} \tilde{H} &= \frac{1}{2} \delta \sigma_z - \frac{1}{4} \Omega \{ e^{i\phi_0} (\sigma_x + i\sigma_y) + e^{-i\phi_0} (\sigma_x - i\sigma_y) \} \\ &= \frac{1}{2} \delta \sigma_z - \frac{1}{2} \Omega (\sigma_x \cos \phi_0 - \sigma_y \sin \phi_0) \end{aligned} \quad (\text{B.63})$$

We thus identify

$$\begin{aligned} \hat{\mathbf{n}} &= \frac{1}{\Omega_{\text{eff}}} \begin{pmatrix} -\Omega \cos \phi_0 \\ \Omega \sin \phi_0 \\ \delta \end{pmatrix} \\ a &= \frac{1}{2} \Omega_{\text{eff}} t \end{aligned} \quad (\text{B.64})$$

where $\Omega_{\text{eff}}^2 := \Omega^2 + \delta^2$ is the effective Rabi frequency.

Our propagator is thus

$$\begin{aligned} U(t) &= \mathbb{1} \cos \left(\frac{1}{2} \Omega_{\text{eff}} t \right) + \frac{i}{\Omega_{\text{eff}}} (\delta \sigma_z - \Omega \sigma_x \cos \phi_0 + \Omega \sigma_y \sin \phi_0) \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) \\ &= \mathbb{1} \cos \left(\frac{1}{2} \Omega_{\text{eff}} t \right) + \frac{i}{\Omega_{\text{eff}}} (\delta \sigma_z - \Omega e^{i\phi_0} \sigma_+ - \Omega e^{-i\phi_0} \sigma_-) \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) \end{aligned} \quad (\text{B.65})$$

which, in matrix form, is

$$U(t) = \begin{pmatrix} \cos \left(\frac{1}{2} \Omega_{\text{eff}} t \right) + i \frac{\delta}{\Omega_{\text{eff}}} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) & -i \frac{\Omega}{\Omega_{\text{eff}}} e^{i\phi_0} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) \\ -i \frac{\Omega}{\Omega_{\text{eff}}} e^{-i\phi_0} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) & \cos \left(\frac{1}{2} \Omega_{\text{eff}} t \right) - i \frac{\delta}{\Omega_{\text{eff}}} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) \end{pmatrix} \quad (\text{B.66})$$

Well, dear reader, wasn't that a lot of fun with Pauli matrices?

B.6.1 π -qubits

To read out a π -qubit, we wish to use the composite pulse sequence $\pi/2$ -delay- $\pi/2$ to perform a π -pulse on one qubit while doing nothing (up to a phase change) to a second qubit (see section 2.6.2.1). We will now use the propagator of equation (B.66) to determine how long the appropriate delay should be. For convenience, we will assume the radiation's phase at $t = 0$ to be 0. Our propagator reduces to

$$U(t) = \begin{pmatrix} \cos \left(\frac{1}{2} \Omega_{\text{eff}} t \right) + i \frac{\delta}{\Omega_{\text{eff}}} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) & -i \frac{\Omega}{\Omega_{\text{eff}}} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) \\ -i \frac{\Omega}{\Omega_{\text{eff}}} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) & \cos \left(\frac{1}{2} \Omega_{\text{eff}} t \right) - i \frac{\delta}{\Omega_{\text{eff}}} \sin \left(\frac{1}{2} \Omega_{\text{eff}} t \right) \end{pmatrix} \quad (\text{B.67})$$

We consider the following sequence of events. First, we apply the radiation for a pulse area $\theta := \Omega_{\text{eff}} t$. Next, we turn the microwave radiation off for a period τ , so that $\Omega \rightarrow 0$ and $\Omega_{\text{eff}} \rightarrow \delta$. Finally, we turn the radiation back on for a further pulse area θ . The complete sequence is described by the propagator

$$U_{\text{T}} = U_{\theta} U_{\tau} U_{\theta} \quad (\text{B.68})$$

where

$$U_\theta = \begin{pmatrix} \cos \frac{1}{2}\theta + i \frac{\delta}{\Omega_{\text{eff}}} \sin \frac{1}{2}\theta & -i \frac{\Omega}{\Omega_{\text{eff}}} \sin \frac{1}{2}\theta \\ -i \frac{\Omega}{\Omega_{\text{eff}}} \sin \frac{1}{2}\theta & \cos \frac{1}{2}\theta - i \frac{\delta}{\Omega_{\text{eff}}} \sin \frac{1}{2}\theta \end{pmatrix} \quad (\text{B.69})$$

and

$$U_\tau = \begin{pmatrix} \cos \frac{1}{2}\tau\delta + i \sin \frac{1}{2}\tau\delta & 0 \\ 0 & \cos \frac{1}{2}\tau\delta - i \sin \frac{1}{2}\tau\delta \end{pmatrix} \quad (\text{B.70})$$

We assume that the magnetic field is resonant with the transition we wish to drive and $\theta = \pi/2$ for this transition, resulting in a π -pulse. For the other transition, however, we want the diagonal elements of U_T to vanish. These elements are given by

$$U_{10} = U_{01} = 2i \frac{\Omega}{\Omega_{\text{eff}}^2} \sin\left(\frac{\theta}{2}\right) \left\{ \delta \sin\left(\frac{\theta}{2}\right) \sin\left(\frac{\delta\tau}{2}\right) - \Omega_{\text{eff}} \cos\left(\frac{\theta}{2}\right) \cos\left(\frac{\delta\tau}{2}\right) \right\} \quad (\text{B.71})$$

which vanish when

$$\tan\left(\frac{1}{2}\delta\tau\right) = \frac{\Omega_{\text{eff}}}{\delta} \cot\left(\frac{\theta}{2}\right) \quad (\text{B.72})$$

We will be interested in the case where $\theta \simeq \frac{\pi}{2}$, so that

$$\tan\left(\frac{1}{2}\tau\delta\right) \simeq \frac{\Omega_{\text{eff}}}{\delta} \quad (\text{B.73})$$

If $\delta \ll \Omega$ this condition becomes

$$\tau = \frac{\pi}{\delta} \quad (\text{B.74})$$

as one might expect.

B.6.2 Ramsey experiments

When $\theta = \frac{\pi}{2}$ and $\delta \ll \Omega$, so that $\Omega_{\text{eff}} \sim \Omega$, the propagator of (B.68) describes a Ramsey experiment. In this case

$$U \sim \begin{pmatrix} i \sin\left(\frac{\delta\tau}{2}\right) & -i \cos\left(\frac{\delta\tau}{2}\right) \\ -i \cos\left(\frac{\delta\tau}{2}\right) & -i \sin\left(\frac{\delta\tau}{2}\right) \end{pmatrix} \quad (\text{B.75})$$

The population undergoes oscillations, returning to its original state (up to a global phase) after a time of $\delta\tau = 2n\pi$. Note that δ is in angular units. For $\delta = 2\pi f$, the population returns to its original state when $\tau = \frac{n}{f}$.

B.6.3 Ramsey experiments as a phase noise probe

An important use of Ramsey experiments is probing the phase fluctuations that accumulate due to noise in the frequency of the transition or radiation field. We consider the case where $\delta \ll \Omega$ and treat the two $\pi/2$ -pulses as instantaneous, occurring at times t_0 and t_1 , so that the second $\pi/2$ -pulse has a phase of $\delta\phi := \phi(t_0 + \tau) - \phi(t_0)$ relative to the first. In this case our propagator is

$$U = \frac{1}{2} \begin{pmatrix} 1 - e^{i\delta\phi} & -i(1 + e^{i\delta\phi}) \\ -i(1 + e^{-i\delta\phi}) & 1 - e^{-i\delta\phi} \end{pmatrix} \quad (\text{B.76})$$

We consider applying this propagator to a qubit initially in the state $|0\rangle$. The populations oscillate as a function of $\delta\phi$ according to

$$P_0 = \frac{1}{2} (1 + \cos \delta\phi) \quad (\text{B.77})$$

We now imagine that the relative phases of the radiation and qubit fluctuate so that $\phi(t) = \phi_0(t) + \phi_n(t)$ where ϕ_0 represents the average evolution of the phase and $\phi_n(t)$ is a zero-mean fluctuation. In this case

$$\begin{aligned} P_0 &= \frac{1}{2} [1 + \cos(\delta\phi_0 + \delta\phi_n)] \\ &= \frac{1}{2} [1 + \cos \delta\phi_0 \cos \delta\phi_n - \sin \delta\phi_0 \sin \delta\phi_n] \end{aligned} \quad (\text{B.78})$$

The peak population occurs when $\delta\phi_0 = 0$ for which

$$\begin{aligned} \langle P_0^{\max} \rangle &= \frac{1}{2} [1 + \langle \cos \delta\phi_n \rangle] \\ &= \frac{1}{2} \left[1 + \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n)!} \langle \delta\phi_n^{2n} \rangle \right] \end{aligned} \quad (\text{B.79})$$

Equation (B.79) allows us to relate the contrast of observed Ramsey fringes to the statistics of the noise process ϕ_n .

We will generally be concerned with systems for which the timescale of the fringe decay is orders of magnitude longer than the correlation time of the frequency fluctuations. In this case, one may imagine the integrated phase noise as the result of a series of independent increments. From the central limit theorem we thus expect it to be approximately Gaussian under most conditions.

We take advantage of the important result that all odd moments of zero-mean Gaussian processes vanish, with the even moments given by

$$\langle \delta\phi^{2n} \rangle = \sigma^{2n} (2n - 1)!! \quad (\text{B.80})$$

where $\sigma^2(\tau) := \langle \delta\phi(\tau)^2 \rangle$ is the phase noise variance and $(2n - 1)!!$ denotes the double factorial³. Thus

$$\begin{aligned} \langle P_0^{\max} \rangle &= \frac{1}{2} \left[1 + \sum_{n=0}^{\infty} \frac{1}{n!} (-\sigma^2/2)^n \right] \\ &= \frac{1}{2} \left[1 + e^{-\frac{1}{2}\sigma^2(\tau)} \right] \end{aligned} \quad (\text{B.81})$$

If the detuning undergoes white fluctuations then the resulting phase noise is a Wiener diffusion process, characterised by a linearly increasing variance. Thus, the Ramsey contrast decays exponentially as

$$\langle P_0^{\max} \rangle = \frac{1}{2} \left[1 + e^{-\tau/T_2} \right] \quad (\text{B.82})$$

where T_2 is the coherence time, which is related to the noise variance by

$$\sigma(\tau)^2 = 2\tau/T_2 \quad (\text{B.83})$$

³The double factorial of $2n - 1$ is defined as the product of every odd number between $2n - 1$ and 1. It may be related to the normal factorial by the identity $(2n - 1)!! = \frac{(2n)!}{2^n n!}$.

A discussion of phase noise resulting from bandwidth limited Gaussian frequency fluctuations can be found in [BP03].

Over shorter time-scales one will often find that the detuning fluctuations are non-Gaussian. However, so long as the accumulated phase noise is small, one can simplify the problem by making the small noise approximation. Assuming that $\delta\phi_n \ll 1$ radian, we have

$$\langle P_0(\tau) \rangle \simeq 1 - \frac{1}{4} \langle \delta\phi_n^2(\tau) \rangle \quad (\text{B.84})$$

It is useful to relate this expression to the autocorrelation function for the detuning fluctuations, $\delta(t)$. This may be accomplished as follows

$$\begin{aligned} \delta\phi_n^2(t_1) &= \int_0^{t_1} \int_0^{t_1} \delta(t) \delta(t') dt dt' \\ &= \int_0^{t_1} \int_0^{t_1} \delta(t) \delta(t + \tau) dt dt' \\ &= \int_0^{t_1} d\tau \int_0^{t_1-\tau} dt \delta(t) \delta(t + \tau) + \int_{-t_1}^0 d\tau \int_{-\tau}^{t_1} dt \delta(t) \delta(t + \tau) \\ &= \int_0^{t_1} d\tau \int_0^{t_1-\tau} dt \delta(t) \delta(t + \tau) + \int_0^{t_1} d\tau \int_{\tau}^{t_1} dt \delta(t) \delta(t - \tau) \\ &= 2 \int_0^{t_1} d\tau \int_0^{t_1-\tau} dt \delta(t) \delta(t + \tau) \end{aligned} \quad (\text{B.85})$$

where we have defined $\tau := t' - t$.

Thus,

$$\langle P_0(\tau) \rangle \simeq 1 - \frac{1}{2} \int_0^{t_1} d\tau \int_0^{t_1-\tau} dt G(\tau) \quad (\text{B.86})$$

where we have used the standard definition of the autocorrelation function

$$G(\tau) := \langle \delta(t + \tau) \delta(t) \rangle \quad (\text{B.87})$$

The autocorrelation is directly related to the power spectral density of frequency fluctuations by the Wiener-Khinchin theorem. We note that for times much shorter than the correlation time scale of the detuning fluctuations

$$\langle P_0(\tau) \rangle \simeq 1 - \frac{1}{4} t_1^2 \langle \delta^2 \rangle \quad (\text{B.88})$$

as might be expected.

We see that on short timescales, the decay in fringe contrast is dominated by the coherence of the fluctuations. For longer times, we expect the random walk to take over leading to exponential decay. Related discussions of these points may be found in, for example, [Lan06], [Mon11] and [MSB⁺11].

C

Modulation theory

In this appendix I present an account of some basic aspects of modulation theory. For a more complete treatment, see e.g. [Rob84]. Standard treatments of modulation theory tend to focus on phase modulation (PM), while leaving amplitude modulation (AM) as a secondary concern. The reason for this is that in communications systems information is largely encoded in the phase of a carrier. Thus, small degrees of phase noise can be catastrophic to signal integrity, whilst moderate amounts AM can have little to no effect. However, in analysing the effect of noise in the RF trapping potential, we will find that AM is by far our main concern. The emphasis of this section reflects the unusual level of care we will need to take over AM during our trap RF design.

The appendix consists of four sections, the first two of which cover the topics of AM and PM. The third section covers the connection between AM and PM. The final section covers the relationship between noise and modulation sidebands.

C.1 Amplitude modulation

We begin with the signal

$$V(t) = V_0 \sin \Omega t \quad (\text{C.1})$$

and introduce a modulation such that

$$\begin{aligned} V_0 \mapsto V_0(t) &= V_c + V_m \sin \Omega_m t \\ &= V_c (1 + m \sin \Omega_m t) \end{aligned} \quad (\text{C.2})$$

where $m := V_m/V_c$ is the *modulation index*. The modulated signal is

$$V(t) = V_c (1 + m \sin \Omega_m t) \sin \Omega t \quad (\text{C.3})$$

The result of the modulation is thus a sinusoid at Ω multiplied by an envelope varying sinusoidally between $V_c (1 \pm m)$ at an (angular) frequency of Ω_m .

Using the trigonometric identity,

$$\sin x \sin y = \frac{1}{2} \cos (x - y) - \frac{1}{2} \cos (x + y) \quad (\text{C.4})$$

we can re-write (C.3) as

$$V(t) = V_c \left\{ \sin \Omega t + \frac{1}{2} m \cos(\Omega - \Omega_m) t - \frac{1}{2} m \cos(\Omega + \Omega_m) t \right\} \quad (C.5)$$

Assuming a load impedance R , we can convert this to the time-averaged power equation

$$\begin{aligned} P &= P_C + P_{\text{LSB}} + P_{\text{USB}} \\ &= \frac{V_c^2}{2R} \left(1 + \frac{m^2}{4} + \frac{m^2}{4} \right) \end{aligned} \quad (C.6)$$

where P_C is the power in the carrier wave at Ω , P_{LSB} and P_{USB} are the powers in the lower and upper sidebands at $\Omega \mp \Omega_m$ respectively and we have used the orthogonality relationship

$$\begin{aligned} \langle \sin \Omega_1 t \sin \Omega_2 t \rangle &:= \lim_{T \rightarrow \infty} \left\{ \frac{1}{T} \int_0^T dt' \sin \Omega_1 t' \sin \Omega_2 t' \right\} \\ &= \begin{cases} \frac{1}{2} & : \Omega_1 = \Omega_2 \\ 0 & : \Omega_1 \neq \Omega_2 \end{cases} \end{aligned} \quad (C.7)$$

Thus, the fraction of the power in the sidebands is given by

$$\frac{P_{\text{LSB}}}{P_C} = \frac{P_{\text{USB}}}{P_C} = \frac{m^2}{4} \quad (C.8)$$

C.1.1 Modulation at multiple frequencies

An important feature of amplitude modulation is that it is linear. Thus, if we add two modulations at different frequencies, such that

$$\begin{aligned} V(t) &= V_c (1 + m_1 \sin \Omega_{m_1} t + m_2 \sin \Omega_{m_2} t) \sin \Omega t \\ &= V_c \left\{ \sin \Omega t + \frac{1}{2} m_1 \cos(\Omega - \Omega_{m_1}) t - \frac{1}{2} m_1 \cos(\Omega + \Omega_{m_1}) t \right. \\ &\quad \left. + \frac{1}{2} m_2 \cos(\Omega - \Omega_{m_2}) t - \frac{1}{2} m_2 \cos(\Omega + \Omega_{m_2}) t \right\} \end{aligned} \quad (C.9)$$

then the power spectrum is given by

$$\begin{aligned} P &= \frac{V_c^2}{2R} \left(1 + \frac{m_1^2}{4} + \frac{m_1^2}{4} + \frac{m_2^2}{4} + \frac{m_2^2}{4} \right) \\ P &= P_C + P_{\text{LSB}_1} + P_{\text{USB}_1} + P_{\text{LSB}_2} + P_{\text{USB}_2} \end{aligned} \quad (C.10)$$

where we have again used the orthogonality relationship to remove cross-terms involving products of different frequencies. Thus, the sidebands from multiple modulations at different frequencies simply add.

We can similarly calculate the RMS of fractional amplitude jitter resulting from the modulations by

$$\begin{aligned} \left\langle \left(\frac{\Delta V}{V_c} \right)^2 \right\rangle &:= \left\langle \left(\frac{V_0(t) - V_c}{V_c} \right)^2 \right\rangle \\ &= \frac{1}{2} (m_1^2 + m_2^2) \\ &= \frac{2(P_{\text{SB}_1} + P_{\text{SB}_2})}{P_c} \end{aligned} \quad (C.11)$$




Figure C.1: Carrier with power P_c and frequency f_0 modulated by a continuous noise spectrum with long-term average power spectral density at a frequency offset from the carrier by f given by $N_0(f)$. Here, N_0 is shown as being flat (white) apart from very close to the carrier wave. We imagine the noise spectrum as arising from pairs of upper and lower sidebands due to discrete modulations with bandwidth δf .

where $P_{\text{SB}_i} = P_{\text{LSB}_i} = P_{\text{USB}_i}$ ($i = 1, 2$) is the power in one of the sidebands resulting from modulation i .

In the case of modulation at k different frequencies, this becomes

$$\left\langle \left(\frac{\Delta V}{V_c} \right)^2 \right\rangle = 2 \sum_{n=1}^k \frac{P_{\text{SB}_n}}{P_c} \quad (\text{C.12})$$

C.1.2 Modulation by a continuous spectrum

We now turn to the case of a signal modulated by a continuous noise spectrum, as illustrated by figure C.1. For the time being, we assume that the noise represents only amplitude modulation with no phase modulation present and hence that it is symmetric around the carrier. We denote the long-term average noise power spectral density (PSD) at a frequency offset from the carrier by f by $N_0(f)$.

It is useful to imagine the noise spectrum as the result of a series of modulation sidebands. We thus discretize the spectrum into regions of bandwidth δf and represent the regions between $f_0 \pm f_m$ and $f_0 \pm (f_m + \delta f)$ by the sidebands due to a modulation with frequency f_m .

Our previous analysis assumed that the modulation sidebands represented sinusoids with well-defined amplitudes and phases. However for a noise quantity the amplitude and phase are not constant, but fluctuate over time. Despite this, during a time $\delta t \ll \frac{1}{\Delta f}$ neither the amplitude nor the phase of the signal

can change appreciably, so the sidebands at any instant may be well represented by a sinusoidal signal of the form¹

$$v = \sqrt{2n_0} \sin(\Omega t \pm \Omega_m t - \phi) \quad (\text{C.13})$$

where n_0 and ϕ are the instantaneous power and phase of the wave. These quantities will vary from one time period to another, with values of n_0 distributed according to a Rayleigh distribution with mean $N_0(f) \delta f$, while ϕ has a uniform probability distribution over the interval $[0, 2\pi]$. When our interest is in the long term mean power contribution, rather than the power measured in any given time interval, we may write

$$V = \sqrt{2N_0} \sin(\Omega t \pm \Omega_m t) \quad (\text{C.14})$$

Bearing this in mind, we see from equation (C.8) that the index of modulation due to noise in the interval between $f_0 \pm f_m$ and $f_0 \pm (f_m + \delta f)$ is

$$m^2 = 4 \frac{P_{SB}}{P_c} := 4 \frac{N_0(f_m) \delta f}{P_c} \quad (\text{C.15})$$

where P_{SB} is the power in either the upper or lower noise sideband.

We can calculate the total amplitude jitter due to modulations with frequencies in the range $f_{\min} < f_m < f_{\max}$ using equation (C.12). In the limit $\delta f \rightarrow 0$, this sum becomes an integral and we find the important relationship

$$\left\langle \left(\frac{\Delta V}{V_c} \right)^2 \right\rangle = 2 \int_{f_{\min}}^{f_{\max}} \frac{N_0(f)}{P_c} df \quad (\text{C.16})$$

For the special case of white amplitude noise ($N_0(t) = N_0$) over a bandwidth $\Delta f := f_{\max} - f_{\min}$ this reduces to

$$\left\langle \left(\frac{\Delta V}{V_c} \right)^2 \right\rangle = \sqrt{\frac{2N_0 \Delta f}{P_c}} \quad (\text{C.17})$$

If the system has a first-order low-pass response with $f_{3\text{dB}} \ll f_{\max}$ then we can use equation (C.17) with $\Delta f = 1.571 f_{3\text{dB}}$, where the correction factor arises from the integral of the Lorentzian lineshape of the low-pass response (see e.g. [MC93] for a derivation).

C.1.3 The Wiener-Khinchin theorem

A further quantity of fundamental interest is the signal's autocorrelation function, defined by

$$\begin{aligned} G(\tau) &:= \langle V(t + \tau)V(t) \rangle \\ &:= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T V(t + \tau)V(t) dt \end{aligned} \quad (\text{C.18})$$

¹Here, we have assumed without loss of generality that the load impedance is 1Ω .

The autocorrelation function is closely related to the PSD, which we recall is defined as [Gar90]

$$\begin{aligned} P(f) &:= \lim_{T \rightarrow \infty} \frac{1}{T} |V_T(f)|^2 \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \int_0^T V(t) V^*(t') e^{-2\pi i f(t-t')} dt dt' \end{aligned} \quad (\text{C.19})$$

where

$$V_T(f) := \int_0^T e^{-2\pi i f t} V(t) dt \quad (\text{C.20})$$

Defining $\tau := t' - t$, we have

$$\begin{aligned} P(f) &= \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \int_0^T V(t) V(t + \tau) e^{2\pi i f \tau} dt d\tau \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \left\{ \int_0^T e^{2\pi i f \tau} d\tau \int_0^{T-\tau} dt V(t) V(t + \tau) + \int_{-T}^0 e^{2\pi i f \tau} d\tau \int_{-\tau}^T dt V(t) V(t + \tau) \right\} \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \left\{ \int_0^T e^{2\pi i f \tau} d\tau \int_0^{T-\tau} dt V(t) V(t + \tau) + \int_0^T e^{-2\pi i f \tau} d\tau \int_0^{T-\tau} dt V(t + \tau) V(t) \right\} \\ &= 2 \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \cos(2\pi f \tau) d\tau \int_0^{T-\tau} dt V(t) V(t + \tau) \\ &= 2 \int_0^\infty \cos(2\pi f \tau) G(\tau) d\tau \end{aligned} \quad (\text{C.21})$$

Since G is an even function, we may re-write this as

$$P(f) = \int_{-\infty}^\infty e^{2\pi i f \tau} G(\tau) d\tau \quad (\text{C.22})$$

taking the inverse Fourier transform,

$$G(\tau) = \int_{-\infty}^\infty e^{-2\pi i f \tau} P(f) df \quad (\text{C.23})$$

Equations (C.22) & (C.23), known as the Wiener-Khinchin relations (see, e.g. [Gar90] or [SRS06]), demonstrate that the autocorrelation function and power spectral density are a Fourier transform pair. This means that a signal's autocorrelation function contains all possible information about its statistical properties.

We can apply these relations to the fluctuating (noise) component of the signal to show that

$$\left\langle \frac{\Delta V(t + \tau) \Delta V(t)}{V_c^2} \right\rangle = \int_{-\infty}^\infty e^{-2\pi i f \tau} \frac{N_0(f)}{P_c} df \quad (\text{C.24})$$

For $\tau = 0$, and assuming a symmetric noise spectrum, this becomes

$$\left\langle \left(\frac{\Delta V}{V_c} \right)^2 \right\rangle = 2 \int_{f_{\min}}^{f_{\max}} df \frac{N_0(f)}{P_c} \quad (\text{C.25})$$

in agreement with equation (C.16).

For non-zero τ , this tells us about the coherence time of the fluctuations. We will generally be interested in single-pole filtered white noise, in which case

$$\begin{aligned} \left\langle \frac{\Delta V(t + \tau) \Delta V(t)}{V_c^2} \right\rangle &= \frac{N_0}{P_c} \int_{-\infty}^{\infty} e^{-2\pi i f \tau} \frac{1}{1 + (f/f_{3\text{dB}})^2} df \\ &= \frac{1}{2\tau_c} \frac{N_0}{P_c} e^{-|\tau|/\tau_c} \end{aligned} \quad (\text{C.26})$$

where $\tau_c := 1/(2\pi f_{3\text{dB}})$ is the correlation time of the noise.

For white noise, we find

$$\begin{aligned} \left\langle \frac{\Delta V(t + \tau) \Delta V(t)}{V_c^2} \right\rangle &= \frac{N_0}{P_c} \int_{-\infty}^{\infty} e^{-2\pi i f \tau} df \\ &= \frac{N_0}{P_c} \delta(\tau) \end{aligned} \quad (\text{C.27})$$

C.1.4 Coherent demodulation

In order to measure the amplitude noise spectrum of a device, we will usually use a non-linear device to coherently demodulate the signal. In this section we demonstrate the important result that when we coherently demodulate a signal, the voltages due to each sideband add in phase at the detector output. Thus, the baseband power due to a pair of noise sidebands is four times that which would be measured due to a single sideband. This is a significant correction factor when making coherent AM measurements.

We begin with the standard form of an amplitude modulated signal

$$V(t) = V_c \left\{ \sin \Omega t + \frac{1}{2} m \cos(\Omega - \Omega_m) t - \frac{1}{2} m \cos(\Omega + \Omega_m) t \right\} \quad (\text{C.28})$$

We next consider applying this signal to a detector, whose output at any instant is proportional to the instantaneous square of the signal voltage. We assume that the detector's output has a low-pass filter response, which removes any terms with frequencies close to the carrier or its harmonics. We will also assume that that $m \ll 1$ so that we can ignore terms of second order in m . Thus,

$$\begin{aligned} V^2(t) &= V_c^2 \left[\sin^2 \Omega t + m \sin \Omega t \{ \cos(\Omega - \Omega_m) t - \cos(\Omega + \Omega_m) t \} \right] \\ &= V_c^2 \left[\sin^2 \Omega t + m \{ \sin \Omega t \cos \Omega t \cos \Omega_m t + \sin^2 \Omega t \sin \Omega_m t \} \right. \\ &\quad \left. - m \{ \sin \Omega t \cos \Omega t \cos \Omega_m t - \sin^2 \Omega t \sin \Omega_m t \} \right] \\ &= V_c \sin^2 \Omega t [1 + m \sin \Omega_m t + m \sin \Omega_m t] \end{aligned} \quad (\text{C.29})$$

Now, $\sin^2 \Omega t = \frac{1}{2} (1 - \cos 2\Omega t)$. However, the term at 2Ω will be filtered out, leaving only the D.C. term. Thus

$$V^2(t) = \frac{1}{2} V_c^2 [1 + m \sin \Omega_m t + m \sin \Omega_m t] \quad (\text{C.30})$$

The three terms in (C.30) represent the powers in the carrier and two sidebands. The important thing to notice here is that the voltages due to the two sidebands add in phase. Thus, the voltage we will measure at the output of the detector will be twice that which we would have measured with only a single sideband. The measured baseband power will thus be increased by a factor of 4 (6 dB) compared with a single sideband.

This result also applies to coherent demodulation of a PM signal. A full derivation of the corresponding result can be found in [Rob84], which I will not repeat here. For a discussion of some methods of coherent AM and PM detection, see [HP174].

C.2 Phase modulation

We now consider a signal whose phase is sinusoidally modulated such that

$$V(t) = V_0 \sin \{ \Omega t + h \sin(\Omega_m t) \} \quad (\text{C.31})$$

where h is the modulation index, which is equal to the peak phase deviation in radians. Since the signal's phase, ϕ , is related to its instantaneous frequency, $\Omega(t)$, by $\dot{\phi} = \Omega$ this phase modulation is equivalent a modulation of the signal's frequency by

$$\Omega(t) = \Omega + h\Omega_m \cos \Omega_m t \quad (\text{C.32})$$

Our aim will be to re-cast the time-dependent modulation of (C.31) in terms of frequency-domain sidebands. We begin by using the trigonometric identity

$$\sin(x + y) = \sin x \cos y + \cos x \sin y \quad (\text{C.33})$$

to re-write (C.31) as

$$V(t) = V_0 \sin(\Omega t) \cos \{ h \sin(\Omega_m t) \} + V_0 \cos(\Omega t) \sin \{ h \sin(\Omega_m t) \} \quad (\text{C.34})$$

Next, we use the Bessel function identities²

$$\begin{aligned} \cos(z \sin x) &= \sum_{k=-\infty}^{\infty} J_k(z) \cos(kx) \\ \sin(z \sin x) &= \sum_{k=-\infty}^{\infty} J_k(z) \sin(kx) \end{aligned} \quad (\text{C.35})$$

where J_k is the k^{th} Bessel function of the first kind. Substituting these into (C.34) to get

$$\begin{aligned} V(t) &= V_0 \sin(\Omega t) \sum_{k=-\infty}^{\infty} J_k(h) \cos(k\Omega_m t) \\ &+ V_0 \cos(\Omega t) \sum_{k=-\infty}^{\infty} J_k(h) \sin(k\Omega_m t) \\ &= V_0 \sum_{k=-\infty}^{\infty} J_k(h) [\sin(\Omega t) \cos(k\Omega_m t) + \cos(\Omega t) \sin(k\Omega_m t)] \\ &= V_0 \sum_{k=-\infty}^{\infty} J_k(h) \sin[(\Omega + k\Omega_m)t] \end{aligned} \quad (\text{C.36})$$

²See, for example, [RHB02] pp 574.

Where, in the last line, we have re-used the trigonometric identity of (C.33). Next, we use the identity $J_{-k}(x) = (-1)^k J_k(x)$ to re-write this as

$$\begin{aligned} V(t) &= V_0 J_0(h) \sin(\Omega_c t) \\ &+ V_0 \sum_{k=1}^{\infty} J_k(h) \sin[(\Omega + k\Omega_m)t] \\ &+ V_0 \sum_{k=1}^{\infty} (-1)^k J_k(h) \sin[(\Omega - k\Omega_m)t] \end{aligned} \quad (\text{C.37})$$

we thus get a carrier and a series of sidebands, separated by the modulation frequency. The power spectrum is given by

$$\begin{aligned} P(t) &= P_{\text{carrier}} + \sum_{k=1}^{\infty} P_{\text{USB}_k} + \sum_{k=1}^{\infty} P_{\text{LSB}_k} \\ P(t) &= P_c J_0^2(h) + P_c \sum_{k=1}^{\infty} J_k^2(h) + P_c \sum_{k=1}^{\infty} J_k^2(h) \end{aligned} \quad (\text{C.38})$$

where, $P_c = \frac{V_0^2}{2R}$, is the power of the carrier and, P_{USB_k} and P_{LSB_k} are the powers of the k^{th} upper and lower sidebands respectively.

We note that the modulation of equation (C.37) is non-linear – if we double the modulation index, we do not double the sideband amplitude. However, in many cases we will be concerned with small modulations, where $h \ll 1$. In this case, one may expand the Bessel functions to first order so that

$$\begin{aligned} \lim_{x \rightarrow 0} J_k(x) &= \lim_{x \rightarrow 0} \left[x^k \sum_{m=0}^{\infty} \frac{(-1)^m x^{2m}}{2^{2m+k} m! (k+m)!} \right] \\ &= \lim_{x \rightarrow 0} \left[\frac{x^k}{2^k k!} - \frac{x^{k+2}}{2^{k+2} (k+1)!} + O(x^{k+4}) \right] \end{aligned} \quad (\text{C.39})$$

Hence

$$\lim_{x \rightarrow 0} J_k^2(x) = \lim_{x \rightarrow 0} \left[\frac{x^{2k}}{(2^k k!)^2} + O(x^{2k}) \right] \quad (\text{C.40})$$

$$\begin{aligned} J_0^2(x) &\sim 1 \\ J_1^2(x) &\sim \frac{x^2}{4} \end{aligned} \quad (\text{C.41})$$

where “ $\sim$ ” should be read as “asymptotically equal to” and we have neglected all but the lowest order terms in x .

We thus obtain the low-modulation limit of (C.38)

$$\begin{aligned} P(t) &= P_{\text{carrier}} + P_{\text{USB}} + P_{\text{LSB}} \\ &= P_c \left(1 + \frac{h^2}{4} + \frac{h^2}{4} \right) \end{aligned} \quad (\text{C.42})$$

This is linear and is identical in form to equation (C.6). We note that the power in the sidebands is given by

$$\frac{P_{\text{LSB}}}{P_c} = \frac{P_{\text{USB}}}{P_c} = \frac{h^2}{4} \quad (\text{C.43})$$

For the rest of this appendix, we will assume this small modulation approximation.

C.2.1 PM at multiple frequencies

We can calculate the RMS phase fluctuations due to a pair of modulations with frequencies Ω_{m_1} and Ω_{m_2} and modulation indices h_1 and h_2 as follows

$$\begin{aligned}
 (\delta\phi)^2 &= \langle (\phi(t) - \Omega t)^2 \rangle \\
 &= \langle (h_1 \sin \Omega_{m_1} t + h_2 \sin \Omega_{m_2} t)^2 \rangle \\
 &= \frac{1}{2} (h_1^2 + h_2^2) \\
 &= 2 \frac{P_{SB_1} + P_{SB_2}}{P_c}
 \end{aligned} \tag{C.44}$$

In the case of modulation at k different frequencies, this becomes

$$(\delta\phi)^2 = 2 \sum_{n=1}^k \frac{P_{SB_n}}{P_c} \tag{C.45}$$

Similarly, the RMS frequency fluctuation due to a the pair of modulations is given by

$$\begin{aligned}
 (\delta\Omega)^2 &= \langle (\Omega(t) - \Omega)^2 \rangle \\
 &= \langle (h_1 \Omega_{m_1} \cos \Omega_{m_1} t + h_2 \Omega_{m_2} \cos \Omega_{m_2} t)^2 \rangle \\
 &= \frac{1}{2} \left\{ (h_1 \Omega_{m_1})^2 + (h_2 \Omega_{m_2})^2 \right\} \\
 &= 2 \frac{P_{SB_1} \Omega_{m_1}^2 + P_{SB_2} \Omega_{m_2}^2}{P_c}
 \end{aligned} \tag{C.46}$$

In the case of modulation at k different frequencies, this becomes

$$(\delta\Omega)^2 = 2 \sum_{n=1}^k \frac{P_{SB_n} \Omega_{m_n}^2}{P_c} \tag{C.47}$$

C.2.2 Modulation by a continuous spectrum

Following the above treatment for AM noise, we now consider a signal modulated by a continuous PM noise spectrum. We find the modulation index due to noise in the interval δf around $\Omega + \Omega_m$ is

$$h(\Omega_m)^2 = 4 \frac{P_{SB}}{P_c} = 4 \frac{\delta f N_0}{P_c} \tag{C.48}$$

The RMS phase fluctuation due to modulations with frequency between $f_{\min}$ and $f_{\max}$ is given by

$$(\delta\phi)^2 = \frac{2}{P_c} \int_{f_{\min}}^{f_{\max}} N_0(f) df \tag{C.49}$$

for white phase noise,

$$(\delta\phi)^2 = \frac{2N_0\Delta f}{P_c} \tag{C.50}$$

Similarly, the frequency jitter due to a noise power density $N_0(f)$ is

$$(\delta\Omega)^2 = \frac{2}{P_c} \int_{f_{\min}}^{f_{\max}} (2\pi f)^2 N_0(f) df \tag{C.51}$$

We note that the frequency jitter due to a given noise power density increases with the modulation frequency. Physically, this represents the fact that as the modulation frequency increases, it takes a larger frequency shift to produce a given phase deviation since the phase has less time to accumulate over the course of a cycle.

C.2.3 PM Wiener-Khinchin relations

We now wish to derive Wiener-Khinchin relations for the autocorrelation function for phase fluctuations, defined as

$$\begin{aligned} G_{\delta\phi}(\tau) &:= \langle \delta\phi(t+\tau)\delta\phi(t) \rangle \\ &= \langle [\phi(t+\tau) - \Omega \times (t+\tau)] [\phi(t) - \Omega t] \rangle \end{aligned} \quad (\text{C.52})$$

We begin, as usual, with the case of a signal modulated at a single frequency, for which

$$\begin{aligned} G_{\delta\phi}(\tau) &= \langle [h \sin \Omega_m (t+\tau)] [h \sin \Omega_m t] \rangle \\ &= \frac{1}{2} h^2 \langle \cos \Omega_m \tau - \cos \Omega_m (2t+\tau) \rangle \\ &= \frac{1}{2} h^2 \cos \Omega_m \tau \end{aligned} \quad (\text{C.53})$$

For a signal modulated at a finite number of frequencies, this becomes

$$G_{\delta\phi}(\tau) = \sum_{i,j} \langle [h_i \sin \Omega_{m_i} (t+\tau)] [h_j \sin \Omega_{m_j} t] \rangle \quad (\text{C.54})$$

all terms involving products of different frequencies average to zero, leaving

$$G_{\delta\phi}(\tau) = \sum_i \frac{1}{2} h_i^2 \cos \Omega_{m_i} \tau \quad (\text{C.55})$$

In the limit of modulation by a continuous (but symmetric) spectrum, this becomes

$$\begin{aligned} G_{\delta\phi}(\tau) &= \frac{2}{P_c} \int_0^\infty \cos(2\pi f \tau) N_0(f) df \\ &= \frac{1}{P_c} \int_0^\infty (e^{2\pi i f \tau} + e^{-2\pi i f \tau}) N_0(f) df \\ &= \frac{1}{P_c} \int_0^\infty e^{-2\pi i f \tau} N_0(f) df + \frac{1}{P_c} \int_{-\infty}^0 e^{-2\pi i f \tau} N_0(f) df \\ &= \int_{-\infty}^\infty e^{-2\pi i f \tau} \frac{1}{P_c} N_0(f) df \end{aligned} \quad (\text{C.56})$$

Taking the inverse Fourier transform, we find

$$\frac{1}{P_c} N_0(f) = \int_{-\infty}^\infty G_{\delta\phi}(\tau) e^{2\pi i f \tau} d\tau \quad (\text{C.57})$$

Equations (C.56) & (C.57) are our Wiener-Khinchin relations for phase modulation. By comparing them to (C.24), we see that

$$G_{\delta\phi} = \frac{1}{V_c^2} \langle \Delta V(t+\tau) \Delta V(t) \rangle \quad (\text{C.58})$$

In other words, the autocorrelation function for phase fluctuations is the same as that for voltage fluctuations, normalised to the carrier amplitude.

For $\tau = 0$, (C.56) reduces to

$$\langle (\delta\phi)^2 \rangle = 2 \int_0^\infty \frac{1}{P_c} N_0(f) df \quad (\text{C.59})$$

in agreement with (C.49). When N_0 has a single-pole spectrum, $G_{\delta\phi}$ becomes

$$\begin{aligned} G_{\delta\phi}(\tau) &= \frac{N_0}{P_c} \int_{-\infty}^{\infty} e^{-2\pi i f \tau} \frac{1}{1+(f/f_{3\text{ dB}})^2} df \\ &= \frac{1}{2\tau_c} \frac{N_0}{P_c} e^{-|\tau|/\tau_c} \end{aligned} \quad (\text{C.60})$$

where $\tau_c := 1/(2\pi f_{3\text{ dB}})$ is the noise correlation time.

For frequency fluctuations, we have

$$G_{\delta f}(\tau) := \langle [f(t+\tau) - f_c] [f(t) - f_c] \rangle \quad (\text{C.61})$$

For modulation at a single frequency

$$\begin{aligned} G_{\delta f}(\tau) &= (hf_m)^2 \langle \cos[2\pi f_m(t+\tau)] \cos 2\pi f_m t \rangle \\ &= \frac{1}{2} (hf_m)^2 \cos(2\pi f_m \tau) \end{aligned} \quad (\text{C.62})$$

For modulation by a continuous spectrum we have

$$\begin{aligned} G_{\delta f}(\tau) &= 2 \frac{1}{P_c} \int_0^\infty N_0(f) f^2 \cos(2\pi f \tau) df \\ &= \int_{-\infty}^{\infty} \frac{1}{P_c} N_0(f) f^2 e^{-2\pi i f \tau} df \end{aligned} \quad (\text{C.63})$$

This result is unsurprising, given that

$$\begin{aligned} G_{\delta f}(\tau) &= \frac{1}{(2\pi)^2} G_{\delta\dot{\phi}}(\tau) \\ &= \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} |\delta\dot{\phi}(f)| e^{-2\pi i f \tau} df \\ &= \int_{-\infty}^{\infty} f^2 |\delta\phi(f)|^2 e^{-2\pi i f \tau} df \\ &= \int_{-\infty}^{\infty} \frac{1}{P_c} N_0(f) f^2 e^{-2\pi i f \tau} df \end{aligned} \quad (\text{C.64})$$

C.2.4 Frequency multiplication

In RF systems, it is often necessary to increase the frequency of a signal by frequency multiplication. This has the effect of multiplying the signal's phase by a factor n . For a phase modulated signal, this results in

$$V_c \sin(\Omega t + h \sin \Omega_m t) \xrightarrow{\times n} V_c \sin(n\Omega t + nh \sin \Omega_m t) \quad (\text{C.65})$$

We thus see that frequency multiplication by a factor of n increases the PM modulation index by the same factor. From the previous results, we see that this increases the power in the PM sidebands by a factor of n^2 ($20 \log_{10} n$ dB). We note that the transformation of (C.65) has no any effect on any AM present. In practice, frequency multiplication may increase the level of AM or, in so far as the multiplier acts as limiter, may suppress it entirely. Such effects are, however, not fundamental to the process and are purely an artefact of the particular multiplier used.

C.3 The relationship between AM and PM

Thus far we have treated AM and PM as independent phenomena. We have shown that modulated signals may be expressed as

$$\begin{aligned} AM : V(t) &= V_c \left\{ \sin \Omega t + \frac{1}{2} m \cos (\Omega - \Omega_m) t - \frac{1}{2} m \cos (\Omega + \Omega_m) t \right\} \\ PM : V(t) &= V_c \left\{ \sin \Omega t - \frac{1}{2} h \sin (\Omega - \Omega_m) t + \frac{1}{2} h \sin (\Omega + \Omega_m) t \right\} \end{aligned} \quad (C.66)$$

These equations make clear the important point that the only difference between AM and PM is the phase relationship³ between the sidebands and the carrier – if we removed the carrier from a modulated waveform then we would not be able to say whether the modulation represented AM or PM. Indeed, by re-inserting a carrier with appropriate phase we would be able to convert the modulation into pure PM, pure AM or a mixture of the two.

We note that pure AM or pure PM will always produce a modulation spectrum which is symmetric about the carrier. Any asymmetry can only result from a combination of AM and PM with a definite phase relationship. This will not happen for most noise processes, where the fluctuations have no definite phase.

It is useful to note that in the case of high-quality quartz oscillators, the AM and PM spectra are usually very closely related in the f^0 (white phase noise⁴) and f^{-1} (flicker PM), with significant differences only appearing very close to the carrier in the f^{-3} (flicker FM) and f^{-4} (random walk FM) parts of the spectrum [NNW94]. The measurements in [NNW94] found that the AM and PM noise of 5 MHz quartz crystal oscillators was similar for frequencies above 3 Hz. For 100 MHz oscillators, the PM and AM were similar above approximately 300 Hz. This close connection suggests that they share a common origin, which is most likely the sustaining amplifier used in the oscillator circuit.

³While the sidebands and carrier are at different frequencies and hence are, strictly speaking, uncorrelated they nonetheless have a well-definite phase relationship with each other. [Rob84] terms this relationship “conformability”.

⁴For a discussion of the different regions of the PM spectrum of a typical oscillator see, e.g. [LSL84]

C.4 Modulation due to noise

We begin by considering a single noise sideband superimposed on top of a carrier with detuning Ω_n

$$V(t) = V_c \sin \Omega t + V_n \sin [(\Omega + \Omega_n) t + \phi] \quad (\text{C.67})$$

here, $V_n := \sqrt{2N_0\delta f}$ is the noise amplitude (where N_0 and δf are the noise PSD and bandwidth) and ϕ is its phase. We may re-express this as

$$\begin{aligned} V(t) &= V_c \sin \Omega t + V_n [\sin \Omega t \cos (\Omega_n t + \phi) + \cos \Omega t \sin (\Omega_n t + \phi)] \\ &= \sin \Omega t [V_c + V_n \cos (\omega_n t + \phi)] + V_n \cos \Omega t \sin (\Omega_n t + \phi) \end{aligned} \quad (\text{C.68})$$

We now use the trigonometric identity

$$a \sin x + b \cos x = \sqrt{a^2 + b^2} \cdot \sin (x + \theta) \quad (\text{C.69})$$

where (when $a \geq 0$)

$$\begin{aligned} \tan \theta &= \frac{b}{a} \\ &= \frac{V_n \sin (\Omega_n t + \phi)}{V_c + V_n \cos (\Omega_n t + \phi)} \\ &\sim \frac{V_n}{V_c} \sin (\Omega_n t + \phi) \end{aligned} \quad (\text{C.70})$$

so

$$\theta \sim \frac{V_n}{V_c} \sin (\Omega_n t + \phi) \quad (\text{C.71})$$

Since $V_n \ll V_c$,

$$\sqrt{a^2 + b^2} \sim V_c + V_n \cos (\omega_n t + \phi) \quad (\text{C.72})$$

Defining the modulation indices $m = h = \frac{V_n}{V_c}$ equation (C.68) becomes

$$\begin{aligned} V(t) &= V_c [1 + m \cos (\omega_n t + \phi)] \sin [\Omega t + h \sin (\Omega_n t + \phi)] \\ &= V_c [1 + m \cos (\omega_n t + \phi)] \left\{ \sin \Omega t - \frac{1}{2} h \sin (\Omega - \Omega_n) t \right. \\ &\quad \left. + \frac{1}{2} h \sin (\Omega + \Omega_n) t \right\} \end{aligned} \quad (\text{C.73})$$

Dropping terms of order mh , this becomes

$$\begin{aligned} V(t) &= V_c \sin \Omega t \\ &\quad + V_c \left\{ -\frac{1}{2} h \sin (\Omega - \Omega_n) t + \frac{1}{2} h \sin (\Omega + \Omega_n) t \right\} \\ &\quad + V_c \left\{ \frac{1}{2} m \cos (\Omega - \Omega_n) t - \frac{1}{2} m \cos (\Omega + \Omega_n) t \right\} \end{aligned} \quad (\text{C.74})$$

We thus see that we can represent a single noise sideband as a combination of four modulation sidebands consisting of a pair of upper and lower PM and AM sidebands. From our previous results, we see that the power in each sideband is given by $\frac{1}{4}N_0\delta f$. Thus, the power is evenly shared between the four sidebands.

In the more usual case where, rather than a single noise sideband, we have a symmetric noise spectrum around the carrier, our noise consists of pairs of uncorrelated sidebands with detunings $\pm\Omega_n$. In this case, we can combine the

modulation sidebands due to the upper and lower noise sidebands. Since the phases of the two noise sidebands are uncorrelated, the phases of the modulation sidebands they produce are also uncorrelated. Thus, when we superpose them, their powers add. Thus, the power in each of the four modulation sidebands is equal to half of the single-sideband noise power.

The fundamental noise floor of any system is set by the thermal noise of $N_0 = k_B T = -174 \text{ dBm/Hz}$ [Rob84]. This produces an AM and PM noise floor of -177 dBm/Hz .

D

Electronic Circuits

D.1 7400 limiting amplifier

In this section we describe design and preliminary characterisation of a prototype limiting amplifier intended to remove AM noise from a signal source (see section 6.5.3). The circuit is based on the use of a 7400 series inverter as a saturated amplifier, which converts an input sine wave into a square wave with constant amplitude.

The RF signal is capacitively coupled into the input of an inverter that has been biased into its high-gain linear region. Further inverters then act to ensure a sharp rising edge that is not limited by the input signal. The output square wave is then matched to $50\ \Omega$ using a π network¹. Once matched to $50\ \Omega$, the square wave is filtered using a somewhat half-heartedly constructed nodal capacitor-coupled bandpass filter to produce our final signal.

The circuit is powered from a relatively low-noise regulator constructed using a transistor, OpAmp and voltage reference. At DC, the OpAmp feeds back on the regulator's output voltage through R_F to maintain a constant voltage level. At AC, feedback is through C_F . R_x isolates the feedback network from the decoupling and filtering capacitors to ensure loop stability, creating a pole at 320 Hz with the $10\ \mu\text{F}$ capacitor.

At 40 MHz, the amplifier puts out approximately 3 dBm of power when driven with an input power of $\gtrsim -10\ \text{dBm}$. This low power level is mainly due to 7 dB of loss in the filter, which was not tuned to optimise transmission.

One may obtain a crude estimate of the regulator's impact on the overall amplifier noise by assuming its noise is directly present on the RF output. At the inverter's output, there is a 2.5 V amplitude sine wave (plus harmonics, which have yet to be filtered), so $25\ \text{nV}/\sqrt{\text{Hz}}$ of noise corresponds to a noise power density of $-160\ \text{dBc}/\text{Hz}$. We measured the regulator noise using a low-noise amplifier based on an AD8099 OpAmp in the non-inverting configuration. The amplifier has a gain of 1,000 into a high-impedance load and 3 dB bandwidths of 10 Hz and 185 kHz. We measure a noise voltage of $1\ \mu\text{V}$ in this bandwidth, equivalent to an integrated noise power of $-130\ \text{dBc}$. Even if this noise were all

¹For a discussion of matching networks, see for example the Motorola application note [Dav93].

contained in a bandwidth of 100 Hz around the carrier, it would correspond to a noise power density of -150 dBc/Hz, which would be sufficient for our present purposes².

It is hard to estimate the noise contribution due to the final inverter. We have, however, measured its static performance when driven to its “on” and “off” states using a DC input. We measure an integrated noise of $\lesssim 1\text{ }\mu\text{V}$ when the amplifier is in its off (low-voltage) state and $1.3\text{ }\mu\text{V}$ when in the on state. Based on these numbers, we expect the noise density to be extremely low over our bandwidth of interest – if the $1.3\text{ }\mu\text{V}$ of noise were purely white, it would correspond to a density of $\lesssim -170\text{ dBc/Hz}$ compared with a 2.5 V amplitude carrier. However, this number has limited significance since at 40 MHz the inverter spends a significant fraction of the RF cycle in the linear region between supply rails, where it is liable to introduce a significant level of noise.

A characterisation of the AM noise of this amplifier at 40 MHz is given in section 6.5.3.

²Clearly, this measurement is somewhat unsatisfactory and one would like to know the noise spectral density in this regime. That particular measurement, however, will have to wait for another day.


D.2 RF levelling loop


The frequency of the trap's radial modes depends linearly on the amplitude of the RF trapping voltage. For the microwave-driven two qubit gate, we require our 3 MHz radial modes to be stable to $\lesssim 10$ Hz (section 6.7), equivalent to a fractional RF amplitude stability of 3 ppm.


To achieve this, we use a directional coupler (MiniCircuits ZFDC-20-3-S+) to pick-off a sample of the RF power before the resonator. We then monitor this power and feed back to the control voltage on a voltage-variable attenuator (MiniCircuits ZX73-2500+) using the circuit shown below. This is based around an Analog Devices ADL5511 power detector and a PI loop constructed from OpAmps. The directional coupler and circuit are contained in the same box and temperature stabilised to ~ 1 mK.


To characterise the loop's stability, we measured the RF output power using a second power meter (AD8361, not temperature stabilised). Using this detector, we measured RF amplitude drifts of $\lesssim 30$ ppm over a period of several hours. Since the second detector was not temperature stabilised, this should be considered an upper limit on the drift.


We also measured the loop's temperature coefficient by measuring the output power at both its 34.38°C normal operating temperature and at 32°C . We observed a change in the RMS RF amplitude of $+4 \times 10^{-4}$ between these temperatures, corresponding to a temperature coefficient for the levelling loop of $170 \text{ ppm}/^\circ\text{C}$. Given that our box is stabilised to ~ 1 mK, we expect this should be sufficient to meet our stability target.














E

Atomic structure of $^{43}\text{Ca}^+$

E.1 Atomic constants

At various points in this thesis we make use of accurate calculations of ground-level transition energies, which are calculated using measured values of atomic constants. For convenience, the values of the used constants are collected in table E.1. All values of fundamental constants assume the CODATA 2012 values which may be found, for example, on the NIST website. A list of other $^{40}\text{Ca}^+$ and $^{43}\text{Ca}^+$ atomic constants may be found in [Szw09].

In any discussion of laser intensities, the following definition of the saturation intensity was assumed

$$I_0 := \frac{\hbar\omega^3\Gamma}{6\pi c^2} \quad (\text{E.1})$$

where ω is the transition's frequency (radians per second), Γ is its natural linewidth (radians per second) and c is the speed of light. Notice that according to this definition, when a cycling transition is driven resonantly with one saturation intensity, the rates of stimulated and spontaneous emission are equal.

E.2 Properties of the ground level

E.2.1 Expansion coefficients

In chapter 2, we demonstrated that at any magnetic field we can express states in the $4S_{\frac{1}{2}}$ level as a superposition of at most two of the high-field angular-

Symbol	Value	Notes	Reference
g_I	1.315 348 (9)	Nuclear g-factor	This work
g_J	2.002 256 64 (9)	Landé g-factor	[TPR ⁺ 03]
E_{HFS}	3 225 608 286.4 (3) Hz	Hyperfine structure splitting	[ABG ⁺ 94]

Table E.1: Atomic constants to do with $^{43}\text{Ca}^+$ used in this work.




Figure E.1: The $4S_{1/2}$ ground level of $^{43}\text{Ca}^+$ at 146 G. Despite being in the intermediate field, we label to states as $|F, M_F\rangle$ (see section 2.3). When discussing microwave transitions between the manifolds, I use $|M_n\rangle$ to refer to the state in the higher-energy $F = 3$ level and $|M_m\rangle$ to refer to the state in the $F = 4$ manifold. The corresponding matrix element is R_{nm} . The 146 G clock-transition between the states $|4, 0\rangle$ and $|3, 1\rangle$ is shown in blue. The green transitions are used to transfer the ion from the $|4, +4\rangle$ “stretched” state to the qubit after optical pumping. The π transitions shown in brown are used to improve the fidelity of optical pumping.

momentum eigenstates $|M_I, M_J\rangle$. We can thus write

$$|F, M\rangle_{n,m} = \alpha_{n,m} |M_I = M + \frac{1}{2}, M_J = -\frac{1}{2}\rangle + \beta_{n,m} |M_I = M - \frac{1}{2}, M_J = +\frac{1}{2}\rangle \quad (\text{E.2})$$

where $M := M_I + M_J$.

These coefficients are calculated by diagonalising the interaction Hamiltonian

$$\begin{aligned} H_I &= \mathbf{A}\mathbf{I} \cdot \mathbf{J} - \boldsymbol{\mu}_J \cdot \mathbf{B} - \boldsymbol{\mu}_I \cdot \mathbf{B} \\ &= A \left\{ I_z J_z + \frac{1}{2} (I_+ J_- + I_- J_+) \right\} + g_J \mu_B B J_z - g_I \mu_N B I_z \end{aligned} \quad (\text{E.3})$$

The result of this diagonalisation is tabulated in tables E.2 and E.3 for the magnetic fields 146.0942 G and 287.7827 G, corresponding to the two clock-transitions between the states $|4, 0\rangle \Leftrightarrow |3, +1\rangle$ and $|4, +1\rangle \Leftrightarrow |3, +1\rangle$ respectively. For each value of M apart from the $M = \pm 4$ stretched-states, there are two energy eigenstates. I use the subscript m to denote the state with lower energy,

corresponding to the low-field $F = 4$ manifold, and the subscript n to denote the higher-energy state corresponding to $F = 3$.

M	M_J	M_I	α_m	α_n	M_J	M_I	β_m	β_n
-4	-1/2	-7/2	1				0	
-3	-1/2	-5/2	0.948 3	0.317 5	+1/2	-7/2	0.317 5	-0.948 3
-2	-1/2	-3/2	0.890 6	0.454 8	+1/2	-5/2	0.454 8	-0.890 6
-1	-1/2	-1/2	0.825 5	0.564 5	+1/2	-3/2	0.564 5	-0.825 5
0	-1/2	+1/2	0.750 3	0.661 1	+1/2	-1/2	0.661 1	-0.750 3
+1	-1/2	+3/2	0.661 0	0.750 4	+1/2	1/2	0.750 4	-0.661 0
+2	-1/2	+5/2	0.549 7	0.835 4	+1/2	3/2	0.835 4	-0.549 7
+3	-1/2	+7/2	0.396 4	0.918 1	+1/2	5/2	0.918 1	-0.396 4
+4			0		+1/2	7/2	1	

Table E.2: Expansion coefficients for the $4S_{\frac{1}{2}}$ level at 146.094 2 G.

M	M_J	M_I	α_m	α_n	M_J	M_I	β_m	β_n
-4	-1/2	-7/2	1				0	
-3	-1/2	-5/2	0.957 6	0.288 0	+1/2	-7/2	0.288 0	-0.957 6
-2	-1/2	-3/2	0.909 6	0.415 5	+1/2	-5/2	0.415 5	-0.909 6
-1	-1/2	-1/2	0.854 1	0.520 2	+1/2	-3/2	0.520 2	-0.854 1
0	-1/2	+1/2	0.788 2	0.615 4	+1/2	-1/2	0.615 4	-0.788 2
+1	-1/2	+3/2	0.707 1	0.707 1	+1/2	+1/2	0.707 1	-0.707 1
+2	-1/2	+5/2	0.601 1	0.799 2	+1/2	+3/2	0.799 2	-0.601 1
+3	-1/2	+7/2	0.445 5	0.895 3	+1/2	+5/2	0.895 3	-0.445 5
+4			0		+1/2	+7/2	1	

Table E.3: Expansion coefficients for the $4S_{\frac{1}{2}}$ level at 287.782 7 G

E.3 Matrix elements

We define the matrix elements of the magnetic dipole Hamiltonian by

$$R_{nm} := (-1)^{Q+1} \langle n | \mu_Q | m \rangle \quad (\text{E.4})$$

Using this definition, we define the Rabi frequency on a transition to be

$$\hbar\Omega := B_{-Q} R_{M3,M4} \quad (\text{E.5})$$

Notice that according to this definition, a π -pulse takes a duration of π/Ω .

	M_m	M_n	Frequency (GHz)	df/dB_0 [d^2f/dB_0^2] (MHz/G)	R_{nm} (μ_B)
a	-4	-3	3.589 033 180	2.519 894	1.342 420
b	-3	-3	3.543 000 725	2.237 093	-0.602 802
c	-3	-2	3.495 825 733	1.940 390	+1.195 577
d	-2	-3	3.495 742 031	1.939 817	-0.204 417
e	-2	-2	3.448 567 039	1.643 113	+0.810 866
f	-2	-1	3.400 062 158	1.330 090	-1.040 756
g	-1	-2	3.399 978 455	1.329 517	-0.363 393
h	-1	-1	3.351 473 573	1.016 493	+0.932 839
i	-1	0	3.301 519 669	0.684 932	-0.876 791
j	0	-1	3.301 435 966	0.684 359	-0.528 271
k	0	0	3.251 482 062	0.352 797	+0.993 060
l	0	+1	3.199 941 077	0 [2.416 kHz/G²]	-0.702 127
m	+1	+0	3.199 857 374	-0.000 573	+0.702 254
n	+1	+1	3.148 316 389	-0.353 370	-0.993 034
o	+1	+2	3.095 026 842	-0.730 728	-0.514 410
p	+2	+1	3.094 943 139	-0.731 301	+0.887 366
q	+2	+2	3.041 653 592	-1.108 659	+0.919 344
r	+2	+3	2.986 424 605	-1.514 736	+0.308 504
s	+3	+2	2.986 340 903	-1.515 309	-1.085 679
t	+3	+3	2.931 111 915	-1.921 385	-0.728 641
u	+4	+3	2.873 631 427	-2.362 039	+1.299 654

Table E.4: Properties of transitions in the $4S_{\frac{1}{2}}$ ground-level of $^{43}\text{Ca}^+$ at 146.094 2 G. The qubit transition is shaded in blue, transitions used to prepare the qubit from the stretched-state are shaded green and transitions used in microwave enhanced optical pumping are shaded brown.

	M_m	M_n	Frequency (GHz)	df/dB_0 [d^2f/dB_0^2] (MHz/G)	R_{nm} (μ_B)
a	-4	-3	3.949 771 857	2.569 912	+1.355 692
b	-3	-3	3.867 368 475	2.337 131	-0.552 243
c	-3	-2	3.781 465 945	2.086 061	+1.233 104
d	-2	-3	3.781 301 064	2.085 488	-0.169 444
e	-2	-2	3.695 398 534	1.834 419	+0.756 705
f	-2	-1	3.605 294 276	1.560 347	-1.099 756
g	-1	-2	3.605 129 395	1.559 774	+0.305 995
h	-1	-1	3.515 025 138	1.285 701	-0.889 435
i	-1	0	3.420 035 468	0.982 944	+0.953 011
j	0	-1	3.419 870 588	0.982 371	+0.453 176
k	0	0	3.324 880 918	0.679 614	-0.971 138
l	0	+1	3.224 113 796	0.340 094	+0.789 034
m	+1	0	3.223 948 915	0.339 521	+0.616 028
n	+1	+1	3.123 181 794	0 [2.514 kHz/G²]	-1.001 026
o	+1	+2	3.015 435 346	-0.388 298	+0.601 736
p	+2	+1	3.015 270 465	-0.388 871	-0.799 970
q	+2	+2	2.907 524 018	-0.777 169	-0.961 755
r	+2	+3	2.791 111 000	-1.233 169	-0.379 132
s	+3	+2	2.790 946 119	-1.233 742	-1.012 843
t	+3	+3	2.674 533 102	-1.689 743	-0.798 542
u	+4	+3	2.546 787 218	-2.246 218	-1.267 363

Table E.5: Properties of transitions in the $4S_{\frac{1}{2}}$ ground-level of $^{43}\text{Ca}^+$ at 287.7827 G. The qubit transition is shown in bold.

Bibliography

- [ABG⁺94] F. Arbes, M. Benzing, Th. Gudjons, F. Kurth, and G. Werth. Precise determination of the ground state hyperfine structure splitting of ^{43}Ca II. *Z. Phys. D*, 31(1):27–30, 1994.
- [Aga76] G. S. Agarwal. Exact Solution for the Influence of Laser Temporal Fluctuations on Resonance Fluorescence. *Phys. Rev. Lett.*, 37(21):1383–1386, 1976.
- [Aga78] G. S. Agarwal. Quantum statistical theory of optical-resonance phenomena in fluctuating laser fields. *Phys. Rev. A*, 18(4):1490, 1978.
- [AGH⁺11] D. T. C. Allcock, L. Guidoni, T. P. Harty, C. J. Ballance, M. G. Blain, A. M. Steane, and D. M. Lucas. Reduction of heating rate in a microfabricated ion trap by pulsed-laser cleaning. *New J. Phys.*, 13(12):123023, December 2011.
- [AHB⁺13] D. T. C. Allcock, T. P. Harty, C. J. Ballance, B. C. Keitch, N. M. Linke, D. N. Stacey, and D. M. Lucas. A microfabricated ion trap with integrated microwave circuitry. *Appl. Phys. Lett.*, 102(4):044103, 2013.
- [AHJ⁺11] D. T. C. Allcock, T. P. Harty, H. A. Janacek, N. M. Linke, C. J. Ballance, A. M. Steane, D. M. Lucas, R. L. Jarecki, S. D. Habermehl, M. G. Blain, D. Stick, and D. L. Moehring. Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap. *Appl. Phys. B*, 107(4):913–919, November 2011.
- [AKB⁺07] L. Aolita, K. Kim, J. Benhelm, C. F. Roos, and H. Häffner. High-fidelity ion-trap quantum computing with hyperfine clock states. *Physical Review A*, 76(4):040303, October 2007.
- [ALH⁺13] D. P. L. Aude Craik, N. M. Linke, T. P. Harty, C. J. Ballance, D. M. Lucas, A. M. Steane, and D. T. C. Allcock. Microwave control electrodes for scalable, parallel, single-qubit operations in a surface-electrode ion trap. *Appl. Phys. B*, November 2013.
- [All11] D. T. C. Allcock. *Surface-Electrode Ion Traps for Scalable Quantum Computing*. PhD thesis, 2011.
- [ASS⁺10] 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(5):053026, May 2010.
- [BB02] D. J. Berkeland and M. G. Boshier. Destabilization of dark states and optical spectroscopy in Zeeman-degenerate atomic systems. *Phys. Rev. A*, 65(3):033413, February 2002.
- [BCOR09] S. Blanes, F. Casas, J. A. Oteo, and J. Ros. The Magnus expansion and some of its applications. *Physics Reports*, 470(5-6):151–238, January 2009.

- [BDL⁺00] 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, 2000.
- [BDS⁺03] M. D. Barrett, B. DeMarco, T. Schaetz, D. Leibfried, J. Britton, J. Chiaverini, W. M. Itano, B. Jelenković, J. D. Jost, C. E. Langer, T. Rosenband, and D. J. Wineland. Sympathetic cooling of $^9\text{Be}^+$ and $^{24}\text{Mg}^+$ for quantum logic. *Phys. Rev. A*, 68(4):042302, October 2003.
- [BEM⁺11] G. R. Brady, A. R. Ellis, D. L. Moehring, D. Stick, C. Highstrete, K. M. Fortier, M. G. Blain, R. A. Haltli, A. A. Cruz-Cabrera, R. D. Briggs, J. R. Wendt, T. R. Carter, S. Samora, and S. A. Kemme. Integration of fluorescence collection optics with a microfabricated surface electrode ion trap. *Appl. Phys. B*, 103(4):801–808, June 2011.
- [Ber84] M. V. Berry. Quantal phase factors accompanying adiabatic changes. *Proc. R. Soc. Lond. A*, 392:45–57, 1984.
- [BHI⁺91] J. J. Bollinger, D. J. Heinzen, W. M. Itano, S. L. Gilbert, and D. J. Wineland. A 303-MHz Frequency Standard Based on Trapped Be^+ Ions. *IEEE Trans. Instr. Meas.*, 40(2):126–128, 1991.
- [BKM⁺14] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. Logic gates at the surface code threshold: Superconducting qubits poised for fault-tolerant quantum computing. 2014.
- [BKR⁺07] J. Benhelm, G. Kirchmair, U. Rapol, T. K. Körber, C. F. Roos, and R. Blatt. Measurement of the hyperfine structure of the $S_{12}\text{-}D_{52}$ transition in $^{43}\text{Ca}^+$. *Phys. Rev. A*, 75(3):032506, March 2007.
- [BKRB08a] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. Experimental quantum-information processing with $^{43}\text{Ca}^+$ ions. *Phys. Rev. A*, 77(6):062306, June 2008.
- [BKRB08b] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nature Physics*, 4(6):463–466, April 2008.
- [BMB⁺98a] D. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Laser-Cooled Mercury Ion Frequency Standard. *Phys. Rev. Lett.*, 80(10):2089–2092, March 1998.
- [BMB⁺98b] 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), 1998.
- [BP03] S. Brouard and J. Plata. Internal-state dephasing of trapped ions. *Phys. Rev. A*, 68(1):012311, July 2003.
- [Bur10] A. H. Burrell. *High Fidelity Readout of Trapped Ion Qubits*. PhD thesis, May 2010.
- [BWC⁺11] K. R. Brown, A. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland. Single-qubit-gate error below 10^{-4} in a trapped ion. *Phys. Rev. A*, 84(3):1–4, September 2011.

-
- [CBWT12] Z. Chen, J. G. Bohnet, J. M. Weiner, and J. K. Thompson. General formalism for evaluating the impact of phase noise on Bloch vector rotations. *Phys. Rev. A*, 86(3):032313, September 2012.
 - [CN65] P. Carruthers and M. M. Nieto. Coherent States and the Forced Quantum Oscillator. *Am. J. Phys.*, 33(7):537, 1965.
 - [CN97] I. L. Chuang and M. A. Nielsen. Prescription for experimental determination of the dynamics of a quantum black box. *J. mod. Optics*, 44(11/12):2455–2467, 1997.
 - [CP54] H. Y. Carr and E. M. Purcell. Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments. *Phys. Rev.*, 94(3):630–638, 1954.
 - [CS13] J. Chiaverini and J. M. Sage. Insensitivity of Ion Motional Heating Rate to Trap Material over a Large Temperature Range. 2013.
 - [CZ95] J. I. Cirac and P. Zoller. Quantum Computations with Cold Trapped Ions. *Phys. Rev. Lett.*, 74(20):4091–4094, 1995.
 - [Dav93] F. Davis. AN267 - Matching Network Designs with Computer Solutions. Technical report, 1993.
 - [Deh75] H. Dehmelt. Proposed $10^{14}\delta\nu > \nu$ laser fluorescence spectroscopy of Ti^+ mono-ion oscillator II. *Bulletin of the American Physical Society*, 20(1):60, 1975.
 - [Deu85] D. Deutsch. Quantum theory , the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London A*, 400:97–117, 1985.
 - [DGB⁺13] N. Daniilidis, S. Gerber, G. Bolloten, M. Ramm, A. Ransford, E. Ulin-Avila, I Talukdar, and H. Häffner. Probing surface electric field noise with a single ion. pages 1–13, 2013.
 - [DM96] S. Dyrting and G. J. Milburn. Quantum chaos in atom optics: using phase noise to model continuous momentum and position measurement. *Quantum Semiclass. Opt.*, 8, 1996.
 - [DS13] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: an outlook. *Science*, 339(6124):1169–74, March 2013.
 - [EA05] J. Emerson, R. Alicki, and K. yczkowski. Scalable noise estimation with random unitary operators. *J. Opt. B: Quantum Semiclass. Opt.*, 7(10):S347–S352, October 2005.
 - [Ebe76] J. H. Eberly. Atomic Relaxation in the Presence of Intense Partially Coherent Radiation Fields. *Phys. Rev. Lett.*, 21:1387, 1976.
 - [ECMG13] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta. Investigating the limits of randomized benchmarking protocols. 2013.
 - [Fey82] R. P. Feynman. Simulating Physics with Computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, June 1982.
 - [FSL⁺95] P. T. H. Fisk, M. J. Sellars, M. A. Lawn, C. Coles, A. G. Mann, and D. G. Blair. Very High Q Microwave Spectroscopy on Trapped $^{171}\text{Yb}^+$ Ions : Application as a Frequency Standard. *IEEE Trans. Instr. Meas.*, 44(2):113–116, 1995.

- [Gar90] C. W. Gardiner. *Handbook of Stochastic Methods*. Springer-Verlag, second edition, 1990.
- [GKW⁺04] S. Groth, P. Kruger, S. Wildermuth, R. Folman, T. Fernholz, J. Schmiedmayer, D. Mahalu, and I. Bar-Joseph. Atom chips: Fabrication and thermal properties. *Appl. Phys. Lett.*, 85(14):2980, 2004.
- [GRB⁺01] S. Gulde, D. Rotter, P. Barton, F. Schmidt-Kaler, R. Blatt, and W. Hoger-vorst. Simple and efficient photo-ionization loading of ions for precision ion-trapping experiments. *Appl. Phys. B*, 73(8):861–863, December 2001.
- [Gro97] L. K. Grover. Quantum Mechanics Helps in Searching for a Needle in a Haystack. *Phys. Rev. Lett.*, 79(2):325–328, 1997.
- [HBD⁺05] P. C. Haljan, K-A Brickman, L. Deslauriers, P. Lee, and C. Monroe. Spin-Dependent Forces on Trapped Ions for Phase-Stable Quantum Gates and Entangled States of Spin and Motion. *Physical Review Letters*, 94(15):153602, April 2005.
- [HBHB10] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt. Trapped-ion probing of light-induced charging effects on dielectrics. *New J. Phys.*, 12(9):093035, September 2010.
- [HCD⁺12] D. Hayes, S. Clark, S. Debnath, D. Hucul, I. Inlek, K. Lee, Q. Quraishi, and C. Monroe. Coherent Error Suppression in Multiqubit Entangling Gates. *Physical Review Letters*, 109(2):1–5, July 2012.
- [HCW⁺12] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland. 100-Fold Reduction of Electric-Field Noise in an Ion Trap Cleaned with *In Situ* Argon-Ion-Beam Bombardment. *Phys. Rev. Lett.*, 109(10):103001, September 2012.
- [HHWW06] A. Hati, D. A. Howe, F. L. Walls, and D. K. Walker. Merits of PM noise measurement over noise figure: a study at microwave frequencies. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 53(10):1889–94, October 2006.
- [HL11] W. M. Haynes and D. R. Lide. *CRC Handbook of Chemistry and Physics*. 92 edition, 2011.
- [HNNH06] A. Hati, C. W. Nelson, and D. A. Howe. LOW PHASE NOISE AMPLIFIER AND OSCILLATOR USING FEED-FORWARD TECHNIQUE AT 10 GHz. *IEEE Int. Freq. Contr. Symp.*, 2006.
- [HP174] Spectrum Analysis: Noise Measurements - Application Note 150-4. Technical Report April, 1974.
- [HSM⁺13] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret. Quantum Back-Action of an Individual Variable-Strength Measurement. *Science*, 339(6116):178–81, January 2013.
- [IBBW95] W. M. Itano, J. C. Bergquist, J. J. Bollinger, and D. J. Wineland. Cooling methods in ion traps. *Phys. Scr.*, pages 166–181, 1995.

- [IBR⁺09] W. M. Itano, J. C. Bergquist, T. Rosenband, D. J. Wineland, D. Hume, C.-W. Chous, S. R. Jefferts, T. P. Heavner, T. E. Parker, S. A. Diddams, and T. M. Fortier. Frequency Measurements of Al^+ and Hg^+ Optical Standards. *Proc. 2009 ICOLS Conf.*, pages 117–124, 2009.
- [Jam98] D. F. V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Appl. Phys. B*, 66(2):181–190, 1998.
- [Jan] H. A. Janacek. *Simulating trapped-ion qubits using the optical Bloch equations*. PhD thesis.
- [JBT⁺09] M. Johanning, A. Braun, N. Timoney, V. Elman, W. Neuhauser, and C. Wunderlich. Individual Addressing of Trapped Ions and Coupling of Motional and Spin States Using rf Radiation. *Phys. Rev. Lett.*, 102(7):073004, February 2009.
- [Jef53] C. D. Jeffries. *Phys. Rev.* *Phys. Rev.*, 90(6):1130–1130, 1953.
- [KBZ⁺09] G. Kirchmair, J. Benhelm, F. Zähringer, R. Gerritsma, C. F. Roos, and R. Blatt. Deterministic entanglement of ions in thermal states of motion. *New Journal of Physics*, 11(2):023002, February 2009.
- [Kei07] B. C. Keitch. *A Quantum Memory Qubit in Calcium-43*. PhD thesis, 2007.
- [KHG⁺12] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. J. Martin, L. Chen, and J. Ye. A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity. *Nature photonics*, 6(September):2–7, 2012.
- [KLR⁺08] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. Blakestad, J. Jost, C. E. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Phys. Rev. A*, 77(1):1–7, January 2008.
- [KPS⁺12] A. Khromova, Ch. Piltz, B. Scharfenberger, T. F. Gloger, M. Johanning, A. F. Varón, and C. Wunderlich. Designer Spin Pseudomolecule Implemented with Trapped Ions in a Magnetic Gradient. *Phys. Rev. Lett.*, 108(22):220502, 2012.
- [KWM⁺98] B. E. King, C. S. Wood, C. J. Myatt, Q. A. Turchette, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland. Cooling the Collective Motion of Trapped Ions to Initialize a Quantum Register. *Phys. Rev. Lett.*, 81(7):1525–1528, August 1998.
- [Lan06] C. E. Langer. *High Fidelity Quantum Information Processing with Trapped Ions*. PhD thesis, 2006.
- [LBD⁺05] P. J. Lee, K-A Brickman, L. Deslauriers, P. C. Haljan, L-M Duan, and C. Monroe. Phase control of trapped ion quantum gates. *J. Opt. B: Quantum Semiclass. Opt.*, 7(10):S371–S383, October 2005.
- [LDM⁺03] D. Leibfried, B DeMarco, V. Meyer, D. M. Lucas, M. Barrett, J Britton, W. M. Itano, B Jelenković, C. E. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422(6930):412–5, March 2003.
- [Lev86] M. H. Levit. Composite Pulses. *Prog. Nucl. Magn. Reson. Spectrosc.*, 18:61–122, 1986.

- [LFP98] J. Li and E. Ferre-Pikal. Review of PM and AM Noise Measurement Systems. *Microwave and Millimeter Wave Technology Proceedings, 1998*, pages 197–200, 1998.
- [LGA⁺08] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibbrandt, K. R. Brown, and I. L. Chuang. Suppression of Heating Rates in Cryogenic Surface-Electrode Ion Traps. *Phys. Rev. Lett.*, 100(1):013001, January 2008.
- [Lin76] G. Lindblad. On the Generators of Quantum Dynamical Semigroups. *Commun. math. Phys.*, 48:119–130, 1976.
- [LKH⁺11] D. M. Lucas, B. C. Keitch, J. P. Home, G. Imreh, M. J. McDonnell, D. N. Stacey, D. J. Szwer, and A. M. Steane. A long-lived memory qubit on a low-decoherence quantum bus. pages 1–4, 2011.
- [LOJ⁺05] C. E. Langer, R. Ozeri, J. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. Blakestad, J. Britton, D. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland. Long-Lived Qubit Memory Using Atomic Ions. *Phys. Rev. Lett.*, 95(6):2–5, August 2005.
- [LRH⁺04] D. M. Lucas, A. Ramos, J. 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(1):1–13, January 2004.
- [LSL84] A. L. Lance, W. D. Seal, and F. Labaar. Phase Noise and AM Noise Measurements in the Frequency Domain. *Infrared and Millimeter Waves*, 11:239–289, 1984.
- [LSU73] O. Lutz, A. Schwenk, and A. Uhl. Nuclear Magnetic Resonance Studies of ⁴³Ca. *Z. Nat. A*, 28:1535, 1973.
- [Mag54] W. Magnus. On the Exponential Solution of Differential Equations for a Linear Operator. *Commun. Pure Appl. Math.*, VII:649–673, 1954.
- [MC93] C. D. Motchenbacher and J. A. Connelly. *Low-Noise Electronic System Design*. John Wiley & Sons, Inc., 1993.
- [MGE11] E. Magesan, J. M. Gambetta, and J. Emerson. Scalable and Robust Randomized Benchmarking of Quantum Processes. *Phys. Rev. Lett.*, 106(18):180504, May 2011.
- [MK13] C. Monroe and J. Kim. Scaling the Ion Trap Quantum Processor. *Science*, 339(6124):1164–9, March 2013.
- [Mon11] T. Monz. *Quantum information processing beyond ten ion-qubits*. PhD thesis, 2011.
- [MSB⁺11] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt. 14-Qubit Entanglement: Creation and Coherence. *Phys. Rev. Lett.*, 106(13):130506, March 2011.
- [MSJ00] G. J. Milburn, S. Schneider, and D. F. V. James. Ion Trap Quantum Computing with Warm Ions. *Fortschr. Phys.*, 48(9-11):801–810, September 2000.
- [MSW⁺04] M. J. McDonnell, J.-P. Stacey, S. C. Webster, J. Home, A. Ramos, D. M. Lucas, D. N. Stacey, and A. M. Steane. High-Efficiency Detection of a Single Quantum of Angular Momentum by Suppression of Optical Pumping. *Phys. Rev. Lett.*, 93(15):153601, October 2004.

-
- [MSW⁺08] 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):2–5, May 2008.
- [MW01] F. Mintert and C. Wunderlich. Ion-Trap Quantum Logic Using Long-Wavelength Radiation. *Phys. Rev. Lett.*, 87(25):257904, November 2001.
- [NNW94] L.M. Nelson, C.W. Nelson, and F. L. Walls. Relationship of AM to PM Noise in Selected RF Oscillators. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 41(5):680–684, 1994.
- [NRJ09] X. R. Nie, C. F. Roos, and D. F. V. James. Theory of cross phase modulation for the vibrational modes of trapped ions. *Phys. Lett. A*, 373(4):422–425, January 2009.
- [NVG⁺13] R. Noek, G. Vrijsen, D. Gaultney, E. Mount, T. Kim, P. Maunz, and J. Kim. High Speed, High Fidelity Detection of an Atomic Hyperfine Qubit. 2(1):2–5, 2013.
- [OIB⁺07] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. E. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Phys. Rev. A*, 75(4):042329, April 2007.
- [OLA⁺08] C. Ospelkaus, C. E. Langer, J. M. Amini, K. R. Brown, D. Leibfried, and D. J. Wineland. Trapped-Ion Quantum Logic Gates Based on Oscillating Magnetic Fields. *Phys. Rev. Lett.*, 101(9):1–4, August 2008.
- [Ols72] L. Olschewski. Messung der magnetischen Kerndipolmomente an freien ⁴³Ca-, ⁸⁷Sr-, ¹³⁵Ba-, ¹³⁷Ba-, ¹⁷¹Yb- und ¹⁷³Yb-Atomen mit optischem Pumpen. *Z. Physik*, 249(3):205–227, 1972.
- [OWC⁺11] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland. Microwave quantum logic gates for trapped ions. *Nature*, 476(7359):181–4, August 2011.
- [OYM⁺07] S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe. Manipulation and detection of a trapped Yb⁺ hyperfine qubit. *Phys. Rev. A*, 76(5):052314, November 2007.
- [PCZ97] J. Poyatos, J. Cirac, and P. Zoller. Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate. *Phys. Rev. Lett.*, 78(2):390–393, January 1997.
- [Poz05] D. M. Pozar. *Microwave Engineering*. Addison-Wesley Publishing Company, Inc., third edition, 2005.
- [RGP⁺12] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. a. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms. *Phys. Rev. B*, 86(10):100506, September 2012.
- [RHB02] K. F. Riley, M. P. Hobson, and S. J. Bence. *Mathematical Methods for Physics and Engineering*. Cambridge University Press, second edition, 2002.

- [RKS⁺06] M. Riebe, K. Kim, P. Schindler, T. Monz, P. O. Schmidt, T. K. Körber, W. Hänsel, H. Häffner, C. F. Roos, and R. Blatt. Process Tomography of Ion Trap Quantum Gates. *Phys. Rev. Lett.*, 97(22):220407, December 2006.
- [Rob84] W.P. Robins. *Phase Noise in Signal Sources*. Peter Peregrinus Ltd., 1984.
- [Roo08] C. F. Roos. Ion trap quantum gates with amplitude-modulated laser beams. *New Journal of Physics*, 10(1):013002, January 2008.
- [Sch52] E. Schrödinger. Are There Quantum Jumps? Part II. *Brit. J. Phil. Sci.*, 3:233–242, 1952.
- [SFF99] M Sampietro, L Fasoli, and G Ferrari. Spectrum analyzer with noise reduction by cross-correlation technique on two channels. *Rev. Sci. Instrum.*, 70(5):2520–2525, 1999.
- [SFZ99] E Solano, R.L. de Matos Filho, and N Zagury. Deterministic Bell states and measurement of the motional state of two trapped ions. *Phys. Rev. A*, 59(4):2539–2543, 1999.
- [Sho97] P.W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509, October 1997.
- [SKDB10] G. Shu, N. Kurz, M. R. Dietrich, and B. B. Blinov. Efficient fluorescence collection from trapped ions with an integrated spherical mirror. *Phys. Rev. A*, 81(4):042321, April 2010.
- [SKHR⁺03] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt. Realization of the Cirac-Zoller controlled-NOT quantum gate. *Nature*, 422(6930):408–11, March 2003.
- [SLA⁺08] 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, April 2008.
- [SM98] S. Schneider and G. J. Milburn. Decoherence in ion traps due to laser intensity and phase fluctuations. *Phys. Rev. A*, 57(5):3748–3752, May 1998.
- [SM99] A. Sørensen and K. Mølmer. Quantum Computation with Ions in Thermal Motion. *Phys. Rev. Lett.*, 82(9):1971–1974, March 1999.
- [SM00] A. Sørensen and K. Mølmer. Entanglement and quantum computation with ions in thermal motion. *Phys. Rev. A*, 62(2):022311, 2000.
- [SNM⁺] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt. A quantum information processor with trapped ions.
- [SPC⁺12] A. P. Sears, A. Petrenko, G. Catelani, L. Sun, H. Paik, G. Kirchmair, L. Frunzio, L. I. Glazman, S. M. Girvin, and R. J. Schoelkopf. Photon shot noise dephasing in the strong-dispersive limit of circuit QED. *Phys. Rev. B*, 86(18):180504, November 2012.
- [SRS06] B. Schiek, I. Rolfes, and H. Siweris. *Noise in High-Frequency Circuits and Oscillators*. Wiley-Interscience, 2006.

- [SRS13] SRS. *FS725 Rubidium Frequency Standard Operation and Service Manual*. 2013.
- [SRW⁺] R. C. Sterling, H. Rattanasonti, S. Weidt, K. Lake, P. Srinivasan, S. C. Webster, M. Kraft, and W. K. Hensinger. Two-dimensional ion trap lattice on a microchip. pages 1–7.
- [SSS⁺13] K. Saeedi, S. Simmons, J. Z. Salvail, P. Dluhy, H. Riemann, N. V. Abrosimov, P. Becker, Hansj-Joachim Pohl, J. J. L. Morton, and M. L. W. Thewalt. Room-Temperature Quantum Bit Storage Exceeding 39 Minutes Using Ionized Donors in Silicon-28. *Science*, 342:28–31, 2013.
- [Ste97] A. M. Steane. The ion trap quantum information processor. *Appl. Phys. B*, 642:623–642, 1997.
- [Ste03] A. M. Steane. How to build a 300 bit, 1 giga-operation quantum computer. *Quant. Inf. Comp.*, 0(0):1–13, 2003.
- [Szw09] D. J. Szwed. *High Fidelity Readout and Protection of a $^{43}\text{Ca}^+$ Trapped Ion Qubit*. PhD thesis, University of Oxford, 2009.
- [TKK⁺00] Q. A. Turchette, D. Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Phys. Rev. A*, 61(6):063418, May 2000.
- [TPR⁺03] G. Tammasso, T. Pfeil, G. Revalde, G. Werth, P. Indelicato, and J. P. Desclaux. The g_J -factor in the ground state of Ca^+ . *Eur. Phys. J. D*, 25(2):113–121, August 2003.
- [TTM⁺00] Q. A. Turchette, Q. A. Te, C. J. Myatt, B. E. King, C. A. Sackett, D. Kielpinski, W. M. Itano, C. Monroe, and D. J. Wineland. Decoherence and decay of motional quantum states of a trapped atom coupled to engineered reservoirs. *Phys. Rev. A*, 62(5):1–22, 2000.
- [UBV⁺09] Hermann Uys, Michael J. Biercuk, Aaron P. VanDevender, Nobuyasu Shiga, W. M. Itano, and J. J. Bollinger. Prolonging qubit coherence: dynamical decoupling schemes studied in a Penning ion trap. *Proceedings of SPIE*, 7225:72250R–72250R–9, 2009.
- [Wal92] W.F. Walls. Cross-Correlation Phase Noise Measurements. *Proc. IEEE Frequency Control Symposium, 1992*, pages 257–261, 1992.
- [Wal97] F. L. Walls. Suppressed Carrier Based PM and AM Noise Measurement Techniques. *Proc. IEEE Frequency Control Symposium, 1997*, 1997.
- [WCF⁺88] F. L. Walls, A. J. D. Clements, C. M. Felton, M. A. Lombardi, and M. D. Vanek. Extending the range and accuracy of phase noise measurements. *Proc. 42 Ann. Symp. on Freq. Contr.*, pages 432–441, 1988.
- [WEL⁺07] 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. E. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Fluorescence during Doppler cooling of a single trapped atom. *Phys. Rev. A*, 76(5):053416, November 2007.
- [WHL⁺11] Shannon X. Wang, Guang Hao Low, Nathan S. Lachenmyer, Yufei Ge, Peter F. Herskind, and Isaac L. Chuang. Laser-induced charging of microfabricated ion traps. *Journal of Applied Physics*, 110(10):104901, 2011.

- [WL11] D. J. Wineland and D. Leibfried. Quantum information processing and metrology with trapped ions. *Laser Phys. Lett.*, 8(3):175–188, March 2011.
- [WM08] D. F. Walls and G. J. Milburn. *Quantum Optics*. Springer, second edition, 2008.
- [WMI⁺98] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259, May 1998.
- [WOC⁺12] U. Warring, C. Ospelkaus, Y. Colombe, K. R. Brown, J. M. Amini, M. Carsjens, D. Leibfried, and D. J. Wineland. Microwave Near-Field Quantum Control of Trapped Ions. *arXiv preprint*, pages 1–12, 2012.
- [WOC⁺13] U. Warring, C. Ospelkaus, Y. Colombe, R. Jördens, D. Leibfried, and D. J. Wineland. Individual-Ion Addressing with Microwave Field Gradients. *Phys. Rev. Lett.*, 110(17):173002, April 2013.
- [Woo80] G. K. Woodgate. *Elementary Atomic Structure*. Oxford University Press, 1980.
- [Wün91] A Wünsche. Displaced Fock states and their connection to quasiprobabilities. *Quantum Opt.*, 3(6):359–383, December 1991.