

Sideband Cooling to the Quantum
Ground State in a Penning Trap

Joseph Francis Goodwin

Thesis submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

Department of Physics · Imperial College London

March 2015

Declaration

This thesis is my own work. Wherever I have used material from other sources, those sources are acknowledged and referenced.

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Abstract

For 35 years, laser-cooled trapped ions have been at the frontier of progress in quantum computing, quantum simulation and precision measurement, and remain one of the most valuable tools in these fields to this day. Most of these experiments are predicated upon or benefit from the ability to place ions in the motional quantum ground state, a technique that was first demonstrated in radiofrequency ion traps 25 years ago. For a range of crucial experiments that are impossible to conduct in radiofrequency traps or are not well suited to this architecture, Penning traps provide an important alternative. However, the performance of Penning traps had been limited by the fact that ground-state cooling was yet to be achieved in such a system.

This thesis reports the first demonstration of resolved-sideband cooling in a Penning trap, for $^{40}\text{Ca}^+$ ions cooled with light at 729-nm, achieving a ground state occupation of 99% in one dimension. Demonstrations of the coherent manipulations possible at this level of confinement are presented. The ion heating rate is measured and although higher than might be expected given the unusually large ion-electrode distance remains amongst the lowest reported in any trap to date. Achieving this result required the development of a number of new experimental systems and major upgrades to the stability and reliability of the experiment, the details of which are also given.

The thesis also presents theoretical work into the use of two-dimensional Coulomb crystals in a Penning trap as a resource for quantum information. Using the symmetries of the crystal, we find that it is possible to engineer complex entangled states, specifically two small quantum error correcting codes, using a very small number of global entangling pulses. Efficient entanglement protocols such as these are vital for the implementation of useful quantum error correction.

Acknowledgements

I would like to thank my two supervisors, Danny Segal and Richard Thompson for their support, guidance and friendship throughout my MRes and PhD. It has been a pleasure to work in such a friendly and warm environment and I am very happy to be continuing in the group next year.

I am also very grateful to my colleagues in the lab over the years: Sean Donnellan, Sandeep Mavadia, Dan Crick, Shailen Bharadia and Graham Stutter. Constructing a large and complicated experiment from predominantly ‘home-made’ apparatus is at turns arduous and deeply satisfying and I am lucky to have been able to learn from so many talented people in the process. To Dan, I am particularly indebted for turning my fear of electronics into a clumsy enthusiasm, although I doubt I’ll ever ‘get it’ like he does.

I would do Graham an injustice if I did not single him out for special praise and thanks. To find someone you work really well with is brilliant and this is made all the better when they are a great friend. I hope the rest of my career can be as productive and enjoyable as the last two years has been.

Further thanks goes to all the MSci, MSc, MRes and summer students who have worked with us, several of whom have made vital contributions to the experiment. I would particularly like to thank Sarah Woodrow for her excellent work on several experimental control systems.

I am indebted to the technical staff in the Physics department for their help and advice with countless projects and designs. Particular thanks go to Brian Willey, who has expertly machined all of the parts of the experiment that *aren’t* held together by blutack and string (we did the rest). Thank you also to Valerijus Gerulis, Bandu Ratnasekara and Susan Parker for their help with electronics.

Thank you to all those in our ‘foster’ group, the Centre for Cold Matter, for a healthy trade in technical advice and borrowed equipment, not to mention reliably entertaining lunchtime conversation, pub sessions and softball. A big thank you also Judith Baylis and Sanja Maricic for navigating the murky waters of academic bureaucracy.

My PhD has been funded as part of the Controlled Quantum Dynamics CDT and it has been both fun and stimulating to be part of this wider group of researchers and students.

I am grateful to the 12 other students in the first cohort of the DTC, who made that ‘pioneering’ year so much fun and who have provided good company and friendship since then. Particular mention goes to Ben Brown with whom I have enjoyed collaborating on theoretical work in the last year, along with Howard Dale. I am also grateful to all those who have made the CDT work so well, especially Danny Segal (again), Terry Rudolph, Myungshik Kim, Lilian Wanjohi, Miranda Smith and Nic Harrigan.

On a more personal note, I would like to thank all my friends and family for their support throughout this PhD, and for doing an excellent job of pretending not be bored when I insisted on giving them detailed updates on my progress. I owe a very big thank you to my parents and grandparents, who have been there through the highs and the lows and have always provided me with great encouragement. Finally, I would like to thank Siobhan for all her love and support, and for making my life outside physics a happy one.

List of Publications

1. S. Mavadia, J. F. Goodwin, G. Stutter, S. Bharadia, D. R. Crick, D. M. Segal, and R. C. Thompson, “Control of the conformations of ion Coulomb crystals in a Penning trap,” *Nat. Comms.*, vol. 4, no. 2571, p. 2571, 2013.
2. S. Mavadia, G. Stutter, J. F. Goodwin, D. R. Crick, R. C. Thompson, and D. M. Segal, “Optical sideband spectroscopy of a single ion in a Penning trap,” *Phys. Rev. A*, vol. 89, p. 032502, Mar 2014.
3. J. F. Goodwin, G. Stutter, R. C. Thompson, and D. M. Segal, “Sideband cooling an ion to the quantum ground state in a Penning trap with very low heating rate,” *arXiv:1407.6121*, 2014.
4. J. F. Goodwin, B. J. Brown, G. Stutter, H. Dale, R. C. Thompson, and T. Rudolph, “Trapped ion quantum error correcting protocols using only global operations,” *arXiv:1407.1858v1*, 2014.

Contents

List of Figures	15
List of Tables	19
1 Introduction	21
1.1 Description of Thesis and Context	26
2 Theoretical Background	29
2.1 Penning Traps	29
2.1.1 Classical motion in a Penning trap	30
2.1.2 Quantised motion in a Penning trap	34
2.2 Laser-Ion Interactions	38
2.2.1 Interaction with free atom	38
2.2.2 Interaction with trapped atom	42
2.3 Laser Cooling	46
2.3.1 Doppler cooling	47
2.3.2 Relaxation heating	49
2.3.3 Doppler cooling in Penning traps	53
2.3.4 Sideband cooling	58
3 Experiment Overview	65
3.1 Introduction & Aims	66
3.2 Calcium Electronic Structure	66
3.2.1 Effects of magnetic field	68
3.2.2 Transition strengths and saturation intensities	71
3.3 Overview of Apparatus	73
3.4 Trap, Magnet and Associated Systems	75
3.4.1 Overview	75
3.4.2 Magnetic field inhomogeneity	79
3.4.3 Trap supply electronics	82
3.4.4 Trap compensation	83

3.5	Doppler Cooling Lasers and Locking Systems	86
3.5.1	Overview	86
3.5.2	Cooling laser locks	87
3.5.3	Repump laser system	93
3.5.4	Laser lock evaluation	94
3.5.5	Fibre EOM	99
3.6	729-nm Laser, Locking and Amplifier Systems	101
3.6.1	Overview	101
3.6.2	ULE cavity	103
3.6.3	Tapered amplifier	113
3.7	Experimental Control System	120
4	Sideband Cooling in a Penning Trap	125
4.1	Spectroscopy on the $S_{1/2} \leftrightarrow D_{5/2}$ Transition	126
4.1.1	Spectroscopy pulse sequence	127
4.1.2	Beam geometry	129
4.1.3	Choice of transition	129
4.2	Doppler Cooling	130
4.2.1	Axial motion	130
4.2.2	Radial motion	135
4.2.3	Coherent manipulation near the Doppler limit	136
4.3	Sideband Cooling	137
4.3.1	Pulse sequence	137
4.3.2	Optimising the quench laser parameters	138
4.3.3	Optimising the sideband cooling frequency	138
4.3.4	Effects of radial motion	139
4.3.5	Ground state cooling	141
4.4	Heating Rate	146
5	Global Entangling Gates in a Penning Trap	149
5.1	Overview	149
5.2	Quantum Error Correction	150
5.2.1	Practical quantum codes	151
5.3	Geometric Phases and the Driven Quantum Harmonic Oscillator	154
5.4	Spin-Dependent Forces in a Penning Trap	157
5.5	Dynamics of Two-Dimensional Coulomb Crystals	163
5.6	Global Unitary Operations on a Coulomb Crystal	166
5.6.1	Application to quantum error correction	169
5.6.2	Performance under noise	173
5.6.3	Key experimental considerations	175

<i>Contents</i>	13
6 Conclusion	177
Appendices	181
A Transition Matrix Elements	183
B Excitation Lineshapes	189
C Artefacts of Radial Motion	193
D Miscellaneous Electronic Circuits	197
E CAD Renders	203
Bibliography	209

List of Figures

2.1	Radial trajectories in a Penning trap.	32
2.2	Rabi oscillations	40
2.3	Energy distributions after relaxation heating	51
2.4	Motional sideband spectra after relaxation heating	52
2.5	Intensity gradient cooling geometry	54
2.6	Radial trajectory in the presence of axialisation	57
2.7	Cascade configuration for broadening a metastable level	61
2.8	Sideband coupling strength versus phonon number	62
2.9	High- n population trapping for first, second and third red sidebands	63
2.10	Theoretical radial sideband spectra at two trap frequencies	64
3.1	Electronic structure of a $^{40}\text{Ca}^+$ ion	67
3.2	Overview of experimental subsystems and their interactions	73
3.3	SPECTRAP electrode structure and imaging optics	77
3.4	Trap breadboard laser and imaging optics	78
3.5	Thermal-induced drift of carrier transition frequency	80
3.6	Magnetic field gradient at ion before and after adjustment	81
3.7	Simplified schematic of Arduino-based MCA	85
3.8	Locking and control systems for Doppler cooling and repumping lasers	88
3.9	Scanning cavity peak detector principle	90
3.10	Input pulses to scanning cavity Arduino controller	91
3.11	Detection and feedback systems for scanning cavity lock	91
3.12	Exploded view of scanning cavity components	93
3.13	Section view of assembled scanning cavity	94
3.14	397-nm beatnote with and without lock	95
3.15	Linewidth of 397-nm laser vs sampling time	96
3.16	Allan deviations for 397-nm laser, locked and unlocked	98
3.17	Effect of second-order Zeeman splitting on 866-nm transitions	100
3.18	Schematic layout of 729-nm laser system	102

3.19	Temperature dependence of ULE cavity length	105
3.20	Section view of ULE cavity cooler superstructure	107
3.21	Section view of cavity cooler showing insulation	108
3.22	Zero-crossing measurement of ULE cavity	111
3.23	Tapered amplifier optics	114
3.24	Exploded view of tapered amplifier mount	117
3.25	Section view of tapered amplifier mount	118
3.26	Tapered amplifier output beam profile	119
3.27	Screenshots of spectroscopy control and analysis software	122
3.28	Layout of AOM RF control system	124
4.1	Laser pulse sequence for electron-shelving spectroscopy	127
4.2	Structure of a generic spectroscopy experiment	129
4.3	Axial motional sideband spectra after Doppler cooling	132
4.4	Spectrum of carrier after Doppler cooling	134
4.5	Radial motional sideband spectrum after Doppler cooling	135
4.6	Rabi oscillation on carrier after Doppler cooling	137
4.7	Laser pulse sequence for sideband cooling	138
4.8	Measuring the quench laser broadening	139
4.9	Optimising the sideband cooling frequency with ‘dip’ spectrum	140
4.10	Radial artefacts in saturated axial spectrum	141
4.11	Sideband cooling at $f_z = 405$ kHz	142
4.12	Axial sideband spectrum after Doppler cooling with $f_z = 389$ kHz	143
4.13	Red and blue axial sidebands after ground state cooling	144
4.14	Rabi oscillation on carrier after ground state cooling	144
4.15	Spectroscopy of carrier after ground state cooling	145
4.16	Heating rate measurement on red and blue sidebands	146
4.17	Heating versus delay time for SPECTRAP	147
4.18	Comparison with heating rates in other traps	148
5.1	Quantum circuit for formation of pentagon graph state.	152
5.2	Phase space trajectory of driven quantum harmonic oscillator	155
5.3	Optical dipole force lasers for $^9\text{Be}^+$	157
5.4	Experimental geometry for producing transverse optical dipole forces	158
5.5	Optical dipole force versus Raman laser detuning	162
5.6	Vibrational mode frequencies and inter-ion distance for 6 ion planar crystal	165
5.7	Optical dipole force and readout lasers and microwave control for $^9\text{Be}^+$	166
5.8	Phase space trajectories of two basis states driven on the centre of mass mode	167
5.9	Protocol for producing the five qubit code with global pulses	170
5.10	Behaviour of encoded information under realistic noise	174

B.1	Excitation lineshapes for a range of pulse times and Rabi frequencies	191
D.1	Trap supply: AC and DC compensation circuit	198
D.2	Trap supply: Endcap supply and isolatable pulse-down circuit.	199
D.3	Scanning cavity peak detector circuit	200
D.4	866-nm and 854-nm laser piezo PI lock circuit	201
E.1	Quarter-section and exterior views of assembled scanning cavity	204
E.2	Exterior view of complete ULE cavity cooler	205
E.3	Cross-section of ULE cavity cooling system with insulation	206
E.4	Assembled tapered amplifier and optics on temperature controlled mount .	207
E.5	Assembled tapered amplifier mount with acrylic dust cover	208

List of Tables

3.1	First and second order Zeeman shifts of $^{40}\text{Ca}^+$ levels at $B = 1.85\text{ T}$	69
3.2	Saturation intensities of relevant transitions in $^{40}\text{Ca}^+$	72
5.1	Errors, syndromes and the corresponding corrections to the final teleported state, for 5RC and 5QC	171
5.2	Phases acquired by register basis states after each stage of composite entangling pulse	172

1

Introduction

“Taking the attitude that the pursuit of as basic an ideal as a single atomic particle at rest in space is a thoroughly worthwhile intellectual endeavour, we are undertaking experiments along these lines.”

– Neuhauser, Hohenstatt, Toschek and Dehmelt, 1978

Writing these words in July 1978, as the introduction to the joint-first¹ reported observation of laser cooling [2], the Toschek group clearly realised the importance of their work, but could surely not have foreseen that 36 years later the method would remain at the very foundation of experimental atomic physics, nor have envisioned the panoply of complex applications that would spring forth from this source. The primary motivation for both of these pioneering groups was the improvement of frequency standards. The accuracy of atomic clocks in the late 1970s was limited by the short interaction times and Doppler broadening associated with thermal atomic beams and cold, trapped ions provided a potential solution to both these problems.

Despite the stated aim of their investigations, at this point the Toschek group were working with 50 or so trapped ions, which remained at temperatures many times higher than the Doppler limit. It was not long before both these limitations were overcome: two years later the same group succeeded in trapping and cooling a single ion [3], and in 1989 the group at the National Institute of Standards (NIST) Boulder reported the first laser cooling of a single ion to the quantum ground state of its motion [4], extending this to three dimensions in 1995 [5] to realise the ‘basic ideal’ described by Neuhauser.

¹A similar result from David Wineland and colleagues at NIST was published a month prior [1], but had been submitted to Physical Review Letters a single day later.

The years since have seen an explosion of applications and further achievements in ion trapping. Quantum information science, born of work in the 1980s by Feynman, Deutsch and others [6, 7], leapt to the fore in 1994 with the discovery of Shor's algorithm [8] and trapped ions were rapidly recognised as an ideal architecture for the implementation of these ideas. In 1995 Cirac and Zoller described a basic protocol for quantum computation, using motional modes of a Coulomb crystal to mediate the interaction between the constituent ion qubits [9]. Further theoretical efforts towards the end of the decade, by Mølmer and Sørensen [10], Milburn [11] and others, proposed quantum gates that did not require individual-ion addressing and were more tolerant of residual motional excitation. The first demonstrations of two-qubit gates appeared in the early 2000s, with controlled-not and controlled-phase gates realised in late 2002 by groups from Innsbruck and NIST respectively [12, 13].

Since these early demonstrations, trapped ions have led the march in quantum computing, with demonstrations of deterministic entanglement of up to 14 qubits [14], quantum error correcting protocols [15, 16] and ever greater coherence times and gate fidelities [17, 18]. Trap technology engineering has developed into a discipline in its own right, with an increasing number of experiments conducted on traps fabricated with lithography techniques adapted from those of the semiconductor industry [19].

In recent years, quantum simulation – where a well-controlled quantum system is used to model the Hamiltonian of one less easily accessed – has been the focus of much research, spurred on by the increasing likelihood that it may be here, rather than in gate-model universal quantum computation, that the first useful results appear. Quantum simulations in both 'digital' [20] and 'analogue' [21] formats are under development and the leading efforts are already approaching the boundaries of classical verification [22].

The original aim to improve frequency standards has also yielded results. While trapping and cooling allow for more accurate measurement of a microwave clock transition, the real benefits come from moving to a much higher frequency optical transition where the same absolute frequency uncertainty translates to a much lower absolute error [23]. The current world record remains with NIST, where expertise in quantum logic gates led to the development of quantum logic spectroscopy, allowing clocks to be built using ions where direct cooling or read-out is impossible. The group use an aluminium ion mechanically coupled to a magnesium ion, marrying the high stability of the aluminium clock state with the more conveniently manipulated magnesium to produce a clock with fractional error of $< 10^{-17}$ [24]. At this level, the time dilation due to the Earth's gravitational field can be observed by moving the ion a few centimetres up or down [25].

The vast majority of these techniques require or are improved by ground-state cooling. Sideband cooling in particular has become the technique which underpins the leading quantum information experiments worldwide. Furthermore, in recent years sideband cooling has been applied to a wide range of physical systems including neutral atoms in optical

lattices [26] and optical dipole traps [27], and even mesoscopic micromechanical oscillators [28]. Yet, despite the fact that the first demonstration of laser cooling at NIST was performed in a Penning trap, the extension to ground state cooling in such a system had not been realised in the nearly four decades since. Indeed, until the last two years, only two groups continued to use Penning traps in laser cooling experiments (the Imperial group and the Bollinger group at NIST). The primary cause of this lack of interest is not hard to diagnose: the experimental overhead for work in a Penning trap is considerably higher than for RF traps. From the fact that the entire experiment must be conducted in the heart of a large and expensive superconducting magnet which requires constant cryogenic cooling and severely restricts optical access to the trap, to the greatly increased number of lasers needed to cool and repump levels with large Zeeman splittings, the amount and complexity of equipment for a given task is multiplied. The apparently complex dynamics and subtleties of cooling the radial motion present a further barrier.

So why bother? Aside from pure curiosity and breadth of research, there is a range of applications in which Penning traps excel. The two most important benefits follow immediately from the definitional properties of the trap – that there are no radiofrequency fields² and that there is a strong axial magnetic field. The lack of RF means that large two- and three-dimensional Coulomb crystals can be produced free from the perturbations of coherent micromotion. Such structures could provide a convenient means of accessing large registers of qubits in a compact geometry and also show great promise for quantum simulations [29], as will be discussed below. The presence of the large magnetic field makes the trap ideal, and indeed the only option, for experiments where such a field is necessary for other reasons, such as cyclotron frequency measurements for precision mass spectrometry [30] and g -factor measurements [31]. Penning traps have a rich history in both fields, where they have been valued for the high degree of environmental isolation they provide. However it is only recently that the possibilities of improving such measurements with the addition of laser cooling are being explored. A secondary advantage of the magnetic field is that the large Zeeman splitting – while increasing the complexity of the laser systems – allows for near-perfect state readout with frequency selection alone. Finally, Penning traps have the advantage that respectable trapping frequencies can be achieved in a very large trap, reducing the effects of ion heating due to the proximity of electrode surfaces. As will be shown in this thesis, world-class heating rates can be achieved without the need to make special efforts to reduce them. The potential for unparalleled levels of isolation could make the technique of great interest to studies in quantum thermodynamics.

Sideband cooling in a Penning trap has remained one of the primary aims of the Imperial ion trapping group since the group began working with calcium ions in 2002, and it is testament to the challenges involved that it has taken this long. In the time it has taken

²With the possible exception of an axialisation or rotating wall potential, which are both at least two orders of magnitude weaker than the trapping potential of an RF trap.

to demonstrate the technique, the motivation for sideband cooling in a Penning trap has changed somewhat. In 2002, trapped ion quantum information processing was still in its infancy and there was still some uncertainty as to which architectures would prove successful. RF traps had been plagued by problems with micromotion [32] and Penning traps provided an attractive alternative route. These problems have largely been overcome since and with the enormous progress seen over the last decade, it no longer seems likely that Penning traps will provide direct competition where RF traps do so well – performing gate operations on strings of ions, or shuttling ion registers on-chip. However, two other applications have since risen in prominence and have led to an increased general interest in Penning traps as well as experiments in two new groups.

Firstly, with the surge of interest in quantum simulation, the ability of the Penning trap to store large two-dimensional crystals has attracted the attention of many theorists, who point to the triangular lattice that naturally forms as the ideal test-bed for studies of quantum magnetism and spin frustration. Indeed, a recent demonstration of a Penning trap quantum simulator by the Bollinger group at NIST [21] made headlines worldwide and set the benchmark for progress in the field. The simulator is based on a Coulomb crystal of approximately 200 beryllium ions arranged in a single plane and phase locked to a rotating quadrupole potential that permits stroboscopic imaging of the crystal to resolve the individual ions. An Ising-type interaction is simulated, with two ground state hyperfine levels of each beryllium ion representing the spin-up and spin-down states. The effect of the external magnetic field B is modelled by a microwave field directly coupling the spin states, while the spin-spin interaction term J_{ij} is provided by a spin-dependent optical dipole force. The optical dipole force couples the spins via an off-resonant excitation of the transverse vibrational modes of crystals. By controlling the detuning from all the modes the effective range of the interaction can, in principle, be varied from infinite to $1/d^3$, modelling dipole-dipole coupling. The variation of the qubit population under the external field and spin-spin interactions of different strengths and ranges then provides a means to simulate the effects of magnetism in a variety of regimes, from antiferromagnetic to ferromagnetic. In this experiment, the strength of the spin-spin interaction remained low to simulate true dipole-dipole coupling and to go further will require this to be boosted. The interaction strength is limited in part by the fact that the crystal modes remain in a thermal state, and could be greatly increased if the ions were ground-state cooled. It should also be noted that the Bollinger group are no longer alone in this endeavour - the new Penning trap experiment undertaken by the Biercuk group in Sydney aims to use the addressing of single qubits to extend the architecture demonstrated at NIST and produce universal quantum simulation.

A rather different application is the subject of research by the (C.) Ospelkaus group, which proposes to employ the quantum logic spectroscopy techniques developed for aluminium-ion optical clocks to make a measurement of the proton and anti-proton g -factors, as part of

the CERN ‘BASE’ collaboration [33]. The scheme involves bringing a ground-state cooled beryllium ion into close proximity with the proton, sympathetically cooling it to the ground state. This beryllium ion can then also be used for indirect read-out of the proton spin-state. Accurate measurements of the cyclotron and Larmor frequencies of a proton demand a large magnetic field, making the use of a Penning trap unavoidable. This is a challenging but exciting experiment, and one that is predicated on the ability to ground-state cool in a Penning trap.

1.1 Description of Thesis and Context

The aim of my PhD research was to sideband cool an ion in a Penning trap. I had the privilege of joining the group just as the current experiment was beginning. During the first few months of my PhD, the existing students (S. Mavadia and S. Donnellan) and post-doc (D. Crick, later joined by S. Bharadia) were building up the optical systems from scratch and redesigning the trap. My first year was spent finishing the work I had started in my MRes project, designing and building scanning transfer cavity systems to stabilise the Doppler cooling and repump lasers. For the following six months I worked on a variety of other systems and learnt about the day-to-day running of the experiment. During this period we also recorded the data that would be used for a paper describing the structure of Coulomb crystals [34], and performed the first resolved sideband spectroscopy in a Penning trap.

By mid-2012, the previous students and post-docs had left the group and the running of the experiment fell to me and fellow student G. Stutter. The first six months of this period were spent making wide-ranging improvements to the reliability and stability of the experiment, in particular minimising drifts to the spectroscopy laser and magnetic field. These changes allowed us to record the data that was then used in our paper presenting the results of the resolved-sideband spectroscopy experiments [35]. The primary aim from this point onwards was to perform sideband cooling and we spent the following nine months increasing the cooling laser power, extending the experimental control system and developing techniques to cool at much higher trap frequencies, in the process rebuilding almost every system of optics. We managed to sideband cool by the end of October 2013, and spent some time observing the coherence of our system and measuring the trap heating rate. This work has formed the basis of our paper [36], currently under review.

Having completed these experimental goals in January 2014, I then spent the following five months developing a theoretical proposal for realising quantum error correcting codes in Penning traps. This work was conducted in collaboration with theorists from the Imperial quantum information group (B. Brown, H. Dale and T. Rudolph) and has led to a further paper, currently under review [37].

It should be noted that much of the work described in this thesis was conducted in close collaboration with G. Stutter and that none of the major experimental results could have been achieved without his efforts as well my own, even where not explicitly mentioned in the accompanying text.

Chapter 2 provides the theoretical background for the thesis. The first section covers the classical and quantum theory of Penning traps. The second section derives the basic models used for the interaction of light with free and trapped atoms. The third section discusses laser cooling in the weak-binding and strong-binding regimes, giving special focus to a review of those issues important to cooling in Penning traps and beyond the Lamb-Dicke regime, such as the cooling of the magnetron motion, Fock-state population trapping and radiation pressure relaxation heating.

Chapter 3 is an amalgam of several rather large sections describing the overall working of the experiment at the time of writing. I begin by describing the details of the experiment as applied to calcium-40 ions and some of the important considerations when working in a Penning trap, mainly associated with the breakdown of the first order perturbative treatment of the magnetic field. I give an overview of the total working of the experiment and the manner in which the four loosely-defined experimental subsystems interact and continue by describing each of these subsystems in turn, highlighting those elements I have developed or designed. I cover: the trap, the magnet and associated electronics; the lasers used for Doppler cooling and repumping and their locking systems; the 729-nm spectroscopy and sideband cooling laser, its locking system and amplifier; and the FPGA-based experimental control system and PC-based spectroscopy analysis software.

In Chapter 4, I present the major experimental results of this thesis, namely the one-dimensional ground-state cooling of a single ion in a Penning trap. The chapter begins with a description of the experimental spectroscopy technique, before continuing to present a range of Doppler cooling results demonstrating the gradual improvements made in the run-up to sideband cooling. I then present the sideband cooling results, including a description of the experimental techniques used to tune and optimise the performance, as well as demonstrations of the coherent manipulation that the technique makes possible. The final section concludes with a measurement of the heating rate in our trap, which we have found to be amongst the lowest yet recorded in any trap.

Chapter 5 is somewhat self-contained with respect to the rest of the thesis, and covers the theoretical work I undertook in the final six months of my PhD. The work focused on the applications of two-dimensional Coulomb crystals in Penning traps and the possibilities of engineering complex entangled states (in particular, quantum error correcting codes) by using a very small number of global entangling operations. The first three sections of the chapter provide an overview of quantum error correction, geometric phase manipulation and the creation of spin-dependent forces using polarisation gradients. The remaining two sections describe a basic technique for generating graph states on a two-dimensional crystal and demonstrate its application to the production of two simple codes.

I conclude the thesis in Chapter 6 and discuss the outlook for the coming months and years.

2

Theoretical Background

2.1 Penning Traps

The charge of an atomic ion provides a convenient and powerful ‘handle’ with which to manipulate it. All we need to do to trap one is arrange a field that will provide an electrostatic force on the ions of charge e (which we assume to be positive), directed towards the centre of the trap. The lowest order, spherically symmetric field that satisfies this is given by

$$\nabla\Phi(\mathbf{r}) = \mathbf{E}(\mathbf{r}) = \frac{-E_0\mathbf{r}}{e}. \quad (2.1)$$

This requires an electric potential of the form

$$\Phi = (\alpha x^2 + \beta y^2 + \gamma z^2), \quad (2.2)$$

with the constants $\alpha, \beta, \gamma = \frac{E_0}{2e}$. However, elementary electrostatics forbids this form of potential, as to satisfy Laplace’s equation in the absence of free charges,

$$\nabla^2\Phi = 0, \quad (2.3)$$

α, β and γ cannot all be positive. Dropping the spherical symmetry we had desired, but maintaining a cylindrical symmetry around the z axis, we find that a potential of the form

$$\Phi = \frac{\Phi_0}{2d_0^2} (2z^2 - r^2) \quad (2.4)$$

is valid, where $r^2 = x^2 + y^2$ and d_0 is the characteristic dimension of the trap.¹ This potential may be produced by placing electrodes at the equipotential surfaces of Φ , which here take the form of three hyperboloids of revolution of 1- and 2-sheets, known as the ‘ring’ and ‘endcaps’ of the trap. If the potential between the electrodes is Φ_0 , then choosing the inner radius, r_0 , and axial separation, $2z_0$, of the electrodes such that

$$r_0 = \sqrt{2} z_0 = d_0 \quad 2d_0^2 = 2z_0^2 + r_0^2 \quad (2.5)$$

produces a 3-dimensional quadrupole or saddle potential that provides harmonic trapping along the z axis but is repulsive in the radial plane.

In practice, these requirements on the geometry of the trapping electrodes may be relaxed to a great degree [38] while still maintaining a potential that is very close to quadratic in the region around the centre of the trap. Cylindrical [39] and linear [40] geometries are common, while the last decade has seen an increasing number of traps produced with all the electrodes lying in a single plane or on two opposing planes, permitting trapping of ions on chip [41, 42]. The trap used throughout this thesis consists of a stack of rings of identical radius; for more details see Section 3.4.

This is the best we can do with a static electric field, but there are two common techniques to overcome this problem. The most popular approach, known as the radiofrequency (RF) or Paul trap, superimposes an additional oscillating potential between the two electrodes. This causes the polarity of the saddle to rapidly flip between the original configuration and one that is attractive in the radial direction and repulsive in the axial. For certain ratios of the static and periodic fields, trapping is possible.² Alternatively, by maintaining the static quadrupole potential and adding a powerful magnetic field along the axis of the trap, the magnetic forces on the ion (as it is drawn towards the ring electrode) will cause it to curl into an orbit, preventing its escape from the sides of the trap. This is known as a Penning trap and is the focus of our group’s work.

2.1.1 Classical motion in a Penning trap

For a quadrupole electric potential combined with an axial magnetic field B_z the force on a positive ion is

$$\mathbf{F} = e(\mathbf{v} \times \mathbf{B} - \nabla\Phi) = e \left[\left(v_y B_z + \frac{\Phi_0}{d_0^2} \right) \hat{\mathbf{x}} + \left(-v_x B_z + \frac{\Phi_0}{d_0^2} \right) \hat{\mathbf{y}} + \frac{2\Phi_0}{d_0^2} \hat{\mathbf{z}} \right]. \quad (2.6)$$

¹Throughout this chapter, r will be defined in this way as the distance from the origin in the radial plane alone - some texts use ρ instead.

²RF traps have been covered comprehensively in many texts (e.g. [43, 44]), are not used in our experiment and will not be discussed further here.

Substituting the cyclotron frequency

$$\omega_c = \frac{eB_z}{m} \quad (2.7)$$

the equations of motion become:

$$\begin{aligned} \frac{d^2x}{dt^2} - \omega_c \frac{dy}{dt} - \frac{\omega_z^2 x}{2} &= 0 \\ \frac{d^2y}{dt^2} + \omega_c \frac{dx}{dt} - \frac{\omega_z^2 y}{2} &= 0 \\ \frac{d^2z}{dt^2} + \omega_z^2 z &= 0. \end{aligned} \quad (2.8)$$

As the magnetic field is axial, there is no component of the Lorentz force in this direction. Therefore, the ion's axial motion is decoupled from the radial motion and it performs simple harmonic oscillation in the static, quadratic electric potential with frequency

$$\omega_z = \sqrt{\frac{2e\Phi_0}{md_0^2}}. \quad (2.9)$$

The x and y motions are coupled and can be solved by combining them into a single differential equation of a complex variable, $s = x + iy$

$$\frac{d^2s}{dt^2} + i\omega_c \frac{ds}{dt} - \frac{1}{2}\omega_z^2 s = 0. \quad (2.10)$$

Substituting the trial solution $s = e^{-i\omega t}$ gives the characteristic equation

$$\omega^2 - \omega_c \omega + \frac{1}{2}\omega_z^2 = 0, \quad (2.11)$$

with solutions

$$\begin{aligned} \omega_+ &= \frac{1}{2} \left(\omega_c + \sqrt{\omega_c^2 - 2\omega_z^2} \right) = \frac{\omega_c}{2} + \omega_1 \\ \omega_- &= \frac{1}{2} \left(\omega_c - \sqrt{\omega_c^2 - 2\omega_z^2} \right) = \frac{\omega_c}{2} - \omega_1. \end{aligned} \quad (2.12)$$

The faster oscillation ω_+ is known as the modified cyclotron frequency while the slower ω_- is the magnetron frequency; the radial motion is a linear combination of these two modes. Not all solutions are stable, as no real solutions exist when the discriminant of ω_1 becomes negative, setting a limit on the axial confinement for a given magnetic field:

$$\begin{aligned} 2\omega_z^2 &\leq \omega_c^2 \\ \Phi_0 &\leq \frac{ed_0^2 B_z^2}{4m}. \end{aligned} \quad (2.13)$$




Figure 2.1: Radial trajectories in a Penning trap. Left panel: $R_m = R_c$, $\omega_m = \omega_c/4.52$. Right panel: $R_m = 4R_c$, $\omega_m = \omega_c/7$.

Conceptually, if the axial confinement is too tight, the ion will be ‘squeezed’ out radially and no stable orbit will be possible.

A charged particle moving in a magnetic field undergoes pure cyclotron motion, but the addition of an electric field at an angle to \mathbf{B} gives rise to an additional $\mathbf{E} \times \mathbf{B}$ drift, which acts perpendicular to both fields. In the Penning trap, the electric field has radial symmetry and so the average force is perpendicular to the radius vector, causing the cyclotron motion’s centre to drift in a second, slower orbit. This is the origin of the magnetron motion, while the modified cyclotron frequency is the frequency of the drifting cyclotron motion observed from the laboratory frame, as $\omega_+ = \omega_c - \omega_-$. The overall radial motion follows an epitrochoidal path. Two possible radial trajectories are shown in Figure 2.1.

Certain bounds exist on the range of values these eigenfrequencies can take. The axial frequency ω_z can be higher or lower than ω_+ but is always greater than ω_- . At the limit of stability, the magnetron and modified cyclotron motions become degenerate at the half the true cyclotron frequency, $\frac{\omega_c}{2}$. As the axial frequency decreases, it becomes lower than the modified cyclotron when $\omega_z = \omega_+ = \frac{2\omega_c}{3}$. Approaching $\omega_z = 0$, the magnetron frequency approaches zero and the modified cyclotron reaches the normal cyclotron frequency, as no electric field is present and the only force is due to $\mathbf{v} \times \mathbf{B}$.

The trajectories of the ions are given by the general solution to the radial equation

$$s = R_- e^{-i(\omega_- t + \phi_-)} + R_+ e^{-i(\omega_+ t + \phi_+)} \quad (2.14)$$

where the radii of the orbits R_- and R_+ and their phases ϕ_- and ϕ_+ are determined by the position and momentum of the ion upon its creation in the trap.

Taking real and imaginary parts, and the solution of the axial equation, we find the

trajectories in Cartesian coordinates:

$$\begin{aligned}
 x &= R_- \cos(\omega_- t + \phi_-) + R_+ \cos(\omega_+ t + \phi_+) \\
 y &= -R_- \sin(\omega_- t + \phi_-) - R_+ \sin(\omega_+ t + \phi_+) \\
 z &= R_z \cos(\omega_z t + \phi_z).
 \end{aligned} \tag{2.15}$$

The energy spectrum of the Penning trap is found by evaluating the Hamiltonian, choosing the symmetric gauge so that the vector potential $A = B_z(\frac{-y}{2}, \frac{x}{2}, 0)$ and thus

$$\begin{aligned}
 H &= \frac{(\mathbf{p} - e\mathbf{A})^2}{2m} + e\Phi \\
 &= \frac{1}{2m}p^2 + \left(\frac{eB_z}{m}\right)^2 \frac{mr^2}{8} - \frac{eB_z}{2m}(p_y x - p_x y) + \frac{m\omega_z^2}{4}(2z^2 - r^2) \\
 &= \frac{1}{2m}p^2 + \frac{m}{2}\omega_1^2 r^2 - \frac{\omega_c}{2}L_z + \frac{m}{2}\omega_z^2 z^2,
 \end{aligned} \tag{2.16}$$

where $\omega_1 = \sqrt{\omega_c^2 - 2\omega_z^2}/2$ and L_z is the z component of the angular momentum about the trap centre. Substituting the relations for x , y and z and their derivatives, the total energy of the ion can be expressed as the sum of contributions from the axial, modified cyclotron and magnetron modes:

$$H = \frac{1}{2}mR_z^2\omega_z^2 + m\omega_1(R_+^2\omega_+ - R_-^2\omega_-). \tag{2.17}$$

It is now clear that the magnetron motion is somewhat unusual as the orbit has a negative total energy. Its radius increases as it dissipates energy to the environment making the magnetron motion unstable, though the timescale over which this occurs via synchrotron radiation (years) is so slow that for practical purposes it can be considered metastable [43]. However, as will be discussed in Section 2.3.3, this aspect of Penning trap motion presents a difficulty when trying to laser cool the ions towards the zero point motion in each mode, as energy must be removed from the cyclotron and axial modes whilst being added to the magnetron mode.

Another curious property of the motion in this field is highlighted by writing the total cycle-averaged kinetic energy

$$\langle E_k \rangle = \frac{m\omega_z^2 R_z^2}{4} + \frac{m\omega_+^2 R_+^2}{2} + \frac{m\omega_-^2 R_-^2}{2}. \tag{2.18}$$

Comparing the system with $(R_+ = R_1, R_- = 0)$ and $(R_- = R_1, R_+ = 0)$ it is apparent that an ion can have two alternative, stable modes of motion along the *same* radial trajectory, each with a different frequency and kinetic energy, but the same potential energy.

Rotating frame

It is often useful to transform to a frame rotating with the same handedness as the motion at a frequency of $\omega_c/2$. In this frame the potential is transformed to that of a two-dimensional harmonic oscillator and the modified cyclotron and magnetron motions are left and right handed circular motions with frequencies $\pm\omega_1$ [34, 45]. The source of the degenerate orbits mentioned above is now obviously the left and right handed orbits of similar radius. Returning to the Hamiltonian of Equation 2.16, we can separate the contribution of the radial motion as

$$H_r = \frac{1}{2m}p_r^2 + \frac{m}{2}\omega_1^2 r^2 - \frac{\omega_c}{2}L_z. \quad (2.19)$$

The first two terms describe an isotropic two-dimensional harmonic oscillator with frequency ω_1 , correctly describing the dynamics in the rotating frame. As this frame rotates at $\omega_c/2$, the cross term, $-\omega_c\hat{L}_z/2$, simply represents the transformation back to the laboratory frame.

2.1.2 Quantised motion in a Penning trap

To achieve the coherence times necessary for quantum information experiments, or the stability required for examining Coulomb crystals, the ions must be cooled towards the zero-point motional energy, and a quantum description of their motion is necessary.

In the classical treatment we started with the equations of motion, solved for the trajectories, and then returned to substitute these for the canonical variables of the Hamiltonian to derive the energy of the system. Rather than substituting the classical trajectories, quantization proceeds directly from the use of Hermitian operators in place of the canonical variables of the Hamiltonian. The Schrödinger equation for stationary states (eigenstates) of the system is then

$$\left[\frac{1}{2m} (\hat{\mathbf{p}} - e\hat{\mathbf{A}})^2 + e\hat{\Phi} \right] |\psi\rangle = E |\psi\rangle. \quad (2.20)$$

To calculate the wavefunctions of the quantized system, the Schrödinger equation may be solved directly. However, unless solving time-dependent problems, a more powerful and intuitive approach is to generate the eigen-spectrum of the system with creation and annihilation operators similar to those used in the one dimensional quantum harmonic oscillator. Before this can be performed, the contributions of each motional mode to the Hamiltonian need to be separated.

The Hilbert space of a composite quantum system is given by the tensor product of the Hilbert spaces of the components, which for our system (ignoring the spin of the ion) is

$$\mathcal{H} = \mathcal{H}_+ \otimes \mathcal{H}_- \otimes \mathcal{H}_z, \quad (2.21)$$

with wavefunctions given by

$$|\psi\rangle = |\psi_+\rangle \otimes |\psi_-\rangle \otimes |\psi_z\rangle. \quad (2.22)$$

Because in the case of the ideal Penning trap the three modes are uncoupled, operators on the system are linear combinations of operators that act on the single component spaces. This means that if the Hamiltonian of the system is split into three components, we can treat each separately, and the total energy of the system will be the sum of the component energies.

Expanding the Hamiltonian in terms of the kinetic momentum $m\hat{\mathbf{r}} = \hat{\mathbf{p}}_r - e\hat{\mathbf{A}}$ and noting that $m\hat{z} = \hat{p}_z$ as the magnetic field is axial, the axial and radial contributions are readily isolated:

$$\hat{H} = \frac{m}{2} \left[\hat{\mathbf{r}}^2 - \frac{\omega_z^2}{2} \mathbf{r}^2 \right] + \left[\frac{1}{2m} \hat{p}_z^2 + \frac{m}{2} \omega_z^2 \hat{z}^2 \right] = \hat{H}_r + \hat{H}_z. \quad (2.23)$$

H_z is the Hamiltonian for a 1-dimensional quantum harmonic oscillator. Defining creation and annihilation operators

$$\begin{aligned} \hat{a}_z^\dagger &= \sqrt{\frac{m}{2\hbar\omega_z}} \left(z\omega_z - i\frac{p}{m} \right) \\ \hat{a}_z &= \sqrt{\frac{m}{2\hbar\omega_z}} \left(z\omega_z + i\frac{p}{m} \right), \end{aligned} \quad (2.24)$$

the axial Hamiltonian can be written as

$$H_z = \hbar\omega_z \left(\hat{a}_z^\dagger \hat{a}_z + \frac{1}{2} \right), \quad (2.25)$$

generating the energy spectrum for the oscillator

$$E_z = \hbar\omega_z \left(k + \frac{1}{2} \right). \quad (2.26)$$

We now aim to quantize the radial motion in a similar manner. Following the treatment by Brown and Gabrielse [46], we define two new operators

$$\begin{aligned} \mathbf{V}^+ &= \dot{\mathbf{r}} + \omega_- (\hat{\mathbf{z}} \times \mathbf{r}) \\ \mathbf{V}^- &= \dot{\mathbf{r}} + \omega_+ (\hat{\mathbf{z}} \times \mathbf{r}) \\ \mathbf{V}^{\pm 2} &= V_x^{\pm 2} + V_y^{\pm 2}. \end{aligned} \quad (2.27)$$

It can be shown that these operators rotate in \mathbf{r} space with frequencies ω_+ and ω_- respectively:

$$\dot{\mathbf{V}}^\pm = \omega_\pm \hat{\mathbf{z}} \times \mathbf{r}. \quad (2.28)$$

The radial Hamiltonian can therefore be expressed as

$$\hat{H}_r = \frac{m}{4\omega_1} \left(\omega_+ \mathbf{V}^{+2} - \omega_- \mathbf{V}^{-2} \right). \quad (2.29)$$

From the canonical commutation relations, we find the Cartesian operator commutation relations (with i, j representing radial coordinates x or y):

$$\begin{aligned} [r_i, r_j] &= 0 \\ [r_i, \dot{r}_j] &= \frac{i\hbar}{m} \delta_{ij} \\ [\dot{r}_i, \dot{r}_j] &= \frac{i\hbar}{m} \omega_c (1 - \delta_{ij}). \end{aligned} \quad (2.30)$$

Using these, the commutation relations for the new operators can be evaluated

$$\begin{aligned} [V_i^+, V_j^-] &= [\dot{r}_i, \dot{r}_j] + \omega_+ [\dot{r}_i, r_i] (1 - \delta_{ij}) - \omega_- [r_j, \dot{r}_j] (1 - \delta_{ij}) + \omega_- \omega_+ [r_i, r_j] \\ &= \frac{i\hbar}{m} (\omega_c - \omega_+ - \omega_-) (1 - \delta_{ij}) \equiv 0 \\ [V_x^+, V_y^+] &= \frac{2\omega_1}{m} ([y, p_y] - [p_x, x]) \\ &= \frac{2i\hbar\omega_1}{m} \\ [V_x^-, V_y^-] &= \frac{-2i\hbar\omega_1}{m}. \end{aligned} \quad (2.31)$$

The first commutation relation shows that the two operators are independent and thus the modes have been successfully uncoupled. The others show that for each mode, V_x and V_y are a canonically conjugate pair of operators representing independent observables of the system. Using these relations, we can rewrite the modified cyclotron Hamiltonian as

$$\hat{H}_+ = \hbar\omega_+ \left(\hat{a}_+^\dagger \hat{a}_+ + \frac{1}{2} \right) \quad (2.32)$$

where the cyclotron creation and annihilation operators are

$$\begin{aligned} \hat{a}_+^\dagger &= \sqrt{\frac{m}{4\hbar\omega_1}} (V_x^+ - iV_y^+) \\ \hat{a}_+ &= \sqrt{\frac{m}{4\hbar\omega_1}} (V_x^+ + iV_y^+). \end{aligned} \quad (2.33)$$

Similarly, the magnetron Hamiltonian is written as

$$\hat{H}_- = -\hbar\omega_- \left(\hat{a}_-^\dagger \hat{a}_- + \frac{1}{2} \right) \quad (2.34)$$

with

$$\begin{aligned}\hat{a}_-^\dagger &= \sqrt{\frac{m}{4\hbar\omega_1}} \left(V_x^- - iV_y^- \right) \\ \hat{a}_- &= \sqrt{\frac{m}{4\hbar\omega_1}} \left(V_x^- + iV_y^- \right).\end{aligned}\tag{2.35}$$

By comparison with the axial Hamiltonian, the radial modes - under this coordinate transform - behave as two uncoupled 1D harmonic oscillators with quanta of $\hbar\omega_-$ and $\hbar\omega_+$.

The total energy of the system is then given by

$$E_{nlk} = \hbar\omega_+ \left(n + \frac{1}{2} \right) - \hbar\omega_- \left(l + \frac{1}{2} \right) + \hbar\omega_z \left(k + \frac{1}{2} \right).\tag{2.36}$$

As expected, the magnetron motion has a ladder of ‘negative’ energy states and the other two motions have normal harmonic oscillator spectra. Transitions between these states can be driven via the application of an oscillating electric field of appropriate geometry between the trap electrodes, set at the frequency of the relevant transition [47]. Alternatively, as will be shown in the following section, it is also possible for optical transitions between different electronic levels to simultaneously produce changes in the harmonic oscillator state.

2.2 Laser-Ion Interactions

2.2.1 Interaction with free atom

We now consider the interaction of a monochromatic laser field with a two-level atom. This provides a good approximation to the dynamics under a single driving field that is much closer to resonance with the two levels in question than with any others. Denoting the upper and lower levels of the transition $|e\rangle$ and $|g\rangle$ respectively, with energies taken symmetrically around zero, the Hamiltonian of the unperturbed atom is

$$\mathbf{H}_0 = \frac{\hbar\omega_0}{2} |e\rangle\langle e| - \frac{\hbar\omega_0}{2} |g\rangle\langle g| = -\frac{\hbar\omega_0}{2} \sigma_z, \quad (2.37)$$

where ω_0 is the transition frequency between the two levels and σ_z is the Pauli- z operator. A general electronic state of the atom under this approximation is written as

$$|\psi(t)\rangle = a_g(t) |g\rangle + a_e(t) |e\rangle. \quad (2.38)$$

We can write the electric field due to a monochromatic, plane laser field with wavevector k_L , frequency ω_L and unit polarization vector $\hat{\epsilon}_L$ as

$$\mathbf{E}(t) = \frac{\hat{\epsilon}_L E_0}{2} \left(e^{i(\mathbf{k}_L \cdot \mathbf{r} - \omega_L t)} + e^{-i(\mathbf{k}_L \cdot \mathbf{r} - \omega_L t)} \right). \quad (2.39)$$

When driving a single atom on an electric dipole (E1) transition, we can approximate the spatial exponent by the first order term of its Taylor expansion $e^{i\mathbf{k}_L \cdot \mathbf{r}} = 1 + i\mathbf{k}_L \cdot \mathbf{r} + \dots \sim 1$. This is known as the dipole approximation and ignores the small spatial variance of the laser field over the atom. The interaction term of the Hamiltonian is then

$$\mathbf{H}_i = \frac{\hbar\Omega_0}{2} \left(e^{-i\omega_L t} + e^{i\omega_L t} \right) (\sigma_+ + \sigma_-), \quad \mathbf{H} = \mathbf{H}_0 + \mathbf{H}_i \quad (2.40)$$

where $\sigma_{\pm} = \hat{\sigma}_x \pm i\hat{\sigma}_y$ are Pauli matrices and Ω_0 is the Rabi frequency, which quantifies the strength of the coupling. The definition of the Rabi frequency depends on which term of the multipole expansion of the Hamiltonian dominates the interaction. For the electric dipole transition considered here, the Rabi frequency is given by $\Omega_0 = \boldsymbol{\mu}_{eg} \cdot \boldsymbol{\epsilon}_L E_0 / \hbar$ where $\boldsymbol{\mu}_{eg} = e \langle e | \hat{\mathbf{r}} | g \rangle$ is the matrix element of the atomic dipole operator coupling the ground and excited states.³

Solving the Schrödinger equation for 2.40 is made more straightforward by transforming to the interaction representation, a frame rotating at the transition frequency. The

³For the work in the strong-binding (resolved sideband) regime described later in this chapter, we use dipole-forbidden, electric quadrupole (E2) transitions. Here the amplitude associated with the leading order term of $e^{i\mathbf{k}_L \cdot \mathbf{r}}$ goes to zero and we must consider the next term in the expansion. A more detailed discussion and calculation of matrix elements for both electric dipole and quadrupole transitions is given in Appendix A.

Schrödinger equation becomes

$$i\hbar \frac{\partial}{\partial t} |\psi(t)_I\rangle = \mathbf{H}_I |\psi(t)_I\rangle \quad (2.41)$$

for

$$|\psi(t)_I\rangle = e^{i\mathbf{H}_0 t/\hbar} |\psi(t)_S\rangle = c_g(t) |\tilde{g}\rangle + c_e(t) |\tilde{e}\rangle \quad \mathbf{H}_I = e^{i\mathbf{H}_0 t/\hbar} \mathbf{H}_i e^{-i\mathbf{H}_0 t/\hbar} \quad (2.42)$$

where the eigenstates now contain the ω_0 time dependence as $|\tilde{g}, \tilde{e}\rangle = e^{\pm i\omega_0 t/2} |g, e\rangle$.

The Hamiltonian in the interaction representation is

$$\mathbf{H}_I = \frac{\hbar}{2} \begin{pmatrix} 0 & \Omega_0 e^{i\delta t} \\ \Omega_0 e^{-i\delta t} & 0 \end{pmatrix}, \quad (2.43)$$

where $\delta = \omega_L - \omega_0$ is the detuning of the laser from resonance and the rotating wave approximation has been applied to remove the $e^{i(\omega_0 + \omega_L)}$ terms which oscillate very fast on the timescale of the other dynamics ($\omega_0 + \omega_L \gg \delta, \Omega_0$). The transformed Schrödinger equation gives the evolution of the state amplitudes as

$$\dot{c}_g(t) = \frac{i\Omega_0}{2} e^{i\delta t} c_e(t) \quad \dot{c}_e(t) = \frac{i\Omega_0}{2} e^{-i\delta t} c_g(t). \quad (2.44)$$

The solution of this system of equations for initial conditions of $c_g(0) = 1$, $c_e(0) = 0$ is

$$\begin{aligned} c_g(t) &= e^{+i\delta t/2} \left[\cos\left(\frac{\Omega t}{2}\right) - i \frac{\delta}{\Omega} \sin\left(\frac{\Omega t}{2}\right) \right] \\ c_e(t) &= e^{-i\delta t/2} \left[-i \frac{\Omega_0}{\Omega} \sin\left(\frac{\Omega t}{2}\right) \right] \end{aligned} \quad (2.45)$$

where $\Omega = \sqrt{\Omega_0^2 + \delta^2}$ is the generalised Rabi frequency.

The excited state population evolves as

$$|c_e(t)|^2 = \frac{\Omega_0^2}{\Omega^2} \sin^2\left(\frac{\Omega t}{2}\right). \quad (2.46)$$

On resonance ($\delta = 0$) the population undergoes full amplitude Rabi oscillations between the two eigenstates, at frequency Ω_0 . As the laser is detuned from resonance the oscillation amplitude decreases while the frequency increases, producing a time-averaged resonance lineshape that is Lorentzian with amplitude 0.5 and FWHM = $2\Omega_0$.

Spontaneous decay

We have so far considered the two levels of the atom to be absolutely stable with all transitions between them mediated by the driving laser. However, all states also interact

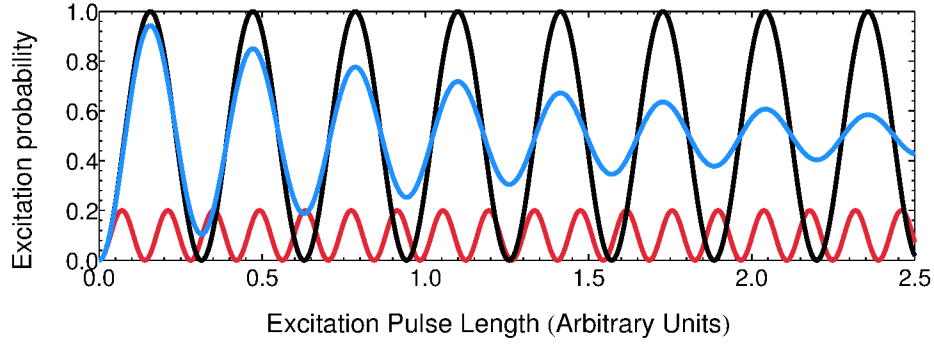


Figure 2.2: Rabi oscillations on resonance ($\delta = 0$, black), off-resonance ($\delta = 2 \times \Omega_0$, red) and in the presence of spontaneous decay ($\delta = 0$, $\Gamma = 0.05 \times \Omega_0$, blue).

with vacuum fluctuations due to the zero-point energy of the electromagnetic field and these can also cause any excited state to spontaneously decay to a lower level. Because this involves the interaction of the two-level system with an external environment of which we have only limited information, the above approach, considering the amplitudes of the states, breaks down. To tackle this problem we introduce the density matrix which in the interaction picture is given by

$$\boldsymbol{\varrho} = \begin{bmatrix} \varrho_{gg} & \varrho_{ge} \\ \varrho_{eg} & \varrho_{ee} \end{bmatrix} = \begin{bmatrix} |c_g|^2 & c_g c_e^* \\ c_e c_g^* & |c_e|^2 \end{bmatrix} \quad (2.47)$$

where the second equality applies only when there is no interaction with the environment. The diagonal elements represent the population in the ground and excited states, while the off-diagonal elements describe the coherences between those populations. Including the effect of spontaneous decay [48], the density matrix evolves as

$$i\hbar \frac{d\boldsymbol{\varrho}}{dt} = [\mathbf{H}_I, \boldsymbol{\varrho}] + \begin{bmatrix} \Gamma \varrho_{ee} & -\frac{\Gamma}{2} \varrho_{ge} \\ -\frac{\Gamma}{2} \varrho_{eg} & -\Gamma \varrho_{ee} \end{bmatrix} \quad (2.48)$$

for an upper state with spontaneous emission rate Γ , leading to the optical Bloch equations

$$\begin{aligned} \dot{\varrho}_{gg} &= \frac{i\Omega_0}{2} (e^{-i\delta t} \varrho_{ge} - e^{i\delta t} \varrho_{eg}) + \Gamma \varrho_{ee} \\ \dot{\varrho}_{ee} &= \frac{i\Omega_0}{2} (e^{i\delta t} \varrho_{eg} - e^{-i\delta t} \varrho_{ge}) - \Gamma \varrho_{ee} \\ \dot{\varrho}_{ge} &= \frac{i\Omega_0}{2} e^{i\delta t} (\varrho_{gg} - \varrho_{ee}) - \frac{\Gamma}{2} \varrho_{ge} \\ \dot{\varrho}_{eg} &= \frac{i\Omega_0}{2} e^{-i\delta t} (\varrho_{ee} - \varrho_{gg}) - \frac{\Gamma}{2} \varrho_{eg}. \end{aligned} \quad (2.49)$$

These can be solved numerically, and the dynamics take the form of damped Rabi oscillations that decay towards their equilibrium value, as shown in Figure 2.2.

We consider the steady state solution after the Rabi oscillations have reached their equilibrium value. Setting the time derivatives to zero and solving by substitution we find

$$\varrho_{ee} = \frac{\Omega_0^2/4}{\Gamma^2/4 + \Omega_0^2/2 + \delta^2}. \quad (2.50)$$

The power-broadened Lorentzian linewidth $\Gamma_B = \sqrt{2\Omega_0^2 + \Gamma^2}$ has contributions from both the natural linewidth of the transition and the Rabi linewidth described in the previous section. Equation 2.50 can be written in terms of the saturation parameter, s , as

$$\varrho_{ee} = \frac{1}{2} \left(\frac{s}{1+s} \right) \quad \text{for} \quad s = \frac{\Omega_0^2/2}{\Gamma^2/4 + \delta^2}. \quad (2.51)$$

In this form it is immediately apparent that population inversion is impossible in two-level systems without coherent control. As $s \rightarrow \infty$, $\varrho_{ee} \rightarrow 0.5$. The saturation parameter at $\delta = 0$ provides a useful measure of the laser intensity,

$$s_0 = \frac{2\Omega_0^2}{\Gamma^2} = \frac{I}{I_{sat}}. \quad (2.52)$$

On resonance at $I = I_{sat}$, $1/4$ of the population is in the excited state, and the linewidth is $\sqrt{2} \times \Gamma$.

Doppler broadening

We now consider the effect of classical motion of the atom upon the interaction. In the rest frame of an ion moving with velocity \mathbf{v} , the laser frequency appears Doppler shifted as

$$\omega_L \rightarrow \omega_{L,Dopp} = \omega_L - \mathbf{k}_L \cdot \mathbf{v}. \quad (2.53)$$

Equation 2.50 can therefore be rewritten in the presence of atomic motion as

$$\varrho_{ee}(\mathbf{v}) = \frac{\Omega_0^2/4}{\Gamma_B^2/4 + (\delta + \mathbf{k}_L \cdot \mathbf{v})^2}. \quad (2.54)$$

If the velocity was constant the line would be shifted uniformly away from resonance. However, if an atom is allowed to thermalise (e.g. through interaction with a hot environment, or via laser cooling) then repeated measurement will sample a Boltzmann distributed range of energies, giving a Gaussian velocity distribution in the propagation direction of the laser

$$D(v) = \frac{e^{-v^2/v_0^2}}{v_0\sqrt{\pi}} \quad (2.55)$$

where $v_0 = \sqrt{2k_B T/m}$ is the width of the velocity distribution in one dimension for temperature T .

To calculate the effect on the lineshape of the transition, we average the steady state

population over this distribution for a laser propagating in the z direction:

$$\langle \varrho_{ee} \rangle_v = \frac{\Omega_0^2/4}{v_0\sqrt{\pi}} \int_{-\infty}^{\infty} \frac{e^{-v_z^2/v_0^2}}{\Gamma_B^2/4 + (\delta + k_L v_z)^2} dv_z \quad (2.56)$$

This integral is the convolution of Gaussian and Lorentzian profiles, known as a Voigt profile. In the limit $v_z \rightarrow 0$ we recover the original Lorentzian of Equation 2.50. At the other extreme, if $k_L v_z \gg \Gamma_B$, the Lorentzian component becomes sharp compared to the width of the Gaussian and the integral becomes - up to a constant - the convolution of the Gaussian with a delta function at $v_z = \delta/k_L$. We therefore find a Gaussian lineshape with (half-)width $k_L v_0$. Under normal environmental conditions, Doppler broadening is the dominant contributor to the linewidth - for $^{40}\text{Ca}^+$ at $T = 300$ K, the 397-nm transition is broadened to a full width of $\sim 2\pi \times 1$ GHz, over 40 times greater than the natural linewidth of $\Gamma \sim 2\pi \times 21$ MHz.

While this analysis of atomic motion provides a good model for free or weakly-trapped atoms at relatively high temperatures, the classical approximation breaks down as quantisation of the energy becomes significant. This is especially important as the quantum number becomes low enough that the energetic limit presented by the quantum ground state starts to introduce asymmetry to the spectrum. Furthermore, to properly describe the Doppler cooling of trapped ions - or to describe sideband cooling at all - we must consider the effects of the trap potential on the laser-ion interaction.

2.2.2 Interaction with trapped atom

For a two-level atom trapped in a harmonic potential, the unperturbed Hamiltonian becomes⁴

$$\mathbf{H}_0 = \frac{\hat{p}^2}{2m} + \frac{m\omega_z^2 \hat{z}^2}{2} + \frac{\hbar\omega_0}{2} \boldsymbol{\sigma}_z. \quad (2.57)$$

Assuming now that the laser has wavevector $\mathbf{k}_L = (0, 0, k_z)$, the interaction term of the Hamiltonian is

$$\mathbf{H}_i = \frac{\hbar\Omega_0}{2} \left(e^{i(k_L \hat{z} - \omega_L t)} + e^{-i(k_L \hat{z} - \omega_L t)} \right) (\boldsymbol{\sigma}_+ + \boldsymbol{\sigma}_-). \quad (2.58)$$

We now proceed by writing \hat{z} and \hat{p} in terms of the creation and annihilation operators as defined in 2.24 so that

$$\begin{aligned} \hat{z} &= \sqrt{\frac{\hbar}{2m\omega_z}} (\hat{a} + \hat{a}^\dagger) = z_0 (\hat{a} + \hat{a}^\dagger) \\ \hat{p} &= i\sqrt{\frac{\hbar m\omega_z}{2}} (\hat{a} - \hat{a}^\dagger) = \frac{i\hbar}{2z_0} (\hat{a}^\dagger - \hat{a}) \end{aligned} \quad (2.59)$$

⁴For the majority of this chapter we will consider only one-dimensional harmonic motion, which in a Penning trap applies only to the axial motion ω_z . In Section 2.3.3 we will show how this can be generalised to include the cyclotron and magnetron motion.

and define the Lamb-Dicke parameter:

$$\eta = k_L z_0 = \sqrt{\frac{\hbar k_L^2}{2m\omega_z}}. \quad (2.60)$$

The Lamb-Dicke parameter is the product of the wavenumber and the ground state wavefunction spread, z_0 , quantifying the confinement of the atom relative to the wavelength of the light. It is this gradient of the electric field with respect to the harmonic oscillator wavefunction's extent which determines the probability of a spontaneous decay leading to a change in phonon number, as is hinted at by the alternative definition of η^2 as the ratio of the emission recoil energy to the phonon energy.

The laser-ion Hamiltonian can then be reformulated as

$$\begin{aligned} \mathbf{H}_0 &= \hbar\omega_z(\hat{a}^\dagger\hat{a} + \frac{1}{2}) + \frac{\hbar\omega_0}{2}\boldsymbol{\sigma}_z \\ \mathbf{H}_i &= \frac{\hbar}{2} \left(\Omega_0 e^{i(\eta(\hat{a}^\dagger + \hat{a}) - \omega_L t)} + \Omega_0 e^{-i(\eta(\hat{a}^\dagger + \hat{a}) - \omega_L t)} \right) (\boldsymbol{\sigma}_+ + \boldsymbol{\sigma}_-). \end{aligned} \quad (2.61)$$

Transforming to the interaction representation as per Equation 2.42 and making the rotating wave approximation leads to

$$\mathbf{H}_I = \begin{pmatrix} 0 & \Omega_0 e^{-i(\eta(\tilde{a}^\dagger + \tilde{a}) - \delta t)} \\ \Omega_0 e^{i(\eta(\tilde{a}^\dagger + \tilde{a}) - \delta t)} & 0 \end{pmatrix}. \quad (2.62)$$

Note that the interaction representation transform on the ladder operators is equivalent to transforming them to their Heisenberg representation [49] with $\hat{a} \rightarrow \tilde{a} = e^{-i\omega_z t}\hat{a}$. The total state of an atom with oscillator states n' may now be expressed as

$$|\psi_I\rangle = \sum_{n'} (c_{g,n'} |g, n'\rangle + c_{e,n'} |e, n'\rangle) \quad (2.63)$$

and applying the Schrödinger equation (2.41) yields

$$\sum_{n'} \dot{c}_{g,n'} |g, n'\rangle + \dot{c}_{e,n'} |e, n'\rangle = \frac{-i\Omega_0}{2} \sum_{n'} \left(c_{e,n'} e^{i\delta t} e^{-i\eta(\tilde{a}^\dagger + \tilde{a})} |g, n'\rangle + c_{g,n'} e^{-i\delta t} e^{i\eta(\tilde{a}^\dagger + \tilde{a})} |e, n'\rangle \right). \quad (2.64)$$

Pre-multiplying by $\langle g, n|$ we find

$$\begin{aligned} \dot{c}_{g,n} &= \frac{-i\Omega_0}{2} e^{i\delta t} \sum_{n'} c_{e,n'} \langle n| e^{-i\eta(\tilde{a}^\dagger + \tilde{a})} |n'\rangle \\ &= \frac{-i\Omega_0}{2} e^{i\delta t} \sum_{n'} \left[c_{e,n'} \langle n| \sum_{k=0}^{\infty} \frac{(-i\eta)^k}{k!} \left(\hat{a}^\dagger e^{-i\omega_z t} + \hat{a} e^{i\omega_z t} \right)^k |n'\rangle \right]. \end{aligned} \quad (2.65)$$

Expanding the term containing the \hat{a}, \hat{a}^\dagger operators we find a series with combinations of

the operators up to $(\hat{a})^k$ and $(\hat{a}^\dagger)^k$ showing that the exponential term provides coupling between $|n\rangle$ and every other state of the harmonic oscillator. Because the raising and lowering operators are associated with oscillation frequencies of $\mp\omega_z$, terms with p raising operators and q lowering operators (which couple state $|n\rangle$ to state $|n+p-q\rangle = |n+m\rangle$) have an oscillation frequency of $(q-p)\omega_z = -m\omega_z$. The total frequency of the term is then $e^{i(\delta-m\omega_z)t} = e^{i\Delta t}$. If the driving laser has narrow linewidth ($\Delta_{\omega_L} \ll \omega_z$), it may be tuned close to resonance ($\Delta = 0$) with one of these ‘sideband’ frequencies and the interaction with all other sidebands will be insignificant. We can therefore ignore all terms in the sum over n' except $|n+m\rangle$ and Equation 2.65 simplifies to

$$\dot{c}_{g,n} = \frac{-i}{2} e^{i\Delta t} \Omega_0 \langle n | e^{-i\eta(\hat{a}^\dagger + \hat{a})} | n+m \rangle c_{e,n+m}. \quad (2.66)$$

Similarly, by premultiplying by $\langle e, n |$ we find

$$\dot{c}_{e,n} = \frac{-i}{2} e^{-i\Delta t} \Omega_0 \langle n | e^{i\eta(\hat{a}^\dagger + \hat{a})} | n+m \rangle c_{g,n+m}. \quad (2.67)$$

These two equations are clearly of the same form as Equations 2.44, and lead to Rabi oscillations but with a modified Rabi frequency of

$$\Omega_{n,n+m} = \Omega_0 \langle n+m | e^{i\eta(\hat{a}^\dagger + \hat{a})} | n \rangle. \quad (2.68)$$

This matrix element can be evaluated [50] to give

$$\langle n+m | e^{i\eta(\hat{a}^\dagger + \hat{a})} | n \rangle = \left(\frac{n!}{(n+m)!} \right)^{\text{sgn}(m)/2} e^{-\frac{\eta^2}{2}} (\eta)^{|m|} L_n^{|m|}(\eta^2), \quad (2.69)$$

where the associated Laguerre polynomial is defined as

$$L_n^{|m|}(\eta^2) = \sum_{k=0}^n (-1)^k \binom{n+|m|}{n-k} \frac{\eta^{2k}}{k!}. \quad (2.70)$$

$\Omega_{n,n+m}$ describes the excitation on the m^{th} sideband of an atom in the $|n\rangle$ Fock state. In many cases the trapped ion will be in a thermal state, which in the terminology of quantum mechanics is a mixed state with a diagonal density matrix. To calculate the average excitation probability on resonance with a sideband m , for a thermal state with average phonon number \bar{n} , the excitations due to each Fock state must be combined incoherently, weighted by their Boltzmann distributed occupation probabilities $P(n)$ [49]:

$$P_{e,m}(t) = \sum_{n=0}^{\infty} P(n) \sin^2 \left(\frac{\Omega_{n,n+m} t}{2} \right), \quad P(n) = \frac{\bar{n}^n}{(\bar{n}+1)^{n+1}}. \quad (2.71)$$

The excitation spectrum summed over m will be a series of evenly spaced sidebands. Although it is not immediately apparent in this form, in the classical limit of $m \ll \bar{n}$,

the heights of these sidebands will follow a Gaussian envelope. This has been shown analytically by Wineland and Itano [50], or can be readily demonstrated via numerical evaluation of the sideband heights. The width of the Gaussian envelope is $k_L \sqrt{2\hbar\omega_z \bar{n}/m}$, equivalent to that found for the classical Doppler broadened model described in 2.2.1, as demanded by the correspondence principle.

Lamb-Dicke regime approximation

If the confinement of the ion, parameterized by the spread of its harmonic oscillator wavefunction $z_0 \sqrt{2\bar{n} + 1}$, is small relative to the wavelength of the laser, the ion is said to be in the Lamb-Dicke regime. Precisely, the condition

$$\eta \sqrt{2\bar{n} + 1} \ll 1 \quad (2.72)$$

must be satisfied. Under these conditions only the carrier and first order sidebands at $\omega_0 - \omega_z$ and $\omega_0 + \omega_z$ couple significantly to the laser field.⁵ A much simpler expression of Equation 2.68 can now be derived by Taylor expanding as

$$\langle n + m | e^{i\eta(\hat{a}^\dagger + \hat{a})} | n \rangle = \langle n + m | (1 + i\eta(\hat{a}^\dagger + \hat{a}) + (i\eta(\hat{a}^\dagger + \hat{a}))^2 + O\eta^3) | n \rangle. \quad (2.73)$$

With this approximation, we can readily calculate the effective Rabi frequencies for transitions on the carrier, red and blue sidebands:

$$\begin{aligned} \Omega_C &= \Omega_0(1 + \eta^2 \bar{n}) \\ \Omega_R &= \Omega_0 \eta \sqrt{\bar{n}} \\ \Omega_B &= \Omega_0 \eta \sqrt{\bar{n} + 1}. \end{aligned} \quad (2.74)$$

⁵Hereafter the ‘first red sideband’ and ‘first blue sideband’, respectively.

2.3 Laser Cooling

Early ion trapping experiments cooled the ions by bringing them into contact with a thermal bath, either in the form of a cryogenic buffer gas, or by coupling certain trap electrodes with a resistive load held at cryogenic temperatures [51]. In a Penning trap, the buffer gas method fails because the negative total energy of the magnetron motion causes this to rapidly increase in radius (henceforth ‘heat’, despite the fallacy of this description) when exposed to a thermal source [52]. Resistive cooling can be applied in a more controlled way such that only, e.g., the axial motion is cooled, but the same problem with magnetron heating will occur if the cyclotron motion is cooled, as the two radial modes cannot be decoupled. The magnetron motion may be cooled indirectly by coupling it to the axial motion with a tilted quadrupole field [46]. However, it remains impossible to reduce the temperature below that of the cryogenic load resistor or buffer gas.⁶ While achieving 4 K temperatures is therefore fairly straightforward, moving beyond this limit would require the entire experiment to be conducted on the cold plate of a helium-3- or dilution-refrigerator.

I will now show how the interactions defined above can be used to cool a trapped ion to temperatures far lower than is possible with resistive or buffer gas cooling, and indeed to the ground state itself. Laser cooling can be considered in two regimes - the ‘weak-binding’ regime where $\omega_z \ll \Gamma$ and the ‘strong-binding’ regime where $\omega_z \gg \Gamma$.

In the weak-binding regime, the motion of the ion is well approximated by a classical oscillator in a thermal state. We cannot easily limit absorption to a single sideband⁷ but we can weight absorption towards the red sidebands by detuning the laser from resonance. At high temperatures where the Doppler broadening is much larger than Γ this is a highly effective method of cooling, but as the broadening becomes insignificant, the lack of frequency selectivity means diffusive heating processes grow to balance the cooling rate.

In the strong-binding regime, the ions’ spectrum is characterised by an envelope of discrete, well-resolved sidebands. The scattering rate is lower due to the necessarily narrower linewidth of the transition, so cooling rates are limited and it is essential to provide some sort of pre-cooling stage, generally by Doppler cooling. However, absorption on a particular red sideband can be ensured via frequency selection, allowing much lower temperatures to be reached; for reasonably high trap frequencies (typically $\omega_z > 2\pi \times 100$ kHz) it is possible to cool the ion to the quantum ground state of the motion. The final energy is limited only by environmental heating rates, off-resonant interactions with other sidebands and the single photon recoil limit itself.

⁶In fact, the equilibrium temperature of the indirectly cooled magnetron motion is approximately that of the load multiplied by the ratio ω_-/ω_z , and therefore somewhat lower. This is because it is the oscillator phonon numbers that come into equilibrium rather than the temperatures.

⁷EIT cooling [53] targets the red sideband by using a narrow spectral feature produced by the interference of the cooling laser with another laser coupled to a third, ancillary level. While this arguably allows sideband-specific absorption while in the weak-binding regime, the method is significantly more complicated than Doppler cooling.

I begin this section by giving a simple description of Doppler cooling, to illustrate the general principle and highlight some important features that effect the final temperature. I provide a calculation of the ‘relaxation’ heating that occurs as the beam is switched off, an important consideration in low-frequency traps. Laser cooling in a Penning trap is significantly complicated by the properties of the radial motion, which is unstable under conventional Doppler cooling. I briefly discuss two alternative techniques used in our experiment that allow effective cooling of the radial motion, via either an inhomogeneous cooling beam or an additional quadrupolar ‘axialisation’ field. Finally I describe the principles of resolved sideband cooling and discuss some important considerations when applying the technique to the Penning trap.

2.3.1 Doppler cooling

Section 2.2.1 described the manner in which a cloud of atoms, or weakly trapped ions, interacted with a near-resonant laser. It was shown that when the temperature was large enough to significantly broaden the transition, only a small fraction of the population would interact with the laser at any one time - those whose velocities provided a Doppler shift equal and opposite to the detuning of the laser from resonance. In this manner, by detuning to the red of the transition (decreasing the wavelength), the laser can be made to interact only with ions that were moving towards it at the moment of interaction. Similarly, by blue detuning the interaction will only occur with ions moving away.

Each time a photon is absorbed, the ion receives a momentum impulse of $\hbar\mathbf{k}_L$, in the direction of laser propagation. The absorbed photon will be re-emitted, either spontaneously by interaction with vacuum field modes, or by stimulated emission back into the driving field. In the latter case, the recoil momentum of the ion is $-\hbar\mathbf{k}_L$ and there is no net effect on its motion. If the emission is spontaneous, it will be isotropic⁸ and over many scattering events, the recoil momentum will average to zero, leaving a net momentum change per scatter of $\hbar\mathbf{k}_L$. Because the scattering only occurs for those ions moving towards the incoming beam (for red detuning), the effect is to gradually decrease their momentum, and so the temperature of the ensemble.

I now give a basic derivation of cooling by such a mechanism. While this approach does not capture the full dynamics, especially regarding the interaction with the radial motion, it provides a concise and transparent description of the general properties and behaviour. The results of a more precise treatment of laser cooling in a Penning trap are discussed in 2.3.3.

In the weak-binding limit (unresolved sidebands), scattering is essentially instantaneous on the timescale of the harmonic oscillator period and we can model the cooling process as

⁸The actual radiation pattern will depend on the type of transition, but isotropy is a reasonable approximation for development of the basic argument.

that of a free atom. The radiative force on the ions is given by the product of the scattering rate (Equation 2.54) and the average momentum transferred per event:

$$F(v_z) = \hbar k_L \Gamma \varrho_{ee} = \frac{\hbar k_L \Gamma \Omega_0^2}{\Gamma_B^2 + 4(\delta + k_L v_z)^2}. \quad (2.75)$$

For a given velocity class, the force is clearly greatest when the detuning exactly balances the Doppler shift, and is therefore

$$F_{max} = \frac{\hbar k_L \Gamma \Omega_0^2}{\Gamma_B^2} = \frac{\hbar k_L \Gamma}{2} \frac{I}{I + I_{sat}}, \quad (2.76)$$

increasing to an asymptotic value of $F_{max} = \hbar k_L \Gamma / 2$ at high intensity.

If we want to work out the cooling rate, we need the rate of change of energy, which can be calculated as

$$\begin{aligned} \left(\frac{dE}{dt} \right)_{cool} &= \langle F v_z \rangle = \int_{-\infty}^{\infty} F v_z D(v_z) dv_z \sim F v_{z\delta} D(v_{z\delta}) \Gamma_B \\ &= \frac{\hbar k_L \Gamma \Gamma_B}{2 v_0 \sqrt{\pi}} \left(v_{z\delta} e^{-(v_{z\delta}/v_0)^2} \right) \end{aligned} \quad (2.77)$$

where the integral over the velocity distribution $D(v_z)$ has been approximated by a window of width Γ_B around the population at the resonant velocity $v_{z\delta} = -\delta/k_L$, as only this narrow velocity class interact significantly with the laser. This cooling rate is maximised for $v_{z\delta} = v_0/\sqrt{2}$, where v_0 is the Doppler width as defined in Section 2.2.1.

To cool hot clouds of ions, the laser should be detuned by $1/\sqrt{2}$ of the width of the Doppler broadened line. As the cloud cools, the distribution will narrow, until the Lorentzian becomes the dominant contributor to the width. In this limit, the laser will interact with all of the population simultaneously and the above assumptions do not hold. To proceed we instead Taylor expand the force around zero velocity, giving a superposition of a constant force which displaces the cloud from the centre of the trap and a velocity dependent frictional force:

$$F(v) \sim F(0) + \left. \frac{dF}{dv} \right|_{v_z=0} v_z + \dots = \frac{\hbar k_L \Gamma \Omega_0^2}{\Gamma_B^2 + 4\delta^2} - \frac{8\hbar k_L^2 \Gamma \delta \Omega_0^2}{(\Gamma_B^2 + 4\delta^2)^2} v_z. \quad (2.78)$$

The frictional force is maximized for a detuning of $\delta = \Gamma_B/2$. Evaluating for the cooling rate as in Equation 2.77 then gives

$$\left(\frac{dE}{dt} \right)_{cool} = \frac{\hbar k_L \Gamma \Omega_0^2}{\Gamma_B^2 + 4\delta^2} \langle v_z \rangle - \frac{8\hbar k_L^2 \Gamma \delta \Omega_0^2}{(\Gamma_B^2 + 4\delta^2)^2} \langle v_z^2 \rangle = -\frac{8\hbar k_L^2 \Gamma \delta \Omega_0^2}{(\Gamma_B^2 + 4\delta^2)^2} \left(\frac{k_B T}{m} \right) \quad (2.79)$$

as the average velocity of the harmonic oscillation is zero, and the kinetic energy per mode is $k_B T/2$.

Cooling occurs more rapidly for transitions with shorter lifetimes, as more photons are

scattered per second. The process does not continue without limit however, because the discrete nature of the momentum kicks upon absorption and emission cause a Brownian motion-like diffusion which heats the ion. Although the momentum averages to zero, the square of the momentum (and thus the energy), makes a random walk in phase space over time. Assuming a one-dimensional random walk, where emission occurs either directly towards or away from the laser, the rate of change of kinetic energy due to these kicks is given by

$$\left(\frac{dE}{dt}\right)_{heat} = \frac{1}{2m} \frac{d}{dt} \langle p^2 \rangle = \frac{(\hbar k_L)^2}{2m} \Gamma_{ee} = \frac{(\hbar k_L)^2 \Gamma \Omega_0^2}{2m(\Gamma_B^2 + 4\delta^2)}. \quad (2.80)$$

A second heating process of similar magnitude, due to the fluctuations of the absorption process, also contributes [54,55], and so equilibrium is reached at the Doppler limit

$$T_D = \frac{\hbar(\Gamma_B^2 + 4\delta^2)}{8\delta k_B} = \frac{\hbar\Gamma_B}{2k_B} \quad (2.81)$$

where choosing $\delta = \Gamma_B/2$ also minimizes the final temperature. This is minimised by using an intensity well below saturation, so that $\Gamma_B \sim \Gamma$. The natural linewidth depends on the choice of species and transition - generally, broad E1 transitions are selected to allow rapid cooling down to this limit, and provide $\Gamma \sim 2\pi \times 20$ MHz, leading to $T_D \sim 0.5$ mK. Beyond this, alternative techniques such as sideband cooling (Section 2.3.4) are employed.

The Doppler limit result given in 2.81 is a crude approximation, assuming that all emission is along the z -axis, and considering the effects of only a single cooling laser. In reality, the atom will emit in nearly all directions, with an amplitude in each set by the matrix element for the transition (See Appendix A). This will act to reduce the diffusion upon emission and somewhat decrease the Doppler limit. On the other hand, to stably trap an ion, cooling must be performed in three dimensions and in the Penning trap involves both axial and radial cooling beams. Emitted photons from the radial beam will also have a component of momentum along the z -axis, causing increased axial heating and a higher Doppler limit. The relative contribution of these two effects depends on the intensities of the two beams and the choice of cooling transition.

2.3.2 Relaxation heating

To reach the lowest temperatures, consideration must also be given to the manner in which any Doppler cooling process is completed. Radiation pressure from the cooling laser causes the equilibrium position of the ion to be displaced in the direction of the beam. If the beam is then switched off abruptly, the amplitude of the ion's motion will be altered. Depending upon the phase of the mechanical motion at this instant⁹ the amplitude may be increased or decreased with equal probability. However, because the oscillator energy is proportional

⁹For the purposes of this derivation we consider the ion's motion to be classical, which is a good approximation above the Doppler limit at the trap frequencies used in our experiment.

to the square of the displacement, on average the ion energy increases.

As shown in Equation 2.78, the cooling beam provides a continuous radiation pressure force

$$\mathbf{F}_\gamma = \hbar \mathbf{k}_L \Gamma \varrho_{ee} = \hbar \mathbf{k}_L \gamma_s, \quad (2.82)$$

on the harmonically trapped ion, where γ_s is the spontaneous scattering rate. This force causes a displacement of

$$\mathbf{D}_0 = -\frac{\hbar \mathbf{k}_L \gamma_s}{m \omega_z^2}, \quad (2.83)$$

increasing the static potential energy by

$$U_0 = \frac{1}{2} m \omega_z^2 D_0^2 = \frac{\hbar^2 k_L^2 \gamma_s^2}{2 m \omega_z^2}. \quad (2.84)$$

The potential energy of the oscillation in the displaced frame with motional phase ϕ_0 and amplitude z_0 is, as usual,

$$U_i = \frac{m \omega_z^2 z^2}{2} = \frac{m \omega_z^2 z_0^2}{2} \cos^2(\phi_0). \quad (2.85)$$

If the laser is turned off abruptly, the centre of the ion's oscillation will change instantaneously, with $z \rightarrow z + D_0$. The kinetic energy is initially left unchanged as $E_{k,f} = E_{k,i} = \frac{m \omega_z^2 z_0^2 \sin^2(\phi_0)}{2}$, while the potential energy is shifted, leading to a final energy of

$$\begin{aligned} E_f = E_{k,f} + U_f &= \frac{m \omega_z^2 z_0^2 \sin^2(\phi_0)}{2} + \frac{m \omega_z^2 (z + D_0)^2}{2} \\ &= \frac{m \omega_z^2 z_0^2 \sin^2(\phi_0)}{2} + \frac{m \omega_z^2 z_0^2 \cos^2(\phi_0)}{2} + \frac{m \omega_z^2 D_0^2}{2} + m \omega_z^2 D_0 z_0 \cos(\phi_0) \\ &= E_i + U_0 + \Delta E \cos(\phi_0). \end{aligned} \quad (2.86)$$

The final energy can vary by $\pm \Delta E$, depending on the phase of the mechanical motion at the instant of relaxation and if we take this phase angle to be random we find

$$\int_{E_i + U_0 - \Delta E}^{E_i + U_0 + \Delta E} P(E_f) dE_f = \int_{\pi}^0 P(\phi_0) d\phi_0 = \int \frac{1}{\pi} \frac{d\phi_0}{dE_f} dE_f \quad (2.87)$$

and thus

$$P(E_f) = \frac{1}{\pi \sqrt{(\Delta E)^2 - (E_i + U_0 - E_f)^2}} = \frac{1}{\pi \sqrt{2U_0 E_i - (E_i + U_0 - E_f)^2}}, \quad (2.88)$$

which integrates to unity over E_f for any initial energy E_i . We can therefore use this mapping to calculate the effect of beam turn-off relaxation on a thermally distributed ion

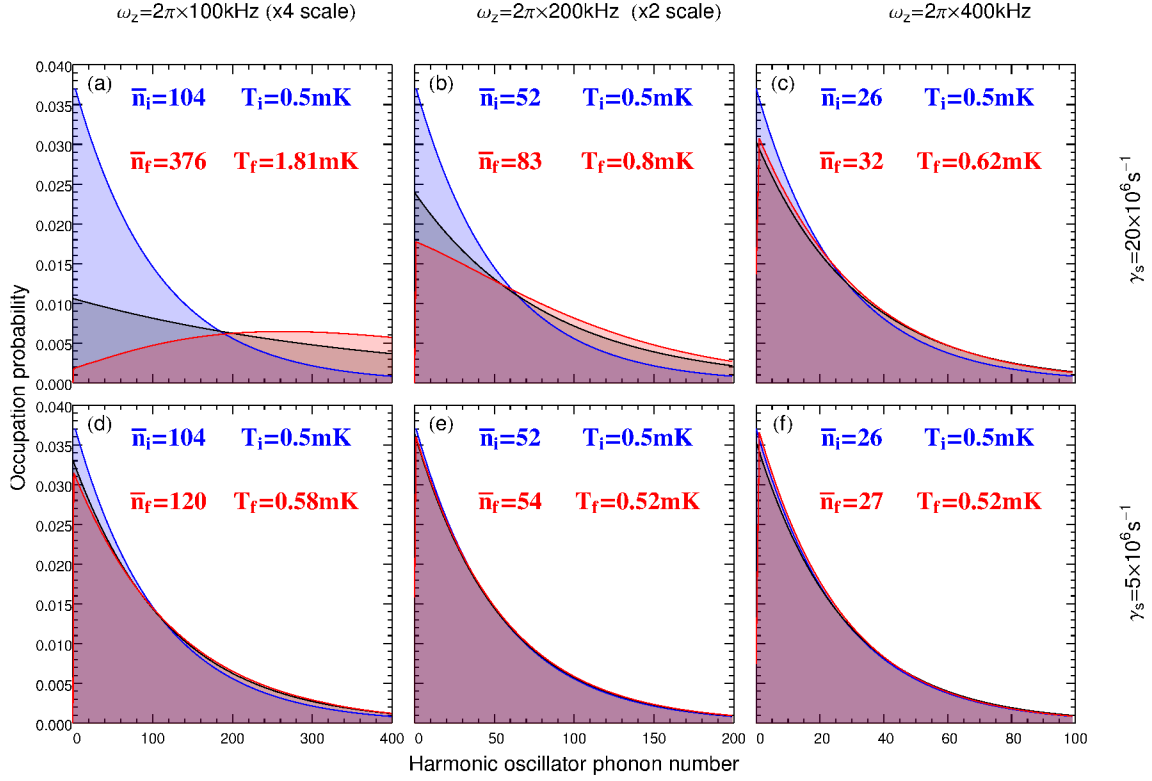


Figure 2.3: Distribution of ion energy in phonon units before and after relaxation heating for a range of scattering rates and trap frequencies. The ion is initially at the Doppler limit, which we assume to be 0.5 mK, shown in blue. After relaxation heating the distribution is flattened and shifted upwards in energy (red), increasing the average phonon number. An equivalent thermal distribution with this new \bar{n} is shown in black. While there is very little change in the distribution at $f_z = 400$ kHz, the effect can be significant at high scattering rates and low trap frequencies.

with temperature T as

$$P(E_f) = \int_0^\infty \frac{e^{-E_i/k_B T} dE_i}{k_B T \pi \sqrt{2U_0 E_i - (E_i + U_0 - E_f)^2}}. \quad (2.89)$$

From Equation 2.84 it is clear the heating is more significant for low trap frequencies, and is minimised by reducing the scattering rate. However, to determine the exact energy distribution that results – which is non-thermal – we must numerically integrate Equation 2.89. Figure 2.3 shows the resulting energy distributions of an ion cooled to the Doppler limit before and after relaxation heating, for a range of scattering rates and trap frequencies.

To put these examples into perspective, consider a typical scattering rate when laser cooling $^{40}\text{Ca}^+$ at 397 nm. When $s_0 = I/I_{\text{sat}} = 1$ on resonance, the saturation parameter at the optimum Doppler cooling frequency will be $s = 0.5$. Assuming that the $D_{3/2} \leftrightarrow P_{1/2}$




Figure 2.4: Motional sideband spectra for the non-thermal distribution produced after relaxation heating (red) and for a thermal distribution of equal \bar{n} . The left hand plot corresponds to panel (b) in Figure 2.3. Despite the distinct difference in energy spectra the sideband spectrum is almost indistinguishable from a thermal spectrum of the same average energy. The right hand plot corresponds to panel (a), in which the relaxation heating effect dominates the final temperature. There is now a clear suppression of the central sidebands, due to the fact that population in the lower Fock states has been significantly reduced. The spectrum begins to take on a bimodal structure commensurate with that associated with a coherent state, as would be expected when the initial distribution of energies is small compared to the relaxation heating.

repump transition is saturated, the total population in the $D_{3/2}$ manifold will be double that in the $P_{1/2}$ states. The ratio of the populations in the three manifolds is thus $N_{S_{1/2}} : N_{P_{1/2}} : N_{D_{3/2}} = 5 : 1 : 2$ and the spontaneous scattering rate is given by

$$\gamma_s = \frac{\Gamma}{8} = \frac{2\pi \times 21 \text{ MHz}}{8} = 16.5 \times 10^6 \text{ s}^{-1}. \quad (2.90)$$

The high and low intensity cases given in 2.3 are therefore just above ($s_0 \sim 1.5$) and well below ($s_0 \sim 0.1$) saturation respectively.

Even when the heating is significant, the sideband spectrum appears fairly similar to the thermal case. Figure 2.4 shows the strengths of the motional sidebands after relaxation heating (red), compared to a thermal distribution of identical \bar{n} (blue). In the extreme case (Figure 2.3(a)) when relaxation heating dominates the final temperature, the central sidebands are suppressed due to the preferential depopulation of the low- n Fock states and the spectrum takes on an increasingly bimodal structure. However in most cases, even when the contribution of relaxation heating is significant (Figure 2.3(b)), the differences from a thermal spectrum are small, making diagnosis of the effect difficult.

In summary, if the lowest Doppler cooling temperatures are required, minimising the effect of relaxation heating is crucial when working at low trap frequencies. This may be achieved either by using very low cooling laser intensities at all times, or by ramping down the intensity immediately before turn-off.

2.3.3 Doppler cooling in Penning traps

In Section 2.3.1, we considered a basic model of Doppler cooling that made several assumptions: that the cooling was performed in the weak binding regime, that the ion's motion was a one-dimensional harmonic oscillator and that emission only occurred along the oscillator axis. In a Penning trap this is a particularly limited approximation, as the radial motion is not simple harmonic in the laboratory frame. Unlike in a 3-dimensional, anisotropic harmonic oscillator (a reasonable approximation to most RF traps) where each mode can be addressed by an appropriately angled beam, the two radial modes of a Penning trap cannot be geometrically decoupled. This is a significant problem as to reduce the amplitude of the magnetron motion we must increase its energy ('blue detune'), while to cool the cyclotron we must reduce its energy as normal ('red detune'); these two requirements cannot be satisfied simultaneously with a conventional Doppler cooling laser. If a laser is set to Doppler cool the cyclotron motion, it will act to increase the radius of the magnetron orbit, limited only by the extent of the beam.

Intensity gradient cooling

One solution is to use a cooling beam that is inhomogeneous at the location of the ion. This can be easily realised by using a Gaussian beam offset from the ion, as shown in Figure 2.5, providing an intensity gradient parameterised by y_0 . We make the assumption that the ion completes many cyclotron orbits over the course of one magnetron orbit ($\omega_+ \gg \omega_-$) and that the laser is red detuned so that frequency selection over the cyclotron orbit causes this mode to be cooled. We arrange the beam so that the intensity is higher on the side of the magnetron orbit where the ion is moving in the same direction as the laser wave vector, increasing the probability of scattering on this side and thus increasing the energy of the magnetron motion on average. In our experiment we use one such beam in the radial plane and one standard beam along the trap axis that provides conventional Doppler cooling. In this manner all three modes can be cooled simultaneously. To calculate the rates and limits due to this process, we must consider a more comprehensive model of the scattering process that takes the inhomogeneity into account. If the limits are to be accurate, we must also account for scattering of photons from the axial beam in the radial plane and vice-versa. A suitable model was first described by Wineland and Itano [55] and has been discussed in detail in the thesis of S. Mavadia [56]. Here I give a brief overview of the key results from these documents, skipping some considerable algebra between each step.

The derivation starts with the equations 2.15, and considers the change in the mode amplitudes (R_z , R_+ , and R_-) after a scattering event with wave vectors at absorption and emission of \mathbf{k} and \mathbf{k}_s respectively. It is assumed that the scattering process acts to change the velocity abruptly as $\Delta\mathbf{v} = \hbar(\mathbf{k} - \mathbf{k}_s)/m$ while the positions remain initially unchanged.




Figure 2.5: A laser beam with Gaussian profile is offset from the centre of the trap so that the ion sits at the point of steepest intensity gradient. The gradient is parameterised by the offset y_0 at the x -axis intercept of the tangent at this point. For a beam of waist w , $y_0 = w/\sqrt{2}$. (After Wineland [55])

This leads to mode amplitude shifts per scattering event of

$$\begin{aligned}\Delta R_z^2 &= \left(\frac{\Delta v_z}{\omega_z} \right)^2 - \frac{2\Delta v_z R_z}{\omega_z} \sin(\omega_z t + \theta_z) \\ \Delta R_-^2 &= \frac{\Delta v_x^2 + \Delta v_y^2}{4\omega_1^2} + \frac{R_-}{\omega_1} [\Delta v_x \sin(\omega_- t + \theta_-) + \Delta v_y \cos(\omega_- t + \theta_-)] \\ \Delta R_+^2 &= \frac{\Delta v_x^2 + \Delta v_y^2}{4\omega_1^2} - \frac{R_+}{\omega_1} [\Delta v_x \sin(\omega_+ t + \theta_+) + \Delta v_y \cos(\omega_+ t + \theta_+)],\end{aligned}\tag{2.91}$$

where θ_z , θ_+ and θ_- are the phases of the motion at the time of scattering.

We can then calculate the rate at which the mode amplitudes change by multiplying these values by the number of scattering events per second:

$$\frac{dR_i^2}{dt} = \frac{I(r)}{\hbar\omega_L} \sigma(\omega_L, \mathbf{v}) \Delta R_i^2\tag{2.92}$$

where the $I(y)$ is the beam intensity and $\sigma(\omega_L, \mathbf{v})$ is the absorption cross-section. For the radial beam, $I(r) \sim I_1 \frac{y}{y_0}$ while for the axial beam $I(r) = I_2$. Averaging uniformly over the phases of the motion θ^{10} and the emitted photon wave vector distribution \mathbf{k}_s , we find

¹⁰The weighting towards absorption at certain velocities or orbital phases is encapsulated by the functions $I(r)$ and $\sigma(\omega_L, \mathbf{v})$; beyond these factors scattering is a randomly occurring event.

average cooling rates of

$$\frac{d\langle R_z^2 \rangle}{dt} = \gamma_{s2} \left\{ \frac{\hbar^2 k_L^2 \left[1 + (1 + \gamma_{s1}/\gamma_{s2}) f_{sz} \right]}{m^2 \omega_z^2} - \frac{8\hbar k_L^2 \delta \langle R_z^2 \rangle}{m(\Gamma^2 + 4\delta^2)} \right\} \quad (2.93a)$$

$$\frac{d\langle R_-^2 \rangle}{dt} = \gamma_{s1} \left\{ \frac{\hbar^2 k_L^2 \left[1 + (1 + \gamma_{s1}/\gamma_{s2}) (f_{sx} + f_{sy}) \right]}{4m^2 \omega_1^2} + \frac{4\hbar k_L^2 \delta \omega_- \langle R_-^2 \rangle}{m\omega_1(\Gamma^2 + 4\delta^2)} - \frac{\hbar k_L \langle R_-^2 \rangle}{2m\omega_1 y_0} \right\} \quad (2.93b)$$

$$\frac{d\langle R_+^2 \rangle}{dt} = \gamma_{s1} \left\{ \frac{\hbar^2 k_L^2 \left[1 + (1 + \gamma_{s1}/\gamma_{s2}) (f_{sx} + f_{sy}) \right]}{4m^2 \omega_1^2} - \frac{4\hbar k_L^2 \delta \omega_+ \langle R_+^2 \rangle}{m\omega_1(\Gamma^2 + 4\delta^2)} + \frac{\hbar k_L \langle R_+^2 \rangle}{2m\omega_1 y_0} \right\} \quad (2.93c)$$

where γ_{si} , $i = 1, 2$ is the scattering rate due to the radial and axial beams respectively and $f_{si} = \int P_s(\hat{\mathbf{k}}_s) \hat{\mathbf{k}}_s^2 d\Omega$ is a spatial factor associated with the fraction of energy emitted into each direction i for emission pattern $P_s(\hat{\mathbf{k}}_s)$. The first term in each expression represents the diffusion upon emission and heats all three modes. The second term represents the frequency-selective Doppler cooling, which cools the axial and cyclotron modes but heats the magnetron. The third term in the radial equations represents the effect of the intensity gradient, which acts to heat the cyclotron while cooling the magnetron. Note the cyclotron and magnetron dynamics are identical up to a factor of ω_{\pm} and the polarity of the second and third terms.

For the $\Delta m = \pm 1$ Doppler cooling transitions used in our experiment, the spatial factors are $f_z = 0.4$, $f_x = f_y = 0.3$, and we set $\delta = \Gamma/2$ and $\gamma_{s1} = \gamma_{s2}$ to reach to lowest overall energies. To calculate the limiting temperatures, these rate equations are set to zero, and solved for $\langle R_i^2 \rangle$, before using $T = 2E_k/k_B$ and Equation 2.18 to convert to temperature. We find:¹¹

$$\begin{aligned} T_z &= \frac{9\hbar\Gamma}{20k_B} \\ T_- &= \frac{11\hbar\Gamma\omega_-}{20\omega_1 k_B \left(\frac{\Gamma}{2k_L y_0 \omega_-} - 1 \right)} \\ T_+ &= \frac{11\hbar\Gamma\omega_+}{20\omega_1 k_B \left(1 - \frac{\Gamma}{2k_L y_0 \omega_+} \right)}. \end{aligned} \quad (2.94)$$

Note that axial Doppler temperature is slightly lower than that given by the basic derivation

¹¹Beware that the original derivation of this result by Wineland & Itano [55] contains a typographic mistake in equation 61c, which implies a factor of two increase in T_+ . The same mistake also appears in the version given in the thesis of S. Mavadia [56], where 3.72c, 3.74c, 3.75c and 3.76c are erroneous. The results stated here are the correct limits.

in Equation 2.81, due to accounting for the the emission pattern and diffusion from both beams correctly. The relative temperatures of the radial modes depend on the intensity gradient which acts to cool the magnetron, but heats the cyclotron mode. Cooling both modes simultaneously is only possible while the following condition is satisfied:

$$\omega_- < \frac{\Gamma}{2k_L y_0} < \omega_+. \quad (2.95)$$

The higher the axial frequency, the narrower this range (as at the stability limit, $\omega_+ = \omega_-$) and, in practice, the less stable the ion becomes to perturbations such as background gas collisions. Unlike in the axial case, the equilibrium temperatures also depend on the intensity gradient and trap frequency. The magnetron temperature, in the right conditions, can be far lower than the Doppler limit, with $T_- \rightarrow 0$ as $\omega_- \rightarrow 0$. However, as the trap frequency increases the minimum intensity gradient required to cool this mode becomes ever steeper and the temperature reached with a beam of fixed gradient rapidly rises. The cyclotron temperature is less sensitive to the intensity gradient, but note that even in the ‘best case’ for this mode (at $\omega_1 = \omega_c/2$ and $y_0 \rightarrow \infty$) the final temperature is $T_+ > 1.1$ mK, over a factor of two higher than the axial Doppler limit.

Axialisation

While the use of an inhomogeneous radial beam to cool the magnetron motion is a conceptually elegant approach and is readily achievable using nothing more than standard Doppler cooling lasers, the range of trap frequencies that support stable cooling is severely limited. An alternative approach is to avoid cooling the magnetron directly, instead coupling it dynamically to another mode of motion that can be effectively cooled, indirectly cooling the magnetron motion in the process. Typically this involves the use of an additional radiofrequency field applied to the trap electrodes, to couple the magnetron to either the cyclotron or axial modes. Coupling to the axial mode involves a somewhat more complicated field geometry¹² unsuited to our trap, so we take the radial approach, known as azimuthal quadrupolar excitation or ‘axialisation’ [57]. Axialisation has been studied in some depth by a number of groups including our own and my discussion here follows that of R. Hendricks [45, 58], where I will again skip any algebra unimportant to the development of the main ideas.

Axialisation applies a quadrupole RF field of amplitude V_{ax} in the radial plane at the true cyclotron frequency ω_c , with potential

$$\phi_{ax} = \frac{V_{ax}}{2d_0^2} (x^2 - y^2) \cos(\omega_c t). \quad (2.96)$$

Because $\omega_c = \omega_+ + \omega_-$, this field acts to couple the levels $\langle n+1, l | \phi_{ax} | n, l+1 \rangle$, where I

¹²A tilted quadrupole RF field [46].




Figure 2.6: Radial trajectory in the presence of axialisation, in the laboratory frame.

have used the conventions of Equation 2.36.

The geometric aspect is best understood in a rotating frame. In the polar coordinates in the frame rotating at ω_r such that $\theta' = \theta - \omega_r t$ and $r' = r$, the axialisation potential becomes

$$\phi'_{ax} = \frac{V_{ax} r'^2}{4d_0^2} [\cos(2\theta' + (2\omega_r - \omega_c)t) + \cos(2\theta' + (2\omega_r + \omega_c)t)]. \quad (2.97)$$

Choosing the frame described in Section 2.1.1 where $\omega_r = \omega_c$, we find that one term oscillates at twice the cyclotron frequency and the other is static. Because the two modes are degenerate in frequency in this rotating frame, the static field is effectively resonant with the transition while the other term is far off-resonant and can be ignored via a rotating wave approximation. The axialisation potential in the rotating frame is then

$$\phi'_{ax} = \frac{V_{ax}}{4d_0^2} (x'^2 - y'^2), \quad (2.98)$$

where x' , y' are the Cartesian coordinates in the rotating frame. The effect of axialisation is thus to break the cylindrical symmetry of the potential in this frame, ‘squeezing’ it in one direction while relaxing it in the other. Because the degeneracy of the system has been broken, left- and right-handed circular motions are no longer eigenmodes of the system and the system energy will oscillate between the magnetron and cyclotron degrees of freedom. This is analogous to the action of a birefringent crystal on circularly polarised light, as in an optical wave-plate. A typical trajectory in the presence of axialisation, in the laboratory frame, is given in Figure 2.6. Under axialisation, the average radii of the two modes become equal (i.e. $\langle R_+ \rangle = \langle R_- \rangle$). In steady state, energy will continue to oscillate between the two modes - the amplitude of this oscillation is determined by the phase of the axialisation drive relative to the two modes - if the mode phases are symmetric around that of the drive, the resultant linear oscillation is aligned with the axis of the quadrupole and the system is stable. This is equivalent to sending light into a wave-plate with linear

polarisation that is aligned to either the fast or slow axis of the crystal. With the two modes coupled, if one experiences an external damping force, such as laser cooling, the other will cool with it.

The treatment in [45] describes cooling in the presence of axialisation and a radial intensity gradient, but no diffusive heating is included in the model which prevents a cooling limit from being calculated. To properly account for the final stages of cooling, one would expect to follow a similar derivation as given in Section 2.3.3 for the intensity gradient cooling method, starting with the equations of motion and calculating the change to the mode amplitudes after each scattering event. Axialisation modifies the radial equations of motion (2.8) to give

$$\begin{aligned}\ddot{x} &= \frac{1}{2}\omega_z^2 x + \omega_c \dot{y} - \frac{eV_{ax}}{md_0^2} \cos(\omega_c t) x \\ \ddot{y} &= \frac{1}{2}\omega_z^2 y - \omega_c \dot{x} + \frac{eV_{ax}}{md_0^2} \cos(\omega_c t) y.\end{aligned}\tag{2.99}$$

Unfortunately, unlike Equations 2.8, these cannot be easily solved to give the analytic trajectories required to proceed with deriving the rate equation. However the equations can be numerically integrated, and if a stochastic scattering process is included in the numeric model it is straightforward to simulate the effects of Doppler cooling on the coupled system. Simulations of this sort have been described in the MSci thesis of H. Nejad [59] and show that most effective approach is to combine axialisation with an offset radial beam. If the beam is not offset, cooling of both modes is still possible but the necessary axialisation voltage is significantly increased and the final temperature is slightly higher.

Axialisation somewhat spoils the purity of the Penning trap, as an oscillating field is introduced to an environment that has until that point been completely static. This will introduce RF micromotion, but because the axialisation field is so weak compared to that used for providing the trapping potential in RF traps, the deleterious effects are much diminished. These drawbacks are generally far outweighed by the fact that the technique works efficiently at all trap frequencies, allowing trapping in regimes not accessible when using intensity gradient methods alone.

2.3.4 Sideband cooling

If the ion provides a narrow enough transition to realise the strong-binding regime, much lower temperatures can be reached ($\omega_z > \Gamma$) using a cooling laser with narrow enough linewidth to selectively drive individual sidebands.¹³ Sideband cooling uses frequency selection to ensure that absorption occurs predominantly on a red sideband, while spontaneous emission occurs over a range of sidebands with branching ratios determined by the effective Lamb-Dicke parameter. Each scattering process thus reduces the energy by an

¹³Narrow linewidth processes can also be engineered in species lacking a suitable transition by using Raman lasers to couple pairs of hyperfine ground state sub-levels

average of one phonon and the cooling proceeds until the ion approaches the ground state of motion. The limiting value of the ion energy is determined by the effect of off-resonant interactions with the carrier and blue sidebands which provide a heating rate that becomes significant as $n \rightarrow 0$ and the red sideband Rabi frequency becomes small.

Sideband cooling limit

We now consider the dynamics of the cooling process of an ion trapped with axial frequency ω_z as it approaches a steady state at low energies. We consider a two level system driven by a narrow linewidth laser. The rate of scattering is the product of the decay rate and the sum of the populations excited via each sideband,

$$R = \Gamma \varrho_{ee} = \sum_{m=-\infty}^{\infty} \frac{\Gamma |\Omega_{n-m,n}|^2}{4\delta_m^2 + 2|\Omega_{n-m,n}|^2 + \Gamma^2}, \quad (2.100)$$

where $\delta_m = \omega_L - \omega_0 - m\omega_z$ is the detuning from the m^{th} sideband. We assume that the sidebands are well resolved ($\Omega_0, \Gamma \ll \omega_z$) and that $\bar{n} \ll 1$ so that the average occupation number is equal to the population of the first excited harmonic oscillator level, $\bar{n} \sim n_1$ and the sideband strengths are well described by the Lamb-Dicke regime approximations (Equation 2.74). When the laser is tuned to the first red sideband, the scattering rate becomes

$$R = \frac{\Gamma \Omega_0^2 \eta^2 n_1}{2\Omega_0^2 \eta^2 n_1 + \Gamma^2} + \frac{\Gamma \Omega_0^2}{4\omega_z^2} + \frac{\Gamma \Omega_0^2 \eta^2 n_0}{16\omega_z^2}, \quad (2.101)$$

where the three terms correspond to absorption on the first red sideband, carrier and first blue sideband respectively.¹⁴

The spontaneous decay can also change the motional state, but we have no control over this process which depends only on the effective Lamb-Dicke parameter. Expanding the scattering rate in terms of decays on the first blue, carrier and first red sidebands we find

$$R \sim \frac{\Gamma \Omega_0^2 \eta^2 n_1}{2\Omega_0^2 \eta^2 n_1 + \Gamma^2} + \frac{\Gamma \Omega_0^2 \tilde{\eta}^2 n_1}{4\omega_z^2} + \frac{\Gamma \Omega_0^2}{4\omega_z^2} + \frac{\Gamma \Omega_0^2 \tilde{\eta}^2 n_0}{4\omega_z^2} + \frac{\Gamma \Omega_0^2 \eta^2 n_0}{16\omega_z^2} + O(\eta^4), \quad (2.102)$$

where, for reasons that will become apparent, we have allowed for different Lamb-Dicke parameters for decay ($\tilde{\eta}$) and absorption.

Overall, to $O(\eta^2)$, we have two scattering processes that lead to the loss of one phonon (absorption on the red sideband with decay on the carrier, and absorption on the carrier with decay on the blue sideband), two processes that lead to the gain of a phonon (absorption on the blue sideband with decay on the carrier and absorption on the carrier with decay on the red sideband) and one process that does not change the phonon number (absorption and decay on the carrier). The resonant cooling process far outweighs the off-

¹⁴Absorption on the second red sideband occurs at a rate that both decreases with \bar{n} and is proportional to η^4 and can therefore be ignored.

resonant cooling process (typically $\omega_z/\Gamma \sim 10$), which is therefore ignored. Also note that the heating processes scale with Ω_0^2 , while the cooling rate starts to saturate as the Rabi frequency approaches the decay rate and cooling is thus most effective when $\Omega_0\eta < \Gamma$.

The rate equations for the two participating oscillator levels are therefore

$$\begin{aligned}\dot{n}_0(t) &= \frac{\Omega_0^2\eta^2}{\Gamma}n_1(t) - \frac{\Gamma\Omega_0^2(4\tilde{\eta}^2 + \eta^2)}{16\omega_z^2}n_0(t) \\ \dot{n}_1(t) &= -\dot{n}_0(t).\end{aligned}\tag{2.103}$$

Solving these coupled equations gives an exponential cooling process with rate $W = \Omega^2\eta^2/\Gamma$ that reaches a steady state phonon number of

$$\bar{n}_{limit} = n_1(t \rightarrow \infty) = \left(\frac{\Gamma}{2\omega_z}\right)^2 \left[\left(\frac{\tilde{\eta}}{\eta}\right)^2 + \frac{1}{4}\right].\tag{2.104}$$

As in the Doppler cooling case, there exists a trade-off between cooling rate and final temperature, which both increase with the linewidth Γ . In principle, it is possible to reach the strong binding regime using a dipole transition, but to cool significantly below the Doppler limit requires impractically high trap frequencies. Electric quadrupole (E2) transitions provide a much narrower alternative – in $^{40}\text{Ca}^+$ the 729-nm transition has a linewidth of $\Gamma \sim 1\text{ s}^{-1}$ – that allow cooling to the ground state even in low frequency traps. However with such a low linewidth only a single phonon can be scattered per second, much lower than the background heating rate in most traps.

The solution is to use a ‘quench’ laser to weakly couple the upper level to a third, ancillary level that can decay rapidly back to the ground state. In this experiment we use the ‘cascade’ configuration described by Marzoli [60] and shown in Figure 2.7. Providing the ancillary level has a low lifetime and the quench laser remains far below saturation, the ancillary level may be adiabatically eliminated. This leaves a quasi-two-level system, with an increased linewidth due to the alternative decay route for excited population. The exact linewidth of the upper level may be controlled by varying the intensity and detuning of the quench laser. For an ancillary level with decay rates to the ground state and upper level of Γ_{ag} and Γ_{ae} respectively, coupled by a quench laser with Rabi frequency Ω_q and detuning δ_q , the effective linewidth is

$$\tilde{\Gamma} = \frac{\Omega_q^2}{(\Gamma_{ag} + \Gamma_{ae})^2 + 4\delta_q^2}\Gamma_{ae}.\tag{2.105}$$

The linewidth may thus be tuned to provide the best performance for the trap and experiment in question, or even altered during the cooling process to optimise the initial rate and final temperature.

As emission occurs at ω_{ag} , it is this transition rather than the main cooling transition that




Figure 2.7: ‘Cascade’ configuration for broadening a metastable level $|e\rangle$ via weak coupling to a short-lived ancillary state $|a\rangle$.

determines the Lamb-Dicke parameter $\tilde{\eta}$. However, the decay occurs into three dimensions, decreasing the chance of an axial phonon number change occurring and effectively reducing the value of $\tilde{\eta}$. Assuming isotropic emission, $\tilde{\eta} = \frac{z_0 k_{aq}}{\sqrt{3}}$. At the same time, the absorption process driven by the quench laser can also lead to a change in n , which might properly be included in this treatment. However, the quench detuning δ_q is typically much larger than the trap frequency ω_z , which suppresses absorption on either sideband by a factor of $O(\eta^2)$, meaning this process can usually be ignored.

Sideband cooling beyond the Lamb-Dicke regime

Sideband cooling is generally employed once the ion has already been cooled to the Lamb-Dicke regime, as detailed in the treatment above. At higher phonon numbers decay on high order sidebands becomes more likely, providing a diffusive heating process that makes cooling on the first red sideband inefficient. Furthermore, if the trap frequency is low and the initial phonon number is high, cooling on the first red sideband can lead to significant population trapping in high Fock states. For an ion in a Fock state, the coupling strength on each sideband is given by Equation 2.69. Plotting this as a function of the initial state n (Figure 2.8) the Rabi frequency of each sideband repeatedly passes through zero – at certain values of n the strength of the interaction becomes vanishingly weak. For $^{40}\text{Ca}^+$ trapped at a frequency of $\omega_z = 2\pi \times 200$ kHz, $\eta = 0.217$ and the first red sideband of the 729-nm has near-zero amplitude at $n = 78$. When sideband cooling on the first red sideband, any population that starts in states $n \geq 78$ will be gradually cooled towards this level, at which point the laser ceases to interact with the ion and cooling stops, leaving the population trapped in this high- n state. Any population initially below this ‘dark’ state will be cooled to the ground state as intended. Because the first red sideband disappears both for the population in $n = 0$ and for that trapped at $n = 78$, this could make the ion appear – to a




Figure 2.8: Normalised coupling strength versus phonon number, for carrier and first, second and third red sidebands of $^{40}\text{Ca}^+$ ion trapped at $\omega_z = 2\pi \times 200$ kHz and probed with light at 729 nm. The coupling strength (specifically, the square of the Rabi frequency given in 2.68) for each sideband is zero for certain values of n . Reproduced from [61] with permission.

limited measurement – to be in the ground state while significant population remains at relatively high energy. The Doppler limit for a 200 kHz trap is $\bar{n} = 47$. A thermal distribution of this energy has $\sim 20\%$ of its population with $n \geq 78$, limiting the final average energy to $\bar{n} \geq 15$. This becomes a severe problem at lower trap frequencies, as the Lamb-Dicke parameter increases, reducing the value of n at which the first dark state occurs. This is exacerbated by the fact that \bar{n} at the Doppler limit is also higher. For a trap frequency of $\omega_z = 2\pi \times 100$ kHz, $\eta = 0.3$ and $\bar{n}_{Dopp} \sim 94$. The 65% of the population initially above $n = 40$ will become trapped in this state, limiting the final energy to $\bar{n} > 26$. Figure 2.9 shows a simulation, produced by F. Beck [61], of population trapping for an ion initialized in a Fock state and cooled on each of the first three red sidebands.

Efficient cooling from outside the Lamb-Dicke regime is certainly possible (providing the Lamb-Dicke parameter is reasonably small¹⁵) and has recently been demonstrated in an RF trap [64]. However the protocol is slightly more complicated as cooling must proceed in multiple stages - initially cooling on a higher order red sideband before progressively stepping towards the first red sideband to complete the sequence. This both increases the cooling rate, as more phonons are removed per absorption, and prevents population being trapped in high- n levels. For extreme situations where population is spread above multiple dark states of several sidebands, more complex pulse sequences would have to be used. An in-depth discussion of this topic is given in reference [61].

¹⁵For large intrinsic Lamb-Dicke parameters ($\eta > 1$), the cooling process becomes very inefficient and it will generally be impossible to reach the ground state [62, 63]

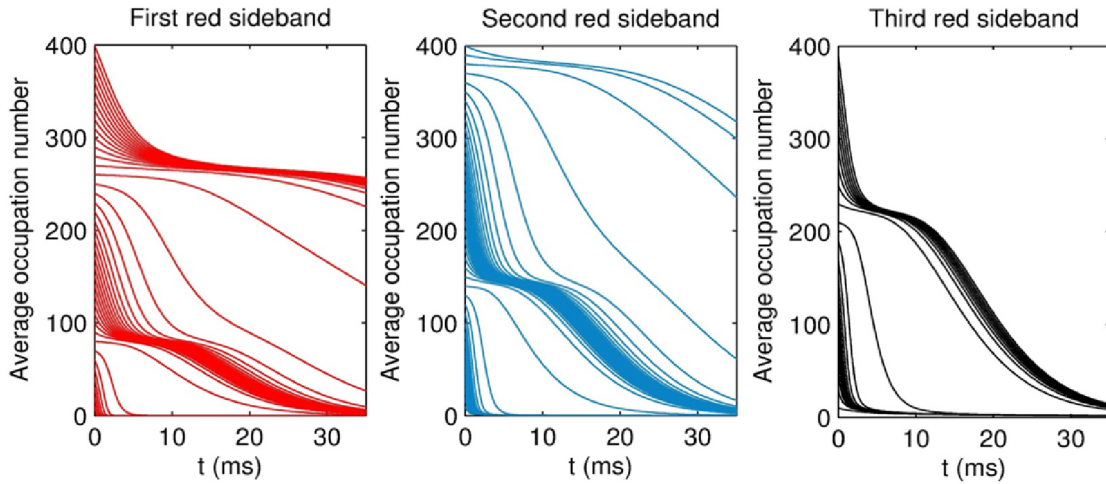


Figure 2.9: Average phonon number versus cooling time for ions initially in Fock states $n = \bar{n}$ and cooled on the first, second and third red sidebands. Population that is initialised in a higher energy state than the lowest ‘dark’ state for the sideband becomes trapped in the dark state. However, the dark states for different sidebands lie at different n , allowing composite cooling sequences to be designed that alternate between sidebands to drive all population to the ground state. Note that the population is not trapped indefinitely in the dark states as a slow decay occurs due to off-resonant coupling with other sidebands. Reproduced from [61] with permission.

Sideband cooling the radial motion

Sideband cooling presents an attractive alternative to the variously imperfect methods of cooling the magnetron motion discussed in Section 2.3.3. Because a particular sideband can be targeted, it should be possible to choose a transition that simultaneously removes energy from the cyclotron mode while adding energy to the magnetron, moving both towards their motional ground states.

At low trap frequency and magnetron temperature, the comb of magnetron sidebands associated with each cyclotron sideband is distinct from other combs and the process should be relatively straightforward. However, if the frequencies or temperatures are higher and the combs overlap, it will be difficult if not impossible to drive a specific sideband (see Figure 2.10).

The benefits of reducing the trap frequency are somewhat offset by the great increase in the oscillator quantum number at the Doppler limit, and in any case both radial modes will remain far beyond the Lamb-Dicke regime after Doppler cooling. If the radial motion could be pre-cooled to Lamb-Dicke confinement, then cooling to the ground state would be no harder than when cooling two modes of a conventional harmonic trap with a single laser as demonstrated in, e.g. [65]. Unfortunately this is impossible to achieve at the trap frequencies available in this experiment and cooling will therefore be slow and will need to proceed via a number of steps.




Figure 2.10: Schematic representation of radial spectra at two different trap frequencies. At 169 kHz the magnetron combs from neighbouring cyclotron sidebands are distinct, allowing a single sideband to be targeting for cooling without excessive off-resonant heating. At 236 kHz the combs overlap, making this kind of discrimination much harder. Reproduced from [61] with permission.

3

Experiment Overview

This chapter describes the current experiment, in principle and practice. I begin the chapter with a review of the basic aims and approach, highlighting the particularities of applying the methods to calcium ions. I then give an overview of the apparatus at the largest scale, to introduce the roles of the various sub-systems and show how they are combined to form the overall experiment. The final four sections of the chapter cover each of these sub-systems in turn: the magnet, the trap and its associated systems; the Doppler cooling lasers and locks; the narrow-linewidth 729-nm laser and locks; and finally the experimental control systems, centred on an FPGA and desktop interface.

The relative emphasis given to these different sections does not reflect their relative importance or complexity. Indeed, certain critical systems will be only briefly mentioned as they have been covered in detail in previous group theses, mainly those of S. Mavadia [56] and S. Donnellan [66]. A great deal of the experimental work during this thesis has focussed on the development of a wide range of electronic circuits and electro-optical systems. Many of these systems are also provided with local logical control or computer interface via a number of dedicated Arduino microcontrollers. Certain circuits are quite readable from a simple schematic and these are included in Appendix D. Others – in particular those that involve an Arduino or have a number of discrete, interconnected boards – are less easily understood out of context. To avoid confusion I have refrained from including the schematics for such circuits or the associated controller code, but all are available on request from the Imperial ion trap group.

3.1 Introduction & Aims

The aim of the work described in this thesis was to apply the techniques of sideband cooling to a single ion in a Penning trap. On one hand, this has required a number of specialised systems to be built and developed with this express aim. On the other hand, another major strand of work has necessarily focused on the improvement of our Doppler cooling techniques, as effective and reliable pre-cooling is essential to successful sideband cooling, especially when working at low trapping frequencies.

The trap is situated in the centre of a superconducting solenoid, capable of a field of approximately 2.5 T but currently running at 1.85 T. Calcium ions are trapped – anything from a single ion to a cloud of several thousand in number – and exposed to two beams of laser light, incident on the trap in directions nominally parallel with and perpendicular to the field (the ‘axial’ and ‘radial’ beams respectively). For appropriate choice of parameters, these two beams can be used to Doppler cool both the axial and radial modes of the ion’s motion. Once the ion is close to the Doppler limit, another very narrow-linewidth laser can be used to sideband cool the ion towards the motional quantum ground state. This laser can, in principle, be directed in either the axial or radial directions (or both), but cooling has only been attempted in the axial direction to date.

3.2 Calcium Electronic Structure

$^{40}\text{Ca}^+$ lends itself well to laser cooling experiments. It has dipole transitions suitable for Doppler cooling on the edge of the visible spectrum, that do not require specialist optics or detectors. Some repumping is necessary to prevent population trapping, but is easily achieved with lasers in the near-IR. It also has a very narrow electric quadrupole transition on the edge of the visible spectrum that is well suited to sideband cooling and qubit manipulation. Crucially, all the necessary laser frequencies are readily accessible with diode lasers, providing a cost-effective and relatively low maintenance solution to running an experiment that requires a large number of transitions to be addressed simultaneously.

The electronic structure of $^{40}\text{Ca}^+$ is shown in Figure 3.1, along with the transition wavelengths and decay rates¹ relevant to the experiments described in this thesis. Doppler cooling is performed on the 397-nm transition coupling the $S_{1/2}$ and $P_{1/2}$ manifolds. The $P_{1/2}$ states have a natural linewidth of 22.5 MHz, and can also decay at 866 nm to the $D_{3/2}$ manifold with a branching ratio of 0.064. The $D_{3/2}$ states are metastable, with a lifetime of ~ 1 s, and must be repumped at 866 nm² to avoid population becoming trapped here and lost from the Doppler cooling cycle.

The 729-nm electric quadrupole transition between the ground state and the metastable

¹Lifetimes taken from [67–69].

²Alternatively the $D_{3/2}$ states may be repumped at 850 nm via the $P_{3/2}$ levels, but we do not use this method in the current experiment.




Figure 3.1: Low-energy electronic structure of a $^{40}\text{Ca}^+$ ion. Transitions addressed in this experiment are given by solid lines, labelled with the transition wavelength and the decay rate of the upper level at that wavelength. Other strong E1 transitions that are not addressed in this experiment are marked by wiggly lines. The Zeeman splittings of each manifold are also given.

$D_{5/2}$ levels provides direct access to the strong-binding regime, as the $D_{5/2}$ state has a natural linewidth of 0.14 Hz. These two levels can also potentially be used as a qubit, in which case a 729-nm laser can be used to manipulate the qubit state. To empty any population shelved in the upper manifold an 854-nm laser can be used to drive population into the $P_{3/2}$ state, which then decays predominantly ($BR = 0.935$) directly to the ground state with a wavelength of 393 nm. This laser can also be applied at low Rabi frequency to provide a weak ‘quench’ coupling and artificially broaden the linewidth of the $D_{5/2}$ states when sideband cooling (see Section 2.3.4).

3.2.1 Effects of magnetic field

The picture is significantly complicated by the strong magnetic fields used in the Penning trap, which lead to Zeeman splittings of the energy levels. At 1.85 T, the Zeeman splittings are about two orders of magnitude smaller than the spin-orbit coupling and can be treated as a perturbation. The Hamiltonian interaction term of an atom in magnetic field is given by

$$\hat{H}_{mag} = \hat{\mu} \cdot \mathbf{B}, \quad (3.1)$$

where $\hat{\mu} = \mu_B(g_L\hat{L} + g_S\hat{S}) \cdot \mathbf{B}/\hbar$ is the magnetic moment of the atom. By perturbation theory the first and second order energy shifts of an electronic level n are given by

$$E_n^{(1)} \sim \frac{\mu_B B_z}{\hbar} \left\langle n^{(0)} \left| (\hat{L}_z + 2\hat{S}_z) \right| n^{(0)} \right\rangle \sim \mu_B g_j B_z m_j \quad (3.2a)$$

$$E_n^{(2)} \sim \sum_{m \neq n} \frac{\mu_B^2 B_z^2 \left| \left\langle m^{(0)} \left| (\hat{L}_z + 2\hat{S}_z) \right| n^{(0)} \right\rangle \right|^2}{\hbar^2 (E_n^{(0)} - E_m^{(0)})}, \quad (3.2b)$$

where μ_B is the Bohr magneton and g_j is the Landé g-factor. The calculation of the matrix elements in the second order shift are straightforward via a Clebsch-Gordan expansion of the eigenstates and is given in, e.g. [70]. The effects of these two shifts on the calcium energy levels at 1.85 T are summarised in Table 3.1.

First order Zeeman effect

The effect of the familiar first-order shift is to lift the degeneracy of the m_j sub-levels of each manifold, producing an even comb of states with splittings linear in B . Because the splittings (of the order 10 GHz) are far bigger than the linewidth of the transitions driven by the 397-nm, 866-nm and 854-nm lasers, each sub-level must be driven by a dedicated laser frequency. Until recently this meant a large number of laser sources were necessary to cool ions like calcium in high fields. However, with the bandwidth of electro-optical modulator technology now in the 10’s of GHz, it is possible to produce the entire comb from a single seed frequency, significantly simplifying experiments such as ours.

Manifold	m_j sub-level	1 st Order Shift	2 nd Order Shift
		$E^{(1)}/h$ (GHz)	$E^{(2)}/h$ (MHz)
$S_{1/2}$	$\pm \frac{1}{2}$	± 25.92	0
$P_{1/2}$	$\pm \frac{1}{2}$	± 8.62	-22
$P_{3/2}$	$\pm \frac{3}{2}$	± 51.82	0
	$\pm \frac{1}{2}$	± 17.27	+22
$D_{3/2}$	$\pm \frac{3}{2}$	± 31.05	-58
	$\pm \frac{1}{2}$	± 10.35	-87
$D_{5/2}$	$\pm \frac{5}{2}$	± 77.71	0
	$\pm \frac{3}{2}$	± 46.63	+58
	$\pm \frac{1}{2}$	± 15.54	+87

Table 3.1: First and second order Zeeman shifts of $^{40}\text{Ca}^+$ levels at $B = 1.85$ T.

While there is an overhead in the complexity of the laser or modulator systems needed to drive such split transitions, the large Zeeman splitting provides one distinct advantage over RF traps by permitting resonant frequency addressing for state preparation or readout. In principle, these operations can then be performed with 100% fidelity, as off-resonant coupling to other levels in the manifold is vanishingly small. By contrast, state selection in RF traps must be achieved via the polarisation of the beams and achieving high fidelities on these operations is much harder. The polarisation purity requirements of experiments in Penning traps are often quite mild.³

Second order Zeeman effect

At 1.85 T, the second order Zeeman shifts are approximately three orders of magnitude smaller than the first order splittings, but because they are equal to or greater than the natural linewidth they remain an important consideration. The second order effect occurs via the interaction of sub-levels of identical L and m_j , but different j (e.g. $D_{3/2, +3/2}$ and $D_{5/2, +3/2}$), and has the effect of shifting those pairs of levels away from one another symmetrically. The shift on the sub-levels of $D_{3/2}$ and $P_{1/2}$ is always down in energy, while the $D_{5/2}$ and $P_{3/2}$ sub-levels shift up in energy. The resultant shifts in the transition frequencies are neither constant, nor directly proportional to the m_j value of the levels involved, making it impossible to perfectly address all transitions simultaneously with a comb of evenly spaced frequencies such as that produced by a modulator. The potential ramifications of this constraint are discussed in more detail in Section 3.5.

³In some cases, such as the phase gates described in Chapter 5, beam polarisations remain very important.

j-mixing

The transition from $P_{1/2} \rightarrow D_{5/2}$ is strongly forbidden⁴ and under normal conditions the upper state can be considered to be stable against this decay. However, at high fields the same interaction that causes the second-order Zeeman shift also leads to a perturbation of the eigenstates of the system, with the first order correction given by

$$|n^{(1)}\rangle = \sum_{m \neq n} \frac{\mu_B B_z \langle m^{(0)} | (\hat{L}_z + 2\hat{S}_z) | n^{(0)} \rangle}{\hbar (E_n^{(0)} - E_m^{(0)})} |m^{(0)}\rangle. \quad (3.3)$$

As the matrix elements involved are identical to those in Equation 3.2b, the same pairs of states (with equal m_j) that interact to produce the second order Zeeman shift are now mixed, so the $P_{1/2}$ levels mix with their partners in the $P_{3/2}$ manifold while the $D_{3/2}$ levels mix with their partners in the $D_{5/2}$ manifold. Because the transitions $P_{1/2} \rightarrow D_{3/2}$ and $P_{3/2} \rightarrow D_{5/2}$ are dipole-allowed, the small admixtures of these other states make the $P_{1/2} \rightarrow D_{5/2}$ decay allowed, albeit weakly.

It is immediately apparent from the matrix element that the decay rate due to this ‘*j*-mixing’ effect is both proportional to B^2 and inversely proportional to the square of the fine structure splitting ΔE_{FS}^2 . The rate has been calculated both perturbatively [70,71] and exactly [66] (valid for all field strengths), and for moderate fields it can be shown that

$$\Gamma_{j\text{-mix}} \sim \Gamma_{397} \times 4.2 \times 10^{-7} (B/T)^2. \quad (3.4)$$

At 1.85 T and under typical Doppler cooling intensities, this corresponds to a shelving rate of $\sim 25 \text{ s}^{-1}$ and the ion will remain trapped in these dark states for 95% of the time. In practice this means that it is essential to repump on the 854-nm transition continuously while Doppler cooling. Although *j*-mixing decays only occur to the central four states in the manifold, decay to the $m_j = \pm 5/2$ states from $P_{3/2}$ can occur during repumping, so all six states must be driven.

Although *j*-mixing occurs for any ion with D-states, among the species commonly used in ion trap experiments it is only a significant concern in calcium - and to some extent strontium - where the fine-structure splitting, proportional to $1/n^3$, is small. During Doppler cooling the presence of *j*-mixing only increases the number of states that must be repumped and the absolute rate of the process is of little concern. However because the 854-nm laser must be blocked during spectroscopy or coherent manipulation on the 729-nm transition, *j*-mixing limits the fidelity of the state detection used in these techniques. For this reason calcium is not an ideal candidate for Penning trap applications at high fields, though at 1.85 T the disadvantages are arguably balanced by the convenience of easily accessible laser frequencies.

⁴The transition requires magnetic quadrupole or electric octopole coupling to proceed directly.

3.2.2 Transition strengths and saturation intensities

The cooling or repumping lasers addressing a lower level may drive transitions to a range of upper levels to similar effect. However, the strength of a particular transition will depend on the polarisation and wavevector of the driving field (see Appendix A) and the geometry of the trap and beam optics will determine which options are most convenient. As we only have optical access parallel with and perpendicular to the magnetic field, we consider only the transition amplitudes for beams in these directions. For technical reasons it is also advantageous for us to use only linearly polarised beams where possible. From Equation 2.52 and Appendix A, we find the saturation intensities of an electric dipole and electric quadrupole transition to be

$$\begin{aligned}
 I_{sat}^{E1} &= \frac{2c^2\hbar^2\varepsilon_0\alpha k_L^3\Gamma}{3e^2(2j'+1)} \left[\begin{pmatrix} j & 1 & j' \\ -m_j & \Delta m_j & m_{j'} \end{pmatrix} c^{(\Delta m_j)} \cdot \hat{\mathbf{e}}_L \right]^{-2} \\
 I_{sat}^{E2} &= \frac{2c^2\hbar^2\varepsilon_0\alpha k_L^3\Gamma}{15e^2(2j'+1)} \left[\begin{pmatrix} j & 2 & j' \\ -m_j & \Delta m_j & m_{j'} \end{pmatrix} \hat{\mathbf{k}}_L \cdot c^{(\Delta m_j)} \cdot \hat{\mathbf{e}}_L \right]^{-2}.
 \end{aligned} \tag{3.5}$$

The saturation intensities for the transitions of interest in the axial and radial directions for linearly polarised beams are given in Table 3.2.

It is necessary for us to cool both axially and radially with the same laser frequencies, therefore the Doppler cooling lasers must address the σ transitions as the π transition cannot be driven axially. The 397-nm transition saturation intensity is equivalent to $P_{sat} \sim 10\mu\text{W}$ over a Gaussian beam with a waist $w_0 = 50\mu\text{m}$ typical of those used in this experiment.

The 866-nm and 854-nm lasers must empty all four and six states of their targetted manifolds so vertically polarised radial beams are not an option. However, either a horizontally polarised radial beam or axial beam work equally well for these repumper lasers. The saturation intensities of the transitions are much lower, typically of the order 200 nW per transition for a $50\mu\text{m}$ beam waist.

As will be discussed in Section 3.6, several technical considerations influence the choice of transition for resolved-sideband spectroscopy and cooling, but it is immediately apparent from Table 3.2 that the π transitions are not accessible with the current apparatus and the δ transitions are only accessible with a radial 729-nm beam.

Transition		Type	Saturation Intensity (Wm^{-2})		
			Axial Linear	Radial Linear (H)	Radial Linear (V)
397 nm	$\pm\frac{1}{2} \rightarrow \pm\frac{1}{2}$	π	–	–	2600
	$\pm\frac{1}{2} \rightarrow \mp\frac{1}{2}$	σ	2600	2600	–
866 nm	$\pm\frac{1}{2} \rightarrow \pm\frac{1}{2}$	π	–	–	12
	$\pm\frac{1}{2} \rightarrow \mp\frac{1}{2}$	σ	70	70	–
	$\pm\frac{3}{2} \rightarrow \pm\frac{1}{2}$	σ	23	23	–
854 nm	$\pm\frac{1}{2} \rightarrow \pm\frac{1}{2}$	π	–	–	14
	$\pm\frac{3}{2} \rightarrow \pm\frac{3}{2}$	π	–	–	21
	$\pm\frac{1}{2} \rightarrow \mp\frac{1}{2}$	σ	58	58	–
	$\pm\frac{1}{2} \rightarrow \pm\frac{3}{2}$	σ	170	170	–
	$\pm\frac{3}{2} \rightarrow \pm\frac{1}{2}$	σ	28	28	–
	$\pm\frac{5}{2} \rightarrow \pm\frac{3}{2}$	σ	17	17	–
729 nm	$\pm\frac{1}{2} \rightarrow \pm\frac{1}{2}$	π	–	–	–
	$\pm\frac{1}{2} \rightarrow \mp\frac{1}{2}$	σ	2.9×10^{-6}	–	2.9×10^{-6}
	$\pm\frac{1}{2} \rightarrow \pm\frac{3}{2}$	σ	1.4×10^{-6}	–	1.4×10^{-6}
	$\pm\frac{1}{2} \rightarrow \pm\frac{5}{2}$	δ	–	1.2×10^{-6}	–
	$\pm\frac{1}{2} \rightarrow \mp\frac{3}{2}$	δ	–	5.8×10^{-6}	–

Table 3.2: Saturation intensities of relevant transitions in $^{40}\text{Ca}^+$ for linearly polarised beams propagating in the axial ($\mathbf{k}_L \parallel \mathbf{B}$) and radial ($\mathbf{k}_L \perp \mathbf{B}$) directions.

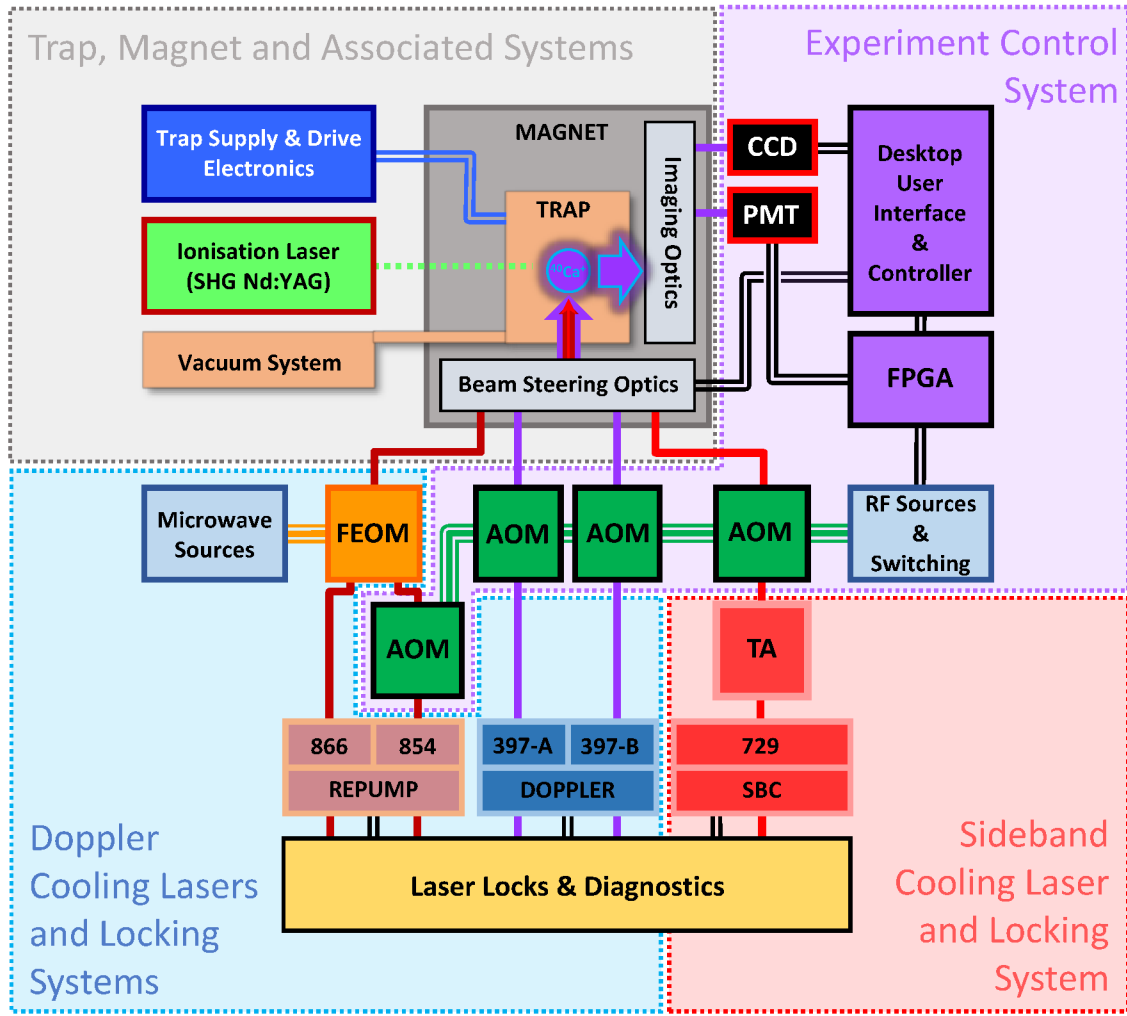


Figure 3.2: Schematic overview of experimental subsystems and their interactions. Optical paths are shown in solid purple, red and brown, and in dashed green in the case of the pulsed ionisation laser. Black, double-line paths represent data and control signals, while green and orange triple-line paths represent RF and microwave signals, respectively.

3.3 Overview of Apparatus

The experimental apparatus can be broadly divided into four subsystems and the four remaining sections of this chapter will discuss each in more detail. Particular focus will be given to the elements of these systems that I have designed or updated for the work in this thesis.

The ion trap and superconducting magnet form the heart of the experiment and along with their associated systems are described in Section 3.4. These systems include the electronics used to control the trap voltages, compensation and axialisation drive, the oven and pulsed Nd:YAG laser used to load the trap, and the optics in and immediately around the trap vacuum system used for beam focusing, steering and fluorescence collection.

The laser systems used for Doppler cooling currently consist of two 397-nm diode lasers, an 866-nm diode laser and an 854-nm diode laser. These lasers are locked to a series of optical cavities that provide good stability and narrow (~ 1 MHz) linewidth on all timescales while allowing the laser frequencies to be adjusted under lock. Both repump lasers pass through a single fibre electro-optical modulator (FEOM) with 20 GHz bandwidth, which provides sidebands close to all the necessary repump transition frequencies. These lasers, their locks and the FEOM are described in Section 3.5.

Resolved-sideband spectroscopy and cooling is performed with an ultra-narrow linewidth (< 1 kHz) diode laser at 729 nm. This laser is locked via the Pound-Drever-Hall (PDH) technique to a temperature-controlled, high finesse cavity made of ultra-low-expansion glass (ULE) and amplified in two stages, with a second injection-locked diode followed by a tapered amplifier. These systems are described in Section 3.6.

On the most basic level, the experimental procedure involves applying a sequence of different laser frequencies to the ion while monitoring the emitted fluorescence with a pair of photomultiplier tubes (PMT). The lasers are switched by individual acousto-optic modulators driven by a range of attenuatable RF sources. Precisely timed TTL switching of these RF sources is performed by a field programmable gate array (FPGA) which also counts photons detected by the PMTs. Experimental sequences can be programmed and uploaded to the FPGA using a custom .NET interface on a desktop PC, which also receives and analyses spectroscopy data sent back from the FPGA while the experiment is running. This computer also provides remote control of motorised mirror mounts that steer the beams in the trap, and monitors the output of a CCD camera that is sometimes used for diagnostics or imaging of larger crystals of ions. These experimental control and monitoring systems are described in Section 3.7.

3.4 Trap, Magnet and Associated Systems

The experiments described in this thesis have all been conducted in a cylindrical electrode Penning trap known as ‘SPECTRAP’. SPECTRAP was originally designed by Manuel Vogel for use in heavy-ion experiments at GSI Darmstadt [72], but was designed to also be compatible with the superconducting magnet at Imperial. The Imperial trap has been gradually modified by previous members of the group, notably Shailen Bharadia and Sandeep Mavadia, who also designed the optical systems for the cooling beams and for the imaging of the trapped ions. The trap has now remained under vacuum and in the bore of the magnet for the last three years, and no major modifications have been made to the mechanics or optics during this period. For an in-depth discussion of the trap mechanics and original vacuum system design, the reader is referred to reference [73]. The current optical system is described in detail in reference [56]. In this section I review only those details necessary to understand the improvements that I have made in the last two years.

3.4.1 Overview

Trap

The trap, shown in Figure 3.3, consists of a concentric stack of 7 circular oxygen-free copper electrodes, of internal diameter 21.6 mm, supported by insulating sapphire balls. The outermost two electrodes were designed for capturing high-energy ions at GSI, and are earthed throughout the current experiment. The next electrodes are the end-caps of the trap and are held at an equal potential of up to 366 V, the stability limit of the trap at 1.85 T. These provide the axial confinement. The next two electrodes are narrow compensation electrodes, that allow the shape of the field to be modified, or the centre of the potential shifted in the vertical direction. The central electrode in the stack is the ring, which is held at a potential close to ground. The ring is split into four segments, each of which has a 4 mm aperture providing radial optical access to the trap; the radial cooling beam passes through two diametrically opposed segments, while fluorescence is collected through the perpendicular pair. The potentials on the split ring segments can be individually controlled. This allows the centre of the trapping potential to be shifted in the radial direction to compensate for stray potentials and also provides a means to apply an RF ‘axialisation’ potential to couple the radial modes.

Vacuum system and optics

The trap is held in a vacuum chamber with windows above and below (via a 45° mirror) the stack of electrodes providing optical access for an axial beam, and four windows at the side of the can aligned with the ring electrode apertures. Unfortunately none of the radial windows are AR coated, meaning that the background levels of scattered light are much

higher than would be expected in such a large and unobstructed trap.

The vacuum system supports the trap and is in turn bolted to an optical breadboard (Figure 3.4) which also carries the optics used for combining and steering the cooling and spectroscopy beams. The repump lasers (866nm and 854nm) and 729-nm laser reach this breadboard through a pair of optical fibres, one for the axial beam and one for the radial. The two 397-nm cooling lasers are carried by a second pair of fibres. The light from the fibres is combined on dichroic mirrors before being reflected by pair of steering mirrors and focussed into the trap. One of the radial steering mirrors is attached to a mount with piezo inch-worm drives⁵ that allows the position to be adjusted precisely (although not very repeatably) from desktop software. This allows easy adjustment of the radial beam offset. The waists of the axial and radial 397-nm beams at the position of the ion are $w_0 \sim 45 \mu\text{m}$ and $w_0 \sim 55 \mu\text{m}$ respectively.

Trap positioning

This breadboard sits on a tripod of sharpened 6 mm hex bolts resting on a solid aluminium base plate. The base plate can be raised and lowered using a ‘Genie Lift’ crane, before being clamped to the solid aluminium legs that support the superconducting magnet. The positioning provided by the crane is relatively crude ($\pm 3 \text{ mm}$) and after the base plate has been clamped to the magnet legs, the final position of the trap can be adjusted using the tripod of bolts. These provide vertical alignment to a precision of approximately $50 \mu\text{m}$, as well as the ability to tilt the vacuum system within the bore. This tilt allows the radial position of the trap centre to be adjusted (to $\sim 0.5 \text{ mm}$ precision), but with significant crosstalk with the angle of the trap axis within the field. Once the trap position has been finalised, the breadboard is clamped to the base plate at the points of the tripod. This trap positioning system is awkward, inaccurate, somewhat hysteretic and would be quite unsatisfactory for precision work. Nonetheless, as will be shown in Section 3.4.2, we have used these fairly crude adjustments of the ion’s position within the field to make significant improvements to the practical performance of the experiment.

Ion loading

Loading is provided by photoionisation of a beam of neutral calcium emitted by an atomic oven situated at the bottom of the trap, off-axis. Photoionisation occurs as a three-photon process driven by the second harmonic of a pulsed Nd:YAG laser at 1064nm. The YAG beam enters the trap from above, through the upper window, in the opposite direction to the other lasers. This process generally loads several ions at once, although some may initially lie on very large stable orbits within the trap and take several minutes to cool sufficiently to produce detectable fluorescence. To load single ions (or any other small

⁵Newport Agilis AG-M100N




Figure 3.3: Schematic of SPECTRAP electrodes and imaging optics. The axial cooling beam propagates upwards, while the radial beam propagates into the page. Diagram reproduced from [56] with permission.




Figure 3.4: Layout of final beam optics and imaging optics on breadboard (bottom) and position within bore (top). Diagram reproduced from [56] with permission.

number) near-deterministically, a small cloud is initially loaded and cooled before the endcaps are grounded for a period of $5 - 100 \mu\text{s}$. With the confinement removed, the resultant Coulomb explosion forces the ions out towards the ends of the trap where some are lost, preferentially those in large orbits. When the end-cap potential raises again, the ions return to the centre of the trap and are recooled by the Doppler lasers. By repeating this process while monitoring the either the fluorescence or the image on a electron-multiplied charge-coupled device camera (emCCD), one can observe the number of ions as it is reduced in a relatively controllable way. The discrete electronic circuit used to ‘pulse down’ the trap end-caps in this manner was originally designed by D. Crick. The integrated trap supply I recently designed (Section 3.4.3) includes an updated version of this circuit.

Two key improvements I have made to the trap system since 2012 have been the reduction of the magnetic field gradient at the position of the ion and the development of new trap supply electronics. Reducing the field gradient (Section 3.4.2) has made the transition frequencies of the ion more stable, paving the way for precision spectroscopy. The new supply integrates a number of previously separate systems and allows for independent AC and DC compensation of the trapping potential, making it possible to use the axialisation drive at high trap voltages (Sections 3.4.3 & 3.4.4)

3.4.2 Magnetic field inhomogeneity

An ideal solenoid, with length much greater than its diameter, provides a uniform field within its enclosed volume. However the coils of our superconducting magnet have a length of approximately 300 mm and surround a bore of radius 52 mm, making the field far from uniform. The field is cylindrically symmetric, increases towards the geometric centre in the axial sense and decreases towards the centre in the radial sense, making the strength of the field near the centre quadrupolar, to first order.

Because the field is so high, even very slight fractional variations in the Zeeman splittings can lead to problematic shifts in the transition frequencies. The linewidth of the 729-nm spectroscopy laser is less than 1 kHz and ideally movement within the magnetic field should not cause broadenings or slow drifts greater than this. The differential Zeeman splitting of the $S_{1/2,+1/2} \rightarrow D_{5/2,-1/2}$ transition is 41 GHz at 1.85 T.⁶ To prevent drifts on the kilohertz level, the field at the position of the ion must not be allowed to vary by more than 1 part in 40 million.

The trap and vacuum system are fixed securely to the breadboard, but the entire system has a degree of flex, making it susceptible to acoustic vibrations. Variations in the temperature of the laboratory (as much as 2 K over the course of a day), can also cause the breadboard and vacuum system to expand, moving the trapped ion within the magnetic

⁶This was used for many of our initial experiments with resolved-sideband spectroscopy and cooling. More recently we have used the $S_{1/2,-1/2} \rightarrow D_{5/2,-3/2}$ transition, which has half the field sensitivity.




Figure 3.5: Measured thermally induced drift of carrier transition frequency over several hours. The Nd:YAG flashlamps had been switched off for several hours and the power supply fan was blowing cool air towards the trap. The fan was switched off at point ‘A’ and the transition frequency changed rapidly. At point ‘B’, the flashlamps were switched on at the same time as the fan and the frequency stabilised. At point ‘C’ the fan and lamps were switched off again and at point ‘D’ the fan was switched on without the lamps, returning the system to its original state.

field. The steel⁷ and aluminium used in the vacuum system and breadboard both have coefficients of thermal expansion close to $\alpha = 20 \times 10^{-6} \text{ K}^{-1}$.

Two processes are likely to contribute to movement within this field. The first is linear expansion of the arm of the vacuum system that extends into the magnet bore. This has a length of approximately 40 cm and thus experiences a vertical displacement of $8 \mu\text{mK}^{-1}$. With the steel components, the relatively low thermal conductivity also means that significant temperature differentials can develop between two sides of one object. Naturally this will only occur if the source of heat is directional, but this was indeed the case in our system, where the exhaust fans of our main electronics rack and the Nd:YAG power supply were all directed towards the magnet and trap. Such differentials would then cause the components of the vacuum system to bend, leading to deflections of the unconstrained end that contains the trap.

The vacuum system, breadboard and magnet cryostat form a rather complex mechanical and thermal structure and the thermal environment provided by the room, its occupants and the equipment is also not easily modelled. However, in empirical terms, strong correlations have been seen in the drift of the ion transition frequencies and the variation of the thermal load. When large pieces of equipment exhausting near the trap were switched off

⁷Type 304 & 316 stainless steels, which have very low magnetic permeability [73].

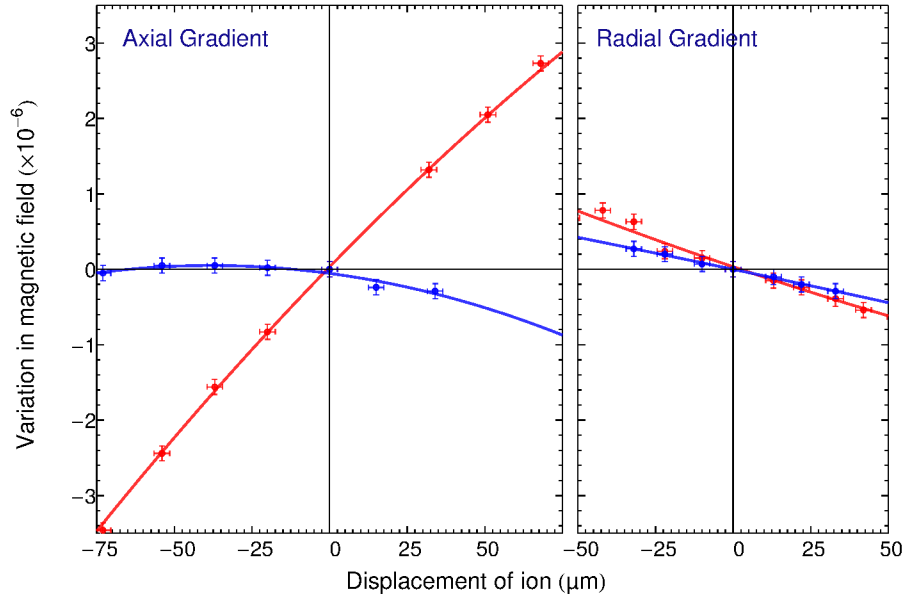


Figure 3.6: Axial and radial fractional variation of the magnetic field in the vicinity of the ion, before (red) and after (blue) adjustment of the trap position to minimise the gradient. The axial gradient is reduced from 77 mTm^{-1} to 9.2 mTm^{-1} , bringing the trap within $40 \mu\text{m}$ of the field centre in the axial direction. The radial gradient is reduced from 26 mTm^{-1} to 16 mTm^{-1} .

the carrier frequency was seen to vary rapidly (Figure 3.5), with total shifts of several hundred kilohertz observed. Whether this was due to a greater total variation in temperature or due to differential flexing across various components is not clear, but it is very likely that the thermal movement due to these localised sources would be considerably greater than the $\sim 16 \mu\text{m}$ shift due to linear expansion under the variations of the laboratory temperature.

These major drifts were significantly reduced by repositioning of some equipment and the provision of a skirt around the base of the magnet to prevent movement of air around the trap and its supports. However, not all equipment could be moved and the macroscopic temperature variations of the lab cannot be easily improved, so minimisation of the effects of the inevitable residual movement in the field was essential.

The spatial variation of the magnetic field was measured by displacing the ion a known distance and measuring the resultant shift of the carrier transition frequency using 729-nm resolved-sideband spectroscopy. The ion's position was moved in the axial and radial directions via the application of DC offsets on the compensation and ring electrodes. This displacement was measured via emCCD camera imaging at known magnification $m = 3.0 \pm 0.1$ [34]. The initial variation of the magnetic field in the axial and radial directions is shown in red in Figure 3.6. The gradient of the axial field shows that 100 kHz shifts would occur for displacements of $50 \mu\text{m}$, making thermal movement within this inhomogenous field consistent with the level of drift that was observed.

The field increased as the ion was displaced upwards, suggesting that the centre was located above the initial position of the trap. The trap breadboard was gradually raised 7 mm using the tripod of hex bolts and the inhomogeneity monitored. The gradient was shown to change sign, before the trap was slightly lowered to bring it back to the centre of the field. The field variation at the final position is shown in blue in Figure 3.6, demonstrating that the gradient has been reduced by a factor of 8. The centre of the field appears to lie approximately $40\mu\text{m}$ below the ion – the limitations of the current positioning system and variations due to thermal expansion prevent more precise alignment.

The radial position was adjusted in a similar way. Here the gradient could only be reduced by a factor of 1.5, although the initial position was closer to the centre of the field. It is considerably more difficult to find a minimum in two dimensions than one, particularly with a positioning mechanism that does not allow the movement in each cardinal direction to be decoupled, and this level of homogeneity was found to be sufficient for our purposes.

3.4.3 Trap supply electronics

Another important recent development was the consolidation of several previously disparate systems involved with the trap supply, endcap pulse-down circuitry and axialisation drive into a single, low-noise unit that I designed. The new supply takes inputs from a high (end-cap) voltage supply, a low voltage supply, a pulse-down logic signal and a radiofrequency source for the axialisation drive.

The system contains a pulse-down circuit that drops the end-cap voltage to zero (see Section 3.4.1), controlled by the pulse-down logic input. The transistors in this circuit introduce significant noise to the endcaps if they are left connected during the operation of the experiment. To avoid this the entire circuit is bypassed after the loading phase by a pair of relays that route the high voltage via a 20 Hz low pass RC filter to the endcaps. These relays are controlled by a toggle switch on the front panel.

The low voltage supply is used to power a pair of regulators which in turn supply a LTC6655 low noise voltage reference. This reference is used to supply six linear potentiometers which are used to provide DC offsets to the ring segments and compensation electrodes. The RF axialisation input is sent via four potentiometers to four unity gain amplifiers, currently TL071, but due to be upgraded to a lower noise model. Two of these amplifiers are inverting and two non-inverting, providing the necessary oscillating quadrupole potential when connected to the ring segments. The potentiometers on the input allow the amplitude of the axialisation drive on each segment to be varied. The entire axialisation system and its amplifiers can be isolated from the rest of the circuit by a set of relays controlled by a single toggle switch. The DC compensation voltages are low-pass filtered at 20 Hz and passively summed with the AC axialisation voltages before being connected to the ring electrodes.

Schematics for the compensation and pulse down circuits can be found in Appendix D.

3.4.4 Trap compensation

In an ideal trap, the centre of the trapping potential remains fixed as the trapping voltage is varied, while any additional RF fields have a null that coincides perfectly with the equilibrium position of the ion, making it immune to micromotion to first order. Naturally, these conditions are not always achieved with a real trap and the trap potentials must be adjusted to compensate for the imperfections.

DC compensation of trap

Over a period of time, we have observed that stray potentials develop in the trap that cause a displacement of the ion equilibrium position at low fields, and vary over the course of weeks. While the exact source is not confirmed, a correlation has been seen between periods of intensive heating of the atomic oven and large changes in these potential, and it is hypothesised that these may be due to a build up of adsorbed calcium atoms on the copper electrodes. Because the work functions of calcium and copper are 2.9 eV and 4.26 eV respectively, the complete coating of an electrode in calcium would increase the voltage of that electrode by 1.36 V, although a partial coating would of course effect less of a change.

To compensate for these stray fields, we supply additional DC voltages to the ring electrodes, which are nominally held at ground. The ion position is recorded using the emCCD at a very high trap voltage, where the stray fields have a very small effect on the equilibrium position – this should be the geometric centre of the trap. The trap voltage is then reduced, which in general leads to the ion moving away from its original position, and the DC compensation voltages are adjusted to return the ion to the high voltage position. When the ion is properly compensated, the voltage can be varied from just above zero to the stability limit without any visible shift in the ion position. Naturally, this does not correct for any displacement along the optical axis of the imaging system. To correct for stray potentials in this direction, the compensation voltages are initially adjusted to ensure the fluorescence does not vary significantly, which implies that the ion is not moving across the width of the cooling beams. To perform a more accurate compensation, the carrier frequency (which varies spatially due to the gradient of the magnetic field) can be measured at a range of voltages, and the compensation adjusted to keep this constant.

An uncompensated ion has an equilibrium position which varies to first order as a function of trap voltage. Noise on the trap voltage can therefore cause small movements of the ion and lead to transition frequency noise which may broaden lines or introduce spectral artefacts. Proper compensation reduces these effects, and also ensures that the beam pointing does not need to be adjusted when the trap frequency is changed. The DC compensation process needs to be repeatedly periodically, as the stray potentials on the electrodes vary over time.

AC compensation of axialisation drive

After DC compensation, the ion should remain in the centre of the trap at all voltages. If an axialisation potential is applied to the ring electrodes, the RF null should be aligned with this position. In practice, several factors may effect the validity of this assumption. If, for example, one of the trap electrodes were mechanically misaligned, this would be compensated for in DC while the AC potential would remain offset – though in general this is expected to be a very small effect in such a large trap. Secondly, if the capacitances of the circuits connecting the individual ring segments vary then the amplitude of the RF signal on each electrode will be unequal. As the wires supplying each segment take quite different paths it is quite possible that significant variations could occur. Finally, the electronics supplying the axialisation drive – the inverting and non-inverting amplifiers in the trap supply – may simply not supply perfectly equal amplitudes to each segment.

If any of these factors affects the axialisation potential, the effect will be the same – to shift the RF null away from the equilibrium position of the ion and introduce micromotion at the axialisation frequency. At low axialisation voltages this is not a huge problem, but to provide sufficient coupling of the radial modes at high end cap voltages the axialisation voltage must be increased [59], increasing the micromotion in turn and preventing the secular motions from cooling efficiently. For this reason, axialisation of single ions had generally not been successful at trap potentials over ~ 50 V. To compensate the axialisation potential, the drive voltage is gradually increased until either visible micromotion is seen on the emCCD image, or the fluorescence level drops. The amplitudes on the four segments are then adjusted to return the ion to stable behaviour before the drive amplitude is increased further and the process repeated. When correctly compensated it is possible to apply the maximum 10 V of axialisation with no visible micromotion and only a very small drop in fluorescence. The ion can then be trapped and cooled at axial frequencies approaching the stability limit of $f_z \sim 500$ kHz, which requires an axialisation amplitude of 5 – 7 V. By contrast, cooling under intensity gradient methods alone (with the current optical system with a radial waist of $w \sim 50 \mu\text{m}$) should be impossible for axial frequencies higher than 220 kHz (see Equation 2.95).⁸

MCA

An alternative method of monitoring the micromotion due to a coherent drive is to perform an RF-photon correlation measurement. Typically a time-to-amplitude converter (TAC) is used to measure the time between the RF drive trigger and the detection of a photon by the PMT, and the resultant pulses are binned with a multi-channel-analyser (MCA). If

⁸The fact that we have successfully used intensity gradient cooling at frequencies higher than 250 kHz implies the presence of a gradient of $y_0 < 15 \mu\text{m}$ and it is likely that in these cases the ion was aligned to a steep, non-Gaussian artefact or that saturation in the central region of the beam meant the gradient at the edges was steeper than expected.

micromotion is present, then the scattering rate will be related to the phase of this coherent motion which is driven directly by the RF source.

An EG&G Ortec 567 TAC/SCA is used to provide the TAC signal and event trigger, but the matching MCA computer card is no longer working. I designed a cheap ($<£100$) alternative based on an Arduino, which allows the TAC amplitude to be recorded at rates of up to $90\,000\text{ s}^{-1}$. The MCA, shown in Figure 3.7 consists of an Intersil HA5351 sample-and-hold and an Analog Devices AD976AN 16-bit parallel ADC. The sample-and-hold holds the TAC pulse for long enough for the ADC to sample the amplitude (the pulse from the EG&G card is $\sim 1\mu\text{s}$ long). The Arduino then reads the digital signal from the ADC and after every 2000 samples sends the accumulated TAC height readings to a desktop PC via USB. There they are received by an interface written in Processing, which bins the readings and adds them to a histogram spanning the range of possible TAC lengths.

If correlations are present these should appear as peaks on the timing histogram and when the detection rate is much lower than the drive frequency one expects to see a periodic signal. A Fast Fourier Transform can then be applied to the histogram data to extract the contributing frequencies. This spectrum can be used to monitor the degree of micromotion – as the compensation is improved the size of the peak at the axialisation frequency shrinks.

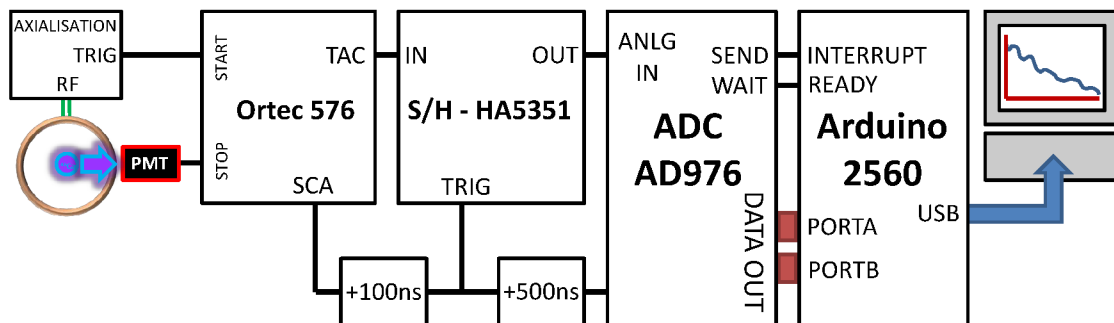


Figure 3.7: Simplified schematic of Arduino-based multi-channel analyser used for RF-photon correlation measurements

3.5 Doppler Cooling Lasers and Locking Systems

3.5.1 Overview

The Doppler cooling and repumping laser systems are shown in Figure 3.8. Below I give an overview of each part of this system and the performance required for the current experiments; for more details see Sections 3.5.2 and 3.5.3.

397-nm lasers

The 397-nm Doppler cooling laser frequencies are provided by a pair of Toptica DL100 external-cavity diode lasers (ECDL) in Littrow configuration and fitted with gallium nitride laser diodes.⁹ Each ECDL has an output power of approximately 15 mW. Observation of the optical beat note of the outputs of these two unlocked ECDL provided a linewidth of ~ 10 MHz over 10 ms, and drifts of many hundreds of megahertz per day. To approach the Doppler limit for a transition with $\Gamma_B \sim 2\pi \times 22$ MHz, the linewidth and drift should be constrained to 1-2 MHz via a locking system. Short term stability is provided by Pound-Drever-Hall (PDH) locks [74] to a pair of low-finesse ($\mathcal{F} \sim 30$) optical cavities with lengths that are adjusted via piezo actuators. These cavities reduce the short-term linewidth but are susceptible to thermal drift. Such drift is prevented by comparing both lasers to a frequency-stabilised HeNe reference laser via a scanning transfer cavity. The error signal from the transfer lock is used to stabilise the low-finesse cavities, which then provide stability to the lasers on long and short timescales.

854-nm and 866-nm lasers

The repump lasers are provided by a pair of Littrow-configuration ECDLs, based on the design by Arnold [75] and built in-house. The 854-nm ECDL has an output power of approximately 20 mW and the 866-nm has an output power of > 50 mW. Both lasers pass through a fibre electro-optic modulator, which produces combs of sidebands that allow the single wavelength sources to drive many Zeeman transitions simultaneously. The linewidths of these diode systems have not been determined via a beatnote as we only have one of each wavelength, but can be inferred from the width of dark resonances observed when scanning the lasers relative to the Doppler cooling. These suggest a width of order 1 MHz. The stability requirements for these lasers in the current experiment are not very stringent (of the order 10 MHz), but it is nonetheless essential to prevent long-term drifts. Light from both repumper lasers is sent to a second scanning transfer cavity, allowing them to be locked to the same HeNe reference as the cooling lasers. The error signals from this transfer lock are fed back via P-I controllers to the ECDL piezos.

⁹Nichia NDV4814T

Lock controller

The transfer locks are controlled by individual Arduino microcontrollers, which communicate via a serial protocol with a master Arduino in a system control box. This box allows the set point frequencies of each of the four lasers to be controllably moved while remaining under lock, and also displays information regarding the 397-nm lock status and relative laser frequencies. The serial communication and lock control system was developed by two MSci students, S. Woodrow and J. Aryaman, under my supervision; more details are available in reference [76].

Auxilliary systems

The 854-nm laser and both 397-nm lasers can be shuttered with double-passed AOMs, controlled by an FPGA - these systems are described in Section 3.7. The FPGA also monitors logic signals from both the PDH locks to ensure spectroscopic runs are performed under the expected conditions. Laser wavelengths are monitored via 3 channels of a Fizeau wavemeter¹⁰, the use of which is provided by the Imperial College Centre for Cold Matter. The two repump lasers share one channel, the Doppler cooling lasers share a second and a third is reserved for the 729-nm laser. Mechanical shutters provide switching between the lasers on shared channels.

3.5.2 Cooling laser locks

PDH lock

The PDH locking system was developed primarily by G. Stutter and will be discussed in detail in his forthcoming thesis. I give a very brief outline here.

Each laser is locked to a dedicated cavity with finesse ~ 30 and a free spectral range of 1.5 GHz. The cavity spacers are machined from Schott Zerodur low expansion glass and the cavity length can be adjusted with a tripod of piezoelectric transducers that support one of the mirrors. The current through each laser diode is modulated at ~ 10 MHz to produce a pair of weak FM sidebands, which provide the PDH error signal¹¹ when the light from the laser is coupled into the cavity and the reflected signal mixed with the oscillator driving the FM.

The PDH error signal is used to provide fast feedback to the laser current to remove noise up to a bandwidth of ~ 0.5 MHz. A further integrator stage feeds back to the ECDL grating piezo to prevent the accumulation of a frequency offset. The integration is performed digitally by an Arduino microcontroller, which also monitors a photodiode measuring laser transmission through the optical cavity. If the transmission drops below a fixed threshold

¹⁰High Finesse WS-6

¹¹Technically, these are ‘dither’ locks rather than Pound-Drever-Hall, as the laser modulation frequency is smaller than the cavity linewidth but though the performance is not quite as high, the principle and apparatus are identical. The principles of and differences between PDH and dither locking are described in [56].

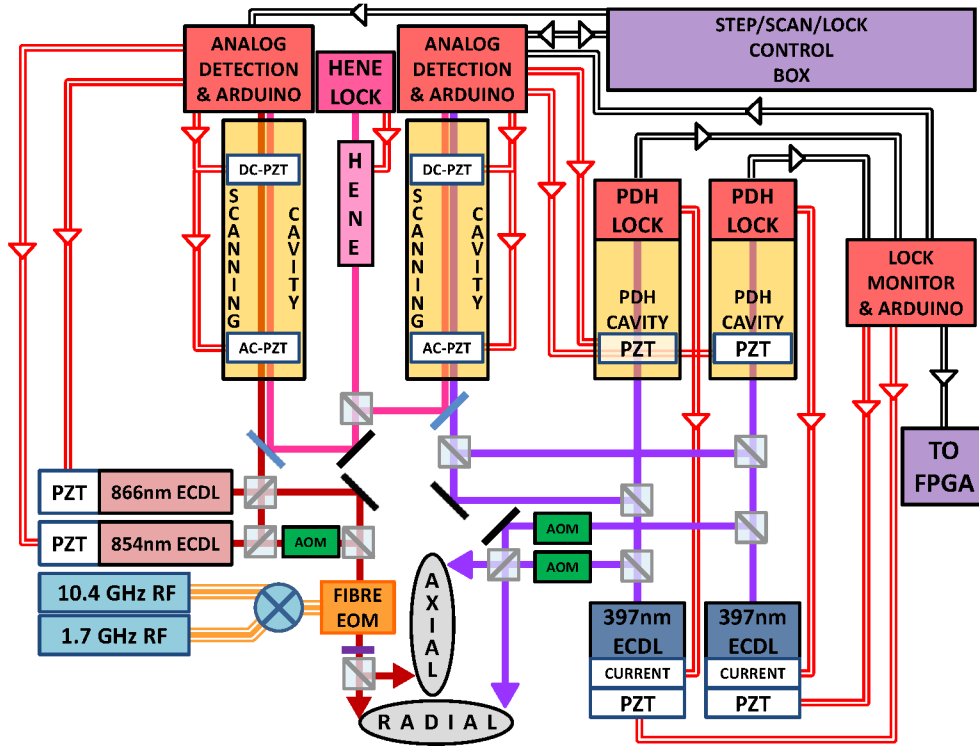


Figure 3.8: Locking, modulation and control systems for 397-nm Doppler cooling lasers and 866-nm and 854-nm repump lasers. For a description of each subsystem, see the main text in this section (3.5.1).

(usually half the amplitude observed on resonance), the Arduino stops integrating and holds the integrated value until transmission is recovered. This helps prevent runaway feedback during sharp perturbations such as knocks to the table or even a blocked beam. The logic signal conditioned on the cavity transmission is also used by the FPGA pulse controller to monitor the status of the 397-nm locks.

The cavity length-adjustment piezos are connected to a piezo driver which allow both manual adjustment and an external control voltage input. While locked to the PDH cavities only, the manual potentiometer allows more accurate tuning of the laser frequencies than can be achieved by using the ECDL piezos directly. When long term stability is required, feedback from the scanning cavity system can be applied to the control voltage input.

Scanning cavity principle

Development of the scanning cavity systems for the cooling and repump lasers was the primary focus of my MRes research, and continued throughout the first year of my PhD. The basic principles, theory and considerations that went into the design are covered in detail in my MRes thesis [77].

The aim of the transfer lock is to transfer the stability of one ‘master’ laser onto one or more ‘slave’ lasers at a different wavelength. The master laser can be any conveniently

stabilised source, and the only limit to the range of wavelengths that can be spanned by the lock is the reflectivity of the cavity mirrors at the master and slave wavelengths. The lock works by coupling the master and slave lasers into the same optical cavity. Both of the cavity mirrors are mounted on piezo transducers. One piezo is small with short throw and low capacitance to allow fast scanning. The second piezo is a stacked device with long throw to allow compensation of the cavity's thermal expansion. A periodic signal¹² is then sent to the small piezo, which causes the length of the cavity to repeatedly scan over a given range.

As the cavity comes into resonance with each laser frequency in turn, that frequency is transmitted, producing a peak in the transmission spectrum. By using a dielectric beam splitter at the output, the signal due to lasers of different wavelength (e.g. master and slave) can be recorded on separate detectors.¹³ The signal from each photodiode is sent to an analog circuit which outputs a logical pulse whenever a peak is detected. A functional representation of the peak detector circuit is given in Figure 3.9 and a full schematic is given in Appendix D. If the scan range is set to cover slightly more than one free spectral range of the master laser, the combined signal from the peak detector outputs and the scan signal trigger will look something like that in shown in Figure 3.10.

The master laser is assumed to be stable, so to stabilise the slave laser(s), the system feeds back to adjust the slave frequency so as to maintain the cavity length difference between the master and slave transmission peaks. If the scan range and frequency are stable, then this equates to keeping t_1 constant. Two secondary effects must also be compensated for. The first is the thermal expansion of the cavity, which can be detected as a variation in time t_0 . This is not a problem until one of the master laser peaks leaves the scan range at which point the lock will malfunction. To prevent this, t_0 is held constant by feeding back to the large cavity mirror piezo. This ensures the equilibrium mirror separation remains constant despite changes in length of the cavity spacer. Secondly, if the scan amplitude or frequency changes, the time t_1 will change without a change in slave frequency. A third lock therefore monitors the time t_2 , which corresponds to the fixed free spectral range of the cavity at the master wavelength. Any changes in this spacing are compensated by feeding back to the amplitude modulation input of the scan function generator.¹⁴

A schematic version of the scanning cavity, detector and feedback system for a single slave laser is shown in Figure 3.11. The locks are implemented by sending the timing signals from the peak detectors and function generator to a pair of Arduino microcontrollers¹⁵ which control the locking system and provide analog feedback voltages via two 16-bit

¹²In general this will be a triangle or saw-tooth wave to provide a linear time-length response; at higher scan frequencies a sinusoid drive works better, but requires more complicated signal processing

¹³Typically photodiodes are sufficient, although the transmission of our cavity at 397-nm is so poor that we use a PMT to detect this light.

¹⁴Alternatively, all timing values may be divided by t_2 to provide appropriate normalisation - this technique is employed in the system used for the repump laser lock.

¹⁵Arduino MEGA2560

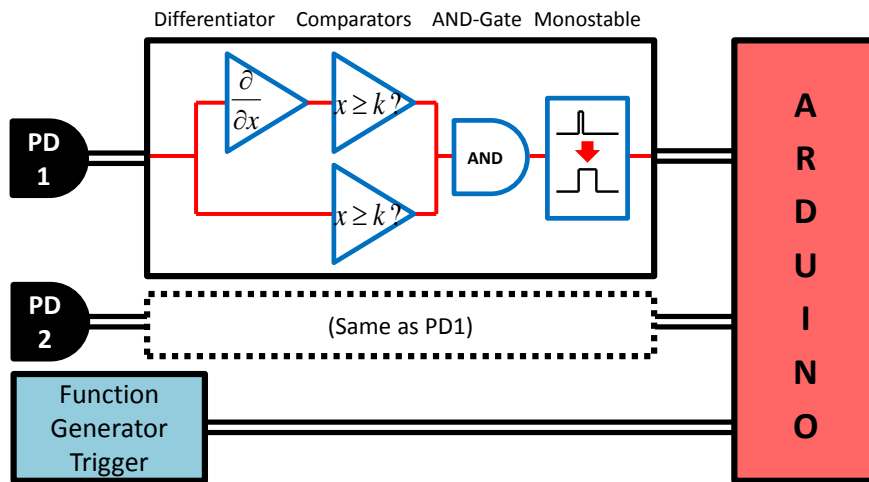


Figure 3.9: Peak detector principle. An analog signal from a photodiode is sent to two parallel signal paths. The first differentiates the signal, giving an output that crosses zero at every peak and trough. This is then sent to a comparator, which outputs a logical high whenever the differentiated signal drops below zero, so a rising logical edge indicates a maximum. The second parallel path is immediately sent to a second comparator, which takes as the other input a tuneable threshold voltage. If the raw photodiode signal exceeds this voltage, the comparator outputs logical high. The outputs of the comparators are sent to an AND gate, ensuring that the output only rises when a peak is detected above a certain threshold voltage, preventing triggering on low level noise. Finally, a monostable multivibrator extends the output pulse from the AND gate to a length long enough to be detected by the Arduino controller electronics. The Arduinos take inputs from two such peak detectors, corresponding to the peaks produced by the master and slave lasers, plus an input from the scan signal generator trigger.




Figure 3.10: Typical sequence of pulses recorded by the Arduino controlling the scanning cavity lock. The scan voltage is shown in black, while the yellow, blue and red traces represent signals from the function generator trigger, master and slave peak detectors respectively. The timing of the pulses is recorded on the forward scan, which are used to update the feedback voltages on the outward scan. Only a single slave peak is shown here for clarity; in principle any number of slave frequencies can be locked providing their order is known and their transmission peaks do not overlap. We generally lock two slave lasers to each cavity.

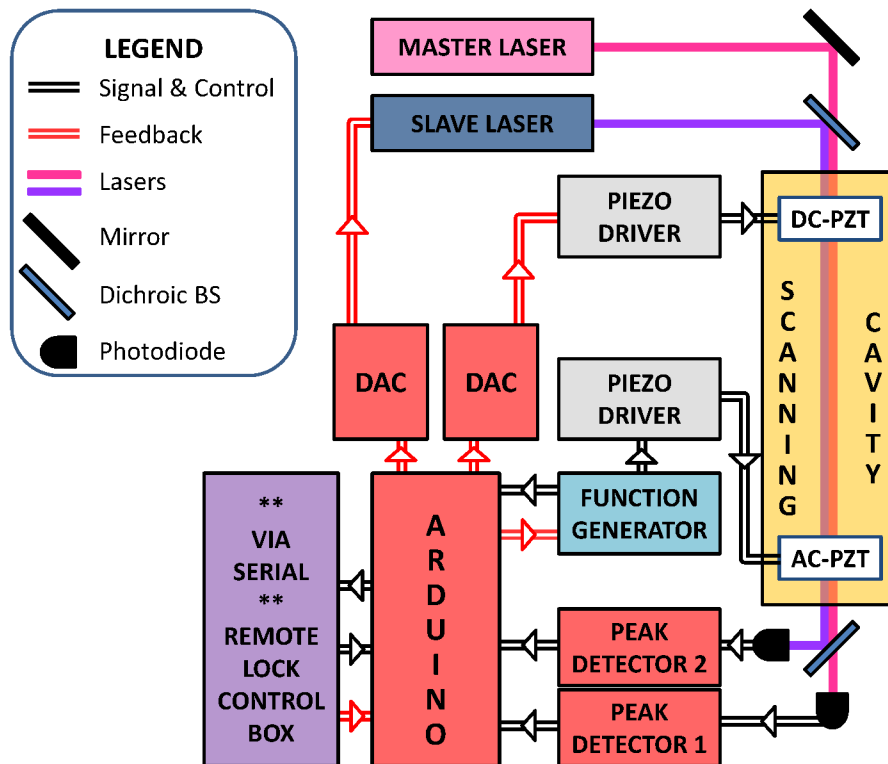


Figure 3.11: Simplified detection and feedback systems for a scanning cavity transfer lock.

parallel DACs¹⁶ per Arduino. In the current system, one Arduino ('master lock') provides feedback to the cavity length and piezo scan amplitude, while the second Arduino ('slave lock') provides feedback to the two slave lasers by varying the piezo voltage of the low-finesse cavities to which they are PDH locked. The slave lock Arduino also communicates via serial with the remote control box, sending information on the current status and set points of the slave locks and receiving user-generated commands to adjust these setpoints in fixed steps up or down in frequency so the laser frequencies can be adjusted while remaining locked.

HeNe reference laser

The master laser used for both the 397-nm and repump laser transfer cavities is a two-mode polarisation-stabilised 2 mW Helium-Neon laser. I based the polarisation-stabilisation system on a design by D. Crick [70]¹⁷ and added a low-pass filter to remove switch-mode power supply noise from the laser output. The HeNe linewidth was measured by beating the laser against another similar polarisation locked HeNe on a photodiode and observing the signal on a spectrum analyser. The locked linewidth is ~ 0.5 MHz with a drift of less than 0.5 MHz/day.

Optics

The cavity is arranged in a confocal configuration, so the mode structure consists of two degenerate sets (the even and odd transverse modes) half a free spectral range apart. Both lasers are mode matched as closely as possible to the TEM₀₀ mode so that coupling to the set of even degenerate modes is suppressed by a factor of ~ 20 dB and each laser produces one transmission peak per free spectral range.

The cavity mirrors are custom built by CVI Melles-Griot, with double-stacked dielectric coatings specified to have a reflectivity of $R > 0.995$ at 633 nm and 397 nm, suggesting a finesse of $\mathcal{F} > 600$. However, the performance has never met that of the specification, with measured finesses of $\mathcal{F}_{633} < 400$ and $\mathcal{F}_{397} < 60$. Transmission through the cavity is also very poor, particularly at 397nm for which less than 10^{-4} of the incident light is transmitted. This suggests that high absorption is more of an issue than low reflectivity itself. It is not known why this is so high, as the mirror surface appears undamaged under basic inspection and a fault during manufacturing is possible.

We have also found that the mirrors are not rotationally symmetric with respect to the polarisation of the incoming light and that in general the transmission peak splits into two polarisation components, separated by 30 MHz or more. To prevent this, both mirrors must be aligned to the polarisation axis of the incoming light simultaneously, which we achieve by differentially rotating the cavity end caps and adjusting a half-wave plate while

¹⁶Linear Technology LTC1597

¹⁷The technique was first described by Balhorn et al [78].




Figure 3.12: Structure of the scanning cavity. The spacer (green) is machined from a solid cylinder of brass with diameter 6cm. The mirrors (blue) are supported on two piezo transducers (white). The piezos are held in finely-threaded endcaps (purple) which screw into the cavity spacer, allowing adjustment of the cavity length to a precision of $< 50\mu\text{m}$. Lock rings (dark green) hold the endcaps in position. The AR-coated windows (red) and thread-cover sleeves (orange) are designed to provide a hermetic seal once epoxied in position so that the cavity can be evacuated via the KF flange (cyan) fitted to the top.

observing the structure of the transmitted light. We do not currently understand the source of this effect.

Mechanics

The exploded scanning cavity structure is shown in Figure 3.12. The scanning cavity lock will compensate for thermal expansion, so it was not necessary to machine the spacer from a low expansion material such as Zerodur or quartz. Brass was chosen as it is easy to machine and will hold a fine thread for accurate adjustment of the cavity length. The dimensions of the cavity (diameter 6cm) and the fact that it is milled from solid brass make it massive and very rigid. While these are of course desirable properties for an optical cavity, such substantial structure is unnecessary for this basic system and could be replaced with a more lightweight alternative with little ill effect. The cavity was originally designed to be evacuated to prevent dispersive drifts due to atmospheric variations. Unfortunately the home-made piezo voltage feedthroughs leak, preventing the current system from being evacuated. Fixing this will probably require us to design new endcaps for the cavity.

3.5.3 Repump laser system

The scanning cavity system used for locking the 866-nm and 854-nm lasers is almost identical to that used for the cooling lasers. The main difference is that in the case of the repumpers, the feedback from the transfer lock is applied directly to the ECDL piezos, rather than to the piezo of a second (PDH) cavity. I designed a simple electronic circuit




Figure 3.13: Cross-sectional view of the assembled scanning cavity.

to provide manual control and external feedback to the piezos of up to three lasers, one channel of which is currently unused. Three pairs of front-panel potentiometers provide manual course and fine adjustment of the piezo voltage over ± 40 V, and are used for initial tuning of the laser frequencies. An external input for each channel allows a feedback signal to be applied when locking the laser; each input can be toggled on or off with a front-panel switch. The feedback signal is passed to a circuit with variable-gain proportional and integral stages based on TL074 op-amps, the output of which is combined with the manual pot voltage in a TL071 summing amplifier. An internal potentiometer allows the internal setpoint to be trimmed so that it matches the zero-error voltage of the external signal, allowing the system to be adapted to other feedback sources in the future. A schematic of this piezo lock system can be found in Appendix D.

3.5.4 Laser lock evaluation

The stability of an oscillator is generally measured in one of two ways. The linewidth measures the spread of the spectral intensity of the beam. The Allan variance measures the variance of the frequency when sampled over a varied period of time. Optical frequency oscillations are too fast to measure directly but can be mixed down by combining with a second oscillator of similar frequency to give a beat note at the difference frequency. Alternatively, the laser can be sent into an optical cavity and the frequency measured against a reference, as in the scanning cavity system.

397-nm lock performance

To measure the stability of the 397-nm locks I have used both of these techniques. By tuning the 397-nm lasers to frequencies approximately 100 MHz apart and combining the beams

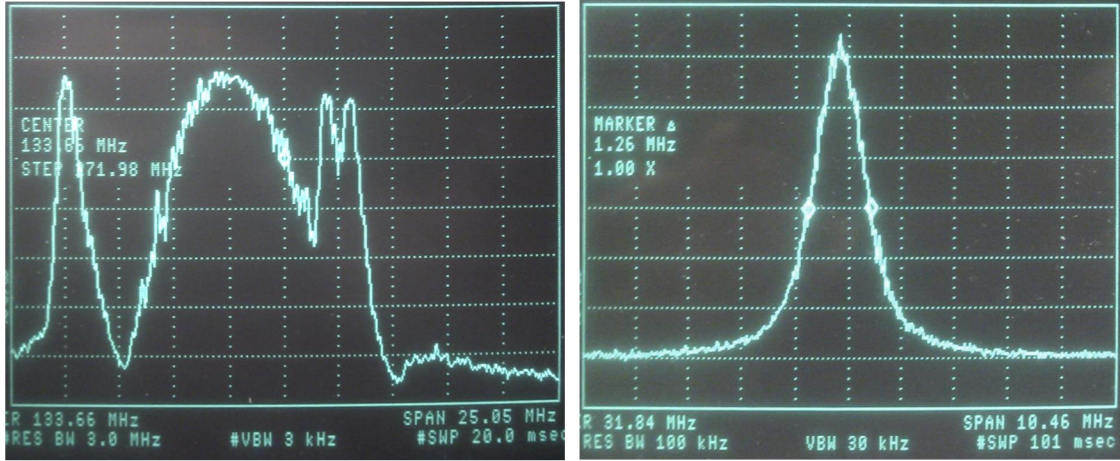


Figure 3.14: Spectrum analyser traces for a 397-nm beatnote, with the lasers unlocked (left) and PDH locked to the low-finesse cavities (right), reducing the single laser linewidth from ~ 10 MHz to < 1 MHz. Note the change of sweep time and span.

on a photo-diode, a beat note is produced which can be measured with a spectrum analyser, oscilloscope or counter. The spectrum analyser and Fourier transform of the oscilloscope provide a linewidth measurement. Figure 3.14 shows a spectrum analyser sweep with the lasers unlocked (left panel) and after PDH locking (right panel). Before locking, the beat note frequency varies over a range of 15 MHz during the ~ 12 ms that the 20 ms analyser sweep interacts with it. The fact that the peak is fragmented suggests that some of the noise processes contributing to this width are not much higher frequency than the 50 Hz sweep rate - consistent with acoustic noise on the ECDL mechanics. After locking, this low frequency noise is completely removed, and the overall linewidth of the beatnote is reduced to 1.26 MHz measured over the ~ 20 ms interaction time of the 100 ms sweep. Assuming that the noise properties of the lasers are identical, the linewidth per laser when locked is then < 1 MHz. By varying the sweep time and observing how the linewidth changes, we can learn how the noise processes and lock work on different timescales. A plot of this type is shown in Figure 3.15 where the linewidth given is the assumed single laser linewidth. With the lasers unlocked the measured linewidth rapidly increases with increasing sampling time, while when locked it remains around 1 MHz for sampling times approaching 1 s. Note that the locked and unlocked linewidths approach a similar value for very short sampling times, as the PDH lock has limited bandwidth.

The Allan deviation¹⁸ arguably provides more information about the nature of the laser noise, but can only measure noise processes that are slower than the frequency sampling time. To measure the Allan deviation, the photo-diode beat note signal was sent to a microwave counter (EIP 545A) and the reading from the counter logged via GPIB every 100 ms for 12 hours. The Allan deviations shown in Figure 3.16 show the performance of

¹⁸The Allan deviation and Allan variance are related as $\sigma_y = \sqrt{\sigma_y^2}$, in the same way as the standard deviation and variance.

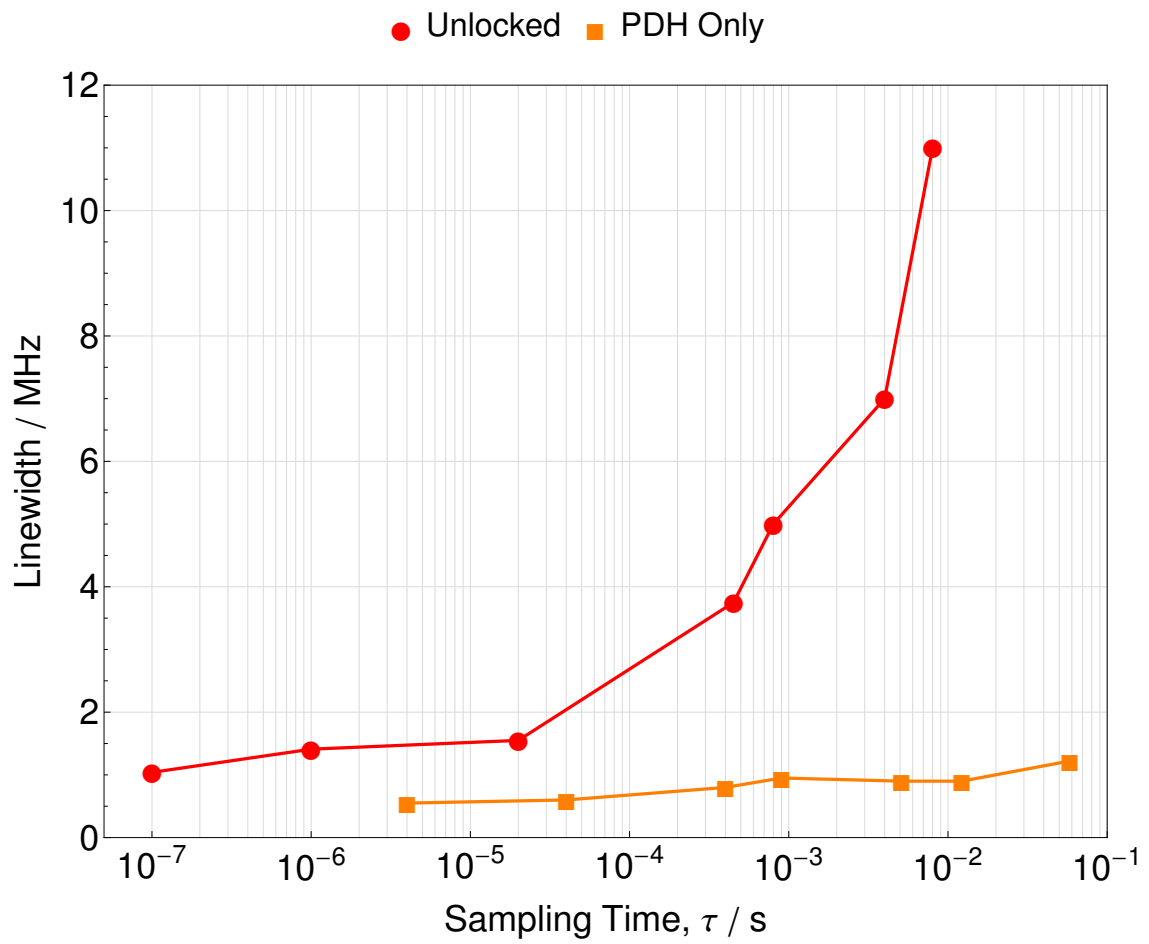


Figure 3.15: Linewidth of 397-nm lasers measured from two-laser beatnote. Red data are the unlocked laser linewidth, orange data are the linewidth when PDH-locked.

the laser while unlocked, PDH locked only, and with both locks operational.¹⁹

It is clear that PDH locking the laser drastically reduces the frequency noise on timescales up to about 10 s, as was seen through the linewidth measurements. However, over longer timescales, thermal drift of the cavities cannot be prevented and the frequency becomes little more stable than that of the unlocked laser. Using the scanning cavity lock to prevent this long term drift of the cavity length works, maintaining the Allan deviation below 100 kHz over timescales of many hours. However, it is clear that on timescales of up to ~ 10 s the scanning cavity actually increases the frequency noise. This noise remains significantly lower than the fast linewidth of the laser and is not a major problem but the technical performance could probably be improved by reducing the feedback bandwidth of the scanning cavity system. Note that the measured Allan variances of the locked lasers are much lower than the linewidth measured in Figure 3.15. This is because the major contribution to this linewidth is likely to be noise at frequencies of > 1 MHz which does not vary between even the fastest Allan variance sampling periods shown here and is therefore not detected.

Because we are measuring the difference frequency between two lasers, no common mode noise will be detected. Because the lasers lie on the same optical table and their PDH cavities share a single vacuum can, it is likely that some of the noise will be correlated. In principle, the scanning cavity can be used to measure this noise, but because it is part of the feedback loop, does not provide an independent measurement of the performance. Furthermore, certain processes that can lead to drifts, such as atmospheric pressure variations, will affect both lasers identically while locked to the scanning cavity. Our best means of detecting absolute frequency drift remains the ion itself. At the beginning of an experiment, the laser frequencies are set by performing crude spectroscopy using the step control box. Repeating this spectroscopy process after several hours, we sometimes observe (predominantly common-mode) drifts of several MHz. As I have derived in [77], the expected frequency variation due to dispersion shifts with atmospheric pressure is $\sim 5 \text{ MHz mbar}^{-1}$ (with 24 hour pressure variations of $\sim 10 \text{ mbar}$), so such shifts are to be expected. To prevent this drift reliably there is no option but to evacuate the transfer cavities.

866nm/854nm lock performance

The performance of the 866-nm and 854-nm laser locks has not been measured directly. As we only have one laser at each wavelength it is not possible to produce a measurable beat note. It would be possible to measure the position relative to the HeNe with the scanning

¹⁹Note that the unlocked data is not measured with a beat note, but instead uses the measured offset from the HeNe reference laser, as measured by the scanning cavity system. This allows the measurement to be extended to much smaller sampling times while the two techniques have been shown to produce nearly identical results over longer sampling times. The scanning cavity technique is not appropriate for measuring the performance while locked as the system will by default give zero variance while locked.

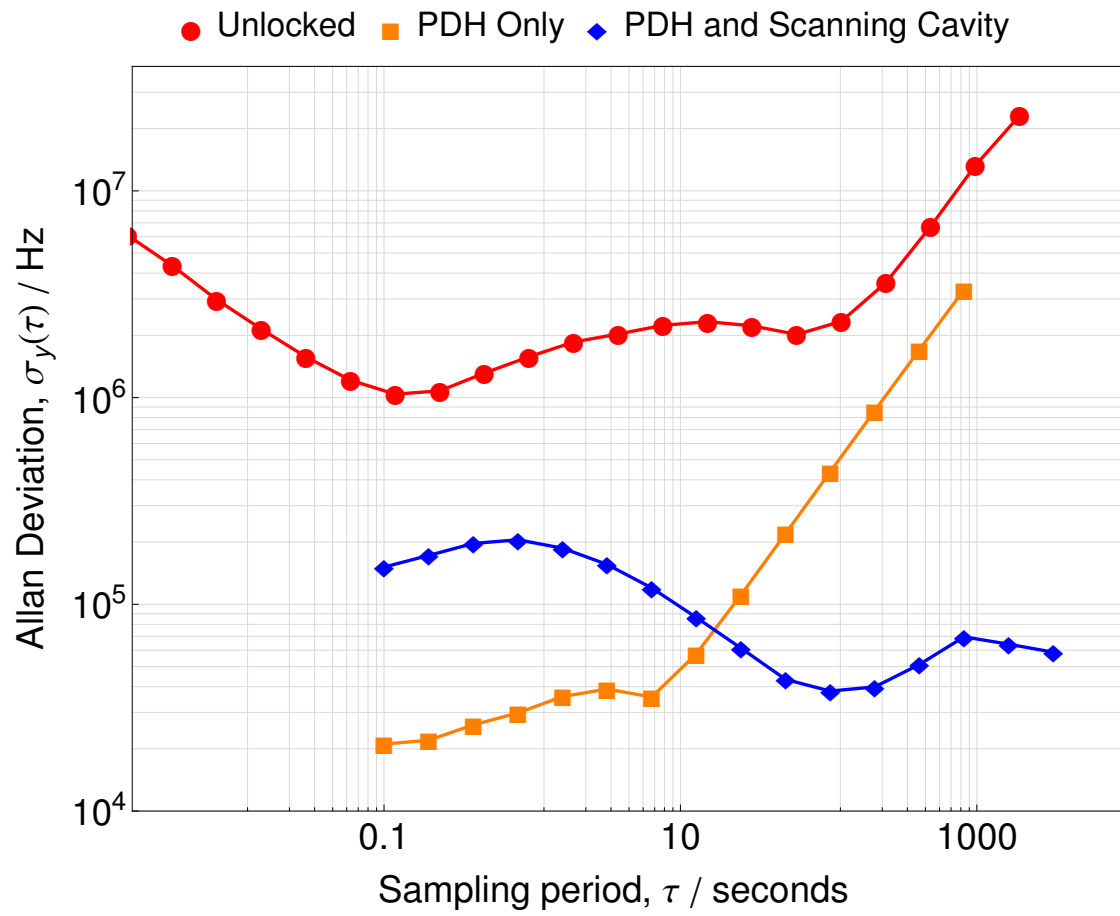


Figure 3.16: Allan deviations for 397-nm laser while unlocked, while PDH-locked only, and with both locks in operation.

cavity but, as discussed above, this will remain close to zero while the lock is operating successfully. From a practical point of view, we observe no noticeable drift on this system throughout the day. This is partly because our means of identifying such a drift are much more crude. The expected drift is also an order of magnitude smaller than with the 397-nm transfer lock as the dispersion profile of air between the visible and near-IR is not as steep as between the visible and near-UV, and the daily frequency variations will not be significant for our experiments.

3.5.5 Fibre EOM

The system for modulating the repump lasers is built around an EOSpace PM-0K5-20-PFA-PFA-850 fibre-pigtailed waveguide electro-optical modulator, which has a modulation bandwidth of 20 GHz. The modulator is driven with two frequencies simultaneously, at ~ 10.4 GHz and ~ 1.7 GHz, with the modulation indices set to maximise the second and first order sidebands respectively. Because the differential (first order) Zeeman splittings of the transitions at 866 nm and 854 nm are related by straightforward integer ratios, the comb of sidebands produced when both lasers are sent through the same modulator is able to drive all necessary repump transitions simultaneously. The technique and apparatus involved has been described in detail in the thesis of S. Donnellan [66].

One important point that has not been covered in these treatments is the significance of the second-order Zeeman shift, as derived in 3.2.1. The second order shift breaks the symmetry of the differential transition splittings around the unperturbed value, as shown in Figure 3.17. Considering the σ transitions at 866 nm at 1.85 T, the $D_{3/2,\pm 3/2} \rightarrow P_{1/2,\pm 1/2}$ transitions experience a differential second order Zeeman shift of 36 MHz, while the $D_{3/2,\pm 1/2} \rightarrow P_{1/2,\mp 1/2}$ transitions experience a shift of 65 MHz. This means that the pair of σ_+ transitions is split by 29 MHz less than in the first order approximation, while the pair of σ_- transitions is split by 29 MHz more. This makes it impossible to address all of the transitions simultaneously with a regularly spaced comb of sidebands, as is produced by the EOM. The effect for the 854 nm transitions is about three times larger.

Fortunately, at this field strength the degree of detuning from the transitions is quite slight and these problems can be overcome by simply applying more power to the lasers. At high fields, this will become an increasingly problematic issue and may necessitate the use of more individually tuneable source lasers. This behaviour also makes precise tuning of the repump lasers difficult, as there is not one resonant frequency but two or more, probably ruling out EIT techniques using 397-nm/866-nm dark states with the current apparatus.

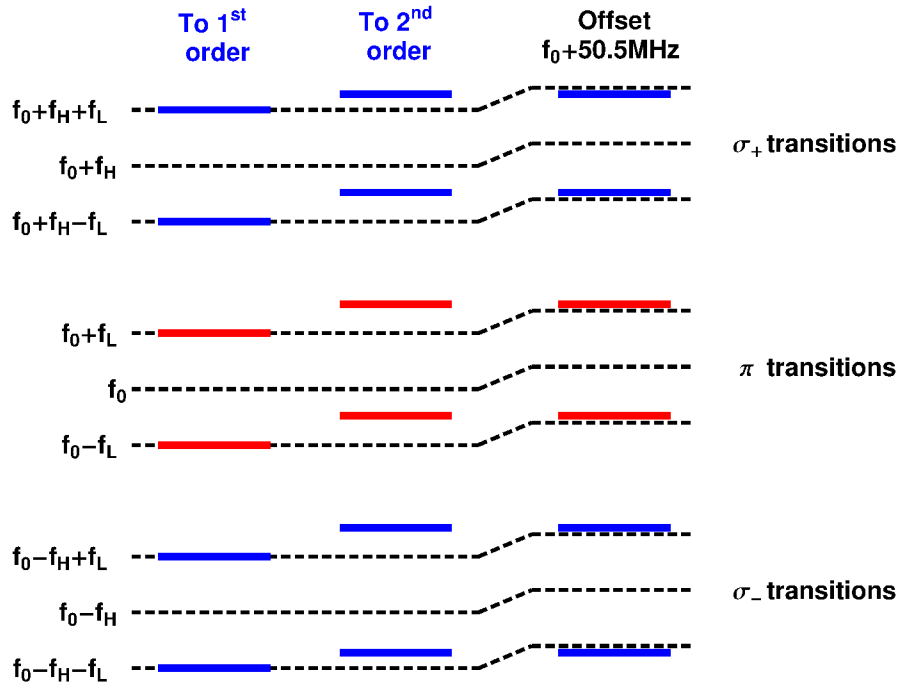


Figure 3.17: Effect of second order Zeeman splitting on 866-nm transitions at 1.85 T (not to scale). The transition frequencies at relative to the zero-field value f_0 are shown in blue (σ_{\pm} transitions) and red (π transitions). The dashed lines represent the even comb of frequencies produced by the fibre-EOM when driven with two modulating frequencies, $f_H \sim 10.4\text{GHz}$ and $f_L \sim 1.7\text{GHz}$. The second order Zeeman effect breaks the symmetry of the comb around f_0 . Shifting the centre frequency can compensate for some of the shift but the transitions can no longer all be addressed perfectly at the same time.

3.6 729-nm Laser, Locking and Amplifier Systems

Calcium provides a transition at 729 nm with a linewidth of ~ 1 Hz that is ideal for sideband-resolved spectroscopy, cooling and coherent manipulation. However to make use of such a transition requires a laser of suitably low linewidth and high stability. Quite how exacting these requirements are depends on the exact application. By definition, any work in the strong binding regime requires the linewidth to be much lower than the trap frequency, which in our experiment is usually a few hundred kHz. For sideband cooling or coherent manipulation, the tolerable level of detuning from resonance depends on the Rabi frequency and the condition $\delta \ll \Omega$ must be met to ensure effective excitation. It follows that linewidth and power output of the laser cannot each be considered in isolation when setting the specifications for such a laser system.

The 729-nm laser in our laboratory was originally built by D. Crick, who developed the main laser and locking system. The original system performed well, but over time two major modifications became necessary, and my work to implement these is the subject of this section. The first was to remove the residual thermally-induced drift from the laser frequency, which involved building a new cooler and temperature controller for the reference cavity (Section 3.6.2). The second was to build a tapered amplifier system to boost the power by over 10 dB (Section 3.6.3). Both removing the drifts (see also Section 3.4.2) and increasing the Rabi frequency were essential in reaching a level of stability that would allow us to reliably drive a given sideband without constant recalibration.

3.6.1 Overview

The system at the time of writing is shown in Figure 3.18. The master laser is an external-cavity diode laser (ECDL) in Littrow configuration, driven by a current of ~ 70 mA and providing an output power of ~ 20 mW. This laser is stabilised with a Pound-Drever-Hall lock to a cavity with a finesse of $\mathcal{F} = 60000$ and linewidth of 25 kHz. The stability of this cavity is the main limitation on the stability of the laser, and Section 3.6.2 describes my work to improve this. The PDH technique itself is described in detail in numerous sources [74, 79] and the details of its application to this system are covered in [56].

After optical isolation, a portion of the light is sent to the PDH lock and frequency modulated with an electro-optical modulator at 10 MHz to produce the necessary sidebands, before being sent into the reference cavity. The reflected light from the reference cavity is mixed with the RF modulation signal and low-pass filtered (‘loop filter’ in Figure 3.18) to generate the PDH error signal. This signal is used to provide fast feedback to the laser diode current via a bias-tee. The signal is also fed to an Arduino microcontroller which provides slow feedback to the ECDL grating piezo voltage via a 16-bit DAC. The Arduino also monitors the amplitude of light transmitted through the cavity, as measured by a photodiode at the far end. If this drops below a threshold value (indicating the laser has

become unlocked), the Arduino stores the integrated piezo voltage and instead scans the piezo at ~ 150 Hz in an attempt to recover the lock, at which point normal operation resumes. With this system in place the lock is very robust and tolerates substantial sharp knocks to the enclosure or optical bench or having the beam blocked for seconds at a time. The bandwidth of the lock is approximately 2 MHz.

The power of the 729-nm laser is then boosted with a second 729-nm slave diode, which is injection locked with the stabilised master diode light. The master light is injected in the side-port of an optical isolator placed in front of the free-running diode. Because this diode is not placed in an ECDL, the light injected into the front facet will dominate the gain, provided that the free-running wavelength is close to the injected wavelength (this is adjusted with the current and temperature of the slave diode). The status of this injection lock is monitored by sending mixing light from the slave with pre-injection light from the master on a photodiode and observing the beat signal. For balanced amplitudes the beat note should appear at DC and zero amplitude, exhibiting phase-sensitive excitations if the relative path length of the master and slave sources is minutely varied (e.g. by gently tapping the breadboard). The injection lock boosts the power of the laser by a factor of three, from approximately 8 mW pre-injection to 24 mW post-injection. In the original system, this light was then sent through the double-pass AOM used for switching and frequency-sweeping of the laser and through a single mode fibre to be combined with the other lasers. The final power at the position of the ion was approximately 3.3 mW. Due to the relatively large beam waist at the trap this led to unacceptably low Rabi frequencies and it was necessary to build a further stage of amplification in the form of a tapered amplifier, which boosted the power at the ion to ~ 40 mW. The tapered amplifier system is described in Section 3.6.3.

3.6.2 ULE cavity

The high-finesse cavity that provides the reference for the PDH lock is the most critical element in the operation of a narrow-linewidth laser. Even if (in fact, particularly if) the PDH lock is perfect, if the cavity length is not stable then nor will be the laser frequency. The most important factors when considering the linewidth²⁰ of the laser are acoustic vibrations of the cavity and any optical fibre the laser passes through (where acoustic stresses modulate the refractive index of the fibre and introduce phase noise), as well as microscopic Brownian fluctuations of the cavity mirrors. For groups trying to reach to the lowest linewidths for metrology purposes – at the level of 1 Hz or less – these are the most critical considerations and great efforts are made to optimise mirror materials and

²⁰Linewidth is only a meaningful description of laser noise if combined with a timescale over which it is measured. A spectroscopic measurement of the linewidth takes approximately 15 seconds. The linewidths later extracted from the coherence time of a Rabi oscillation are measured over a fraction of a millisecond. For simplicity, any effects that manifest themselves over timescales longer than the measurement of a single sideband are referred to as ‘drifts’.

design cavities isolated from acoustic vibrations. However, early trials with our laser system demonstrated that the linewidth was under 1 kHz using only crude vibration isolation in the form of a bicycle inner tube placed under the 729-nm system breadboard. This is sufficiently narrow for our purposes and no further measures have been employed to reduce this further.

Of much greater concern to our work were the slow frequency drifts that plagued our early experiments. These made it impossible to record accurate and repeatable spectra, let alone to effectively drive a single sideband for cooling or coherent manipulation purposes. The biggest contribution to the drifts was the movement of the trap within the field (Section 3.4.2) but with this corrected, drifts in the cavity length become important.

The cavity in our experiment is made from Corning Ultra-Low Expansion (ULE) glass and manufactured by REO²¹, with a design finesse of 250 000 that has since reduced to $\sim 60\,000$. ULE glass is guaranteed to have a coefficient of thermal expansion (CTE) of $0 \pm 30 \text{ ppb}/^\circ\text{C}^{-1}$ between 5°C and 35°C , exceptionally stable by any standard and over 200 times lower than BK7. This performance is considerably enhanced by the fact that the CTE of ULE glass passes through zero at a point close to room temperature. The glass expands as the temperature is either increased or decreased from the ‘zero-crossing’ point and at this temperature, T_{zc} , the cavity exhibits no thermal response to first order.

The value of T_{zc} depends on the precise mix of the glass constituents and the conditions in which it cools, and varies slightly from batch to batch. Generally the zero-crossing temperature lies between 5°C and 35°C but some samples have been produced with T_{zc} significantly lower [81]. This value has considerable practical implications for the complexity of the cavity temperature stabilisation system. For temperatures a couple of degrees or more above room temperature, the lock may operate via heating alone [82], while for T_{zc} at or below the ambient conditions a bipolar temperature controller is necessary. For temperatures below the dew point (typically 12°C to 16°C depending on atmospheric conditions) condensation presents a further problem and the cold elements of the apparatus have to be held in a dry, hermetically sealed environment or placed under vacuum [83].²² In most situations, a value of T_{zc} above room temperature is very desirable.

Initial temperature dependence

The original design for the 729-nm laser placed the ULE in a vacuum chamber, continuously pumped by an ion pump and temperature controlled with a monopolar heater and proportional-integral (PI) lock. By changing the lock point while monitoring the AOM frequency necessary to compensate for the frequency shift, the data in Figure 3.19 were pro-

²¹www.reo.com For more details on this cavity and PDH locking in general, see the thesis of J. Sudbery [80].

²²Alternatively, composite cavities can be constructed with mirrors, of e.g. fused silica, mounted in such a way that the difference in CTE leads to a shift in the effective zero crossing [84], bringing this above room temperature.

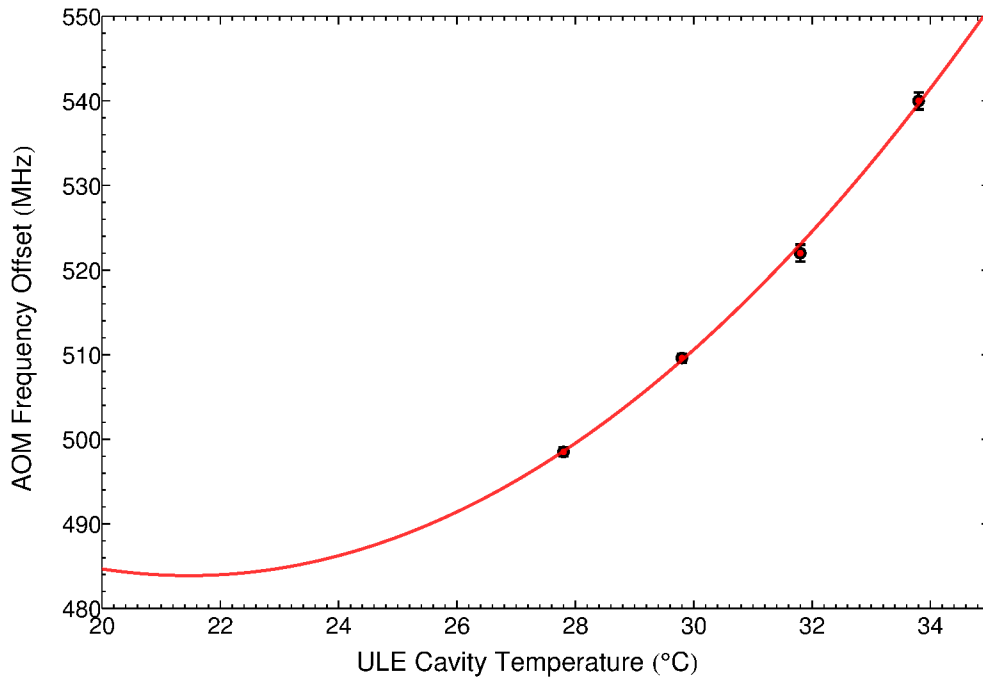


Figure 3.19: Temperature dependence of ULE cavity resonance frequency, given by the AOM offset frequency needed to keep the 729-nm laser resonant with a calcium transition. Decreasing offset frequency corresponds to decreasing cavity length.

duced²³. The fitted line gives a frequency temperature dependence of $df/dT = 5 \text{ MHz K}^{-1}$ at 28 °C. The original temperature lock had a stability of 2 mK, leading to thermal drifts of up to 10 kHz, higher than could be tolerated for any of our applications, and further cooling was clearly necessary.

We were faced with a dilemma: if it turned out that the cavity needed to be cooled well below the dew point a heat-shielded vacuum system would be necessary which would be complex, expensive and would have required rebuilding the rest of the laser optical system; if on the other hand the zero crossing point lay only a little below room temperature a much simpler unsealed system would work. The fitted line suggested a zero-crossing temperature of 21.5 °C but this was very poorly constrained and we knew the real value could deviate greatly from this. We found some evidence that other cavities produced by REO in the same year (such as some used in experiments at the National Physical Laboratory) had zero crossing points of between 14 °C and 17 °C and decided to gamble on producing the simple system – a gamble that paid off, though by a rather narrow margin.

Cavity cooler design

The design we settled on could not be recommended to anyone trying to cool a cavity with unknown T_{zc} for the reasons described above. However, while crude, we have found

²³Note that the discrepancy between these values and those given in [56], which are the same data, is due to a electronic thermometer which had not been calibrated. These are the correct values.

it to provide very effective performance when locking the cavity at temperatures $> 15^{\circ}\text{C}$. We have also used it to lock the cavity at $T < 10^{\circ}\text{C}$ for test purposes, but the rapid accumulation of condensation on the cold elements make it unsuitable for long term operation at such low temperature.

The principle behind the design was to retrofit the existing vacuum chamber (centred around a CF63/45 4-way cross) with an aluminium and copper superstructure that would provide good thermal contact to a powerful thermoelectric cooler (TEC). The superstructure consists of three aluminium risers machined with circular apertures designed to clamp flushly around the three CF63 flanges of the vacuum chamber, which provides the largest regular surface for thermal contact. The three risers are connected below the vacuum system by an acrylic base plate, which provides good mechanical stability while providing electrical isolation from the breadboard. A further layer of neoprene rubber is placed between the baseplate and breadboard to improve thermal isolation. The baseplate is mirrored by a oxygen-free copper cold plate that connects the risers above the cavity, and a large TEC is situated in the middle of this plate. The sides of the TEC device are potted to avoid vapour ingress that could cause problems when operating below the dew point. Thermal contact between component surfaces throughout the design is improved with thermal compound. The mechanical structure of the cooler and how it attaches to the vacuum system is shown in Figure 3.20.

This cooler necessarily emits a considerable amount of heat and will only work efficiently if the TEC hot plate is kept as close to ambient temperature as possible. The existing 729-nm laser system was housed in a large particle board enclosure finished in a fetching hardwood-effect laminate trim²⁴, to prevent air currents and acoustic noise from reaching the laser and cavity. Dumping the localised heat from the TEC directly into this volume would introduce internal air currents and prevent efficient cooling, so it was deemed necessary to dissipate this heat outside the enclosure. This is achieved by extending the risers to sufficient height that the cold plate lies just below the enclosure lid. A passive, heatpipe-based CPU heatsink²⁵ connected to the TEC hot plate can then protrude through a close fitting aperture in the lid. The majority of the heat is dissipated outside the enclosure without compromising the isolation from air currents that it was designed to provide. In practice, the system performance is improved by introducing some very gentle forced air cooling to the heatsink, via a low-noise, low-rpm 140mm fan²⁶ which draws air upwards through the heatpipe fins. To avoid introducing vibrations to the cavity, the fan is mounted on equipment strutting above the optical breadboard with no physical contact to the laser system.

In principle, this system could then be temperature locked to an error signal derived

²⁴IKEA Expedit shelving carcass, walnut finish.

²⁵Nofan CR-95C.

²⁶Noctua NF-P14.




Figure 3.20: Cross-section of cavity cooling system superstructure. The original vacuum chamber lies in the middle. One CF45 arm leads to the UHV valve on the right, the other leads to the left to a 55 l s^{-1} ion pump (not shown). The cavity sits in the main body of the vacuum chamber, resting on two Viton O-rings (not shown) for thermal and mechanical isolation. Three aluminium risers clamp to the CF63 flanges of the vacuum canister and are connected above and below by copper and acrylic plates respectively. The copper plate provides thermal coupling to a TEC, which dissipates heat through a large heatpipe heatsink.




Figure 3.21: Cross-section of cavity (taken orthogonal to that in Figure 3.20) showing TEC cooler and heatsink (top), heater tape surrounding inner vacuum can (middle, orange stripes) and internal and external insulation (textured white areas). The ULE cavity is shown in the centre and windows on the left and right conflat flanges (cyan) provide optical access through the cavity to the CCD camera and photodiode beyond. Two further images of the complete system are provided in Appendix E.

from the resistance of a thermistor placed near the cavity. Unfortunately, the considerable thermal distance between the TEC and the region to be stabilised – approximately 40cm of copper, aluminium and stainless steel and the joins between them – severely limits the bandwidth of a lock due to the time it takes for the heat to conduct over this distance. Better performance can be achieved by using a differential heating and cooling system, with the addition of a length of silicon rubber heater tape²⁷ coiled directly around the vacuum tube containing the cavity. This system allows the cooled superstructure to provide a well-coupled cold sink against which the heater can act differentially. Because the heater is much closer to the cavity itself, it provides a much higher bandwidth of lock than the TEC could alone. Further, because the TEC provides a pre-stabilised environment, the input noise to the heater lock is much lower, improving the output performance in turn.

While the cavity chamber itself is evacuated, the cooling superstructure remains exposed to the laboratory environment, and we compensate with a large amount of insulation, as shown schematically in Figure 3.21. The central vacuum system is first tightly wrapped in thick neoprene lagging, fastened with cable ties, before the remaining voids are filled with shaped neoprene and blocks of closed-cell foam. The outside of the system is fitted tightly with closed-cell foam on all sides, leaving only a small aperture to the window at each end for optical access to the cavity.

Both the heating and cooling systems are locked using home-built PI controllers, with output stages provided by a TIP142 bipolar transistor. The output stage of the heater controller is a straightforward emitter-follower, while the TEC controller uses the TIP142 (with emitter sense resistor) in the feedback loop of a CA2372 op-amp to provide current sensing. This provides more linear control of a device such as a TEC, where the resistance varies with thermal load and temperature. The TEC system (and, by definition, the heater) is monopolar but this provides adequate performance due to the large temperature gradient between the room and cooler structure. Each error signal is generated by a thermistor circuit supplied a REF102 voltage references and measured with an INA129 instrumentation amplifier. An internal potentiometer allows the set point temperature to be adjusted. The thermistors were calibrated against a mercury thermometer in the laboratory and found to be identical; this calibration is used to convert set point voltage to temperature.

The TEC system is locked to a thermistor embedded in one of the aluminium risers, and cools the superstructure to a temperature somewhat lower than the target cavity temperature. The high ratio of thermal conductivity to volumetric heat capacity of the aluminium and copper, combined with the high cooling power (~ 200 W) of the TEC mean that this lock responds rapidly to perturbations. The temperature at the position of the thermistor stabilises to $\Delta T \pm 2$ mK within a couple of hours. With the available TEC supply current (5.5A) the cooler is capable at holding the entire system at temperatures of less than 5 °C, with the performance likely limited by heat ingress through the base plate and

²⁷Omega SRFG-142/2.

ion pump, which could not be effectively insulated with the system in its current location. The high thermal mass and low surface area and thermal conductivity of the ion pump help prevent this from coupling thermal noise to the insulated interior on anything but the longest timescales.

In normal operation, the copper cold plate is locked to a temperature of $\sim 14^\circ\text{C}$ which is generally slightly below the dew point, depending on the atmospheric conditions. Left unchecked, this causes a slow but steady build-up of condensation around the TEC which provides a thermal bridge between the hot and cold plates and prevents the cooler from working. Our solution for preventing this is very low-tech: a square of cotton material²⁸ placed over the cold plate and cut to fit closely to the TEC wicks moisture away from the insulated interior to a frilly edge which drapes over the exterior faces in a manner akin to a Valance sheet. Moisture continuously condenses on the cold surface but as droplets of liquid water form they are drawn away by capillary action to the edge, where they can re-evaporate.

The heater system is locked to a network of 3 thermistors spaced around the vacuum canister exterior and heats the central region of the system differentially against the stabilised ‘environment’ provided by the TEC cooler. In this configuration, locking requires a differential of approximately 1°C . Due to the low thermal conductivity of the stainless steel, the time constant of this lock is considerably higher and equilibrium is reached after approximately 24 hours. The final stability was measured to be $\Delta T \pm 0.2\text{ mK}$, an order of magnitude improvement on the performance of the original heater lock. The temperature of the cavity itself can only be inferred from spectroscopic measurements, which suggest that it reaches equilibrium considerably slower than the vacuum system around it - we have observed continuous drifts for up to two days after a major temperature change. This is to be expected as the cavity sits on a pair of Viton o-rings which provide further thermal isolation from the vacuum canister.

Measurement of T_{zc}

We used the new cavity temperature controller to extend the data of Figure 3.19 to temperatures as low as 12°C . These data are plotted in Figure 3.22, where the fit gives a zero crossing temperature of $T_{zc} = 17.4 \pm 0.5^\circ\text{C}$, consistent with that observed in the ‘sister’ cavities that had been used at NPL. With the range of recent improvements to the experiment it would now be possible to constrain this value far more precisely, but this has not been necessary as the expected drifts are now considerably lower than the linewidth: at 18°C , $df/dT = 0.3\text{ MHz }^\circ\text{C}^{-1}$, which for a temperature lock of $\Delta T \sim 0.2\text{ mK}$ leads to drifts of $\Delta f \sim 60\text{ Hz}$.

It is interesting to note that if the cavity is brought within 1 mK of the zero-point (once

²⁸Generously donated by Thorlabs.




Figure 3.22: Temperature dependence of ULE cavity resonance frequency, given by the AOM offset frequency needed to keep the 729-nm laser resonant with a calcium transition. Decreasing offset frequency corresponds to decreasing cavity length. The cavity length can be seen to increase either side of the ‘zero-crossing’ temperature $T_{zc} = 17.4^\circ\text{C}$, at which the cavity exhibits no thermal expansion to first order.

this has been identified to this precision), the frequency drifts due to thermal expansion will be $\sim 0.25\text{ Hz}$. This is a fractional variation below 10^{-15} , equivalent to a length variation of significantly less than the charge radius of a proton over a cavity of length 0.1 m . Naturally this refers to the length of the cavity averaged over the area of the mirror – thermal fluctuations will cause local displacements that are far higher – but it is still a remarkable number, made only the more so by the fact that such stability is now almost routinely achieved by groups in standards labs around the world [85].

ULE creep contraction

While its thermal sensitivity is incredibly low, ULE exhibits a continuous creep as the material ages, typically contracting at a rate of $\frac{\Delta L}{L} \sim 1.5 \times 10^{-16}\text{ s}^{-1}$ [86]. This process is predicted to occur for all glasses due to vibrational relaxation effects [87], but is only a significant problem when thermal expansion is sufficiently low, as in ULE. The theory in reference [87] predicts that the aging rate is lower at lower temperatures and that ULE samples with a lower T_{zc} will therefore exhibit a slower rate of contraction when monitored at this temperature. A recent measurement [81] of the contraction rate of a sample of ULE with unusually low $T_{zc} = -3.3^\circ\text{C}$ found a creep rate of $-5.5 \times 10^{-17}\text{ s}^{-1}$, the lowest yet

observed, providing support for the theory. Assuming a laser frequency of 410 THz one therefore expects to observe a drift of between 2 and 10 kHz/day.

Since suppressing more significant variations in the system (such as movement within the field) we are yet to observe any such distinctive linear drift in our system, though the periods of observation without other perturbations have been limited and there has been no concerted effort to make such a measurement. One contributing factor may be the age of our cavity, which was manufactured c. 2000, nearly 15 years ago. Both the principle of the vibrational relaxation mechanism and some anecdotal evidence [88, 89] suggest that the creep rate of low expansion glasses such as ULE and Schott Zerodur reduces with a time constant of several years and it is certainly likely that the drift rate of our cavity spacer is lower than it was immediately after production. More importantly, the reference against which we measure the laser frequency – that of the calcium transition we probe – also exhibits drift, due to the slow decay of the magnetic field. This rate is not precisely known but the manufacturer specifications estimate a fractional decay of $\sim 10^{-7}$ /day. Depending on the transition, a decrease in the Zeeman splitting may increase or decrease the transition frequency. In the case of the $S_{1/2,-1/2} \rightarrow D_{5/2-3/2}$ transition used in recent experiments, the frequency varies with the field as $\Delta f_0 \sim -11 \text{ GHz T}^{-1}$. The field decay would therefore be expected to introduce a drift of -2 kHz/day , similar in magnitude and sign to that of the laser and leading to at least partial cancellation of the observed effect.

Cavity mode considerations

The 729-nm AOM has an RF bandwidth of $\sim 50 \text{ MHz}$ with good first order diffraction efficiency, leading to a tuneable range of 100 MHz after double pass. The ULE cavity free spectral range is 1.5 GHz and so it is necessary to lock to a variety of cavity transverse modes to span the frequency spectrum with the AOM. With a length of $L = 0.1 \text{ m}$ and symmetric mirrors with radii of curvature $R = 0.35 \text{ m}$, the cavity g parameter is $g = 1 - L/R = 0.714$, leading to a transverse mode spacing of 369 MHz. Unfortunately this is not far from precisely $1/4$ of the free spectral range and higher order modes do not efficiently span the free spectral range. Over 70% of the frequency spectrum is not covered by the first four transverse modes while 6% can only be reached with a transverse mode of order 40 or higher.

Locking to high order transverse modes is quite possible, but the mode matching conditions become ever harder to meet, reducing the spectral prominence of that mode compared to its lower-order neighbours. This makes locking to such modes relatively unstable as small perturbations may move the laser into the capture region of a stronger mode. This issue is of practical importance when choosing a transition for spectroscopy, as some are conveniently locked to a low order $\text{TEM}_{0,2}$ mode while we have only succeeded in driving some while locked to $\text{TEM}_{0,\sim 21}$ modes, which are quite impractical for everyday use. An AOM with higher RF bandwidth would ameliorate this situation.

3.6.3 Tapered amplifier

Principle

A tapered amplifier (TA) is a semiconductor optical amplifier consisting of a region of gain medium with AR coated input and output facets, into which seed light can be injected and amplified as it traverses the chip. The 'tapered' element describes the fact that the gain material flares transversely (in one dimension) over the length of the amplifier to ensure that the optical power density does not reach a level where it can damage the AR coating of the output facet. The narrow tip of the taper must remain of the correct size to couple to a ridge waveguide that forms the input to the chip and ensures that the light entering the tapered section is spatially single mode. The combination of these two features leads to better beam quality than would be possible with broad-area diode lasers, while offering output powers much higher than a Fabry-Perot diode can provide. The AR coated facets mean that the TA cannot lase free-running and if driven without a seed will generate a large amount of amplified spontaneous emission (ASE). For this reason, any seed light of suitable wavelength will be perfectly amplified by the device, maintaining its spectral purity. Some residual ASE remains, but this is generally more than 30 dB below the signal and will be further suppressed far from the output due to its different spatial behaviour, especially if coupling into a single mode fibre.

Precautions

Due to the high optical powers involved tapered amplifiers can very quickly destroy themselves under the wrong conditions, but there are a few basic precautions that will prevent the majority of accidents. The most obvious is that, as with all diodes, the current limit must be respected. The diode also dissipates a significant amount of heat, and must be well heat-sunk and actively cooled with a TEC to avoid overheating. Unintended optical feedback is undesirable in most laser applications, but in tapered amplifiers it can be catastrophic, as light entering the broad front facet and being amplified towards the narrow input facet will easily reach intensities high enough to damage the AR coating and destroy the chip. Although a tapered amplifier will amplify even a small amount of seed light, it is best to maintain an input power high enough to begin to saturate the gain medium – this way any back-reflections into the front facet will find themselves competing for gain with a much larger amount of forward-propagating light and will not be amplified. For the same reason it is also necessary to use very good optical isolators, ideally double-stage, especially if the output beam is going to be coupled into a fiber or optical cavity, and for great care to be taken during alignment of any optics situated pre-isolator (which, of course, must also be AR coated). Finally, the high output power means that any particle of dust entering the beam near the output can be optically tweezed towards the highest intensity region. This draws them onto the output facet itself which can cause a localised build up of power and




Figure 3.23: Schematic depiction of input and output optics for a tapered amplifier. The beam width in the vertical axis is shown in red while the width in the horizontal axis is superimposed in dashed blue.

destroy the AR coating, so the chip must be well protected from the environment around it with appropriate covers.

Optics

Providing these rules are observed, tapered amplifiers are quite straightforward to use. Approximately 20 mW of seed light from a laser of choice is focused into the ridge-waveguide input, a similar process to coupling a single-mode fibre. Current is sent through amplifier chip and laser light emerges from the output, with virtually identical spectral properties to the seed but with 10 dB to 20 dB higher power. The main difficulty is dealing with the spatial properties of the output light, which is highly astigmatic, has differing horizontal and vertical divergence, and is emitted across a source considerably extended in one dimension (the width of the taper).

A simplified version of the optics is shown in Figure 3.23. The tapered amplifier used here is a Toptica TA-0735-0500-4, an Eagleyard chip with non-catalogue centre wavelength. Light is focussed into the input facet using a high-NA aspheric lens AL1.²⁹ The output is collimated using a second aspheric AL2.³⁰ The extremely high NA of the collimating lens is required to capture the light from the fast vertical axis which emerges from the front facet with a divergence half angle of $\Theta = 33^\circ$. The slower horizontal axis diverges with a half angle of $\Theta = 14^\circ$ and is much less broad by the time it reaches the aspheric, leading to an elliptical profile in the lens plane. The situation is complicated by the fact that there is ~ 0.6 mm of astigmatism between these two axes, with the slow axis source situated somewhat behind the fast. This means that if the aspheric acts to collimate the fast axis, it will focus the slow axis, and this will need to be corrected with a cylindrical lens, CL2. However we can take advantage of this fact to simultaneously correct the aspect ratio of the beam. The slow axis is allowed to come to a focus and diverge out beyond, and the cylindrical lens of appropriate focal length is placed at the position where both horizontal

²⁹Thorlabs C220-TME-B, $f = 11$ mm, $NA = 0.25$.

³⁰Thorlabs C330TME-B, $f = 3.1$ mm, $NA = 0.68$.

and vertical widths become equal.³¹

The width of the emitting region in the horizontal direction ($d = 190\mu\text{m}$) means that the slow axis is spatially multimode and the beam quality is much lower, with striping and other artefacts often visible. This width also causes problems for collimation using aspheric lenses, which are generally suited for on-axis use only, and in fact may perform worse than spherical lenses off-axis. The lens used in our setup is of particularly short focal length, which in hindsight was not the best choice as the apparent width of the source is so great, and indeed the slow axis beam quality is noticeably poor. The short focal length also makes the effects of the astigmatism more significant. In future, it would probably be better to use a lens of similar NA but much larger focal length to make the output facet behave more like a point source.

Electronics

The tapered amplifier requires a current supply and temperature controller. I have built both using off-the-shelf core modules built into custom systems, providing a reliable and cost-effective solution. The temperature controller consists of a Wavelength Electronics MPT2500 module, which provides 2.5A of bipolar drive to the TEC. The input to the temperature lock is provided by a thermistor embedded in the TA mount. As the TA does not rely on any parametric lasing or internal modes, the temperature must only be kept constant to avoid misalignment of the optics and the requirements on the lock are quite modest.

The current controller is built around a Thorlabs LD3000R module, which provides up to 2.5A of voltage-controlled current. This is mounted in a box with a front panel potentiometer and calibrated display controlling and monitoring the current output. An internal potentiometer allows the user to set a current limit. An important addition to the current controller is a custom latched interlock circuit which cuts the current if the temperature near the amplifier rises above a threshold. This ensures that if the temperature controller fails unexpectedly, the current will shut down before the chip can overheat. This consists simply of a second thermistor connected to the current controller, where a comparator monitors the voltage across it relative to the threshold value set by an internal potentiometer. If the threshold is not met, the comparator output drops and switches the state of a set-reset flip-flop, the buffered output of which opens a relay between the current module and tapered amplifier. Once triggered, the interlock must be reset manually with a front-panel toggle switch which returns the flip-flop to its original state, closing the relay.

The final part of the drive electronics is a small protection circuitry board which is attached to the TA mount immediately before the chip itself. This consists of a $1\mu\text{F}$ capacitor,

³¹Alternatively, the aspheric lens can be set to collimate the slow axis, which will leave the fast axis divergent. In this case, the fast axis will need to be focused with one cylindrical lens before being collimated with a second, negative cylindrical lens. The method described above is usually a better option.

a large Schottky diode and a 12V Zener diode, all placed in parallel with the TA chip with the diodes in reverse bias. The Zener provides over-voltage protection, the Schottky diode provides reverse current protection and the capacitor helps protect against fast, transient ESD events that the relatively slow Zener cannot respond quickly enough to protect against.

Mechanical design

My first attempt at producing a mount was based on the design described in [90], which had been used to varying degrees of success by several other experiments at Imperial. The design allows for the two aspherics to be moved along the optical axis with respect to the tapered amplifier, but does not provide any transverse travel. If the tapered amplifier output happens to be perfectly aligned with the axis, then this is a relatively satisfactory solution (although we also found that the optics had some play within the collimation tubes which made precise alignment harder). However, the tolerance of the emitter position on the tapered amplifiers themselves is surprisingly large, with as much as $200\mu\text{m}$ variance between individual chips. We found that our chip needed to be mounted significantly higher than was possible on that mount and as such the beam quality was particularly poor.

After repeated attempts to improve the alignment, the original chip was broken, either via mechanical impact or dust on the output facet. Rather than rebuild the mount to fit the next chip, I decided upon a new design that would allow complete 5-dimensional positioning of the input and output lenses for optimal coupling, as shown in Figure 3.24 and described in the accompanying caption. The key feature of the design is the lens mounting system which consists of no more than two vee-grooves into which the lenses can be epoxied. Each lens can be held on an arm attached to a 5-axis mount and adjusted to optimise the coupling through the isolator and into a fibre. With the lenses perfectly positioned, each can then be carefully epoxied in place, permanently fixing them in the correct position. A detailed view of the lens arrangement is shown in Figure 3.25. An important addition not shown in these two diagrams is an acrylic cover which fits tightly over the aluminium base and prevents the ingress of dust to the tapered amplifier. Currently, the cover has small holes cut for the input and output beams – these will ultimately be fitted with AR coated windows to seal the unit. Renders of the assembled mount and cover are given in Appendix E.

Performance

With the optics fixed in position, the tapered amplifier is now essentially a turn-key system. The temperature controller remains on permanently while the current is switched on once the seed light is present, with the system is ready to use within seconds. The only maintainance that has been required in over a year of operation has been occasional re-




Figure 3.24: Exploded view of new tapered amplifier mount design. The mount is centred on the plate shown in grey, which is clamped on top of the cold plate of a TEC, heatsunk to the optical bench via the solid support below. The plate also holds both temperature lock thermistors and provides support for the TA protection circuitry. The tapered amplifier itself mounts to the component shown as a translucent red cuboid, the detailed form of which is not shown as it is a proprietary part. The brown cover above protects the top of the chip, while the two covers shown in green fit into the inset grooves and protect the front and back of the chip from damage. These two covers also provide vee-grooves which support the input and output aspheric lenses. The lenses are held in position by a multi-axis stage while being aligned, before a small amount of epoxy resin is placed between the lens and vee-groove to fix them permanently. A more detailed view of the covers and lenses is given in Figure 3.25.




Figure 3.25: Detailed cross-section view of the protective covers and lens placement for the new tapered amplifier mount. The chip itself is mounted to a proprietary component that sits below the cover and is not shown. As in Figure 3.23, the fast axis rays are shown in red with the slow axis in blue. Note that the astigmatism between the two axes is a significant fraction of the depth of the chip.




Figure 3.26: Output profile of tapered amplifier beam after isolator, recorded with a neutral-density-filtered CMOS camera (Thorlabs DCC1545M). Circular diffraction artefacts from the isolator are visible across the profile and the beam appears far from the Gaussian ideal. Nonetheless, we have achieved $> 55\%$ coupling efficiency to a single mode fibre, indicating a lower M^2 value than the profile suggests. Pixel size is $5.2 \times 5.2 \mu\text{m}$.

optimisation of the coupling to the input waveguide.

Seeded with 24 mW of light, the tapered amplifier reaches its maximum rated output power of 500 mW with 1150 mA of current, some way below its maximum of 1570 mA. We achieve a little over 70% transmission through the double-stage isolator³² and 55% transmission through the single-mode inter-room fibre, providing a maximum of 200 mW pre-AOM. The AOM double pass efficiency varies with the RF frequency used but is typically 50% including losses on the accompanying optics. The final power on the trap breadboard is ~ 55 mW, reducing to an estimated 40 mW at the position of the ion. This provides a 12x increase on the original power, increasing the Rabi frequency by a factor of 3.5.

The output beam profile, measured one metre from the TA, is shown in Figure 3.26. The beam looks quite far from a gaussian and it is slightly surprising that the fibre coupling efficiency is as high as it is. This profile would probably be improved by using a longer focal length lens on the output, which would make the horizontal extent of the source less significant.

³²Thorlabs IOT-5-730-MP -60 dB

3.7 Experimental Control System

Spectroscopy and sideband cooling experiments are controlled by a field-programmable gate array (FPGA) system,³³ which provides accurate and synchronised switching of many TTL outputs at once. These outputs are used to control a variety of RF switches and attenuators that provide power to the AOMs that switch each laser frequency. The FPGA also records fluorescence data from the PMTs in the form of photon counts. The original FPGA system was developed by two MSci students, S. Patel and S. Zeeman, supervised by D. Crick and S. Mavadia. This system worked well for the preliminary studies, but as the complexity and volume of the experiments grew it became necessary to make major modifications to the software and hardware, which I implemented in collaboration with a visiting summer student, S. Woodrow.

The original system is described in references [56,91,92] and the basic principles remain unchanged from those described there. A graphical user interface, coded in C# in the .NET framework, provides desktop PC control of the FPGA. Pulse sequences are constructed in the software before being compiled to a binary format that can be uploaded to the FPGA, where the instructions are stored in the SDRAM. The desktop software can then trigger the FPGA to run the uploaded pulse sequence, while also controlling the frequency of the function generator used to drive the 729-nm AOM, via GPIB.

Once running, the FPGA steps through the pulse sequence, storing PMT count data for certain periods of the pulse sequence in the onboard SRAM. The FPGA communicates with the control software whenever it reaches a point where the RF frequency must be altered and uploads the stored PMT data to the PC. When the RF frequency shift is confirmed and the PMT data has finished uploading, the FPGA wipes the SRAM and repeats the pulse sequence. The fluorescence data that is uploaded from the FPGA is stored in a file, and is also used to generate a live display of the spectrum as it is recorded.

The original system could only control five devices via FPGA TTL logic, and a single RF source over GPIB. Furthermore, the spectra were limited to continuous scans in fixed steps which meant that broad but sparsely populated spectra took an unnecessarily long time. The data files produced were also simple lists of fluorescence data, meaning that all information associated with each spectrum had to be recorded by hand and re-entered when the spectrum was later analysed. This was very impractical and error-prone once large amounts of data were being analysed.

³³Altera Cyclone II, DK-CYCII-2C20N

Summary of upgrades to FPGA system

The most important modifications that I have made to the system are as follows:

- **Interleaved spectra**

One of the most important new features is the ability to record multiple spectra at the same time. The spectra are interleaved on a reading-by-reading level, which ensures that no systematic variation of uncontrolled experimental parameters can occur between the various spectra. This allows the effects of changing certain parameters to be properly evaluated, e.g that of sideband cooling duration on final phonon number, or of a wait period on the heating of an ion.

- **Windowed and fixed frequency spectra**

There are many instances when it is useful to record spectra at fixed probe laser frequencies (such as a Rabi oscillation) or to record spectra over a series of windows centred on the features of interest, skipping the regions between. Provision for both these methods is included in the new software. As the windows of interest are almost always centered on sidebands, the software includes a straightforward interface to set up axial or radial spectra with windows at the correct positions for a user-selected trap voltage or frequency.

- **Metadata**

The file containing the fluorescence data is now prefaced by a set of metadata in a fixed format. The metadata contains all the information known to the spectroscopy program, such as the frequency steps or windowing pattern, the nominal carrier frequency, time, date, number of repeats, as well as optional information entered by the user when prompted at the start of the spectrum. This means that files can later be opened and analysed without error, as the frequency parameters of each spectrum are recorded unambiguously.

- **Data structure**

Linked to the inclusion of metadata was a complete rewriting of the data structure of the analysis program. The fluorescence data includes metadata but beyond this remains an anonymous list of readings. Upon loading into the spectroscopy viewer software, the metadata is used to immediately identify each reading with a particular spectrum, frequency and repeat. Each spectrum in the experiment then consists of a series of 'Data Point' data structures, each of which contains all the relevant fluorescence readings, plus the AOM frequency and a range of other higher-level information calculated from these fundamentals. This makes manipulation of complex interleaved and/or windowed spectra fast and entirely automated.

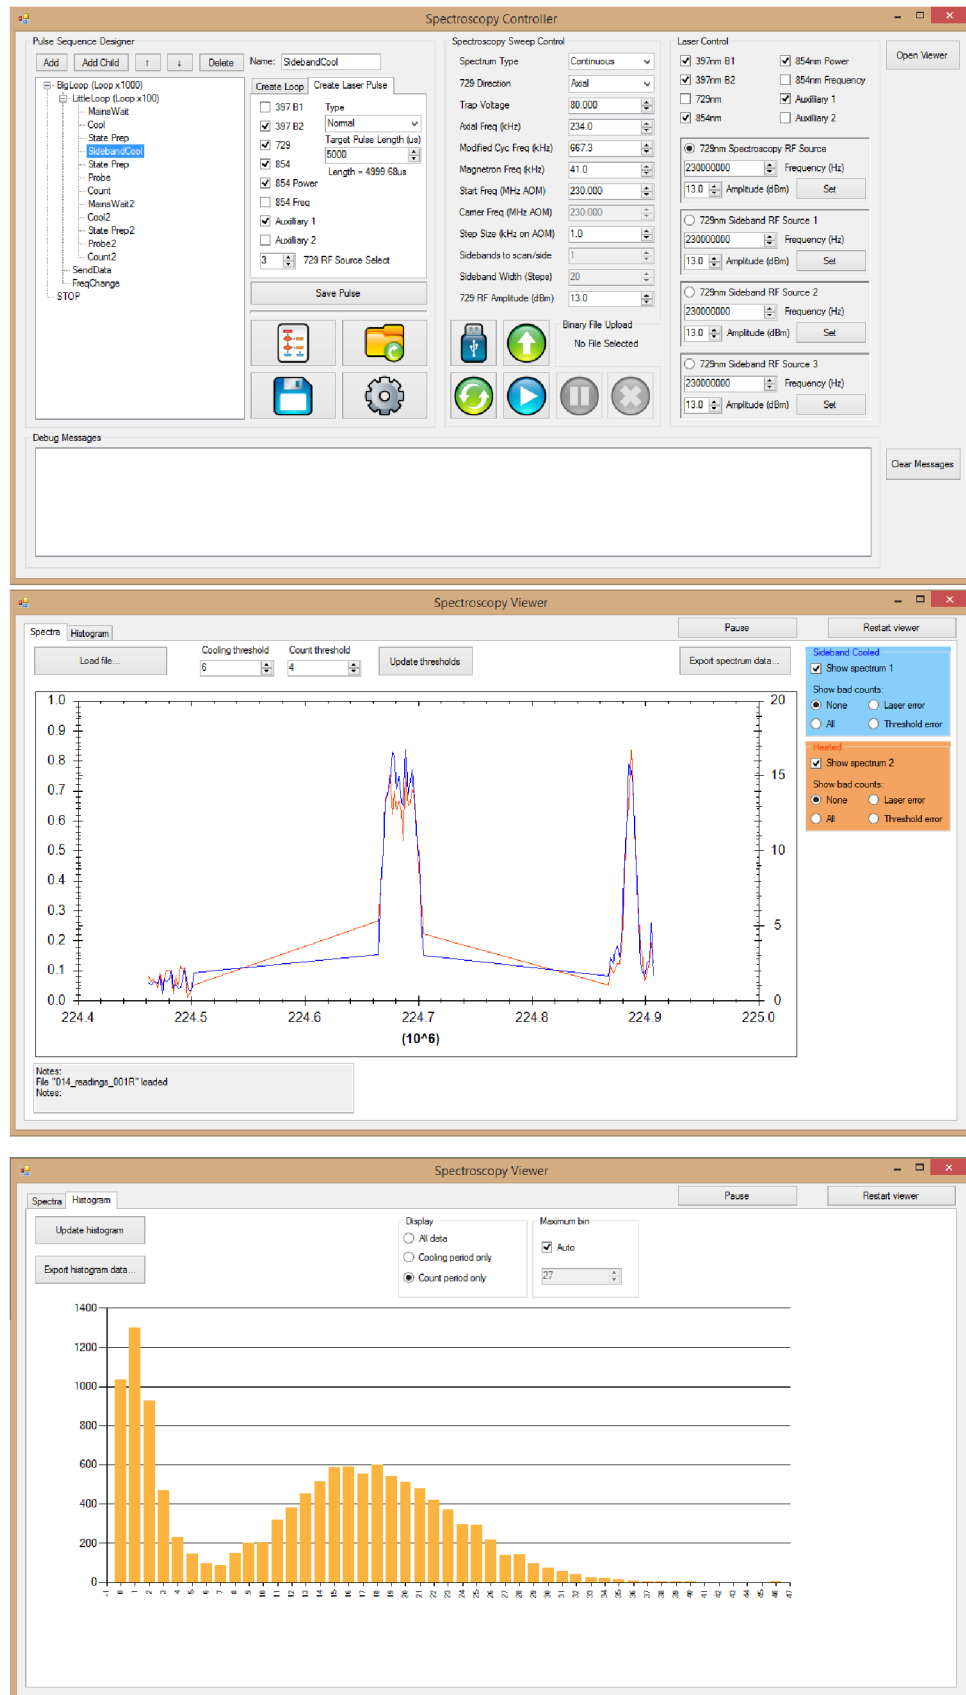


Figure 3.27: Screenshots of the upgraded FPGA pulse sequence designer and spectroscopy controller (top) and spectroscopy output viewer showing interleaved spectra (middle) and fluorescence histogram (bottom).

- **Improved viewer functionality**

The new data structure has allowed for a considerably more powerful and flexible manipulation of spectra, both live and in post-analysis. Multiple spectra may be loaded to the viewer, both those that have been interleaved and those recorded at different times. Multiple files may be combined into single spectra, or displayed as separate entities. Individual spectra may be toggled on or off, along with the ‘bad count’ records that denote readings that do not satisfy certain thresholds or logic conditions. The fluorescence data can be viewed as a histogram, where the user can select between data recorded during Doppler cooling phases and readout phases of the pulse sequence.

- **More outputs**

The five original TTL outputs of the FPGA did not provide enough logical control to perform sideband cooling, so the system was extended to provide ten outputs. These now control the switches for the two 397-nm AOMs, the 729-nm AOM and the 854-nm AOM, two-bit source-switching for the 729-nm AOM, attenuator switching for the 854-nm and 397-nm AOMs, plus two auxilliary bits for future use.³⁴ The GPIB control was also extended to up to four function generators, allowing for multi-stage sideband cooling or other complex manipulations.

- **Improved hardware**

The changes to the number of outputs naturally required a physical alteration to the interface hardware of the FPGA. The original hardware, assembled on prototyping board, had also become temperamental and prone to intermittent faults, so all the auxilliary electronics were rebuilt from scratch. These consist of a mains-phase trigger system, a set of 3.3 V to 5 V logic buffers for every output, and an LED array displaying the current status of the FPGA outputs. The RF switching system was upgraded and expanded at the same time; the new system is shown in Figure 3.28.

³⁴The pulse sequences passed to the FPGA consist of sequences of 32-bit instructions, all bits of which were used in the original implementation. To provide bandwidth for the five extra logic bits, the pulse timing bits had to be reduced from 24 to 19. It was decided to limit the minimum pulse length multiple to 640 ns to preserve the ability to provide long single pulses (0.34 s), however this could easily be switched to return the timing precision to its original 20 ns value. In this case pulses longer than 10 ms would be constructed from several concatenated short pulses, though this should lead to no ill effects beyond inconvenience.

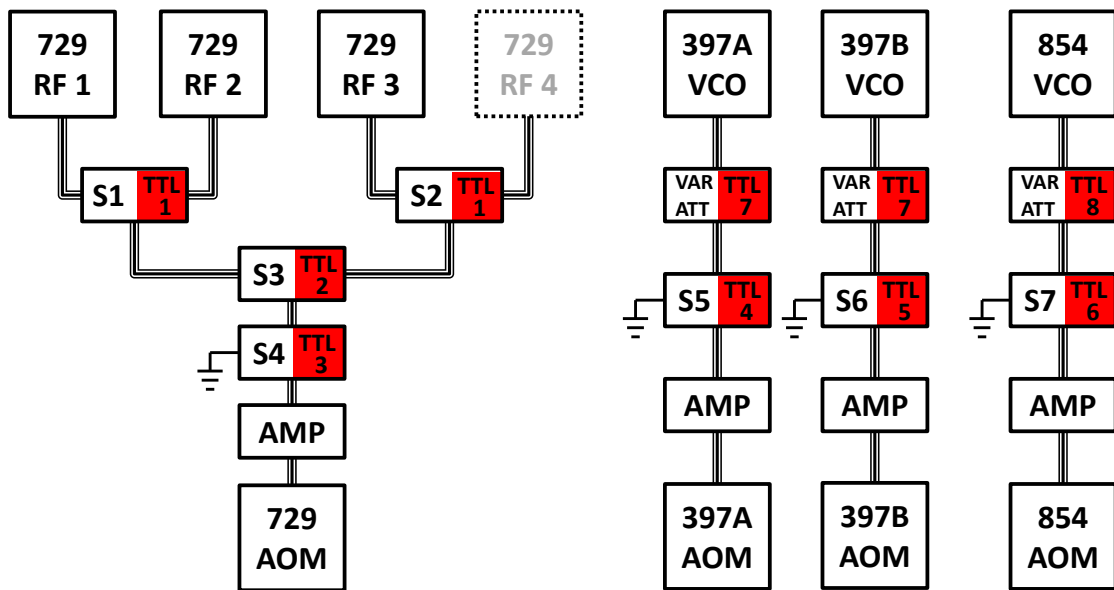


Figure 3.28: Layout of AOM RF system. Eight buffered TTL outputs from the FPGA control a series of switches and variable attenuators. The 729-nm sources are three HP8643A signal generators while the 397-nm and 854-nm AOMs are driven by Mini-Circuits voltage-controlled-oscillators (397: ZX95-200, 854: ZOS-100). Switches S1-S3 select the RF source for the 729-nm AOM, while S4-S7 switch the RF to each AOM on or off. TTL7 and TTL8 control analog switches to feed one of two control voltages to each attenuator to switch between high or low laser power. The control voltages are set manually.

4

Sideband Cooling in a Penning Trap

In this chapter I present the main experimental results of this thesis, namely the ground state cooling of an ion in a Penning trap.

While the sideband cooling technique was perfected in RF traps many years ago [4, 93], it had yet to be realised in a Penning trap due to the increased technical complexities associated with any work in this system. Indeed, while sideband cooling is the most important result of this thesis, the great majority of my work has focussed on achieving the prerequisite conditions necessary to perform such a task.

To sideband cool effectively, the system must start in a low phonon number state, ideally meeting the Lamb-Dicke criterion of $\eta\sqrt{2\bar{n}+1} \ll 1$. In Sections 3.5 and 3.4 I have described the alterations made to the apparatus to allow both reliable Doppler cooling to temperatures just above the Doppler limit and the increased trap frequencies that decrease the average phonon number at these temperatures. Naturally, sideband cooling also requires the ability to continuously drive a single sideband of a very narrow transition, effectively demanding the condition $\Omega_0^2 \gg \Delta^2$, that the Rabi frequency be much larger than any detuning from the sideband due to uncontrolled drifts. In Section 3.6 I have described the work to both increase the laser intensity and decrease the frequency drifts, which has allowed this condition to be comfortably met.

I begin this chapter with a brief review of electron-shelving spectroscopy, the most important diagnostic tool in the experiment. The spectra produced over the past year demonstrate how the changes listed above have allowed us to improve the axial temperature, phonon number and stability of our ion after Doppler cooling. I show that despite the huge improvements, the phonon number remains too high to allow coherent control. I also present the results of spectroscopy of the radial motion of the ion after Doppler cooling,

which we have not yet attempted to sideband cool.

In the following section, I discuss the protocol and techniques used for resolved sideband cooling, and highlight some important considerations in a Penning trap. I then demonstrate the cooling of the axial motion of a single ion to its quantum ground state with 99% probability, and show the increased coherence at this temperature in the form of Rabi oscillations.

Finally, I present a measurement of the ion heating rate of SPECTRAP, and show it to be one of the lowest measured to date, with a phonon heating rate of $\dot{n} = 2.3 \pm 0.1 \text{ s}^{-1}$ and inferred frequency-scaled electric field spectral noise density of $\omega_z S_E(\omega_z) = 1.47 \pm 0.06 \times 10^{-8} \text{ V}^2 \text{ m}^{-2} \text{ Hz}^{-1} \text{ s}^{-1}$, a result of some interest given the unusually large electrode dimensions of the trap.

4.1 Spectroscopy on the $S_{1/2} \leftrightarrow D_{5/2}$ Transition

In $^{40}\text{Ca}^+$, successful Doppler cooling depends primarily on the parameters of the 397 nm cooling lasers. However, while measurement of the Doppler-broadened linewidth of the cooling transition does provide information about the ion's motion, as temperatures approach the Doppler limit this broadening becomes very small ($\sim 0.5 \text{ MHz}$) compared with the natural linewidth (21 MHz) making accurate measurements become difficult. Furthermore, if multiple motional modes are present (such as in the radial motion, or in any motion of a multi-ion crystal) then the temperatures of each mode will contribute to the same broadening effect and will not be resolved. For this reason, spectroscopy in the strong-binding regime provides an incredibly useful tool for the characterisation of the ion's motion after Doppler cooling, allowing for accurate, mode-resolved temperature measurement at the Doppler limit - and beyond - as well as the ability to distinguish thermal and non-thermal states. The ability to resolve sidebands spectroscopically is also the foundation for resolved-sideband cooling (Section 4.3) and many of the necessary technical capabilities are common to both techniques.

The linewidth of the $S_{1/2} \leftrightarrow D_{5/2}$ transition is approximately 1 Hz and, if probed with 729-nm radiation of suitably narrow linewidth,¹ the absorption spectrum of a harmonically trapped ion will consist of a carrier surrounded by a number of equally spaced, resolved sidebands. However, traditional spectroscopy looks for some sort of modulation in the fluorescence and this is essentially impossible to detect when only a single photon is scattered each second. The solution is to detect the state of the ion indirectly, by applying a pulse of 729-nm light before probing on the $S_{1/2} \leftrightarrow P_{3/2}$ Doppler cooling transition and looking for fluorescence. If the 729-nm pulse has succeeded in exciting the ion to the $D_{5/2}$ state, it is stuck outside the cooling cycle and no fluorescence will be observed. Each experiment gives a binary outcome ('bright' or 'dark'), but by performing repeated Bernoulli trials at each

¹The 729-nm system described in Chapter 3.6 has a linewidth of $< 1 \text{ kHz}$




Figure 4.1: Laser pulse sequence for electron-shelving spectroscopy of Doppler cooled ion. After cooling, the ion is optically pumped into the correct $S_{1/2}$ sub-level by switching off one of the cooling lasers. N.B. Pulse lengths not to scale

729-nm probe frequency, the probability of absorption for that frequency can be deduced. This technique is known as *electron-shelving* and was first used for resolved-sideband spectroscopy by workers at NIST [94], based on principles originally pioneered by the Dehmelt group [95]. A detailed description of our own spectroscopy system can be found in the thesis of S. Mavadia [56]; for ease of reference I provide a brief overview here.

4.1.1 Spectroscopy pulse sequence

The pulse sequence used for basic spectroscopy is shown in Figure 4.1. Such sequences are controlled by the FPGA hardware and desktop user interface described in Section 3.7.

The sequence is triggered by a electrical mains phase detector. Because the 50 Hz mains currents form one of largest sources of noise in the lab, both within electronic systems themselves and in the form of AC magnetic fields, performing spectroscopy experiments at the same point in each mains cycle significantly reduces the effective linewidth of the system.

The first step in the sequence is to Doppler cool the ion. Once the initial cooling of the ion from room temperature has occurred, Doppler cooling only need occur for a fraction of a microsecond per mains cycle to keep the ion cold. However we choose to use a much longer Doppler cooling period (typically 5-8ms) during which the fluorescence photons from the ion are counted. This ensures that if the ion undergoes a collision with a background gas molecule it is rapidly recooled and also provides a reference of the fluorescence level before shelving. If the counts during this period are less than a certain threshold value, the reading is ignored. At this stage, the control software also records the status of the various laser locks. If any are found to be unlocked the reading is again ignored, as the

probe frequency or ion temperature could have changed from the expected value.

The ion is prepared for spectroscopy by switching off one of the Doppler cooling lasers, which optically pumps the system into the appropriate $S_{1/2}$ sub-level. One notable advantage of using a Penning trap is that this process is achieved through frequency, rather than polarisation, selection due to the large Zeeman splitting of the ground state sub-levels. This provides essentially perfect state preparation fidelity, avoiding the difficulties of preparing and maintaining very pure laser polarisations through the optical system. The remaining 397-nm cooling laser and 866-nm repump laser remain on throughout the experiment. However, the 854-nm laser is switched off after the state preparation step to avoid removing population from the shelving level during the spectroscopy that follows.

A pulse of light at 729nm is applied to the ion, shelving population in the $D_{5/2}$ level if the laser is resonant with any sideband. The length of the probe pulse is adjusted to ensure a significant amount of shelving occurs without saturation. In general, the higher the intensity of the beam, the shorter the probe time and the broader the apparent spectral features. In the semi-classical picture this broadening can be understood as the Fourier transform limit of the rectangular pulse, and at low saturations the quantum dynamics show the same behaviour. These ideas are discussed further in Appendix B, where I also provide the expected excitation lineshapes for a range of temperatures, Rabi frequencies and pulse times.

To perform read-out the ion is then probed with both lasers at 397nm, but the 854nm laser remains switched off to preserve any shelved population. The fluorescence during this 5-8ms period is counted and if it is below a certain threshold value the ion is considered to have been shelved by the 729nm pulse. Beyond any uncontrolled noise in the frequency and intensity of the probe laser, two important sources of error may affect the recorded spectrum. The first is spontaneous decay from the excited state before read-out is complete. As the lifetime of the state is ~ 1 s and the read-out time is < 10 ms this will produce a $\sim 0.5\%$ reduction in the apparent signal. The second effect occurs because of magnetically induced j -mixing decays (Section 3.2.1) on the $S_{1/2} \rightarrow D_{5/2}$ transition, producing $D_{5/2}$ shelving without 729nm excitation and increasing the apparent signal.

j -mixing decays occur with a fixed branching ratio (approximately 10^{-6} at $B = 1.85$ T) and will occur, on average, after a certain number of photons have been scattered. The effects of j -mixing are minimised when the second threshold is set to a low value. Being able to do so without losing the true shelving signal at the same time relies on the background being low. The higher the fluorescence collection efficiency, the fewer photons that must be scattered during read-out and the lower the j -mixing rate in this period. More details on how optimum thresholds are calculated is given in [56]. Under typical experimental conditions, j -mixing provides an average background of $\sim 2\%$.²

After the detection period is completed, the entire pulse sequence is then repeated 100-

²If the ion has already been excited, j -mixing cannot occur and so the background scales with $1 - P_{excite}$

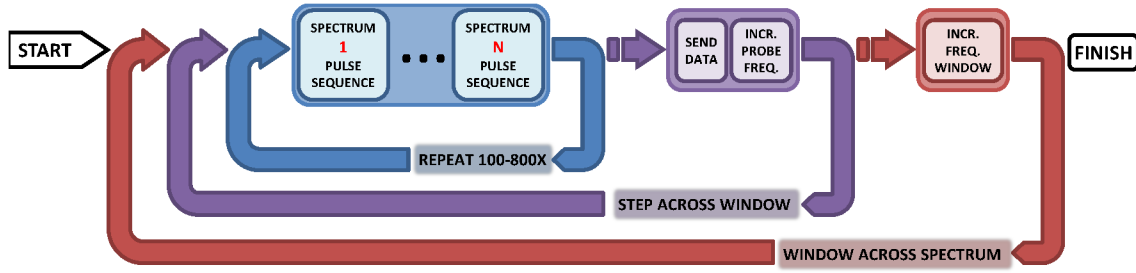


Figure 4.2: Structure of a generic spectroscopy experiment. Multiple spectra are recorded simultaneously and interleaved on a reading-by-reading basis. After a number of Bernoulli trials to determine the excitation probability, the frequency is incremented. If using a windowed spectra, the probe is applied at a range of frequencies around one sideband before the window is moved to the next sideband and the process repeated.

800 times to produce an estimate of the probability of excitation. The frequency of the 729nm laser is then incremented³ and the process repeated. Two variations on a ‘standard’ spectrum are regularly used. The first involves ‘windowing’ the spectrum so that only a small range of frequencies around each sideband is scanned. This greatly reduces the time taken to record a spectrum especially when features are fine and sparse, but care must be taken that the unrecorded areas contain no (perhaps unexpected) features of interest. The second technique is employed when a comparison is required between two spectra, typically before and after sideband cooling. The fairest comparison is made when both spectra are recorded at the same time, and in this case we interleave readings for the two (or more) spectra at the finest level so the pulse sequence for each spectrum follows the previous one sequentially before the entire set is repeated 100-800 times and the laser frequency is incremented (Figure 4.2).

4.1.2 Beam geometry

Our experiment allows for the 729-nm laser to propagate along either of the existing beam paths, either parallel (axial) or perpendicular (radial) to the magnetic field. In each case, the spectra recorded will be determined by the motion of the ion in the direction of propagation. Of course, if the beam is not perfectly aligned with the trap axis, a small component of the nominally orthogonal motion may also be visible on the spectrum.

4.1.3 Choice of transition

As with state-preparation, the large Zeeman splittings mean that transitions may be selected via the frequency of the 729-nm laser alone, without needing to worry about polarisation purity. However, whether a transition can be driven at all does depend on the polarisation- and wave-vectors of the beam, as discussed in Appendix A and [56]. For an

³If producing a Rabi oscillation rather than a spectrum, the frequency remains constant while the length of the 729nm probe period is increased.

electric quadrupole transition like $S_{1/2} \leftrightarrow D_{5/2}$, π -transitions ($\Delta m_j = 0$) cannot be excited by light exactly parallel or perpendicular to the quantisation direction and thus cannot be driven in our trap. σ^\pm -transitions ($\Delta m_j = \pm 1$) can be driven by light at all incident angles, providing the polarisation is correct, while δ^\pm -transitions ($\Delta m_j = \pm 2$) cannot be driven by beams parallel to the quantisation axis and are only accessible for radial spectroscopy. Different transitions also have different fundamental strengths (governed by the angular momentum combinatorics captured in the Clebsch-Gordan coefficients - see Table 3.2) and different magnetic field sensitivities (proportional to the differential Zeeman splitting of the transition - see Table 3.1).

This choice should simply be a matter of selecting a strong transition with low field sensitivity, but there may also be pragmatic considerations. One issue discussed in Section 3.6.2 is whether the modes of the 729-nm ULE cavity near the transition frequency are suitable for the PDH lock. Our initial work used the $S_{1/2}(m_j = \frac{1}{2}) \leftrightarrow D_{5/2}(m_j = -\frac{1}{2})$ transition as this could be accessed while locked to the TEM_{0,2} mode. High order modes are harder to mode match and provide a weaker error signal, but with careful alignment it has been possible to lock to modes as high as TEM_{0,21}. We now use the $S_{1/2}(m_j = -\frac{1}{2}) \leftrightarrow D_{5/2}(m_j = -\frac{3}{2})$ transition, which is twice as strong and with half the field sensitivity but requires us to lock to the TEM_{0,12} mode and requires more frequent alignment.

4.2 Doppler Cooling

4.2.1 Axial motion

After Doppler cooling, the ion will be in a thermal state, providing care has been taken to minimise the effects of relaxation heating (Section 2.3.2). The Doppler limit for the axial motion of $^{40}\text{Ca}^+$ cooled on the $S_{1/2} \leftrightarrow P_{1/2}$ transition with equal axial and radial cooling laser intensities and low saturation is (Equation 2.94)

$$T_{Dopp} = \frac{9\hbar\Gamma}{20k_B} = 454\mu\text{K}. \quad (4.1)$$

Outside the Lamb-Dicke regime, and providing the spectroscopy transition is not saturated, the sideband spectrum of a thermal ion will follow an approximately Gaussian envelope (Section 2.2.2). We can relate the $1/e^2$ half-width of this envelope, f_0 , to the temperature of the ion via

$$T = \frac{mf_0^2\lambda^2}{2k_B} \quad (4.2)$$

where λ is the wavelength of the probe laser (729nm). We see that an ion at the Doppler limit will produce a spectrum of width $f_{0,Dopp} \sim 600\text{ kHz}$. Naturally, cooling at the wrong frequency or with high saturation parameters will increase this limit.

Once motional sidebands have been observed, sideband cooling is – in principle – simply a matter of tuning the laser to the centre of one of the red sidebands, and waiting.

However, if the laser and sideband drift apart, cooling will stop. Furthermore, if the initial temperature is too high (or the spontaneous decay rate of the transition is too low, see Section 4.3.2), cooling may occur at a rate lower than that at which the ion is heated by external sources. Our initial axial spectra displayed both of these deleterious characteristics. Drifts of tens of kHz/hour combined with low Rabi frequencies made it impossible to keep the laser driving a single sideband for long enough to cool. Meanwhile, high temperatures and low trap frequencies meant that sideband cooling would have required multiple stages, cooling first on a high order red sideband before moving in steps towards the first red sideband to reach the ground state.⁴ This would have been necessary both to ensure a reasonable rate of cooling and to avoid trapping significant amounts of the population in high- n states (Section 2.3.4).

Figure 4.3 shows a series of axial spectra recorded over the course of 2013, highlighting several stages of improvement to these characteristics over this period.

The very first spectra recorded (see reference [56]) were heavily saturated which, combined with huge drifts (several hundred kHz/day) and instabilities, made it impossible to make any reliable temperature measurements. Initial improvements (primarily increasing the laser power while significantly reducing the probe time) provided the first relatively repeatable spectra, an early example of which is shown in Figure 4.3(a). The spectrum is recorded at an axial frequency of 117 kHz, which yielded the best performance at the time, and the spectrum has been ‘windowed’ around each sideband to minimize the total measurement time. The low probe laser power and large number of sidebands give a very low Rabi frequency on each sideband, leading to very narrow peaks ($\text{FWHM} \sim 3 \text{ kHz}$). The combination of narrow peaks and significant frequency drifts mean that the spacing appears irregular when viewed at large scale, preventing statistical fitting without an overly convoluted model. The solid line - a comb of Lorentzians under a Gaussian envelope - has therefore been ‘hand-fitted’, providing a temperature of approximately $T \sim 1.3 \text{ mK}$ and a average phonon number of $\bar{n} \sim 230$.

We make the assumption that the major contributing effects to the temperature are the power broadening of the transition and the relaxation heating caused by the abrupt termination of the associated radiation pressure (Section 2.3.2). The temperature excess above the Doppler limit is then consistent with a saturation parameter of $s_0 = 0.5$ which, taking into account population trapped in the $D_{3/2}$ manifold, corresponds to a scattering rate of $\gamma_s = 16 \times 10^6 \text{ s}^{-1}$. Due to the low trap frequency, the relaxation heating is the dominant effect, responsible for an increase of $\Delta\bar{n} \sim 130$, while the power broadening is quite weak, with $\Gamma_B \sim 1.2 \times \Gamma$. This significant coherent excitation makes the energy distribution significantly non-thermal and the spectrum should show a small amount of suppression of the central sidebands, approximately half-way between the two examples given in Figure 2.4. No obvious central dip is observed in (a), but given the subtlety of

⁴Such techniques have been demonstrated, but only at much lower phonon numbers [64].

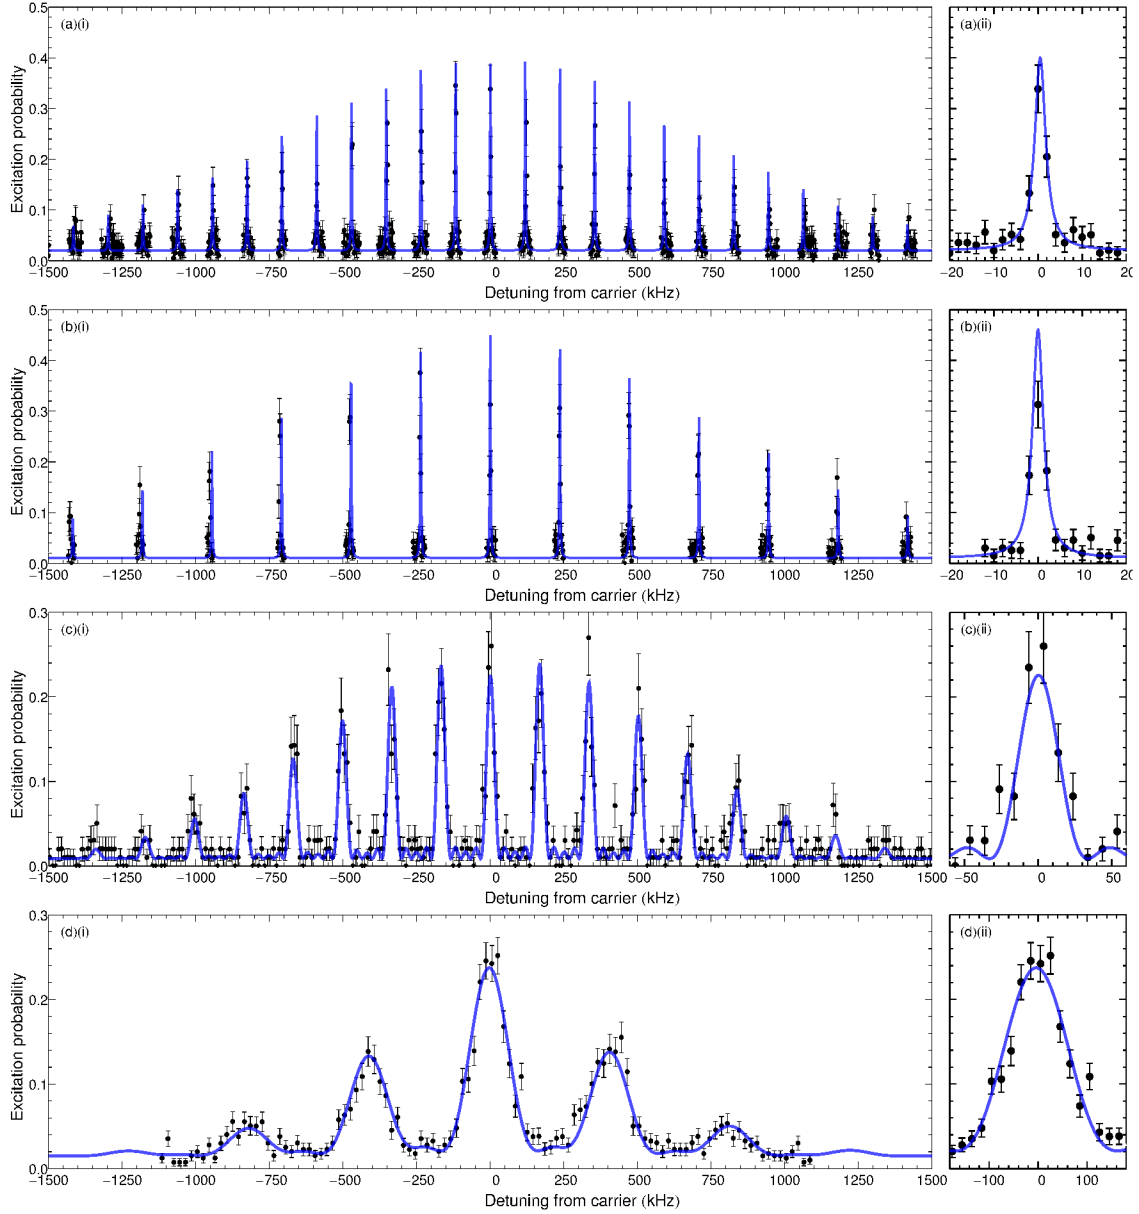


Figure 4.3: Axial motional sideband spectra after Doppler cooling. The top two frames show spectra ‘windowed’ around each sideband. Frequency drifts during these spectra prevent statistical fitting without an overly convoluted model and the solid lines are hand-fitted combs of Lorentzians under a Gaussian envelope. The data in the bottom two frames are fitted to a comb of sinc^2 sidebands with heights determined by the Rabi dynamics of a thermally distributed ion. The panels (a-d)(ii) of the right column show each carrier in detail. (a) early spectrum recorded at a trap frequency of $f_z = 117$ kHz. The temperature is $T \sim 1.3$ mK, and the phonon number is $\bar{n} \sim 230$. (b) spectrum produced after DC compensation of the trap, at a frequency of 234 kHz, with $T \sim 1.5$ mK, and $\bar{n} \sim 130$. (c) spectrum produced after a range of stability and intensity upgrades, with $f_z = 165.4$ kHz, $T = 755 \pm 55 \mu\text{K}$ and $\bar{n} = 95 \pm 7$. (d) a recent spectrum, after the introduction of axialisation to provide stability at a much higher trap frequency of $f_z = 405$ kHz. The fit gives $\bar{n} = 24 \pm 2$, equivalent to a temperature of $T = 470 \pm 40 \mu\text{K}$.

the effect and particularly poor quality of the data it would be unwise to place too much significance on this.

The first major improvement was seen after better DC compensation of the trap (Section 3.4.4) and an improvement in our radial cooling technique. If stray potentials develop (e.g. when electrodes become coated in calcium from the oven), the equilibrium position of the ion can vary strongly with trap frequency. Intensity gradient cooling at high trap frequencies requires careful alignment of the gradient with the ion position, and if this position is not constant this can be hard to achieve. Removing the unnecessary repump frequencies from the radial beam (these can be provided by the axial beam alone) also improved our ability to set the beam position correctly via the behaviour of the fluorescence signal. These changes allowed us to reach somewhat higher trap frequencies, as shown in Figure 4.3(b) which shows an axial spectrum at 234 kHz, with $T \sim 1.5$ mK and a average phonon number of $\bar{n} \sim 130$. This spectrum has also been ‘hand-fitted’ for the reasons described above.

Making the same assumptions as for the previous spectrum, this temperature implies a saturation parameter of $s_0 = 3$ and scattering rate of $\gamma_s = 28 \times 10^6 \text{ s}^{-1}$. Despite the high scattering rate, the higher trap frequency greatly reduces the effect of the relaxation heating, and this and the power broadening make a similar contribution to the final temperature. The final energy distribution is relatively close to thermal, and the sideband spectrum is not expected to show significant deviation from this model.

Figure 4.3(c) shows a spectrum at $f_z = 165.4$ kHz after a range of improvements were made to the laser frequency stability and intensity. Finding the zero-crossing point of the ULE cavity and improving the temperature controller (Section 3.6.2) removed all significant drift from the laser. Moving the trap to the centre of the magnetic field, where the gradient approaches zero, and improving the mechanical stability (Section 3.4.2) greatly reduced the frequency drift due to movements in the field. After these changes, the frequency variations were reduced to < 2 kHz/hour with a long term drift of < 2 kHz/day. Meanwhile the addition of the tapered amplifier (Section 3.6.3) increased the 729-nm Rabi frequency by a factor of ~ 3.3 , broadening the observed peaks and further reducing the sensitivity to frequency drifts. Finally, the cooling laser scanning cavity lock was improved, and the new control interface made reliable laser frequency selection possible (Section 3.5), reducing the achievable temperatures. The data in Figure 4.3(c) are fitted to a comb of sinc^2 sidebands with heights determined by the Rabi dynamics of a thermal distribution, which is a good approximation to the full dynamics as the spectrum is recorded far below probe saturation. The parameters from the fit provide $\bar{n} = 95 \pm 7$, equivalent to $755 \pm 55 \mu\text{K}$, with a Rabi frequency of $\Omega_0 = 2\pi \times (104 \pm 2) \text{ kHz}$. This temperature would be expected with a saturation parameter of $s_0 \sim 0.35$ and scattering rate of $\gamma_s \sim 13 \times 10^6 \text{ s}^{-1}$. As the trap frequency is low, relaxation heating would dominate, but the effect is slight and the sideband spectrum should appear thermal, as is observed.




Figure 4.4: Spectrum of carrier after Doppler cooling with 800 repeats. The solid line is a fit to the dynamics of a thermally distributed population with $\bar{n} = 25 \pm 5$.

The final, and arguably most significant, improvement was the use of axialisation to mediate the radial cooling, which required the careful compensation of the RF null position to avoid micromotion (Section 3.4.4). This provided stable trapping right up to the theoretical stability limit of the trap ($f_z \sim 500$ kHz). The increased stability also meant the Doppler cooling laser intensities could be reduced well below saturation, reducing the temperatures to the Doppler limit. Because of the low intensity and high trap frequency, relaxation heating became insignificant. Adjustments to the 729-nm laser optics on the trap breadboard decreased the beam waist at the ion, increasing the Rabi frequency by a further factor of 2.5. The result is shown in Figure 4.3(d), where the solid line is a fit to the theory as described for (c). The fit gives a Rabi frequency of $\Omega_0 = 2\pi \times (264 \pm 5)$ kHz and $\bar{n} = 24 \pm 2$, equivalent to $470 \pm 40 \mu\text{K}$ ($1.04 \times T_{\text{Dopp}}$). At this trap frequency, relaxation heating is insignificant. Such a low temperature suggests a very low saturation parameter $s_0 \sim 0.1$, but could also be explained by relatively weak scattering from the radial beam, which reduces the expected Doppler limit (see Equation 2.93a). The ion now lies just outside the Lamb-Dicke regime, and can be cooled to the ground state on the first red sideband directly, leaving only 0.15% of the population trapped in high- n dark states.

The improvements to the stability of the system have also allowed more precise spectroscopy to be performed. Figure 4.4 shows a scan over the carrier after Doppler cooling, with 800 repeats per frequency to reduce the statistical noise. The intensity of the probe laser has been reduced to better determine the carrier frequency. The data are an excellent fit to the theoretical dynamics of a thermally distributed population with $\bar{n} = 25$ and $\Omega_0 = 2\pi \times 70$ kHz.




Figure 4.5: Radial motional sideband spectrum after Doppler cooling at a low axial trap frequency.

4.2.2 Radial motion

The axial motion of an ion in a Penning trap is well described by a one-dimensional simple harmonic oscillator, and is similar to the secular motion of an ion in a Paul trap. However, as has been discussed in Chapter 2.1.1, the radial motion is quite unique, consisting of a pair of magnetron and modified cyclotron orbital motions in a single plane, producing a more complicated, epistrochoidal motion. Because these two modes lie in the same plane, it is impossible to measure their individual temperatures via Doppler spectroscopy. For this reason, sideband resolved spectroscopy is a particularly useful tool in analysing this behaviour, and indeed the following radial spectrum provides the first mode-resolved temperature measurement of the radial motion in a Penning trap.

Figure 4.5 shows the spectrum recorded when the probe laser is incident in the radial plane, for an ion trapped with an axial frequency of 58.5 kHz giving magnetron and modified cyclotron frequencies of $f_m = 2.4$ kHz and $f_{c'} = 705.9$ kHz. The spectrum consists of a carrier modulated by the motion at the two radial frequencies. Because both motions have high n , the envelopes of the sideband heights for each mode are well approximated by Gaussians. The comb of magnetron sidebands around each cyclotron sideband follow a Gaussian profile, the width of which provides the magnetron temperature. The height of these Gaussians is in turn modulated by a Gaussian envelope, the width of which gives the cyclotron temperature. The fit here gives temperatures of $T_m = 4.3 \pm 0.1 \mu\text{K}$ and $T_{c'} = 5.3 \pm 0.1$ mK, equivalent to phonon numbers of $n_{c'} = 78$ and $n_m = 5400$.

This spectrum was recorded using intensity gradient cooling only, with no axialisation drive. Under these conditions, the magnetron temperature limit depends on the steepness of the intensity gradient, parameterised by $1/y_0$ (see Section 2.3.3), and for optimal detuning is given by equation 2.94. From the measured temperature, we can infer a maximum value of $y_{0,max} = 190 \mu\text{m}$. This is fairly gentle gradient and substituting this into the cyclotron temperature equation gives no significant deviation from the temperature for a uniform beam, which is $T_c(y_0 \rightarrow \infty) = 1.1$ mK for this trap frequency. This is nearly five times lower than the observed temperature.

Such an imbalance in mode temperatures can be caused by a reduced Doppler cooling rate, which increases the cyclotron temperature while decreasing the magnetron temperature (as Doppler selection acts to heat the magnetron). For low saturation, Doppler cooling is optimised when $\delta = \Gamma/2$, but if this is not set correctly the cooling term in equation 2.93c will be reduced. The observed temperatures are produced if the gradient is $y_0 = 125 \mu\text{m}$ and the cooling term is reduced to 0.42 of its maximum value. For sub-saturation intensities, this occurs at a detuning of $\delta = 2\Gamma$, a full 30 MHz from the optimum value. It is therefore more likely that the reduction in cooling rate is at least in part due to saturation effects, which are not captured in equation 2.93c.

While these effects could explain the increased cyclotron temperature, this has not yet been tested and remains one of the primary concerns of our ongoing research. Because the radial spectrum is so wide, with complex structure and low magnetron Rabi frequency, spectroscopy is very time consuming and the drifts of both the Doppler cooling lasers and spectroscopy system have stymied our attempts to make systematic studies of the behaviour. However, it is hoped that the recent improvements to these systems may allow some progress to be made in the near future. Alternatively, using axialisation to control the radial motion may provide more stability and repeatability at the expense of an increased minimum cyclotron temperature.

4.2.3 Coherent manipulation near the Doppler limit

To understand the importance of sideband cooling, I now consider our ability to coherently control the quantum state of the ion after Doppler cooling. By probing the ion axially, on-resonance with the carrier transition and gradually increasing the probe time, we observe Rabi oscillations (Figure 4.6). After Doppler cooling, the axial phonon number was in an approximately thermal distribution with $\bar{n} = 25 \pm 3$, determined by fitting to a full spectrum as in Figure 4.12. Because this is a statistical mixture of large number of different Fock states, each contributing a different Rabi frequency, the oscillations rapidly lose contrast.⁵ The corresponding fit to this dephased mixture provides an independent measure of the phonon number, $\bar{n} = 31 \pm 2$, which is slightly higher than spectroscopically measured value. This difference could be due to the presence of laser noise (see Section 4.3.5), which provides a secondary dephasing mechanism to the Rabi oscillations and leads to an apparently raised temperature. It is clear that coherent control is not possible at such a high phonon number - to improve this significantly requires cooling far below the Doppler limit.

⁵It is important to remember that this is an incoherent mixture, rather than a superposition, and the ion is only ever in a single Fock state at any time. However, our lack of knowledge of *which* Fock state it is in at a given time prevents us from performing reliable coherent manipulations.




Figure 4.6: Rabi oscillation on the carrier of the $S_{1/2}(m_j = +\frac{1}{2}) \leftrightarrow D_{5/2}(m_j = -\frac{1}{2})$ transition after Doppler cooling. The fit is to the summed Rabi dynamics of a thermal distribution of Fock states and gives $\Omega_0 \sim 2\pi \times 27$ kHz. The oscillations rapidly lose contrast as $\bar{n} \sim 30$.

4.3 Sideband Cooling

I now describe the first ground state cooling of an ion in a Penning trap. Throughout this section I refer to the cooling of the axial motion only. While sideband cooling of the radial motion should be possible, and indeed would provide an ideal solution to cooling the unstable magnetron mode, the initial quantum numbers are far higher and off-resonant heating effects will be more significant.

4.3.1 Pulse sequence

The ion is initially Doppler cooled, before one of the two 397-nm lasers is switched off to optically pump population into the $S_{1/2}(m_j = -\frac{1}{2})$ sub-level. Sideband cooling is performed by applying the 729-nm laser to the first red axial sideband of the $S_{1/2}(m_j = -\frac{1}{2}) \leftrightarrow D_{5/2}(m_j = -\frac{3}{2})$ transition. The $D_{5/2}$ state has decay rate of 0.895 s^{-1} [67], making cooling directly from this level impractically slow and we use a weak 854-nm ‘quench’ laser to increase the scattering rate by emptying the $D_{5/2}$ level via $P_{3/2}$. Provided the saturation parameter of the quench laser is small, the $P_{3/2}$ level can be adiabatically eliminated and the system behaves like a virtual two level system with a controllable linewidth set by the properties of the quench laser and upper level [60]. We provide this quench by leaving the 854nm laser turned on during the sideband cooling sequence, reduced to 25% of full power to limit the effective linewidth to an estimated 25 – 50 kHz. We also reduce the 729-nm laser to 50% of its maximum intensity during cooling, providing a Rabi frequency on the first red sideband of $\sim 2\pi \times 7$ kHz in $n = 1$, reducing the off-resonant absorption which limits the final temperature. The second 397-nm laser and 866-nm repump lasers are applied throughout the sequence to ensure that population decaying on the $P_{3/2} \leftrightarrow S_{1/2}$ transition is optically pumped into the correct ($m_j = -\frac{1}{2}$) ground state sub-level. After sideband cooling, electron shelving spectroscopy is performed as described in 4.1.




Figure 4.7: Laser pulse sequence for sideband cooling. The 854nm laser is applied at reduced power during the sideband cooling period to increase the spontaneous decay rate from the $D_{5/2}$ level. 729nm ‘A’ is fixed at the AC-Stark shifted frequency of the first red sideband, with power reduced by 50% to minimise the final temperature. 729nm ‘B’ is stepped sequentially across the spectrum, as in Figure 4.1. A second state preparation period ensures the ion population is optically pumped from the sideband cooling cycle to the correct ground state for spectroscopy.

4.3.2 Optimising the quench laser parameters

The quench laser parameters are initially adjusted by performing spectroscopy with it switched on, as shown in Figure 4.8. The power and detuning of the quench laser are adjusted to produce an effective linewidth of 25 – 50 kHz, which is measured spectroscopically. The presence of the quench laser leads to an AC Stark shift of the $D_{5/2}$ level of up to ~ 80 kHz. The greater the detuning of the 854-nm laser from resonance, the greater the AC Stark shift (for the same quench rate) but the lower the sensitivity to frequency drifts. Because the second-order Zeeman effect means that the evenly-spaced 854-nm laser frequencies cannot be resonant with all the transitions simultaneously (Sections 3.5 and 3.2.1), accurate and repeatable tuning of the laser is difficult. Combined with the fact that the laser is not locked to an absolute frequency, this means that the exact detuning varies from day to day, changing the Stark shift.

4.3.3 Optimising the sideband cooling frequency

To cool effectively, the cooling laser must be resonant with the AC Stark shifted frequency of the first red sideband. It is therefore essential that the shifted frequency must be measured before we attempt to sideband cool.

The AC Stark shifted frequency is most easily determined spectroscopically. By tuning




Figure 4.8: Carrier after Doppler cooling probed with (blue) and without (orange) the 854-nm quench laser applied during spectroscopy. The probe time is $60\mu\text{s}$ for the unquenched spectrum and 2ms for the quenched spectrum, the much longer pulse being needed to provide significant shelving with the decreased upper state lifetime. Here, the presence of the quench laser causes an AC Stark shift of the $D_{5/2}$ level of $\sim 65\text{ kHz}$.

the probe laser frequency to the top of the first red sideband and instead scanning the cooling frequency, the spectrum shows a notch where sideband cooling occurs, the centre of which is the optimum frequency for the cooling laser. Figure 4.9 shows such a ‘dip’ spectrum at higher and lower cooling laser powers. At low power, the notch is narrower allowing for more accurate determination of the optimum frequency, while the higher power spectrum confirms that cooling to or near the ground state is occurring. The optimum cooling frequency in this case is found to be -414 kHz , implying an AC Stark shift of 25 kHz due to the 854 nm quench laser.

4.3.4 Effects of radial motion

If the radial temperature is high, the effects of the motion can be seen when probing the ion with an axial 729-nm beam in the form of sidebands at the magnetron, modified cyclotron and cyclotron frequencies (the last due to axialisation micromotion). There are two possible mechanisms that could produce such an effect.

The first is a simple misalignment of the probe beam with the trap axis, so a small projection of the radial motion causes modulation proportional to the velocity in the direction of the laser. At the axialisation cooling limit 2.3.3, the magnetron and cyclotron modes have equal quantum number and thus modulation indices, producing combs with equal numbers of sidebands.

The second relates to the inhomogeneity of the magnetic field across the diameter of the radial orbit. As the field varies, so does the Zeeman splitting and so periodic motion within




Figure 4.9: Height of the first red sideband against the detuning of the sideband cooling laser from the carrier, for an axial frequency of 389 kHz. The height after 21 ms of cooling at high and low 729 nm laser powers are shown in red and orange respectively, while the height after Doppler cooling only is shown for reference, in blue.

an inhomogeneous field will produce a comb of sidebands at the modulating frequencies, which are again those of the radial modes. Although in this case the modulation index depends on the field gradient and radius of the orbit rather than the velocity, the spectral effect is identical and at $\bar{n}_m = \bar{n}_{c'}$ the magnetron and modified cyclotron combs would have the same numbers of sidebands. Note also that in this case the sidebands will not be true ‘radial’ sidebands (with absorption associated with a change in radial quantum number) but FM sidebands of the axial spectrum.

To determine the source of these radial artefacts, the ion was sideband cooled on the second red sideband to well within the Lamb-Dicke regime. A spectrum was then measured at high radial temperature, produced by slightly misaligning the trap compensation and reducing the axialisation amplitude, shown in Figure 4.10. The probe pulse length was chosen to significantly saturate the carrier transition such that the other features would be enhanced. When the depth of modulation in the spectrum is measured and compared to the magnitude expected due to each mechanism (see Appendix C for details), it becomes apparent that such large artefacts could only be produced by a probe beam misalignment of several milliradians, and that the field gradient does not make a significant contribution.

Although the artefacts become much smaller when the radial temperature is reduced to its normal limit, they still present a significant problem when attempting to sideband cool. At certain trap voltages, the first axial red sideband coincides nearly exactly with one or more radial sidebands. In the case of $\omega_z = 2\pi \times 405$ kHz, the first red sideband lies very close to the radial sideband at $\omega_m - \omega_{c'} = -415$ kHz. As the ion approaches the ground state and the red sideband is reduced to a similar amplitude to the radial artefact, absorption occurs




Figure 4.10: Heavily saturated spectrum around the carrier after sideband cooling on the second red sideband. Because the probe laser is not perfectly aligned with the trap axis, a projection of the radial motion is seen, providing peaks at combinations of ω_c , $\omega_{c'}$ and ω_m as well as the carrier and axial red and blue sidebands that would be expected. Note the overlap between the first axial red sideband at $-\omega_z = -405$ kHz and the radial peak at $\omega_m - \omega_{c'} = -415$ kHz. In this example, the radial temperatures are unusually high due to poor compensation and insufficient axialisation drive amplitude, but the features remain present on a much smaller scale under normal conditions. For details of the fit, see Appendix C

increasingly on the radial feature rather than the red sideband. Because this absorption process does not reduce the axial phonon number, the net effect of one cooling cycle is to heat the ion, and this heating process will limit the minimum temperature than can be reached. Furthermore, even if the ion is somehow placed in the ground state, the radial artefact will appear to be a residual red sideband and prevent an accurate temperature measurement from being made. Figure 4.11 shows the limit of sideband cooling at an axial frequency of $\omega_z = 2\pi \times 405$ kHz, with a distinctive peak at 409 kHz. This lies almost exactly between the red sideband frequency and the radial artefact and it is possible that it is a product of both the aforementioned effects.

This kind of problematic coincidence can be avoided by simply reducing the beam misalignment until the artefact disappears. Unfortunately, due to a region of Nd:YAG-induced optical damage on the trap window, it is not possible to perfectly align the axial beam without causing large amounts of scatter from the Doppler cooling lasers. Without being able to sideband cool the radial motion to the ground state, we must resort to the more pragmatic option of choosing trap frequencies that do not produce this sort of undesirable coincidence, as will be demonstrated in the following section.

4.3.5 Ground state cooling

The ion is trapped with an axial frequency of 389 kHz and Doppler cooled for 5 ms. From the observed spectrum we measure an average phonon number of $\bar{n} = 24 \pm 1$, corresponding to




Figure 4.11: After cooling at an axial frequency of 405 kHz, the average phonon number appears to remain at $\bar{n} = 0.09$. It is not clear whether this is due to poor cooling, or whether the remaining peak is actually a coincident radial sideband. This can be avoided by reducing the radial temperature, providing more accurate alignment of the probe laser, or simply by choosing a trap frequency at which such coincidences do not occur.

an axial temperature of 0.45 mK, equal to the expected Doppler cooling limit for matched axial and radial cooling beam intensities at the ion position. Figure 4.12 shows a typical axial sideband spectrum after this step. We do not measure a radial temperature directly during these experiments, but previous work suggests this is several times higher than the axial temperature (see Section 4.2.2). Note that for this trap frequency, the 729-nm axial Lamb-Dicke parameter is $\eta_{729} = 0.155$ and the ion remains somewhat outside the Lamb-Dicke regime ($\eta_{729}^2(2\bar{n} + 1) \sim 1.18$) after Doppler cooling, though not far enough to prevent us from cooling efficiently on the first red sideband.

It is important to note at this point that not all the population can be cooled to ground state by cooling on the first order sideband. As has been discussed in Section 2.3.4, the sideband coupling strength vanishes close to $n = 150$ for this value of η , leaving the population initially in oscillator states higher than this trapped at this energy. For $\bar{n} = 24$, approximately 0.0017 of the population remains trapped at this level.⁶ Although the fraction of the population trapped is very small, the high oscillator number at which it is trapped significantly raises the average phonon number, and prevents $\bar{n} < 0.26$. However, this value is somewhat misleading as the energy distribution is strongly non-thermal and produces very different behaviour to that of a thermal distribution with the same \bar{n} (which would imply much smaller population in the $n = 0$ state). In the experiments described in this thesis, the population trapped at $n = 150$ is expected to produce effects that are nearly two orders of magnitude below that of the statistical noise and other errors. For the rest of this chapter I ignore the trapped population and for convenience define \bar{n} to describe the average energy of the thermal population only.

⁶To prevent this from occurring, two stage sideband cooling is necessary. This has recently been demonstrated for our system and will be discussed in the thesis of G. Stutter.




Figure 4.12: Motional sideband spectrum after Doppler cooling at $\omega_z = 2\pi \times 389$ kHz. The solid line is a fit to the full Rabi dynamics of a thermally distributed population, giving $\bar{n} = 24 \pm 1$, equivalent to a temperature of 0.45 ± 0.02 mK. The peaks between the main sidebands are slightly higher than would be expected from the theory. This excess could be due to a projection of the magnetron motion. This is at a frequency $\omega_m = 2\pi \times 131$ kHz, coincident with the first Rabi sinc side-lobes, and would be visible due to a misalignment of the probe laser from the trap axis.

Figure 4.13 shows the first red and blue sidebands after 21 ms of sideband cooling.⁷ The blue sideband is fitted to a Rabi sinc profile on a constant background. Because the blue sideband Rabi dynamics in $n = 0$ are identical to the red sideband dynamics in $n = 1$, and $\bar{n} \ll 1$, the red sideband is correctly described by the sum of a uniform background B and the blue profile scaled by \bar{n} :

$$P_{RSB} = p_0 B + p_1 \frac{\Omega_{BSB}^2}{\Omega_{BSB}^2 + \Delta^2} \sin^2 \left(\sqrt{\Omega_{BSB}^2 + \Delta^2} \times \frac{t}{2} \right). \quad (4.3)$$

The fit shows the average phonon number after sideband cooling to be $\bar{n} = 0.013 \pm 0.009$. For an effective linewidth of $\Gamma = 2\pi \times 50$ kHz and Lamb-Dicke parameters for absorption and emission of $\eta_{729} = 0.156$ and $\eta_{393,em} = 0.167$ respectively,⁸ the theoretical limit (Equation 2.104) for sideband cooling at this trap frequency is $\bar{n} = 0.006$, consistent with the measured result. The most straightforward way to reduce this temperature further would

⁷Such a long cooling period is not necessary and crude measurements of the dynamics (not shown) suggest the cooling limit is reached in ~ 4 ms. However, in this experiment we were primarily interested in achieving the best possible ground state confinement and chose a cooling time that would guarantee we reached the cooling limit.

⁸As discussed in Section 2.3.4, the calculation of the Lamb-Dicke parameter $\bar{\eta}$ is complicated by the fact that emission occurs into 3-dimensions, which decreases the effective Lamb-Dicke parameter. The value given here assumes an isotropic decay, and is thus a factor of $\sqrt{3}$ smaller than the parameter for absorption at 393 nm, $\eta_{393,abs} = 0.290$.




Figure 4.13: (a) First red and (b) first blue sidebands after 21 ms of cooling on the first red sideband. The ratio of the sideband heights above the background shows that the average phonon number is $\bar{n} = 0.013 \pm 0.009$.

be to reduce the quench laser intensity, although we are already beginning to approach the limits of what we can measure without using much higher statistics across the spectrum.

Coherent manipulation in the ground state




Figure 4.14: Rabi oscillation on the carrier after ground state cooling. The Rabi frequency is initially $\Omega_0 \sim 2\pi \times 41$ kHz but laser intensity noise causes an abrupt shift at $70 \mu\text{s}$ and the fit to the remainder of the plot gives $\Omega_0 \sim 2\pi \times 39$ kHz. The overall visibility decays exponentially with a coherence time of $\tau = 0.7 \pm 0.1$ ms.

The motivation for ground state cooling is to decrease the thermal dephasing during coherent manipulation. Figure 4.14 shows Rabi oscillations on the carrier after ground state cooling, with $\bar{n} \sim 0.01$. Compared to those measured near the Doppler limit (Figure 4.6), the oscillations are now well defined, with a coherence time of $\tau = 0.7 \pm 0.1$ ms, approximately consistent with our spectroscopically measured linewidth of $\Delta\nu = 0.6 \pm 0.4$ kHz [35]. Note that the fit to the first half of the data gives a Rabi frequency of $\Omega_0 = 2\pi \times 41$ kHz




Figure 4.15: High-resolution scan of the Rabi dynamics of a $100\mu\text{s}$ probe pulse on the carrier after sideband cooling to the ground-state. The solid line is a fit to the theory from which we deduce a pulse area on resonance of 3.25π

and the second half $\Omega_0 = 2\pi \times 39\text{ kHz}$, with a noisy discontinuity between the two. Several other smaller discontinuities are also visible. This behaviour is due to uncontrolled intensity noise on the 729-nm probe laser. Because the laser's power is boosted twice, by an injection locked slave laser followed by a tapered amplifier, it is affected by mechanical drifts in both of these systems as well as the variations in the master diode intensity, fibre coupling and beam pointing that are normally present. As well as the abrupt changes in intensity, there is likely to be a continuous variation on shorter timescales due to acoustic noise, which will reduce the coherence during ion-laser interactions. To overcome these issues, the group are currently developing a power-noise-eating feedback system for this laser and are also improving the beam pointing stability into the trap.

Coherence can also be observed during spectroscopy. Figure 4.15 shows the Rabi dynamics as the probe laser is detuned from the carrier transition. The probe time is $100\mu\text{s}$ and the Rabi sinc fit gives a pulse area on resonance of 3.25π .

For a Gaussian beam with waist $w = 50\mu\text{m}$ and a power of 40 mW at the ion, the Rabi frequency should be $\Omega_0 \sim 2\pi \times 370\text{ kHz}$, while the highest observed in our experiments has been 55 kHz. The reason for this large discrepancy is not yet known. In principle, the intensity at the ion could be lower than we have calculated, but we have no reason to believe the waist size is larger than we have previously measured and observe an appropriate amount of power transmitted through the trap. It seems most likely that the problem is instead one of insufficient spectral power - i.e. that the 1 kHz linewidth of the 729-nm laser only contains 15% of the total power in the beam, while the rest is spread across a broad range of frequencies at a much lower spectral density. It is possible this could occur if much of the noise on the master laser occurred at frequencies beyond the bandwidth of the 729-nm PDH lock, although a malfunctioning injection lock should also be ruled out. Given the degree of possible improvement, understanding the source of this




Figure 4.16: (a) First red and (b) first blue sidebands of a ion sideband cooled to the ground state. The blue spectrum was measured immediately after sideband cooling, while the red spectrum was measured after a 100ms delay during which heating occurs. The ratio of the sideband heights above the background shows that the average phonon number before (after) the delay is $\bar{n} = 0.067 \pm 0.011$ ($\bar{n} = 0.333 \pm 0.024$), implying a heating rate of $\dot{\bar{n}} = 2.7 \pm 0.3 \text{ s}^{-1}$.

discrepancy is a priority for future work.

4.4 Heating Rate

The heating rate of the trap is measured by cooling to the ground state and inserting a delay period before spectroscopy is performed, during which no cooling is applied. The delayed spectroscopy measurements are interleaved with a set of prompt measurements, providing a pair of spectra showing the temperatures before and after the heating delay, shown in Figure 4.16. The change in the ratio of the red and blue sideband heights then allows the change in phonon number during this delay to be determined. The results of several heating rate measurements, with delays of either 50 ms and 100 ms, are summarised in Figure 4.17. The heating rate of $\dot{\bar{n}} = 2.3 \pm 0.1 \text{ s}^{-1}$ is very low, as would be expected from such a large trap.

Figure 4.18 shows a comparison of our trap heating rate to those of a variety of other traps, plotted against the distance to the nearest electrode. Here I have taken the usual approach of plotting the phonon heating rate in terms of an inferred noise spectral density,

$$S_E(\omega_z) = 4m\hbar\omega_z\dot{\bar{n}}/e^2, \quad (4.4)$$




Figure 4.17: The heating rate is determined by measuring the phonon number after a variable delay between cooling and spectroscopy. The fitted line gives a heating rate of $\dot{n} = 2.3 \pm 0.1 \text{ s}^{-1}$ at 389 kHz.

scaled by the trap frequency as $\omega_z S_E(\omega_z)$. Note that I have not included any cryogenically cooled traps, nor those that have undergone any in-situ surface treatment. For ease of comparison, I have also only considered traps with 3-dimensional electrode structures. The SPECTRAP heating rate, shown in the lower right corner, is comparable to those presented in [64] and amongst the lowest reported to date in any type of trap.

There has been a great deal of research over the last decade into the underlying mechanism of heating in ion traps, and in particular the ‘anomalous’ heating due to the proximity of electrode surfaces [96]. It has become increasingly apparent in recent years that the dominant heating mechanism varies between traps and that this behaviour cannot necessarily be encapsulated in a single power-law, such as the oft-quoted $1/d^4$ electrode distance scaling [97]. With this proviso, I note that - empirically - the data in Figure 4.18 appear to display a steep power law scaling, of approximately $1/d^3$.

A recent measurement of the heating rate of two-dimensional crystals in a Penning trap with $d = 2 \text{ cm}$ [98] provided an indirect constraint on the single ion anomalous heating rate of $\dot{n} < 5 \text{ s}^{-1}$ at a trap frequency of 795 kHz, equivalent to a rate of $\dot{n} = 21 \text{ s}^{-1}$ at 389 kHz. The SPECTRAP result provides a constraint that is nearly an order of magnitude lower, despite the lower ion-electrode distance of $d = 1.08 \text{ cm}$. However, if the scaling suggested by the data in Figure 4.18 was valid for all d , an even lower heating rate might be expected in our trap. It is not yet clear whether this increased heating is due to the presence of a scale-independent process such as soft collisions with background gas, or due to technical noise and EM pickup on the trap electrodes, and we are conducting further experiments to ascertain the source, the first of which is to determine the spectrum of heating rate versus trap frequency.




Figure 4.18: The reported heating rates of a variety of linear, ring and needle traps, shown as a frequency-scaled, noise spectral density plotted against the distance to the nearest electrode. The SPECTRAP heating rate reported in this chapter is shown in the lower right corner. No uncertainty data is available for the results marked with “X” and no error bars are shown for these points. [Data sources: DesLauriers [99], Diedrich [4], Roos [93], Schulz [100], Stick [41], Poulsen [64], Benhelm [101], Blakestad [102], DeVoe [103], Monroe [5], Tamm [104]]

5

Global Entangling Gates in a Penning Trap

5.1 Overview

The demonstration of sideband cooling in a Penning trap described in the previous chapters represents an important milestone and opens the door to a wide range of new experiments. However, considering the great advances that have been made in RF traps in the last 20 years, one might fairly ask what advantages a Penning trap would provide in most situations. In those experiments where a strong magnetic field is required for other reasons, the use of a Penning trap is unavoidable and sideband cooled ions will provide the foundation for quantum logic spectroscopy in precision g -factor measurements, such as that proposed by the Ospelkaus group [33]. In most conventional quantum information experiments, involving linear chains or shuttled ions, the benefits are harder to see. A Penning trap could provide unparalleled heating rates and perfect state preparation, but these advantages are likely more than offset by the fact that reaching high trap frequencies is challenging, the laser systems more complicated and the route to scalability less obvious than for microfabricated traps.

Where Penning traps truly excel is in offering the ability to manipulate 2- or 3-dimensional Coulomb crystals. One topic which has engendered a great deal of recent interest is in the use of two-dimensional crystals as an analogue system for the quantum simulation of magnetism [21, 29, 105] via the generation of effective spin-spin interactions with optical dipole force excitation of collective motional modes. Such a system could also be of great use to quantum computation, where the array of ions would provide a convenient and

compact register of qubits, however the methods for implementing quantum gates and operations on this type of architecture have not been described in theory or practice.

This chapter covers theoretical work conducted in the final 6 months of my PhD, exploring the possibilities of engineering quantum resources with this system. I describe a novel technique for producing complex entangled states across a Coulomb crystal by driving many-ion, spin-asymmetric, geometric phase gates on vibrational modes that couple inhomogeneously to the ions in the crystal. By tailoring the parameters of these phase gates, a wide range of unitaries can be produced. The key advantage of this approach is that all laser and microwave pulses can be applied globally, avoiding the need for single ion addressing - a challenging technique in any trap, and particularly difficult for ions in a rapidly rotating crystal. This chapter will review the basic physics behind this technique and outline its application to the production of two simple quantum error-correcting codes using a two-dimensional Coulomb crystal, concluding with a discussion of the expected performance of the codes and the necessary experimental considerations. This work forms the basis for a new experiment that the Imperial group plans to conduct over the following years.

5.2 Quantum Error Correction

Quantum information is notoriously delicate, and any interaction with the environment can cause a qubit to decohere or decay, introducing errors. In classical computation, it is possible to correct against random errors by copying the information multiple times and taking a ‘majority vote’ upon readout - providing that less than half the bits receive errors, the information will remain intact. However, the no-cloning theorem means that quantum information cannot be copied and more subtle techniques must be found to protect against errors. The fundamental idea behind a quantum error correcting code (QECC) is to *spread* one qubit of logical information across several physical qubits. Once encoded, ‘syndrome’ measurements can be made that diagnose errors that have occurred on the code state (up to a certain maximum number) so they can be corrected immediately or compensated for in later operations.

The endeavour towards large-scale quantum technologies has recently seen impressive experimental efforts in the realisation of small quantum error-correcting codes (QECCs) using ion traps [15, 16, 20] and superconductors [106–108]. Given sufficient control over the composite physical systems, coherence times of encoded quantum states will increase exponentially with the size of the code. The ultimate goal of research in this field is the construction of arbitrarily large codes, but the experimental challenges are considerable and a more modest intermediate aim is the development of small error-correcting codes that can maintain quantum coherence for time scales longer than those of their composite parts.

The simplest QECCs are repetition codes, analagous to the classical repetition code. An example is the 5 qubit repetition code in the conjugate basis ($|+\rangle = \frac{1}{\sqrt{2}}[|0\rangle + |1\rangle]$, $|-\rangle = \frac{1}{\sqrt{2}}[|0\rangle - |1\rangle]$), which encodes a logical state as

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle \Rightarrow \alpha|++++\rangle + \beta|-----\rangle = |\Psi_L\rangle. \quad (5.1)$$

Pauli- \hat{Z} dephasing errors such as

$$(\hat{I} \otimes \hat{I} \otimes \hat{Z} \otimes \hat{I} \otimes \hat{I})|\Psi_L\rangle = \alpha|++-++\rangle + \beta|--+-\rangle \quad (5.2)$$

can occur on up to two physical qubits while leaving the logical information intact. If we make a parity measurement between each adjacent pair of physical qubits ($\hat{Z}_1\hat{Z}_2$, $\hat{Z}_2\hat{Z}_3$ etc), the four bits of classical information obtained tells us which qubits, if any, have suffered errors. However, if $n > 3$ or more qubits have suffered errors, the state is indistinguishable from one with $5 - n$ physical errors and a logical \hat{Z} error (taking $\alpha \leftrightarrow \beta$), and the code fails. It is also clear that any Pauli- \hat{X} or $-\hat{Y}$ bit-flip errors, which take $|+\rangle \rightarrow |+\rangle$ and $|-\rangle \rightarrow -|-\rangle$ will alter the logical information in a way that cannot be detected, although in certain trapped ion systems the bit-flip error rate is low enough that it can be ignored.

The simplest code that protects against any single Pauli error is the 5 qubit code,¹ which has codewords

$$\begin{aligned} |0_L\rangle &= \frac{1}{4} [|00000\rangle + |01001\rangle + |10010\rangle + |00101\rangle + |01010\rangle + |10100\rangle \\ &\quad - |01111\rangle - |11110\rangle - |11101\rangle - |11011\rangle - |10111\rangle \\ &\quad - |00011\rangle - |00110\rangle - |01100\rangle - |11000\rangle - |10001\rangle] \\ |1_L\rangle &= \frac{1}{4} [|00001\rangle + |00010\rangle + |00100\rangle + |01000\rangle + |10000\rangle \\ &\quad + |00111\rangle + |01110\rangle + |11100\rangle + |11001\rangle + |10011\rangle \\ &\quad - |01011\rangle - |10110\rangle - |01101\rangle - |11010\rangle - |10101\rangle - |11111\rangle]. \end{aligned} \quad (5.3)$$

5.2.1 Practical quantum codes

Although (5.1) looks rather similar in form to a classical repetition code, it is in fact a many-body entangled state, as are all quantum error correcting codes. Producing multipartite entanglement in trapped ions remains a matter of considerable difficulty that has only been achieved by a handful of groups worldwide, and even the simplest error correcting codes

¹It is straightforward to show that the 5 qubit code allows the detection of one error (and is corrupted by any more), but this requires the use of the stabiliser formalism, which is beyond the scope of this chapter and covered thoroughly by, e.g. [109].




Figure 5.1: Initial product state and quantum circuit for formation of pentagon graph state.

sit at the frontier of experimental quantum information science.

The traditional approach to producing complex entanglement has been to proceed piecewise, entangling a pair of ions at a time. There exist a number of different entangling gates, and methods of implementing them. One of the most common - and of most relevance to the work I will discuss - is the two-qubit controlled-phase gate described by

$$C\phi = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{i\phi} \end{pmatrix}, \quad (5.4)$$

which applies a phase shift of ϕ to the qubits if they are in the state $|1\rangle \otimes |1\rangle$. Often $\phi = \pi$ is chosen, and the gate becomes the controlled-Z:

$$CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad (5.5)$$

While the CZ gate (combined with the usual set of single-qubit gates) is universal, of particular interest to quantum error codes is its use in producing *graph states*. These are a class of entangled states where the qubits are represented by vertices V , joined by edges E , representing the entanglement between them. Precisely, the graph state $|\Psi_G\rangle$ corresponding to the mathematical graph $G(V, E)$ is given by

$$|\Psi_G\rangle = \prod_{(a,b) \in E} CZ_{(a,b)} |+\rangle^{\otimes V}. \quad (5.6)$$

In practical terms, the qubit register is placed in the product state $|+\rangle_1 \otimes |+\rangle_2 \otimes \dots \otimes |+\rangle_N$ and controlled-Z gates are applied between some or all of these qubits. For a relevant example, consider the ‘pentagon’ graph state given in Figure 5.1

Written in the computational basis, the initial product state is the sum of all binary terms from 0 to $2^5 - 1$. Applying the ring of CZ gates has the effect of introducing a π -radian phase shift for each neighbouring pair of $|1\rangle$ states in the ring, so for example

$$\begin{aligned}
 |00010\rangle &\rightarrow + |00010\rangle \\
 |00110\rangle &\rightarrow - |00110\rangle \\
 |00101\rangle &\rightarrow + |00101\rangle \\
 |10011\rangle &\rightarrow + |10011\rangle.
 \end{aligned} \tag{5.7}$$

Applying this to the whole 2^5 -term superposition gives the logical $|+_L\rangle$ state of the 5 qubit code. As I will show later in the chapter, more complicated graph states allow arbitrary qubits to be teleported back and forth to this code.

The theoretical description of such a code is certainly straightforward, and indeed the difficulties in the realisation of QECCs in ion traps are technical. First among these is the experimental finesse necessary to perform two-ion entangling gates on an array of many ions, as any spectator ions must be decoupled and recoupled from the entangling interaction. One approach is to physically shuttle [15] the ions back and forth between trapping zones, so that only two interact with the entangling laser beam at once. This requires careful design and advanced manufacturing techniques for the trap and drive electronics to prevent heating of the ions during shuttling. Alternatively, the ions can be left in fixed positions, continually exposed to the interaction, and spectroscopically decoupled. This can be achieved either by ‘hiding’ ions in ancillary electronic states [16], or by using spin-echo-type dynamic decoupling techniques [20]. Both spectroscopic methods require the repeated laser addressing of individual ions.

For a quantum code to be useful, the fidelity must be higher than it would have been on an unencoded qubit at the time the information is used or retrieved. One important factor is the rate of decoherence of the physical qubits: once the probability of error exceeds a certain (code-dependent) threshold, all codes will fail. However, another often-overlooked factor is the loss of fidelity associated with the encoding process itself [110, 111]. Encoding typically consists of a large number of quantum gates, each of which introduces some error to the system - the encoding sequence for even a small code can involve well over 100 discrete operations [16]. In this case, even the very modest goal of maintaining 90% fidelity immediately after encoding requires that each of these individual pulses have a fidelity of $> 99.9\%$. Finding efficient entangling protocols is therefore of the utmost importance in producing practical codes.

In this remainder of this chapter I propose an alternative to the local circuit approach for realising many-body entangled states. I describe an experimental protocol to prepare two QECCs using global entangling operations alone, avoiding the need to decouple subsets of the qubit register. Specifically, I consider a six-ion Coulomb crystal in a Penning trap [34]

(Figure 5.4), and provide a protocol to encode an arbitrary quantum state to either the five-qubit code (5QC), or the five-qubit repetition code (5RC). Syndrome information can be extracted while teleporting information from the code, allowing the correction of errors that may have occurred while the logical state was stored.

In general, it is difficult and often impossible to find global unitary operations that realise a certain multipartite entangled state from a given input state. However, many target states share a common symmetry with the crystal architecture.² This simplifies protocols to realise desired states, as all the terms in the state that are equivalent up to the symmetry evolve identically.

The global pulse sequences I describe enable the execution of entire error-correcting protocols, including syndrome extraction, using a very low number of operations compared with other trapped-ion protocols [15, 16, 20]. The pulse sequences found produce 5RC(5QC) using only 14(18) discrete operations, while the number of operations required to perform these protocols with two-qubit entangling gates is typically an order of magnitude greater. Conventional non-demolition stabilizer measurements [112, 113] are even more taxing; the 5QC stabilizers require several hundred operations per cycle. This reduction in operational complexity should substantially reduce the total gate noise introduced to the system.

Before describing this scheme, I will review some of the principles behind producing many-ion phase gates in a Penning trap.

5.3 Geometric Phases and the Driven Quantum Harmonic Oscillator

The Hamiltonian of a forced quantum harmonic oscillator is

$$\hat{H} = \hat{H}_0 + \hat{H}_i = \hbar\omega_0 \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) + F(t)x_0 \left(\hat{a} + \hat{a}^\dagger \right), \quad (5.8)$$

where ω_0 is the natural frequency of the oscillator, $F(t) = F_0 \cos(\omega t)$ is a periodic driving force and x_0 is the spread of the ground state wavefunction. Moving into a frame rotating at ω_0 (the interaction representation), the interaction Hamiltonian becomes

$$\begin{aligned} \hat{H}_I &= e^{iH_0 t/\hbar} \hat{H}_i e^{-iH_0 t/\hbar} \\ &= F_0 x_0 \cos(\omega t) \left[\hat{a}^\dagger e^{i\omega_0 t} + \hat{a} e^{-i\omega_0 t} \right] \\ &= \frac{F_0 x_0}{2} \left[\hat{a}^\dagger \left(e^{i(\omega+\omega_0)t} + e^{i(\omega-\omega_0)t} \right) + \hat{a} \left(e^{-i(\omega+\omega_0)t} + e^{-i(\omega-\omega_0)t} \right) \right]. \end{aligned} \quad (5.9)$$

²In the case of the codes I will describe, this is a five-fold rotational symmetry.




Figure 5.2: Phase space trajectory of driven quantum harmonic oscillator in interaction picture, showing relation of loop area to acquired phase. Coordinates in units of $F_0 x_0 / \delta \hbar$.

Making the rotating wave approximation we find

$$\hat{H}_I = \frac{F_0 x_0}{2} \left[\hat{a}^\dagger e^{i\delta t} + \hat{a} e^{-i\delta t} \right] \quad (5.10)$$

where $\delta = \omega - \omega_0$ is the detuning of the driving force from resonance. In general, $\hat{H}_I(t)$ does not commute with itself at different times, and the time evolution operator contains higher-order terms reflecting this [114]:

$$\hat{U}(t) = \exp \left[\frac{-i}{\hbar} \left(\int_0^t \hat{H}_I(t') dt' + \frac{1}{2} \int_0^t dt' \int_0^{t'} dt'' [\hat{H}_I(t'), \hat{H}_I(t'')] + \dots \right) \right]. \quad (5.11)$$

However, for a short enough time Δt , the Hamiltonian can be considered constant and the expression simplifies to

$$U(\Delta t) = e^{-iH_I \Delta t / \hbar} = e^{\Delta \alpha \hat{a}^\dagger - \Delta \alpha^* \hat{a}} = \hat{D}(\Delta \alpha) \quad (5.12)$$

where

$$\Delta \alpha = \frac{-i F_0 x_0}{2 \hbar} e^{i\delta t} \Delta t \quad (5.13)$$

and $\hat{D}(\Delta \alpha)$ is the displacement operator. This describes a displacement in a phase space parameterised by α , associated with the normalised position and momentum via

$$\alpha = \frac{x}{2x_0} + \frac{ip}{p_0}. \quad (5.14)$$

Using the Baker–Campbell–Hausdorff identity

$$e^{\hat{A}}e^{\hat{B}} = e^{(\hat{A}+\hat{B})}e^{\frac{1}{2}[\hat{A},\hat{B}]}, \quad (5.15)$$

the effect of two sequential small displacements is

$$\hat{D}(\Delta\alpha)\hat{D}(\Delta\beta) = \hat{D}(\Delta\alpha + \Delta\beta)e^{i\text{Im}\{\Delta\alpha\Delta\beta^*\}} \quad (5.16)$$

and thus

$$\prod_{i=1}^N \hat{D}(\Delta\alpha_i) = \hat{D}\left(\sum_{i=1}^N \Delta\alpha_i\right) \exp\left[i\text{Im}\left\{\sum_{j=2}^N \Delta\alpha_j \left(\sum_{k=1}^{j-1} \Delta\alpha_k\right)^*\right\}\right]. \quad (5.17)$$

In the infinitesimal limit $\Delta\alpha \rightarrow d\alpha$, $\Delta t \rightarrow dt$, the evolution after time t can then be written in terms of a displacement in phase space and an acquired phase:

$$\begin{aligned} \hat{U}(t) &= \hat{D}\left(\int_0^t \frac{d\alpha}{dt'} dt'\right) \exp\left[i\text{Im}\left\{\int_0^t (\alpha(t'))^* \frac{d\alpha}{dt'} dt'\right\}\right] \\ &= \hat{D}\left(\int_0^t \frac{-iF_0x_0}{2\hbar} e^{i\delta t'} dt'\right) \exp\left[i\text{Im}\left\{\int_0^t \frac{i}{\delta} \left(\frac{F_0x_0}{2\hbar}\right)^2 (1 - e^{i\delta t'}) dt'\right\}\right] \\ &= \hat{D}\left(\frac{F_0x_0}{2\delta\hbar} [1 - e^{i\delta t}]\right) \exp\left[i\frac{t}{\delta} \left(\frac{F_0x_0}{2\hbar}\right)^2\right] = \hat{D}(\alpha(t))e^{i\Phi(t)}. \end{aligned} \quad (5.18)$$

The harmonic oscillator is therefore driven around a circular trajectory through phase space (Figure 5.2), with radius

$$r_0 = \frac{F_0x_0}{\delta\hbar}, \quad (5.19)$$

returning to its original motional state every $t = 2\pi n\delta$. In doing so, the quantum harmonic oscillator wavefunction acquires a phase of

$$\Phi = \frac{\pi}{2} \left(\frac{F_0x_0}{\hbar\delta}\right)^2, \quad (5.20)$$

equal to twice the area enclosed by the trajectory in the x - p phase space.

It would be easy to consider geometric phases a curiosity of little importance, as the global phase of a wavefunction can never be observed directly. However, the *relative* phases between terms of a quantum system are vital to their description. Most interestingly, if phase shifts can be produced between terms of a multi-qubit system conditioned on their joint spin state, then it is possible to produce entanglement, as described in Section 5.2.1. If we want to introduce geometric phases conditioned on the spins of the system, we must first be able to produce a spin-dependent force.




Figure 5.3: Schematic of optical dipole force lasers for ${}^9\text{Be}^+$. At typical magnetic field strengths, the detuning of the Raman lasers from the $S_{1/2} \rightarrow P_{3/2}$ transition, Δ_R , is several 10s of GHz, while the difference frequency between the upper and lower beams is $\mu_R \approx 1\text{MHz}$.

5.4 Spin-Dependent Forces in a Penning Trap

I will now describe how off-resonant AC Stark shifts can be used to produce spin-dependent forces on ions in a Penning trap. The technique described is based on the methods of Britton *et al* [21] with an important extension to the use of unbalanced forces (spin-up vs spin-down). This extra degree of freedom greatly increases the utility of the technique to our application, namely the engineering of precise multi-qubit unitaries.³

The principles of this interaction are not specific to a particular species of ion, but a suitable example is beryllium, shown in Fig. 5.3, where the qubit is provided by the two Zeeman sublevels of the $S_{1/2}$ ground state of the ion.⁴ The method described relies on the Zeeman splitting being large and is thus only suited for implementation in a Penning trap, where these levels are separated by $\sim 100\text{GHz}$ at typical magnetic field strengths. Other techniques have been described for producing similar forces in RF traps [114].

We consider the two Raman laser frequencies shown in green in Figure 5.3. This pair of frequencies acts simultaneously on both the $S_{1/2}$ qubit levels $|\uparrow\rangle, |\downarrow\rangle$, coupling them off-resonantly with levels in the $P_{3/2}$ manifold. The frequency of the Raman pair is chosen such that $|\uparrow\rangle$ is coupled to a point between the $P_{3/2}(m_j = +1/2)$ and $P_{3/2}(m_j = +3/2)$ sublevels, with a detuning from the upper level of $\Delta_R \sim 10\text{s of GHz}$. The $|\downarrow\rangle$ state is then coupled to a point between the $P_{3/2}(m_j = -1/2)$ and $P_{3/2}(m_j = -3/2)$ sublevels. The Raman beams will produce AC Stark shifts on the qubit levels due to their off-resonant interaction with the levels in the upper manifold; the influence of each sublevel on this shift will depend

³The work in [21] focuses on the generation of short range spin-spin interactions for quantum simulation.

⁴It is assumed that the nuclear spin has been optically pumped to the $m_I = +3/2$ level.




Figure 5.4: A periodic spin-dependent force is generated on the ions of a Coulomb crystal by a pair of off-resonant Raman laser beams with frequency difference μ_R . The trap magnetic field is B , aligned along the z axis. The Raman beams lie in the $x-z$ plane, with an angular separation of $\theta/2$, above and below the crystal plane.

upon the polarisation and wavevector of the beams.

The pair of Raman beams lie in the $x-z$ plane, arranged at angles $\pm\theta_R/2$ slightly above and below the $x-y$ plane of the Coulomb crystal (Figure 5.4), so the crystal lies in the region where the two beams interfere. The polarisations of the two beams can be adjusted such that the differential AC Stark shift of the qubit levels (i.e. the shift of the transition frequency ω_0) due to each individual beam is zero, while their interference produces a one-dimensional optical lattice with wavevector $\mathbf{k}_R = (0, 0, k_R)$, the polarisation gradient of which provides a spin-dependent *optical dipole force* (ODF). Additionally, if a frequency difference $\mu_R = \omega_U - \omega_L$ is introduced between the lasers, the optical lattice will scan across the crystal, leading to a periodic, transverse ODF.⁵ If this frequency difference is tuned close to one of the normal mechanical modes of the crystal, the ODF can be made to excite collective vibrations of the ions.

I now derive this force, considering each of six allowed transitions between the qubit levels and the $P_{3/2}$ manifold as an effective two-level system and summing their contributions to give the total Stark shift. Modelling each laser as a uniform, plane wave with linear polarisation, the electric fields of the upper and lower beams are

$$\begin{aligned}\mathbf{E}_U(\mathbf{r}, t) &= \hat{\mathbf{e}}_U E \cos(\mathbf{k}_U \cdot \mathbf{r} - \omega_U t) \\ \mathbf{E}_L(\mathbf{r}, t) &= \hat{\mathbf{e}}_L E \cos(\mathbf{k}_L \cdot \mathbf{r} - \omega_L t).\end{aligned}\tag{5.21}$$

The wavevectors are

$$\mathbf{k}_U = k \left(\cos \frac{\theta_R}{2}, 0, -\sin \frac{\theta_R}{2} \right) \quad \mathbf{k}_L = k \left(\cos \frac{\theta_R}{2}, 0, \sin \frac{\theta_R}{2} \right), \tag{5.22}$$

⁵Transverse to the plane of the ion crystal, so in the direction of the magnetic field/trap axis.

and the polarisations are given by

$$\hat{\epsilon}_U = \begin{pmatrix} \sin \frac{\theta_R}{2} \cos \phi_p \\ \sin \phi_p \\ \cos \frac{\theta_R}{2} \cos \phi_p \end{pmatrix} \quad \hat{\epsilon}_L = \begin{pmatrix} -\sin \frac{\theta_R}{2} \cos \phi_p \\ \sin \phi_p \\ \cos \frac{\theta_R}{2} \cos \phi_p \end{pmatrix}, \quad (5.23)$$

where ϕ_p is the angle of the polarisation vector from the $x - z$ plane.

The Hamiltonian for the two-level system perturbed by the Raman beams is

$$\mathbf{H} = \frac{\hbar\omega_{ge}}{2} (|e\rangle\langle e| - |g\rangle\langle g|) + \frac{\hbar\boldsymbol{\mu}_{ge} \cdot (\mathbf{E}_U + \mathbf{E}_L)}{2} (|e\rangle\langle g| + |g\rangle\langle e|), \quad (5.24)$$

where ω_{ge} is the transition frequency and $\boldsymbol{\mu}_{ge}$ is the dipole transition matrix element. Writing this in terms of the Rabi frequencies $\Omega_{U,L} = (\boldsymbol{\mu}_{ge} \cdot \hat{\epsilon}_{U,L}) E/\hbar$, and using the transformation

$$\mathbf{H} \Rightarrow \tilde{\mathbf{H}} = e^{i\omega_U t \sigma_z / \hbar} \mathbf{H} e^{-i\omega_U t \sigma_z / \hbar} \quad (5.25)$$

to move into a frame rotating at ω_U , the field interaction Hamiltonian is given by

$$\tilde{\mathbf{H}} = \frac{\hbar}{2} \begin{pmatrix} \Delta & \Omega_U e^{-i(\mathbf{k}_U \cdot \mathbf{r})} + \Omega_L e^{-i(\mathbf{k}_L \cdot \mathbf{r} - \mu_R t)} \\ \Omega_U e^{i(\mathbf{k}_U \cdot \mathbf{r})} + \Omega_L e^{i(\mathbf{k}_L \cdot \mathbf{r} - \mu_R t)} & \Delta \end{pmatrix}, \quad (5.26)$$

where $\Delta = \omega_{ge} - \omega_U$ and the rotating wave approximation has been used to remove terms oscillating at frequencies $\sim 2\omega_U$. Solving for the eigenvalues of this Hamiltonian gives the energies of the upper and lower levels,

$$\begin{aligned} \tilde{E}_{\pm} &= \pm \frac{1}{4} \sqrt{\Delta^2 + |\Omega_U|^2 + |\Omega_L|^2 + \Omega_U \Omega_L^* e^{i(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t)} + \Omega_U^* \Omega_L e^{-i(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t)}} \\ &\approx \pm \frac{\Delta}{2} \pm \frac{|\Omega_U|^2 + |\Omega_L|^2 + 2 \operatorname{Re}(\Omega_U^* \Omega_L) \cos(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t)}{4\Delta}. \end{aligned} \quad (5.27)$$

The first term is the unperturbed energy of the levels in the field interaction representation, while the second represents the AC Stark shift produced by the Raman beams, the sign of which depends on the detuning Δ . The AC Stark shift due to each allowed transition can be found given the Rabi frequencies $|\Omega_{U,L}|$ and detuning from the transition resonance, Δ .

In general, the Rabi frequency of an electric dipole transition with $q = m_j - m'_j$ is

$$\Omega_0 = \frac{e|E|}{\hbar} \sqrt{\frac{3(2j'+1)A_{12}}{4\alpha k_{12}^3 c}} \left[\begin{pmatrix} j & 1 & j' \\ -m_j & q & m'_j \end{pmatrix} \mathbf{c}^{(q)} \cdot \hat{\epsilon} \right] \quad (5.28)$$

where e is the electronic charge, A_{12} the Einstein A coefficient, k_{12} the transition wavevector, α the fine structure constant, c the speed of light and $j(j')$ the ground(excited) state angular momentum. The final, geometric term gives contains the $3 - j$ symbol and the Racah tensor $\mathbf{c}^{(q)}$ that defines the polarisation dependence of the transition (see Appendix A

for details). The Rabi frequency of a given transition $m_j \rightarrow m'_j$ due to the upper or lower beam is then

$$\Omega_{U,L} \{m_j \rightarrow m'_j\} = |E| \kappa_{(m_j \rightarrow m'_j)} \mathbf{c}^{(\mathbf{m}_j - \mathbf{m}'_j)} \cdot \hat{\mathbf{e}}_{U,L} \quad (5.29)$$

where the $3-j$ symbol has been combined with other constants as $\kappa_{(A \rightarrow B)}$. The Rabi frequencies for the allowed transitions are then

$$\begin{aligned} \Omega_{U,L} \{-1/2 \rightarrow -1/2\} &= |E| \kappa_A \times \left(\cos \frac{\theta_R}{2} \cos \phi_p \right) \\ \Omega_{U,L} \{-1/2 \rightarrow -3/2\} &= |E| \kappa_B \times \frac{-1}{\sqrt{2}} \left(\pm \sin \frac{\theta_R}{2} \cos \phi_p - i \sin \phi_p \right) \\ \Omega_{U,L} \{-1/2 \rightarrow 1/2\} &= |E| \kappa_C \times \frac{-1}{\sqrt{2}} \left(\pm \sin \frac{\theta_R}{2} \cos \phi_p + i \sin \phi_p \right) \\ \Omega_{U,L} \{1/2 \rightarrow 1/2\} &= |E| \kappa_D \times \left(\cos \frac{\theta_R}{2} \cos \phi_p \right) \\ \Omega_{U,L} \{1/2 \rightarrow -1/2\} &= |E| \kappa_E \times \frac{-1}{\sqrt{2}} \left(\pm \sin \frac{\theta_R}{2} \cos \phi_p - i \sin \phi_p \right) \\ \Omega_{U,L} \{1/2 \rightarrow 3/2\} &= |E| \kappa_F \times \frac{-1}{\sqrt{2}} \left(\pm \sin \frac{\theta_R}{2} \cos \phi_p + i \sin \phi_p \right). \end{aligned} \quad (5.30)$$

Combining (5.27) and (5.30), we find the total shift of the upper qubit level to be

$$\begin{aligned} \Delta_{ACSS}(\uparrow) &= \frac{|E|^2}{4} \left\{ \left[2 \cos^2 \frac{\theta_R}{2} \cos^2 \phi_p (1 + \cos(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t)) \right] \times \left(\frac{\kappa_A^2}{\Delta_A} \right) + \right. \\ &\quad \left[\left(\sin^2 \phi_p + \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) + \left(\sin^2 \phi_p - \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) \cos(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t) \right] \times \left(\frac{\kappa_B^2}{\Delta_B} + \frac{\kappa_C^2}{\Delta_C} \right) \Big\} \\ &\quad (5.31) \end{aligned}$$

and the lower level to be

$$\begin{aligned} \Delta_{ACSS}(\downarrow) &= \frac{|E|^2}{4} \left\{ \left[2 \cos^2 \frac{\theta_R}{2} \cos^2 \phi_p (1 + \cos(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t)) \right] \times \left(\frac{\kappa_D^2}{\Delta_D} \right) + \right. \\ &\quad \left[\left(\sin^2 \phi_p + \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) + \left(\sin^2 \phi_p - \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) \cos(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t) \right] \times \left(\frac{\kappa_E^2}{\Delta_E} + \frac{\kappa_F^2}{\Delta_F} \right) \Big\}. \end{aligned} \quad (5.32)$$

We note that the AC Stark shifts due to individual horizontally polarised (driving σ_{\pm} transitions) and vertically polarised beams (driving π transitions), incident in the radial

plane are

$$\begin{aligned}
 A_{\uparrow} &= \frac{|E|^2}{4} \left(\frac{\kappa_A^2}{\Delta_A} \right) & (\pi) \\
 B_{\uparrow} &= \frac{|E|^2}{8} \left(\frac{\kappa_B^2}{\Delta_B} + \frac{\kappa_C^2}{\Delta_C} \right) & (\sigma_{\pm}) \\
 A_{\downarrow} &= \frac{|E|^2}{4} \left(\frac{\kappa_D^2}{\Delta_D} \right) & (\pi) \\
 B_{\downarrow} &= \frac{|E|^2}{8} \left(\frac{\kappa_E^2}{\Delta_E} + \frac{\kappa_F^2}{\Delta_F} \right) & (\sigma_{\pm}).
 \end{aligned} \tag{5.33}$$

We consider the constant and time dependent parts of (5.31) and (5.32) separately. Taking the constant part, the time-independent differential AC Stark shift is given by

$$\begin{aligned}
 \Delta_{ACSS}^0 &= \Delta_{ACSS}^0(\uparrow) - \Delta_{ACSS}^0(\downarrow) \\
 &= 2 \left[\cos^2 \left(\frac{\theta_R}{2} \right) \cos^2 \phi_p \right] (A_{\uparrow} - A_{\downarrow}) + 2 \left[\sin^2 \left(\frac{\theta_R}{2} \right) \cos^2 \phi_p + \sin^2 \phi_p \right] (B_{\uparrow} - B_{\downarrow}).
 \end{aligned} \tag{5.34}$$

Experimentally, it is beneficial to choose a polarisation that nulls this shift, to avoid dephasing due to laser intensity noise and unwanted σ_z rotations. This can be achieved providing $(A_{\uparrow} - A_{\downarrow})$ and $(B_{\uparrow} - B_{\downarrow})$ have opposite signs and

$$\theta_R < \arccos \left(\left| \sqrt{\frac{4(B_{\downarrow} - B_{\uparrow})}{(A_{\uparrow} - A_{\downarrow}) + (B_{\downarrow} - B_{\uparrow})}} \right| \right). \tag{5.35}$$

The angle ϕ_p that makes $\Delta_{ACSS} = 0$ depends on the detuning of the Raman beams within the $P_{3/2}$ manifold. The angle is symmetric about zero so we are free to choose positive or negative solutions $\pm\phi_p$, and as will be shown, it is this freedom that allows us to produce the ODF. Choosing $\phi_p(U) = +\phi_p$ and $\phi_p(L) = -\phi_p$, the time-dependent parts of (5.31) and (5.32) are, respectively

$$\begin{aligned}
 \Delta_{ACSS}^T(\uparrow) &= \left[\cos^2 \frac{\theta_R}{2} \cos^2 \phi_p A_{\uparrow} - \left(\sin^2 \phi_p + \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) B_{\uparrow} \right] 2 \cos(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t) \\
 \Delta_{ACSS}^T(\downarrow) &= \left[\cos^2 \frac{\theta_R}{2} \cos^2 \phi_p A_{\downarrow} - \left(\sin^2 \phi_p + \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) B_{\downarrow} \right] 2 \cos(\mathbf{k}_R \cdot \mathbf{r} - \mu_R t)
 \end{aligned} \tag{5.36}$$




Figure 5.5: Plot of optical dipole force ratio $R = F_{\uparrow}/F_{\downarrow}$ at different detunings Δ'_R of the Raman frequencies from resonance with the $P_{3/2, +1/2}$ level at $B = 4.5$ T (blue). Also shown in orange is the polarisation $\pm\phi_p$ needed to null the time-averaged differential qubit shift at each detuning. The force ratio approaches 2:1 as the Raman lasers approach resonance with the $P_{3/2, +1/2}$ or $P_{3/2, -1/2}$ levels. However, near resonance the polarisation angle required to null the differential shift asymptotically approaches 90° , making the system very sensitive to imperfectly polarised light, and for robust performance a smaller ratio must be chosen.

leading to spin-dependent optical dipole forces of

$$\begin{aligned}
 F_{\uparrow}(t) &= -2\mathbf{k}_{\mathbf{R}} \left[\cos^2 \frac{\theta_R}{2} \cos^2 \phi_p A_{\uparrow} - \left(\sin^2 \phi_p + \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) B_{\uparrow} \right] \sin(\mathbf{k}_{\mathbf{R}} \cdot \mathbf{r} - \mu_R t) \\
 F_{\downarrow}(t) &= -2\mathbf{k}_{\mathbf{R}} \left[\cos^2 \frac{\theta_R}{2} \cos^2 \phi_p A_{\downarrow} - \left(\sin^2 \phi_p + \sin^2 \frac{\theta_R}{2} \cos^2 \phi_p \right) B_{\downarrow} \right] \sin(\mathbf{k}_{\mathbf{R}} \cdot \mathbf{r} - \mu_R t).
 \end{aligned}
 \tag{5.37}$$

Though each linear polarisation nulls the qubit shift, the interference between two crossed polarisations produces a polarisation gradient along the optical lattice, such that the polarisation is periodically circular. While the time-averaged Stark shift remains zero, the spin states experience a periodic differential force. In the configuration shown in 5.3, $\text{sgn}(A_{\uparrow}) \equiv \text{sgn}(A_{\downarrow})$ and $\text{sgn}(B_{\uparrow}) \equiv \text{sgn}(B_{\downarrow})$, so $\text{sgn}(F_{\uparrow}) \equiv -\text{sgn}(F_{\downarrow})$. The ratio of these forces, $R = F_{\uparrow}/F_{\downarrow}$ can varied between approximately -0.5 and -2 through adjustment of the detuning Δ_R . The ability to set non-unity force ratios is critical to the protocols described later in this chapter.

A plot of the force ratio that results as a function of detuning from the $P_{3/2, +1/2}$ level is given in Figure 5.5, which also shows that as the ratio increases, the system becomes increasingly sensitive to errors in the polarisation chosen to null the qubit shift.

5.5 Dynamics of Two-Dimensional Coulomb Crystals

While the oscillating quadrupole axialisation field provides the coupling between the radial modes necessary to laser cool the radial motion, when working with large crystals of ions it is more useful to introduce a *rotating* quadrupole in the radial plane. The effect of this ‘rotating wall’ drive is to increase the confinement in one direction in the radial plane, while decreasing it in the other, squeezing the crystal along the quadrupole axis. As the field rotates, the crystal is pulled with it, phase locked to the RF drive.⁶

It is normal to consider the dynamics of the Coulomb crystal in a frame rotating with the crystal, in which both the radial and axial potentials are harmonic. Following the derivation in Wang et al [105], we describe a crystal of N ions, with the coordinates of the j^{th} ion in the rotating frame given by \mathbf{r}_j^R , with the Lagrangian

$$L^R = \sum_{j=1}^N \left[\frac{1}{2} m |\dot{\mathbf{r}}_j^R|^2 - \frac{eB_{eff}[\Omega]}{2} (\dot{x}_j^R y_j^R - \dot{y}_j^R x_j^R) - e\phi_j^R \right], \quad (5.38)$$

where e is the elementary charge, m is the mass of a single ion, Ω the frequency of the rotating wall, ϕ_j^R the effective potential energy in the rotating frame and $B_{eff}[\Omega] = B_z - 2\Omega m/e$ an effective magnetic field, comprising the axial trap field in the laboratory frame, B_z , and a contribution due to the Coriolis force in the rotating frame.

At equilibrium in the rotating frame, both the kinetic energy term and effective Lorentz energy term are zero. The equilibrium position of the ions in the crystal is therefore the configuration that minimizes the effective potential

$$\begin{aligned} e\phi_j^R = & eV_0 (z_j^R)^2 + \frac{1}{2} (eB_z\Omega - m\Omega^2 - eV_0) \left[(x_j^R)^2 + (y_j^R)^2 \right] \\ & + eV_W \left[(x_j^R)^2 - (y_j^R)^2 \right] + \frac{k_e e^2}{2} \sum_{k \neq j} \frac{1}{(R_{jk}^0)^2}, \end{aligned} \quad (5.39)$$

where k_e is Coulomb’s constant, $R_{jk}^0 = |\mathbf{r}_j - \mathbf{r}_k|$ is the distance between ions j and k , V_0 is the static electric potential (Φ_0/d_0^2 , where d_0 is the characteristic dimension of the trap) and V_W is the rotating wall potential, which is generally small and will be ignored for the purposes of this calculation.

With the equilibrium positions determined from minimisation of Equation 5.39, the small-amplitude dynamics can be captured by expanding the Lagrangian around these coordinates in terms of small displacements and velocities, with $\mathbf{r}_j = \mathbf{r}_j^0 + \delta\mathbf{r}_j$ and $\dot{\mathbf{r}}_j = \delta\dot{\mathbf{r}}_j$. If the trap axis is aligned to the magnetic field, then the Lorentz force terms only apply to the radial modes and the axial and radial dynamics can be decoupled.

⁶Providing the RF drive amplitude is high enough. A balance must be struck to find a drive amplitude high enough to prevent phase slippage but low enough to avoid significant distortion of the crystal geometry.

Expanding (5.38) up to second-order (the first order terms are zero around an equilibrium point), it can be shown that the Lagrangian which governs the axial motion of the ions is given by

$$L_{ph}^A = \sum_{j=1}^N m \delta \dot{z}_j^2 - \frac{1}{2} \sum_{j,k=1}^N K_{jk}^A \delta z_j \delta z_k, \quad (5.40)$$

where K_{jk}^A is the axial stiffness matrix:

$$K_{jk}^A = \begin{cases} 2eV_0 - \sum_{l=1, l \neq j}^N \frac{k_e e^2}{(R_{lj}^0)^3} & j = k \\ \frac{k_e e^2}{(R_{lj}^0)^3} & j \neq k \end{cases}. \quad (5.41)$$

The equation of motion in the axial direction can therefore be written

$$m \delta \ddot{z}_j + \sum_{k=1}^N K_{jk}^A \delta z_k = 0 \quad j = 1, 2 \dots N. \quad (5.42)$$

As the trap is harmonic, we make the ansatz solution

$$\delta z_j(t) = b_j^\nu \cos(\omega_\nu t), \quad (5.43)$$

where b_j^ν is the N -element mode eigenvector to be found. Substituting (5.43) into (5.42) provides the eigenvector equation

$$\sum_{k=1}^N \left[m \omega_\nu^2 \delta_{jk} - K_{jk}^A \right] b_k^\nu = 0 \quad j, \nu = 1, 2 \dots N \quad (5.44)$$

The mode eigenvectors are therefore identical to the eigenvectors of the stiffness matrix, but with eigenvalues $\omega_{z\nu}$ related to those of the stiffness matrix, $\lambda_{z\nu}$ by

$$\omega_{z\nu} = \sqrt{\lambda_{z\nu}/m}. \quad (5.45)$$

For the examples given later in this chapter, I will be considering a planar crystal consisting of a pentagon of five ions surrounding a central ion. This forms naturally in a Penning trap when the axial confinement is sufficiently high. Solving Equation 5.44 for such a configuration, we find six transverse modes including two degenerate pairs, shown in the




Figure 5.6: Transverse mode frequencies (ω_1 : red, ω_2 : blue, ω_3 : orange, ω_4 : pink) and inter-ion distance (dashed black) vs crystal rotation frequency, at the $1 \rightarrow 2$ layer threshold.

matrix \mathbf{A}^T in order of decreasing frequency

$$\begin{pmatrix}
 1 & 0 & 0 & 0 & 0 & -5 \\
 1 & (G-2) & 0 & -(G+1) & (1-G) & 1 \\
 1 & 2(1-G) & G & 2G & 0 & 1 \\
 1 & (G-2) & -G & -(G+1) & (G-1) & 1 \\
 1 & 1 & -1 & 1 & -1 & 1 \\
 1 & 1 & 1 & 1 & 1 & 1
 \end{pmatrix}$$

ω_1

ω_2

ω_3

ω_4

where $G = (1 + \sqrt{5})/2$ and the mode vectors are shown un-normalised for clarity. Note the modes of a rotationally symmetric crystal are fixed only up to a rotational permutation of the ions - the exact choice of basis is arbitrary.

These mode vectors are constant while the equilibrium configuration remains unchanged, but the eigenfrequencies and lattice constant of the crystal are functions of the axial confinement and rotation frequency. As the axial confinement is relaxed, the planar crystal reaches a threshold where it becomes energetically unfavourable for ions to remain in one plane. This threshold voltage increases with crystal rotation frequency.

At the threshold, the mode eigenfrequencies are maximally separated, with the lowest frequency mode approaching zero frequency. Conversely, for very high axial confinement, inter-ion interactions become small compared to the external potential and the modes become near-degenerate. As will be described in the following section, frequency addressing of modes is often necessary and in this situation confinement close to the threshold is desirable. The mode frequencies and inter-ion distance for a 6 ion crystal just above the monoplanar threshold are given in Figure 5.6.




Figure 5.7: Schematic of optical dipole force and readout lasers for ${}^9\text{Be}^+$. Readout is performed on the closed-cycle $S_{1/2, +1/2} \rightarrow P_{3/2, +3/2}$ transition, although off-resonant processes can lead to decay to the lower qubit level, which is repumped by a second laser. This laser is also used to optically pump population to $|\uparrow\rangle$ during state preparation. Microwaves at the qubit frequency ($\omega_0 = 28 \text{ GHz T}^{-1}$) provide for collective local rotations of the qubits in the register. For details of the optical dipole force Raman lasers, see Section 5.4.

5.6 Global Unitary Operations on a Coulomb Crystal

I now consider a qubit register in a two-dimensional ion Coulomb crystal. A basis state of the qubit register of N ions is written in the spin basis as

$$|s\rangle_R \equiv |s_1\rangle \otimes |s_2\rangle \otimes \cdots \otimes |s_N\rangle \quad (5.46)$$

for $s_j = \uparrow, \downarrow$. Arbitrary register states are then written as

$$|\chi\rangle_R = \sum_s r_s |s\rangle_R \quad (5.47)$$

with $\sum_s |r_s|^2 = 1$.

Projective readout of the physical qubits in the computational basis is given by fluorescence from a laser resonant with a suitable dipole transition, for instance $S_{1/2} \rightarrow P_{3/2}$ (Figure 5.7). This transition is also used for Doppler cooling the crystal. A microwave field at the qubit frequency ω_0 permits global Pauli X and Y rotations, while Pauli Z rotations are effected by applying the laser normally used for readout, detuned far from resonance, producing an AC Stark shift of the qubit levels.

It is possible to exploit the coupling between transverse mechanical modes of the Coulomb crystal and spin-dependent optical dipole forces (ODFs) to realise complex unitary operations on the qubit register. I will show that it is possible to generate global unitary




Figure 5.8: Trajectories taken through phase space for two spin states under the application of a uniform force tuned near the centre of mass mode. The generalised force applied to the mode when in state $s_2 = |\uparrow\uparrow\uparrow\uparrow\rangle$ is 5/3 that in $s_1 = |\downarrow\downarrow\downarrow\downarrow\rangle$ and s_2 acquires a phase 25/9 times greater.

operations of the form

$$U(\phi) = \sum_s e^{i\phi_s} |s\rangle_R \langle s|_R, \quad (5.48)$$

where ϕ is a vector of phases ϕ_s determined by the spin-mode couplings.

Writing the set of mode eigenvectors as a matrix $\mathbf{A} = \sum_i |i\rangle \langle \mathbf{a}_m|$ with elements A_{mi} , we can write the coordinates of the ions, r_i , as a superposition of generalised modal coordinates q_m , and similarly decompose any forces f_i on the individual ions into generalised forces Q_m acting on each mode m :

$$q_m = \sum_i A_{mi} r_i \quad Q_m = \sum_i A_{mi} f_i. \quad (5.49)$$

Given a uniform ODF beam across the crystal, we express the forces on the whole register as a matrix \mathbf{F} , with elements F_{is} representing the forces on ion i for each register basis state $|s\rangle_R$ in $|\chi\rangle_R$. The product of the modal transformation matrix \mathbf{A} and this new dipole force matrix gives the spin-mode coupling matrix \mathbf{M} with elements

$$M_{ms} = \sum_i A_{mi} F_{is}. \quad (5.50)$$

The matrix \mathbf{M} defines the generalised force acting on each normal mode for each state $|s\rangle_R$ of the register. The driven system will evolve in phase space, as described in Section 5.3, but the phase space coordinates are now the generalised mode coordinates q_m and their conjugate momenta p_m , and the driving forces are the corresponding generalised forces, Q_m . All the modes are orthogonal and contribute distinct phases associated with their individual excitations. However, in the case of a symmetric crystal, many of the modes will have a degenerate partner of the same frequency. Though orthogonal, a pair of degenerate modes will be driven simultaneously, so it is proper to consider the Pythagorean sum of the generalised forces on each pair and the corresponding rows in \mathbf{M} are combined in quadrature. This is physically equivalent to assuming each mode is independently driven and then summing the acquired phases for the pair.

We apply a set of pulses near-resonant with each mode such that $\mu_R = \omega_m - \delta$, with pulse areas proportional to the elements v_m , where we use the freedom to adjust the detuning δ , pulse duration and intensity to ensure the loops close. The geometric phases acquired by each basis state in $|\chi\rangle_R$ are proportional to the square of the generalised forces (see equation 5.20 and Figure 5.8), and given by

$$\phi_s = \sum_m (M_{ms})^2 v_m, \quad (5.51)$$

determining the unitary of Eqn. (5.48).

5.6.1 Application to quantum error correction

I now show how this technique can be used to prepare and read-out the five-qubit code (5QC) and the five-qubit repetition code (5RC), in a six-ion register. The 5QC protects the encoded information against arbitrary Bloch sphere rotations suffered by a single physical qubit of the register. Alternatively, the 5RC protects against dephasing channel errors acting on up to two physical qubits, but cannot correct for spin-flip noise. In many ion qubits, the T_1 lifetime is several hundred seconds. In these cases, the 5RC provides sufficient protection against a realistic noise model.

I consider the example of a qubit register of six ions in a crystal in the configuration shown in Fig. 5.9(a). The qubit of the central ‘hub’ ion is initialised in the logical state

$$|\psi\rangle_H = \alpha |\downarrow\rangle_H + \beta |\uparrow\rangle_H. \quad (5.52)$$

The five ‘code’ ions in the ring around the hub will encode the quantum information. They are initially prepared in the state $|+++++\rangle$, which can be written as

$$|S\rangle_C = \sum_s |s\rangle_C. \quad (5.53)$$

Given the ability to prepare the initial state

$$|\chi_0\rangle_R = |\psi\rangle_H \otimes |S\rangle_C \quad (5.54)$$

and to measure qubits of the register in the Pauli- X basis, we require only two global unitary operations, U_{ring} and U_{spokes} , to encode and read out the 5QC, and only U_{spokes} to encode and read out 5RC. The unitaries $U_{\text{spokes}}(U_{\text{ring}})$ will globally apply a controlled phase gate between each of the qubits in the register that share a grey edge in Fig. 5.9(b)(5.9(c)). I summarise the error-correcting protocol in the caption of Fig. 5.9.

Both protocols require the ability to prepare and measure the hub qubit independently of the code qubits and vice versa. This could either be achieved by physically addressing the ions with a focussed beam, or with a uniform beam but with two isotopes of the same ion species used for hub and ring. Here the isotopic shift would be sufficient to allow resonant frequency addressing during state preparation and readout, without significantly affecting the coupling to the off-resonant entangling pulses.⁷

Given suitable control to prepare the initial state and the ability to perform the necessary measurements, we also require suitable ϕ_s to perform U_{ring} and U_{spokes} . There are five free parameters with which to engineer the desired unitary, the pulse areas, P_i , applied to each of the four modes, and the force ratio, R . In general, it is difficult to find suitable pulse sequences, but by exploiting the freedom to include a global phase and to find elements of

⁷Naturally this requires an ion species with several stable isotopes, such as magnesium. Beryllium would not be suitable in this case.




Figure 5.9: The protocol. (a) The crystal is initialised in the product state, $|\chi_0\rangle_R$. (b) Unitary U_{spokes} is performed. This entangles the hub qubit to the 5RC codespace. (c) The central ion is measured in the Pauli- X basis to teleport the logical information onto the blue ring ions, completing the encoding of 5RC. To prepare 5QC, U_{ring} is also applied immediately after the measurement, mapping codewords of 5RC to 5QC. (d) (Omitted for 5RC) To begin readout for 5QC, U_{ring} is applied to map 5QC codewords to 5RC codewords. (e) The central qubit is then prepared in the $|+\rangle = \frac{1}{\sqrt{2}}(|\uparrow\rangle + |\downarrow\rangle)$ state and U_{spokes} is performed once more, entangling the codespace to the hub. (f) Measuring the code qubits teleports the quantum information, up to a Pauli error, back to the central blue ion. As the Pauli- X measurements performed commute with the stabilizers of 5RC, the fluorescence pattern of the measured bright and dark ions, shown in red and yellow, provides syndrome information. This information is used to apply the appropriate Pauli correction to the central ion, recovering the initial state $|\chi_0\rangle$. Full details of the possible (correctable) errors, syndromes and corrections are given in Table 5.1.

Projected Error on 5RC	Fluorescence Syndrome	Single-Qubit Pauli Correction
<i>IIII</i>	□□□□	<i>I</i>
<i>IIII</i>	■□□□	<i>X</i>
<i>IIZI</i>	□□■□	<i>I</i>
<i>IIZI</i>	■□■□	<i>X</i>
<i>ZIII</i>	■□□■	<i>I</i>
<i>ZIII</i>	□■□■	<i>X</i>
<i>IZIZ</i>	□■□■	<i>I</i>
<i>IZIZ</i>	■□■□	<i>X</i>
Projected Error on 5QC	Fluorescence Syndrome	Single-Qubit Pauli Correction
<i>IIII</i>	□□□□	<i>I</i>
<i>IIII</i>	■□□□	<i>X</i>
<i>IIZI</i>	□□■□	<i>I</i>
<i>IIZI</i>	■□■□	<i>X</i>
<i>IYII</i>	■□□■	$-Y$
<i>IYII</i>	□■□■	iZ
<i>IXII</i>	□■□■	<i>Z</i>
<i>IXII</i>	■□■□	iY

Table 5.1: Errors, syndromes and the corresponding corrections to the final teleported state, for 5RC and 5QC. Syndromes are given as sequences of bright (□) and dark (■) ions around the ring during the final readout stage.

ϕ_s modulo 2π it is possible to find a very large number of sparsely distributed solutions for these two unitaries and we are free to choose solutions that behave well under predicted experimental noise.

One solution, (P_1, P_2, P_3, P_4, R) , for U_{spokes} is

$$(3.1250, 2.6043, 2.6043, 0, 1.4000), \quad (5.55)$$

where the pulse areas are normalised such that a centre-of-mass pulse $P_1 = 1$ would produce a π radian phase shift on the $|\downarrow\rangle_H \otimes |\downarrow\downarrow\downarrow\downarrow\rangle_C$ state. Additionally we find pulses

$$(10.9862, 7.6768, 19.6502, 10.9861, 0.6737) \quad (5.56)$$

will produce U_{ring} . It would also be possible to vary the ratio of forces for each pulse, but for the present purposes it is sufficient to keep this constant throughout the sequence.

Table 5.2 shows how each basis state evolves under the application of the series of pulses for U_{spoke} , from the initial product state in the leftmost column, towards the target graph state in the rightmost column. The columns list the absolute phases and phase angles after each pulse. The final row lists the fidelity of each intermediate state relative to the target.

Basis $ s_i\rangle$ + $Perms$	Initial Phase/ π ϕ_i	Phase after ω_1 Pulse/ π ϕ_i (mod 2π)	Phase after ω_2 Pulse/ π ϕ_i (mod 2π)	Phase after ω_3 Pulse/ π ϕ_i (mod 2π)	Target Phase/ π ϕ_i
$ 0\rangle 0000\rangle$	0	3.125	3.125	3.125	0
$ 0\rangle 0001\rangle$	0	1.125	2.125	3.125	0
$ 0\rangle 0011\rangle$	0	0.125	2.743	3.125	0
$ 0\rangle 0010\rangle$	0	0.125	0.507	3.125	0
$ 0\rangle 0011\rangle$	0	0.125	2.743	3.125	0
$ 0\rangle 0101\rangle$	0	0.125	0.507	3.125	0
$ 0\rangle 0111\rangle$	0	1.125	2.125	3.125	0
$ 0\rangle 1111\rangle$	0	3.125	3.125	3.125	0
$ 1\rangle 0000\rangle$	0	1.125	1.125	1.125	0
$ 1\rangle 0001\rangle$	0	0.125	1.125	2.125	1
$ 1\rangle 0011\rangle$	0	0.125	2.743	3.125	0
$ 1\rangle 0010\rangle$	0	0.125	0.507	3.125	0
$ 1\rangle 0011\rangle$	0	1.125	3.743	4.125	1
$ 1\rangle 0101\rangle$	0	1.125	1.507	4.125	1
$ 1\rangle 0111\rangle$	0	3.125	4.125	5.125	0
$ 1\rangle 1111\rangle$	0	6.125	6.125	6.125	1
Fidelity	0.5	0.25	0.634	1	-

Table 5.2: Phases ϕ_i acquired by basis states $|s_i\rangle$ after each stage of the composite entangling pulse. Note that the final state is equivalent to the target up to a global phase of $\epsilon = 1.125\pi$. The fidelity with the target state at each stage is shown on the final row. Note that all cyclic permutations of the ring qubits are grouped as these acquire the same phase.

5.6.2 Performance under noise

These results show that we can execute both QECC protocols using a low number of pulses. Short pulse sequences are advantageous as errors in the protocol are minimised. By numerically simulating both protocols in the presence of environmental noise and noisy operations it is possible to estimate the experimental requirements necessary to realise improved quantum coherence compared with the coherence of a single qubit.⁸

I now consider the decoherence of the encoded state over some time t , normalised to the coherence time of the qubit, in the case where the protocol is performed using imperfect unitary operations. The 5RC protocol is numerically simulated using the imperfect U_{spokes}^σ unitary, where the optimal pulse sequence (5.55) for U_{spokes} is disrupted by replacing pulse areas P_j and force ratio R with $\epsilon_j(\sigma)P_j$ and $\epsilon_R(\sigma)R$ respectively, where $\epsilon_j(\sigma)$ and $\epsilon_R(\sigma)$ are a random variables chosen from a normal distribution with standard deviation σ and unity mean.

Additionally, environmental noise is modelled by allowing the spin of each individual ion ρ to dephase via the map

$$\gamma_t(\rho) = (1 + e^{-t})\rho/2 + (1 - e^{-t})Z\rho Z/2, \quad (5.57)$$

where Z is the Pauli-Z matrix. The model assumes that each application of U_{spokes}^σ is performed in time $t_U = 5 \cdot 10^{-4}$, equivalent to a $50 \mu\text{s}$ gate for qubits with 100 ms coherence time. This is modelled by applying the dephasing map for an additional t_U before each instantaneous application of U_{spokes}^σ .

We measure decoherence using the average fidelity, \mathcal{F} , calculated with ~ 500 random samples of the initial state $|\psi\rangle_H$ and noisy unitary U_{spokes}^σ for each given t . We plot

$$\mathcal{F}(\rho_H^f) = \left| \text{Tr}[\rho_H^f |\psi\rangle \langle \psi|_H] \right| \quad (5.58)$$

where ρ_H^f is the state that is read out from the code after the indicated Pauli correction has been applied. Note that the simulation models the ‘round-trip’ process, including both noisy encoding and decoding. The results for 5RC are shown in Fig. 5.10.

To quantify the improvement in coherence time attained using the code, the high-fidelity time, τ , is defined to be the time the prepared code maintains $\mathcal{F} > 0.99$. This is then compared to τ_{1Q} , the high-fidelity time of a single qubit. The inset of Fig. 5.10 shows τ/τ_{1Q} for 5RC as a function of σ . We observe scaling of the form

$$\tau/\tau_{1Q} = \tau_0(2 - e^{\sigma^2/\alpha}) \quad (5.59)$$

where $\tau_0 = \tau/\tau_{1Q}$ is the code-dependent optimal coherence time improvement that is observed in the noiseless $\sigma = 0$ case. For 5RC, we find $\alpha \sim 5.4 \cdot 10^{-4}$, and $\tau_0 \sim 6.92$.

⁸These simulations were performed in collaboration with B. Brown and H. Dale.




Figure 5.10: Average fidelity as a function of time with imperfectly prepared 5RC states. Here, simulations with pulse noise of $\sigma = 0, 0.005, 0.01$ and 0.015 are shown in red, yellow and green and blue respectively. Also plotted, in blue, is the fidelity of a single qubit under the dephasing map. The dashed line marks 0.99 fidelity. The inset shows the ratio of the ‘round-trip’ (encode-wait-decode) high-fidelity (>0.99) time of the encoded system to that of an unencoded physical qubit, τ/τ_{1Q} , plotted as a function of σ . A fit to the data yields $\tau/\tau_{1Q} \sim \tau_0(2 - \exp(1858\sigma^2))$.

We perform a similar simulation for 5QC, using disrupted U_{ring} and U_{spoke} unitaries. In this case we consider a situation where both T_1 and T_2 decays occur, modelling the decoherence via the depolarising map

$$\xi_t(\rho) = (1 + 3e^{-t})\rho/4 + (1 - e^{-t})(X\rho X + Y\rho Y + Z\rho Z)/4. \quad (5.60)$$

Under these conditions we find $\alpha \sim 3.0 \cdot 10^{-5}$, and $\tau_0 \sim 2.45$.

We extrapolate the present data to the value $\tau/\tau_{1Q} = 1$ to obtain a threshold pulse noise, σ_{th} , of $\sigma_{\text{th}} \sim 0.0038$ for 5QC, and $\sigma_{\text{th}} \sim 0.018$ for 5RC. Producing gates that meet these threshold values is feasible with existing technology.

It may also be possible to make further improvements to the coherence time of the code by repeating the protocol periodically, teleporting the stored information back and forth between the hub qubit and the code and thus suppressing errors before significant decoherence has occurred. The limitation of this scheme is that the information is unprotected for the period that it remains on the hub, which is likely to be at least 1 ms while readout of the ring is performed. A more robust scheme would teleport the information between two similar codes so that it remains protected throughout. Our ongoing work investigates whether this may be possible using multiple concentric rings of a larger planar crystal.

5.6.3 Key experimental considerations

The errors on the pulse area and force ratio given in the previous section appear as single values, but may be introduced by a range of physical variables, such as the laser intensity, laser frequency, and beam polarisations. Many of these considerations are shared with conventional two qubit gate experiments, and this experiment would benefit from the accumulated knowledge regarding the control of common noise sources whilst also suffering from similar limitations. Of course, these problems are somewhat mitigated as the scheme uses far fewer pulses than protocols based on two-qubit gates, reducing the threshold fidelity requirements per operation. The scheme also avoids individual ion addressing, one of the most notoriously challenging techniques.

In certain aspects Penning traps reliably outperform RF trap experiments, and errors on state preparation and read-out (with suitably high NA optics) are likely to be vanishingly small, while electromagnetic interference at the position of the trap is typically suppressed by -10 dB due to the Meissner effect of the superconducting solenoid. The heating rate in Penning traps should also be very low. Although in large crystals this is thought to be limited by soft collisions with background gas [98], with sufficient care very pure vacuums can be produced and, ultimately, cryogenic temperatures can eliminate this problem entirely.

The particular arrangement of lasers used to generate the optical dipole force across a large crystal comes replete with its own challenges. Probably the most important decision is in choosing the beam crossing angle, which once engineered into the trap cannot easily be altered. This angle sets the wavevector of the lattice and puts bounds on several other aspects of the performance. Fast gates require strong optical dipole forces, and to generate these without introducing significant off-resonant scattering from the Raman beams requires a steep intensity gradient and thus a large wavevector. However, very tight lattices require lower temperatures or higher frequency traps to ensure Lamb-Dicke confinement within the lattice. Another significant difficulty lies in ensuring precise alignment of the lattice wavefronts with the plane of the ions, and satisfying this condition becomes harder with increasing wavevector. Fortunately, work tackling such issues is already underway in the groups at NIST and Sydney, and the original NIST experiment [21] has proven that such methods can be realised.

One final problem that relates more uniquely to the experiments described in this chapter is the unwanted off-resonant excitation of spectator modes when driving entangling pulses.⁹ For high gate speeds (essential if a useful error-correcting code is to be produced), the optical dipole force must be detuned as far as possible from the target mode, as $t = 2\pi n\delta$. However, as the detuning increases, the relative suppression of neighbouring modes is decreased and some degree of excitation is unavoidable. For certain choices of detuning, it is possible to find periods over which all modes (generally only two or three contribute

⁹The work conducted at NIST generally involved large detunings from all the modes with the express intention of driving them uniformly.

significantly) return to the initial motional state simultaneously. It is then essential to ensure that these compound pulses are properly accounted for in the overall evolution. For small crystals or modest gate speeds this is quite straightforward, but the complexity rapidly increases with the number of ions, and may make larger codes impractical. This remains a topic of ongoing research.

6

Conclusion

The work presented in this thesis has described the development of an experiment to allow ground-state cooling of single $^{40}\text{Ca}^+$ ions in a Penning trap. Earlier work, which formed the basis of two previous theses [56, 66], laid the foundation for these experiments, in the form of a working trap, imaging system, Doppler cooling, repump and spectroscopy lasers, and an experimental control system. However, to make this system suitable for sideband cooling and other more advanced work required almost every experimental subsystem to be upgraded or overhauled and several new ones to be introduced. I have introduced versatile locking systems for the cooling and repump lasers, suppressed drifts in the spectroscopy laser frequency and trap magnetic field, increased the intensity of the spectroscopy laser by over an order of magnitude, extended and rationalised the experimental control and analysis system and consolidated several disparate systems into one low-noise trap supply.

All of these have helped systematise an experimental process that can easily become chaotic, by increasing the reliability, repeatability and speed of experimental runs and by providing the accurate temperature measurements necessary to judge the efficacy of changes made. With these new systems in place, we have been able to develop techniques to reduce the ion temperature and increase the trap frequency. The most important of these techniques was the proper AC and DC compensation of the trapping potential, allowing intensity gradient cooling to be combined with an axialisation potential in order to cool both radial modes at all trap frequencies up to the fundamental stability limit. With the radial cooling less critically reliant on the geometry and saturation of the cooling beam, the cooling laser powers could also be reduced well below saturation, allowing us to cool to the Doppler limit and paving the way for cooling in the strong-binding regime.

The primary aim of this work was to demonstrate resolved-sideband cooling and in that respect it has been a resounding success. We have used sideband cooling to place a single $^{40}\text{Ca}^+$ ion in the axial motional ground state with 99% probability, a technique which we now employ on a regular basis for more advanced experiments. We have made a series of basic demonstrations of the coherent control that is possible without thermally-induced dephasing. We have also measured the heating rate of SPECTRAP and found it to imply an electric field spectral noise density of $\omega_z S_E(\omega_z) = 1.47 \pm 0.06 \times 10^{-8} \text{ V}^2 \text{ m}^{-2} \text{ Hz}^{-1} \text{ s}^{-1}$, the second-lowest reported to date. This is a result of some interest given the unusually large electrode dimensions of the trap, although it is actually slightly higher than would be expected for a trap of this size if applying a $1/d^3$ or $1/d^4$ scaling.

In the coming months there are a number of further experiments that we plan to conduct, with some already underway. The source of the heating rate will be investigated with the aim of reducing it further. Initial indicators point towards resonant heating due to the radial motion, which could be produced by the non-orthogonal modes that occur when the trap axis is misaligned with the magnetic field. Two other important demonstrations will be to extend the sideband cooling technique to the two radial modes of a single ion, as well as to the axial modes of a two-ion crystal. As has been discussed at length in this thesis, radial cooling is one of the most troublesome aspects of work in a Penning trap and simultaneous sideband cooling of the magnetron and modified cyclotron motions avoids many of the compromises involved in other techniques. Ongoing experiments by G. Stutter have produced exciting preliminary results, though sideband cooling has yet to be combined with the use of an axialisation potential. If this could be achieved, it should be possible to cool to the three-dimensional motional ground state. As a corollary, this would help confirm (or otherwise) the suspected source of the measured heating rate.

While much progress can be made on these experiments with the current apparatus, as the complexity and delicacy of the operations increases, so do the demands upon the performance of the experimental systems. In particular the existing Doppler cooling laser lock system, despite its improvements, remains the most unreliable aspect of the experiment, mainly due to intermittent creep in the PDH cavity piezos and environmental drifts in the scanning cavity. Replacing the PDH lock cavities and evacuating the 397-nm scanning cavity will greatly increase the amount of useful experimental time per week.

Because we work at such low trap frequencies compared to many groups, ground-state cooling of the radial motion and multi-ion crystals will certainly require sideband cooling on more than one red sideband to avoid population trapping. While this is quite straightforward in principle, the complexity of the RF system will soon reach the limit of the existing switchable-analogue-source approach. To this end, a more versatile radiofrequency source will be necessary and work is already underway to develop a direct digital synthesis system which can provide as many phase-controlled RF frequencies as necessary.

Another target for improvement is the 729-nm laser. The most pressing need is for

closed-loop intensity stabilisation in the form of an AOM noise-eater, which should greatly reduce the dephasing and sharp discontinuities observed in Rabi oscillations and will pave the way for high-fidelity quantum control. As has been discussed in 4.3.5, the observed Rabi frequency is also far lower than would be expected for the intensity of the laser, suggesting that the spectral purity of the light is not as high as was inferred from the sideband linewidth and coherence time. Characterising and improving this is another aim for the immediate future.

The final chapter of this thesis described a very different strand of theoretical investigation, exploring the possibilities of producing an error-correctable quantum memory on the qubit register provided by a two-dimensional Coulomb crystal. The outcomes of this work were very positive and showed that by exploiting the symmetries of such crystals, it is possible to efficiently produce certain complex entangled states using global entangling operations alone. While a variety of proposals exist for using global gates in RF traps, these examples have always called for additional local unitaries to produce any entanglement more complicated than the GHZ state. By avoiding these local operations, our proposed pulse sequences offer the potential for repetitive quantum error correction with a greatly reduced number of discrete gates. Continued research by an MSc student under my supervision, Y. Seis, has helped provide analytic insights into our numerical results and suggests that more complex and efficient schemes may yet be possible [115].

A major new direction of work for the group will focus on implementing these and similar ideas using optical dipole forces in two-dimensional qubit arrays, and plans are currently underway to construct a new experiment. These techniques will benefit from an increased trap frequency and are thus likely to necessitate a move to a lighter ion such as beryllium or magnesium, both of which are also unaffected by the j -mixing that makes high-fidelity readout of calcium ions difficult in strong fields. However, this new experiment is likely to take several years to reach maturity and calcium will provide the ideal test-bed for a number of the methods we later hope to employ. In particular, phase-locked control and imaging of crystals via a rotating wall underpins our future work but will be most easily tested in a system that is otherwise operational and well understood. This will require the number of ring segments in the existing trap to be increased, a task we hope to perform in mid-2015. The disassembly of the trap will also provide the opportunity to remove sources of scatter and improve the beam steering systems.

This thesis has described the first demonstration of ground-state cooling in a Penning trap. This important result provides the foundation for a range of interesting future work in quantum information and metrology. Further theoretical work has also explored the possibility of engineering quantum error-correcting codes on two-dimensional qubit arrays in Penning traps, which may provide an architecture for the efficient realisation of these and other complex entangled states.

Appendices

A

Transition Matrix Elements

The Hamiltonian of an electron in a electromagnetic field (with longitudinal and transverse electric field $\mathbf{E}_{\parallel} = -\nabla\phi$ and $\mathbf{E}_{\perp} = -\frac{\partial\mathbf{A}}{\partial t}$ and magnetic field $\mathbf{B} = \nabla \times \mathbf{A}$) is

$$H = \frac{(\mathbf{p} + e\mathbf{A})^2}{2m} - e\phi + \frac{e}{m}\mathbf{S} \cdot \mathbf{B}, \quad (\text{A.1})$$

where \mathbf{S} is the spin operator. Expanding the kinetic momentum term, and considering only the interaction term with linear dependence on the transverse field (the non-linear term can be ignored at low intensities), we find

$$H_i = \frac{e}{2m}(\mathbf{p} \cdot \mathbf{A} + \mathbf{A} \cdot \mathbf{p}) + \frac{e}{m}\mathbf{S} \cdot \mathbf{B}. \quad (\text{A.2})$$

Choosing the Coulomb gauge $\nabla \cdot \mathbf{A} = 0$ and using the fact that $\mathbf{p} = -i\hbar\nabla$, this becomes

$$H_i = \frac{e}{m}\mathbf{A} \cdot \mathbf{p} + \frac{e}{m}\mathbf{S} \cdot \mathbf{B}. \quad (\text{A.3})$$

For the following steps I drop the spin interaction term for clarity, but will re-introduce this later when considering magnetic dipole coupling.

We now consider the case of a laser incident upon our electron and model the transverse field as a plane wave of the form

$$\mathbf{A} = \frac{\hat{\epsilon}_L E_0}{2\omega_L i} \left(e^{i(\mathbf{k}_L \cdot \mathbf{r} - \omega_L t)} + e^{-i(\mathbf{k}_L \cdot \mathbf{r} - \omega_L t)} \right). \quad (\text{A.4})$$

To find the Rabi frequency used in Section 2.2.1, we need to calculate the matrix element

$$(H_i)_{eg} = \langle e | \mathbf{H}_i | g \rangle = \int \psi_e^* \hat{H}_i \psi_g \, d\mathbf{r}, \quad (\text{A.5})$$

which in the case of the plane wave is given by

$$(H_i)_{eg} = \frac{eE_0}{2m\omega_L i} \left[\langle e | \left(\hat{\epsilon}_L \cdot \mathbf{p} e^{i(\mathbf{k}_L \cdot \mathbf{r} - \omega_L t)} + \text{c.c.} \right) | g \rangle \right]. \quad (\text{A.6})$$

We now make the important step of Taylor expanding the spatial part of the exponent as $e^{i(\mathbf{k} \cdot \mathbf{r} - \omega_L t)} = 1 + i\mathbf{k} \cdot \mathbf{r} + \dots$ to give

$$(H_i)_{eg} = \frac{eE_0}{2m\omega_L i} \left[\left(\langle e | \hat{\epsilon}_L \cdot \mathbf{p} | g \rangle + i \langle g | (\hat{\epsilon}_L \cdot \mathbf{p})(\mathbf{k}_L \cdot \mathbf{r}) | e \rangle \right) e^{-i\omega_L t} + \text{h.o.} + \text{c.c.} \right]. \quad (\text{A.7})$$

Electric dipole - E1

Providing the wavelength of the light is much larger than the size of the atom or ion, we can make the dipole approximation, treating the electromagnetic field amplitude as constant over its extent and terminating the expansion at the first term. Considering this electric dipole (E1) term alone, and using the fact that the position operator is time-independent in the Schrödinger picture to write $\mathbf{p} = \frac{m}{i\hbar}[\mathbf{r}, \mathbf{H}_0]$, we find

$$\begin{aligned} (H_i)_{eg}^{E1} &= -\frac{eE_0}{2\hbar\omega_L} \left[\hat{\epsilon}_L \cdot \langle e | [\mathbf{r}, \mathbf{H}_0] | g \rangle e^{-i\omega_L t} + \text{c.c.} \right] \\ &= -\frac{eE_0}{2\hbar\omega_L} \left[\hat{\epsilon}_L \cdot \left(\langle e | \mathbf{r} \mathbf{H}_0 | g \rangle - \langle e | \mathbf{H}_0 \mathbf{r} | g \rangle \right) e^{-i\omega_L t} + \text{c.c.} \right] \\ &= -\frac{eE_0}{2\hbar\omega_L} \left[\hat{\epsilon}_L \cdot \left(\frac{-\hbar\omega_0}{2} \langle e | \mathbf{r} | g \rangle - \frac{\hbar\omega_0}{2} \langle e | \mathbf{r} | g \rangle \right) e^{-i\omega_L t} + \text{c.c.} \right] \\ &= \frac{eE_0\omega_0}{2\omega_L} \hat{\epsilon}_L \cdot \langle e | \mathbf{r} | g \rangle \left(e^{-i\omega_L t} + e^{i\omega_L t} \right) \\ &\approx E_0 \hat{\epsilon}_L \cdot \boldsymbol{\mu}_{eg} \cos(\omega_L t) \end{aligned} \quad (\text{A.8})$$

and thus

$$\Omega_0^{E1} = \boldsymbol{\mu}_{eg} \cdot \boldsymbol{\epsilon}_L E_0 / \hbar, \quad (\text{A.9})$$

as given in Section 2.2.1.

In this limit, the light-atom interaction Hamiltonian is dominated by the E1 term and the other contributions can be ignored. For optical processes this is usually a very good approximation, and the majority of transitions we will use occur via electric dipole coupling. However, angular momentum and parity selection rules mean that some transitions are forbidden and have vanishing electric dipole matrix elements. In this case, we must look for coupling via a higher order term of the expansion.

Returning to A.7, we split the $i\mathbf{k} \cdot \mathbf{r}$ term in the expansion into a symmetric and anti-symmetric part, using the fact that $\mathbf{UV} = [\mathbf{U}, \mathbf{V}] + \{\mathbf{U}, \mathbf{V}\}$:

$$(H_i)_{eg} = (H_i)_{eg}^{E1} + \frac{eE_0}{4m\omega_L} \left[(\hat{\mathbf{e}}_L \cdot \langle g | [\mathbf{p}, \mathbf{r}] | e \rangle \cdot \mathbf{k}_L + \hat{\mathbf{e}}_L \cdot \langle g | \{\mathbf{p}, \mathbf{r}\} | e \rangle \cdot \mathbf{k}_L) e^{-i\omega_L t} + \text{h.o.} + \text{c.c.} \right]. \quad (\text{A.10})$$

Magnetic dipole - M1

The commutator term, combined with the spin interaction term dropped at A.3, leads to the magnetic dipole (M1) coupling. No magnetic dipole transitions are used in our current experiment, but the coupling is given here for completeness.

Using the vector identity

$$(\mathbf{A} \times \mathbf{B}) \cdot (\mathbf{C} \times \mathbf{D}) = (\mathbf{A} \cdot \mathbf{C})(\mathbf{B} \cdot \mathbf{D}) - (\mathbf{A} \cdot \mathbf{D})(\mathbf{B} \cdot \mathbf{C}), \quad (\text{A.11})$$

we write

$$\begin{aligned} (H_i)_{eg}^{M1} &= -\frac{eE_0}{2m\omega_L} \left[(\hat{\mathbf{e}}_L \times \mathbf{k}_L) \cdot \langle g | (\mathbf{r} \times \mathbf{p}) | e \rangle \cos(\omega_L t) \right] + \frac{e}{m} \langle g | \mathbf{S} | e \rangle \cdot \mathbf{B} \cos(\omega_L t) \\ &= \frac{eE_0}{2mc} \left[(\hat{\mathbf{k}}_L \times \hat{\mathbf{e}}_L) \cdot \langle g | \mathbf{L} | e \rangle \cos(\omega_L t) \right] + \frac{e}{m} \langle g | \mathbf{S} | e \rangle \cdot \mathbf{B} \cos(\omega_L t), \end{aligned} \quad (\text{A.12})$$

where I have used $\mathbf{k}_L = \hat{\mathbf{k}}_L \omega_L / c$. Noting that $(\mathbf{k}_L \times \hat{\mathbf{e}}_L)$ points in the direction of the laser \mathbf{B} -field, with amplitude $B_0 = E_0/c$ and time dependence $\cos(\omega_L t)$, the first term in the above expression can be seen as describing the interaction of the magnetic field with the electron angular momentum, under the assumption that the \mathbf{B} -field does not vary significantly over the atom (hence M1). The magnetic dipole interaction Hamiltonian is therefore

$$H_i^{M1} = \frac{e}{2m} (\mathbf{L} + 2\mathbf{S}) \cdot \mathbf{B} \cos(\omega_L t), \quad (\text{A.13})$$

and thus

$$\Omega_0^{M1} = \frac{eB_0}{2m\hbar} (\mathbf{L} + 2\mathbf{S}) \cdot (\hat{\mathbf{k}}_L \times \hat{\mathbf{e}}_L). \quad (\text{A.14})$$

Electric quadrupole - E2

The anti-commutator term in A.10 leads to the electric quadrupole (E2) interaction. Following a similar procedure to that used in A.8,

$$\begin{aligned}
 (H_i)_{eg}^{E2} &= \frac{eE_0}{4m\omega_L} \left[(\hat{\epsilon}_L \cdot \langle g | \{ \mathbf{p}, \mathbf{r} \} | e \rangle \cdot \mathbf{k}_L) (e^{i\omega_L t} - e^{-i\omega_L t}) \right] \\
 &= \frac{eE_0}{4i\hbar c} \left[\hat{\epsilon}_L \cdot (\langle g | [\mathbf{r}, \mathbf{H}_0] \mathbf{r} | e \rangle + \langle g | \mathbf{r} [\mathbf{r}, \mathbf{H}_0] | e \rangle) \cdot \hat{\mathbf{k}}_L (e^{i\omega_L t} - e^{-i\omega_L t}) \right] \\
 &= \frac{eE_0\omega_0}{4ic} \left[\hat{\epsilon}_L \cdot \langle g | \mathbf{r} \mathbf{r} | e \rangle \cdot \hat{\mathbf{k}}_L (e^{i\omega_L t} - e^{-i\omega_L t}) \right] \\
 &= \frac{eE_0\omega_0}{2c} \hat{\epsilon}_L \cdot \langle g | \mathbf{r} \mathbf{r} | e \rangle \cdot \hat{\mathbf{k}}_L \sin(\omega_L t)
 \end{aligned} \tag{A.15}$$

and thus

$$\Omega_0^{E2} = \frac{eE_0\omega_0}{2\hbar c} \langle g | \mathbf{r}_i \mathbf{r}_j | e \rangle (\hat{\mathbf{k}}_L)_j (\hat{\epsilon}_L)_i, \tag{A.16}$$

where the summation over i and j is implied.

Evaluation of matrix elements

For anything but the simplest atoms calculation of the wavefunctions is difficult, making explicit evaluation of E1 and E2 transition matrix elements unfeasible in everyday work. Fortunately, the Wigner-Eckart theorem allows us to calculate the matrix element for any given transition, given knowledge of another transition between states with the same initial and final total angular momentum quantum numbers. Providing we can calculate or measure one matrix element $\langle n, j, m_j | H_i | n', j', m'_j \rangle$, we can easily extrapolate the values for all other matrix elements with the same n, j, n', j' based on their Clebsch-Gordan coefficients alone. Most usefully of all this means if we measure the decay rate from one manifold of states to another we can calculate the matrix elements between any given pair of sublevels.

A clear and concise description of this procedure from the starting point of Equations A.9 and A.16 has been given by D.F.V. James [116]; I summarise the results for the electric dipole and electric quadrupole transitions here.

In general, the Rabi frequency of E1 and E2 transitions with $\Delta m_j = m_j - m'_j$ are

$$\Omega_0^{E1} = \frac{e|E|}{\hbar} \sqrt{\frac{3(2j'+1)A_{12}}{4\alpha c k_{12}^3}} \left[\begin{pmatrix} j & 1 & j' \\ -m_j & \Delta m_j & m'_j \end{pmatrix} \mathbf{c}^{(\Delta \mathbf{m}_j)}_i (\hat{\epsilon}_L)_i \right] \tag{A.17}$$

$$\Omega_0^{E2} = \frac{e|E|}{\hbar} \sqrt{\frac{15(2j'+1)A_{12}}{4\alpha c k_{12}^3}} \left[\begin{pmatrix} j & 2 & j' \\ -m_j & \Delta m_j & m'_j \end{pmatrix} \mathbf{c}^{(\Delta \mathbf{m}_j)}_{ij} (\hat{\epsilon}_L)_i (\hat{\mathbf{k}}_L)_j \right], \tag{A.18}$$

where e is the electronic charge, A_{12} the Einstein A coefficient, k_{12} the transition wavevector, α the fine structure constant, c the speed of light and $j(j')$ the ground(excited) state angular momentum.

The final, geometric term gives contains the 3- j symbol and the Racah tensor $\mathbf{c}^{(\Delta m_j)}$. The 3- j symbol provides the Clebsch-Gordan coefficient and thus determines the line strength of a particular transition relative to the total manifold decay rate A_{12} . The Racah tensor, which is rank-1 in the case of the electric dipole transition and rank-2 for the electric quadrupole transition, defines the polarisation dependence.

In Cartesian coordinates, the rank-1 Racah tensors are

$$\mathbf{c}^{(1)} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \\ 0 \end{pmatrix} \quad \mathbf{c}^{(0)} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \quad \mathbf{c}^{(-1)} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \\ 0 \end{pmatrix}, \quad (\text{A.19})$$

while the rank-2 Racah tensors are

$$\mathbf{c}^{(2)}_{ij} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i & 0 \\ -i & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad \mathbf{c}^{(-2)}_{ij} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i & 0 \\ i & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (\text{A.20})$$

$$\mathbf{c}^{(0)}_{ij} = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 2 \end{pmatrix} \quad (\text{A.21})$$

$$\mathbf{c}^{(1)}_{ij} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 & -1 \\ 0 & 0 & i \\ -1 & i & 0 \end{pmatrix} \quad \mathbf{c}^{(-1)}_{ij} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & i \\ 1 & i & 0 \end{pmatrix}. \quad (\text{A.22})$$

The rank-1 tensors are identical to $\hat{\mathbf{e}}_L$ for light of linear ($\Delta m_j = 0$ [π]) and right and left circular ($\Delta m_j = \pm 1$ [σ_{\pm}]) polarisation and it is to these field polarisations that the atom will respond. In fact, in the case of a dipole transition, the *only* geometric consideration is the projection of the light polarisation - the angle at which the beam is incident makes no intrinsic difference, although does indirectly affect the maximum possible overlap of the polarisation vector with $\mathbf{c}^{(\Delta m_j)}$.¹

The rank-2 tensors are less familiar but simple to use. Unlike electric dipole transitions, electric quadrupole transitions depend on both the polarisation and wavevector of the laser, as the coupling is to the vector gradient of the field rather than the scalar magnitude.

The emission patterns for E1 and E2 transitions are given in many texts, including [56].

¹I.e. a Gaussian beam propagating in the z direction can have no component of polarisation in this direction and so has zero coupling to π transitions.

B

Excitation Lineshapes

When driving a narrow transition with a narrow laser, as we do when performing 729-nm electron-shelving spectroscopy, a range of excitation lineshapes may be observed. Such lineshapes are governed by straightforward atomic physics but are summarised here for a range of commonly used parameters for ease of reference.

Assuming the natural and laser linewidths are much smaller than the Rabi frequency of the driving laser, the lineshape is governed by the Rabi frequency and the probe pulse time alone. A range of lineshapes of a carrier excitation are shown in Figure B.1. I consider two cases: an ion in the ground state (blue) and an ion sampled from a thermal distribution with $\bar{n} = 25$ (red).

At low saturation, where $\Omega_0 t \lesssim 1$, both the ground state and thermal ions produce a lineshape that is almost exactly $P_0 \text{sinc}^2(\pi f t_{\text{pulse}})/(\pi f t_{\text{pulse}})$, where P_0 is a scaling parameter. This function, shown in black in Figure B.1 is the squared Fourier transform of a rectangular pulse – the lineshape can be understood as Fourier limited as per the classical theory.

As the level of saturation increases, the lineshape diverges from this classical limit, changing in different ways depending on whether Ω_0 or t_{pulse} is increased.

In the case of the ion in the ground state, the effect of increasing saturation is to produce a comb of sinc^2 -like oscillations with decreasing period, often referred to as a ‘Rabi sinc’ profile. The profile follows a Lorentzian envelope, with width $2\Omega_0$, so the effect of increasing the Rabi frequency is to broaden the spectral feature. Alternatively, increasing the pulse time reduces the oscillation period of the Rabi sinc, so more side-lobes appear within a fixed peak width.

For an ion at the Doppler limit, many Fock states contribute incoherently to the total

dynamics. Each of these Fock states has a different Rabi frequency, so there is partial destructive interference and the oscillations wash out to an average that approaches the Lorentzian envelope profile. At low saturation, the evolution of each Fock state is similar, but as the excitation becomes stronger the dephasing due to the varying Rabi frequencies becomes more significant.

In practice, for long pulse times, the side-lobe feature width is often smaller than the laser linewidth, and the features of an ion in the ground state will be averaged out in a similar way to those of the thermal ion. As a rule of thumb, at low saturation everything looks like a sinc^2 , at high saturation everything looks like a Lorentzian, while for moderate saturations the presence of coherent oscillations differentiates between thermal ions and those in a Fock state.

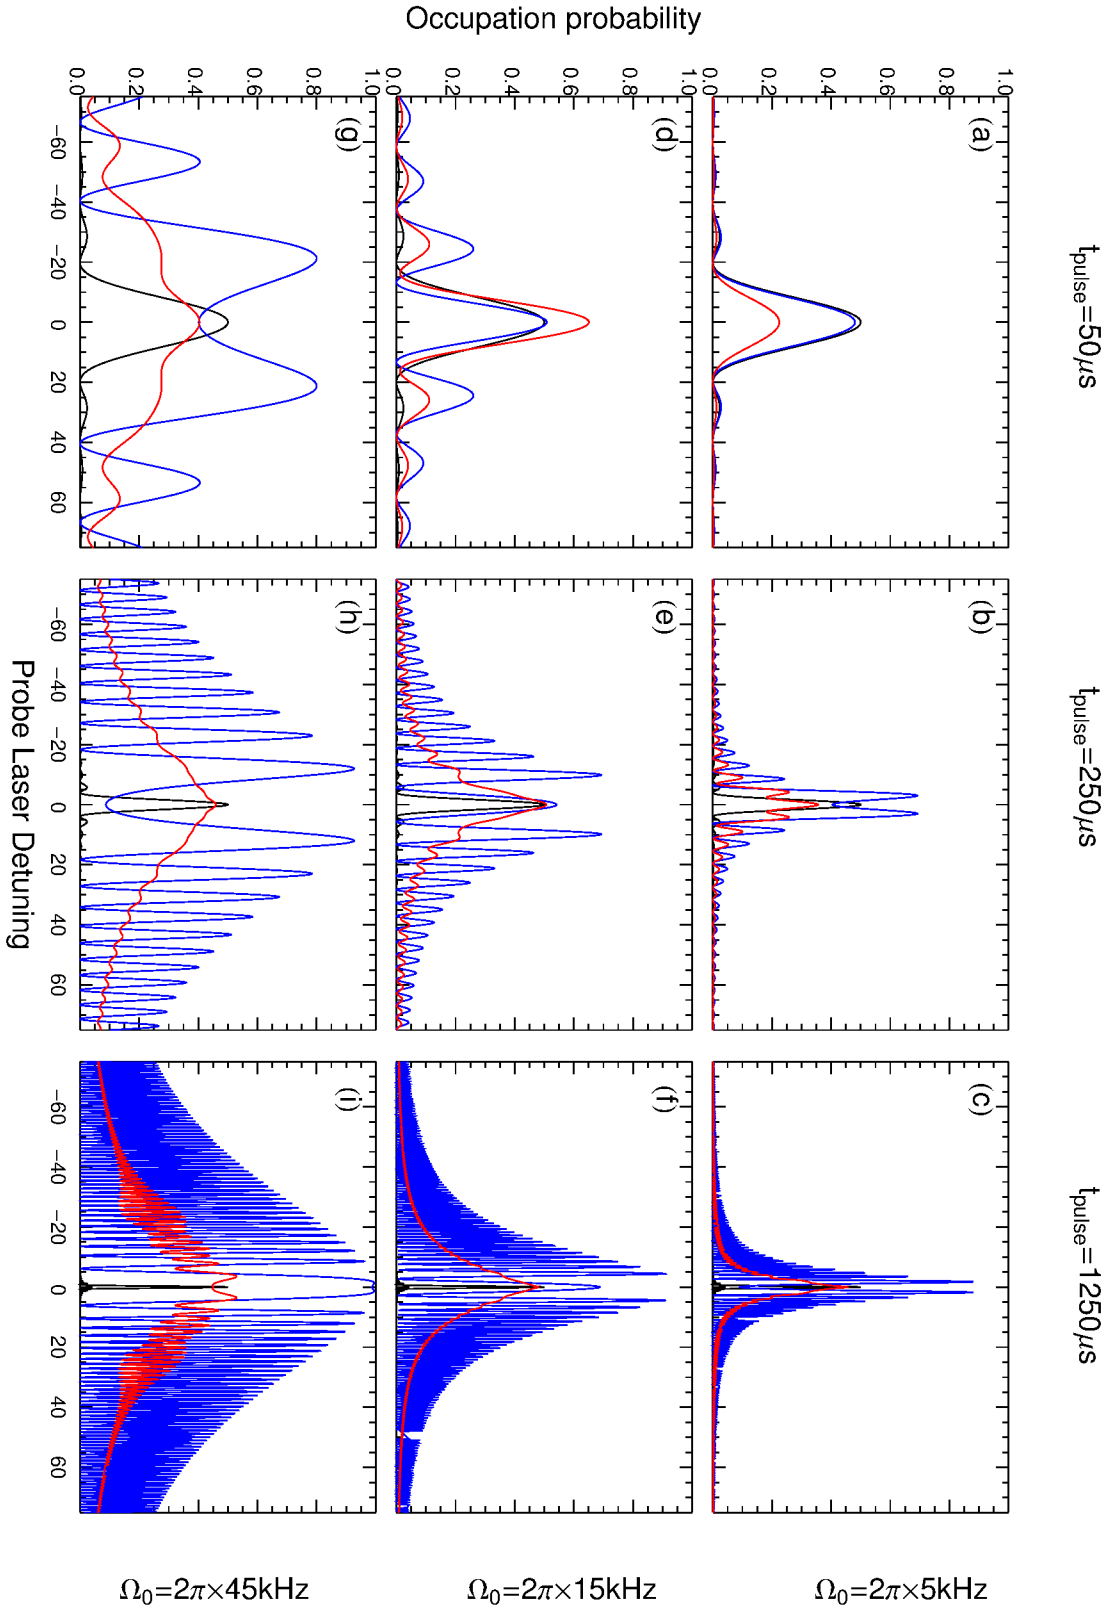


Figure B.1: Lineshapes for a range of pulse times and Rabi frequencies. The blue plots represent the carrier dynamics of an ion in the ground state. The red plots represent the dynamics for an ion in $\bar{n} = 25$, approximately the Doppler limit for the $\sim 400 \text{ kHz}$ trap frequencies used for most of this thesis. The black plots are $0.5 \times \text{sinc}^2$ function that is the squared Fourier transform of a rectangular pulse of length t_{pulse} .

C

Artefacts of Radial Motion

Taking the hypothesis that the artefacts in Figure 4.10 were caused by modulation due to the radial motion in an inhomogeneous field, the spectrum was fitted to the Rabi dynamics of an ion in the Lamb-Dicke regime, driven by a frequency modulated laser. Three frequencies of modulation would be expected, $\omega_a = \omega_c, \omega_{c'}$ and ω_m , due to the axialisation micromotion, modified cyclotron motion, and magnetron motion respectively. For FM side-band numbers z, a, c', m and modulation indices $h_a, h_{c'}, h_m$, Bessel functions $J_\alpha(x)$, and a constant background B due to j -mixing, the fitted model is

$$\begin{aligned}
 P(\uparrow) = & B + (1 - B) \sum_{z=-1}^1 \sum_{a=-1}^1 \sum_{c'=-2}^2 \sum_{m=-4}^4 \frac{[\Omega(z, a, c', m, \bar{n}_z, h_a, h_{c'}, h_m)]^2}{[\omega - \omega'_0(z, a, c, m)]^2 + [\Omega(z...h_m)]^2} \\
 & \times \sin^2 \left[\frac{\sqrt{[\omega - \omega'_0(z...m)]^2 + [\Omega_0(z...h_m)]^2}}{2} t \right], \\
 \Omega(z, a, c', m, \bar{n}_z, h_a, h_{c'}, h_m) = & \left[1 - |z| \left(1 - \eta \sqrt{\bar{n}_z + \frac{1+z}{2}} \right) \right] J_{|a|}(h_a) J_{|c'|}(h_{c'}) J_{|m|}(h_m) \Omega_0, \\
 \omega'_0(z, a, c, m) = & \omega_0 + z\omega_z + a\omega_a + c'\omega_{c'} + m\omega_m.
 \end{aligned} \tag{C.1}$$

The model provides best fit parameters for the phonon number \bar{n}_z and modulation indices h of

$$\bar{n}_z = 1.6 \quad h_a = 0.28 \quad h_{c'} = 0.49 \quad h_m = 0.85, \tag{C.2}$$

The modulation index is defined as

$$h = \frac{\Delta f}{f_m} = \frac{\text{Peak deviation}}{\text{Modulation frequency}} \quad (\text{C.3})$$

leading to peak deviations of

$$\Delta f_a = 275 \text{ kHz} \quad \Delta f_{c'} = 198 \text{ kHz} \quad \Delta f_m = 121 \text{ kHz}, \quad (\text{C.4})$$

If the modulation is due to magnetic field inhomogeneity, these must be equal to the difference in Zeeman shifts across the radii of the respective orbits. The radial field gradient is known from Section 3.4.2 to be approximately 0.025 T m^{-1} , equivalent to $0.5 \text{ kHz } \mu\text{m}^{-1}$ for the $S_{1/2, -1/2} \rightarrow D_{5/2, -3/2}$ transition. This means that the orbital radii would have to be

$$R_a = 400 \mu\text{m} \quad R_{c'} = 550 \mu\text{m} \quad R_m = 240 \mu\text{m}, \quad (\text{C.5})$$

corresponding to temperatures

$$T_a = 7700 \text{ K} \quad T_{c'} = 9000 \text{ K} \quad T_m = 120 \text{ K}. \quad (\text{C.6})$$

Suffice to say, the ion is clearly not this hot¹ and the field gradient cannot explain the presence of radial artefacts on the axial spectrum.

We instead consider whether the artefacts can be explained by component of the probe beam in the radial direction. A Doppler induced frequency shift of Δf corresponds to a velocity of $v = c\Delta f/f$ in the direction of the beam. If the beam is misaligned with the trap axis by an angle θ , an observed velocity of v implies a orbital velocity in the radial plane of $v_{max} = v/\sin\theta$. The modulation indices from the fit therefore correspond to orbital velocities of

$$v_a = \frac{0.14}{\sin\theta} \quad v_{c'} = \frac{0.2}{\sin\theta} \quad v_m = \frac{0.09}{\sin\theta}, \quad (\text{C.7})$$

equivalent to temperatures of

$$T_a = \frac{47 \mu\text{K}}{\sin^2\theta} \quad T_{c'} = \frac{100 \mu\text{K}}{\sin^2\theta} \quad T_m = \frac{20 \mu\text{K}}{\sin^2\theta}. \quad (\text{C.8})$$

If the beam is misaligned by 1° , these temperatures would have to be $\sim 50\text{--}300 \text{ mK}$. This is quite a reasonable result, given that the radial cooling was purposefully tampered with to accentuate the effect of the artefacts. Such temperatures would not increase the orbital size by more than a few micrometres, keeping it well within the extent of the cooling beams. The pixel size of the camera is equivalent (post-magnification) to the radii of orbits with temperatures of this order, so no significant blurring would be seen. Finally, the Doppler

¹We would observe very low fluorescence, if any, due to minimal overlap with the beams, and the orbital radii would be easily visible on any image we could record.

width for a 100 mK motion is approximately 10 MHz, so this sort of temperature would not be very obvious when manipulating the cooling lasers, especially as the axial temperature would still be low. This tells us two things. Firstly, that the source of the radial artefacts is very likely to be a misalignment of the probe beam. Secondly, that while this misalignment persists, such artefacts provide the most accurate readily available indicator of the radial temperature, without recording a radial spectrum directly.

D

Miscellaneous Electronic Circuits

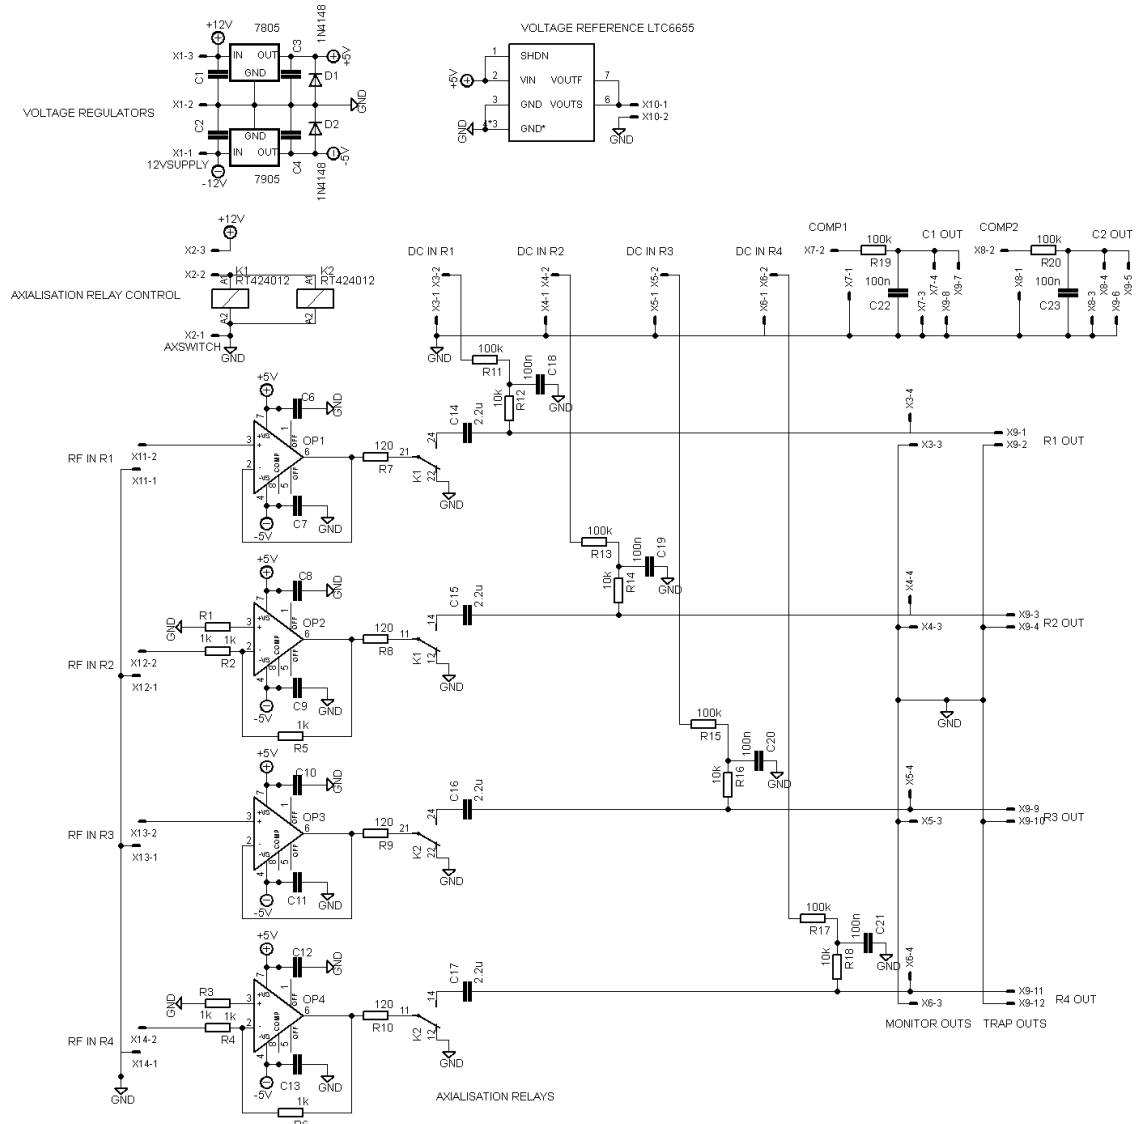


Figure D.1: Trap supply: AC and DC compensation circuit. Front panel potentiometers for AC and DC compensation feed RF IN : R1-R4 and DC IN : R1-R4 plus COMP1 and COMP2. Circuit described in Section 3.4.3.

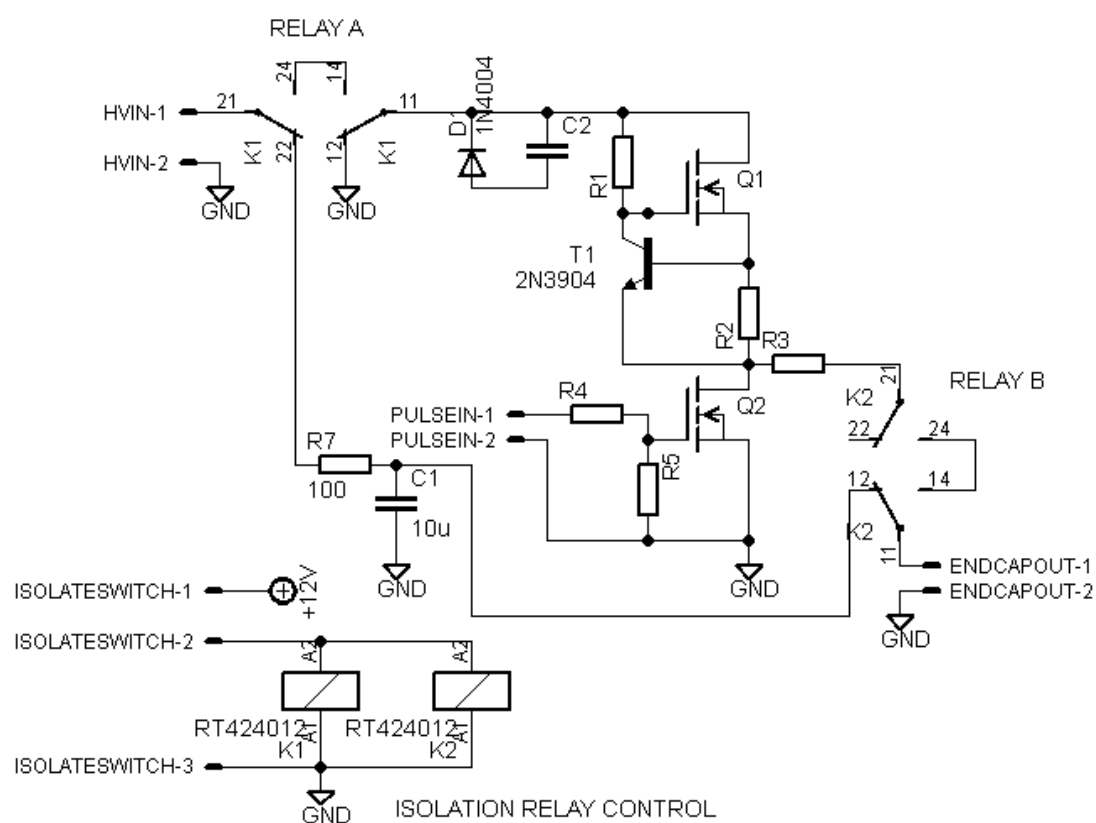


Figure D.2: Trap supply: Endcap supply and isolatable pulse-down circuit. Circuit described in Section 3.4.3.

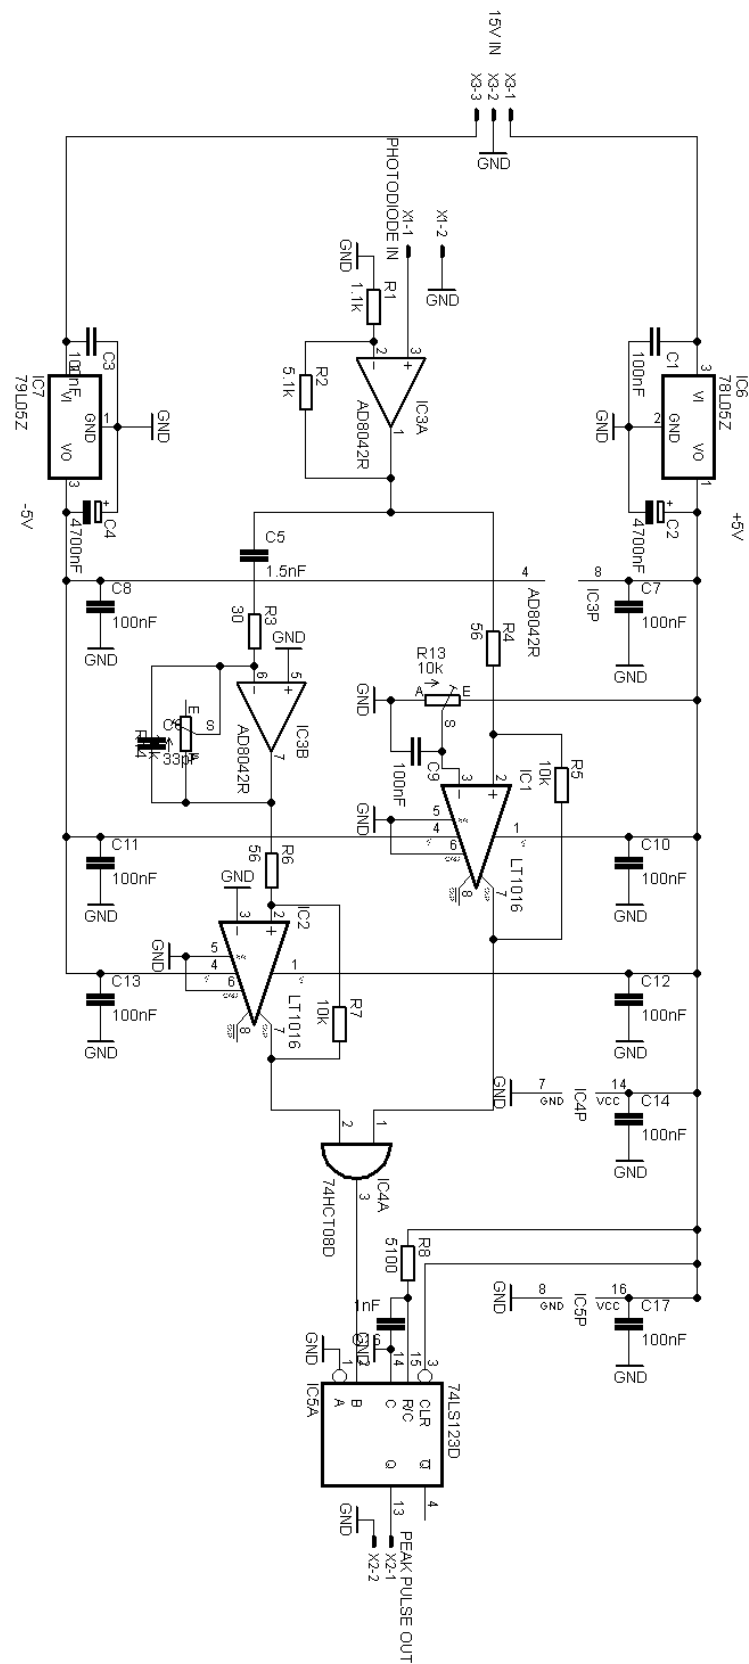


Figure D.3: Scanning cavity peak detector. Some passive component values must be changed to suit cavity scan frequency and finesse. Circuit described in Section 3.5.2.

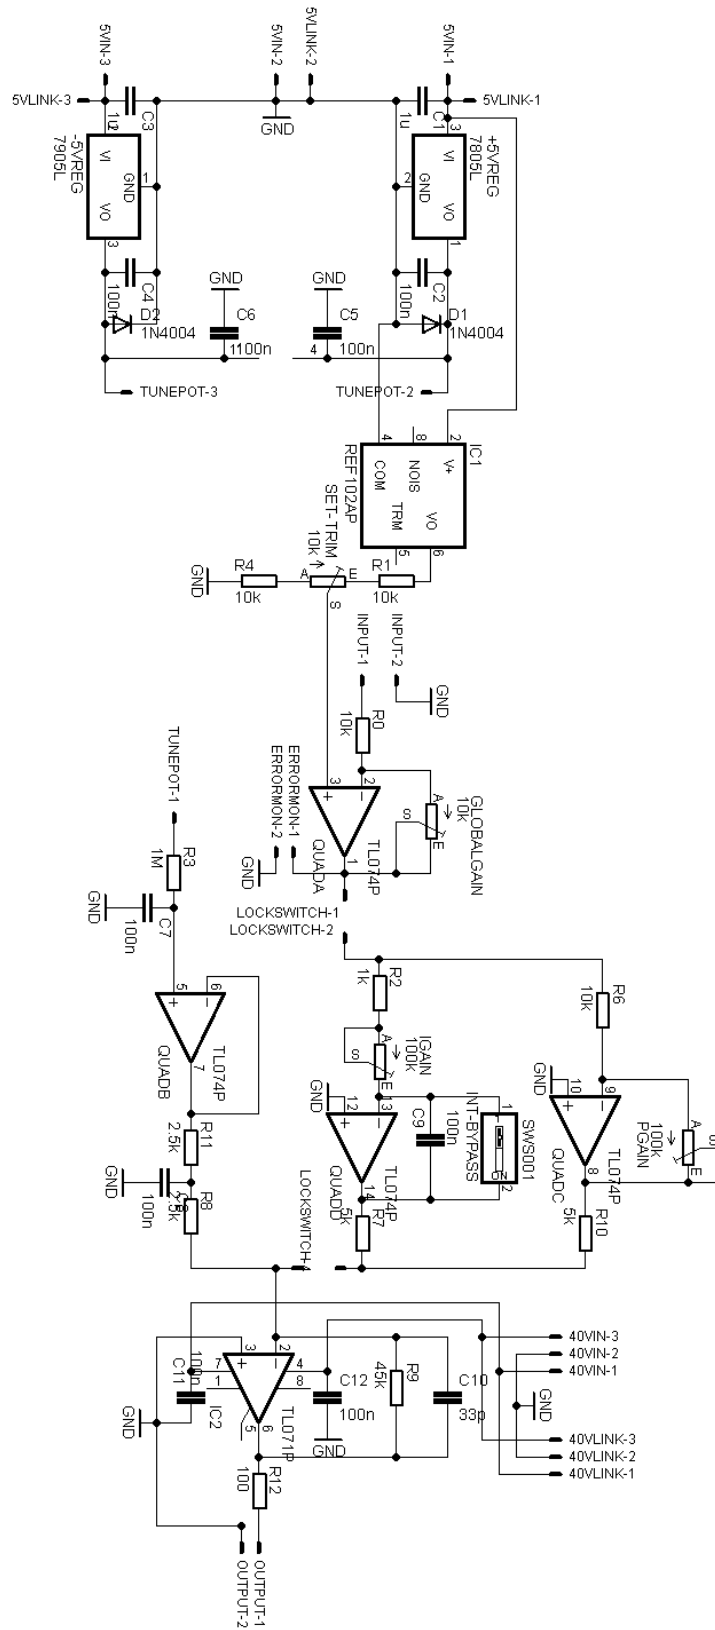


Figure D.4: 866-nm and 854-nm laser piezo PI lock circuit. Circuit described in Section 3.5.3.

E

CAD Renders




Figure E.1: Quarter-section and exterior views of assembled scanning cavity. The exterior view also shows the aluminium vacuum sealing rings which are not currently in use.




Figure E.2: Exterior view of complete ULE cavity cooler superstructure, without insulation.




Figure E.3: Cross-section of ULE cavity cooling system with insulation.




Figure E.4: Assembled tapered amplifier and optics on temperature controlled mount. The base acts as a heatsink for the TEC cooling the upper plate. The top, front and back covers prevent any physical contact with the fragile TA chip.




Figure E.5: Assembled tapered amplifier mount with acrylic dust cover. Note the two holes in the cover to provide optical access for the input and output beams - these should ideally be fitted with AR-coated windows.

Bibliography

- [1] D. J. Wineland, R. E. Drullinger, and F. L. Walls, “Radiation-pressure cooling of bound resonant absorbers,” *Phys. Rev. Lett.*, vol. 40, pp. 1639–1642, Jun 1978.
- [2] W. Neuhauser, M. Hohenstatt, P. Toschek, and H. Dehmelt, “Optical-sideband cooling of visible atom cloud confined in parabolic well,” *Phys. Rev. Lett.*, vol. 41, pp. 233–236, Jul 1978.
- [3] W. Neuhauser, M. Hohenstatt, P. E. Toschek, and H. Dehmelt, “Localized visible Ba^+ mono-ion oscillator,” *Phys. Rev. A*, vol. 22, pp. 1137–1140, Sep 1980.
- [4] F. Diedrich, J. C. Bergquist, W. M. Itano, and D. J. Wineland, “Laser cooling to the zero-point energy of motion,” *Phys. Rev. Lett.*, vol. 62, pp. 403–406, Jan 1989.
- [5] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould, “Resolved-sideband Raman cooling of a bound atom to the 3D zero-point energy,” *Phys. Rev. Lett.*, vol. 75, pp. 4011–4014, Nov 1995.
- [6] R. Feynman, “Simulating physics with computers,” *International Journal of Theoretical Physics*, vol. 21, no. 6-7, pp. 467–488, 1982.
- [7] D. Deutsch, “Quantum theory, the Church-Turing principle and the universal quantum computer,” *Proceedings of the Royal Society of London, A: Mathematical and Physical Sciences*, vol. 400, no. 1818, pp. 97–117, 1985.
- [8] P. Shor, “Algorithms for quantum computation: discrete logarithms and factoring,” in *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pp. 124–134, Nov 1994.
- [9] J. I. Cirac and P. Zoller, “Quantum computations with cold trapped ions,” *Phys. Rev. Lett.*, vol. 74, pp. 4091–4094, May 1995.
- [10] X. Wang, A. Sørensen, and K. Mølmer, “Multibit gates for quantum computing,” *Phys. Rev. Lett.*, vol. 86, p. 3907, 2001.

- [11] G. Milburn, S. Schneider, and D. James, “Ion trap quantum computing with warm ions,” *Fortschritte der Physik*, vol. 48, p. 801, 2000.
- [12] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, “Realization of the Cirac–Zoller controlled-NOT quantum gate,” *Nature*, vol. 422, no. 6930, pp. 408–411, 2003.
- [13] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. Itano, B. Jelenkovic, C. Langer, T. Rosenband, and D. Wineland, “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate,” vol. 422, no. 6930, p. 415, 2003.
- [14] 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.*, vol. 106, p. 130506, Mar 2011.
- [15] L. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland, “Realization of quantum error correction,” *Nature*, vol. 432, no. 7017, pp. 602–605, 2004.
- [16] D. Nigg, M. Müller, E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt, “Quantum computations on a topologically encoded qubit,” *Science*, vol. 345, no. 6194, pp. 302–305, 2014.
- [17] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, “High-fidelity preparation, gates, memory and readout of a trapped-ion quantum bit,” *arXiv:1403.1524*, 2014.
- [18] C. J. Ballance, T. P. Harty, N. M. Linke, and D. Lucas, “High-fidelity two-qubit quantum logic gates using trapped calcium-43 ions,” *arXiv:1406.5473*, 2014.
- [19] M. D. Hughes, B. Lekitsch, J. A. Broersma, and W. K. Hensinger, “Microfabricated ion traps,” *Contemporary Physics*, vol. 52, no. 6, pp. 505–529, 2011.
- [20] B. P. Lanyon, P. Jurcevic, M. Zwerger, C. Hempel, E. A. Martinez, W. Dür, H. J. Briegel, R. Blatt, and C. F. Roos, “Measurement-based quantum computation with trapped ions,” *Phys. Rev. Lett.*, vol. 111, p. 210501, Nov 2013.
- [21] J. W. Britton, B. C. Sawyer, A. C. Keith, C. C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, “Engineered two-dimensional Ising interactions in a trapped-ion quantum simulator with hundreds of spins,” *Nature*, vol. 484, p. 489, 2012.
- [22] R. Islam, C. Senko, W. Campbell, S. Korenblit, J. Smith, A. Lee, E. Edwards, C.-C. Wang, J. Freericks, and C. Monroe, “Emergence and frustration of magnetism with

- variable-range interactions in a quantum simulator,” *Science*, vol. 340, no. 6132, pp. 583–587, 2013.
- [23] P. Gill, G. P. Barwood, H. A. Klein, G. Huang, S. A. Webster, P. J. Blythe, K. Hosaka, S. N. Lea, and H. S. Margolis, “Trapped ion optical frequency standards,” *Measurement Science and Technology*, vol. 14, no. 8, p. 1174, 2003.
- [24] C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband, “Frequency comparison of two high-accuracy Al^+ optical clocks,” *Phys. Rev. Lett.*, vol. 104, p. 070802, Feb 2010.
- [25] C. W. Chou, D. B. Hume, T. Rosenband, and D. J. Wineland, “Optical clocks and relativity,” *Science*, vol. 329, no. 5999, pp. 1630–1633, 2010.
- [26] S. E. Hamann, D. L. Haycock, G. Klose, P. H. Pax, I. H. Deutsch, and P. S. Jessen, “Resolved-sideband Raman cooling to the ground state of an optical lattice,” *Phys. Rev. Lett.*, vol. 80, pp. 4149–4152, May 1998.
- [27] J. D. Thompson, T. G. Tiecke, A. S. Zibrov, V. Vuletić, and M. D. Lukin, “Coherence and Raman sideband cooling of a single atom in an optical tweezer,” *Phys. Rev. Lett.*, vol. 110, p. 133001, Mar 2013.
- [28] J. Teufel, T. Donner, D. Li, J. Harlow, M. Allman, K. Cicak, A. Sirois, J. Whittaker, K. Lehnert, and R. Simmonds, “Sideband cooling of micromechanical motion to the quantum ground state,” *Nature*, vol. 475, no. 7356, pp. 359–363, 2011.
- [29] D. Porras and J. I. Cirac, “Effective quantum spin systems with trapped ions,” *Phys. Rev. Lett.*, vol. 92, p. 207901, May 2004.
- [30] M. Redshaw, W. Shi, J. McDaniel, E. Wingfield, and E. Myers, “Precision mass spectrometry and polarizability shifts with one and two ions in a Penning trap,” in *TCP 2006* (J. Dilling, M. Comyn, J. Thompson, and G. Gwinner, eds.), pp. 421–425, Springer Berlin Heidelberg, 2007.
- [31] K. Blaum, H. Kracke, S. Kreim, A. Mooser, C. Mrozik, W. Quint, C. C. Rodegheri, B. Schabinger, S. Sturm, S. Ulmer, A. Wagner, J. Walz, and G. Werth, “ g -factor experiments on simple systems in Penning traps,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 42, no. 15, p. 154021, 2009.
- [32] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland, “Minimization of ion micromotion in a Paul trap,” *Journal of Applied Physics*, vol. 83, no. 10, pp. 5025–5033, 1998.

- [33] M. Niemann, A.-G. Paschke, T. Dubielzig, M. Carsjens, M. Kohnen, H. Hahn, S. Grondkowski, and C. Ospelkaus, “Quantum logic enabled test of discrete symmetries,” in *IonTech 2*, 2013.
- [34] S. Mavadia, J. F. Goodwin, G. Stutter, S. Bharadia, D. R. Crick, D. M. Segal, and R. C. Thompson, “Control of the conformations of ion Coulomb crystals in a Penning trap,” *Nat. Comms.*, vol. 4, no. 2571, p. 2571, 2013.
- [35] S. Mavadia, G. Stutter, J. F. Goodwin, D. R. Crick, R. C. Thompson, and D. M. Segal, “Optical sideband spectroscopy of a single ion in a Penning trap,” *Phys. Rev. A*, vol. 89, p. 032502, Mar 2014.
- [36] J. F. Goodwin, G. Stutter, R. C. Thompson, and D. M. Segal, “Sideband cooling an ion to the quantum ground state in a Penning trap with very low heating rate,” *arXiv:1407.6121*, 2014.
- [37] J. F. Goodwin, B. J. Brown, G. Stutter, H. Dale, R. C. Thompson, and T. Rudolph, “Trapped ion quantum error correcting protocols using only global operations,” *arXiv:1407.1858v1*, 2014.
- [38] J. R. Castrejón-Pita, H. Ohadi, D. R. Crick, D. F. A. Winters, D. M. Segal, and R. C. Thompson, “Novel designs for Penning ion traps,” *Journal of Modern Optics*, vol. 54, no. 11, pp. 1581–1594, 2007.
- [39] G. Gabrielse, X. Fei, L. A. Orozco, R. L. Tjoelker, J. Haas, H. Kalinowsky, T. A. Trainor, and W. Kells, “Cooling and slowing of trapped antiprotons below 100 meV,” *Phys. Rev. Lett.*, vol. 63, pp. 1360–1363, Sep 1989.
- [40] S. T. Gulde, *Experimental Realization of Quantum Gates and the Deutsch-Jozsa Algorithm with Trapped $^{40}\text{Ca}^+$ Ions*. PhD thesis, Universität Innsbruck, 2003.
- [41] D. Stick, W. K. Hensinger, S. Olmschenk, M. J. Madsen, K. Schwab, and C. Monroe, “Ion trap in a semiconductor chip,” *Nature Physics*, vol. 2, no. 1, p. 39, 2006.
- [42] D. R. Crick, S. Donnellan, S. Ananthamurthy, R. C. Thompson, and D. M. Segal, “Fast shuttling of ions in a scalable Penning trap array,” *Review of Scientific Instruments*, vol. 81, no. 1, pp. –, 2010.
- [43] P. K. Ghosh, *Ion Traps*. Oxford Science Publications, 1995.
- [44] F. G. Major, *Charged Particle Traps*. Springer, 2005.
- [45] R. J. Hendricks, E. S. Phillips, D. M. Segal, and R. C. Thompson, “Laser cooling in the Penning trap: an analytical model for cooling rates in the presence of an axializing field,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 41, no. 3, p. 035301, 2008.

- [46] L. S. Brown and G. Gabrielse, “Geonium theory: Physics of a single electron or ion in a Penning trap,” *Rev. Mod. Phys.*, vol. 58, pp. 233–311, Jan 1986.
- [47] P. B. Schwinberg, J. R. S. Van Dyck, and H. G. Dehmelt, “New comparison of the positron and electron g factors,” *Phys. Rev. Lett.*, vol. 47, pp. 1679–1682, 1981.
- [48] M. Weissbluth, *Photon-Atom Interactions*. Academic Press, 1989.
- [49] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, “Quantum dynamics of single trapped ions,” *Rev. Mod. Phys.*, vol. 75, pp. 281–324, Mar 2003.
- [50] D. J. Wineland and W. M. Itano, “Laser cooling of atoms,” *Phys. Rev. A*, vol. 20, pp. 1521–1540, Oct 1979.
- [51] H. G. Dehmelt and F. L. Walls, “‘Bolometric’ technique for the rf spectroscopy of stored ions,” *Phys. Rev. Lett.*, vol. 21, pp. 127–131, Jul 1968.
- [52] W. M. Itano, J. C. Bergquist, J. J. Bollinger, and D. J. Wineland, “Cooling methods in ion traps,” *Physica Scripta*, vol. 1995, no. T59, p. 106, 1995.
- [53] G. Morigi, J. Eschner, and C. H. Keitel, “Ground state laser cooling using electromagnetically induced transparency,” *Phys. Rev. Lett.*, vol. 85, pp. 4458–4461, Nov 2000.
- [54] S. Stenholm, “The semiclassical theory of laser cooling,” *Rev. Mod. Phys.*, vol. 58, pp. 699–739, Jul 1986.
- [55] W. M. Itano and D. J. Wineland, “Laser cooling of ions stored in harmonic and Penning traps,” *Phys. Rev. A*, vol. 25, pp. 35–54, Jan 1982.
- [56] S. Mavadia, *Motional Sideband Spectra and Coulomb Crystals in a Penning Trap*. PhD thesis, Imperial College London, 2013.
- [57] G. Bollen, R. Moore, G. Savard, and H. Stolzenberg, “The accuracy of heavy-ion mass measurements using time of flight-ion cyclotron resonance in a Penning trap,” *J. Appl. Phys.*, vol. 68, no. 9, p. 4355, 1990.
- [58] R. J. Hendricks, *Spectroscopy and Dynamics of Laser-Cooled Ca^+ Ions in a Penning Trap*. PhD thesis, Imperial College London, 2006.
- [59] H. H. H. Nejad, “Laser cooling in the Penning trap,” Master’s thesis, Imperial College London, 2006.
- [60] I. Marzoli, J. I. Cirac, R. Blatt, and P. Zoller, “Laser cooling of trapped three-level ions: Designing two-level systems for sideband cooling,” *Phys. Rev. A*, vol. 49, pp. 2771–2779, Apr 1994.

- [61] F. Beck, “Theory and simulation of ground-state laser cooling in a Penning trap,” Master’s thesis, Imperial College London, 2014.
- [62] G. Morigi, J. I. Cirac, M. Lewenstein, and P. Zoller, “Ground-state laser cooling beyond the Lamb-Dicke limit,” *EPL (Europhysics Letters)*, vol. 39, no. 1, p. 13, 1997.
- [63] L. Santos and M. Lewenstein, “Dynamical cooling of trapped gases: One-atom problem,” *Phys. Rev. A*, vol. 59, pp. 613–619, Jan 1999.
- [64] G. Poulsen, Y. Miroshnychenko, and M. Drewsen, “Efficient ground-state cooling of an ion in a large room-temperature linear Paul trap with a sub-Hertz heating rate,” *Phys. Rev. A*, vol. 86, p. 051402, Nov 2012.
- [65] C. F. Roos, *Controlling the quantum state of trapped ions*. PhD thesis, Leopold-Franzens-Universität Innsbruck, 2000.
- [66] S. Donnellan, *Towards Sideband Cooling of a Single $^{40}\text{Ca}^+$ Ion in a Penning Trap*. PhD thesis, Imperial College London, 2011.
- [67] A. Kreuter, C. Becher, G. P. T. Lancaster, A. B. Mundt, C. Russo, H. Häffner, C. Roos, W. Hänsel, F. Schmidt-Kaler, R. Blatt, and M. S. Safronova, “Experimental and theoretical study of the 3d ^2D -level lifetimes of $^{40}\text{Ca}^+$,” *Phys. Rev. A*, vol. 71, p. 032504, Mar 2005.
- [68] L. Yang and D. A. Church, “Lifetime measurements on long-lived levels of confined highly charged ions,” *Phys. Rev. Lett.*, vol. 70, pp. 3860–3863, Jun 1993.
- [69] R. Gerritsma, G. Kirchmair, F. Zähringer, J. Benhelm, R. Blatt, and C. F. Roos, “Precision measurement of the branching fractions of the 4p $^2\text{P}_{3/2}$ decay of Ca II,” *The European Physical Journal D*, vol. 50, no. 1, pp. 13–19, 2008.
- [70] D. Crick, *Manipulation of Individual Laser Cooled Ca^+ Ions in Penning Traps*. PhD thesis, Imperial College London, 2009.
- [71] D. R. Crick, S. Donnellan, D. M. Segal, and R. C. Thompson, “Magnetically induced electron shelving in a trapped Ca^+ ion,” *Phys. Rev. A*, vol. 81, p. 052503, May 2010.
- [72] W. Quint, J. Dilling, S. Djekic, H. Häffner, N. Hermanspahn, H.-J. Kluge, G. Marx, R. Moore, D. Rodriguez, J. Schönfelder, G. Sikler, T. Valenzuela, J. Verdú, C. Weber, and G. Werth, “HITRAP: A facility for experiments with trapped highly charged ions,” *Hyperfine Interactions*, vol. 132, no. 1-4, pp. 453–457, 2001.
- [73] S. Bharadia, *Towards Laser Spectroscopy of Highly Charged Ions: Dynamics of $^{40}\text{Ca}^+$ Ions in a Penning Trap*. PhD thesis, Imperial College London, 2011.

- [74] E. D. Black, “An introduction to Pound-Drever-Hall laser frequency stabilization,” *American Journal of Physics*, vol. 69, no. 1, pp. 79–87, 2001.
- [75] A. S. Arnold, J. S. Wilson, and M. G. Boshier, “A simple extended-cavity diode laser,” *Review of Scientific Instruments*, vol. 69, no. 3, pp. 1236–1239, 1998.
- [76] S. Woodrow, “Microprocessor-based improvements to a laser control system for ion trapping,” Master’s thesis, Imperial College London, 2013.
- [77] J. Goodwin, “Development of a fast scanning cavity for laser frequency stabilisation and locking,” Master’s thesis, Imperial College of Science, Technology and Medicine, 2010.
- [78] R. Balhorn, H. Kunzmann, and F. Lebowsky, “Frequency stabilization of internal-mirror helium–neon lasers,” *Appl. Opt.*, vol. 11, pp. 742–744, Apr 1972.
- [79] R. Fox, C. Oates, and L. Hollberg, *Stabilizing Diode Laser To High Finesse Cavities*, vol. 40 of *Experimental Methods in the Physical Sciences*, ch. 1, pp. 1–46. Elsevier Science, 2003.
- [80] J. Sudbery, *Studies of Laser Cooled Calcium Ions in the Penning and Combined Traps*. PhD thesis, Imperial College London, 2003.
- [81] S. Hirata, T. Akatsuka, Y. Ohtake, and A. Morinaga, “Sub-Hertz-linewidth diode laser stabilized to an ultralow-drift high-finesse optical cavity,” *Applied Physics Express*, vol. 7, no. 2, p. 022705, 2014.
- [82] J. Huwer, *Experimental tools for quantum networking operations with single photons and single ions*. PhD thesis, Universität des Saarlandes and ICFO Barcelona, 2013.
- [83] J. Alnis, A. Matveev, N. Kolachevsky, T. Udem, and T. W. Hänsch, “Subhertz linewidth diode lasers by stabilization to vibrationally and thermally compensated ultralow-expansion glass Fabry-Pérot cavities,” *Phys. Rev. A*, vol. 77, p. 053809, May 2008.
- [84] T. Legero, T. Kessler, and U. Sterr, “Tuning the thermal expansion properties of optical reference cavities with fused silica mirrors,” *J. Opt. Soc. Am. B*, vol. 27, pp. 914–919, May 2010.
- [85] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. Martin, L. Chen, and J. Ye, “A sub-40-mHz-linewidth laser based on a silicon single-crystal optical cavity,” vol. 6, no. 10, p. 692, 2012.
- [86] R. Fox, K. Corwin, and L. Hollberg, “Stable optical cavities for wavelength references,” tech. rep., National Institute for Standards and Technology, 2004.

- [87] H. Kobayashi and H. Takahashi, “Structural and vibrational relaxations at spontaneous aging in glass,” *Thermochimica Acta*, vol. 532, no. 0, pp. 172 – 175, 2012. The 21st IUPAC International Conference on Chemical Thermodynamics (ICCT-2010).
- [88] J. Bergquist, W. Itano, and D. Wineland, “Laser stabilization to a single ion,” in *Frontiers in Laser Spectroscopy* (T. Hänsch and M. Inguscio, eds.), Proceedings of the International School of Physics Enrico Fermi, pp. 359–376, 1992.
- [89] A. Ludlow, *The Strontium Optical Lattice Clock: Optical Spectroscopy with Sub-Hertz Accuracy*. PhD thesis, University of Colorado, 2008.
- [90] R. A. Nyman, G. Varoquaux, B. Villier, D. Sacchet, F. Moron, Y. Le Coq, A. Aspect, and P. Bouyer, “Tapered-amplified antireflection-coated laser diodes for potassium and rubidium atomic-physics experiments,” *Review of Scientific Instruments*, vol. 77, no. 3, pp. –, 2006.
- [91] S. Patel, “FPGA controlled pulse spectroscopy of $^{40}\text{Ca}^+$ ions,” Master’s thesis, Imperial College London, 2012.
- [92] S. Zeeman, “FPGA-based laser control for spectroscopy measurements in ion traps,” Master’s thesis, Imperial College London, 2012.
- [93] C. Roos, T. Zeiger, H. Rohde, H. C. Nägerl, J. Eschner, D. Leibfried, F. Schmidt-Kaler, and R. Blatt, “Quantum state engineering on an optical transition and decoherence in a Paul trap,” *Phys. Rev. Lett.*, vol. 83, pp. 4713–4716, Dec 1999.
- [94] J. C. Bergquist, W. M. Itano, and D. J. Wineland, “Recoilless optical absorption and Doppler sidebands of a single trapped ion,” *Phys. Rev. A*, vol. 36, pp. 428–430, Jul 1987.
- [95] W. Nagourney, J. Sandberg, and H. Dehmelt, “Shelved optical electron amplifier: Observation of quantum jumps,” *Phys. Rev. Lett.*, vol. 56, pp. 2797–2799, Jun 1986.
- [96] D. Hite, Y. Colombe, A. Wilson, D. Allcock, D. Leibfried, D. Wineland, and D. Pappas, “Surface science for improved ion traps,” *MRS Bulletin*, vol. 38, pp. 826–833, 10 2013.
- [97] A. Safavi-Naini, P. Rabl, P. F. Weck, and H. R. Sadeghpour, “Microscopic model of electric-field-noise heating in ion traps,” *Phys. Rev. A*, vol. 84, p. 023412, Aug 2011.
- [98] B. C. Sawyer, J. W. Britton, and J. J. Bollinger, “Spin dephasing as a probe of mode temperature, motional state distributions, and heating rates in a two-dimensional ion crystal,” *Phys. Rev. A*, vol. 89, p. 033408, Mar 2014.

- [99] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe, “Scaling and suppression of anomalous heating in ion traps,” *Phys. Rev. Lett.*, vol. 97, p. 103007, Sep 2006.
- [100] S. A. Schulz, U. Poschinger, F. Ziesel, and F. Schmidt-Kaler, “Sideband cooling and coherent dynamics in a microchip multi-segmented ion trap,” *New Journal of Physics*, vol. 10, no. 4, p. 045007, 2008.
- [101] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt, “Experimental quantum-information processing with $^{43}\text{Ca}^+$ ions,” *Phys. Rev. A*, vol. 77, p. 062306, Jun 2008.
- [102] R. B. Blakestad, C. Ospelkaus, A. P. VanDevender, J. H. Wesenberg, M. J. Biercuk, D. Leibfried, and D. J. Wineland, “Near-ground-state transport of trapped-ion qubits through a multidimensional array,” *Phys. Rev. A*, vol. 84, p. 032314, Sep 2011.
- [103] R. G. DeVoe and C. Kurtsiefer, “Experimental study of anomalous heating and trap instabilities in a microscopic ^{137}Ba ion trap,” *Phys. Rev. A*, vol. 65, p. 063407, Jun 2002.
- [104] C. Tamm, D. Engelke, and V. Bühner, “Spectroscopy of the electric-quadrupole transition $2S_{1/2}(f=0) \rightarrow 2D_{3/2}(f=2)$ in trapped $^{171}\text{Yb}^+$,” *Phys. Rev. A*, vol. 61, p. 053405, Apr 2000.
- [105] C.-C. J. Wang, A. C. Keith, and J. K. Freericks, “Phonon-mediated quantum spin simulator employing a planar ionic crystal in a Penning trap,” *Phys. Rev. A*, vol. 87, p. 013422, 2013.
- [106] S. Gladchenko, D. Olaya, E. Dupont-Ferrier, B. Douçot, L. B. Ioffe, and M. E. Gershenson, “Superconducting nanocircuits for topologically protected qubits,” *Nat. Phys.*, vol. 5, p. 48, 2008.
- [107] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, “Realization of three-qubit quantum error correction with superconducting circuits,” *Nature*, vol. 482, p. 382, 2012.
- [108] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffry, 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,” *arXiv:1402.4848*, 2014.
- [109] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, 2000.

- [110] E. Knill, “Quantum computing with realistically noisy devices,” *Nature*, vol. 434, no. 7029, p. 44, 2005.
- [111] A. M. Steane, “Overhead and noise threshold of fault-tolerant quantum error correction,” *Phys. Rev. A*, vol. 68, p. 042322, Oct 2003.
- [112] D. Gottesman, *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology, 1997.
- [113] J. T. Barreiro, M. Müller, P. Schindler, D. Nigg, T. Monz, M. Chwalla, M. Hennrich, C. F. Roos, P. Zoller, and R. Blatt, “An open-system quantum simulator with trapped ions,” *Nature*, vol. 470, no. 7335, pp. 486–491, 2011.
- [114] P. J. Lee, K.-A. Brickman, L. Deslauriers, P. C. Haljan, L.-M. Duan, and C. Monroe, “Phase control of trapped ion quantum gates,” *Journal of Optics B: Quantum and Semiclassical Optics*, vol. 7, no. 10, p. S371, 2005.
- [115] Y. Seis, “Entangling qubits with global pulses,” Master’s thesis, Imperial College of Science, Technology and Medicine, September 2014.
- [116] D. James, “Quantum dynamics of cold trapped ions with application to quantum computation,” *Applied Physics B*, vol. 66, no. 2, pp. 181–190, 1998.