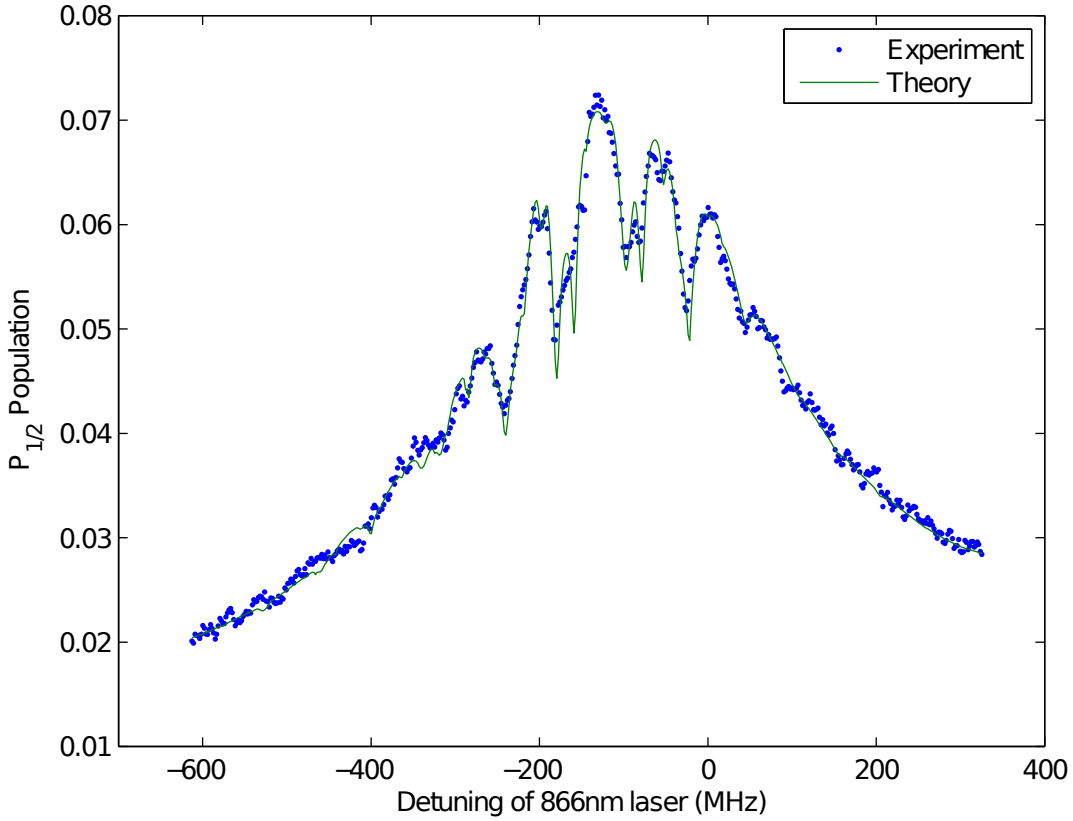

Optical Bloch equations for simulating trapped-ion qubits

Hugo Alexander Janacek

A thesis submitted for the degree
of Doctor of Philosophy

Christ Church
Oxford

Trinity Term 2015

Abstract

Optical Bloch equations for simulating trapped-ion qubits
A thesis submitted for the degree of Doctor of Philosophy
Trinity Term 2015
Hugo Alexander Janacek
Christ Church, Oxford

This thesis describes work on numerical modelling of the $^{43}Ca^{+}$ ion in a Paul trap using the optical Bloch equations. This is a challenging system to study, with many states involved in the internal dynamics. A major outcome is the development of a cooling scheme for the 146.09 gauss atomic clock transitions which makes use of a dark resonance. It is much more effective than methods that avoid coherent effects. The scheme is realised in experiment. Complicated fluorescence data is modelled very well, and predictions for the ion's motional temperature show good agreement with measured values. Data and fits for an ion that has been Doppler cooled below the Doppler limit are presented.

I describe GLOBES, a set of routines that simulates an arbitrary ion in the presence of an arbitrary system of laser beams using the optical Bloch equations. Techniques used to efficiently calculate the steady state, analyse fluorescence scans and solve time-dependent problems for small and large times are discussed. A new routine SILVER IMPER that leapfrogs over the initial dynamics to model the approach to the steady state is introduced.

Doppler cooling in $^{40}Ca^{+}$ is analysed and two extensions made to the basic theory. The 'excursion method' of calculation takes account of the non-linear variation with velocity of the scattering rate. The 'dynamic method' allows for the fact that the ion may not be in equilibrium with the incident radiation during its oscillations, a necessity as the timescale of the external motion is of order the characteristic timescale of the internal motion for standard secular frequencies. This 'dynamic effect' is a general property of trapped ion systems and is also observed in a two-state system.

A two-variable fluorescence scan taken from a four-laser, five-level system in $^{40}Ca^{+}$ is analysed. Techniques to fit large data sets and automatically resolve resonant features are demonstrated. A general treatment of resonant behaviour in three, four and five level pump/probe systems is used to describe the data. This is verified by a second, tailor-made set of scans.

For Rex Hudson
and Spencer Mitchell
who started me down this road.

Acknowledgements

I have made the acquaintance of many wonderful people while working on this thesis. There are too many people to thank here by name. Even if you are not listed here, let me say that I still hugely value everyone who has supported me. Thank you, all of you! You know who you are.

I would like to give a huge thank you to my supervisor, Prof. Andrew Steane. His support during this entire process has been incredible. His knowledge and intuitive understanding of many areas led to many enlightening discussions which never failed to prompt new ideas and avenues that could be explored. In particular, his encouragement and attention to detail and style during the writing of this thesis has been hugely appreciated. Without you, this work would not have been possible.

Thank you too to Prof. Derek Stacey, whose physical intuition and insight is awe-inspiring. I have gained much from working with him, and he was always available with a listening ear and enquiring mind when something interesting came out of the Bloch Equations. His own simulations were of great help and provided assurance of the outputs of my own. He taught me rigour in my work and writing, and I have gained an appetite for fitting many fluorescence scans. Thank you for working through my 'detective stories'!

David Lucas was a huge source of encouragement, especially when I stepped out from the ivory tower and headed into the basement for experimental work. I would like to thank him for all his encouragement and support, and for everything he does to keep the group going behind the scenes.

Many thanks to all those I worked with in the group over the years. I want to give a huge thank you to Dr. Norbert Linke, whose work on background-free detection led to the five-level data that I studied in this thesis. Thank you for your huge and unwavering support of 'banana islands', and for helping me take further data to support my ideas. Without your support, this analysis simply would not have been possible (and I would not have been able to produce a 'banana split'). Thank you too to Dr. Chris Ballance, who was happy for me to take yet more data even as shiny new lasers were being delivered to the lab. Thank you to Dr. David 'D3' Allcock, whose level of knowledge and sheer dedication was hugely inspiring. Thank you also to Dr. Simon Webster, who gave useful advice on technical writing. Dr. Alice Burrell and Dr. David 'Captain Vegetable' Szwer were wonderful members of the office and hugely welcoming when I arrived at the group. Discussions with the Szwer, along with his the-

sis, were a great help when I started writing my simulations. I also very much enjoyed board game night, regularly frequented by Alice, David, Norbert and Simon.

I would like to thank Dr. Tom Harty and Martin Sepiol for providing me with many scans to analyse, and for being so persistent in searching for cold $^{43}Ca^{+}$. Thank you to fellow members of the office Vera Schäfer, Jamie Tarlton, Keshav Thirumalai and Sarah Woodrow for making it such a pleasant place to work. A huge thank you to Diana Prado Lopes Aude Craik for many interesting discussions about atomic physics and Matlab, and for being so enthusiastic about 288 gauss and banana islands. Finally, thank you to visiting students Daniel Shipley, Rustin Nourshargh and Mathieu Nathan Dos Santos for making the office, lab and electronics room great working environments.

In the wider department, thank you to Rob Harris, whose knowledge of all things metal-working is incredible. I learned much from his support in the student workshop. Thanks also to Matt Newport for his superb support after Rob's retirement. Many thanks to Sue Gardner, Monika Porada and Verina Frenzese in the department's office for much administrative help. Thank you too to Axel Kuhn, my college supervisor, with whom I had many interesting discussions about efficient computation.

Dr. Laura Corner has recently provided a chorus of "Have you finished your thesis yet?" every time I have seen her in Oxford. Her advice and encouragement have been a huge help. I have now finished the document - you read it here first!

While working on this thesis, I have also been working professionally with some of the most distinguished choral foundations in the country. I would like to thank everyone in both physics and choral singing for their understanding! I worked with the choir of Christ Church Cathedral for many years, and I have much gratitude for the great work of Dr. Stephen Darlington and Clive Driskill-Smith. Thanks to all those I have worked with over the years, including Michael Hickman, Ghislaine Reece-Trapp, Edward Tambling, Will 'Mad Dog' Robinson and Greg Skidmore. Many thanks to all the alto scholars I worked with: Miachael Hamway, Will 'Bruce' Wayne, Frazer MacDiarmid and Henry 'Kimbauoo' Kimber. Thank you also to James Trickey, who was a really pleasant person to work with. Very many thanks to Alastair Carey, who has been unfailingly supportive and introduced me to many interesting science-fiction films and television series.

In my final year, I worked at Magdalen College. I would like to thank the Informator Choristarum Dan Hyde for his excellent musicianship and for being an inspiring conductor to work for. Thanks to the organists Alexander Berry and Anna Lapwood, who were wonderful company. To all the singers I worked with, including Francis Gush, Seb Lello, James Armitage, Gabriel Bambridge

and Edmund Bridges, thank you for making Magdalen such a welcoming and friendly place to be. I have fond memories of my time at the college.

Many thanks too to everyone working at York Minster and who make it a wonderful place to work. Robert Sharpe is a wonderful person to work for, and he makes great music lists[1]. Thanks to David Pipe and Peter Moger, and to all those I have the pleasure of singing with from day to day, including James Cave, Danny Purtell, Paul Gameson, Josh Madine and Sam Systma.

I have shared my accommodation at 72 Abbey Road with some wonderful and understanding people whilst working on this thesis, including James Potter, Edward Edgcumbe, Tom Chapman (who reintroduced me to James Bond) and Marisa Benoit (who was at the house while I was finishing writing and put up with my being constantly busy). Many thanks to Eileen Head and John Briggs at Christ Church for being so helpful and considerate about my accommodation.

At the start of my final year, I contracted pneumonia with minor complications in my heart and lungs. I have much gratitude for the work of all those at the John Radcliffe Hospital who helped in my recovery. Thanks especially to Dr. Hammersley, Dr. Shah and the nurse on ward, Olga, who did everything they could when I was in hospital, and while I was healing afterwards. Thanks to all of you.

During my spare time, I have indulged in creating artworks and fiction which I post on DeviantArt. Thank you to all those who have viewed and read my work, and all those who have supported me. Thanks to Kyle Doyle, Simon Dixon and Tai Lin, who were encouraging from almost the very outset. Thank you in particular to Andrew Kim, whose work has been the source of much inspiration and who took part in many interesting discussions about art and physics. In addition, very many thanks indeed to Luke and Stan, and everyone in Rangers, who shared their worlds with me when I needed them.

Thanks also to Jeff Vogel, whose excellent company Spiderweb Software produced ways to fill my spare time. Thanks for *Exile*, *Avernum*, *Geneforge* and *Avadon*! Thanks also to all those excellent designers who worked with *Blades of Exile* and *Blades of Avernum* to make further works. In particular, I would like to thank Kelandon for his fantastic creations, including *Bahssikava*, *Exodus* and *The Magic*. Thanks also to LuckySevenDX, raocow and ShadyParadox for creating informative and entertaining little videos to fill the odd spare minute.

Finally I wish to thank everyone in my family, whose support and encouragement throughout this process has been wonderful. Thank you to my parents, to my uncles, aunts and cousins, and to Sophie and Alex who have been incredibly helpful and understanding. Thank you for always being there, and for helping to make this possible.

[1]Today at regular services at the Minster: Finzi's Magnificat. Tomorrow: Haydn with instruments and Lobet den Herrn.

Contents

1 Introduction: Building a quantum computer **1**
- 1.1 Quantum simulation and computation . . . 2
- 1.2 Physical realisations and challenges . . . 3
- 1.3 Quantum computing with trapped ions . . . 6
  - 1.3.1 Scaling trapped ion systems . . . 8
  - 1.3.2 Thesis focus . . . 9
- 1.4 Thesis structure . . . 12

2 Optical Bloch Equations **15**
- 2.1 The differential equations . . . 16
  - 2.1.1 Multiply-connected states . . . 21
- 2.2 Energy eigenstates of an arbitrary ion . . . 24
  - 2.2.1 Operator method . . . 25
- 2.3 Atomic Matrices . . . 27
  - 2.3.1 Stimulated emission rate . . . 27
  - 2.3.2 Spontaneous emission rate . . . 30
  - 2.3.3 Laser Detunings . . . 30
  - 2.3.4 Laser Linewidths . . . 31
  - 2.3.5 Laser Polarizations . . . 32

3 Numerical Methods **33**
- 3.1 The **H** matrix . . . 34
  - 3.1.1 Data structures . . . 35
  - 3.1.2 Creating **H** efficiently . . . 36
  - 3.1.3 Calculations using the **H** matrix . . . 37
- 3.2 Steady-state solutions . . . 38
  - 3.2.1 Solving the Eigensystem . . . 38
  - 3.2.2 Matrix inversion . . . 39
- 3.3 Scanning the steady state . . . 44
  - 3.3.1 Reducing calculation overhead . . . 45
  - 3.3.2 LU preconditioning . . . 45
  - 3.3.3 Improved preconditioners . . . 47
  - 3.3.4 Coherent subsystems . . . 51
  - 3.3.5 Resonance detection . . . 55
  - 3.3.6 Global fitting . . . 57

3.4 Time-dependent solutions . . . 59
3.4.1 Challenges . . . 59
3.4.2 Short times: Product of a matrix exponential . . . 60
3.4.3 Time-varying $\mathbf{H_0}$: Step-size modified Runge-Kutta . . . 62
3.4.4 Large times: Leap-frogging to the steady state . . . 65
3.4.5 Usage of the three routines . . . 71

4 Doppler Cooling in $^{40}Ca^{+}$ 73
4.1 Fitting experimental data . . . 75
4.1.1 Departures from the stationary theory . . . 76
4.1.2 Fitting a scan . . . 81
4.2 Modelling the ion's motion . . . 84
4.3 The steady-state method . . . 85
4.3.1 The Doppler limit . . . 92
4.3.2 Expanding on the steady state . . . 93
4.4 The excursion method . . . 94
4.4.1 Finite heating . . . 96
4.4.2 Fluorescence modification . . . 97
4.4.3 Inadequacies for oscillating ions . . . 100
4.5 The dynamic method . . . 103
4.5.1 The dynamic effect in a two-state system . . . 106
4.5.2 The dynamic effect in $^{40}Ca^{+}$ . . . 115
4.5.3 An excursion method incorporating the dynamic effect . . . 124
4.6 The motional method . . . 124
4.6.1 Dealing with the pole . . . 127
4.6.2 Temperature prediction . . . 131
4.7 Fluorescence of a hot ion . . . 132
4.8 Doppler cooling below the Doppler limit . . . 134
4.8.1 Failure of simple optimisation . . . 135
4.8.2 A minimax approach . . . 137
4.8.3 Minimal temperatures in $^{40}Ca^{+}$ . . . 138

5 Dynamics of a five-level system 142
5.1 The original data . . . 144
5.2 Fitting large, complex data sets . . . 148
5.2.1 Double vetoing . . . 149
5.2.2 Jitter . . . 150
5.2.3 The global optimum . . . 152
5.3 Two-dimensional resonance detection . . . 154
5.4 Resonances in the original data . . . 155
5.5 Experimental analysis of five-level resonances . . . 159

6 A qubit in $^{43}Ca^{+}$ **169**
6.1 Labelling conventions . . . 174
6.2 Atomic clock transitions in $^{43}Ca^{+}$ . . . 176
6.2.1 The 146.09 gauss qubit . . . 182
6.3 Brute-force temperature optimisation . . . 183
6.4 Stretch state cooling . . . 189
6.4.1 Choosing positive or negative M . . . 193
6.4.2 Variation with beam parameters . . . 197
6.4.3 The excursion effect in $^{43}Ca^{+}$ . . . 203
6.5 $\widetilde{\mathbf{F}} = 3$ cooling . . . 205
6.6 850nm/854nm cooling . . . 209
6.7 Fitting fluorescence scans with complicated structure . . . 211
6.7.1 Temperature estimation using the fluorescence curve . . . 220
6.8 Cooling below the Doppler limit . . . 226
6.9 A qubit at 287.78 gauss . . . 235
6.10 Conclusion . . . 236

7 Summary and Outlook **238**
7.1 Further work . . . 243

Bibliography **246**

A Connectivity in an arbitrary system **255**
A.1 A universal input system . . . 255
A.2 Analysing systems of states and beams . . . 256
A.2.1 Detecting connected components . . . 258
A.2.2 Detecting loops . . . 259
A.2.3 Connecting the states . . . 260

B Polarisations in an angled beam **263**
B.1 The polarisation ellipse . . . 263
B.1.1 Optimal polarisations for cooling at $\gamma = \pi/4$ . . . 267
B.2 Polarisations inside the ellipse . . . 267

C The temperature and natural frequency of a two-state system **271**

D Residuals for five-level scans **274**

E Resonant effects in three, four and five-level systems **281**
E.1 The three-state system . . . 282
E.2 The four-state system . . . 288
E.2.1 Differences between three-state and four-state hyperbolae 293
E.2.2 Behaviour with multiple pump beams . . . 296

E.3 The five-state system 299
E.3.1 Differences between four-state and five-state hyperbolae . 304
E.3.2 A bright resonance structure with asymptote η 307
E.3.3 Behaviour with multiple pump beams 308
E.4 Splitting due to Zeeman states 311

F Residuals for $^{43}Ca^{+}$ scans **320**

G Initialisation and Readout **324**
G.1 Initialisation 324
G.2 Readout 327

H Simulation Parameters **333**

1

Introduction: Building a quantum computer

> The scope of the approaches, proposed and underway, to the implementation of quantum hardware is remarkable.
>
> David DiVincenzo

In the foyer of the Clarendon Laboratory in Oxford, there is a display presenting an historic set of data demonstrating the linear relationship between the characteristic X-ray spectra of the elements and their atomic number. This is shown on a plot which was "drawn by H. G. J. Moseley in 1914 in the Electrical Laboratory" [Mos13]. The data points and trends on this plot were added and calculated manually and then the entire plot was drawn by hand.

In the past century, the huge developments in the field of computing have dramatically simplified the process of analysing data. It is now stored, arranged and fitted by a computer on timescales far faster than could be achieved one hundred years ago. However, as ubiquitous as computational simulation has become, there is still a vast set of problems which fundamentally cannot efficiently be tackled by even the most advanced machines currently available to us. Our computers are physical objects constructed on the basis of classical information theory, and they work well for tackling classical physical problems.

However, quantum mechanical problems are a different story.

1.1 Quantum simulation and computation

The problem of simulating quantum mechanical systems was discussed by Feynman in an address made to the MIT Physics of Computation conference held in 1981 [Fey82]. In this address, he concluded that classical simulators could not compute quantum mechanical problems efficiently. For example, let us consider an ensemble of n quantum systems, each of which consists of m energy states. In order to describe this ensemble fully, we need to keep track not only of the individual state of each particle, but also the coherent interaction between each pair of the systems. Modelling the ensemble in a classical sense involves $m^n - 1$ variables. Adding one more particle to the ensemble increases the calculation time by a factor m each time. The calculation time increases exponentially with the size of the problem, so that large problems quickly become intractable. Feynman proposed to solve this problem by using a different type of computer. To quote Norbert Wiener, "the best material model of a cat is another, or preferably the same, cat." Instead of using a classical object for computation, we instead use a *quantum* computer. Consider modelling this example with an artificial duplicate of the original system. Quantum mechanical objects inherently contain phase and coherence information, so adding an extra particle increases the calculation time only by the overhead required to realise the additional quantum system and not the additional coherences. In effect, a quantum mechanical simulator scales as fast as the problem. At first, it was not clear whether it was possible to construct such a simulator. This was an open question for a decade.

David Deutsch [Deu85] considered the idea of quantum simulation and took it in a different direction. Instead of simulating a quantum mechanical object, he considered using the inherent coherent properties of such a system explicitly as a

computer. He proposed a quantum computer and an algorithm that would run on it, an algorithm that runs faster than on a classical computer [DJ92]. Since then, many quantum-mechanical algorithms have been developed that could run on a quantum computer, including Grover's algorithm [Gro97, Nov01], quantum teleportation [RHR+04, BCS+04] and the quantum Fourier transform [Sho97, HS00]. An important example is Shor's factorisation algorithm [Sho97], of particular relevance to cryptography. The standard basis of the vast majority of encrypted transactions at present is the RSA public-key cryptography system. This system relies on the fact that classical computers cannot factorise an n-digit number in polynomial time; the most efficient known algorithms work in sub-exponential time. Information is encoded in a publicly available key which is the product of two large primes, and accessing the information depends on knowing one of the component primes. Decomposing this number into its two parts is unfeasibly slow for a classical computer. Shor's algorithm, however, factorises a number in polynomial time, meaning that such keys could be decomposed very quickly. For example, factoring a number with 10^4 digits, Shor's algorithm is $\sim 10^4$ times faster than the equivalent classical algorithm. However, much more work is needed on practical quantum devices before the RSA system is threatened! For a further list of quantum algorithms, see [Jor15].

1.2 Physical realisations and challenges

Several technologies have been developed to realise a quantum computer. Most make use of binary objects to encode information in an analogous way to classical computers, except that, instead of just two physical states per bit, there is a two-dimensional Hilbert space. These objects are referred to as 'quantum bits' or 'qubits'. Qubits have been constructed using many quantum objects, including photons, neutral atoms, quantum dots, liquid molecule NMR, superconducting systems and trapped ions. For a review of potential qubit systems, see [LJL+10].

In addition to the technology required to realise these qubits, techniques have been developed to carry out computation with them. In analogy to classical computation, the devices used to perform operations on qubits for the purposes of computation are referred to as 'gates'.

Each of these physical realisations of a quantum computer have associated experimental results. In order to distinguish between the level of development of each of these systems, a set of criteria developed by David DiVincenzo [DiV00] is commonly used. These criteria list a set of five key features that must be present for a practical quantum computer. These criteria are:

1. A scalable physical system with well characterised qubits
2. The ability to initialise the state of the qubits to a simple fiducial state such as $|000\ldots\rangle$
3. Long relevant decoherence times, much longer than the gate operation time
4. A 'universal' set of quantum gates
5. A qubit-specific measurement capability

Each of the many physical realisations of a quantum computer fulfil some of these criteria to a greater or lesser degree. However, the first and fourth of these criteria are not met by the majority of current technologies. The current major problems for quantum computing are scaling systems to large numbers of qubits and the errors introduced by quantum gates.

The problem of scaling up a small number of qubits to a larger system is by no means straightforward. To produce a large quantum computer, it is necessary to prepare a large number of quantum objects. However, the requirement of a universal set of quantum gates means that we also need to perform operations that address each one of these qubits in isolation or in pairs. Addressing

each qubit separately while preserving the ensemble as a whole is difficult for large quantum systems. Different qubit technologies may find one or other of these tasks technically challenging. For example, for trapped ions, it is simple to address any ion singly by using a tightly-focused beam. However, two-qubit operations require making use of the modes of vibrational motion in the trap. These motional modes become harder to resolve as the number of ions in the trap increases, so that computing is impractical for chains of approximately ten or more ions (see, for example, [Har12]). By contrast, very many neutral atoms can be trapped in an optical lattice. By manipulating the laser beams that form the lattice, operations can be performed on the entire ensemble of atoms with relative ease. However, because the lattice is formed through interference of optical beams, the atoms are located very close to one another, typically at a distance of half the wavelength of the light forming the lattice. Using conventional techniques, these distances cannot be resolved by an incident laser beam; the atoms cannot be addressed individually by using a tightly-focused beam. In this system, manipulating the ensemble is relatively easy, but addressing individual qubits is hard. Manipulating large numbers of qubits both singly and together is one of the main challenges of quantum computing at present.

The second major challenge is the fidelity of operations on the qubits. This includes both initialisation and readout of the qubits and also logic gate operations. It is not merely the existence of a universal set of quantum gates that is important for the practical implementation of a qubit: the error in the output state of a gate must be sufficiently small (of the order $1/N$ for a sequence of N gates). Coherence in quantum systems is very fragile and easily degraded or destroyed entirely by interactions with the environment, so it might at first seem impractical to produce a useful quantum computer. However, in the mid-1990s Peter Shor and Andrew Steane independently developed the concept of quantum error correction. Quantum error correction achieves for quantum computers what

the far simpler classical error correction routines achieve on conventional computers: extra 'ancilla' qubits distribute the information over more degrees of freedom in such a way that the most likely unwanted departures in the state can be detected and corrected for. There is a tradeoff between the correction power and the ratio of the number of ancilla qubits to physical qubits. The larger the gate error, the more ancilla qubits are required to correct the errors in an algorithm of a given length [Ste03]. This trade-off has often been interpreted in terms of a threshold of the gate error, below which an algorithm of any practical length could be corrected for by making use of a qubit ratio that is no more than a few orders of magnitude. This threshold has been estimated to be between 10^{-2} and 10^{-4} depending on the exact nature of the qubit and the measurement process. Achieving a fidelity of 99.9% on all gate operations is a difficult feat.

1.3 Quantum computing with trapped ions

Trapped ions are good candidates for qubits. Firstly, we can be confident that each atom of a given species will be the same regardless of the physical system used to examine it. Secondly, the electric charge of an ion makes it both relatively easy for an ion to be confined in space using electromagnetic fields and for individual ions to interact with one another. Currently, trapped ions are among the most advanced implementation of a quantum computer [MK13, BW08], and significant steps have been taken towards resolving the problem of scalability and the fidelity of qubit operations. Superconducting Josephson junctions have also made impressive advances in the past several years [DS13, CGC+12].

The techniques and technology for localising charged particles in space have been available for just over half a century. In the 1950s, two distinct types of electromagnetic trap were proposed and constructed. Earnshaw's theorem states that it is impossible to trap a charged particle in space solely using static electric fields. In essence, this is a simple extension of Gauss's law, stating that the diver-

gence of all fields in free space must be zero. Wolfgang Paul and Hans Dehmelt independently proposed two ways around this problem. Dehmelt used a combination of static electric and magnetic fields to trap a particle in what is now referred to as a 'Penning Trap' [JSD78, Pen36]. Paul used a combination of static and oscillating radio-frequency electric fields; this type of trap is now referred to as a 'Paul Trap' [PS53].

Trap technology has improved dramatically in the past several decades, and the process of constructing a trap and trapping single ions is now a standard technique in many laboratories. The ion trap is operated in an ultra-high vacuum environment. Typical pressures are of order 10^{-11} Torr. By heating a small oven, a beam of neutral atoms is produced that passes over the trapping region. Some of the atoms are ionised by radiation incident on the beam. Unionised atoms do not interact with the trapping fields and pass the trap unaffected, whereas the ions are trapped. By performing this process carefully, a chosen number of ions can be localised in space. The experiments described in this thesis make use of a modification of the Paul Trap. Instead of the original design, which uses a three-dimensional array of conductors, the trap is layered down on a two-dimensional surface [AHB+13].

A possible qubit is formed from two energy states of the ion. There are several choices for the types of states used. 'Optical qubits' make use of a state in the ground level and one in an electronically excited level; the frequency of light that excites transitions between these two states is in the optical range. These qubits have the advantage that reading out the state can be achieved with high fidelity. However, the finite lifetime of the upper state provides a fundamental limitation to the fidelities of operations and memory. This in turn leads to the need for narrow band lasers exciting forbidden transitions, increasing the experimental complexity. The type of qubit considered in this thesis is a 'hyperfine' qubit formed of two states in the ground state hyperfine manifold of the ion.

These states have lifetimes that are in practical terms infinite; coherence times of 10 minutes have been measured [BHI+91], and there is no reason why longer coherence times could not be observed in the future.

Initialising the state of an ion is achieved by making use of optical pumping. For the qubit described in this thesis, the readout process involves two stages. In the first phase, a state selective process moves population to a 'shelf' state from which it cannot escape. This process only acts on population in one of the qubit states. In the second phase, transitions are driven among the other levels of the system. If fluorescence is observed, it is inferred that the ion is not in the shelf state and the process records a negative result; the ion will only be dark if all of its population is in the shelf. This allows a determination of the state of the qubit.

Single qubit logic operations are performed by applying laser or microwave pulses to the ion. Multiple qubit operations make use of the Coulomb interaction between ions with the trap which affects the physical motion of the ions. This motion is described by a series of quantised harmonic oscillator modes. Ions are excited from one mode to another in a state-dependent fashion using laser beams. In essence, the gate uses the joint motion degree of freedom to couple the internal states of a chosen number of ions. There are a number of various proposals for two-qubit gates, including the Cirac-Zoller gate [CZ95, SKHR+03], the Mølmer-Sørenson gate [SM99, SM00, SKK+00] and the Wobble gate [LDM+03, Hom06, HML+06]. For a review of ion trapping for quantum computing purposes, see [HHB08].

1.3.1 Scaling trapped ion systems

Much progress has been made on scaling to large numbers of ions and increasing the fidelity of all operations. The number of ions in a single trap that can be manipulated easily is fundamentally limited by the spectral resolution of the

vibrational modes and various technical problems with cooling and controlling large strings of ions. A promising proposed solution is the 'Quantum CCD' architecture [WMI+98]. A large number of ion traps are connected together by a series of electrodes. These electrodes can be used to shuttle ions between the traps by varying the voltages. Different traps can be used for readout, memory and qubit operations, allowing each region to be tailor-made for a specific function. This approach is inherently scalable, but comes with many challenges. The number of regions which need to be accessed by laser beams poses problems, as does the sheer density of electronics, although some of the problems associated with the large number of beams may be alleviated by making use of microwave addressing [CLH+13, Cra16]. In addition, managing and automating the shuttling is a difficult problem. One possible solution to alleviate this problem of complexity is to use a 'Networked Cell' architecture [MK13]. This proposal links quantum CCDs together by using photons. The advantage of such a scheme is that only operations within each CCD need to be high fidelity; even very noisy photonic interaction can be corrected for [NFB14].

1.3.2 Thesis focus

The numerical methods and programmes described in this thesis are of quite general application, but they were developed primarily to investigate the behaviour of the $^{43}Ca^{+}$ under various conditions. We wish to perform high-fidelity gate operations using qubit states chosen from the ground-state hyperfine manifold. We therefore need to Doppler-cool the ion to a temperature low enough for sideband cooling to be effective in bringing the ion to the ground state of the trap. For readout purposes, it must also be possible to obtain high fluorescence from the ion when it is cold. Finally, it is necessary that all this should be done in a magnetic field set at a specific value, for the following reason. The separation of any pair of states depends on the ambient magnetic field. This is subject to

fluctuations, leading to decoherence. A standard method of reducing this effect is to choose states whose separation at some particular value of the field is independent of the field magnitude to first order ([LOJ+05, OWC+11]). In the case of our chosen qubit, this is 146.09 gauss.

To satisfy all these requirements represents a challenge, both experimentally and theoretically. A term diagram of $^{43}Ca^{+}$ is given later in this thesis (see figure 6.1), but for the present purposes we note that in the simplest possible cooling scheme one uses the transition at 397nm between the ground $S_{1/2}$ level and the $P_{1/2}$ level. Spontaneous transitions can occur from the latter to the $D_{3/2}$ and rep umping with 866nm radiation is therefore necessary. The nuclear spin of $^{43}Ca^{+}$ is 7/2; hence, even this simple scheme involves 64 states and $64^2 + 64 = 4160$ coherences. Any calculation must take them all into account; whatever light polarisations are used, population can in principle always spread into all the states. There are two further complications. The large separation between the two groups of hyperfine states in the $S_{1/2}$ (3.1GHz) necessitates the use of two laser frequencies at 397nm. Even within each group, however, the Zeeman effect at 146.09 gauss causes adjacent states in the ground level to be separated by $\sim$50MHz. This is large enough to raise serious doubts concerning the level of fluorescence and efficiency of Doppler cooling one might expect without the deployment of more than two laser frequencies.

These problems were tackled by constructing a suite of numerical routines to study the dynamics of an ion using the optical Bloch equations. This is the 'Global Linear Optical Bloch Equations Suite' (GLOBES). GLOBES was designed to simulate an arbitrary ion in the presence of an arbitrary laser system, provided that LS coupling holds and that, except under exceptional circumstances, no two laser beams address the same transition between any pair of states. GLOBES has the capacity to efficiently optimise any output, such as fluorescence or temperature, allowing the development of new schemes for cooling

an ion and achieving high fluorescence. The routines can also fit very complicated fluorescence scans such as those output by $^{43}Ca^{+}$ at 146.09 gauss and use these fits to predict the temperature of the ion. This allows a direct comparison between the predictions of the simulations and the results of temperature measurements using sideband thermometry. These simulations are efficient in terms of speed and memory usage; optimisations of GLOBES are described in this thesis. The routines can be used on a standard laptop computer; no advanced computing resources are required. A standard ion can be analysed on a machine with no more than 4GB RAM, and calculating the steady state of an ion, even one as complicated as $^{43}Ca^{+}$, typically takes no longer than 0.1s.

GLOBES was used to develop a cooling scheme for our atomic clock transition that is approximately as effective as that for $^{40}Ca^{+}$. While there is no perfect closed set of states in $^{43}Ca^{+}$, there are two sets of states that are approximately closed. Each is centred around a different dark resonance. We make use of one of these dark resonances to cool the ion. Such dark resonance cooling has been previously proposed [RLB96, RAN+02] and recently demonstrated [Har12] for systems without hyperfine structure. However, there is little advantage over conventional cooling routines for such systems. This ion is the first demonstration of dark resonance cooling for a system with hyperfine structure. Using this dark resonance cooling scheme, the temperature of the $^{43}Ca^{+}$ ion was Doppler cooled to $\sim$0.4mK, a temperature 30% lower than the two-state Doppler limit and on order of magnitude lower than would otherwise be possible.

In order to model the process of Doppler cooling in $^{43}Ca^{+}$, two extensions had to be made to the standard theory (see, for example, [WI79, IW82, LBMW03, LL15, Sta03]) to take account of two effects. The 'excursion method' of calculation takes account of the non-linear variation with velocity in the rate the ion scatters photons as it oscillates in the trap. The 'dynamic method' explicitly includes the secular motion of the ion in the trap in the model of the interactions

with the laser beams. As the ion oscillates in the trap, the external timescale of its oscillations may be on a similar scale to the characteristic internal timescale of the atomic dynamics. In such a case, the ion will never be in perfect equilibrium with the radiation at any point during its cycle of external motion, something inherently assumed by the standard theory. The dynamic method takes account of this. This 'dynamic effect' is also observed in the two-state system. Finally, the 'motional method' combines both the excursion and dynamic methods in the same theory. This method provides a very good match to the ion's fluorescence and even accurately models the fluorescence of a very hot ion.

The laser cooling methods described in this thesis have been implemented in the Oxford ion trap quantum computing group. These experiments have demonstrated several promising figures of merit. A coherence time of $50(10)$ s has been recorded [HAB+14], the longest time measured for any single qubit. A combined state preparation and readout error (SPAM) of $6.8(6) \times 10^{-4}$ has been reached and a single qubit gate error of $1.0(3) \times 10^{-6}$ obtained. In addition, recent results show a two qubit gate error of $< 5.0 \times 10^{-2}$ [Sep16]. These results are at least an order of magnitude better than previous work, and they are all obtained in the same experiment. In addition, high fluorescence of 25000 counts per second has been recorded. For this qubit (and assuming the results of networked cell architecture), we have demonstrated all five of the DiVincenzo criteria for this ion and gate operations that are approaching the fault tolerant threshold.

1.4 Thesis structure

Chapter two discusses the optical Bloch equations. The full equations including relaxation are derived and I discuss in what circumstances these equations hold. Notation adopted for properties of the atomic and laser systems, including Rabi frequencies and spontaneous emission rates is presented. The conventions used

to define laser detunings and polarisations are also discussed.

Chapter three describes the simulation itself. I introduce the notation used to refer to the populations and coherences along with the conventions used to store these in data structures. Efficient methods of solving the time-independent problem; calculating steady-state 'scans' in which one laser parameter is varied over a set of data; and solving the time-dependent problem for short and large times and time-dependent Hamiltonians are described. A new routine is described that calculates the approach of a time-dependent problem to the steady state by leapfrogging over the initial fast dynamics.

Chapter four discusses Doppler cooling in the simple $^{40}Ca^{+}$ isotope and uses four models of the effect of the ion's motion on its temperature. Data in the form of fluorescence scans is analysed, and methods used to fit such scans are discussed. Each of the four models is analysed in turn, including their advantages and disadvantages. In addition, an optimisation process is introduced that demonstrates that cooling below the Doppler limit can be achieved in this isotope with high fluorescence rates.

Chapter five describes the analysis of a complex set of data taken in five-level $^{40}Ca^{+}$ illuminated with four laser beams, one of which is much more powerful than the others. Efficient techniques to fit large data sets are mentioned and a method to plot the data that automatically adjusts the step-size around resonance features is introduced. The data is explained using a general treatment of resonant behaviour present in three, four and five level pump/probe systems. Finally, a new set of data is analysed which validates the five-level resonance model.

Chapter six describes work relating to atomic clock transitions in $^{43}Ca^{+}$. All the atomic clock transitions in the five lowest energy levels of this isotope are presented, and those in the ground state are characterised in detail. The results of temperature minimax optimisation are presented. Several cooling schemes

are discussed, including cooling schemes that make use of a dark resonance. Methods of fitting the very complicated fluorescence scans output by this isotope are demonstrated and fits are presented for ions with increasing levels of heating. High fluorescence rates are observed, with populations in the upper energy level of around 7%. A fluorescence scan is analysed for which the ion was Doppler cooled and its temperature measured to be below the Doppler limit for an upper state population of several percent. This result is confirmed by simulation. Data showing variation in the ion's temperature with beam detunings and intensities is also matched very well by the simulation.

2

Optical Bloch Equations

I ... conclude by saying to you: dream your dreams and may they come true!

Felix Bloch

Atomic systems with complex structure have many potential benefits for quantum information processing. To take advantage of these benefits, it is necessary to calculate the behaviour of such systems. Simple methods using a rate equations approach are available and have been used extensively [Szw09, Web05]. However, these methods do not model coherent effects which are a key feature of the behaviour of such ions. To take account of them, one may use the optical Bloch equations.

The optical Bloch questions equations for two-state systems are well documented (see, for example [Lou97]), but treatments for a general atomic system are less so ([MK85] is one example). In order to demonstrate the advantages and limitations of this approach, along with introducing the notation that will be used throughout this thesis, I derive the full form of the equations in detail and describe the information necessary for using them to describe an ion in an arbitrary static $\mathbf{B}$ field.

2.1 The differential equations

We begin with the Schrödinger equation for the density matrix:

$$\dot{\rho} = -\frac{i}{\hbar}[\mathcal{H}, \rho] \tag{2.1}$$

Let us write the Hamiltonian as the sum of the Hamiltonian of an isolated ion, $\mathcal{H}_0$, and a perturbation arising from the interaction of the ion with some radiation, V:

$$\mathcal{H} = \mathcal{H}_0 + V \tag{2.2}$$

A particular element ρ_{jk} of the density matrix obeys the differential equation:

$$\dot{\rho}_{jk} = -\frac{i}{\hbar}\left\{\langle j|\left(\mathcal{H}_0 + V\right)\sum_{m,n}\rho_{mn}|m\rangle\langle n|k\rangle - \langle j|\sum_{m,n}\rho_{mn}|m\rangle\langle n|\left(\mathcal{H}_0 + V\right)|k\rangle\right\} \tag{2.3}$$

The terms involving the unperturbed Hamiltonian can be dealt with simply. By definition:

$$\mathcal{H}_0|s\rangle = E_s|s\rangle \tag{2.4}$$

where E_s is the energy of the unperturbed state $|s\rangle$. These terms contribute an energy term to the differential equation:

$$\dot{\rho}_{jk} = -i\rho_{jk}\frac{E_j - E_k}{\hbar} - \frac{i}{\hbar}\left\{\langle j|V\sum_{m,n}\rho_{mn}|m\rangle\langle n|k\rangle - \langle j|\sum_{m,n}\rho_{mn}|m\rangle\langle n|V|k\rangle\right\} \tag{2.5}$$

From the orthogonality of the states, we can reduce the sum to one only over a single variable:

$$\dot{\rho}_{jk} = -i\rho_{jk}\frac{E_j - E_k}{\hbar} - \frac{i}{\hbar}\sum_{l}\left\{\rho_{lk}\langle j|V|l\rangle - \rho_{jl}\langle l|V|k\rangle\right\} \tag{2.6}$$

Thus, in the absence of a perturbation:

$$\dot{\rho}_{jk} = -i\rho_{jk}\frac{E_j - E_k}{\hbar} \tag{2.7}$$

This can be solved simply:

$$\rho_{jk} = \rho_{jk,0}\exp\left(-i\frac{E_j - E_k}{\hbar}t\right) \tag{2.8}$$

giving exactly what we would expect from an isolated system. However, this solution shows fast oscillations. Dealing with these numerically would be difficult. Instead, then, it is better to make a transformation to a set of parameters which do not involve these oscillations. We introduce the matrix C with elements C_{jk} that obey:

$$C_{jk} = \rho_{jk} \exp\left(+i\frac{E_j - E_k}{\hbar}t\right) \tag{2.9}$$

so that, for a system in the absence of any perturbation, the C_{jk} will not oscillate. We can construct a differential equation for C_{jk}:

$$\dot{C}_{jk} = \dot{\rho}_{jk} \exp\left(i\frac{E_j - E_k}{\hbar}t\right) + \rho_{jk}\left(i\frac{E_j - E_k}{\hbar}\right)\exp\left(i\frac{E_j - E_k}{\hbar}t\right) \tag{2.10}$$

so that:

$$\dot{C}_{jk} = -\frac{i}{\hbar}\sum_l \left\{ C_{lk} \exp\left(i\frac{E_l - E_j}{\hbar}t\right)\langle j|V|l\rangle - C_{jl}\exp\left(i\frac{E_k - E_l}{\hbar}t\right)\langle l|V|k\rangle \right\} \tag{2.11}$$

Let us consider a single laser driving a single-photon transition between a given pair of states. We will generalise to more lasers later. If the laser has an electric field with amplitude $\mathbf{E_0}$, the perturbation matrix elements take the form:

$$\langle m|V|n\rangle = \langle m| - \mathbf{d}\cdot\mathbf{E_0}|n\rangle \cos\left\{\left(\frac{E_m - E_n}{\hbar} + \Delta_{mn}\right)t\right\} \tag{2.12}$$

where $\mathbf{d}$ is the electric dipole operator for the ion and Δ_{mn} is the detuning of the incident radiation from the transition $|m\rangle \leftrightarrow |n\rangle$. Let us define:

$$R_{mn} = \frac{1}{2\hbar}\langle m|\mathbf{d}\cdot\mathbf{E_0}|n\rangle = \frac{\Omega_{mn}}{2} \tag{2.13}$$

$$p_{mn} = \frac{E_m - E_n}{\hbar} + \Delta_{mn} \tag{2.14}$$

We define the rate R_{mn} as half of the Rabi frequency Ω_{mn} as this simplifies the factors of two in the final expressions. The perturbation matrix element can be written in these terms:

$$\langle m|V|n\rangle = -\hbar R_{mn}\left\{\exp\left(ip_{mn}t\right) + \exp\left(-ip_{mn}t\right)\right\} \tag{2.15}$$

so that:

$$\dot{C}_{jk} = -\frac{i}{\hbar}\sum_l \left\{ C_{lk}\left(\exp\left\{i\left(2\frac{E_l - E_j}{\hbar} - \Delta_{jl}\right)t\right\}\hbar R_{jl} + \exp\left\{i\Delta_{jl}t\right\}\hbar R_{jl}\right) + C_{jl}\left(\exp\left\{i\left(2\frac{E_k - E_l}{\hbar} - \Delta_{lk}\right)t\right\}\hbar R_{lk} + \exp\left\{i\Delta_{lk}t\right\}\hbar R_{lk}\right)\right\} \tag{2.16}$$

Applying the rotating wave approximation removes the rapidly oscillating terms:

$$\dot{C}_{jk} = i\sum_l \left\{C_{lk}R_{jl}\exp\left(i\Delta_{jl}t\right) - C_{jl}R_{lk}\exp\left(i\Delta_{lk}t\right)\right\} \tag{2.17}$$

For a complete treatment, we must also consider the relaxation of the ion. We can take this into account at this point as we know how such processes should behave; a relaxation process is a decay process with some decay constant. First we consider radiative decay, characterised by the partial Einstein coefficient A_{jk} between $|j\rangle$ and $|k\rangle$. These terms cause the populations to relax. Hence the diagonal terms now read:

$$\dot{C}_{jj} = \sum_l \left\{iC_{lj}R_{jl}\exp\left(i\Delta_{jl}t\right) - iC_{jl}R_{lj}\exp\left(i\Delta_{lj}t\right) + A_{lj}C_{ll} - A_{jl}C_{jj}\right\} \tag{2.18}$$

Let us define the total radiative decay constant Γ_j for $|j\rangle$ as:

$$\Gamma_j = \sum_{l=1}^{\infty} A_{jl} \tag{2.19}$$

In all that follows, we use the term 'state' to refer to an energy eigenstate of the Hamiltonian. $|j\rangle$ and $|k\rangle$ are two examples. The term 'level' is used to describe groups of states that arise from LS coupling such as, for example, the $S_{1/2}$ and $D_{5/2}$ levels. A_{jl} describes the radiative decay between two states. If state $|j\rangle$ is in level a and $|k\rangle$ in level b then we can also define the full Einstein A-coefficient for transitions between $|j\rangle$ and level b:

$$A(a,b) = \sum_{l\in b} A_{jl} \tag{2.20}$$

where the sum is across all the states in level b. We can therefore write the more familiar expression:

$$\Gamma_j = \sum_{b=1}^{\infty} A(a,b) \tag{2.21}$$

Writing the differential equation in terms of the radiative decay rates:

$$\dot{C}_{jj} = \sum_l \{iC_{lj}R_{jl}\exp(i\Delta_{jl}t) - iC_{jl}R_{lj}\exp(i\Delta_{lj}t) + A_{lj}C_{ll}\} - \Gamma_j C_{jj} \tag{2.22}$$

In order to proceed any further, it will be useful to rotate the system to remove as much of the remaining time dependence in the differential equation as possible. As such, we define:

$$\sigma_{jk} = C_{jk}\exp(-i\Delta_{jk}t) \tag{2.23}$$

The diagonal terms for this matrix now have a differential equation that has no explicit time dependence:

$$\dot{\sigma}_{jj} = \sum_l \{i(R_{jl}\sigma_{lj} - \sigma_{jl}R_{lj}) + A_{lj}\sigma_{jj}\} - \Gamma_j\sigma_{jj} \tag{2.24}$$

Let us now consider the off-diagonal terms of (2.17):

$$\dot{C}_{jk} = i\sum_l \{C_{lk}R_{jl}\exp(i\Delta_{jl}t) - C_{jl}R_{lk}\exp(i\Delta_{lk}t)\} \tag{2.25}$$

We must add relaxation terms into these equations, too. Note that the coherences depend on the amplitude of each state, rather than the populations; for spontaneous emission, the decay rates of the coherences are equal to half that of the populations. The spontaneous terms do not otherwise affect the coherence terms, so the amount of coherent phase between the states is not affected by spontaneous decay. The resulting equations, including the relaxation terms, are:

$$\dot{C}_{jk} = i\sum_l \{C_{lk}R_{jl}\exp(i\Delta_{jl}t) - C_{jl}R_{lk}\exp(i\Delta_{lk}t)\} - \frac{\Gamma_j+\Gamma_k}{2}C_{jk} \tag{2.26}$$

There is another effect which affects these coherence terms. The lasers themselves will have some finite linewidth Γ_L which will reduce the magnitude of

the coherence terms. We can apply this to the off-diagonal terms in the same way as the linewidths of the states:

$$\dot{C}_{jk} = i\sum_l \{C_{lk}R_{jl}\exp(i\Delta_{jl}t) - C_{jl}R_{lk}\exp(i\Delta_{lk}t)\} - \frac{\Gamma_j+\Gamma_k+\Gamma_L}{2}C_{jk} \quad (2.27)$$

Writing this in terms of the rotated matrix σ:

$$\begin{aligned}&(\dot{\sigma}_{jk} + i\Delta_{jk}\sigma_{jk})\exp(i\Delta_{jk}t) = \\ &i\sum_l \{\sigma_{lk}R_{jl}\exp(i(\Delta_{lk}+\Delta_{jl})t) - \sigma_{jl}R_{lk}\exp(i(\Delta_{lk}+\Delta_{jl})t)\} - \\ &\frac{\Gamma_j+\Gamma_k+\Gamma_L}{2}\sigma_{jk}\exp(i\Delta_{jk}t)\end{aligned} \quad (2.28)$$

We note that:

$$\Delta_{jk} = \Delta_{jl} + \Delta_{lk} \quad (2.29)$$

This allows us to remove the oscillating terms and write these equations in a similar form to the diagonal terms:

$$\dot{\sigma}_{jk} = i\sum_l \{\sigma_{lk}R_{jl} - \sigma_{jl}R_{lk}\} - \sigma_{jk}\left\{i\Delta_{jk} + \frac{\Gamma_j+\Gamma_k+\Gamma_L}{2}\right\} \quad (2.30)$$

In conclusion, equations (2.24) and (2.30) are the equations we need to solve. It is more convenient from the point of view of computation to deal with purely real equations. To this end, let us define:

$$\sigma = x + iy, \qquad x_{jk}, y_{jk} \in \mathbb{R} \quad (2.31)$$

This produces the set of real differential equations:

$$\dot{x}_{jj} = \sum_l \{(R_{jl}y_{lj} - y_{jl}R_{lj}) + A_{lj}x_{jj}\} - \Gamma_j x_{jj} \quad (2.32)$$

$$\dot{x}_{jk} = \sum_l \{-R_{jl}y_{lk} + y_{jl}R_{lk}\} - \frac{\Gamma_j+\Gamma_k+\Gamma_L}{2}x_{jl} + \Delta_{jk}y_{jk} \quad (2.33)$$

$$\dot{y}_{jj} = \sum_l \{(R_{jl}x_{lj} - x_{jl}R_{lj}) + A_{jl}y_{ll}\} - \Gamma_j y_{jj} \quad (2.34)$$

$$\dot{y}_{jk} = \sum_l \{R_{jl}x_{lk} + x_{jl}R_{lk}\} - \frac{\Gamma_j+\Gamma_k+\Gamma_L}{2}y_{jl} - \Delta_{jk}x_{jk} \quad (2.35)$$

We can reason out an important property from the equations written in this way. From the definition of the density matrix:

$$y_{jj},\ \dot{y}_{jj} = 0 \tag{2.36}$$

The only way this can be true in general is if:

$$R_{jl} = R_{lj} \tag{2.37}$$

where we have used the first of the symmetry properties of the density matrix:

$$x_{jl} = x_{lj} \tag{2.38}$$

$$y_{jl} = -y_{lj} \tag{2.39}$$

We can therefore write the Optical Bloch Equations as:

$$\dot{x}_{jj} = \sum_l \{2R_{jl}y_{lj} + A_{lj}x_{jj}\} - \Gamma_j x_{jj} \tag{2.40}$$

$$\dot{x}_{jk} = \sum_l \{-R_{jl}y_{lk} + y_{jl}R_{lk}\} - \frac{\Gamma_j + \Gamma_k + \Gamma_L}{2}x_{jl} + \Delta_{jk}y_{jk} \tag{2.41}$$

$$\dot{y}_{jk} = \sum_l \{R_{jl}x_{lk} + x_{jl}R_{lk}\} - \frac{\Gamma_j + \Gamma_k + \Gamma_L}{2}y_{jl} - \Delta_{jk}x_{jk} \tag{2.42}$$

This form of the equations was derived considering only a single laser transition, but the result can be simply generalised to systems with multiple laser transitions provided that it is possible to define a unique value of the laser detuning between every pair of states in the system. I describe such a system as being 'singly-connected', and systems for which unique detuning cannot be defined as 'multiply-connected'. See the following section for further discussion.

In order to proceed further, we need to produce expressions for R_{jk}, A_{jk}, Δ_{jk}, Γ_j, and Γ_L in terms of fundamental atomic properties.

2.1.1 Multiply-connected states

We can use the same process to describe atomic excitation by more than one laser. It is not sufficient simply to add terms to (2.40) - (2.42), however, because

now the procedure of introducing σ in (2.23) breaks down in general. Consider two states that are not connected directly, but by one or more intermediate states. For example, see figure 2.1 and consider one state in the $S_{1/2}$ and another in the $P_{3/2}$. In general, there will in fact be more than one intermediate state by which the two states are connected. To describe this, we simply replace rates and detunings in (2.15) with a summation over the Rabi frequencies and detunings of every possible connection between the two states:

$$\langle m|V|n\rangle = -\hbar \sum_{\alpha} \left(R_{mn}^{\alpha} \left\{\exp\left(ip_{mn}^{\alpha}t\right) + \exp\left(-ip_{mn}^{\alpha}t\right)\right\}\right) \tag{2.43}$$

The superscripts denote different possible connections between the two states. The differential equation can thus be written:

$$\dot{C}_{jk} = i \sum_{l} \left\{C_{lk}R_{jl}^{\alpha} \exp\left(i\Delta_{jl}^{\alpha}t\right) - C_{jl}R_{lk}^{\alpha} \exp\left(i\Delta_{lk}^{\alpha}t\right)\right\} \tag{2.44}$$

The difficulty that now arises is that it is now impossible to simply rotate the system from the C basis by some angle to remove the dependence on the detunings. There is more than one 'detuning' between a given pair of states. In general, then, we cannot write the equations in a simple linear form [SLA+04].

However, all is not lost! For most cases we will want to deal with, all laser connections between pairs of states *will* have the same net detuning. A difference in detuning between different paths will only arise in two situations. Firstly, if the lasers connecting a set of levels in an atomic system form a closed loop with a non-zero net detuning; that is, it is possible to travel from one level and return to it again by tracing a path along the driven transitions such that each transition is passed only once and that a net detuning is accumulated over the path. Such loops are usually not desirable in atomic systems as they produces unnecessary ambiguity about the evolution of the state of the ion. Secondly, any two levels directly connected by one laser must not be directly connected by another laser. In other words, lasers with sidebands cannot strictly be modelled using a linear set of equations. However, we can even get around this

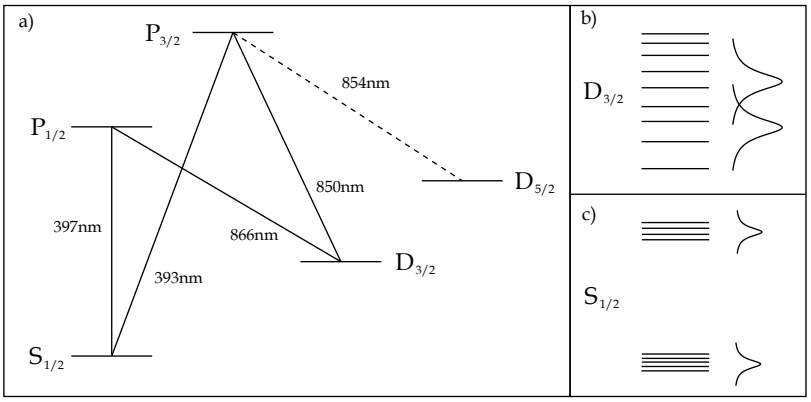

Figure 2.1: A example of systems with multiply-connected states. a) An atomic system involving a closed loop (shown by the solid lines). This cannot be modelled using the linear equations if the net detuning around the loop is non-zero. b) A level connected by a laser with sidebands. The excitation probability is shown schematically next to the states. As the probabilities of the two sidebands overlap, this situation cannot be modelled. c) Another level connected by a laser with sidebands. The hyperfine manifolds and the sideband excitations are well separated. This situation can be modelled.

problem when the sidebands of a laser have such a large frequency splitting that each sideband excites an independent group of states. For example, it is possible to model the use of a laser beam with sidebands to address a level with a large hyperfine splitting.

We shall refer to cases where there is no linear form of the Bloch equations as involving 'multiply-connected states' and cases where the equations are linear have involving 'singly-connected states'. For singly-connected states, we can define a single net detuning between any pair of states, not matter how many intermediate single-photon transitions are required to connect them. The situation thus looks exactly like a pair of levels connected by only a single laser with some net detuning and Rabi frequency. We can therefore use the results for a single laser to describe singly-connected states. See figure 2.1 for two examples of systems with multiply-connected states.

2.2 Energy eigenstates of an arbitrary ion

Let us consider how to determine the energies and energy eigenstates of an arbitrary ion in a static $\mathbf{B}$ field. The Hamiltonian has several contributing factors:

$$\mathcal{H} = \mathcal{H}_{\text{Gross}} + \mathcal{H}_{\text{Fine}} + \mathcal{H}_{\text{Hyperfine}} + \mathcal{H}_{\mathbf{B}} \tag{2.45}$$

namely, the contributions from the gross structure, the fine structure, the hyperfine structure and the interaction of the ion with the $\mathbf{B}$ field. We will take the level structure as given. We are interested in the energies of the hyperfine states relative to the level energies, and thus the Hamiltonian we will consider does not need to involve the gross or fine structure terms:

$$\mathcal{H}_V = \mathcal{H}_{\text{Hyperfine}} + \mathcal{H}_{\mathbf{B}} \tag{2.46}$$

We consider the hyperfine term first. This describes the energy arising from the interaction with the multipole moments of the nucleus. For this treatment, we will consider only the magnetic dipole and the electric quadrupole terms with hyperfine structure constants A_J and B_J respectively [Woo83]:

$$\mathcal{H}_{\text{Hyperfine}} = A_J \mathbf{I}\cdot\mathbf{J} + B_J \frac{3\left(\mathbf{I}\cdot\mathbf{J}\right)^2 + \frac{3}{2}\mathbf{I}\cdot\mathbf{J} - I(I+1)J(J+1)}{2IJ(2I-1)(2J-1)} \tag{2.47}$$

The Hamiltonian describing the interaction with the $\mathbf{B}$ field has a simpler structure. Both the electronic and the nuclear dipole moments $\boldsymbol{\mu}_{\mathbf{J}}$ and $\boldsymbol{\mu}_{\mathbf{I}}$ interact with the field:

$$\mathcal{H}_{\mathbf{B}} = -\boldsymbol{\mu}_{\mathbf{J}}\cdot\mathbf{B_0} - \boldsymbol{\mu}_{\mathbf{I}}\cdot\mathbf{B_0} \tag{2.48}$$

We write the terms using the spin operators:

$$-\boldsymbol{\mu}_{\mathbf{J}}\cdot\mathbf{B_0} = +g_J\mu_B\mathbf{J}\cdot\mathbf{B_0} \tag{2.49}$$

$$-\boldsymbol{\mu}_{\mathbf{I}}\cdot\mathbf{B_0} = -g_I\mu_N\mathbf{I}\cdot\mathbf{B_0} \tag{2.50}$$

where g_J and g_I are the electron and nuclear g-factors and μ_N is the nuclear magneton. The latter is related to the Bohr magneton, μ_B, through the relative

masses of the proton, m_p and the electron, m_e:

$$\mu_N = \frac{m_e}{m_p}\mu_B \tag{2.51}$$

There are two natural choices of basis in which to express the Hamiltonian matrix: the low field basis $|F, M_F\rangle$ and the high field basis $|M_I, M_J\rangle$. For the purposes of this derivation, I will use the high field basis. This has the advantage of naturally separating out the nuclear and electronic angular momenta. From the point of view of numerical simulation, this allows a useful working check of the validity of the simulation. As any laser interaction must only affect the electronic angular momentum and not the nuclear, stimulated transitions must only change the value of M_J and not M_I.

2.2.1 Operator method

We already have everything we need to calculate the energy states of the ion. Doing this in a standard treatment would involve the use of Wigner-3j symbols to express the $\mathbf{I} \cdot \mathbf{J}$ terms. However, calculating such symbols involves a significant number of summations involving factorial functions at non-negligible computational cost. For the evaluation of a scalar product, we can use an operator approach that avoids the need for summing factorial functions and is significantly faster. We write:

$$\begin{aligned} \mathbf{I} \cdot \mathbf{J} &= \mathbf{I}_x\mathbf{J}_x + \mathbf{I}_y\mathbf{J}_y + \mathbf{I}_z\mathbf{J}_z \\ &= \mathbf{I}_z\mathbf{J}_z + \frac{1}{2}\left(\mathbf{I}^+\mathbf{J}^- + \mathbf{I}^-\mathbf{J}^+\right) \end{aligned} \tag{2.52}$$

where:

$$\mathbf{I}^+ = \mathbf{I}_x + i\mathbf{I}_y \tag{2.53}$$

$$\mathbf{I}^- = \mathbf{I}_x - i\mathbf{I}_y \tag{2.54}$$

Then:

$$(\mathbf{I}\cdot\mathbf{J})^2 = \mathbf{I}_z^2\mathbf{J}_z^2 + \frac{1}{2}\left\{\mathbf{I}_z\mathbf{J}_z, \mathbf{I}^+\mathbf{J}^- + \mathbf{I}^-\mathbf{J}^+\right\} + \frac{1}{4}\big\{(\mathbf{I}^+)^2(\mathbf{J}^-)^2 + (\mathbf{I}^-)^2(\mathbf{J}^+)^2 + \left(I^2 + I - M_I - M_I^2\right)\left(J^2 + J + M_J - M_J^2\right) + \left(I^2 + I + M_I - M_I^2\right)\left(J^2 + J - M_J - M_J^2\right)\big\} \tag{2.55}$$

where M_I and M_J are the quantum numbers of the $\mathbf{I}_z$ and $\mathbf{J}_z$ operators. We can now write the Hamiltonian solely in terms of the angular momentum projections on the quantisation axis and the raising and lowering operators. These operators have eigenvalues:

$$\mathbf{I}_z|I, M_I\rangle = M_I|I, M_I\rangle \tag{2.56}$$

$$\mathbf{I}^\pm|I, M_I\rangle = \sqrt{I(I+1) - M_I(M_I \pm 1)}|I, M_I \pm 1\rangle \tag{2.57}$$

We can use these to calculate the matrix elements of the Hamiltonian without using any summations or factorial functions; each matrix element is calculated using only a single product. The method also demonstrates immediately which matrix elements in this Hamiltonian are non-zero; the Hamiltonian only connects high-field states that differ in either angular momentum eigenvalue by no more than two.

An individual calculation of the energy states of an ion using both methods in my simulations takes on average:

- T(Wigner-3j method) = 80ms
- T(Operator method) = 6ms

These timings were determined from 20 runs of a routine that calls the energy state determination function 100 times. We deduce that the operator approach is about an order of magnitude faster than using Wigner-3j symbols.

2.3 Atomic Matrices

We next describe how the four matrices ($R_{mn}, A_{mn}, \Delta_{mn}$ and Γ_L) required for the Bloch equations are calculated.

2.3.1 Stimulated emission rate

The matrix elements for stimulated processes R_{mn} describe the probability of stimulated transitions taking place between two states. In the optical Bloch equations, they control the rate at which stimulated processes drive population between these states. I will therefore refer to R_{mn} and the equivalent A_{mn} for spontaneous processes as stimulated and spontaneous 'rates' of population transfer between states $|m\rangle$ and $|n\rangle$. In equilibrium, the sum of the stimulated and spontaneous rates between all pairs of states will be zero[1].

The stimulated transition rate R_{mn} is defined in (2.13). States at intermediate field only have a single angular momentum quantum number that can be defined. This is the total angular momentum projection M. Clearly this does not give a complete description of the angular momentum of the ion! In order to fully describe the angular momenta of the states, we must adopt a complete angular momentum basis. The suitable bases are those describing systems experiencing a weak or strong $\mathbf{B}$ field, $|J, I, F, M_F\rangle$ and $|J, I, M_J, M_I\rangle$ respectively. Note that we can apply the Wigner-Eckart theorem to the angular momentum basis states but *not* to any intermediate-field states. As with the analysis of the energy states of an ion described in the previous section, we work with the high-field basis states. The intermediate-field states can be written as a sum of these:

$$|m\rangle = \sum_k a_{mk} |J_k, I_k, M_{J_k}, M_{I_k}\rangle \tag{2.58}$$

[1]The R_{mn} directly involve the coherence terms, so the actual rate of population transfer is only of order R_{mn}. A complete expression for the rate of stimulated processes is given by the Rabi frequency in the dressed state; this is a function of all four atomic matrices. See section 4.5.1.

The a_{mk} are given by the calculation of the energy states in the previous section. The stimulated transition rate is the sum of the stimulated rates between the high field basis states:

$$R_{mn} = \sum_k \sum_l a_{mk} a_{ln} S_{kl} \tag{2.59}$$

where we have defined the matrix element:

$$S_{kl} = \frac{1}{2\hbar} \langle J_k, I_k, M_{J_k}, M_{I_k} | \mathbf{d} \cdot \mathbf{E_0} | J_l, I_l, M_{J_l}, M_{I_l} \rangle \tag{2.60}$$

Photons will, to a good approximation, only couple to the electronic angular momentum. As a result, only the J and M_J quantum numbers will change during a transition. We will therefore refer to the states by only these two quantum numbers without loss of generality:

$$S_{kl} = \frac{1}{2\hbar} \langle J_k, M_{J_k} | \mathbf{d} \cdot \mathbf{E_0} | J_l, M_{J_l} \rangle \tag{2.61}$$

Note that M_J is distinct from the total angular momentum projection M which, in this basis, is given by $M = M_J + M_I$.

Each pair of states that is connected by a stimulated transition will interact with a photon with a specific angular momentum projection q, where $q = 0, \pm 1$. The intensity of the driving laser only modulates the overall strength of the transition; it is the dipole moment itself that couples the two states. The inner product of two tensors is defined as [Sob96]:

$$\mathbf{T}^{(k)} \cdot \mathbf{U}^{(k)} = \sum_r \left\{ (-1)^r \, \mathbf{T}^{(k)}_{-r} \mathbf{U}^{(k)}_r \right\} \tag{2.62}$$

We can therefore remove the electric field from the matrix element:

$$S_{kl} = \frac{1}{2\hbar} \left| \mathbf{E}_q \right| (-1)^q \langle J_k, M_{J_k} | \mathbf{d}_{-q} | J_l, M_{J_l} \rangle \tag{2.63}$$

where $|\mathbf{E}_q|$ is the amplitude of the electric field of the laser with polarization q.

Using the Wigner-Eckart theorem, we can write:

$$S_{kl} = \frac{1}{2\hbar} \left| \mathbf{E}_q \right| (-1)^{-J_l + 1 - M_{J_k} + q} \begin{pmatrix} J_l & 1 & J_k \\ M_{J_l} & -q & -M_{J_k} \end{pmatrix} \langle J_k \| \mathbf{d} \| J_l \rangle \tag{2.64}$$

where $\langle J_k \| \mathbf{d} \| J_l \rangle$ is a reduced matrix element. For convenience, I will refer to this in all that follows as $\|\mathbf{d}_{kl}\|$.

Permuting the columns of the 3j-symbol allows the modification to the sign of the term to depend only on the properties of one of the states:

$$S_{kl} = \frac{1}{2\hbar} |\mathbf{E}_q| (-1)^{J_k - M_{J_k} + q} \begin{pmatrix} J_k & 1 & J_l \\ -M_{J_k} & -q & M_{J_l} \end{pmatrix} \|\mathbf{d}_{kl}\| \tag{2.65}$$

The electric field $\mathbf{E}_q$ can be calculated from the fraction f_q of the polarization q in the driving laser:

$$|\mathbf{E}_q|^2 = \frac{2I}{\epsilon_0 c} f_q \tag{2.66}$$

It will be more convenient for the purposes of comparison to deal with the intensity relative to the saturation intensity. There are several definitions of the saturation intensity. We will define it as follows:

$$I_0 = \frac{8\pi^3 \hbar c \Gamma_k}{\lambda^3} \tag{2.67}$$

where Γ_k is the full-width of the level containing state $|k\rangle$. This definition fixes the stimulated and spontaneous transition rates to be equal for a cycling transition on resonance for one saturation intensity [Szw09]. In terms of this:

$$|\mathbf{E}_q|^2 = \frac{16\pi^3 \hbar \Gamma_k}{3\epsilon_0 \lambda^3} \frac{I}{I_0} \tag{2.68}$$

The reduced matrix element of the dipole operator can be written in terms of the Einstein A-coefficient:

$$\|\mathbf{d_{kl}}\|^2 = \frac{6\epsilon_0 \hbar \lambda^3}{16\pi^2} (2J_k + 1) A(k, l) \tag{2.69}$$

where J_k is the angular momentum of the level containing state $|k\rangle$. This arises by treating an ion experiencing a spontaneous transition as a oscillating electric dipole that emits a photon. For good treatments of the radiation emitted by an electric dipole, see [Cor05, Orf14, Dit52]. Using (2.69) and (2.68), the matrix

elements for the high-field states can be written as:

$$S_{kl} = (-1)^{J_k - M_{J_k} + q} \begin{pmatrix} J_k & 1 & J_l \\ -M_{J_k} & -q & M_{J_l} \end{pmatrix} \sqrt{\frac{\pi \Gamma_k}{2} (2J_k + 1) \frac{I}{I_0} A(k,l)} \tag{2.70}$$

Therefore, the matrix elements for stimulated transitions between energy eigenstates at an arbitrary applied magnetic field are:

$$R_{mn} = r_{mn} \sqrt{\frac{I}{I_0}} \tag{2.71}$$

where:

$$r_{mn} = \sum_k \sum_l a_{mk} a_{ln} (-1)^{J_k - M_{J_k} + q} \begin{pmatrix} J_k & 1 & J_l \\ -M_{J_k} & -q & M_{J_l} \end{pmatrix} \sqrt{\frac{\pi \Gamma_k}{2} (2J_k + 1) A(k,l)} \tag{2.72}$$

R_{mn} and r_{nm} are necessarily real.

2.3.2 Spontaneous emission rate

Now that we have the stimulated transition rates, it is a simple matter to determine the spontaneous transition rates between the hyperfine levels. The spontaneous rate must be proportional to the square of the electric dipole matrix element between the two hyperfine states, or equivalently:

$$A_{mn} \propto r_{mn}^2 \tag{2.73}$$

so that[2]:

$$A_{mn} = \frac{r_{mn}^2}{\sum_n (r_{mn}^2)} A(m,n) \tag{2.74}$$

2.3.3 Laser Detunings

In the set of differential equations, the only constraint on the detunings is that the detuning relevant for the transition from state $|m\rangle$ to state $|n\rangle$ through state

[2]Each intermediate-field state is expressed in terms of high-field basis states in the same level. $A(k,l)$ for the high-field and $A(m,n)$ for the intermediate-field are therefore equivalent.

$|p\rangle$ is:

$$\Delta_{mn} = \Delta_{mp} + \Delta_{pn} \tag{2.75}$$

By extension, the detuning between states separated by more than two lasers is given by sum of the detunings of all the lasers making the path between the two states. For two states separated by the chain of lasers:

$$|m\rangle \xrightarrow{\gamma_1} |a\rangle \xrightarrow{\gamma_2} |b\rangle \xrightarrow{\gamma_3} \cdots \xrightarrow{\gamma_N} |n\rangle \tag{2.76}$$

the detuning is given by:

$$\Delta_{mn} = \Delta_{ma} + \Delta_{ab} + \Delta_{bc} + \cdots + \Delta_{(n-1)n} \tag{2.77}$$

Note that only the energies of states $|m\rangle$ and $|n\rangle$ contribute to the detuning; the energies of all the other states cancel in adjacent terms.

Otherwise, the detunings must simply be self-consistent. In all the following, we define a detuning as being positive if the laser driving a transition is blue-detuned relative to it and negative if it is red-detuned. The energies of the states of the ion calculated previously, along with the detuning of each laser relative the centres of mass of the levels involved, give all the information necessary to calculate the detunings for every pair of states.

2.3.4 Laser Linewidths

For pairs of states connected by a single laser, the linewidths are defined in a straightforward way; Γ_L is the linewidth of the laser connecting the states. States connected by multiple lasers need to take account of the linewidths of all the lasers connecting the pair. Each laser in the chain will add its own spread of possible frequencies to the list. The chain of lasers can be modelled as a single laser with a lineshape that is the convolution of all the lineshapes of the contributing lasers. If we assume that the lasers are Lorentzian to a good approximation, the resultant lineshape will also be Lorentzian with a full-width equal to the sum of

the full-widths of all the contributing lasers. So, for the chain :

$$|m\rangle \xrightarrow{\gamma_1} |a\rangle \xrightarrow{\gamma_2} |b\rangle \xrightarrow{\gamma_3} \cdots \xrightarrow{\gamma_N} |n\rangle \tag{2.78}$$

the effective linewidth is:

$$\Gamma_{L,mn} = \Gamma_{L,1} + \Gamma_{L,2} + \Gamma_{L,3} + \cdots + \Gamma_{L,N} \tag{2.79}$$

2.3.5 Laser Polarizations

In order to fully specify the spontaneous and stimulated rates, we need to agree on a convention for describing the polarization of the light. Up to this point, we have only been concerned with how the light changes the angular momentum of the system. We need to relate this to the polarizations of a physical laser, especially the circular polarizations for which there might be some ambiguity. We will use the convention in [HW05] which defines the polarizations of light according to the variations in the total angular momentum projection M:

	Lower-Energy State	Higher-Energy State	q
σ^+	M	M+1	+1
π	M	M	0
σ^-	M	M-1	-1

3

Numerical Methods

I'm busy with numbers.

John Wainwright
Odd John by Olaf Stapledon

The Bloch equations are a system of coupled differential equations. The size of the system scales with the fourth power of the number of atomic states, so the amount of information that needs to be stored and manipulated rises rapidly with the size of the system. In this chapter, approaches to storing and manipulating this information are discussed. All routines are designed to operate efficiently on a standard laptop computer; no advanced computational resources are required. Indeed, it is part of the aim of this work to make that possible. Measurements of the memory requirements of these routines and the speed of calculations were made on a 2.53GHz MacBook Pro (Late 2008) with 4GB RAM.

We consider two forms of solution. For fluorescence information, we need to solve the steady state of the system. In particular, we use the steady state to examine how the fluorescence varies as a function of the laser frequency and other parameters. Fluorescence scans can reveal a large amount of detail about the experimental situation, and reproducing this detail demonstrates that these routines accurately describe the behaviour of an ion. The scans themselves are also useful for the analysis of Doppler cooling.

Secondly, we will be interested in solving the time-dependent problem. This is important when looking at optical pumping problems for qubit initialisation and readout, and for forming more refined models of Doppler cooling. Indeed, it turns out that the steady-state approach fails to describe Doppler cooling for many situations in our experiments on trapped ions.

In section 3.1, we describe ways in which the Bloch equations can be encoded efficiently in terms of memory and calculation time. In section 3.2, we describe the calculation of the steady state, along with approaches for efficiently calculating this for very large systems. In section 3.3, we describe how to quickly generate fluorescence scans. This will be important when fitting processes are employed later in this thesis. In section 3.4, we describe methods for solving the time-dependent problem. Three different approaches are described for integrating over different timescales. In appendix A, we describe numerical methods required for automatically dealing with an arbitrary set of states and lasers. Parameters used to produce the figures throughout this thesis are listed in appendix H.

3.1 The H matrix

The Bloch equations contain a very large amount of information, even for relatively small systems. It is important to encode these equations in an efficient manner, especially when they may be recalculated hundreds or thousands of times during a fitting or optimisation process. In the following, I define the conventions I use to encode the Bloch equations and the notation which will be used throughout this thesis. I then briefly describe the advantage of a sparse storage method.

3.1.1 Data structures

There are four input matrices to be handled: the stimulated matrix R_{jk}, the spontaneous matrix A_{jk}, the detuning matrix Δ_{jk} and the linewidth matrix Γ_L. Each of the four is stored as a square matrix. Each element of the matrix refers to the value of the quantity for a transition from the state $|j\rangle$ to the state $|k\rangle$. State $|j\rangle$ is defined by the first coordinate and state $|k\rangle$ by the second coordinate. In other words, R_{jk} refers to element in the jth row and kth column of the R matrix and refers to the stimulated transition $|j\rangle \rightarrow |k\rangle$.

The coherences are defined by two real parameters. Let x be the real part of the coherence and y be its imaginary part. By definition, the density matrix of any system must be Hermitian and therefore equal to its conjugate transpose:

$$\rho = \rho^{\dagger} \tag{3.1}$$

This symmetry means that complete information about the system is available from only half of the coherence terms. If we know the coherence linking state $|m\rangle$ to state $|n\rangle$, then we also know the coherence linking $|n\rangle$ to $|m\rangle$:

$$x_{mn} = x_{nm} \tag{3.2}$$

$$y_{mn} = -y_{nm} \tag{3.3}$$

The coherence information is stored in the column vector $\mathbf{X}$. This is a concatenation of vectors $\mathbf{x}$ and $\mathbf{y}$ which store the real and imaginary coherence terms respectively:

$$\mathbf{X} = \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} \tag{3.4}$$

$\mathbf{x}$ and $\mathbf{y}$ only store the minimum number of coherences needed for full knowledge of the system. In this case, for coherence terms x_{mn} or y_{mn}, only those terms with $m \geqslant n$ are stored. The elements are stored so that $\mathbf{x}_k = x_{mn}$ with m and n increasing with increasing k. n cycles from 1 to m before each incremental increase in m. Each population term x_{jj} thus occurs when k is a triangular

number. Using this information, we can convert between the indices n and m describing the actual coherence and the index k used for storage:

$$m = \left\lceil \frac{-1+\sqrt{1+8k}}{2} \right\rceil \tag{3.5}$$

$$n = k - \frac{m\,(m-1)}{2} \tag{3.6}$$

and the reverse:

$$k = \frac{m\,(m-1)}{2} + n \tag{3.7}$$

The Bloch equations are stored in a similar way. The equations are encoded in the matrix $\mathbf{H_0}$. The subscript zero is a preparation for the fact that we will modify this matrix in some of the methods of solution to be discussed. The differential equations can then be written as the matrix equation:

$$\dot{\mathbf{X}} = \mathbf{H_0}\mathbf{X} \tag{3.8}$$

$\mathbf{H_0}$ is therefore a modified Liouvillian. This is the primary equation which we will be solving in many different ways when applying the routines described in the rest of this chapter.

3.1.2 Creating H efficiently

From the form of the Bloch equations, it is possible to infer that a large proportion of the elements of $\mathbf{H_0}$ will be zero. This is simply because the matrix describes connections not only between every pair of states in the system, but also between every pair of coherences. Only relatively few states and coherences out of every possible pair will be involved in a stimulated or spontaneous process. It therefore seems natural to encode the matrix using a sparse storage method.

Using sparse storage speeds up the calculation significantly. The gain in speed increases as the number of states in the system increases; see figure 3.1. This speed increase is important for applications where $\mathbf{H_0}$ is calculated a large

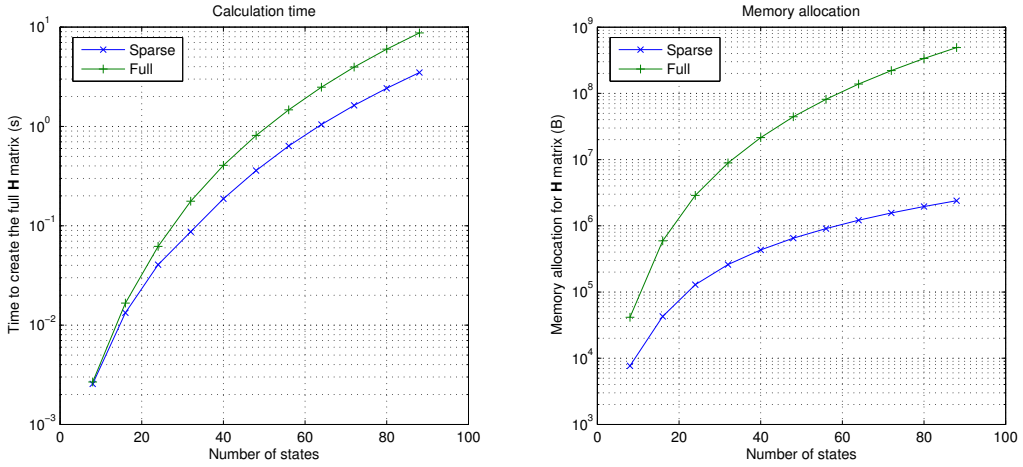

Figure 3.1: The gains of using a sparse storage method. The lefthand plot shows the time taken to construct the **H** matrix whilst the righthand plot shows the memory required to store the matrix. The blue lines show a calculation using a sparse storage method whilst the green show a calculation using a full storage method. The sparse method is always faster, and becomes faster relative to the full calculation for larger numbers of states. It also always requires significantly less memory allocation, and the relative gains likewise increase with increasing numbers of states.

number of times. In addition, such a storage method requires significantly less memory. The number of coherences in a system increases as the square of the number of states, so the **H** matrix can become large very quickly. If n is the number of states, the approximate amount of memory to store the entire matrix is:

$$M_H = 8n^2(n+1)^2 \,\text{bytes} \tag{3.9}$$

The memory requirements therefore reach 1GB for more than approximately 100 states, meaning that computation becomes memory-intensive even for modestly complex systems. By contrast, a sparse system has memory requirements that are smaller by several orders of magnitude, with the result that the matrix can be stored on a standard machine even for very large systems; see figure 3.1.

3.1.3 Calculations using the H matrix

We can solve any dynamic problem we wish by solving the differential equation (3.8). There are two situations to consider: the steady state and the full

dynamics. The steady state is a good approximation when collecting fluorescence with a fixed laser set-up, as long as the motion of the ion is small and slow enough not to introduce significant time-dependence. For fitting and optimization purposes, we will want to find this steady state as quickly as possible. For applications that involve the full dynamics, such as examining pulse sequences and looking at the effects of ion motion, we will need to look at the time evolution of the system directly and therefore solve the differential equation in full. The steady-state case is the simpler, so we consider this first.

3.2 Steady-state solutions

In order to solve for the steady state, we simply need to solve the linear system given by setting the left-hand side of equation (3.8) to zero. This equation can be solved using several different methods. We will consider the standard method of converting the equation to an inhomogenous form in the following section. First, we will examine solving the eigensystem directly.

3.2.1 Solving the Eigensystem

The equation we need to solve is:

$$\mathbf{H_0 X} = \mathbf{0} \tag{3.10}$$

As the $\mathbf{H}$ matrix must necessarily have some non-zero elements, this is an eigenvalue equation. In order to solve for $\mathbf{X}$, we can calculate the eigenvector of $\mathbf{H}$ corresponding to the 0 eigenvalue. The eigenvector will have some arbitrary scaling associated with it. We need it to be scaled so that the populations sum to one:

$$\sum_j x_{jj} = 1 \tag{3.11}$$

This can be a fast method for solving the system. However, there are problems with scalability and automation. For large, sparse matrices, the most efficient

way to solve for a particular eigenvalue is to use an iterative method. The one used in my routines is the Arnoldi method provided by MATLAB [Sor91]. This method uses an initial estimate of the magnitude of the eigenvalue of a required eigenvector as an input to the iterative process. If this initial value is too close to the actual eigenvalue, the iteration may not necessarily converge. This value needs to be chosen so that it is large enough to allow convergence but small enough so that the zero eigenvalue is correctly identified. In addition, the $\mathbf{H_0}$ matrix often contains values spanning at least five orders of magnitude. With this range of values, the routine will occasionally be required to invert matrices that would, but for the difference between several large numbers, be singular. These problems can be avoided by careful choice of the input parameter. However, the value will depend on the size of the system and the inputs to the Bloch equations. The optimal value must be decided manually for each case, meaning that a trial calculation must occur whenever each new problem is tackled.

As a result, my routines use as default a different method to solve the steady state problem. This method is described next. It is faster and more stable; see figure 3.2. Initialisation does not require any external parameter, and the method of solving the problem does not depend on the size of the system.

3.2.2 Matrix inversion

The Bloch equations are necessarily an overdetermined set of linear equations. There is one additional constraint on the equations which is implicit in their construction: namely, the population terms must all sum to 1. For example, if we solve every population equation bar one, the solution to the final population equation is already known through the conservation of probability. As a result, one of the probability equations is unnecessary; all knowledge about the system can be gained entirely from the rest of the system. We therefore do not alter the system of equations if we replace one of the population equations entirely

with some other equation that is consistent with the dynamics of the system. Therefore, we replace one of the population equations with the explicit equation giving that the populations sum to one as per equation (3.11), converting the matrix $\mathbf{H}$ to the matrix $\mathbf{H}'$. The homogeneous equation (3.10) is thus converted to an inhomogeneous equation $\mathbf{H}'\mathbf{X} = \mathbf{V}$ where $V_j = \delta_{ij}$ and the ith equation has been replaced. Instead of dealing with an eigenvalue problem, we are now dealing with a problem of matrix inversion. However, there is one more conversion that is necessary in order to ensure that this problem can be solved. As currently expressed, $\mathbf{H}'$ is singular. The problem is due to the imaginary parts of the population terms, which do not affect the evolution but have still been maintained in this formulation of the Bloch equations:

$$\dot{y}_{jj} = 0 \tag{3.12}$$

The storage method I employ in these routines stores the population terms in amongst the coherence terms; $\mathbf{x}$ and $\mathbf{y}$ store the real and imaginary parts of the upper triangular parts of the density matrix in sequence. This includes the imaginary population terms. We could remove these terms, but this would mean that the numbering system for the populations and coherences would differ between $\mathbf{x}$ and $\mathbf{y}$ and a different function would have to be used to extract particular terms from the two matrices[1]. Instead, we can retain the imaginary population terms and keep the symmetry between $\mathbf{x}$ and $\mathbf{y}$ by replacing the differential equations for the imaginary populations with some dummy equations. As the value of each of these elements must *always* be zero for any system and at any time, replacing the zero with any zero-valued function will not change the dynamics. For example, by replacing (3.12) with:

$$\dot{y}_{jj} = y_{jj} \tag{3.13}$$

[1]Including the $\dot{y}_{jj}$ terms in these routines has a negligible effect on calculation times and memory usage. Doing so makes the structure of the $\mathbf{H}$ matrix and solution vectors more intuitive and faster to interpret.

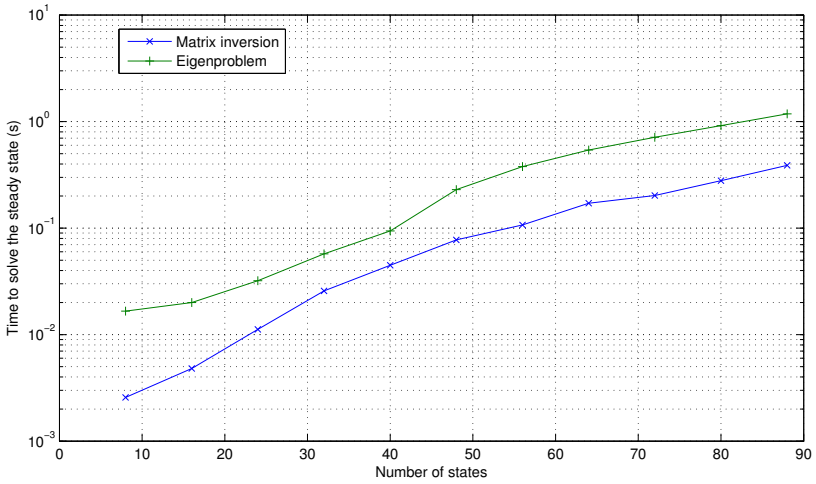

Figure 3.2: The gains in using matrix inversion over solving the eigenproblem. The blue line shows the calculation using matrix inversion whilst the green line shows the calculation that solves the eigenproblem. The method of matrix inversion is consistently several times faster, varying from a factor 2 speed increase at worst to over a factor 6 increase at best. The most significant gains occur for the smallest systems.

we make the system of equations non-singular. We make this conversion, converting $\mathbf{H}'$ to the matrix $\mathbf{H}$:

$$\mathbf{HX} = \mathbf{V} \tag{3.14}$$

This matrix is invertible.

3.2.2.1 Inverting large matrices

For small systems, we can use standard algorithms to invert the $\mathbf{H}$ matrix. However, problems arise when dealing with very large systems. In particular, the memory required to invert a large sparce matrix scales as the number of non-zero elements in the matrix. For a standard atomic system, the number of coherences will scale as the square of the number of states. The scaling of the number of non-zero elements, and therefore the memory requirements, will be between quadratic and quartic. Systems of moderate complexity can be solved on a standard machine, but systems of over a hundred states cannot because they require too much memory. For example, inverting a matrix involving five fine-structure

levels, 144 states and light containing all three polarization components requires 7.0GB of virtual memory.

However, there are other methods of matrix inversion that allow for this system to be solved on machines that do not have the memory capacity to directly invert such large matrices. One approach is to solve the matrix inversion problem iteratively by taking some trial solution and refining it until an approximate solution is found. Such methods can be slower than direct matrix inversion, but they are much less memory intensive.

Not all iterative methods are suitable for the problem under consideration. Several routines have conversion rates for our system that are far slower than in other applications. This may be due to the large range of values found in the **H** matrix. However, there are several methods that do solve the linear system to significant accuracy on useful timescales. I use three different iterative methods in my routines: the quasi-minimal residual method (qmr) [FN91]; the biconjugate gradients stabilised (l) method (bicgstabl) [SF93]; and a combined routine that uses a short initial run of qmr followed by a longer run of bicgstabl. qmr and bicgstabl are routines available in a standard MATLAB distribution. Figure 3.3 presents a demonstration of the effectiveness of these routines for an approximately optimal set of operating parameters in the three and five lowest-lying levels of the $^{43}Ca^{+}$ ion, consisting of 64 and 144 states respectively. The combined routine makes use of the fast initial convergence of qmr and then refines the solution using bicgstabl, a routine which outperforms qmr in solving the system but initially converges more slowly. By using iterative routines, the memory required to solve the linear system is reduced by a factor of 26 in one case and a factor of 2600 in the other. Calculation times are much larger for smaller systems, but for systems with more than 100 states, the calculation times for acceptable errors become comparable to those of direct inversion.

However, the performance of these routines depends on the values in the

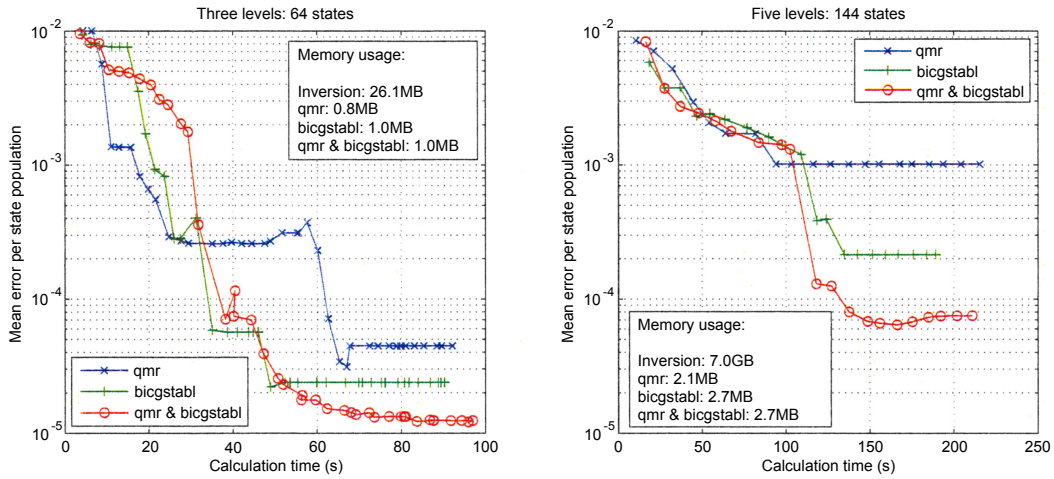

Figure 3.3: The tradeoff in using iterative methods over using matrix inversion. Three methods are shown. The blue curve shows the qmr routine, the green the bicgstab(l) routine and the red a combined method that first uses qmr to obtain fast convergence and then refines using bicgstab(l). The plots show the mean error per state population after the populations have been rescaled so that they sum to one. This rescaling means that the curves using only a single iterative routine do not necessarily decrease monotonically. The lefthand plot shows results for a three-level system in $^{43}Ca^{+}$ including π polarization and the righthand plot shows results for the much larger five-level system. For comparison the mean calculation time for matrix inversion is 0.4s for the three-level case. The five level-system cannot be inverted on the machine used to calculate these timings, but it can be inverted in 360s on a more powerful machine.

solution. The individual routines perform slower and in a less stable manner when the spread of population over the states is less even. Situations where one or a few states are heavily populated are not modelled well. For practical purposes, the routines work most effectively when lasers are tuned close to resonance. See figure 3.4 for a demonstration of how the accuracy of these routines changes for a three-level system as one of the lasers is scanned over resonance. Note that the combined method produces more a more stable output over the entire frequency range, with fewer and smaller high-error spikes. In general, qmr is the best choice for fast convergence when crude errors are acceptable, whereas the combined method is most suitable for situations where high accuracy is required and when a large spread of population configurations are being examined. Whilst qmr can produce extremely small errors under the right conditions, this cannot be guaranteed; even around the optimal regions, high-error

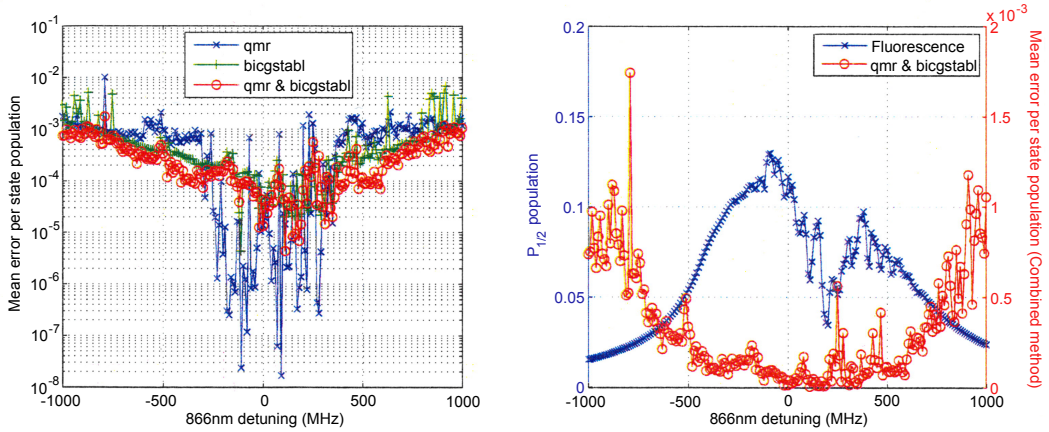

Figure 3.4: The variation in the effectiveness of the three routines for a three-level system containing 64 states as the frequency of one of the lasers is scanned. The routines are set to calculate for the same length of time. All three routines output the lowest errors around the centre of the scan where the laser is most resonant with the transition it is driving. The combined routine does not produce as low errors as the qmr routine, but it is more stable over the entire frequency range.

spikes persist.

3.3 Scanning the steady state

One particularly important application of the steady-state method of solution is to numerically generate fluorescence scans for an ion. The frequency of one of the lasers is scanned and the variation in the fluorescence is examined, where the fluorescence is proportional to the total population in the upper level of the fluorescence transition. Such scans are crucial for linking these numerical simulations with experimental results, especially in complex systems where the structure of such scans is by no means intuitively obvious.

Scanning requires that the steady state problem is solved a large number of times. Over the course of generating a scan, hundreds or thousands of solutions may be calculated. Significantly more will be required as soon as any kind of fitting process is employed. Calculating scans in as short a time as possible is therefore of key importance. Without methods that significantly reduce the calculation time, many of the results of the more complex fitting and optimisation

routines in this thesis would have been intractable.

3.3.1 Reducing calculation overhead

The first and simplest approach that can be used to reduce the calculation times for scans is to note that, across such a scan, only a single variable is being altered. We can reduce the calculation time by separating out the various contributions to $\mathbf{H}$. For instance, in a frequency scan, we create a matrix $\mathbf{H_N}$ that encodes the initial conditions and a matrix $\mathbf{H_C}$ that encodes the change to the equations produced by varying the frequency of the laser being scanned. Owing to the linearity in frequency as shown in equations (2.41) and (2.42), we have:

$$\mathbf{H}(\omega) = \mathbf{H_N} + \omega\mathbf{H_C} \tag{3.15}$$

where ω is the change in frequency of the laser being scanned relative to the frequency described by the initial conditions. Hence we can produce a whole series of $\mathbf{H}$ matrices describing the entire scan without ever recalculating any matrices, sparing the cost of calculating a new $\mathbf{H}$ at every step from scratch. Similar approaches are possible for intensity and linewidth variations to speed up optimization problems. The only variable for which we cannot use a method like this is variation in the $\mathbf{B}$ field; this changes the energies of the states and the entire structure of $\mathbf{H}$, so both need to be calculated at every step.

3.3.2 LU preconditioning

We can use a similar approach to reduce the cost of matrix inversion at every point. $\mathbf{H_C}$ is necessarily a significantly more sparse matrix than $\mathbf{H_N}$. $\mathbf{H}$ therefore has a largely similar structure for any value of ω. Unfortunately, the actual values of the terms in $\mathbf{H}$ will vary significantly with ω; we cannot make any general assumption that ω is small or large, for example. However, the similarity in structure means that we can make use of the output of one inversion to aid subsequent inversions.

Several methods can be employed depending on the exact method being used for inversion. In these routines, matrices are inverted using LU factorisation with full pivoting. For a review of this factorisation method, see [PFTV92]. LU factorisation is essentially equivalent to Gaussian elimination: the matrix is written as the product of a lower-triangular and upper-triangular matrix, both of which allow the system to be solved by using substitutions without doing any actual inversion. At least one permutation matrix is required for the calculation to be stable for the majority of inputs [TI97], and using two permutation matrices converts the matrix into a form which is much simpler to factorize; see figure 3.5. In other words, the $\mathbf{H}$ matrix is factorized as:

$$\mathbf{PHQ} = \mathbf{LU} \tag{3.16}$$

where $\mathbf{P}$ and $\mathbf{Q}$ are the permutation matrices, $\mathbf{L}$ is a lower-diagonal matrix and $\mathbf{U}$ an upper-diagonal matrix. In addition, we use the convention that $\mathbf{L}$ is a unit matrix. The system can then be solved be calculating:

$$\mathbf{X} = \mathbf{Q}\left(\mathbf{PHQ}\right)^{-1}\mathbf{PV} = \mathbf{QU}^{-1}\mathbf{L}^{-1}\mathbf{PV} \tag{3.17}$$

where the inversion of the triangular matrices is trivial.

A large portion of the calculation involves producing the permutation matrices; the routine permutes $\mathbf{H}$ into a more optimal form so that the factorisation is as fast as possible. We can make use of this when inverting a series of matrices with a very similar structure. When we invert the matrix for the first time, we store the permutation matrices $\mathbf{P_0}$ and $\mathbf{Q_0}$. For all subsequent inversions, we skip the step of calculating permutation matrices and instead use the matrices calculated for the first inversion. As the matrices have very similar structures, the matrix $\mathbf{P_0H}(\omega)\mathbf{Q_0}$ will still be approximately optimal with regards LU factorisation. Therefore, the calculation time for inverting all subsequent points will be significantly reduced.

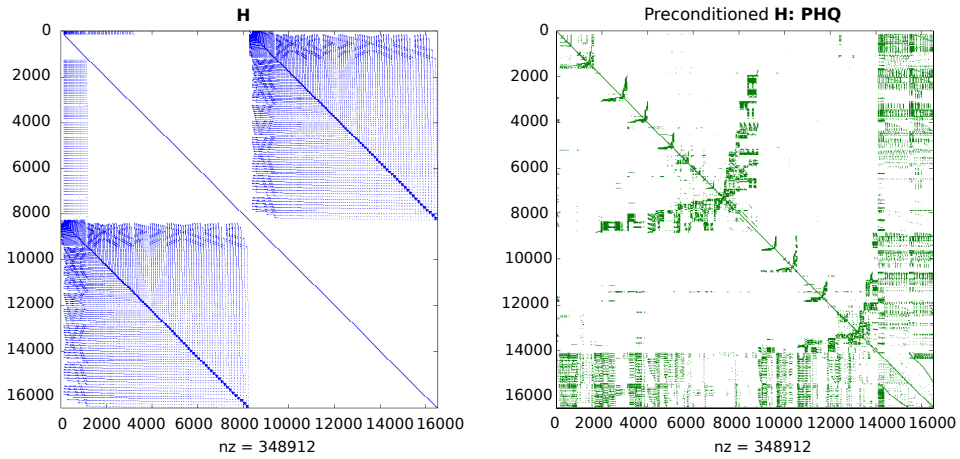

Figure 3.5: Permuting **H**. The lefthand plot shows the sparsity pattern of **H**. The righthand plot shows the sparsity pattern after the preconditioners have been applied. The matrix is now in a bordered, block-diagonal form and is significantly more symmetric. The preconditioned matrix is significantly easier to factorise.

Since this study was conducted, MATLAB has implemented an automatic preconditioning whenever series of similar matrices are inverted. For small systems, the gains of this automatic system are equal to the manual preconditioning used here. However, for large systems, the 'manual' LU preconditioning described here beats the automated preconditioning provided by MATLAB. For a system consisting of 144 states, manual LU preconditioning is faster by a factor of four; see figure 3.6.

3.3.3 Improved preconditioners

In this context, the permutation matrices $\mathbf{P_0}$ and $\mathbf{Q_0}$ are called preconditioners. When it is first formed, **H** is not usually in an optimal form for factorisation. Calculating the preconditioners may take many times longer than the matrix inversion, and there is no guarantee that the preconditioners that are generated will remain useful as **H** varies over the whole range of the scan. If we first partially permute **H** into the form shown in 3.16, the calculation of the preconditioners will occur faster than with the original structure, and the preconditioners produced will be more optimal. If a routine that calculates a preconditioner is given

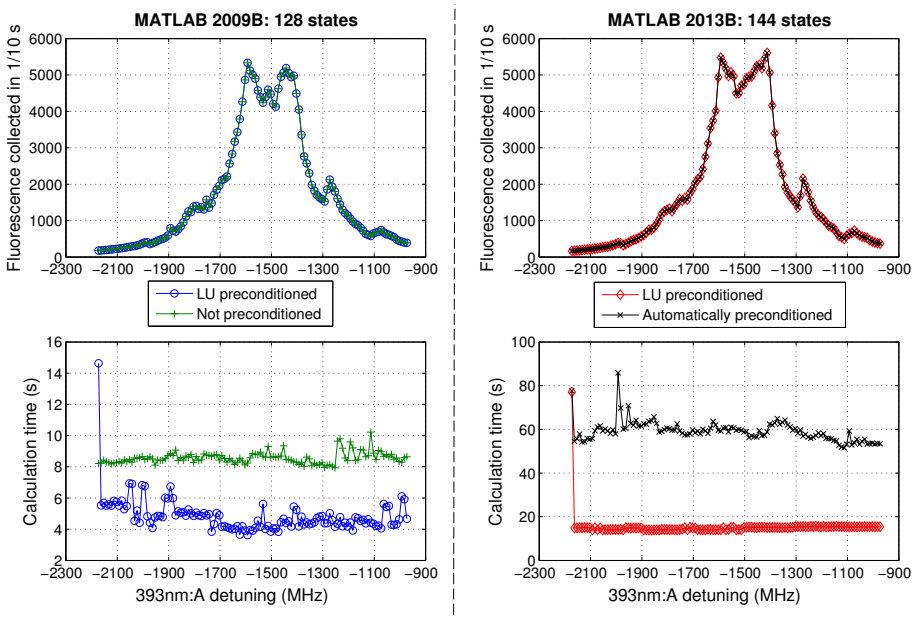

Figure 3.6: The advantages of using LU preconditioning. The lefthand plots show calculations on a version of MATLAB that does not use automatic preconditioning. For smaller systems, the speed increase is a factor of two. Note that the first point, at a detuning of -400MHz, takes longer to calculate due to the creation of the preconditioners. The righthand plots do use automatic preconditioning. The preconditioning calculation takes approximately the same time for both methods, but calculating using the LU preconditioners is significantly faster than the automated preconditioning. The 393nm beam has two frequency components that interact with the system. '393nm:A' refers to the lower-frequency component. See figure 6.1.

a poorly-ordered matrix, it will optimise as best it can but may result in a poor local optimum. On the other hand, if it is given a matrix that is better ordered, the routine will be able to discount poorer local optima and may find a better local optimum.

We use two approaches to calculate this initial permutation. Both of these approaches use the matrix $\mathbf{R}$ to determine a reordering of the states. This is the matrix describing the stimulated transitions between pairs of states. See section 2.3.1 for a treatment of the matrix element R_{mn}. This matrix contains information about the strengths of all transitions in the system that are excited by a laser beam and is a measure of how all the states in the system are 'connected' through stimulated processes. As the matrix is symmetric, we can treat it as an adjacency matrix. $\mathbf{H}$ also contains complete information, but is necessarily not symmetric.

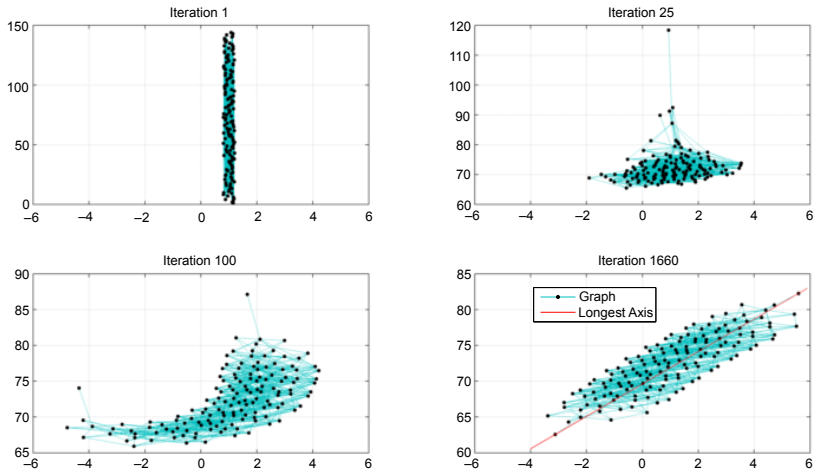

Figure 3.7: A demonstration of the graph-drawing process applied to **R**. The input to the graph-drawing routine is a vertical line of points ordered according to the state number. Each point has a random horizontal displacement. As the routine calculates, the states unfold and reorder themselves along a rectangular structure. The ordering of these points along the longest axis of this structure determines the new ordering of the states.

Any permutation matrix applied to **H** must make some compromise between the structure in the horizontal and vertical dimensions. By applying these approaches directly to **H**, either to the entire matrix or to partitions, calculation times can only be reduced by a small factor; the reordering of the coherences is less optimal than the reordering calculated using **R**.

We explored two approaches. The first approach minimises the bandwidth of **R** using the reverse Cuthill-Mckee algorithm [CM69, Geo71]. It scales as a breadth-first search, so calculates much faster than matrix inversion. The second approach is significantly slower, with calculation times in excess of an order of magnitude slower than matrix inversion. These calculation times range from one to several tens of minutes. This is not a problem, however. The structure of **R** is determined only by the existence of stimulated transitions between states. The magnitude of the transitions, or the energies of the states (so long as the ordering of the states in energy is taken into account), is irrelevant. Therefore, this preconditioner need only be applied once for a given laser configuration. It will then apply for all uses of this configuration, even if the laser intensities or

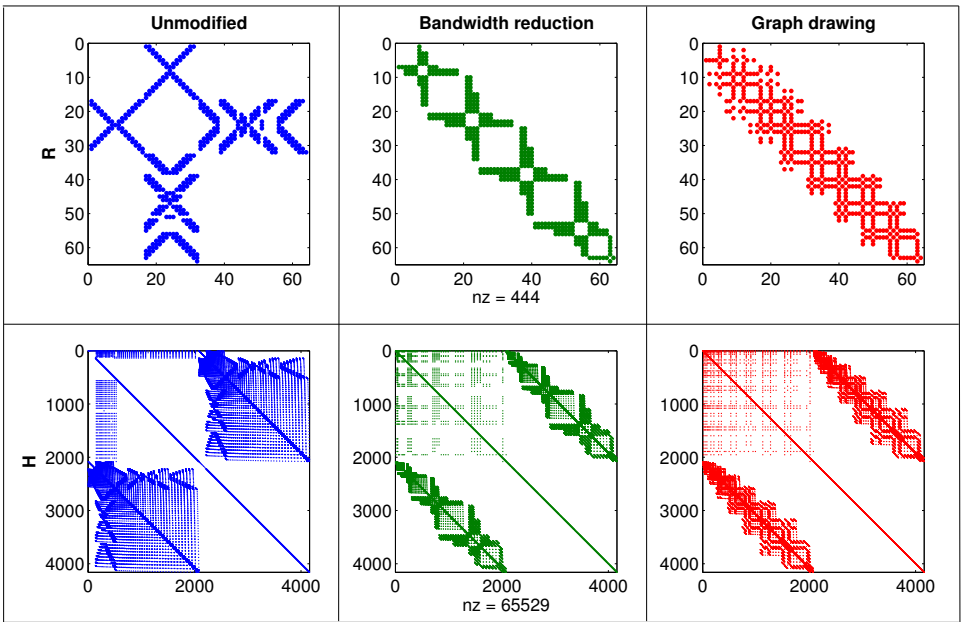

Figure 3.8: A demonstration of the effects of the permutation on $\mathbf{R}$ (upper) and $\mathbf{H}$ (lower). The matrices displayed are, from left to right; unmodified; permuted using the bandwidth-reduction method; and permuted using the graph-drawing method.

the $\mathbf{B}$ field are changed.

This second approach uses graph drawing with the aim of making the symmetries in the structure of $\mathbf{R}$ more prominent. Graph drawing allows for inherent symmetries in a network to display themselves. We use a force-directed graph-drawing method with a variable step-size. Each state is modelled as a particle that feels a Coulomb repulsion from every other particle. Particles that are connected by a stimulated process experience an attractive spring force. The system is then allowed to evolve over time. Energy is damped out of the system to ensure that an equilibrium is eventually reached. The time step is increased (decreased) if the kinetic energy of the system increases (decreases) over a fixed number of steps before damping is taken into account.

The initial input to this routine is arranged so that the points are ordered vertically from top to bottom with a fixed vertical distance between each point. Each point has a small random horizontal displacement. When the graph-drawing routine runs, the graph 'unfolds'. Typically, the states will arrange

themselves along a thin rectangular structure. The permutation is extracted from this by ordering the states according to their coordinate along the longest axis of the rectangle. See figure 3.7.

Figures 3.8 and 3.9 show a demonstration of the preconditioners on $\mathbf{R}$ and $\mathbf{H}$ for these two approaches and the effect on the calculation time for a scan in three-level $^{43}Ca^{+}$. The preconditioners have a broadly similar effect on the structure $\mathbf{R}$, though the graph-drawing process has a slightly larger bandwidth and greater fill-in. The mean improvement in calculation time across the scan for both processes is identical to one part in 500, implying that the preconditioners have significant similarity in the optimal structures they produce. The permutation obtained by graph-drawing leads to smaller variation of calculation time over the course of the scan than the band-width permutation. It rarely slows down the calculation. The band-width permutation, on the other hand, at times slows down the calculation, but it can lead to significant speed increases at certain frequencies. The cause of these speed increases is currently unknown, but it implies that it may be possible to produce a preconditioner that is even more efficient than generated by either of the two methods demonstrated here.

3.3.4 Coherent subsystems

For certain systems of equations, we can reduce calculation times by only calculating a portion of the system. For example, not all of the coherences between every state may be required to give complete information about the populations. An important example of this is when we use lasers that only have σ polarisations. This has the effect of separating the states into two distinct sets of states that are not coherently connected. See figure 3.10 for an example of this in $^{40}Ca^{+}$. This separation is not in itself what we are interested in; the states themselves must all be connected in $\mathbf{H}$ through spontaneous processes. However, the separation also leads to a separation of the coherences into two groups. One group

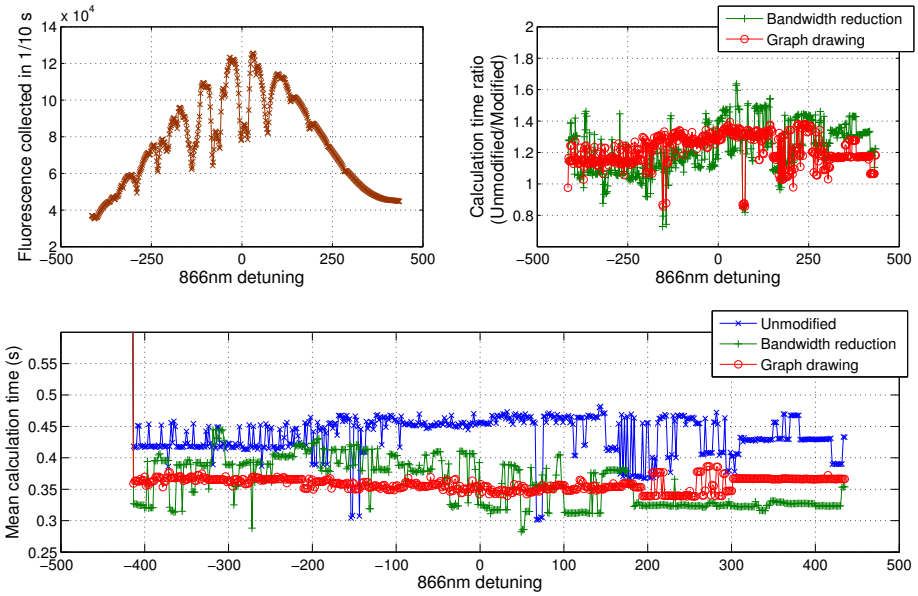

Figure 3.9: The decrease in calculation times when using the two permutation matrices. The scan is an 866nm scan on three-level $^{43}Ca^{+}$, consisting of 64 states. The bandwidth-reduction calculation is shown in green and the graph-drawing calculation in red. The mean increase in speed for both processes is 21.2%. The bandwidth-reduction data implies that it might be possible to perform even better than this.

of coherences links together states within a given set, and the other group links states from both sets. This second group is not meaningful for this calculation; we have already said that the sets of states are not coherently connected, so these coherences can have no bearing on the populations. If we remove all of these coherences from the system, we can calculate using only half of the coherences and still have complete knowledge of the problem. This has the effect of reducing the size of the $\mathbf{H}$ by slightly less than a factor of two in both dimensions (it is not exactly a factor of two because population information is entirely contained in one of the subsets). This leads to a significant reduction in calculation time. For three-level $^{43}Ca^{+}$, calculation times are reduced by over a factor of 2. See figure 3.11.

This effect is most prominent for σ/σ systems, but separation into subsets of coherences occurs whenever some portion of states in the system is not connected coherently to some other portion. This will occur, for example, when

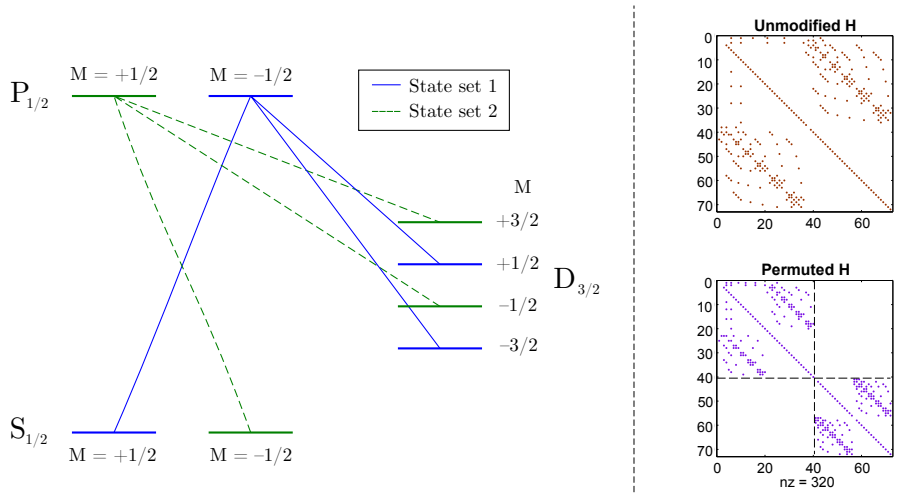

Figure 3.10: A demonstration of the existence of two subsets of coherences in three-level $^{40}Ca^{+}$ illuminated with only σ light. The lefthand plot shows the separation of the states into two sets that do not coherently interact. The terms in $\mathbf{H}$ that would describe such an interaction are still non-zero, but do not couple with any of the populations or non-zero coherences. This allows us to separate out these coherences from the rest of the system using a permutation, as shown in the righthand plots.

looking at optical pumping, or in π/σ systems. The subsystems will not be of equal size in these situations, and the number of coherences that can be neglected will be much smaller than for the σ/σ case. However, a small reduction in calculation times can still be achieved.

We can detect subsets of coherences using a simple approach. If we assume that there are only two subsets of coherences in the system, then the first column of $\mathbf{H}$ contains all the information required. The coherences that coherence 1 connects to form one subset, and all the other states form the second subset. To perform this calculation in a vectorised way, rather than using looping, I use a fast matrix exponential. This approach relies on an analogy with adjacency matrices. The adjacency matrix $\mathbf{G}$ describes elements that are linked by a path of length 1, and its power $\mathbf{G}^n$ describes elements that are linked by a path of length n. Thus, the matrix exponential of a matrix describes elements that are linked by a path of any length as its expansion is a series of powers of the matrix.

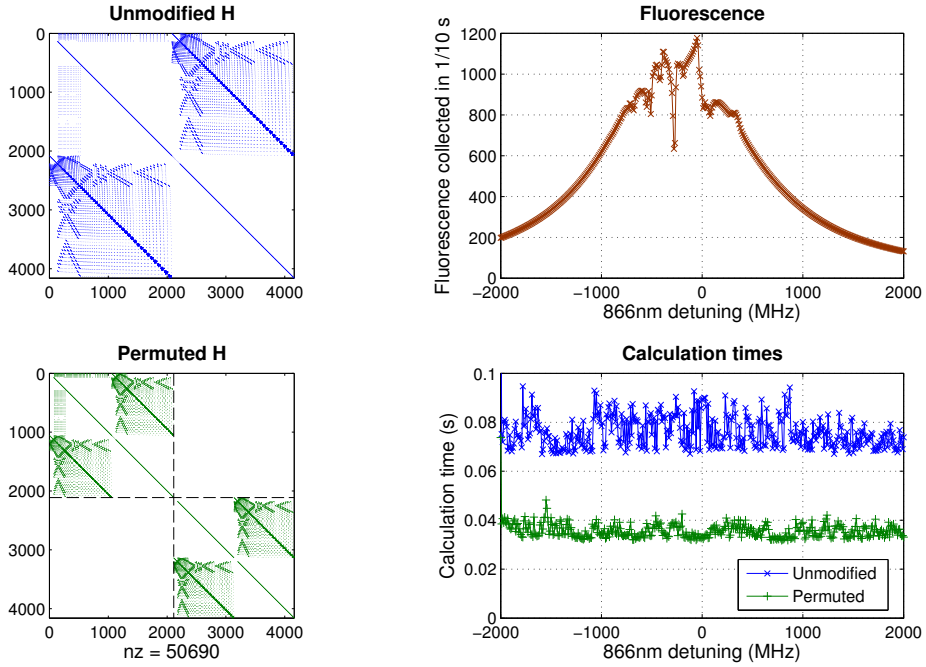

Figure 3.11: The decrease in calculation time produced by using a subset of coherences. The lefthand images show $\mathbf{H}$ before the permutation (top) and after the permutation (lower). The top-left block of the permuted $\mathbf{H}$ is the only portion of the matrix used for calculation. The upper-right image shows the scan and the lower right the decrease in calculation time when using the coherent subset. The mean decrease in calculation time is by a factor of 2.1. The more states there are in a system, the larger this relative decrease in the calculation time will become.

The calculation I use to detect subsets is:

$$\mathbf{Q} = \exp(\mathbf{HC}) \tag{3.18}$$

where $\mathbf{C}$ is a unit vector with a 1 as its first element. The non-zero elements of $\mathbf{Q}$ give the subset of coherences that are connected to the first coherence. This calculation can be performed extremely quickly by using the very efficient EXPMV algorithm written by Awad. H. Al-Mohy and Nicholas J. Higham [CM11]. Another fast algorithm that solves this problem is the depth-first search (see [Deo06, Gou12]). The calculation times for separating the coherences are very small and the search only needs to be performed once for a given set of input parameters, so the differences between these routines are slight for this problem. For example, searching for subsets in a system with 144 states using matrix exponentials takes approximately 0.5s.

3.3.5 Resonance detection

When producing scans, the most important information is contained in the resonance features. Particularly for complex systems such as $^{43}Ca^{+}$, scans may feature many resonances, resulting in a fluorescence structure that is very detailed and hard to predict intuitively. A scanning process with a fixed step-size may miss features that are critical to understanding the behaviour of a system. It is therefore important to use a scanning routine that can adapt its step-size when detail is encountered.

The method I use in these routines is a second-derivative approach. For most cases, it will suffice to only consider the fluorescence curve itself. A set of very few, equally-spaced points is initially taken along the scan, and the derivative S of the curve being considered is approximated at each point. If $R(n)$ is the value of the fluorescence curve at the nth point and $\Delta(n)$ the detuning, the derivative at the nth point is calculated as:

$$\begin{aligned} S(n) = &\frac{\Delta(n+1)-\Delta(n)}{\Delta(n+1)-\Delta(n-1)}\frac{R(n)-R(n-1)}{\Delta(n)-\Delta(n-1)}+ \\ &\frac{\Delta(n)-\Delta(n-1)}{\Delta(n+1)-\Delta(n-1)}\frac{R(n+1)-R(n)}{\Delta(n+1)-\Delta(n)} \end{aligned} \tag{3.19}$$

This is a weighted mean of the derivatives at the midpoints between the nth point and the two neighbouring points. This allows for accurate calculation of the derivatives if the step sizes at adjacent points differ. The relative difference between the derivatives of subsequent points is then calculated:

$$\epsilon(n) = \left| \frac{\{S(n+1)-S(n)\}\{\Delta(n+1)-\Delta(n)\}}{S(n)} \right| \tag{3.20}$$

and compared to an error limit defined by a manually-set tolerance ϵ_u modified by the step size h_0 of the initial detunings Δ_0:

$$\epsilon_0 = \epsilon_u h_0 \tag{3.21}$$

where:

$$h_0 = \Delta_0(n) - \Delta_0(n-1), \qquad n \neq 1 \tag{3.22}$$

If at any point $\epsilon(n) > \epsilon_0$, two additional points are calculated at the midpoints between $\Delta(n-1)$ and $\Delta(n)$ and $\Delta(n)$ and $\Delta(n+1)$, reducing the step-sizes around the point by a factor two.

In (3.20), ϵ depends on the relative magnitude of the difference between the derivatives. This ensures that maxima and minima will have higher values of ϵ and therefore smaller step sizes, ensuring that resonance features are detected and explored fully by the routine. However, using this kind of relative difference has the problem that the size of the difference becomes infinite on points where the derivative is zero. Practically, this means that a routine calculating on relative difference alone will never stop calculating in the vicinity of turning points. ϵ is therefore modified by the step size to avoid this problem as smaller step-sizes are weighted unfavourably. See figure 3.12 for a demonstration of this routine for a relatively simple case in $^{40}Ca^{+}$ with only four resonances along with a scan of $^{43}Ca^{+}$ at 146.1 gauss consisting of very many resonances. In the complex structure, even small resonance features are picked out, whilst areas of the scan with no features are passed over quickly with large step sizes.

This routine can also be used in a more detailed way to examine the resonant structure in a scan. Instead of calculating using the fluorescence curve as an input, we can instead look at all of the state populations simultaneously. At each point, the maximum value of ϵ over all of the curves is used to determine whether the step size should be reduced. This will naturally lead to more points being used compared to calculating using the fluorescence curve, but it ensures that all of the resonant transfer between the states of the system is fully recorded.

By finding the turning points in the population curves, we can identify resonances between pairs of states and link these to the resonant features observed in the fluorescence scan. This allows us to automatically determine which states are involved in producing each feature. This knowledge is important when we want to highlight certain properties of one or more resonances in a scan.

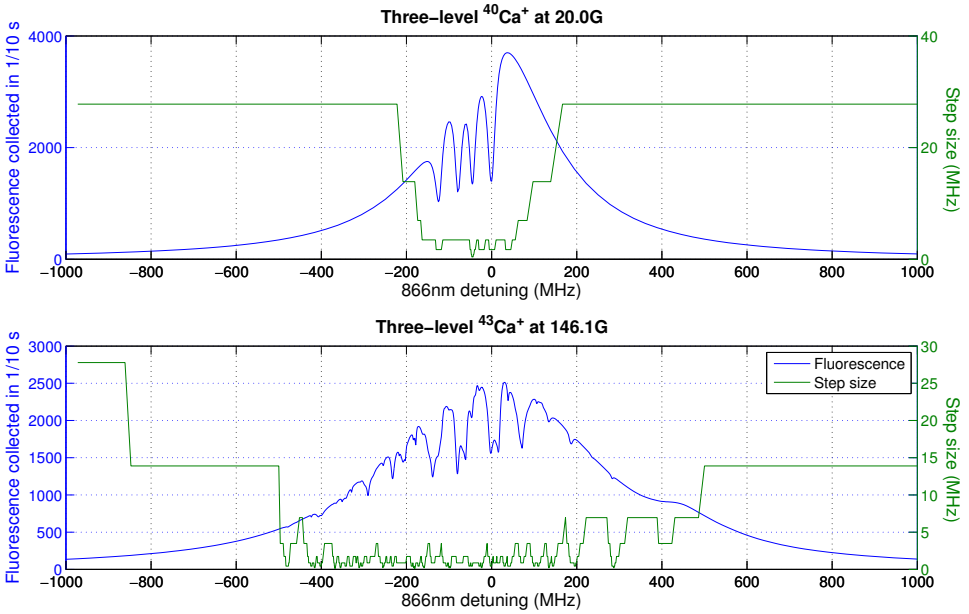

Figure 3.12: A demonstration of the resonance-detection routine. The upper plot shows a scan for three-level $^{40}Ca^{+}$ at 20.0G, consisting of 8 states, and the lower shows a system with the same polarisations in $^{43}Ca^{+}$ at 146.1G, consisting of 64 states. The blue curve shows the fluorescence profile and the green curve the step-sizes. The step-sizes are only reduced in the vicinity of resonance features, including the peak of the overall shape of the resonance. Note that even very small resonance features are visible on the lower profile.

This is particularly the case for selecting and optimising suitable resonances for Doppler cooling.

3.3.6 Global fitting

We will often use these routines to fit experimental fluorescence scans. For many curves that these routines are used to analyse, a significant number of the atomic and laser parameters have been measured to high accuracy. As a result, we can fit the scan using a single least squares fitting process with only a few floated parameters. I will refer to such a fitting process as a 'local fit'. However, there are occasions where we analyse data with ambiguity in several of the parameters, or the data shows a sufficiently complex structure that the output of a fitting process is very sensitive to the initial conditions. In these cases, a local fit produced from the measured parameters is not enough; we need to explore a larger region

of parameter space using a 'global' fitting routine.

The fitting process we use is a fast and computationally simple approach to global optimisation. This approach uses a Monte-Carlo method to explore as much of the parameter space of the ion as possible. The size of the space being explored is limited by the maximum feasible extension of the experimental parameters. We do not need to use the measured parameters to determine these limits; we need only be constrained by what is physically reasonable. For example, we could limit the laser power to be no less than 0 and no greater than the maximum possible power output available from the laser.

Throughout the parameter space, a large number M of trial points is chosen at random. At each point, a fluorescence scan is produced and compared to the data. A small number N of the scans that best agree with the data is chosen and a local fit is performed at each point. These local fits are calculated using the local fitting process used for a regular scan. However, given the large number of scans involved, the entire process is automated and calculated in parallel. The fit that has the greatest agreement with the experimental data is taken as the global optimum. If M and N are chosen well, there may be several local fits that reproduce the optimal parameter set. This routine will only yield information about the nature of the optima if more than one optimum is output; in practice, N should always be greater than one. The larger the size of N, the greater the confidence that the best optimum produced is global, but the longer the calculation time.

This process functions by sampling over parameter space. Those regions that do not show any similar structure to the experimental data are ignored, whilst those that show similar structure are examined more carefully. The resolution of the process obviously improves as M and N increases. As a rule of thumb, I choose M to be equal to the number of points within the frequency scan. Thankfully, for the situations we consider, the variation of the function over most of

parameter space is not extreme, and numbers of points in excess of one thousand are rarely required.

3.4 Time-dependent solutions

We now turn to solution of the dynamical problem:

$$\dot{\mathbf{X}} = \mathbf{H_0 X} \tag{3.23}$$

This is a longer and more challenging computation than the steady state. In particular, we need to solve this equation for a variety of different problems which require information on how $\mathbf{X}$ evolves over vastly differing timescales. In the following, three different methods are proposed to solve this equation efficiently under different conditions. I will use the terms 'calculation time' to refer to the amount of the time a routine takes to run and 'evolution time' to refer to the time parameter in the differential equation.

3.4.1 Challenges

The $\mathbf{H}$ matrix contains a large number of terms that describe all of the possible dynamics of an atomic system. Necessarily, some of the coherent connections between pairs of states will be stronger than others. In general, the elements in $\mathbf{H}$ describing these connections will show a very wide range of values. Differences of four or five orders of magnitude are common. The differential equations therefore encode oscillations between various states with vastly differing frequencies. Furthermore, the smaller terms are vital to the dynamics and cannot be neglected. This contrives to make the Bloch equations a stiff set of equations. The wide range of timescales can lead to instabilities in numerical solutions that either dramatically increase the number of calculation steps required to solve the problem or even introduce oscillations in the error that grow exponentially and eventually destroy the solution. These errors arise from either neglecting or

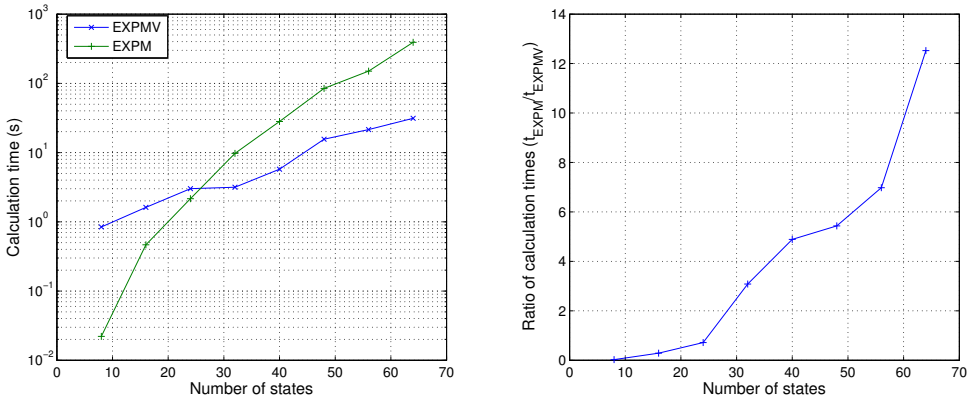

Figure 3.13: The relative merits of EXPM and EXPMV. The left-hand plot shows the calculation time for each routine to integrate to 1μs for systems containing a given number of states. The right-hand plot shows the ratio of the calculation time of EXPM relative to the calculation time of EXPMV. For small systems, EXPM is faster; for large systems, EXPMV is faster.

poorly sampling the highest frequencies in the problem. All the routines we use to solve the time dependent problem will need to be able to deal with the stiff nature of the system.

3.4.2 Short times: Product of a matrix exponential

As $\mathbf{H_0}$ is time-independent, solving the differential equation analytically is simple:

$$\mathbf{X}(t) = e^{\mathbf{H_0}t}\mathbf{X}(0) \tag{3.24}$$

This gives us complete information about the dynamics of the system. However, the difficulty in using this method arises from the cost of calculating the matrix exponential. Although $\mathbf{H_0}$ is highly sparse, its exponential will almost certainly be a very full matrix. The memory required to store a full matrix of the same size as $\mathbf{H_0}$ has already been discussed in section 3.1.2 where it was shown that the storage requirements for a system of over 100 states exceeds 1GB of RAM.

It would be more practical, then, to not form the matrix exponential at all but to instead solve the product of the matrix exponential with the vector $\mathbf{X}(0)$ directly. This removes the requirement of storing a very large full matrix that includes a lot of information that is unnecessary for calculating the time de-

pendence. The routine EXPMV written by Awad. H. Al-Mohy and Nicholas J. Higham [CM11] tackles precisely this problem. The matrix exponential is calculated by using a combination of a scaling and squaring method with a truncated Taylor series approximation. The cost of the routine is determined by the cost of computing a series of products of $\mathbf{H_0}$ with vectors of size $\mathbf{X}(0)$. This avoids needing to perform extensive operations on a matrix of size $\mathbf{H_0}$ and restricts the calculation to the much smaller system of size $\mathbf{X}(0)$. The performance of EXPMV depends on the number of matrix multiplications. For large systems, the benefit of operating on a smaller system will always outweigh the number of multiplications required.

Figure 3.13 shows the benefits of using EXPMV over MATLAB's routine EXPM that calculates the matrix exponential directly. The system is integrated for an evolution time of $1\mu s$. This corresponds to a typical pulse time. For small systems, the overhead of the number of matrix multiplications in EXPMV causes it to perform more slowly than EXPM, although the calculation times of both routines are extremely fast. As the size of the system increases, the calculation rates of EXPMV rapidly surpass those of EXPM, becoming over an order of magnitude faster for even a modestly-sized system of 64 states.

Figure 3.14 shows some characteristics of the EXPMV routine when used to integrate for the three lowest-lying levels in $^{43}Ca^{+}$, consisting of 64 states. The output of the routine is compared to an adaptive Runge-Kutta method and shown to be accurate to at least 1 part in 10^4 for calculation times of less than a minute. The calculation time of the routine scales approximately linearly with the evolution time for pulses of the order $1\mu s$, allowing the integration time for a particular system to be estimated to within 10%.

However, the scaling of the routine means that EXPMV can become slow for very long pulses. Pulses of the order of 100s of μs start taking of the order hours to calculate. If we are to look at such long pulses, we will need to use a

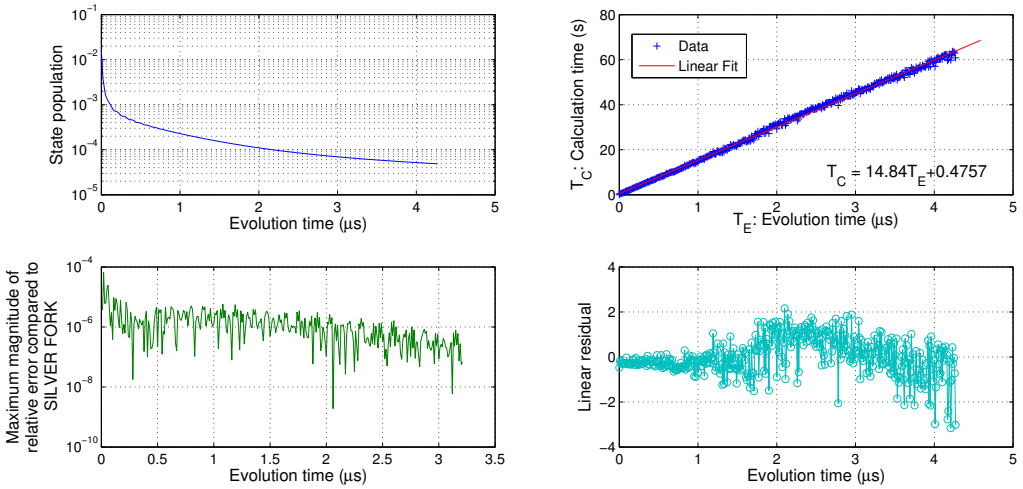

Figure 3.14: Properties of integration using EXPMV. The upper left plot shows the output of the EXPMV routine. The lower left plot shows the error relative to an adaptive Runge-Kutta routine. The upper right plot shows the calculation time required to integrate to a given evolution time along with a linear fit. The linear residuals to this fit are shown in the bottom right.

more efficient method of calculation. Also, this method only works for solving a differential equation that is explicitly time-independent. As soon as we allow $\mathbf{H_0}$ to have some time dependence, we can no longer solve the system analytically in such a simple way. Systems with a time-varying $\mathbf{H_0}$ are important for modelling the ion's motion; see section 4.6. The following two sections describe routines that can deal with time-dependence and long pulses respectively.

3.4.3 Time-varying H_0: Step-size modified Runge-Kutta

If we can no longer assume that $\mathbf{H_0}$ is time-independent, one approach we can adopt is to discretize the differential equation. We solve the system iteratively from the initial solution over a series of time-steps of some duration h. There are many such methods to solve differential equations, but we will focus on a single family of methods in this study. The Runge-Kutta methods calculate the solution at the endpoint of each time-step by using a series of intermediate calculations at various points within the time-step. The number of calculations and the locations within the step used for calculation depend on the method

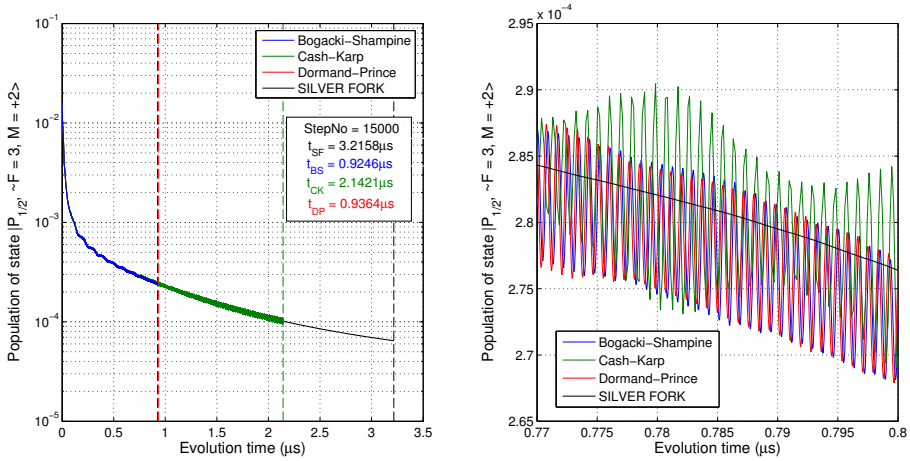

Figure 3.15: Numerical oscillations in the output of embedded Runge-Kutta methods. Four methods are used to the solve the same time dependent problem with the same number of iterations. In the embedded routines, oscillations appear in the solution that do not agree with the output of iterative methods with a fixed step size. The magnitude of the oscillations is small, but their presence dramatically reduces the step sizes. SILVER FORK produces no oscillations, and thus can have longer step sizes. It therefore integrates to a larger evolution time in a fixed number of iterations, outperforming even routines that are lower order.

employed.

There are many Runge-Kutta methods of varying order and complexity. This includes methods that include additional calculations to estimate the truncation error introduced at each time-step and modify the value of h for subsequent time-steps so that some predetermined error bound is not exceeded. However, the stiff nature of the Bloch equations causes problems for such solvers as demonstrated in figure 3.15. The large range of magnitudes of the oscillating terms in $\mathbf{H}$ causes small oscillations in the output of some of the more advanced routines. The routines studied here are the third order Bogacki-Shampine method [BS96], and the fourth order Cash-Karp [CK90] and Dormand-Prince [DP80] methods. These oscillations are numerical artifacts as can be demonstrated by comparison to a non-adaptive fourth-order Runge-Kutta method in which the oscillations do not occur. The main problem with these oscillations is not necessarily the error they introduce in the solution; this is at worst of the order of 1%. The oscillations severely limit the step-size that can be used in the

0										
1/4	1/4									
1/2	1/2									
1/4	0	1/4								
1/2	0	0	1/2							
1/2	0	0	0	1/2						
1	0	0	0	0	0	1				
1/2	1/12	1/6	0	1/6	0	1/12				
3/4	1/12	1/6	0	1/6	0	1/12	0	1/4		
3/4	1/12	1/6	0	1/6	0	1/12	0	0	1/4	
1	1/12	1/6	0	1/6	0	1/12	0	0	0	1/2
	1/6	0	1/3	0	1/3	0	1/6			
	1/12	1/6	0	1/6	0	1/12	1/12	1/6	1/6	1/12

Table 3.1: The Butcher tableau for SILVER FORK.

integration - the method sees an oscillation that, whilst numerical, appears to be a real solution to the integration and therefore must be followed closely. As a result, these routines run more slowly than a method that does not produce these oscillations.

An estimation of the truncation error is, however, a very appealing property of an integrator. The calculation error of such a routine is approximately known at any point, making it useful for applications where high accuracy is required. One example is analysing readout, where high fidelities can require accuracies of 10^{-4} or more [HAB+14, MSW+08]. In the absence of numerical oscillations, such routines will calculate faster than routines with a fixed step-size for a given initial truncation error. Necessarily, the faster dynamics at the outset of any time-dependent problem that reaches a steady state will have structure that varies on smaller timescales than occurs after the transients have mostly died away and the system approaches the steady state more smoothly. The same truncation errors that were achieved during the initial (faster) dynamics can be achieved with larger step-sizes during the decaying (slower) portion of the dynamics. An adaptive routine can therefore integrate to a longer evolution time than a fixed step-size routine for the same truncation error. See figure 3.16.

To maintain this property of adaptive methods, but to avoid the problems of oscillations seen by some of the integrated methods, we return to the remarkably stable fourth-order Runge-Kutta method. Rather than using an embedded method, we use a midpoint-method. We calculate the endpoint of each step *twice* using two fourth-order methods, one of which has half the step-size of the other. In other words, one method calculates the full step and the other calculates two steps by also calculating the midpoint. The relative error between the two routines can be used to estimate the truncation error. This routine estimates the truncation error using the approach described in Numerical Recipes [PFTV92]. This is the routine SILVER FORK (Step-size Iterations Locally Varied and Error Restricted Fourth-Order Runge Kutta). The Butcher tableau for SILVER FORK is shown in table 3.1. This method has the advantage that it does not see the oscillations present in more advanced methods and can significantly outperform them, despite using more computations per step.

SILVER FORK is designed to solve problems where $\mathbf{H_0}$ varies with time. It also has some use for considering simpler systems which do not show this dependence. An adaptive iterative method must necessarily pass over all the oscillating structure of the problem in high detail, and this can provide some insight into the dynamics that a direct solution based on an analytic method might miss. However, calculation times are necessarily much slower. For example, integrating to 100s of μs takes of the order a day. In general, EXPMV is most useful for most situations we are considering for systems where $\mathbf{H}$ has no explicit time dependence.

3.4.4 Large times: Leap-frogging to the steady state

Now let us return to the case of a time-independent Hamiltonian and consider especially the approach to the steady state. EXPMV is a very useful routine, but the linear scaling with system size means that calculations to very large evolu-

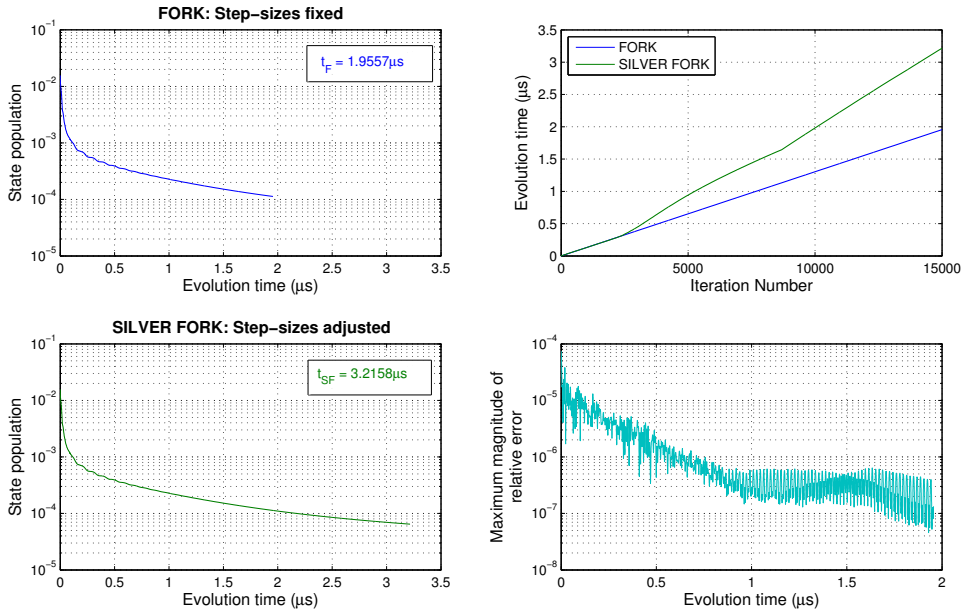

Figure 3.16: The advantage of using an adaptive step size. The upper left image shows the output of FORK (fixed step sizes) and the lower left the output of SILVER FORK for the same number of iterations. The upper right image shows the evolution time reached in a given number of iterations. The lower right image shows the maximum relative error of the population terms of SILVER FORK relative to FORK. These residuals have been processed using a low-pass filter in order to show the behaviour more clearly. The oscillations at evolution times greater than 1ms occur around the bound of the truncation error passed as an input to SILVER FORK. These are most clearly seen when the initial fast dynamics have decayed. For the same number of iterations, SILVER FORK integrates to a larger evaluation time.

tion times can become very time-consuming. Integration to 100s of µs takes of order hours, and analysing long pulses will be needed when considering readout methods. It would be useful, then, to have another method that could quickly analyse the behaviour of a system as it approaches the steady state.

This is the purpose of our final routine. This routine employs an integration method that is deliberately *inaccurate*, Euler's method. This is an iterative method that uses the gradient of $\mathbf{X}$ at each step to determine its value at the following step. We will be using the implicit Euler method, where the n iteration has some step-size h:

$$\mathbf{X_{n+1}} = \mathbf{X_n} + h\dot{\mathbf{X}}_{\mathbf{n+1}}, \tag{3.25}$$

because of its important convergence properties. This is a first-order method, so

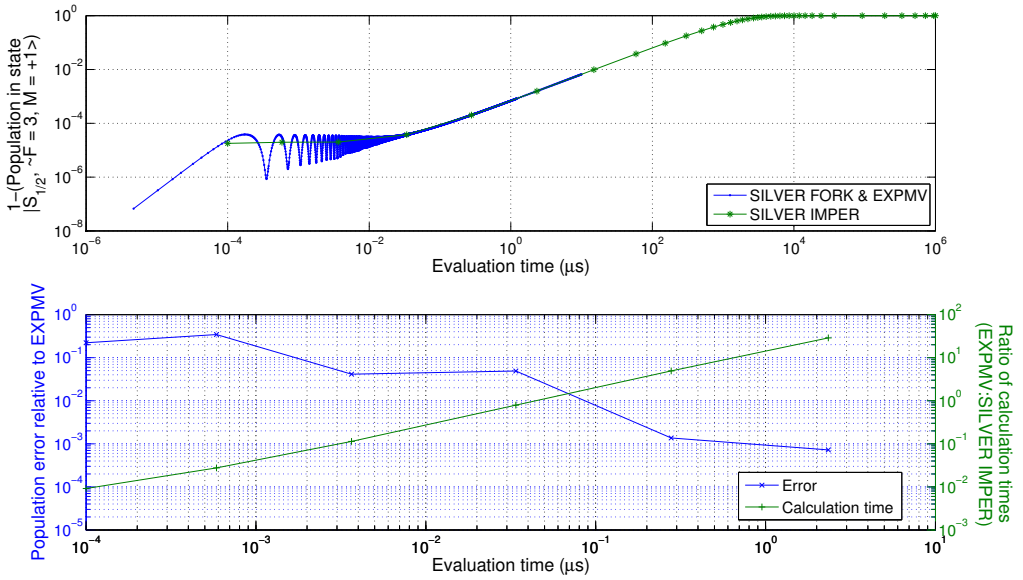

Figure 3.17: Using SILVER IMPER to skip over the initial dynamics. SILVER FORK, EXPMV and SILVER IMPER are used to calculate the error when optically pumping from the state in a qubit that is not resonant with the pumping beam. The system comprises the five lowest-lying levels in $^{43}Ca^{+}$, consisting of 144 states. The state initially has all of the population in the state being examined. The upper plot shows the population that has been pumped out of the qubit state, calculated using a combination of SILVER FORK and EXMPV, shown in blue, and SILVER IMPER, shown in green. The lower plot shows the errors and calculation times of SILVER IMPER relative to EXPMV. The errors are large at small times but quickly reduce in size as the initial oscillations subside. The relative speed of SILVER IMPER increases exponentially with calculation time.

it will in general produce large errors during calculations. However, the implicit Euler method has a nice property for differential equations with no explicit time dependence; it will *necessarily* converge on the steady state at large times. The following section proves this using a modified version of the proof in Numerical Recipes. We can use this property to largely ignore the initial oscillations of the dynamics and skip to the decaying dynamics as the system approaches the steady state. We cannot ignore the initial dynamics completely, of course, as these determine the decaying dynamics. However, we can use orders of magnitude fewer iterations than a Runge-Kutta method. In some cases, even fewer than 10 points can be calculated over the entire range of the oscillating dynamics and still produce an accurate description of the decaying dynamics. See figure 3.17.

We make the most of the advantages of this routine by using an adaptive step size. However, we deliberately don't want to restrict the routine to taking small step sizes during the oscillating dynamics; rather, we want to encourage inaccurate calculation until the decaying dynamics are reached. To achieve this, even though this is a first order routine, we use the same error calculation as SILVER FORK; that is, we use a *fourth-order* estimation of the truncation error. This will necessarily expand the step-sizes of the routine past any error bound we set, encouraging the routine to calculate to large evolution times. This routine is SILVER IMPER (Step-size Iterations Locally Varied and Error Restricted IMPlicit EuleR)[2]

The error parameter in this method Ξ should be chosen with care to ensure that the calculation time is minimised whilst the error in the portion of the dynamics being considered is also minimised. See figure 3.18 for a demonstration of the calculation times and errors produced for a pulse of duration $100\mu s$ for a range of values of Ξ.

3.4.4.1 Proof of Convergence

This proof is an modification of that presented in [PFTV92]. The backwards Euler method solves the differential equation by calculating the gradient at the endpoint of each step rather than at the initial point:

$$\mathbf{X_{n+1}} = \mathbf{X_n} + h\dot{\mathbf{X}}_{\mathbf{n+1}} \tag{3.26}$$

From (3.23), we can replace the gradient term so that:

$$\mathbf{X_{n+1}} = (\mathbf{I} - h\mathbf{H_0})^{-1}\,\mathbf{X_n} \tag{3.27}$$

where we assume that $\mathbf{H_0}$ is time-independent and this modification to (3.26) is therefore exact; the backwards Euler method demonstrates exact convergence

[2] An imper is someone who 'imps', that is, who grafts. From its use in falconry, where wings are grafted onto the tails of birds, the term 'imp' is used figuratively to describe something that is helped to fly higher and farther.

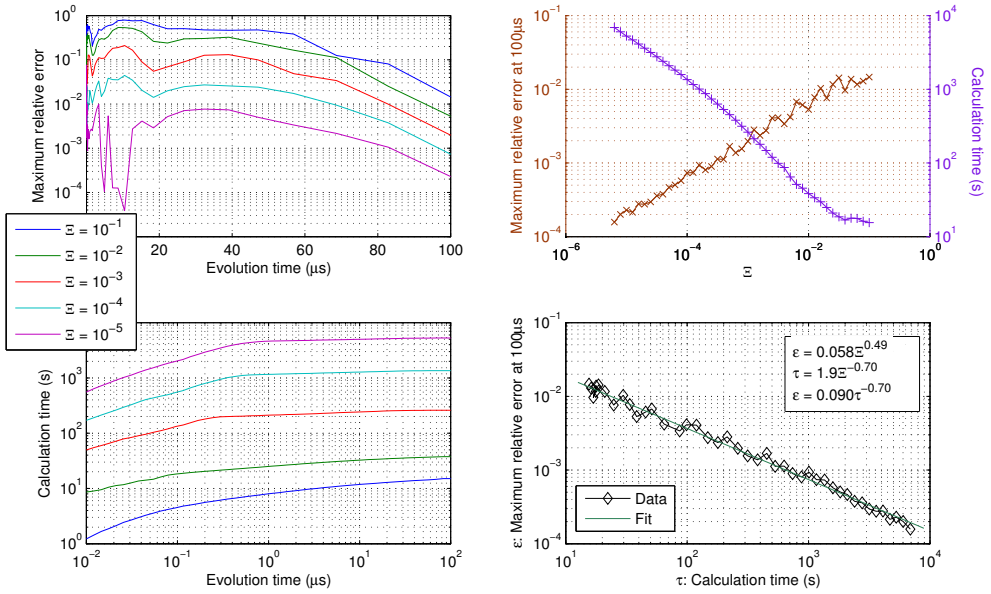

Figure 3.18: The benefits of using SILVER IMPER. The upper left plot shows how the largest relative population error output by SILVER IMPER relative to EXPMV for a three-level system in $^{43}Ca^{+}$ varies with Ξ. This is the same system used to analyse EXPMV and SILVER FORK. The errors dramatically reduce in magnitude as the steady state dynamics are reached. The lower-left plot shows the calculation times. The right-hand plots show the calculation error and calculation time at an evolution time of 100μs for large range of Ξ. Smaller values of Ξ produce successively smaller errors, but increase the calculation time accordingly. Exponential fits for the relations between Ξ, the calculation time and the calculation error are given.

only for a first-order differential equation in which the derivative is not a function of the dependent variable. Let us define:

$$\mathbf{S} = (\mathbf{I} - h\mathbf{H_0})^{-1} \tag{3.28}$$

In order for this method to be stable, we require that the method will reach a steady state. If we define this steady state as $\mathbf{X_V}$ then, at large times:

$$\lim_{n\to\infty} (\mathbf{S}^n\mathbf{X}) = \mathbf{X_V} \qquad \forall \mathbf{X} \tag{3.29}$$

To examine if this is the case, let us briefly return to the original differential equation:

$$\dot{\mathbf{X}} = \mathbf{H_0}\mathbf{X} \tag{3.30}$$

Let us factorise $\mathbf{H_0}$ into its eigenvector and eigenvalue matrices $\mathbf{D}$ and $\mathbf{V}$ respectively:

$$\dot{\mathbf{X}} = \mathbf{DVD}^{-1}\mathbf{X} \tag{3.31}$$

We define:

$$\mathbf{Y} = \mathbf{D}^{-1}\mathbf{X} \tag{3.32}$$

so that equation (3.31) is written as:

$$\dot{\mathbf{Y}} = \mathbf{VY} \tag{3.33}$$

Solving for $\mathbf{Y}$ and writing the solution in terms of $\mathbf{X}$:

$$\mathbf{X} = \mathbf{D}e^{\mathbf{V}t}\mathbf{D}^{-1}\mathbf{X_0} \tag{3.34}$$

where $\mathbf{X_0}$ is some initial set of coherences. Noting that $\mathbf{V}$ is diagonal, convergence can only be achieved when all of the eigenvalues of $\mathbf{H_0}$ associated with the dynamics are nonpositive. If we assume that there is only one steady state, the single zero eigenvalue describes the steady state and will not decay. This ensures that any initial state is eventually projected onto the steady state eigenvector at large times. In addition, any eigenvalues describing terms that do not affect the dynamics need not be negative; terms that remain zero at any value of t can be multiplied by any exponential term without affecting the dynamics. Practically, this means that the eigenvalues of $\mathbf{H_0}$ associated with the y_{jj} terms need not be negative as they do not affect the system. For the rest of this calculation, we will ignore these terms.

Let us now return to $\mathbf{S}$. We note that:

$$\mathbf{S}^n = \mathbf{D_S}\mathbf{V_S}^n\mathbf{D_S}^{-1} \tag{3.35}$$

where $\mathbf{D_S}$ and $\mathbf{V_S}$ are the eigenvalues and eigenvectors of $\mathbf{S}$. The requirement that $\mathbf{S}^n\mathbf{X}$ tend to the steady state is therefore:

$$\lim_{n\to\infty} \left(\lambda_m^S\right)^n = \delta_{mp} \tag{3.36}$$

where λ_m^S are the eigenvalues of $\mathbf{S}$, the pth eigenvalue is the eigenvalue associated with the steady state eigenvector and δ_{mp} is the Kronecker delta. In other words, we require that all of the eigenvalues decay save the steady state eigenvalue. This ensures that the solution will project onto the steady state eigenvector. The eigenvalues λ_m^S of $\mathbf{S}$ are related to the eigenvalues λ_m^H of $\mathbf{H_0}$ as:

$$\lambda_m^S = \frac{1}{1 - h\lambda_m^H} \tag{3.37}$$

Since (3.34) has a unique steady state solution, we must have:

$$\mathrm{sgn}\left(\lambda_m^H\right) = \delta_{mp} - 1 \tag{3.38}$$

so that:

$$\lambda_m^S < 1 \qquad \forall m \neq p \tag{3.39}$$

$$\lambda_p^S = 1 \tag{3.40}$$

which, taken together, reproduces (3.36). In other words, the implicit Euler method will necessarily converge on the steady state at large times, regardless of the step size used.

3.4.5 Usage of the three routines

The three routines EXPMV, SILVER FORK and SILVER IMPER can efficiently solve the time dependent problem for all situations and pulse lengths that are considered in the following sections. See figure 3.19 for a demonstration of the calculation times of the routines for a system where $\mathbf{H}$ is time-independent. For such systems, EXPMV outperforms the other two routines by orders of magnitude for short pulses. As pulses approach the μs length, SILVER IMPER begins to compete, until for pulses of 10s of μs or more EXPMV is outperformed. A particular advantage of SILVER IMPER is that, when the initial oscillating dynamics have fully died away, the routine converges on the steady state extremely quickly. The step-size increases exponentially without increasing the error of the

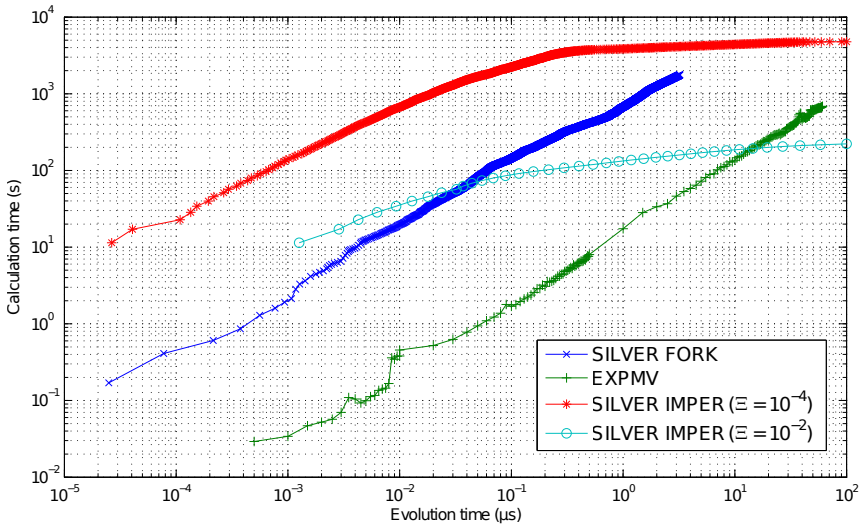

Figure 3.19: A comparison of the three integration routines for a problem with time-independent $\mathbf{H_0}$. The calculation time required to reach a given evolution time is shown. For small and intermediate times, the EXPMV routine always calculates faster. However, at large times, SILVER IMPER outperforms EXPMV.

routine, so that pulses on the timescale of seconds or even longer can be calculated in approximately the same time as pulses on the the timescale of μs. For systems with time-dependent $\mathbf{H}$, we cannot use either EXPMV or SILVER IMPER. Instead, we must use SILVER FORK.

4

Doppler Cooling in $^{40}Ca^{+}$

> We know from general experience that a ship with a moderately deep draught which is steering towards the oncoming waves has to receive, in the same amount of time, more waves with a greater impact compared with a ship that is not moving or is even travelling along on the direction of the waves.
>
> Christian Andreas Doppler

In the previous two chapters, we have described the Bloch equations and numerical methods required to calculate with them efficiently. We now move to using these simulations to analyse and predict the behaviour of an atomic system.

The simplest possible cooling cycle consists of two states. The momentum of photons exciting the atom between the two states is the means by which the atom is slowed (see, for example, [WI79, IW82, LBMW03, LL15]). This is readily generalised to a pair of levels, each of which has several substates (for example, [Sta03]). However, for many atoms, there may not exist a convenient, closed electric dipole transition between a pair of levels. Population may have to be pumped back onto the cooling transition from additional levels. Systems with three or more levels often produce behaviour that cannot be modelled with standard Doppler cooling theory. In particular, the characteristic timescales of repumping process will in general be much slower than that of the main fluores-

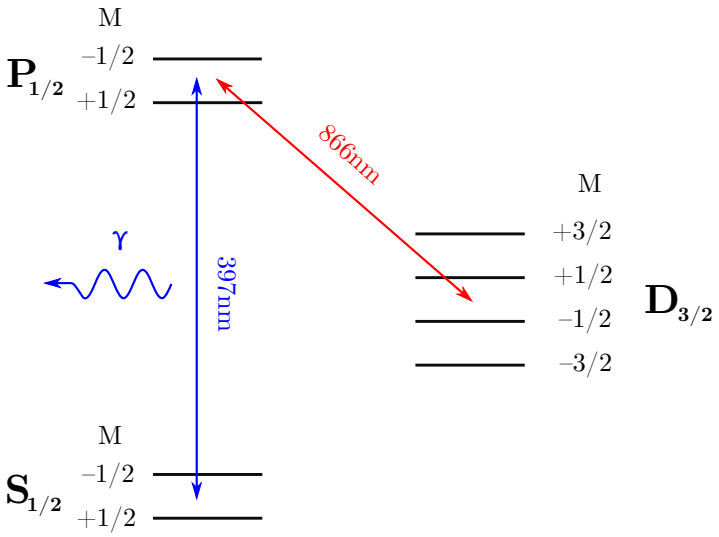

Figure 4.1: A schematic diagram of the levels of $^{40}Ca^{+}$ showing the Zeeman structure. The energy differences between the states are not to scale. Fluorescence is collected from the blue transition.

cence process used for cooling. These timescales may be slow enough that they are no longer small compared to the ion's motion in the trap; in such cases, the ion's secular motion must be treated explicitly. Regimes where the ion's external motion cannot be neglected are also present in two-level and two-state systems, but they are especially prevalent for systems involving repumping.

Ca^{+} is one example of an ion where the simplest cooling scheme requires a repumper. The cooling transition is an S to P transition and population is repumped from a D level. In the final analysis, we will be interested in using the Bloch equations to determine a suitable cooling scheme for the isotope $^{43}Ca^{+}$. However, the benefit and difficulty of using this isotope is that it has a large nuclear spin of $I = 7/2$. The three-level cooling scheme consists of 64 states and 64^2 coherences. Before tackling this system, it will be instructive to look at the much simpler isotope $^{40}Ca^{+}$. This has no nuclear spin, and the relevant three levels consist of only 8 states, reducing the number of coupled equations describing the dynamics by two orders of magnitude. The dynamics are significantly simpler, and it is much easier to interpret the ion's behaviour. The level structure of

	Cold ion	Hot ion
Fast internal dynamics	Steady-state	Excursion
Slow internal dynamics	Dynamic	Motional

Table 4.1: A demonstration of the regions of parameter-space for which the four temperature models described in this chapter are valid. The timescale of the internal dynamics is taken relative to the timescale of the ion's physical oscillations in the trap.

$^{40}Ca^{+}$ is shown in figure 4.1.

In this chapter, I present four increasingly refined analyses of Doppler cooling for an atomic system with a repumper. The final approach incorporates the entire dynamics of the ion and accurately describes the conditions for which an ion is cold or hot. In addition, it reproduces experimental fluorescence scans even in regimes where the ion is too hot to be modelled by the simpler approaches. The four models described are the 'steady-state' method, the 'excursion method', the 'dynamic method' and the 'motional method'. The regions of parameter space for which these four models are valid are summarised in table 4.1. These models are demonstrated through examples. Appendix H lists the parameters used to produce these demonstrations.

First, however, I outline the process of using the simulations to fit experimental data. For details of the apparatus used to take the data shown in this chapter, see [Web05, Szw09]. For a review of the cooling models discussed in this chapter, see [SJS16] (currently in preparation).

4.1 Fitting experimental data

The primary experimental diagnostic we use to study the behaviour of a trapped ion are fluorescence scans. These scans are produced by stepping the frequency of one of the beams through a pre-defined range, keeping all other parameters constant. The variation in the fluorescence of the ion is recorded.

An important feature of such scans are sharp resonant effects. The most

common type of resonance we will encounter are 'dark resonances', which occur when transitions at 397nm and 866nm for a particular pair of states in the $S_{1/2}$ and $D_{3/2}$ levels come into resonance and drive population out of the $P_{1/2}$. The detailed structure of these resonances depends on the atomic and laser parameters along with parameters describing the atomic motion. A combination of the structure of the resonances and the gross structure of the scan can be used to characterise the experimental conditions.

4.1.1 Departures from the stationary theory

A simple and fast method to simulate a fluorescence scan has been described in section 3.2. At each point on the scan, a Liouvillian $\mathbf{H}(\omega)$ is generated and the state populations of the ion are determined by inverting it:

$$\mathbf{X} = \mathbf{H}^{-1}(\omega)\mathbf{V} \tag{4.1}$$

$\mathbf{V}$ is a column vector with first element 1 and every other element 0. The detected fluorescence f is proportional to the population of the upper level of the fluorescence transition:

$$f = \eta A_{397} P(\mathrm{P}_{1/2}) \tag{4.2}$$

where η is the efficiency of the detection system. We can also describe the ion's response through the rate of downward transitions on the cooling transition R where:

$$f = \eta R \tag{4.3}$$

$$R = A_{397} P(\mathrm{P}_{1/2}) \tag{4.4}$$

However, this method of modelling the ion's fluorescence assumes that the ion is stationary. In other words, it assumes that the ion is perfectly cold. For a trapped ion, this is never the case. Fluctuating fields of the trap and the interactions with the laser beams cause the ion to have some finite energy. The ion will

typically oscillate in the trap, and these oscillations cause the frequencies of the beams observed in the frame of the ion to oscillate.

For the purposes of a fluorescence scan, we are interested in the mean fluorescence produced by the ion over a trap cycle. As long as the temperature of the ion is sufficiently low, the fluorescence will only depend linearly on the ion's velocity. The distribution of the fluorescence will vary symmetrically about the value at zero Doppler shift, and the mean fluorescence will be equal to the fluorescence of an ion in equilibrium at zero velocity to a good approximation. We may therefore obtain the ion's mean fluorescence by modelling a stationary ion subjected to lasers which are not Doppler-shifted. However, to take account of the width of the frequency distribution observed by the oscillating ion, we must add additional effective 'linewidths' to these lasers equal to the width of the frequency distribution explored by the ion.

This 'stationary theory' is only valid when second and higher-order derivates of the fluorescence relative to the ion's velocity are small compared to the ion's velocity excursions. I will refer to the modification of the stationary theory due to the inclusion of these higher-order derivatives as the 'excursion effect'. If we assume that all beams are propagating along the same axis $\hat{x}$, the excursion effect will not be prominent when:

$$\frac{R}{\langle v_x^2 \rangle} \gg \left. \frac{d^2 R}{dv_x^2} \right|_{v_x=0} \tag{4.5}$$

where v_x is the component of the ion's velocity along the beam direction and R is calculated at equilibrium. The velocity of the ion can be determined from the ion's temperature using the methods described later in this chapter. One way to interpret this condition is that the ion's temperature needs to be small enough that the ion does not traverse significant structure of the fluorescence curve as it oscillates. If the energy of the ion satisfies the limit over a significant portion of the scan, we can fit the scan using the theory for a stationary ion, if necessary by neglecting those portions of the scan where the ion is hot. Figure 4.2 shows

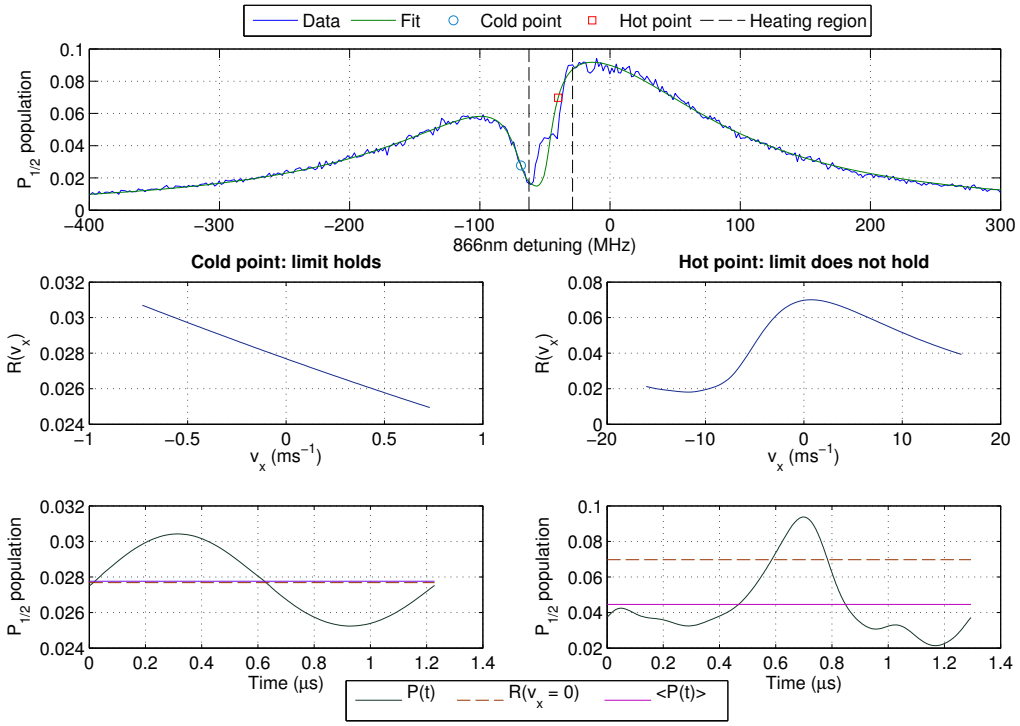

Figure 4.2: A demonstration of the behaviour of the ion for cases when the limit (4.5) does and does not hold. The upper plot shows the fluorescence scan with a fit produced assuming a stationary ion. The region in which the fitted curve departs from the data is marked, and two points at which the ion is either cold and hot are shown. The lefthand plots show the behaviour of the ion at the cold point and the righthand plots the behaviour at the hot point. The middle plots show $R(v_x)$ over the full range of the ion's oscillations at equilibrium. The lower plots show the population of the $P_{1/2}$ as a function of time over one cycle of the ion's motion in the trap. These data are produced using the full calculation method described in section 4.6. For the cold ion, the variations in $R(v_x)$ are linear to a very good approximation over the full extent of the ion's motion. The $P_{1/2}$ population oscillates sinusoidally and the mean population and the population of a stationary ion only differ by of order 0.1%. The values of the parameters appearing in the limit are: $\langle v_x^2 \rangle = 0.267\text{m}^2\text{s}^{-2}$, $R = 0.0277$ and $d^2R/dv_x^2 = +5.537 \times 10^{-4}$. The limit holds in this case. By contrast, the variations in $R(v_x)$ for the hot ion are not linear; the ion traverses significant structure of the fluorescence curve as it moves. The $P_{1/2}$ population oscillates asymmetrically about $R(v_x = 0)$ with a tendency towards lower population. The mean population is thus significantly lower than the population that would be observed in a stationary ion. The values of the parameters in the limit are: $\langle v_x^2 \rangle = 200\text{m}^2\text{s}^{-2}$, $R = 0.0697$ and $d^2R/dv_x^2 = -2.301 \times 10^{-3}$. The limit does not hold in this case.

an example of a fluorescence scan where the stationary theory is only valid for a portion of the scan, and which was fitted by neglecting the region where the ion is hot. The behaviour of the ion when the limit (4.5) holds and when it does not hold is shown.

To provide a reference scale for the temperature of an ion, it is useful to refer

to the limiting temperature of the equivalent two-state system. For a two-state ion, the lowest energy that can be achieved is the 'Doppler limit'. The value of this limit for $^{40}Ca^{+}$ is:

$$E_D = 0.130\mu\text{eV} \tag{4.6}$$

This limit is often expressed in terms of a limiting temperature. Even though this system is not in thermal equilibrium, we can use the distribution of the ion's velocities to define a temperature. The Dopper limit for $^{40}Ca^{+}$ expressed as a temperature is (see, for example, [Sta04, Har13]):

$$T_D = 0.5041\text{mK} \tag{4.7}$$

For a standard cold scan, the energy of the ion will be of order the Doppler limit. This allows us to quantify the size of additional linewidths we will need to apply to the stationary theory. For $\delta_{397} = 2\delta_D$, the additional linewidths are:

$$\delta_{397} \approx 2\text{MHz} \tag{4.8}$$

$$\delta_{866} \approx 1\text{MHz} \tag{4.9}$$

For the purposes of fitting, the sum of the linewidths of both lasers will be taken to be of order 2.5MHz. The value will change from scan to scan due to the dependence of the parameter on the ion's energy.

I will refer to areas of fluorescence scans which are not well fitted by the stationary theory as 'heating regions'. To model the fluorescence of the ion in these regions, and to determine the experimental parameters of scans with large heating regions to better accuracy, we must use another approach. This is discussed in section 4.7.

Figure 4.3 shows an example of three different scans demonstrating different sizes of a heating region. In the upper-left scan, the ion is significantly hotter than the limit (4.5) for a large portion of the scan. In the lower, the ion exceeds the limit in only a small region to the blue of the dark resonance. In the upper-right scan, the ion's energy does not exceed the limit. The two scans with heating

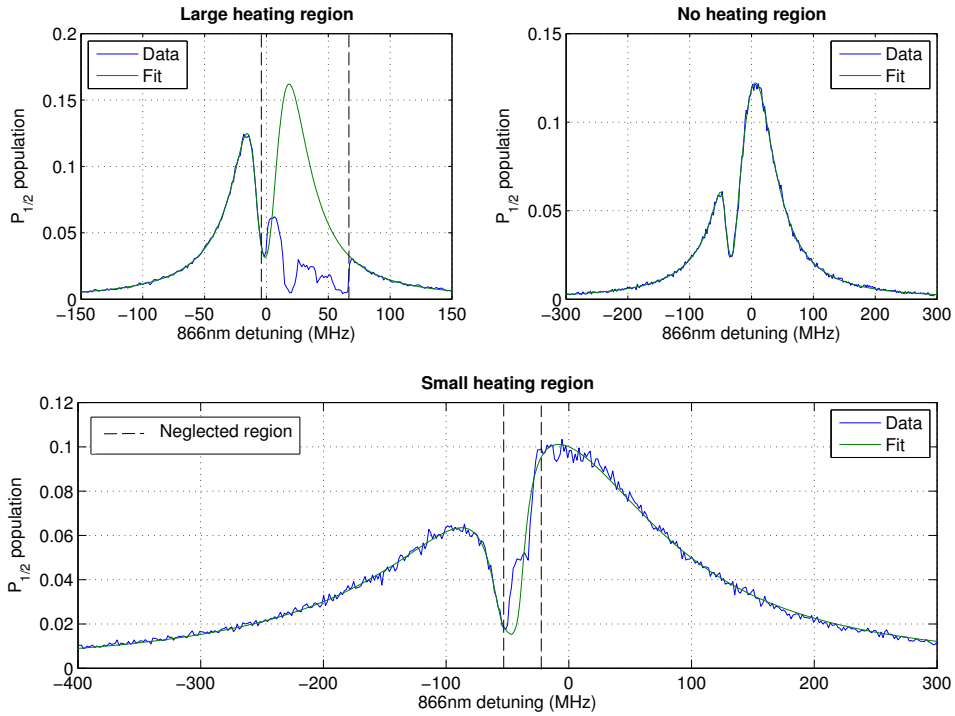

Figure 4.3: A demonstration of the possible sizes of heating regions in $^{40}Ca^{+}$. Each set of data has been fitted using the stationary theory with any heating regions removed. The removed regions are marked by dashed lines. The range of possible curves that fit the data well increases as the size of the neglected region increases. The lower fit still has a small range of possible parameter values; uncertainty arises only in the parameters that describe the depth of the dark resonance, principally the polarisations and the linewidths. By contrast, the parameter uncertainty in the upper-left curve is very large. The height of the principal peak can vary by a factor of two with a negligible change to the residuals outside of the heating region. The stationary theory can only give useful results when the size of heating regions are small.

regions show significant departures from the stationary theory for those regions where the ion is hot, while the upper-right scan fits well over the entire range of the scan. The upper-right scan shows no discrepancies when fitted with the stationary theory and will be used as the basis of a discussion of the fitting process in the following two sections.

It might seem that it would be more rewarding to examine 397nm scans rather than 866nm scans. The 397nm beam is the beam driving the cooling transition, and 397nm photons with their larger Doppler shifts and larger coupling to the ion will dominate the heating and cooling processes. In other words, information about the energy of an ion can be taken directly from a 397nm scan,

but only indirectly from an 866nm scan. However, the strong dependence on the ion's energy means that 397nm scans in actuality contain relatively little useful information. Figure 4.4 shows an 866nm and 397nm scan taken at the same experimental conditions. The 866nm scan fits very well and has no visible heating region. The 397nm scan, by contrast, strongly heats to the blue of the resonance peak. This behaviour is true for all 397nm scans. As the detuning experienced by the ion moves over the resonance peak, the effect of the 397nm switches from cooling the ion to heating it. This change occurs because the direction in which preferential absorption of photons occurs reverses. The dropout of a 397nm scan therefore shows the approximate location of the resonance peak, which is useful information when trying to set the detuning of the 397nm to a particular value in an experiment. However, from the point of view of fitting the data, half of the scan will always be missing. Unless there is a dark resonance in this portion of the scan, there are too few features to be able to identify the experimental parameters to anything but a very crude degree. As a result, the fluorescence scans used for fitting in the rest of this thesis will always be 866nm scans.

4.1.2 Fitting a scan

The standard method we use to fit a scan where the ion is mostly cold is a least-square minimisation approach applied to the steady-state theory. The initial set of parameters input to the fitting process are the measured experimental parameters, although the linewidth assigned to the 397nm beam is increased by 2MHz to take account of the ion's motion as stated above. Parameters are floated that have not been measured to an accuracy of less than of order one percent.

Each laser beam is described by up to 5 parameters: the intensity, two parameters describing the polarization state, the laser detuning and the modified linewidth. The polarisation of the beam is described by the fractions of its energy in the three possible polarisation components, σ^+, π and σ^-. These fractions

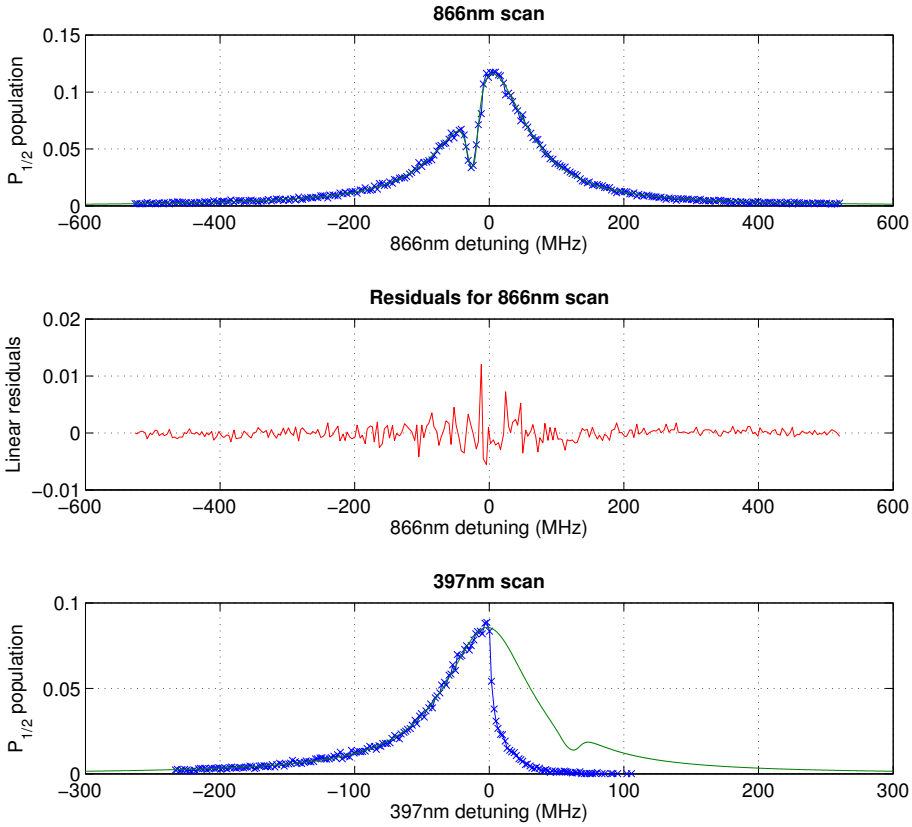

Figure 4.4: An 866nm and 397nm scan at the same experimental conditions. The detuning of the 397nm for the upper scan is -24.6MHz and the detuning of the 866nm for the lower scan is +62.4MHz. The 866nm scan shows no appreciable heating region, whereas the 397nm shows significant heating to the blue of the resonance peak. The 397nm curve contains very little information to identify the experimental parameters, and without a corresponding 866nm scan it would be difficult to make anything but the crudest match to them.

sum to one, so the polarisation of the beam can be described by two orthogonal combinations of these components. In these routines, I will describe the polarisation state of each beam by using the fractions of its sigma components, σ^+ and σ^-. This convention is convient as, in our experiments, the majority of beams have nominally no π fraction and equal fractions of σ polarisation. The two polarisation parameters are constrained within the fitting routine so that their sum is never greater than 1.

Parameters floated can also include calibration factors associated with converting the experimental data to a scan of frequency against state population.

The data takes the form of a voltage on a digital-to-analogue converter (DAC) and the number of counts collected from the ion over some fixed time. The frequency scale is determined from the DAC voltage V by using a calibration C and a scale offset ω_0:

$$\omega = CV + \omega_0 \tag{4.10}$$

The population is determined from the count-rate by:

$$P(\mathrm{P}_{1/2}) = \frac{f - f_0}{\eta \tau A_{397}} \tag{4.11}$$

where f is the detected fluorescence in time τ and f_0 the background fluorescence. The DAC is calibrated by measuring the voltage difference between two atomic features with a well-known frequency separation [Har13]. The collection efficiency has been measured using several different approaches as documented in [Don00] and [SKDB10, Har13]. The background is measured by temporarily shelving the ion in the $\mathrm{D}_{3/2}$ by turning off the 866nm beam, ensuring that the ion does not fluoresce and only background photons are detected. If these calibrations are not known to high accuracy, they are fitted. The parameters that can be fitted are C and ω_0 for the DAC calibration and f_0 and η for the population calibration. Of these parameters, only ω_0 is not measured directly by the current experiments. However, it can be estimated from the condition that the detunings of the 397nm and 866nm must be equal at the bottom of the dark resonance.

It is possible to infer errors in the measured values of the calibration parameters for a particular set of equipment by looking at a series of scans taken at different conditions. The mean fitted value of a given calibration parameter is also a measure of the physical value of the calibration, and if this is significantly different from the measured value, this implies an error in the measured value. From the scans I have analysed, no error in η has been observed. However, the measured values of C for all lasers are consistently 10% larger than output by the fitting routines. Therefore, for the vast majority of the fits in this thesis, these

calibration parameters are not fitted; η is fixed at its nominal value, and C 10% lower than its nominal value.

In order to ensure that the scans are not biased by fitting regions that the steady-state theory does not adequately describe, heating regions are removed before any fitting takes place. There is no hard and fast rule for determining where the ion is hot from the fluorescence curve alone; the fluorescence of the ion is governed by the response of the ion to variations in the frequencies of both beams, and the fluorescence curve only shows the response to a single beam. To identify the heating regions, we must be largely guided by experience. However, we can consult the shape of simulated fluorescence curves close to the nominal experimental parameters to see the sort of features that a cold ion might produce under these conditions. Features in the data that are not matched by these simulations could be evidence that the ion is heating in those regions.

The fitting process we use provides a plot of the residuals of the fit and the norm of the residuals after every fitting iteration. Included with this plot is a call to a function that halts the process with a button press. It is common that a fit to a scan will stagnate, and the norm of the residuals will tend to some fixed value even while the parameter changes at each iteration are still larger than the parameter tolerances. Allowing the user to stop the fitting process when the residuals are acceptable can significantly speed up the fitting process, especially when many scans are being analysed at once.

4.2 Modelling the ion's motion

In the following sections, I describe four models that can be used to determine the temperature of the ion directly.

In the first instance, we ignore the trapping potential and consider only the steady-state of the ion's motion. This standard Doppler cooling theory leads to a prediction that, in the steady state, the motional state is thermal with a

temperature given by the ratio:

$$T \sim R(v_x) \left(\frac{dR(v_x)}{dv_x} \right)^{-1} \tag{4.12}$$

assuming the ion's transitions are dominated by the downward transition rate $R(v_x)$ on a single transition. v_x is the component of the velocity of the ion along the beam path. We will show in later sections that this method is often inadequate to treat cooling of trapped particles, especially when there is a structure such as a dark resonance in the fluorescence profile.

This theory has been used, largely unmodified, since the first days of laser cooling of atoms [LBMW03, LL15]. I present the theory in full for the insights this will provide into the more advanced cooling theories that follow. I include a new and simple derivation of the heating associated with fluctuating forces.

4.3 The steady-state method

Let us consider an ion moving in a trapping potential in one dimension in the presence of several laser beams each with a peak frequency of ω_i. The beams are assumed to be parallel, but each beam may propagate in either direction along this line. Every time the ion absorbs a photon, the component of its momentum along the direction of beam i will change by an amount $\hbar\omega_i/c$. This will either increase the momentum of the ion or decrease it depending on whether the ion is moving away from or towards the laser respectively. If the ion absorbs more photons when it is moving towards the laser than when moving away, it will experience more absorption events that retard it than speed it up. The mean square velocity of the ion will thus be reduced.

Let us set the $\hat{x}$ direction along the direction of propagation the laser $i = 1$. The absorption of photons by the ion will change its momentum according to:

$$\frac{dp_x}{dx} = \sum_i \left\{ R_i(v_x) b_i \frac{\hbar\omega_i}{c} \right\} \tag{4.13}$$

where $R_i(v_x)$ is the rate of downward transitions on the transition excited by the ith laser. We refer to downward transitions, rather than the upward transitions associated with absorption, in order to use the same convention on R for describing both cooling and heating. In the steady state, the rate of downward and upward transitions are equal. p_x and v_x are the ion's momentum and velocity components along $\hat{x}$. b_i is a variable describing the direction of propagation of the ith beam. It has a value of $+1$ if the beam is co-propagating with the beam $i = 1$ and a value of -1 if it is counter-propagating. Each R_i is a function of all the parameters of every laser. They are a function of the ion's velocity through the Doppler effect. We will assume that the velocity of the ion is sufficiently small that we can describe the ion's motion using a Taylor expansion about zero velocity. We expand to first order:

$$\frac{dp_x}{dx} = \sum_i \left\{ \left(R_i(0) + \left.\frac{dR_i}{dv_x}\right|_0 v_x \right) b_i \frac{\hbar\omega_i}{c} \right\} \tag{4.14}$$

The equation of motion of the ion is:

$$m\frac{d^2x}{dt^2} = F(x,t) + \sum_i \left\{ \left(R_i(0) + \left.\frac{dR_i}{dv_x}\right|_0 v_x \right) b_i \frac{\hbar\omega_i}{c} \right\} \tag{4.15}$$

where $F(x,t)$ is the force on the ion from the trap. This force is independent of the cooling and heating imparted to the ion by the laser. I will assume for the purposes of this derivation that the trap is perfectly harmonic with a secular frequency of β. The effect of the constant term in $R_i(0)$ is to offset the centre of the oscillations from $x = 0$. By shifting the coordinate system so that the centre of oscillations is at $x = 0$, we remove the constant term in the differential equation:

$$\frac{d^2x}{dt^2} = -\beta^2 x + \sum_i \left\{ \frac{\hbar\omega_i b_i}{mc} \left.\frac{dR_i}{dv_x}\right|_0 \frac{dx}{dt} \right\}. \tag{4.16}$$

This differential equation describes damped harmonic motion:

$$x = x_0 \exp\left(\gamma_x t\right) \sin\left(\eta' t + \phi\right) \tag{4.17}$$

with a decay constant of the amplitude oscillations:

$$\gamma_x = -\sum_i \left\{ \frac{\hbar\omega_i b_i}{2mc} \frac{dR_i}{dv_x}\bigg|_0 \right\}. \tag{4.18}$$

Note that in order for any one of the beams to induce cooling, we require that the ion absorbs more photons when it is moving opposite to the direction of the beam than when moving in the same direction. In other words, we expect the product $b_i \frac{dR_i}{dv_x}\big|_0$ to be negative. The potential energy of the ion is proportional to x^2, resulting in a decay constant for the energy that is twice as large:

$$\gamma_E = -\sum_i \left\{ \frac{\hbar\omega_i b_i}{mc} \frac{dR_i}{dv_x}\bigg|_0 \right\}. \tag{4.19}$$

If we assume that the energy absorbed by the ion rapidly thermalizes over all three spatial dimensions, the decay constant for the total energy of the ion is:

$$\frac{dE}{dt} = -\frac{\gamma_E}{3} E \tag{4.20}$$

The energy of the ion will also be affected when the ion emits the photons it has absorbed. Let us assume that the ion emits each photon in a random direction. Over an interval dt, we would expect the ion's momentum to vary according to:

$$\mathbf{dp} = -\frac{\hbar}{c} \sum_{n=1}^{dN} \left\{ \omega_n \hat{\mathbf{j}}(n) \right\} \tag{4.21}$$

where $\hat{\mathbf{j}}(n)$ are random unit vectors, ω_n is the frequency of the nth photon and dN is the number of emission events in the interval dt. This change in the momentum leads to an energy change:

$$E + dE = \frac{(\mathbf{p} + \mathbf{dp})^2}{2m} \tag{4.22}$$

Using (4.21) and $E = \mathbf{p}^2/(2m)$:

$$dE = -\frac{\hbar}{mc} \sum_{n=1}^{dN} \left\{ \omega_n \mathbf{p} \cdot \hat{\mathbf{j}}(n) \right\} + \frac{\hbar^2}{2mc^2} \left(\sum_{n=1}^{dN} \left\{ \omega_n \hat{\mathbf{j}}(n) \right\} \right)^2 \tag{4.23}$$

This first term involves the sum of the scalar products of a fixed vector with a series of random vectors. This will tend to zero. We can split this remaining summation into two parts:

$$\left(\sum_{n=1}^{dN}\left\{\omega_n\hat{\mathbf{j}}(n)\right\}\right)^2=\sum_{n=1}^{dN}\left\{\omega_n^2\left|\hat{\mathbf{j}}(n)\right|^2\right\}+\sum_{n=1}^{dN}\sum_{m=1}^{dN}\left\{\omega_n\omega_m(1-\delta_{nm})\hat{\mathbf{j}}(n)\cdot\hat{\mathbf{j}}(m)\right\} \tag{4.24}$$

The double summation involves the scalar product of pairs of random vectors. This term must be zero. Therefore:

$$\left(\sum_{n=1}^{dN}\left\{\omega_n\hat{\mathbf{j}}(n)\right\}\right)^2=\sum_{n=1}^{dN}\left\{\omega_n^2\left|\hat{\mathbf{j}}(n)\right|^2\right\}=\sum_{n=1}^{dN}\omega_n^2 \tag{4.25}$$

We can therefore write:

$$dE=\sum_i\left\{\frac{(\hbar\omega_i)^2}{2mc^2}dN_i\right\} \tag{4.26}$$

where dN_i is the number of photons of the ith laser absorbed in the time interval dt. The number of photons the ion absorbs in the time interval is not a constant but itself varies randomly. To take account of this without the apparatus of stochastic calculus, we consider the arrival time of the photons from a statistical point of view. We assume that the photon density of each beam follows a Poissonian distribution so that, in a time period dt, the number of photons the ion will absorb follows the random variable:

$$\hat{N}_i(dt)\sim\text{Poiss}(R_i(v_x)dt) \tag{4.27}$$

Assuming that the timestep dt is large enough that $R_i(v_x)dt>1$, we can expand this[1] by thinking in terms of $R_i(v_x)dt$ timesteps within dt. The probability of absorbing a photon in each of these timesteps is also Poissonian with $\lambda=1$. We can thus write $\hat{N}_i$ in terms of the sum of $R_i(v_x)dt$ independent Poisson distributions $\hat{r}_{i,n}\sim\text{Poiss}(1)$:

$$\hat{N}_i(dt)=\sum_{n=1}^{R_i(v_x)dt}\hat{r}_{i,n} \tag{4.28}$$

[1]This is equivalent to stating that we consider times longer than the mean time to absorb a single photon.

From equation (4.26), we deduce that the spatial distribution of the photon emissions contributes a mean momentum spread per photon:

$$p_{i,r} = \frac{\hbar\omega_i}{c} \tag{4.29}$$

so that the change in momentum over the time period is:

$$d\hat{p}_i = \frac{\hbar\omega_i}{c} \sum_{n=1}^{R_i(v_x)dt} \hat{r}_{i,n} \tag{4.30}$$

The change in energy arising from the ith photon species is:

$$\hat{E} + d\hat{E}_i = \frac{(\hat{p} + d\hat{p}_i)^2}{2m} \tag{4.31}$$

Using (4.30) and $\hat{E} = \hat{p}^2/(2m)$:

$$d\hat{E}_i = \frac{\hbar\omega_i}{mc}\hat{p} \sum_{n=1}^{R_i(v_x)dt} \hat{r}_{i,n} + \frac{(\hbar\omega_i)^2}{2mc^2} \left\{ \sum_{n=1}^{R_i(v_x)dt} \hat{r}_{i,n} \right\}^2 \tag{4.32}$$

The ratio of the magnitude the first term to the second term is approximately:

$$\frac{1}{\hat{N}_i} \left\{ \hat{p} \Big/ \frac{\hbar\omega_i}{c} \right\} \tag{4.33}$$

where the summations have been taken to be of order $\hat{N}_i$. As the ion is cold, its momentum will only be one or two orders of magnitude larger than the photon momentum. At the Doppler limit, the ratio of the momenta is approximately 20. By contrast, the density of photons that interact with the ion will be very large, typically greater than $10^6\,\mathrm{s}^{-1}$. As long as we consider large enough timescales, (4.33) tends to zero and we can neglect the first term in (4.32), so we have:

$$d\hat{E}_i = \frac{(\hbar\omega_i)^2}{2mc^2} \left\{ \sum_{n=1}^{R_i(v_x)dt} \hat{r}_{i,n} \right\}^2 \tag{4.34}$$

To determine the mean spread, we take the expectation value of this:

$$dE_i = \frac{(\hbar\omega_i)^2}{2mc^2} \left\langle \left\{ \sum_{n=1}^{R_i(v_x)dt} \hat{r}_{i,n} \right\}^2 \right\rangle \tag{4.35}$$

We can expand the square of the sum into two terms:

$$dE_i = \frac{(\hbar\omega_i)^2}{2mc^2}\left\langle \sum_{n=1}^{R_i(v_x)dt} \hat{r}_{i,n}^2 + \sum_{n=1}^{R_i(v_x)dt}\sum_{m=1}^{R_i(v_x)dt} \{\hat{r}_{i,n}\hat{r}_{i,m}\,(1-\delta_{nm})\} \right\rangle \tag{4.36}$$

The $\hat{r}_{i,n}$ are uncorrelated so, by definition:

$$\left\langle \sum_{n=1}^{R_i(v_x)dt}\sum_{m=1}^{R_i(v_x)dt} \{\hat{r}_{i,n}\hat{r}_{i,m}\,(1-\delta_{nm})\} \right\rangle = 0 \tag{4.37}$$

We can evaluate the remaining sum by using the properties of the $\hat{r}_{i,n}$. As they are Poissonian, their mean and variance are equal. From the definition of the variance of a distribution:

$$\sigma^2 = \left\langle \hat{r}_{i,n}^2 \right\rangle - \mu^2 \tag{4.38}$$

For each $\hat{r}_{i,n}$:

$$\sigma = \mu = 1 \tag{4.39}$$

so that:

$$\left\langle \hat{r}_{i,n}^2 \right\rangle = 2 \tag{4.40}$$

and:

$$\left\langle \sum_{n=1}^{R_i(v_x)dt} \hat{r}_{i,n}^2 \right\rangle = 2R_i(v_x)dt \tag{4.41}$$

The combination of randomly directed emissions from the ion and fluctuations of the ith laser will therefore cause the ion to heat at the rate:

$$\frac{dE_i}{dt} = \frac{(\hbar\omega_i)^2}{mc^2}R_i(v_x) \tag{4.42}$$

Therefore, the total heating rate associated with fluorescent emission is:

$$\frac{dE}{dt} = \sum_i \left\{ \frac{(\hbar\omega_i)^2}{mc^2}R_i(v_x) \right\} \tag{4.43}$$

There is an additional term that contributes to heating which this derivation has not yet considered. The trap confining the ions will contribute some anomalous heating of its own due, for example, to noise on the trapping potentials arising

from small surface features on the electrodes [AGH+11]. We can allow for this effect by incorporating an extra term ξ in the equation describing the ion heating:

$$\frac{dE}{dt} = \sum_i \left\{ \frac{(\hbar\omega_i)^2}{mc^2} R_i(v_x) \right\} + \xi \tag{4.44}$$

In the steady state, the heating and cooling rates of the ion shown in (4.20) and (4.44) will equilibrate:

$$\frac{\gamma_E}{3} E = \sum_i \left\{ \frac{(\hbar\omega_i)^2}{mc^2} R_i(v_x) \right\} + \xi \tag{4.45}$$

so that the energy of the ion is:

$$E = -\frac{3\hbar \left(\sum_i \left\{ \omega_i^2 R_i(v_x) \right\} + \frac{mc^2}{\hbar^2}\xi \right)}{c \sum_i \left\{ \omega_i b_i \frac{dR_i(v_x)}{dv_x} \right\}} \cdot \tag{4.46}$$

For the approximation that the trap is perfectly harmonic, we can extract a temperature by drawing an analogy to equipartition, noting that a trapped ion in three dimensions has six degrees of freedom:

$$E = 3k_B T \tag{4.47}$$

giving a temperature:

$$T = -\frac{\hbar \left(\sum_i \left\{ \omega_i^2 R_i(v_x) \right\} + \frac{mc^2}{\hbar^2}\xi \right)}{c k_B \sum_i \left\{ \omega_i b_i \frac{dR_i(v_x)}{dv_x} \right\}} \cdot \tag{4.48}$$

For the majority of atomic systems, we would expect the transition rates to be dominated by a single beam, specifically the cooling beam. If this is the case, we can neglect the summations and simplify this expression considerably:

$$T = -\frac{\hbar\omega}{c k_B} \left(\frac{dR(v_x)}{dv_x} \right)^{-1} \left(R(v_x) + \frac{mc^2}{(\hbar\omega)^2}\xi \right) \cdot \tag{4.49}$$

For a negligible trap heating rate, the temperature of such a system varies according to the ratio:

$$T \sim R(v_x) \left(\frac{dR(v_x)}{dv_x} \right)^{-1} \cdot \tag{4.50}$$

The ratio of the fluorescence to the derivative of the fluorescence therefore gives the variation of the temperature of a two-level ion (as described by [Sta03], for example).

4.3.1 The Doppler limit

For the simplest possible case of an unsaturated transition in a two-level atom with negligible heating rate, we can predict a value for the lowest temperature that can be achieved. The scattering rate is a Lorentzian. For convenience, we write this as a function of frequency rather than velocity:

$$R(\omega) = R_0 \frac{\left(\frac{\Gamma}{2}\right)^2}{\left(\frac{\Gamma}{2}\right)^2 + (\omega - \omega_0)^2} \tag{4.51}$$

where R_0 is the peak scattering frequency at ω_0 and Γ is the full-width at half maximum of the transition. According to (4.49), the temperature of the ion varies as:

$$T(\omega) = -\frac{\hbar}{k_B} \frac{\left(\frac{\Gamma}{2}\right)^2 + (\omega - \omega_0)^2}{\omega - \omega_0} \tag{4.52}$$

At:

$$\omega = \omega_0 - \frac{\Gamma}{2} \tag{4.53}$$

the temperature is minimal:

$$T_{\min} = \frac{\hbar\Gamma}{2k_B} \cdot \tag{4.54}$$

This is the often quoted result for the Doppler Limit [Foo05, Ban06, HW05]. For three-level $^{40}Ca^{+}$, the 397nm photons have the highest energy and dominate the cooling process. By considering the ion as a two-state system with $\Gamma = A_{397}$, we can use (4.54) to calculate a limiting temperature for $^{40}Ca^{+}$:

$$T_{\min} = 0.5041\text{mK} \tag{4.55}$$

giving the Doppler limit for $^{40}Ca^{+}$ quoted in (4.7).

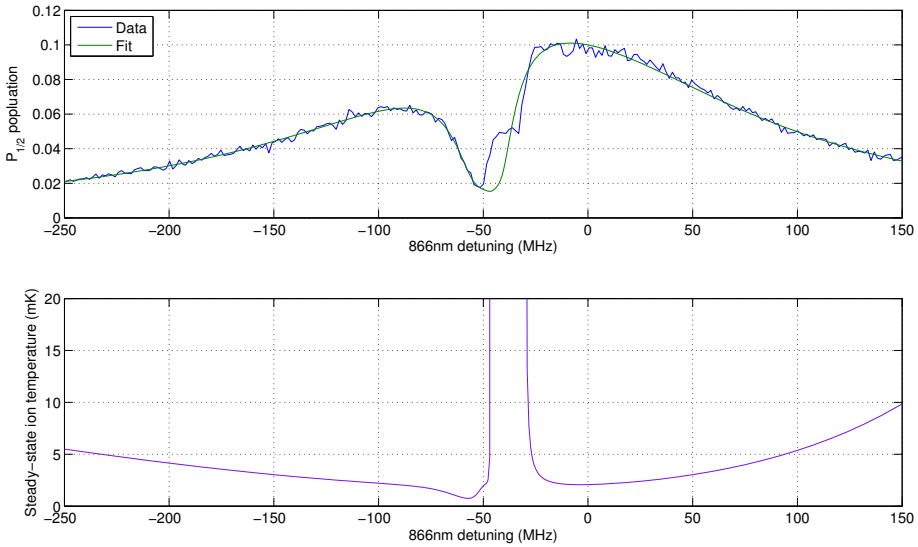

Figure 4.5: The predicted temperatures of the ion over a fluorescence scan using the steady-state method. Around the dark resonance, this method predicts that the ion heats indefinitely; the predicted temperature becomes infinite. If this were the case in the experiment, the ion would escape the trap as soon as these points in the scan were reached. However, there is finite fluorescence in this region, so the ion must still be present in the trap. For more details, see text.

4.3.2 Expanding on the steady state

There is an immediate extension we can make to the steady-state theory to make it more robust without changing the underlying model. The steady-state method adopted a first order Taylor expansion of $R_i(v_x)$ in (4.14). However, the higher order terms may be non-negligible in practice. For example, using only a first order expansion, any regions of the fluorescence curve where the sum of the gradients $\sum_i \left\{ \omega_i b_i \frac{R_i(v_x)}{dv_x}\big|_0 \right\}$ is negative, even if these regions are only a few hertz wide, are predicted to cause the ion to heat indefinitely.

Figure 4.5 shows the predicted temperature of the ion over the central region of the scan with a small heating region shown in figure 4.3. To produce this scan, the frequency of the laser was set to the most negative value and scanned from red to blue. In the centre of the scan, the ion does indeed show heating, and this corresponds to the heating region predicted by the steady-state method. However, this method predicts that the ion will heat *indefinitely* in this region. In

other words, as soon as the ion enters this region, it should become so hot that the ion escapes the trap and is lost. As the ion eventually cools on the blue wing of the scan, this is clearly not the case; the ion instead reaches some large but finite temperature. 'Large' in this context could be anything from a few tens to a few hundred Kelvin.

The steady-state method fails in this region because of the assumption made in (4.14) that the velocity of the ion is small compared to the variation in R. This is, in fact, the same condition described in section 4.1.1 and equation (4.5). When the velocity excursion of the ion is large, we can no longer assume a constant $\sum_i \left\{ \omega_i b_i \frac{R_i(v_x)}{dv_x}\big|_0 \right\}$ over the entire motion of the ion. In order to make a more accurate prediction of the ion's behaviour when it is hot, we need to explicitly consider the oscillations of the ion when calculating the ion's temperature and the regions of the $R_i(v_x)$ that the ion explores during its oscillations. This may be done as follows.

4.4 The excursion method

We keep the same setup as the derivation of the steady-state method in section 4.3. The kinetic energy of the ion's motion along the direction of propagation of the lasers is:

$$E = \frac{1}{2} m v_x^2 \tag{4.56}$$

and its derivative is:

$$\frac{dE}{dt} = v_x \frac{dp_x}{dt} \tag{4.57}$$

where v_x and p_x are the components of the velocity and momentum of the ion along $\hat{x}$ respectively. In a time period dt, the ion will absorb a mean number of photons:

$$dN = \sum_i R_i(v_x) dt \tag{4.58}$$

where the sum is over all of the lasers in the system, as before. The ion will thus experience a change in the magnitude of its momentum:

$$dp_x = \frac{\hbar}{c} \sum_i \{\omega_i b_i R_i(v_x)\}\, dt \cdot \tag{4.59}$$

The change in the energy of the ion is therefore:

$$\frac{dE}{dt} = \frac{\hbar v_x}{c} \sum_i \{\omega_i b_i R_i(v_x)\} \cdot \tag{4.60}$$

However, we no longer make the Taylor expansion we used at the equivalent equation (4.14); we simplify no further.

The calculation of the heating rate in section 4.3 made no assumptions about the time-dependence of the fluorescence curve. As a result, (4.43) is still valid at each instant of time. To determine the energy change of the ion over one period of the ion's motion, we integrate the sum of (4.60) and (4.43):

$$\Delta E = \frac{\hbar}{c} \int_0^{2\pi/\omega_T} \sum_i \left\{ v_x \omega_i b_i R_i(v_x) + \frac{\hbar \omega_i^2}{mc} R_i(v_x) \right\} dt \tag{4.61}$$

where ω_T is the trap's secular frequency along $\hat{x}$. At equilibrium, the change in energy will be exactly balanced by the heating generated by the trap heating rate ξ over one motional period. The value of ΔE will therefore be the negative of this heating rate:

$$\Delta E = -\frac{2\pi}{\omega_T} \xi \tag{4.62}$$

so that:

$$\frac{\hbar}{c} \int_0^{2\pi/\omega_T} \sum_i \left\{ v_x \omega_i b_i R_i(v_x) + \frac{\hbar \omega_i^2}{mc} R(v_x) \right\} dt + \frac{2\pi}{\omega_T} \xi = 0 \cdot \tag{4.63}$$

We will assume that the ion's oscillations are sinusoidal, so that the velocity of the ion varies as:

$$v_x = v_0 \sin(\omega_T t) \cdot \tag{4.64}$$

Hence:

$$\frac{\hbar}{c}\int_0^{2\pi/\omega_T}\sum_i\left\{\omega_i b_i R_i(v_0\sin(\omega_T t))v_0\sin(\omega_T t)+\frac{\hbar\omega_i^2}{mc}R(v_0\sin(\omega_T t))\right\}dt + \frac{2\pi}{\omega_T}\xi = 0\,. \tag{4.65}$$

By solving this integral equation for v_0, we can find the maximum speed of the ion at equilibrium and use this to calculate the temperature of the ion. We will assume as before that the ion has on average equal velocity components along all three spatial dimensions. Making use of equipartition:

$$3k_BT = \frac{1}{2}mv_{\max}^2 = \frac{3}{2}mv_0^2 \tag{4.66}$$

so that the temperature of the ion is:

$$T = \frac{mv_0^2}{2k_B}\,. \tag{4.67}$$

We solve (4.65) by using a standard binary search.

4.4.1 Finite heating

Figure 4.6 compares the temperature predictions for the steady-state and excursion methods for the scan described in subsection 4.3.2. The parameters for the curve were determined by fitting the data removing the region on the high-frequency side of the dark resonance. The curve was fitted using the stationary theory but without modifying the linewidths according to the ion's motion. The excursion method requires no such modification as it treats the oscillations of the ion explicitly. The effect on the fit is that we cannot be confident of the behaviour at the bottom of the dark resonance. However, given that there is clearly heating in this region, we cannot be confident of the predictions of the stationary theory in this region regardless.

At low temperatures, the excursion method reproduces the results of the steady-state method. In the heating region the excursion method makes a more

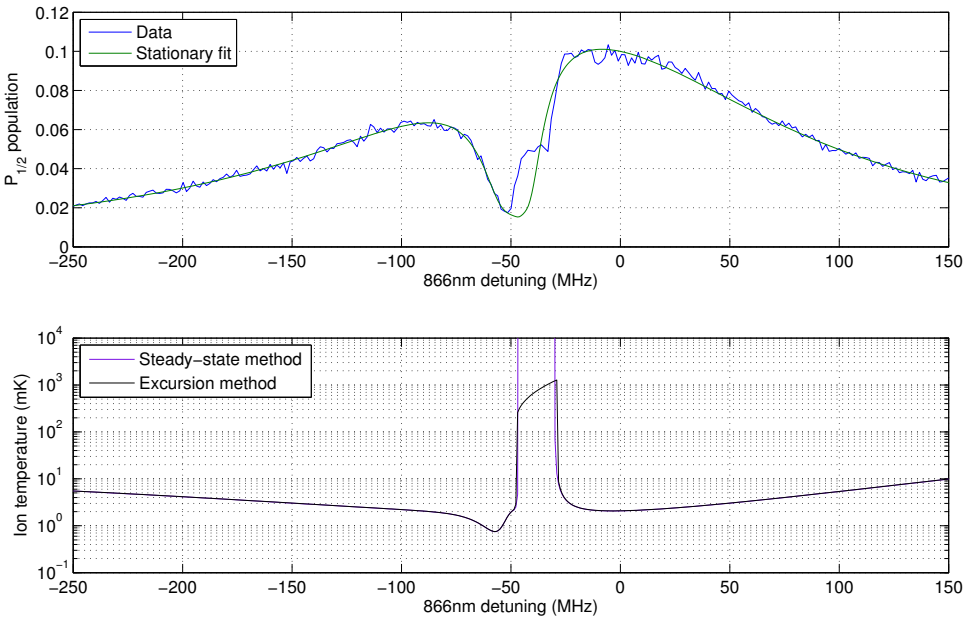

Figure 4.6: The predicted temperature of the ion using the steady-state and excursion methods. The steady-state method predicts infinite heating on the high-frequency side of the dark resonance in contradiction with the data. The excursion method removes this contradiction by predicting a large but finite temperature in this region while tending to the results of the steady-state method when the ion is colder.

robust prediction. Instead of predicting infinite heating, the method predicts a large but finite temperature. The ion will become very hot as it passes through the heating region, but it will still remain trapped.

4.4.2 Fluorescence modification

One immediate consequence of this extension to the steady state method is that we are no longer assuming a stationary ion; rather, we are explicitly treating the ion's motion within the trap. This means that this method can be used to predict the fluorescence of an ion over regions where the stationary theory is inadequate. This model describes the ion in equilibrium as oscillating sinusoidally in the trap. Once the magnitude of the oscillations at equilibrium has been determined from (4.65), this model describes the motion and we can calculate how the fluorescence of the ion varies over one cycle. The average of this will be the signal observed at a detector. Figure 4.7 shows the modified fluorescence pre-

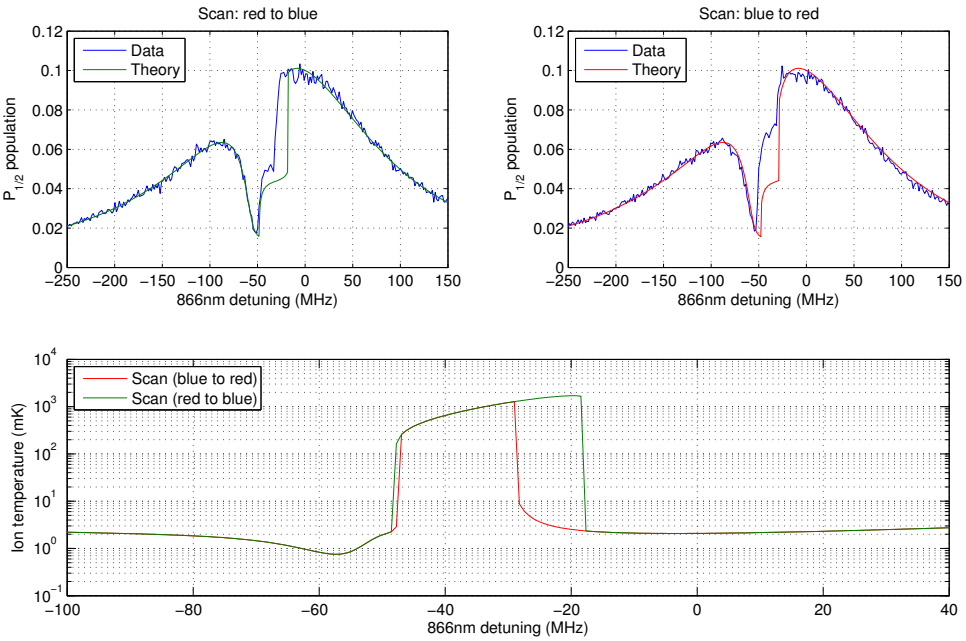

Figure 4.7: A demonstration of the modification to the stationary fluorescence produced by the excursion method. The theory qualitatively matches the data, showing a slanted step in the fluorescence to the high-frequency side of the dark resonance. The excursion effect also predicts hysteresis in this region, although the form this takes does not match the data precisely due to heating effects not present in this model.

dicted by this method, showing a typical slanted 'step' on the high-frequency side of the dark resonance.

This figure also demonstrates another aspect of the underlying physics. The equation being solved is an integral equation, rather than the matrix inversion we used for the steady-state method. This means that there is a possibility for multiple solutions. Fluorescence scans with prominent heating regions are often observed experimentally to have hysteresis effects. The fluorescence over the heating region will be different depending on whether the scan is taken with the 866nm frequency always increasing, from red to blue, or always decreasing, from blue to red. This type of behaviour is predicted by the excursion method.

Figure 4.8 shows how hysteresis can appear in this method. The upper curve shows the value of the integral in (4.65) for a point in the scan shown in figure 4.7 where there are two stable solutions as a function of v_0. The ion will cool when

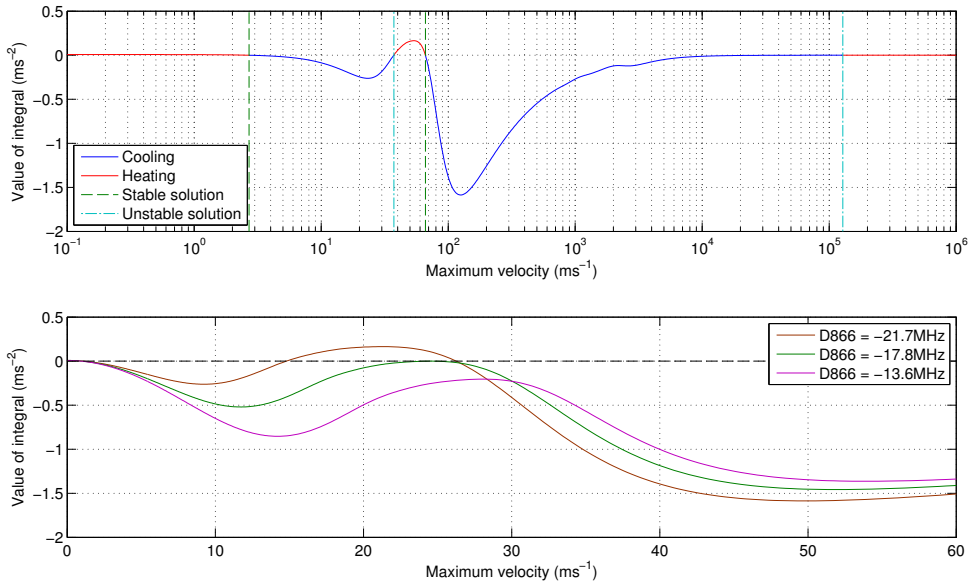

Figure 4.8: A demonstration of multiple solutions predicted by the excursion method. The upper plot shows the value of the integral in (4.65) at an 866nm detuning of -21.7MHz. The two stable solutions are shown. The lower plot shows the behaviour of the function as the 866nm is scanned from this position. There are two stable solutions only when the peak in the function is positive.

the function is negative and heat when the function is positive. Stable solutions occur when the axis is crossed with a negative gradient. The second solution in this system appears because of the positive peak at slightly less than 100ms^{-1}. There will be two solutions for as long as this peak is positive. The height of this peak varies with the detuning of the 866nm. The lower plot shows the behaviour of the integral for three points on the scan: one where there are two solutions, one where there is only a single solution, and the boundary between these two regions.

The existence of two solutions results in hysteresis. The higher-energy solution will only be reached if the ion is in the cooling or heating region immediately surrounding it. If this is not the case, the ion will cool to the lower energy solution (or, if the ion is so hot that it cannot resolve the dark resonance at all, it will heat indefinitely). As the frequency of the 866nm is scanned from red to blue, the ion is already hot at the boundary of the region where multiple solutions appear. As a result, the ion will stabilise at the higher energy solution for

as long as that solution exists. By contrast, when the 866nm is scanned from blue to red, the ion is cold at the boundary of the region and will stabilise to the lower energy solution. The effect of hysteresis is then to change the size of the step in the fluorescence as seen in figure 4.7.

The hysteresis predicted does not match the observations in detail, however. The data has the height of this step changing depending on the scan direction, whilst the theory predicts that only the width of the step will change. This discrepancy is due to the features of heating that are not included in this method, one of which is covered in the following section. Nevertheless, the prediction that hysteresis will occur is valid.

4.4.3 Inadequacies for oscillating ions

The excursion method is an extension of the steady-state method, and can simulate the behaviour of an ion satisfactorily over a larger region of an 866nm scan. However, one feature that it shares with the steady-state method is an assumption about the behaviour of the ion that is not always valid, and can in some cases lead to entirely incorrect predictions of the ion's behaviour. Figure 4.9 shows an example of one such scan, calculated using both the steady-state and excursion methods. This is the scan shown in the top-left of figure 4.3, a curve with a large heating region. The fluorescence calculated using the stationary fluorescence approach is plotted. As with the majority of previous scans, the frequency of the 866nm was scanned from red to blue. Within the heating region, there is a large discrepancy between the stationary theory and the actual fluorescence observed, as we would expect; the stationary theory is not valid in this region. At frequencies to the blue of the red dashed line, the data and simulation come back into agreement, implying that the ion is cold in the blue wing of this scan. However, neither the steady-state nor the excursion methods predict this behaviour. In the blue wing of the scan, both methods predict that

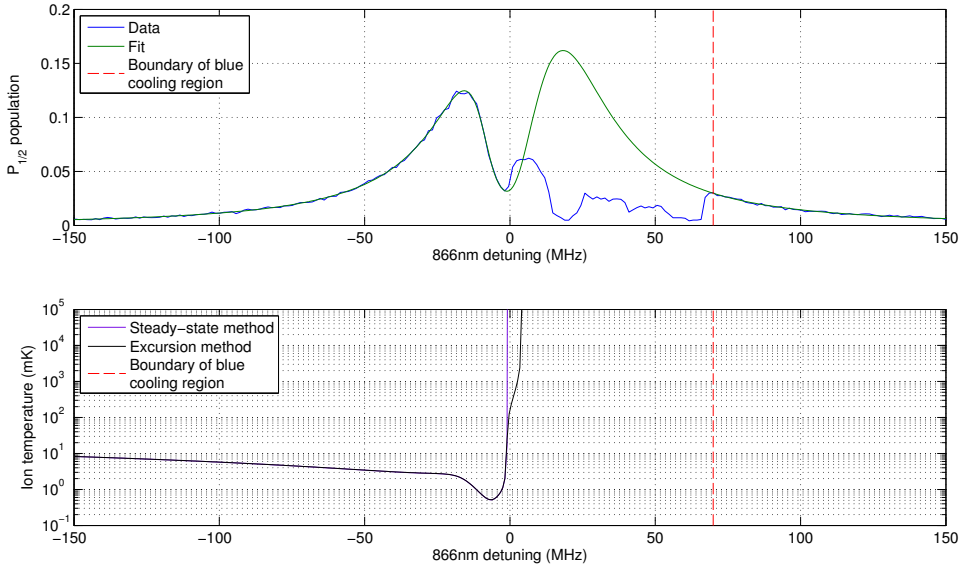

Figure 4.9: The failure of both the steady-state and excursion methods. To the blue wing of the scan, past the red dashed line, the simulation and fit come back into agreement, implying that the ion is cold in this region. However, both temperature models predict that the ion heats indefinitely in this region. In addition, the excursion method fails to predict a finite temperature in the heating region.

the ion heats indefinitely.

This prediction arises from a key assumption in both the steady-state and excursion methods. In both, it is assumed that the ion can adapt to the radiation it experiences on a much faster timescale than the oscillations of the ion in the trap. For the excursion method, we calculate the variation in the fluorescence of the ion over a trap cycle by assuming that the ion is in internal equilibrium at every point, and thus calculate $R(v_x)$ using inversion. However, this may not always be valid. In particular, in the wings of a fluorescence scan, the 866nm beam is far detuned from resonance. Each beam drives population in the system at rates given by the modified Rabi frequency. See section 4.5.1. This frequency is a function of the beam detuning, and increases with an increasing detuning from resonance. In the wings of an 866nm fluorescence scan, the modified Rabi frequencies associated with the 397nm beam will be low, but the corresponding frequencies involving the 866nm beam will be very high; the rate at which pop-

ulation is pumped out of the 866nm transition onto the 397nm transition will be slow. In some cases, this response time may be so slow that it is no longer fast with respect to the oscillation timescale of the motion. When this is the case, we can no longer assume that the ion reaches equilibrium during the cycle.

In an analogous manner to condition (4.5) describing the importance of the excursion effect, we can produce an expression describing the effect of the relative timescales of the ion's internal and external motions. I will refer to the modifying effect of the ion's external motion as the 'dynamic effect'. We characterise the internal timescale of the repumping through the rate of spontaneous transitions on the repumping transition:

$$t(\text{internal}) = \frac{1}{\left\langle P(\mathrm{P}_{1/2}) A(\mathrm{P}_{1/2}, \mathrm{D}_{3/2}) \right\rangle} \tag{4.68}$$

where P refers to the population and A is the Einstein A-coefficient. If the excursion effect is not important, the mean is equivalent to the value of the population at $v_x = 0$. Otherwise, the mean population over the ion's oscillation in the trap must be calculated. The dynamic effect will be negligible as long as the following limit holds:

$$\left\langle P(\mathrm{P}_{1/2}) A(\mathrm{P}_{1/2}, \mathrm{D}_{3/2}) \right\rangle \gtrsim \omega_T \tag{4.69}$$

or, equivalently, that the ion's internal response time is less than its external response time.

The excursion method predicts indefinite heating in the blue wing as the cooling model has the dependence:

$$\frac{dE}{dt} \propto \frac{dR(v_x)}{dv_x} \tag{4.70}$$

even if we allow the derivative to vary over the ion's oscillations. In the blue wing of the scan, the effect of the 866nm will always tend to produce a negative gradient; only in certain situations, such as in the previous scan shown in figure 4.6, will the 397nm counteract this and produce a gradient that is overall positive. If the gradient of $R(v_x)$ is negative throughout the wing, then this

method, in the same manner as the steady-state method, will predict infinite heating throughout the wing.

To better describe a system where the dynamic effect is prominent, I will now describe a model which makes no assumption about the relative timescales of the motion. This is the first model I will describe that explicitly includes the trap secular frequency. In order that the effect of including this behaviour is not obscured by unnecessary complexity, this method will not include the time-dependence of the excursion method; it will be a substantial time-independent modification of the steady-state. The method will assume that, in equilibrium, the ion and its internal dynamics are all oscillating perfectly sinusoidally, but that every internal oscillation has a phase shift relative to the external oscillations. This method is described in the following section.

4.5 The dynamic method

In order to take account of the timescale of the trap motion relative to the timescale of the repumping, we make an extension of the time-independent theory. The differential equation describing the ion's internal dynamics is:

$$\dot{\mathbf{X}} = \mathbf{H}\mathbf{X} \cdot \tag{4.71}$$

Let us write the Liouvillian used in the time-independent theory as $\mathbf{H_S}$ and the solution as $\mathbf{X_S}$. We use the Liouvillian with no modification to the first row of elements, so that:

$$\dot{\mathbf{X}} = \mathbf{H_S}\mathbf{X_S} = \mathbf{0} \cdot \tag{4.72}$$

We now take account of the motion of the ion by making an extension to the Liouvillian. We assume that the ion is oscillating sinusoidally in the trap, and therefore experiencing Doppler shifts of the lasers that vary sinusoidally with the trap frequency:

$$\mathbf{H} \rightarrow \mathbf{H_S} + \sin(\omega_T t)\mathbf{H_D} \tag{4.73}$$

where ω_T is the trap frequency and $\mathbf{H_D}$ describes the Doppler shifts of the lasers. We can write this as:

$$\mathbf{H_D} = \frac{v_{\max}}{c} \sum_i \{b_i \omega_i \mathbf{H_{L,i}}\} \cdot \tag{4.74}$$

$v_{\max}$ is the maximum speed of the ion. The $\mathbf{H_{L,i}}$ have elements either zero or one with the non-zero elements corresponding to the elements of the Liouvillian describing the detuning of the ith laser. b_i describes the direction of propagation of the ith laser, as before. For $^{40}Ca^{+}$, we can write this expression in terms of a characteristic maximum Doppler-shift of the 397nm δ_{397} rather than the maximum velocity of the oscillations, so that:

$$\mathbf{H_D} = \delta_{397} \left(\mathbf{H_{L,397}} \pm \frac{\omega_{397}}{\omega_{866}} \mathbf{H_{L,866}} \right) \cdot \tag{4.75}$$

The sign of the second term is positive if the two beams are co-propagating and negative if they are counter-propagating. We take δ_{397} to be 1MHz, corresponding to a two-level temperature of order the Doppler limit. This term acts only as a reference value and does not alter the final temperature calculated.

We additionally assume that the populations and coherences will also oscillate perfectly sinusoidally. This essentially assumes that each coherence as a function of the Doppler shift has a constant derivative for the regions of the curve explored by the motion. This will be a good approximation for small excursions, or for parts of the fluorescence curve where second and higher order derivatives are negligible. We assume:

$$\mathbf{X} \rightarrow \mathbf{X_S} + \sin(\omega_T t)\mathbf{U} + \cos(\omega_T t)\mathbf{V} \tag{4.76}$$

where $\mathbf{U}$ and $\mathbf{V}$ describe the oscillations that are in phase and in quadrature with the ion's speed oscillations respectively. For the purposes of cooling, the term of interest is $\mathbf{U}$; the in quadrature oscillations will, by symmetry, have no effect on the net energy change over a single cycle. The differential equation

(4.72) is now:

$$\begin{aligned}&\omega_T \left\{\cos(\omega_T t)\mathbf{U} - \sin(\omega_T t)\mathbf{V}\right\} = \\ &\left\{\mathbf{H_S} + \sin(\omega_T t)\mathbf{H_D}\right\}\left\{\mathbf{X_S} + \sin(\omega_T t)\mathbf{U} + \cos(\omega_T t)\mathbf{V}\right\} \cdot\end{aligned} \tag{4.77}$$

We assume that $\mathbf{U}$ and $\mathbf{V}$, along with the Doppler shifts described by $\mathbf{H_D}$, are small. As such, we ignore second-order terms in these variables:

$$\begin{aligned}&\omega_T \left\{\cos(\omega_T t)\mathbf{U} - \sin(\omega_T t)\mathbf{V}\right\} = \\ &\mathbf{H_D X_S}\sin(\omega_T t) + \mathbf{H_S U}\sin(\omega_T t) + \mathbf{H_S V}\cos(\omega_T t) \cdot\end{aligned} \tag{4.78}$$

This must be true for all time, so we can equate the sine and cosine terms:

$$\omega_T \mathbf{U} = \mathbf{H_S V} \tag{4.79}$$

$$-\omega_T \mathbf{V} = \mathbf{H_D X_S} + \mathbf{H_S U} \cdot \tag{4.80}$$

We remove the explicit dependence on the trap frequency by defining:

$$\mathbf{L} = \frac{1}{\omega_T}\mathbf{H_S} \tag{4.81}$$

$$\mathbf{L_D} = \frac{1}{\omega_T}\mathbf{H_D} \tag{4.82}$$

and solve for $\mathbf{U}$ and $\mathbf{V}$:

$$\mathbf{V} = -(\mathbf{I} + \mathbf{L}^2)^{-1}\mathbf{L_D X_S} \tag{4.83}$$

$$\mathbf{U} = \mathbf{L V} \cdot \tag{4.84}$$

For a given maximum speed, $\mathbf{U}$ includes the amplitudes of the population oscillations. From these, we can determine how the fluorescence varies with the ion's speed. If U_i is the amplitude of the oscillations of the total population P_i of the upper level of the ith transition, and A_i is the Einstein-A coefficient of the transition, then:

$$\frac{dR_i(v_x)}{dv_x} = \frac{A_i U_i}{v_{\max}} \tag{4.85}$$

where we have assumed:

$$R_i(v_x) = A_i P_i \cdot \tag{4.86}$$

We can use this and (4.49) to produce a new expression for the temperature:

$$T = \frac{\hbar v_{\max}\left(\sum_i \{\omega_i^2 R_i(v_x)\} + \frac{mc^2}{\hbar^2}\xi\right)}{ck_B \sum_i \{\omega_i b_i A_i U_i\}} \cdot \tag{4.87}$$

For three-level $^{40}Ca^{+}$, we have that $A_{397} \gg A_{866}$ and therefore that $R_{397} \gg R_{866}$. We will neglect the terms in R_{866} and A_{866} and write the temperature as:

$$T = \frac{\hbar\delta_{397}}{k_B A_{397} U_{397}}\left(R_{397} + \frac{mc^2}{\hbar\omega_{397}^2}\xi\right) \cdot \tag{4.88}$$

where U_{397} refers to the magnitude of the in-phase oscillations of the $P_{1/2}$ population and R_{397} refers to the scattering rate on the 397nm transition. This method requires one less matrix inversion than the steady-state method and will calculate on a comparable timescale.

4.5.1 The dynamic effect in a two-state system

The dynamic effect can be observed in any trapped ion system. All that is required is that there are two separate timescales in the problem, that of the atomic dynamics and an external timescale provided by the ion's motion in the trap. The experiments in $^{40}Ca^{+}$ we have been considering thus far involve two different laser beams and eight different states. Transitions between each level and each state will have their own separate timescales, all of which play a role. Before tackling that system, it will be informative to examine the much simpler two-state system. This system allows us to unambiguously examine the modifications produced by the dynamic effect on a single transition.

The two-state system has five parameters that describe its internal and external timescales. These are the spontaneous emission on the transition A, the Rabi frequency $\Omega \equiv 2R$, the detuning of the laser beam Δ, the beam's linewidth Γ and the frequency of the ion's secular motion ω_T. The timescale of the ion's external oscillations is described by ω_T. The ion's internal dynamics, however, are governed by two separate rates: the rate of spontaneous emission and the rate at which the ion responds to the radiation field.

Consider a two-state system in equilibrium. The rate of upward transitions and downward transitions must be equal. The upward stimulated transitions will be exactly matched by downward stimulated and spontaneous transitions. We now abruptly change the frequency of the driving beam and observe the population in the two states as a function of time until a new equilibrium is established. As soon as the frequency of the beam changes, the upward stimulated rate will no longer match the rate of downward transitions; there will be Rabi oscillations and relaxation. The oscillations occur on the timescale of the ion's response to the radiation field. This is given by the modified Rabi frequency $\widetilde{\Omega}$, the natural frequency of the combined ion/photon system. This is in general a function of all the atomic parameters. For a system with $A, \Gamma \ll \Omega$, the modified Rabi frequency is well known (see, for example, [Foo05, Ban06, Sho11]):

$$\widetilde{\Omega} = \sqrt{\Omega^2 + \Delta^2} \tag{4.89}$$

We examine this timescale in more detail in section 4.5.1.1. However, these oscillations are transient and simply shunt population backwards and forwards between the two states. The approach to equilibrium is governed by spontaneous emission, which changes the relative population of the states until a new equilibrium is established. There are therefore two timescales describing the ion's internal motion: a fast timescale given by $\widetilde{\Omega}$ and a slow timescale given by A.

We infer the temperature of the ion from the population oscillations of the higher energy state after dynamic equilibrium is reached. There are two different measures we can use to determine this temperature. As shown in (4.85) and (4.87), we can use the amplitude of the component of the oscillations that is in phase with ion's motion, U. If U is positive, the population oscillations are in phase with the oscillations of the beam's detuning; they are therefore in antiphase with the ion's velocity oscillations, so the ion will absorb more photons when moving towards the beam source than away from it. If U is positive, then,

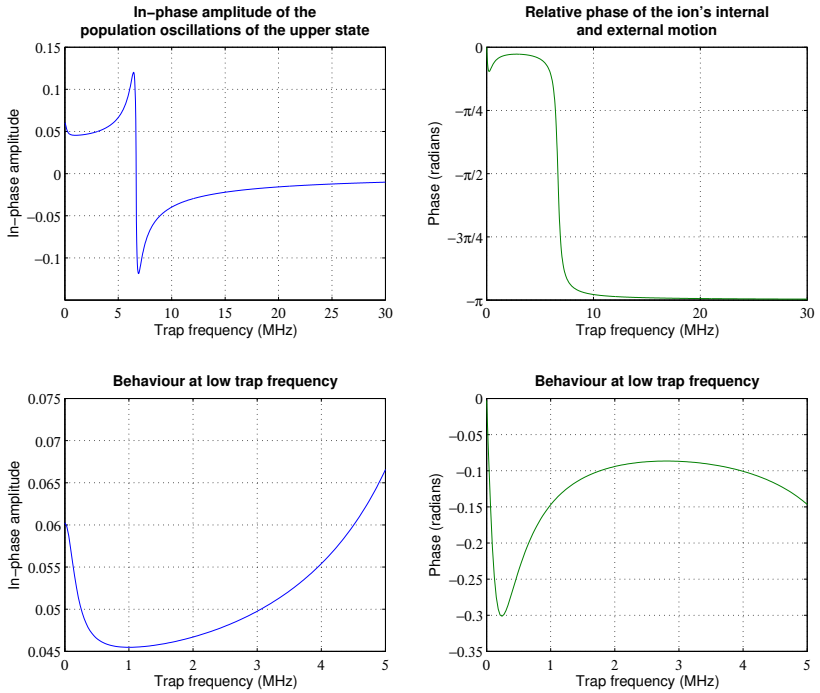

Figure 4.10: The response of a two state system to external harmonic oscillations as a function of the trap secular frequency ω_T. The lefthand curves show U, the in-phase amplitude of the population oscillations of the higher energy state. The righthand curves show ϕ, the relative phase between the ion's internal and external oscillations. The upper curves show the behaviour of U and ϕ at large trap frequency, while the lower curves show the behaviour at low trap frequency. The system responds to the trap oscillations on two different frequency scales. At trap frequencies of order 1MHz, U is damped out while ϕ decreases sharply. At trap frequencies of order 10MHz, resonant behaviour can be seen centred around $\omega_T = 6$MHz. For this system, $\Omega = 20$MHz, $A = 2$MHz, $\Delta = -2$MHz and $\Gamma = 0$MHz. For more details, see text.

the ion will cool, and if it is negative, the ion will heat. The more positive the value of U, the lower the ion's temperature. We can also use U and V to directly calculate the phase of the motion ϕ relative to the oscillations of the detuning:

$$\phi = \text{atan2}(V, U) \tag{4.90}$$

The component of the motion in phase with the detuning ($\phi = 0$) cools the ion, while the component in anti-phase ($\phi = \pm\pi$) heats the ion. In order for the net effect on the ion to be that of cooling, $-\pi/2 < \phi < \pi/2$.

Figure 4.10 shows the response of the ion to external motion as a function of ω_T using the two measures U and ϕ. At zero trap frequency, the behaviour of the ion precisely matches the behaviour predicted by the steady-state theory. In

this limit, the ion can always perfectly track the external motion. As a result, the ion's population oscillations are exactly in phase with its external motion, and their magnitude is given by an ion with a temperature equal to the steady-state prediction.

For non-zero trap frequency, there are two regimes of behaviour corresponding to the two internal timescales of the system. One regime occurs for small trap frequencies comparable with A and the second for larger trap frequencies comparable with $\widetilde{\Omega}$.

For $\omega_T \gtrsim A$, the behaviour is dominated by a characteristic resonant effect. This resonance is centred around $\omega_T = \widetilde{\Omega}$, and occurs due to the ion's external oscillations resonantly driving the ion/photon system. The behaviour of the ion in this regime is similar to a driven harmonic oscillator, with A and Γ acting as damping terms due to their effect of reducing the relative strength of the stimulated transitions compared to spontaneous transitions. As the trap frequency becomes very large compared to this resonant frequency, U tends to zero; the external oscillations are now far too fast for the ion to track, so its internal oscillations are completely damped out.

The system deviates from simple harmonic behaviour at low trap frequencies $\omega_T \lesssim A$. The harmonic oscillator we use to describe behaviour at higher trap frequencies is a one-dimensional problem. The population of the upper state is, however, not a one-dimensional system; it does not uniquely define a point on the Bloch sphere. For example, we can have the same population with different combinations of detuning and Rabi frequency.

Let us consider a point P on the Bloch sphere which represents the ion at rest in equilibrium with the radiation. We now slowly change Δ. The point will trace out a line to some new point Q. By reversing the change in Δ, we bring the ion back along the same line. Now imagine speeding up this movement. The ion will not be able to keep up with the variation in Δ and will not reach Q before

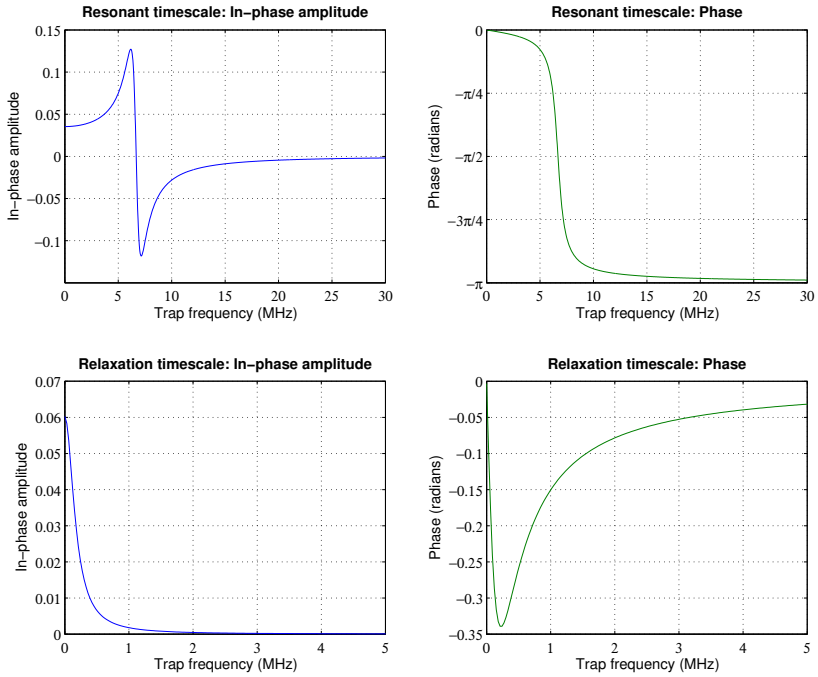

Figure 4.11: The contributions to the response of a two-state system from its two characteristic internal timescales. The upper two plots show the behaviour of a damped harmonic oscillator with parameters fitted to the curves shown in figure 4.10. Over the resonance, the in-phase amplitude changes sign and the phase of the motion changes sharply, qualitatively matching the behaviour of the two-state system. The lower two plots show the behaviour of a two-state system with its $\widetilde{\Omega}$ resonance removed. These plots were generated by taking the output of a standard two-state system with $\widetilde{\Omega} \gg 5\text{MHz}$ and scaling U to match the limits in figure 4.10. The in-phase amplitude drops exponentially as a function of the trap frequency. The phase drops to a minimum at low trap frequency and then tends to zero as the trap frequency increases. By summing these two contributions, the behaviour of seen in figure 4.10 can be approximated.

the change reverses. There is a 'treacle' effect that reduces the amplitude of the ion's oscillations and increases their phase relative to the original oscillations.

This effect is damped out at higher trap frequencies. In a sense, we have two different degrees of freedom in this system which occur on two different timescales: there is the treacle effect for small trap frequencies and the oscillating behaviour at high trap frequencies. Figure 4.11 shows how the response shown in figure 4.10 can be interpreted in terms of the separate contributions to the behaviour of the ion from its two internal timescales.

4.5.1.1 Calculating $\widetilde{\Omega}$

The behaviour of the ion around the resonance $\omega_T = \widetilde{\Omega}$ can be used to directly infer the value of the resonant frequency of the transition. At the resonant trap frequency, $U = 0$ and the function changes sign on either side of the resonance. This zero-crossing is used to identify $\widetilde{\Omega}$. This is true for systems with any number of states; for a given pair of states, the value of U for the individual population oscillations will change sign over the resonant frequency of the transition between them.

For the two-state system, we can use this approach to derive a a more general expression for $\widetilde{\Omega}$ that will be valid for all values of Γ, A, Ω. Expression (4.89) is only valid for $A, \Gamma \ll \Omega$. We produce an analytic expression for $\widetilde{\Omega}$ by solving the Bloch equations analytically using the dynamic method and finding the value of ω_T for which $U = 0$. The expression for non-negligible Γ is long and covered in Appendix C along with other analytic expressions for the two-state system. The natural frequency of a two-state system for $\Gamma \ll A, \Omega$ is:

$$\widetilde{\Omega} = \frac{1}{4}\sqrt{2\sqrt{\left(5A^2 + 4\Delta^2 + 4\Omega^2\right)^2 - 32A^2\Omega^2} - 6A^2 + 8\Delta^2 + 8\Omega^2} \tag{4.91}$$

For a system with $\Omega \ll A$, this simplifies to:

$$\widetilde{\Omega} = \sqrt{\left(\frac{A}{2}\right)^2 + \Delta^2} \tag{4.92}$$

Note that this is an almost identical form to (4.89), except Ω is replaced by $A/2$. This arises from the definitions of A and Ω used in this thesis. A is defined as the rate of change of a state's population while Ω is defined as the rate of change of its amplitude. This introduces a factor of two in the Bloch equations for every term containing A (and Γ) compared to terms containing Ω.

4.5.1.2 Cooling using the internal motion

We can see some of the implications of the dynamic effect for cooling a two-state system by examining its behaviour as a function not only of the trap frequency

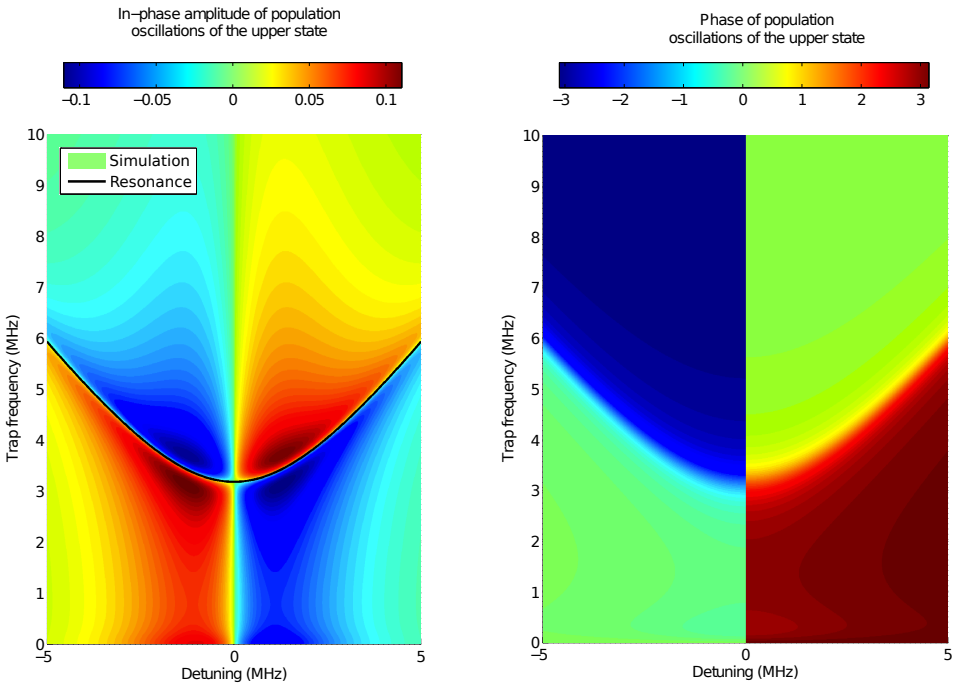

Figure 4.12: The response of a two-state system to external trap oscillations. For this system, $\Gamma \ll \Omega$ and $A < \Omega$. The oscillations of the population of the upper state are examined. The lefthand plot shows U and the righthand plot shows ϕ. The resonance with $\widetilde{\Omega}$ is clearly seen as one arc of a hyperbola centred about the trap frequency axis and is marked on the lefthand plot in black. The parameters used to produce this simulation are $\Omega = 20 \times 10^6 \mathrm{s}^{-1}$, $A = 2.5 \times 10^6 \mathrm{s}^{-1}$ and $\Gamma = 0$. 'Simulation' in the legend refers to the surface described by the colour bar.

but also of the detuning Δ. Figure 4.12 shows the behaviour of the two temperature measures U and ϕ for the two-state system examined in figure 4.10. This figure shows two resonant effects. The resonance at $\omega_T = \widetilde{\Omega}$ is seen as one arc of a hyperbola centred on the trap frequency axis. Another resonant effect occurs about the detuning axis. This, however, doesn't require the dynamic effect to explain. It is a prediction of the steady-state theory, and the same mechanism described in section 4.1.1 that causes heating to the high-frequency side of the peak of 397nm scans in $^{40}Ca^{+}$.

One interesting consequence of the hyperbolic resonance is that it can be used to cool. For a system with light damping, so that $2(A + \Gamma) \lesssim \Omega$, the amplitude of the in-phase oscillations increases in magnitude near the resonance. The temperature is inversely proportional to this amplitude, so there will be a

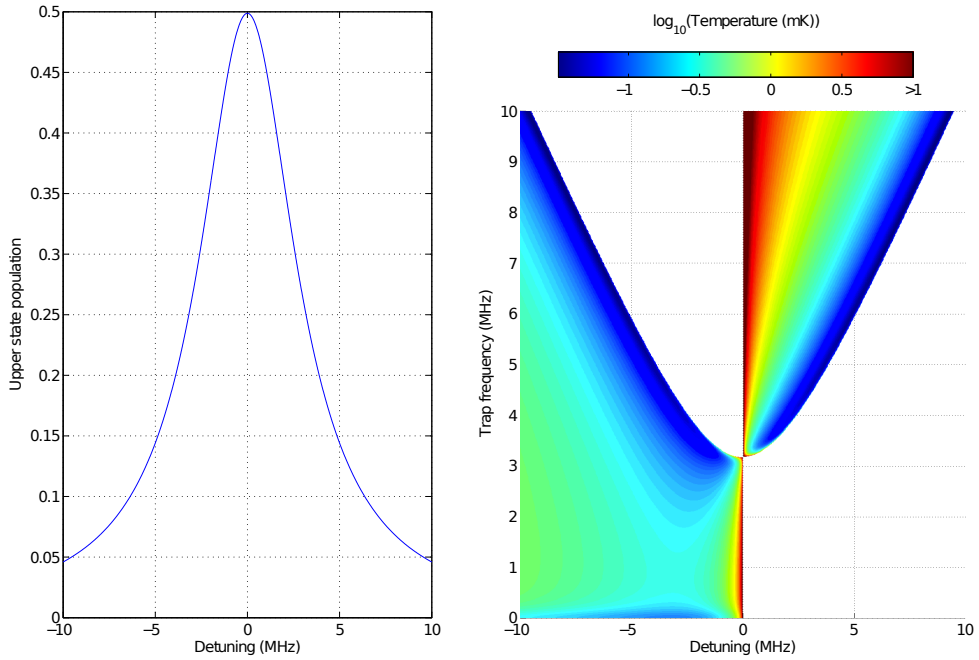

Figure 4.13: A demonstration of optimal cooling around the dynamic resonance for a system with light damping. The lefthand plot shows the fluorescence profile of this system. Because the system is lightly damped, the population of the upper state on resonance is approximately 0.5. The righthand plot shows the temperature of the system as a function of ω_T and Δ. Blank regions show where the ion is predicted to heat. For any fixed value of Δ, the ion's temperature is minimal at some trap frequency close to the $\widetilde{\Omega}$ resonance. This is true even for $\Delta > 0$, a region in which the ion is predicted to heat according to the steady-state theory. For $\Delta < 0$, the optimal temperature is a significant improvement over the temperature at $\omega_T = 0$. For the value of the detuning giving the optimal temperature according to the steady-state theory, the optimal temperature on resonance is an improvement by a factor of 3.5 compared to the steady-state prediction. The parameters used for this simulation are $A, \Gamma = 10^6 \mathrm{s}^{-1}, \Omega = 20 \times 10^6 \mathrm{s}^{-1}$.

point close to the resonance where the temperature is optimal. This will occur regardless of the sign of the amplitude at $\omega_T = 0$, so the resonance can be used to cool at any value of Δ, even if the ion is predicted to heat according to the steady-state theory. For a demonstration of this, see figure 4.13.

Making use of the dynamic resonance can make significantly lower temperatures available compared to the results of the steady-state theory. The amount by which the temperature can be lowered depends on the value of the damping parameters A and Γ. The lower these parameters compared to Ω, the more the optimal temperature will be reduced compared to the steady-state results. To allow the best comparison with the steady-state results, we quantify the tem-

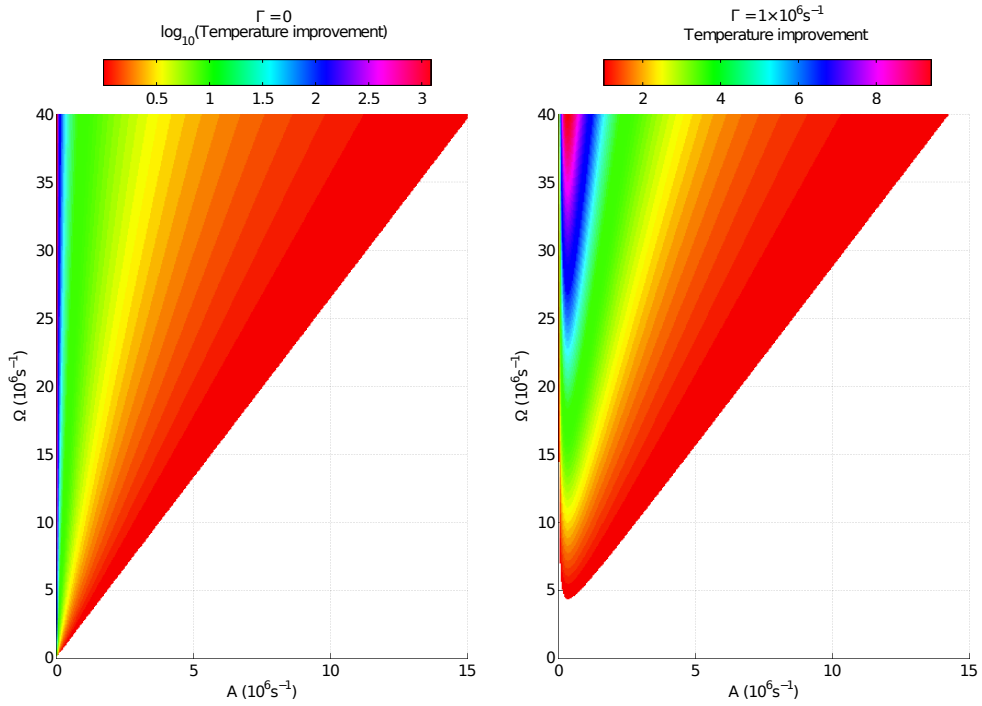

Figure 4.14: A demonstration of the factor Θ by which the optimum temperature of the two-state system can be improved through making use of the dynamic resonance. Both plots show this factor as a function of A and Ω. The lefthand plot shows a simulation with $\Gamma = 0$ and the righthand plot a simulation with $\Gamma = 1 \times 10^6 \mathrm{s}^{-1}$. For a negligible linewidth, Θ can be several orders of magnitude for small A. For an appreciable linewidth, Θ can still approach an order of magnitude for small A. In both cases, the temperature can be lowered by a significant amount even for relatively large values of A.

perature improvement by looking at the function $T(\omega_T, D)$. D is the detuning that minimises the temperature for $\omega_T = 0$. We then define the 'temperature improvement factor' Θ as the ratio of the minimum of this function to the value at zero trap frequency:

$$\Theta = \frac{\min T(\omega_T, D)}{T(0, D)} \,. \tag{4.93}$$

This factor can be large even for dampings that are not at all negligible. For example, for the system shown in figure 4.13, $\Theta = 3.5$. Figure 4.14 shows how this factor varies with A, Γ and Ω. For small values of A and Γ, improvements of several orders of magnitude are possible. Even for dampings as large as $2(A + \Gamma) = \Omega/2$, the temperature is still improved by an appreciable amount. For this example, $\Theta \sim 2$. Note that this means that there is a wide range of parameters

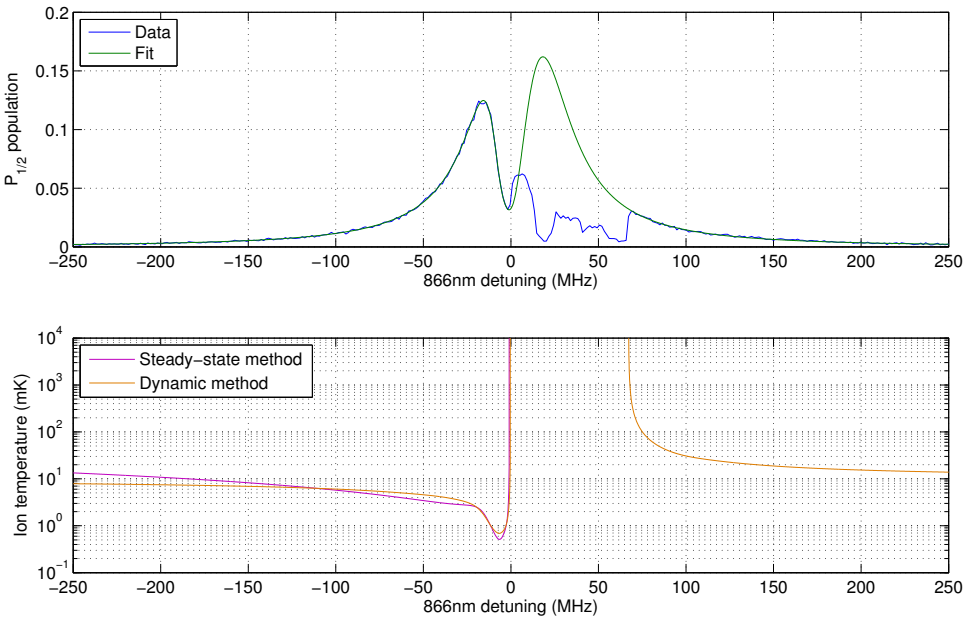

Figure 4.15: A comparison of the temperature predictions of the steady-state and dynamic methods. The dynamic method predicts cooling not only in the red wing but also in the blue wing, in agreement with the data. In addition, the temperatures at large detunings do not rise indefinitely as in the steady state method but converge asymptotically. However, in assuming sinusoidal oscillations, the dynamic method assumes that the ion will always see a fixed gradient and does not allow it to explore the structure of the fluorescence curve. Therefore, it incorrectly predicts infinite heating on the blue side of the dark resonance.

for which it is possible to Doppler cool the two-state system below the Doppler limit.

4.5.2 The dynamic effect in $^{40}Ca^{+}$

Figure 4.15 shows the predictions of the dynamic method for the scan described in section 4.4.3. In this and subsequent plots, the trap frequency used is 816kHz unless otherwise specified. This was the value measured in the trap from which these scans were taken [Web05]. The calculation predicts that the ion will cool in the blue wing of the scan, in agreement with the data. The temperatures in the wing are high due to the small negative value of the 397nm detuning. This also results in the large heating region in the centre of the scan.

The temperatures predicted in the blue wing and the red wing converge

asymptotically to the same value at very large positive and negative detunings. One expects on general grounds that the effect on the ion should be symmetric for very large positive and negative detunings of the repumper. Large 866nm detunings greatly exceed the Zeeman splitting of the $D_{3/2}$ level, and the rempumping rates for positive and negative detunings converge. The repumping rate scales monotonically with the magnitude of the 866nm detuning, directly controlling the 397nm transition rate. Both R_{397} and U_{397} scale in the same manner, resulting in a temperature that is asymptotically independent of the 866nm detuning. This behaviour is not demonstrated by the steady-state and excursion methods. These models both predict that the ion's temperature will increase without limit as the 866nm detuning increases. This is in stark contrast to experimental data.

The dynamic method describes the experiment significantly better than the steady-state method without a significant increase in calculation time. It requires the inversion of two matrices of the same size as the Liouvillian, although one is significantly less sparse as it involves the square of $\mathbf{H}$. The steady-state method requires the inversion of three matrices. As a result, depending on the size and sparseness of the Liouvillian, the dynamic method has calculation times of the same order as the calculation times for the steady-state method, such as for $^{40}Ca^{+}$, or a few times slower for more complex systems, such as for $^{43}Ca^{+}$.

Figure 4.16 demonstrates the portions of the scan shown in figure 4.15 where the dynamic effect is prominent. The location of these regions is calculated using the limit (4.69). For a trap frequency of 816kHz, the dynamic effect is negligible only for a very small portion of the scan where the $P_{1/2}$ population is large and the temperature low. For the vast majority of the scan, the dynamic effect cannot be ignored. The figure also shows how the dynamic effect allows for a finite temperature at large positive 866nm detunings when the steady-state and excursion methods do not. The finite response time of the repumper introduces a phase

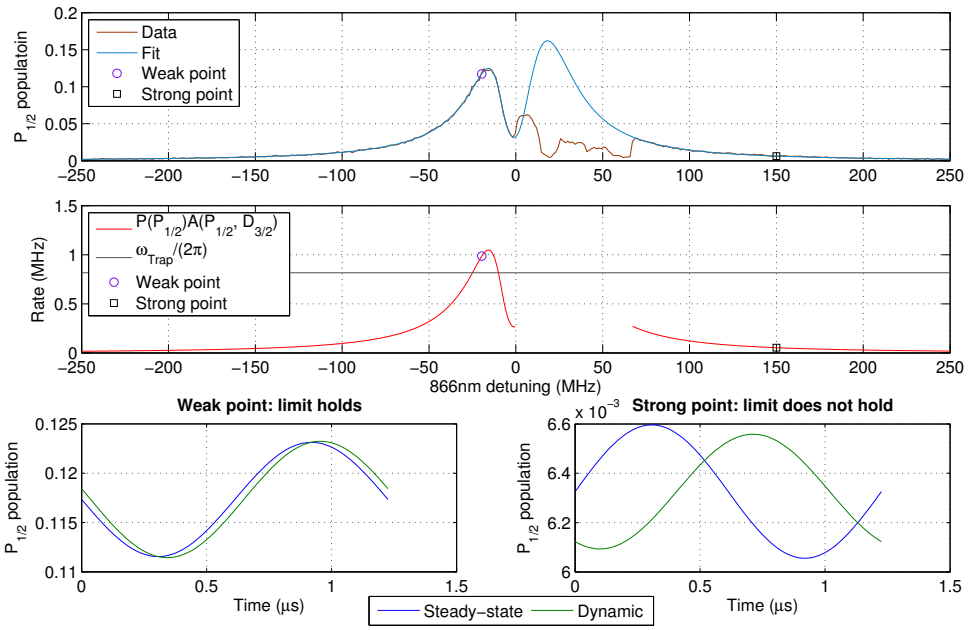

Figure 4.16: A demonstration of the regions in which the dynamic effect is important. The upper plot shows the data and fit. The central plot shows the characteristic internal and external timescales. Only those points where the excursion effect is negligible are shown. When the internal timescale is faster than the external timescale, the dynamic effect will be negligible. This is only true for a very small portion of the scan. The lower two plots show the variation in the $P_{1/2}$ population over one period of the ion's motion using either the steady-state or the dynamic method at the two points highlighted on the upper plots. The steady-state curves are calculated using the temperature given by the dynamic method. The amplitudes of each pair of curves differ slightly due to the dynamic effect. In the lefthand plot, the limit (4.69) holds and the dynamic effect modifies the steady-state result by a very small amount. By contrast, the limit does not hold in the righthand plot. The dynamic effect produces a significant phase lag in the response of the fluorescing population to the oscillating laser frequencies.

lag in the oscillations of the $P_{1/2}$ relative to the ion's external oscillations. Without considering the dynamic effect, a negative gradient dR/dv_x leads to more photons being absorbed when the ion is moving away from the 397nm beam, leading to heating. The lag in the ion's response caused by the dynamic effect shifts the absorption curve in time so that the point at which most photons are absorbed now occurs when the ion is moving towards the 397nm beam, allowing for an equilibrium temperature to be reached. Note that, calculated using the dynamic method, the absolute phase of the $P_{1/2}$ population oscillations is similar in both cases shown. The reason for this will be covered shortly.

Figure 4.17 demonstrates how the strength of the dynamic effect scales with

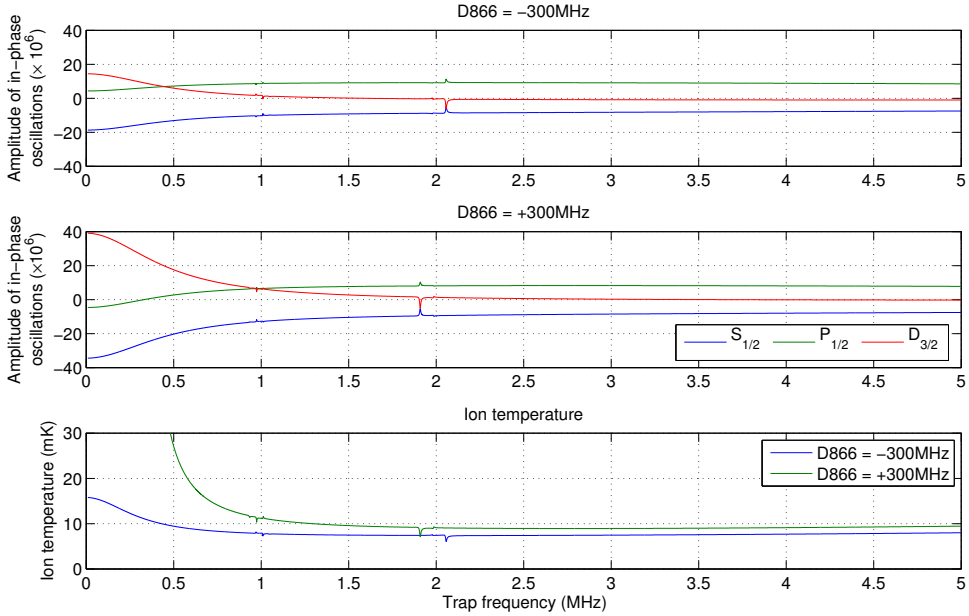

Figure 4.17: A demonstration of how cooling occurs at high 866nm detunings. The upper two plots show the amplitudes of the in-phase oscillations of the populations of the three levels at a large red detuning of -300MHz and a large blue detuning of +300MHz. The lower plot shows the temperature of the ion for these detunings. At zero trap frequency, the dynamic method reproduces the results of the steady-state method; the ion is cold at -300MHz and infinitely hot at +300MHz. However, as the trap frequency increases, the oscillations in the $D_{3/2}$ quickly die to almost zero for both cases. The effect of the 866nm becomes negligible for all but small trap frequencies, and the variations in the populations converge to similar values. This means that the temperatures of the ion at both points also converge. Note that there are small resonance effects in both the oscillations and the temperature curves. These are the resonances at $\omega_T = \widetilde{\Omega}$ for Raman transitions between $D_{3/2}$ states. For more details, see text.

the trap frequency. The temperature measure used is U. The figure shows the variation of the in-phase oscillations of the level populations U_S, U_P and U_D as a function of the trap frequency, and the resulting temperature, for two values of the 866nm detuning taken symmetrically about the centre of the profile. The values are -300MHz and +300MHz. These detunings have been chosen to be large so that no complications will arise from proximity to the dark resonance. As for the two-state system, U_P must be positive for the ion to cool. The greater the amplitude of U_P, the greater the cooling.

This experiment involves three levels in $^{40}Ca^{+}$. In the previous section, we looked at a two-state system and observed the effect on the ion when the dy-

namic effect was prominent. The results of that analysis are little changed when extrapolated to a two-level system. However, a three-level system can show significantly different dynamics. Firstly, there are now two different beams each with an associated timescale, and the dynamic effect can be prominent on either. The characteristic timescale of the 397nm transition is of the order $10^8\mathrm{s}^{-1}$ and that of the 866nm transition of the order $10^6\mathrm{s}^{-1}$. In this study, I have assumed that trap frequencies will be of order 1MHz, so that the dynamic effect will only ever be observed on the 866nm transition and not on the 397nm transition.

To determine the mechanisms by which the dynamic effect on the 866nm transition alters the behaviour of the ion, let us consider two systems: a two-level system comprising only of the $\mathrm{S}_{1/2}$ and $\mathrm{P}_{1/2}$ and the full three-level system. The presence of the third level alters the dynamics compared to the two-level system in two distinct ways.

Firstly, the presence of the $\mathrm{D}_{3/2}$ level will syphon some population out of the 397nm transition; the $\mathrm{P}_{1/2}$ states have transitions to two levels rather than one. The $\mathrm{D}_{3/2}$ states will share in some of the in-phase population oscillations of the $\mathrm{P}_{1/2}$ states, so that U_P is necessarily reduced. In this capacity, oscillations in the $\mathrm{D}_{3/2}$ will always reduce the cooling rate of the system compared to an equivalent two-level system.

The second contribution of the $\mathrm{D}_{3/2}$ level arises from the oscillations in the detuning of the repumper. The 866nm beam contributes directly to the cooling only negligibly. However, it directly controls the transition rate on the 397nm transition, and 397nm photons *do* contribute to the cooling. The 397nm beam is red-detuned, and therefore this beam will always cool the ion more than it heats it. As the ion moves towards the 397nm beam, the 397nm beam is shifted closer to resonance and more absorption will occur. In the steady-state and excursion methods, the 866nm beam will only show the same behaviour when it too is red detuned and will demonstrate the opposite behaviour when it is blue detuned.

In the latter case, the 866nm becomes further detuned, and the effectiveness of the repumping reduced, when the ion moves towards the 397nm beam. This reduces the 397nm transition rate when the ion is moving towards the beam, reducing the effectiveness of the cooling. In some cases, such as the scan considered in this section, the effect can be so pronounced that the 866nm causes the ion to heat when it is sufficiently blue detuned.

At zero trap frequency, the dynamic method reproduces the results of the steady-state method: the ion cools at large negative 866nm detunings, experiences infinite heating at large positive detunings and the detunings of the two beams oscillate in phase with the ion's velocity. For a finite trap frequency, however, the repumping process gets out of phase with the velocity. This is the same mechanism seen in the two-state system for $\omega_T \sim A$. When the timescale of the ion's physical oscillations is comparable to or faster than the timescale of spontaneous emission on the repumper, the 866nm will not be able to induce the full amplitude of the $D_{3/2}$ oscillations that would occur at zero trap frequency. The amplitude of U_D will be reduced and the effect of the 866nm damped. As $\omega_T \to \infty$, $U_D \to 0$. A reduction in the magnitude of the $D_{3/2}$ oscillations will correspondingly reduce the magnitude of the $S_{1/2}$ oscillations as there is less variation in the population repumped onto the 397nm transition.

As the effect of the 866nm is damped out by increasing trap frequency, the effect of both modifications of the cooling caused by the third level is reduced. As $U_D \to 0$, both U_S, U_P and therefore the temperature will tend to constant values independent of the sign of the 866nm detuning. U_P will always increase as U_D decreases due a reduction in the 'syphoning' effect. For positive detunings of the 866nm, the heating caused by the Doppler shift of the 866nm beam is damped out and cooling is recovered. Note that this occurs even at very small values of the trap frequency. This is as expected from the limit (4.69) due to the small populations in the $P_{1/2}$ at large 866nm detunings.

This also explains the similar phases of the population oscillations observed in figure 4.16. At large trap frequencies, the effect of the $D_{3/2}$ oscillations is damped out. The $S_{1/2}$ and $P_{1/2}$ population oscillations are now dominated by the oscillations of the 397nm detuning for which the dynamic effect is negligible. These oscillations therefore show no phase lag relative to the ion's external oscillations. A system for which the dynamic effect is negligible for both beams also shows no phase lag. In this example, the phases of the $P_{1/2}$ oscillations are only approximately similar; the trap frequency is only sufficient to reduce U_D, but not to make it negligible. U_D will only be negligible when:

$$\frac{\omega_T}{2\pi} \gg \left\langle P(\mathrm{P}_{1/2}) A(\mathrm{P}_{1/2}, \mathrm{D}_{3/2}) \right\rangle . \tag{4.94}$$

This is a modification of (4.69).

Figure 4.17 also shows small resonance effects linking the ion's internal and external motion. These resonances occur due to the same mechanisms as the resonances at $\omega_T = \widetilde{\Omega}$ for a two-state system: the external oscillations are driving the ion/photon system at the natural frequency of a driven transition between two of its states. We can group the possible stimulated transitions that can occur in this system into several categories. There are, for example, 397nm transitions and 866nm transitions. However, transitions can also occur between states in the same level. These are two-photon transitions and thus occur at a slower rate than single-photon transitions between levels. Of these transitions, those involving the 866nm beam will be the slowest of all, and the slowest stimulated transitions in this system occur between states in the $D_{3/2}$. The resonant effects seen in the figure are due to these transitions. They are both slow and weak: the effect of these resonances is much less pronounced than the equivalent resonance in a two-state system. Note that the U_D plotted in this figure is the sum of the amplitudes of all the $D_{3/2}$ states, so there is in general no sign change around each individual resonance. Other types of transitions have characteristic frequencies that are too large to be seen on this plot. These large frequencies arise as the

transitions must necessarily involve 397nm photons, and $A_{397} \gg A_{866}$.

The characteristic frequencies of the $D_{3/2} \leftrightarrow D_{3/2}$ transitions are a function of all of the laser parameters. However, we can determine the limiting value of these frequencies when the 866nm beam is far detuned. In this case, the populations of the $S_{1/2}$ and $P_{1/2}$ tend to zero. We can therefore model each $D_{3/2} \leftrightarrow D_{3/2}$ transition as an isolated two-level system. For these systems, $A, \Omega \ll \Delta$, so the characteristic frequency of each transition is simply Δ. This can be demonstrated by taking the limit in (4.91). The limiting value is frequency separation of the two $D_{3/2}$ states[2].

For the **B** field associated with this scan, the states in the $D_{3/2}$ are separated by 0.99MHz. Raman transitions will be driven between pairs these states with $\Delta M = 1$ and $\Delta M = 2$ when the frequency of the ion's external oscillations is approximately 1MHz or 2MHz. These Raman transitions modify the $P_{1/2}$ population and therefore also alter the ion's temperature. There are three pairs of states with $\Delta M = 1$ and therefore three distinct resonances when $\omega_T/(2\pi) \sim$ 1MHz. There are only two pairs of states with $\Delta M = 2$ and only two resonances when $\omega_T/(2\pi) \sim$ 2MHz. The resonances with $\Delta M = 1$ are present only because there is a 3% fraction of π present in the 866nm beam.

Figure 4.18 shows the dependence of the in-phase oscillations of the $D_{3/2}$ population on the trap frequency and 866nm detuning. The predicted temperature is also shown. At low trap frequencies, the primary effect of increasing the trap frequency is to dramatically reduce the width of the heating region. The width of the region becomes constant at trap frequencies greater than approximately 3MHz. This is because, as the heating effect of the 866nm damps down with increasing trap frequency, the ion will stay cold for a greater proportion of the scan. The rising temperature at large negative 866nm detunings is also quickly damped out at low trap frequencies for the same reason. Above a trap

[2]The 866nm detuning does not contribute as the effect on the transitions to and from the $P_{1/2}$ cancels out

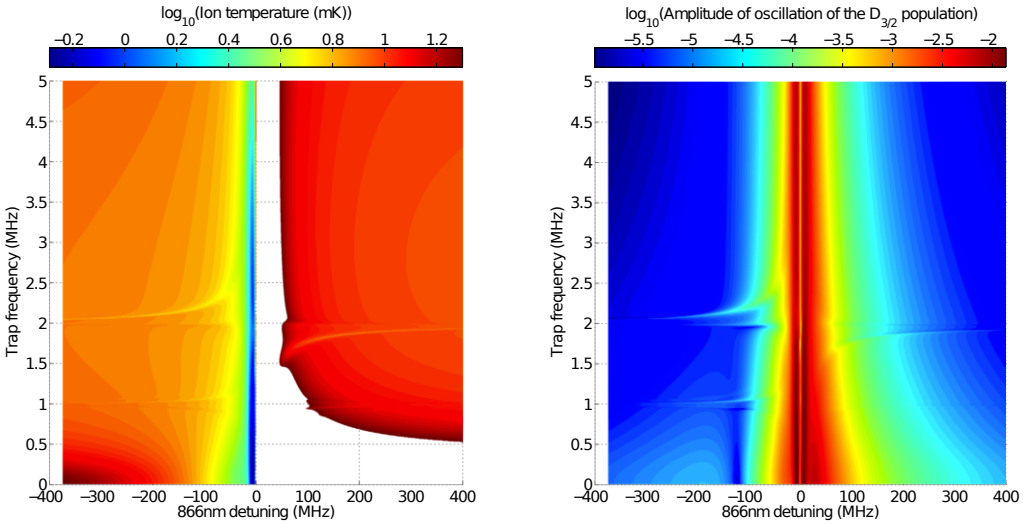

Figure 4.18: A demonstration of the variation in the temperature of the ion and the amplitude of the oscillations of the $D_{3/2}$ population with respect to the trap frequency and 866nm detuning. For both of these plots, a horizontal cross section corresponds to an 866nm scan. Values that are not plotted in the left-hand scan refer to regions where the ion heats indefinitely, or where the temperature is higher than $10^{1.3}$mK. The right-hand scan shows the logarithm of the amplitude of the oscillations to make the small variations clearer. At low trap frequencies, the primary effect of the trap frequency is to determine the width of the heating region. This reduces in size sharply as the trap frequency increases and the effect of the 866nm becomes less prominent. The effects of the sets of resonances at trap frequencies of approximately 1MHz and 2MHz are also visible.

frequency of 3MHz, the behaviour of the system remains approximately constant with the strong oscillations of the $D_{3/2}$ around the dark resonance slowly being damped out. The minimal temperature on the dark resonance also rises slowly with increasing trap frequency. For trap frequencies this large, the effect of the 866nm becomes like that of an increasingly large linewidth, resulting in a smoothing out of sharp features of the temperature curve.

The ion/photon resonances are also visible. The resulting change in temperature due to these resonances is relatively small, but the effect on the $D_{3/2}$ oscillations is larger. The resonances at around 2MHz are more prominent than those around 1MHz as the former requires only σ polarisations whilst the latter relies on the small fraction of π. The variation in the frequencies of the resonances due to the dependence of the characteristic frequencies on the 866nm detuning can also be observed. The frequencies of the resonances converge to

1MHz or 2MHz at large detunings, as predicted.

4.5.3 An excursion method incorporating the dynamic effect

It is worth noting that, at all trap frequencies, the dynamic method still predicts that the ion will become infinitely hot to the high-frequency side of the dark resonance. This prediction is incorrect. It arises from the same flaw as in the steady state method; the method assumes that the gradient of $R(v_x)$ remains constant over the ion's oscillations, something which is not valid when the ion is hot. As with the excursion method, we need to produce an extension to this theory to take account of the ion's excursions over the fluorescence curve as it oscillates in the trap. However, we cannot simply combine the excursion and dynamic methods. The dynamic method assumes a steady state that is perfectly sinusoidal, whereas the steady state produced if we were to allow the gradient to vary would inherently *not* be sinusoidal.

To produce a temperature model that includes the response time of the repumper while also allowing for the detailed variation of the fluorescence curve as the ion oscillates, we need to treat this system more thoroughly. The 'motional method' follows the ion's internal and external oscillations explicitly in a time-dependent way. A differential equation linking the ion's internal and external dynamics over time is derived. This method is discussed in the following section.

4.6 The motional method

For the purposes of this derivation, we will be using a one-dimensional semi-classical treatment of the ion's motion; the internal state of the ion at all times is described by a full quantum-mechanical treatment, but the ion's motion is treated entirely classically. This derivation follows the same reasoning as the

steady-state method but allows coupling between the ion's motion and its internal states.

We begin with the equation of motion of the ion:

$$m\frac{d^2x}{dt^2} = F(x,t) + C(v_x) + H(v_x) \tag{4.95}$$

where $F(x,t)$ is the force exerted by the trapping potential as before, $C(v_x)$ is the force arising from photon absorption and $H(v_x)$ is the effective force arising from photon emission. v_x is the component of the ion's velocity along the direction of the laser $i = 1$. The laser is propagating along positive $\hat{x}$ as before.

We first consider the photon absorption. In a given time period dt, the ion will absorb a certain number of photons:

$$dn = \sum_i R_i(v_x)dt \cdot \tag{4.96}$$

This leads to a change in the ion's momentum:

$$dp_x = \frac{\hbar}{c}\sum_i \{\omega_i b_i R_i(v_x)\}\, dt \cdot \tag{4.97}$$

As with previous methods, we assume that the energy of the ion will rapidly thermalize across all degrees of freedom and that the rate of cooling will therefore be a third of that calculated for a simple one-dimensional case. Both the cooling rate and the effective force on the ion are proportional to the rate at which photons are absorbed. We can model the reduction in the cooling rate by a reduction in dn of the same factor. As a result, reducing the rate of cooling by a factor of three can be modelled by a reduction in the force by a factor of three. The equivalent force on the ion is therefore:

$$C(v_x) = \frac{\hbar}{3c}\sum_i \{\omega_i b_i R_i(v_x)\} \cdot \tag{4.98}$$

The heating of the ion occurs due to the mechanism discussed in detail in section 4.3. Photon emission leads to an increase in the mean square velocity of

the ion. This increase is proportional to the number of absorption events and the recoil velocity $v_{r,i}$:

$$d|v_x|^2 = 2\sum_i \left\{R(\omega)v_{r,i}^2\right\} dt \tag{4.99}$$

where:

$$v_{r,i} = \frac{\hbar\omega_i}{mc} \cdot \tag{4.100}$$

The factor of two arises from the inclusion of the heating caused by the random nature of the photon density in the beam. We can also add the external heating rate ξ of the trap to the model at this point. The heating rate will increase the ion's kinetic energy. In a time period dt, the increase in energy due to heating rate will increase the ion's speed as:

$$\xi.dt = d\left(\frac{m|v_x|^2}{2}\right) = \frac{m}{2}d|v_x|^2 \cdot \tag{4.101}$$

The differential of the mean square velocity depends linearly on the heating rate, so that:

$$d|v_x|^2 = \left(2\sum_i \left\{R_i(v_x)v_{r,i}^2\right\} + \frac{2}{m}\xi\right) dt \cdot \tag{4.102}$$

This leads to a change in the velocity:

$$2v_x dv_x = \left(2\left(\frac{\hbar}{mc}\right)^2 \sum_i \left\{\omega_i^2 R_i(v_x)\right\} + \frac{2}{m}\xi\right) dt \tag{4.103}$$

and the momentum:

$$dp_x = \left(\frac{\hbar^2}{mc^2}\sum_i \left\{\omega_i^2 R_i(v_x)\right\} + \xi\right)\frac{1}{v_x}dt \tag{4.104}$$

producing an equivalent force on the ion:

$$H(v_x) = \left(\frac{\hbar^2}{mc^2}\sum_i \left\{\omega_i^2 R_i(v_x)\right\} + \xi\right)\frac{1}{v_x} \cdot \tag{4.105}$$

The equation of motion of the ion can now be written :

$$\frac{d^2x}{dt^2} = \frac{F(x,t)}{m} + \frac{\hbar}{3mc}\sum_i \left(\left\{b_i + \frac{3\hbar\omega_i}{mcv_x}\right\}\omega_i R_i(v_x)\right) + \frac{\xi}{mv_x} \cdot \tag{4.106}$$

This differential equation describes how an ion oscillating in the trap will behave. For a heating region, the ion's energy will climb exponentially. For a cooling region, the ion will experience decaying oscillations in its motion and will reach an oscillating steady state. The energy of the ion can be obtained from this steady state by extracting the maximum kinetic energy of the ion. This gives a measure of the ion's energy independent of the trapping potential. This is important for situations where the trapping force is more complicated than pure harmonic; a trapping potential that incorporates micromotion, for example, has no well-defined static potential, and thus the energy must be inferred from the velocity. As with the steady-state method, the temperature of the ion can be inferred by analogy to equipartition as shown in 4.66.

4.6.1 Dealing with the pole

Equation (4.106) has a complication that makes numerical solution nontrivial: the equation has a pole around $v_x = 0$. This arises from the statistical treatment of the heating.

The heating term adjusts the ion's velocity by increasing the square of its magnitude according to (4.102). Consider an ion as it approaches the furthest point in its physical oscillations along $\hat{x}$. The velocity of the ion will be small and will tend to zero when the ion reaches its maximum displacement. Because the velocity of the ion varies by only a small amount over this region, the Doppler shifts the ion experiences will not vary significantly and $R_i(v_x)$ will be constant to a good approximation. In this region, we write (4.102) simply as:

$$\frac{d|v_x|^2}{dt} = C \tag{4.107}$$

where C is a constant. In a fixed time interval dt, $d|v_x|^2$ will be equal to the same value regardless of the velocity at the beginning of the interval, so long as that velocity is small. In general, as v_x tends to zero, $|v_x|^2$ tends to zero at a greater rate. If we reduce the magnitude of the velocity at the beginning of the time step,

dv_x must increase in order to produce the same value of $d|v_x|^2$ over the interval. As the ratio of v_x to $|v_x|^2$ increases with no upper limit as v_x tends to zero, dv_x must also increase without bound in order to preserve the constant $d|v_x|^2$. This results in the limiting behaviour around the pole:

$$\lim_{v_x \to 0\pm} \left(\frac{dv_x}{dt} \right) = \pm\infty \cdot \tag{4.108}$$

The differential equation is undefined at the point $v_x = 0$. This is because the heating term relies on the ion having a well-defined direction of travel in order to increase the magnitude of the velocity at each timestep. This is not the case when the velocity is exactly zero.

If it were possible to calculate with perfect precision, this pole would not cause any unphysical behaviour except in the situation where the starting conditions involved a zero velocity. However, problems arise as soon as the precision of computing is limited. Whenever very small velocities are encountered by the routine, the heating term of the differential equation can become orders of magnitude larger than all the other terms. Errors in such a large term can swamp the other terms which are crucial to the ion's behaviour. The effect of this is that, if the pole is not corrected for, the behaviour of the ion will be erratic and unphysical whenever it passes over zero velocity. For example, the most common type of behaviour is that the ion's energy undergoes a large step change, resulting in a step change in the amplitude of the ion's oscillations. In order to calculate using (4.106) avoiding behaviour like this, the large magnitudes of the heating term around the pole must be dealt with.

We solve the differential equation using SILVER FORK as described in subsection 3.4.3. I use two approaches for avoiding the problem posed by the pole. The first method is used for scans where the excursion effect is not prominent and all temperatures are within an order of magnitude of the Doppler limit. In this case, the heating term is modified in a small region around the pole. Instead of following a $1/v_x$ dependence, the magnitude of the term is fixed to

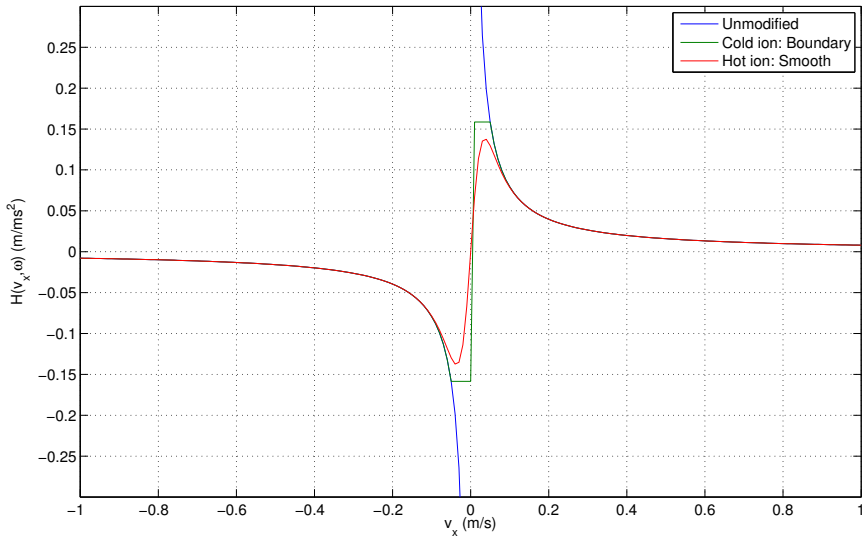

Figure 4.19: A demonstration of the two methods to control the behaviour of the differential equation about the pole. For cold ions, the magnitude of the heating term is limited in a fixed boundary about zero velocity. For hot ions, a faster method is employed that smoothly varies the magnitude of the term about the origin, but results in unphysical behaviour during the calculation.

$1/v_{\max}$ where $v_{\max}$ is small relative to the maximum size of v_x, determined by the recorded history of the ion's oscillations:

$$H(v_x, \omega) \propto \begin{cases} 1/v_x & \text{if } |v_x| \geqslant v_{\max} \\ 1/v_{\max} & \text{if } |v_x| < v_{\max} \end{cases} \tag{4.109}$$

This reduces the heating of the ion. However, because this occurs only during a very small portion of the oscillations of the ion, the effect on the ion's final temperature is not dramatic. As long as $v_{\max}$ is less than approximately 10% of the ion's initial velocity, no effect on the ion's final temperature is observed. If this method fails, the error is obvious: the ion treats the sharp turn in the heating function as a boundary and begins to oscillate entirely within the $1/v_{\max}$ region. This allows the problem to be easily detected and the value of $v_{\max}$ to be corrected if necessary.

This method produces oscillations of the ion that show no distortion around $v_x = 0$. However, the integration can be slow, especially around the boundary. This is not a problem for when the ion is cold; the ion adjusts to its steady-state

temperature quickly. However, for hot ions, the ion may take much longer to reach the steady-state. This means that it can take a long time to calculate temperatures for hot ions. To simulate such systems, a faster but less robust method to avoid the pole is used. Instead of changing the cooling term abruptly around the boundary, it is changed smoothly. The heating term is given an additional dependence:

$$H(v_x, \omega) \propto \frac{1}{v_x} \tanh^2(a_x v_x) \tag{4.110}$$

where:

$$a_x = \frac{1}{v_x} \tanh^{-1}(\sqrt{b_x}) \cdot \tag{4.111}$$

b_x is the value of the modifying tanh function at the 'boundary' velocity $v_{\max}$. For these simulations, I set $b_x = 0.99$ and $v_{\max}$ as 10% of the initial maximum speed during the ion's oscillations.

The $\tanh^2$ term modifies the heating term so that it passes smoothly through the origin at $v_x = 0$. See figure 4.19. The function will calculate more quickly around a smooth boundary than around a step change. The change in the heating also causes distortion around the boundary. The oscillations become slower whenever the boundary is passed. However, even though the physical dynamics of the ion are not accurately reproduced, the variation in the energy of the ion is. The ion's energy change is not altered by this approach, and the final temperature predicted by this method is identical to the method involving a sharp boundary. The disadvantage of this approach is that a failure of the method is much less easy to identify. Rather than a sudden change in the ion's dynamics, the ion's evolution will stall at some point during the calculation, and the only trace of this is the presence of distortions about the boundary $v_x = 0$. When the steady-state is reached, no distortions should be visible. The steady state produced using this method should be checked to ensure that it is consistent.

4.6.2 Temperature prediction

The method under discussion can handle all parameter regimes. Because it includes excursion, there is a possibility that multiple solutions for the ion's dynamics will appear at a given set of parameters depending on the initial conditions. Thus, this method can describe hysteresis. For simple temperature predictions , however, we are only interested in the lowest possible temperature the ion can achieve. We calculate the temperature over a scan using the following method. At each point of the scan, the ion is given a velocity corresponding to 1.5 times the Doppler limit. The steady-state solution, taking account of the detuning of the lasers at this velocity, is taken as the ion's internal state. The ion is then allowed to evolve until its temperature can be inferred.

In order to speed up calculation, we make use of a further property of the ion's evolution. After the ion's initial oscillations, the energy variations of the ion over time converge on the final steady-state exponentially, regardless of the method used to calculate over the pole. As a result, we can calculate partway to the steady state and still produce a value for the steady-state temperature. We do this by looking at, for example, the extent of the ion's velocity oscillations as a function of time. We fit this to an exponential convergence and use the asymptote as a prediction of the final temperature. This can significantly decrease calculation times.

Figure 4.20 shows the results of the motional method. The scan chosen is a simulation with a **B** field of 5G. This has the advantage that the effects of the method on dark resonances are more prominent. The motional method combines the effects of the dynamic and excursion methods. It predicts both cooling at large positive detunings of the repumper and finite temperatures around the dark resonances. The motional method predicts different finite temperatures in these regions than the excursion method. Even when the ion is hot, the relative timescales of the ion's internal and external motion are important, and only a

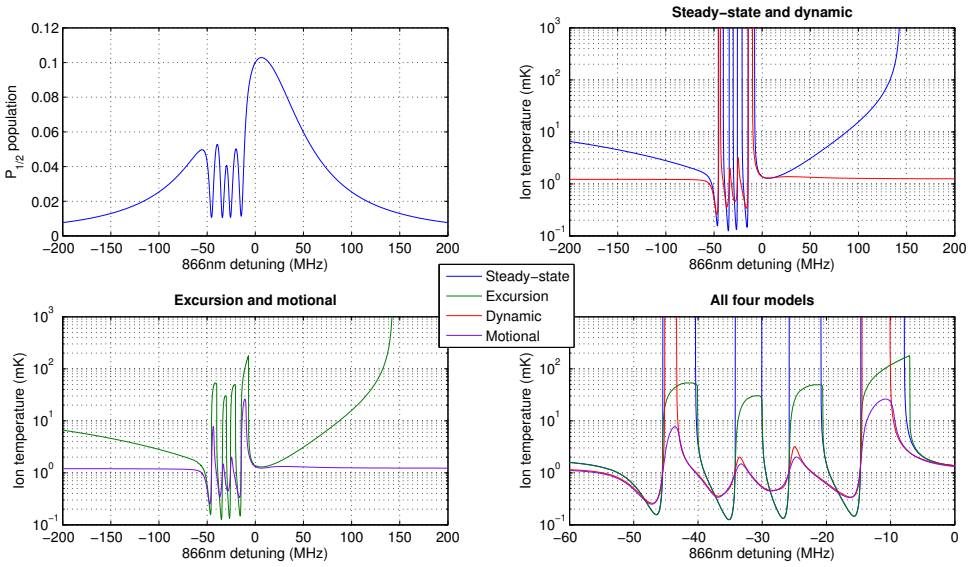

Figure 4.20: A comparison of the results of the four temperature models. The scan shown is a simulation at a **B** field with magnitude 5 gauss to highlight the behaviour about the dark resonances. The steady-state method predicts infinite heating to the blue of each of the dark resonances and to the blue of the entire scan. The excursion method predicts finite temperatures on the dark resonances, but not to the high-frequency side of the scan. The dynamic method predicts finite temperatures to the blue of the scan, but not on all of the dark resonances. The motional method combines the excursion and dynamic effects, predicting both a finite temperature to the blue of the dark resonances and on the high-frequency side of the scan. Note that the temperatures predicted by the motional method on the high-frequency side of the dark resonances are different from those predicted by the excursion method, and that they are similar to the temperatures of the two resonances predicted by the dynamic method to have a finite temperature in these regions.

method that accounts for this will give accurate predictions.

4.7 Fluorescence of a hot ion

There are two advantages of the motional method for describing the behaviour of hot ions. Firstly, the method explicitly calculates the internal dynamics of a hot ion. This means that the fluorescence emitted by even a very hot ion can be simulated. In addition, the method follows the ion's motion in detail so that the output of the simulation depends on the ion's history. This method can model hysteresis of the ion, and it can do so more accurately than the excursion method. To directly simulate a fluorescence scan, the ion is given only an

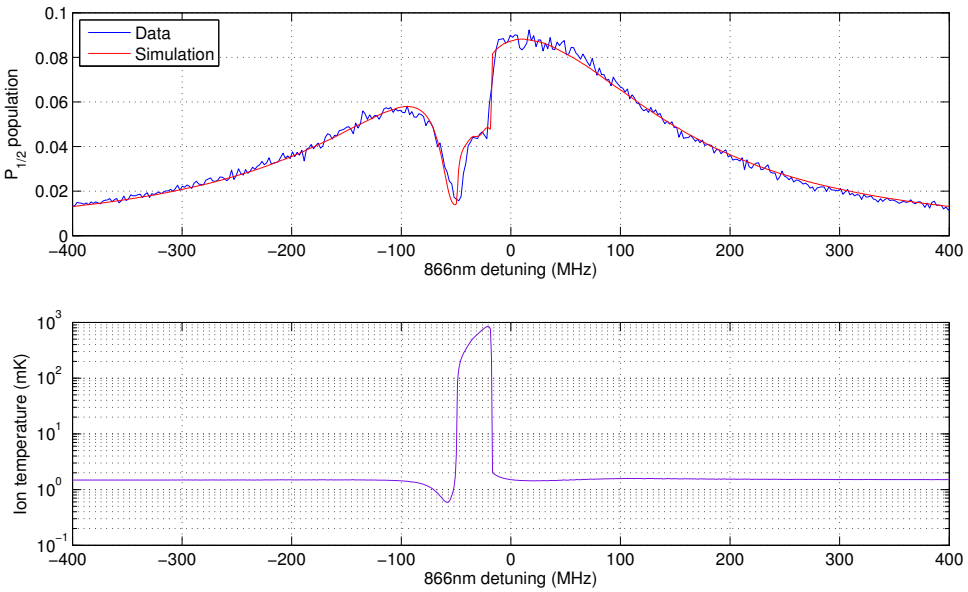

Figure 4.21: A simulation of the fluorescence of a hot ion using the motional method. This scan is the same as that analysed in section 4.4.2, and the one displayed is the scan from red to blue. The observed fluorescence is well matched by this method.

approximate set of initial conditions at the start of each scan. The initial conditions for each subsequent point are taken from the equilibrium conditions of the previous point on the scan. Specifically, the point at which the velocity is most positive at equilibrium is taken as the new starting point. The allows the ion to keep track of its temperature over the course of the scan.

The routines are too slow to allow precise fitting for this sort of effect. However, approximate fits can be performed using a steady-state calculation of the populations and utilising the dynamic method as a guide for the boundaries of any heating regions. Even with these simple fits, the agreement between the outputs of this simulation and experimental data are impressive. See figure 4.21 for an example of the prediction over a small region of heating. This is the scan analysed in section 4.4.2. See figure 4.22 for a simulation of a scan that shows hysteresis over a large portion of the scan.

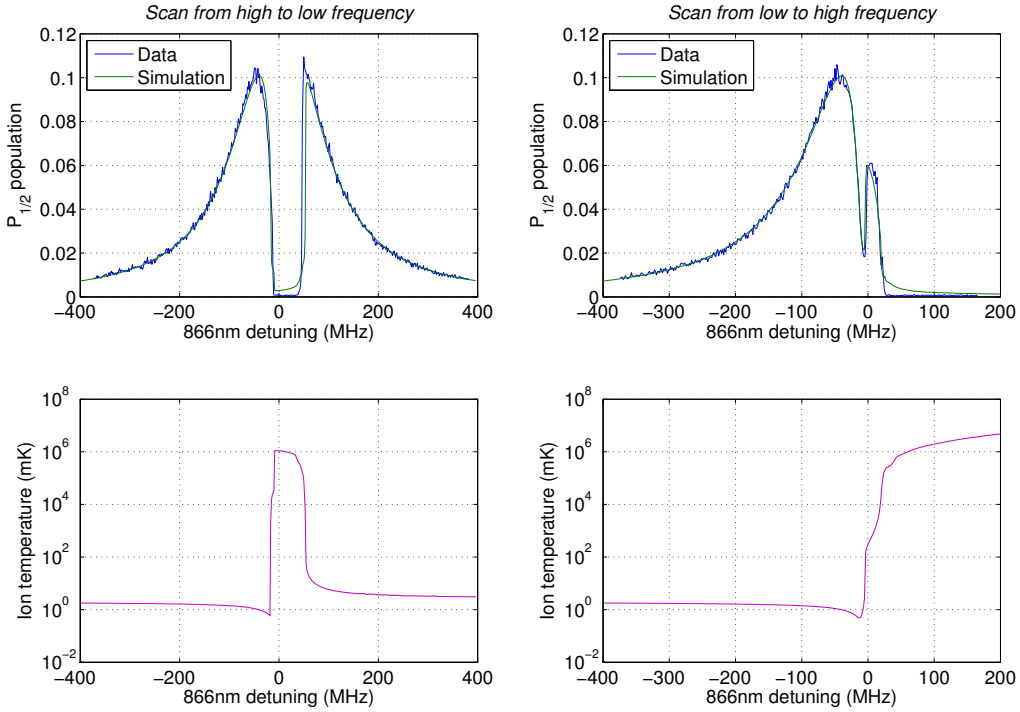

Figure 4.22: Modelling hysteresis using the motional method in fluorescence scans with significant heating. The experimental data shows hysteresis over a large portion of the scan. When scanned from red to blue, the ion is hot in the blue wing. However, when scanned in the other direction, the ion is cold in this wing. This behaviour is well replicated using the motional method. Small discrepancies between the data and simulation occur when the ion is very hot. In these situations, details such as the finite size of the beams and the effects of micromotion become important. These details can be included as refinements to the motional method.

4.8 Doppler cooling below the Doppler limit

In the previous sections, several temperature models have been presented and used to describe the various features of fluorescence scans in $^{40}Ca^{+}$. I now use these temperature models to predict the lowest temperatures that can be achieved for a $^{40}Ca^{+}$ ion under reasonable experimental conditions. It is possible to use Doppler cooling to cool the ion below the Doppler limit, in some cases by a large factor.

Of the models previously discussed, the dynamic method is the most suitable for this study. Calculation times are much faster than the excursion and motional methods and it handles the dynamics of the fast motion of the ion very

well. Discrepancies between the predictions of this theory and experimental data only occur when the ion is hot; we will not be considering such regions when searching for a cold ion.

In the following, I describe the failure of direct optimisation methods for the ion's temperature. I then develop an optimisation process that can calculate minimal values for the temperature of the ion that can be realised experimentally, and I display the results of the process for $^{40}Ca^{+}$.

4.8.1 Failure of simple optimisation

Temperature minimisation using the dynamic method suffers from an immediate drawback. The temperature of $^{40}Ca^{+}$ ion is calculated using the same gradient-based approach as the steady-state method:

$$T = C_1 R(v_x) \left(\frac{dR(v_x)}{dv_x} \right)^{-1} + C_2 \xi \tag{4.112}$$

where C_1 and C_2 are constants and ξ is the trap heating rate. This function has a very prominent local minimum that will dominate direct minimisation of the temperature. This is best demonstrated by looking at a simple two-state system. The scattering rate for the two-state system is Lorentzian. In section 4.3.1 we have shown that the minimal temperature for this system is:

$$T_{min} = \frac{\hbar \Gamma}{2k_B} \tag{4.113}$$

The width of the Lorentzian scales as the square-root of the power in the driving beam through power broadening, and therefore the temperature of the system scales in the same way:

$$T_{min} = D_1 \sqrt{(1 + D_2 I)} \tag{4.114}$$

where D_1 and D_2 are constants that depend on the atomic parameters. The implies that the optimal temperature of the system occurs for $I = 0$. In other words, the temperature of the two-state system is minimal when the interaction of the

ion with the driving laser is minimal! The same result occurs for systems with a larger number of states and levels. In general, a direct optimisation process will conspire to minimise $R(v_x)$ to minimise the temperature.

This result arises as a limit to the cooling process. In particular, this model looks only at the steady state of the ion's oscillations. As $R(v_x)$ decreases, the relative variation of the function over the ion's oscillations becomes large compared to its mean value and the ion is strongly cooled. However, the interaction of the ion with the laser also decreases in strength so that an ion that is hot will take longer to cool to the steady state. As I tends to zero, the temperature of the ion tends to a minimum, but the time required for an ion to cool to that minimum tends to infinity.

The minimal temperature calculated using this approach is of course entirely impractical. In order for a temperature optimum to be useful from an experimental point of view, $R(v_x)$ needs to be sufficiently large that the ion responds on a fast timescale and so that the ion's fluorescence can be observed. In fact, it is beneficial for $R(v_x)$ to be as large as possible. As direct optimisation of the temperature is not useful, this suggests a second approach. Instead of simply minimising the temperature, we perform a coupled optimisation. We simultaneously minimise the temperature whilst maximising $R(v_x)$. Because we are optimising an additional parameter, we have added a new dimension to the problem; rather than producing an optimal point, we will instead produce an optimal line. Each point on the line describes an optimum with a different weighting between the minimisation and maximisation processes, and we can construct this line by performing a series of optimisations varying this weighting.

This is one of several ways the problem could be tackled. One other example includes calculating optima only using minimisation while keeping the laser intensities fixed. However, the minimax approach is particularly useful as it deals directly with the two parameters most important from the point of view

of cooling: the ion's temperature and its fluorescence. Given the experimental requirements, we can directly read off an optimal set of parameters from the optimal line without any conversion or any assumptions based on the effects of the various laser parameters. This minimax approach is described in the following section.

4.8.2 A minimax approach

The core of the minimax approach is the quantity that is being optimised. This incorporates both the temperature and fluorescence, so the output can be optimised using a single optimisation process. The quantity is:

$$Y = \theta + \zeta/P \tag{4.115}$$

where θ is the ion's temperature, ζ is the weighting of the temperature minimisation against the fluorescence maximisation and P the population of the $P_{1/2}$ which is taken to be proportional to the ion's fluorescence.

The optimal line is generated by minimising Y for a series of values of ζ. The minimum value is found using a global optimisation process similar to that described in section 3.3.6 for fitting fluorescence scans. A large number M of initial points is generated with each point taking random values of the parameters which are being floated. The range of these values is taken to be the entire possible range that is feasible to realise experimentally. No bias to any particular region of parameter space is input into the optimisation process; the search is as general as possible. The points are ordered in decreasing value of Y and the N lowest values are chosen. These points have Y optimised locally, and the lowest value of Y of the set of points is taken as the global optimum.

The advantage of this optimisation process over the one used for fitting scans is that each function evaluation involves calculating a single point rather than an entire scan. As a result, each function evaluation is significantly faster, and we can therefore significantly increase the values of M and N to ensure more robust

fitting without compromising on calculation time. For calculations in $^{40}Ca^{+}$, I usually set N to be between 5 and 20, whilst M is set to be several thousand or several tens of thousands.

The output of this optimisation method can be processed further. Y is a general parameter which makes no assumptions about the location of the optimum; every parameter contributes to the minimax process. From an experimental point of view, the laser detunings can be varied simply whilst keeping all other parameters fixed. On a fluorescence scan, the minimal value of Y will necessarily not be the minimal value of the temperature, although it will lie close to it. To better compare optimal sets of parameters produced by this method, the minimax temperature is modified. For each value of ζ, the set of parameters produced by the minimax optimisation is put into an optimisation process that minimises the temperature alone only with respect to the laser detunings. This produces a second optimal line.

The minimal value of Y will occur for low temperatures and large populations. This means that, in general, Y will be minimal around a dark resonance in a three-level system, as we would expect, but offset from the optimal temperature on the dark resonance with respect to the detunings. This final step corresponds to moving the minimax optimum to this point of minimal temperature.

Having discussed the temperature of a $^{40}Ca^{+}$ ion for various fluorescence scans, we now use this approach to see how cold it is possible to cool such an ion.

4.8.3 Minimal temperatures in $^{40}Ca^{+}$

Figure 4.23 shows the two optimal lines for the temperature and fluorescence of $^{40}Ca^{+}$ ions. The parameters floated were the intensities and detunings of the 397nm and 866nm beams along with the magnitude of the $\mathbf{B}$ field. The polarisa-

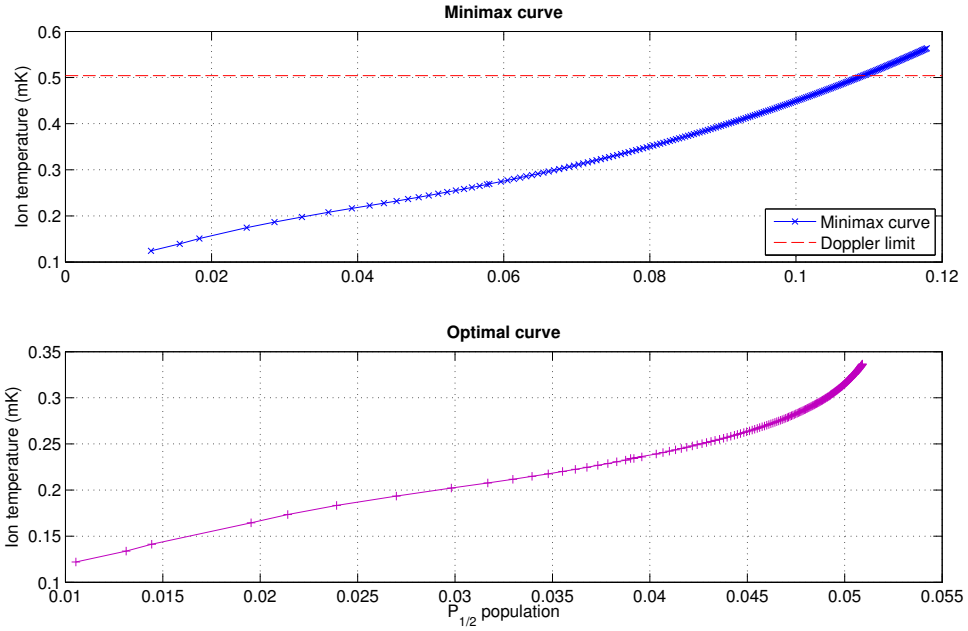

Figure 4.23: The two optimal lines produced by the minimax process. The upper line is the direct result of minimax optimisation. The lower line results from a separate stage that shifts the minimax optimum by altering the value of the 866nm detuning to the point of minimal temperature on the dark resonance. Cooling well below the Doppler limit is feasible, even for large $P_{1/2}$ populations.

tions of both beams were assumed to be equally split between σ^+ and σ^-, corresponding to beams travelling along the field direction. Because this search only involved low temperatures, the large linewidths used in the fitting processes described in section 4.1.1 were not needed; the linewidth of the 866nm was set to 100kHz and that of the 397nm to 500kHz. The effects of the motion on the temperature are incorporated in the dynamic method. The optimal lines show that the predicted temperature is significantly lower than the Doppler limit even when the population in the $P_{1/2}$ is large.

Figure 4.24 shows fluorescence scans of the ion at various points along the optimum line, demonstrating the variation in the optimal set of parameters depending on the weighting of the temperature relative to the population. The two extremes of the line are shown along with a compromise in the centre. For all of these curves, the minimal temperature is well below the Doppler limit.

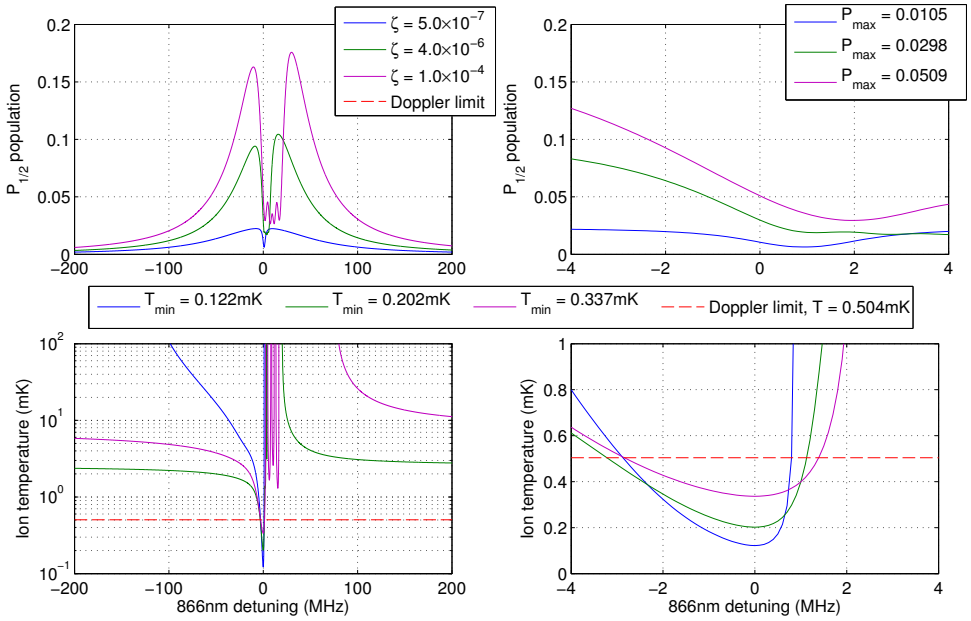

Figure 4.24: A demonstration of three points along the optimal line showing various weightings between minimising the ion's temperature and maximising its fluorescence. The blue and purple curves are the lower and upper extremes of the optimal line shown, and the green curve is a compromise around the middle of the curve. Even for populations as large as 5%, the optimal temperature is significantly below the Doppler limit. The zero of the 866nm scale is centred around the resonance for each plot to aid comparison. The zero corresponds to +12.0MHz, -3.8MHz and -13.5MHz for the curves with $\zeta = 5.0 \times 10^{-7}$, 4.0×10^{-6} and 1.0×10^{-4} respectively.

Depending on the value of the 397nm linewidth, it is possible to do even better than this. The value of this linewidth controls the depth of the dark resonance and the magnitude of the gradient of R in its vicinity. Reducing the linewidth will increase the gradient and allow the ion to cool to a lower temperature. A conservative value of the linewidth was chosen for these calculations. However, it is possible to experimentally realise a linewidth lower than 500kHz. Figure 4.25 takes one point on the minimax line and shows the optimal temperature as a function of the linewidth. If the linewidth is reduced from 500kHz to 100kHz, the optimal temperature is lowered by 60%. For a linewidth of 100kHz, temperatures five times lower than the Doppler limit can be achieved with peak values of the fluorescence over a fluorescence scan in excess of 10%.

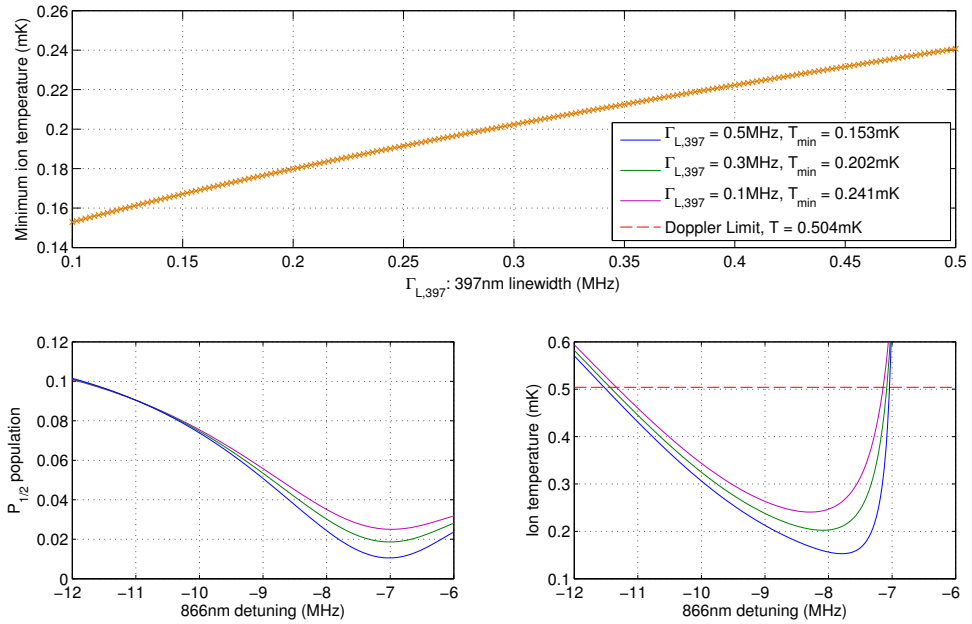

Figure 4.25: A demonstration of the effect on the output of the minimax process of varying the 397nm linewidth. The point chosen has $\zeta = 1.2 \times 10^{-5}$. Reducing the 397nm linewidth so that it is equal to the 866nm linewidth lowers the optimal temperature by 60%, producing temperatures 5 times lower than the Doppler limit even for a point that has a relatively large weighting towards the ion's fluorescence.

5

Dynamics of a five-level system

> While most of my constructions have been based on D'ni designs, I see now that the ones that I have imbued with the power of Five are clearly the most beautiful, the most perfect. And, I believe, the most structurally sound.
>
> Gehn
> *Riven* by Cyan Worlds

In the previous chapter, we treated Doppler cooling in $^{40}Ca^{+}$ and examined the fluorescence and temperatures that can be achieved for a cooling scheme involving the three lowest-energy levels of the ion. However, it is sometimes advantageous to employ additional levels in $^{40}Ca^{+}$. The next simplest closed systems of levels in this isotope involve four and five levels, and these systems are used in several experimental schemes. One such example is the four-level scheme used in high-fidelity readout [MSW+08, Szw09]. Four and five-level systems can show more complex behaviour than the two and three-level systems examined previously. In particular, coherent effects are observed across multiple levels that are non-trivial extensions of three-level behaviour.

This chapter analyses a set of data taken using a background-free cooling and detection scheme that makes use of five levels in $^{40}Ca^{+}$ [LAS+12, Lin12]. The data takes the form of the fluorescence output by the ion as a function of the frequency of two beams. By fitting this surface using the optical Bloch equations,

it is possible to identify the coherent effects that lead to the characteristic features seen in the data. These coherent effects are analysed, and a general approach of describing resonant effects in a five-level system is developed to explain the behaviour. To validate this approach, six additional sets of data were taken and compared to the results predicted by the theory. This study shows the versatility of the numerical routines we have developed; they can efficiently analyse and interpret complicated behaviour in atomic systems.

Bright and dark resonance behaviour in three-level systems has been the basis of many applications (see [FIM05] for a review of one example). Recent work has made use of resonant behaviour in systems with more than three levels. Coherent effects in five-level systems similar to those studied in this chapter have been used in the analysis of many areas. Some examples include electromagnetically-induced transparency in five-level systems [GWW+06], superluminal pulse propagation [LQNG15], atom localisation [QZH+12] and two-qubit photonic gates [ORVT06]. An understanding of resonant behaviour in many-level systems is important for a wide variety of applications.

In sections 5.1 and 5.2, I describe the data and techniques used to fit large-data sets described by a large number of parameters. Section 5.3 describes a method of exploring surfaces using an adaptive step-size to ensure that all resonant features are displayed. Section 5.4 explains the features seen in the data using a general treatment of resonant effects in three, four and five level systems covered in appendix E. This appendix first treats systems with no Zeeman structure, and then examines the modification when each level is split into multiple Zeeman states. Section 5.5 examines the six additional data sets and compares the fitted results to the model of the coherent effects in a five-level system. The parameters used to produce the figures are listed in appendix H.

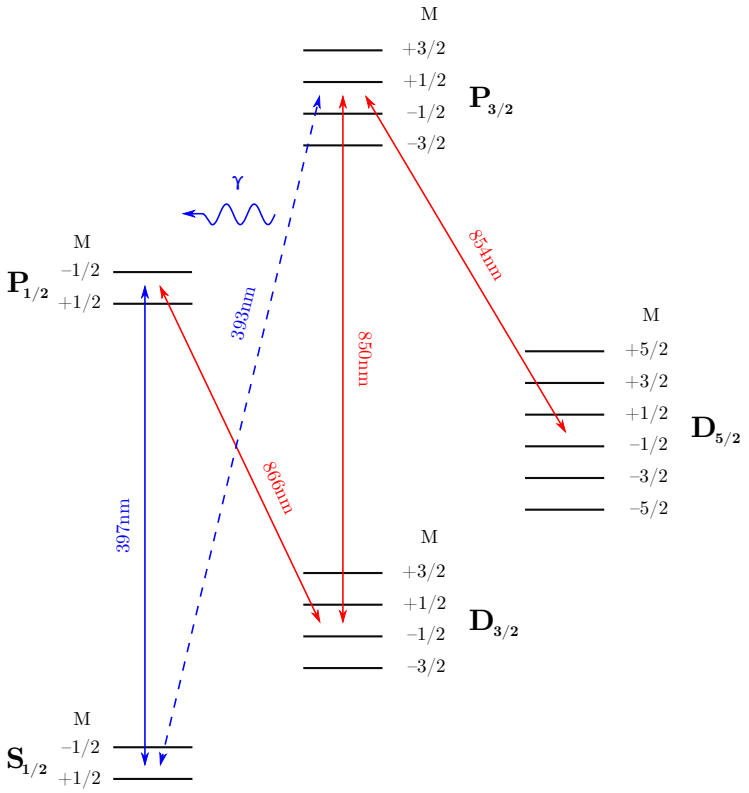

Figure 5.1: A schematic of the five lowest-energy levels of $^{40}Ca^{+}$ being studied including the system of lasers. The energy differences between the Zeeman states and levels are not to scale. The four beams are shown as bold lines and the transition on which fluorescence is collected (393nm) is shown as a dashed line.

5.1 The original data

The background-free detection scheme examined in this chapter is shown schematically in figure 5.1. The scheme involves five levels in $^{40}Ca^{+}$ and four lasers. In a conventional detection scheme, the detected fluorescence is generally on a transition which is subject to laser excitation. As a result, there is a background signal arising from light scattered into the detector from the incident laser beam. The present scheme minimises that scattering by collecting fluorescence on a transition that is not driven by a beam. Light at this frequency is only produced by the atom so, by filtering light arriving at the detector at all other frequencies, scattered light can be almost entirely eliminated.

An advantage of the scheme shown is that the laser beams only drive dipole-allowed transitions. This avoids the need for a high power for very low linewidth laser. However, to observe a strong, non-driven spontaneous transition by driving only dipole-allowed transitions, we have to make use of at least five levels in $^{40}Ca^{+}$ in order to avoid population trapping.

Fluorescence is collected on the 393nm transition. A pair of Semrock FF01-387/11-25 filters is placed before the collection optics to nominally filter out all scattered light. In practice, the other blue beam at 397nm is close enough in frequency to the 393nm beam that a small percentage of this light is admitted by the two filters. The transmission at the two wavelengths is:

$$T_{393} = 91.25\% \tag{5.1}$$

$$T_{397} = 2.2\% \tag{5.2}$$

For more details of the experimental apparatus, see [LAS+12].

In order to maximise the fluorescence, the population in the $P_{3/2}$ level must be maximised. This optimisation was performed empirically. The beam powers were adjusted to produce a maximal signal. The detunings of all the beams were at first set to be approximately close to resonance. In order to avoid any dark states that might occur due to coherent effects in the system, the detunings of 866nm and 850nm beams were then scanned in order to find a set of parameters that maximised the fluorescence. The output of this scan, the fluorescence of the ion as a function of the 866nm and 850nm beam detuning, is the first set of data that we analyse in this chapter. This is shown in figure 5.2. I will refer to this scan as the 'original data'. This scan is composed of a series of frequency scans of the 866nm beam taken at different 850nm frequencies. The frequency difference between each scan was monitored and altered whenever large changes between one 866nm scan and the next were observed. This ensured that detail that was strongly dependent on 850nm was recorded at a higher resolution. In order to

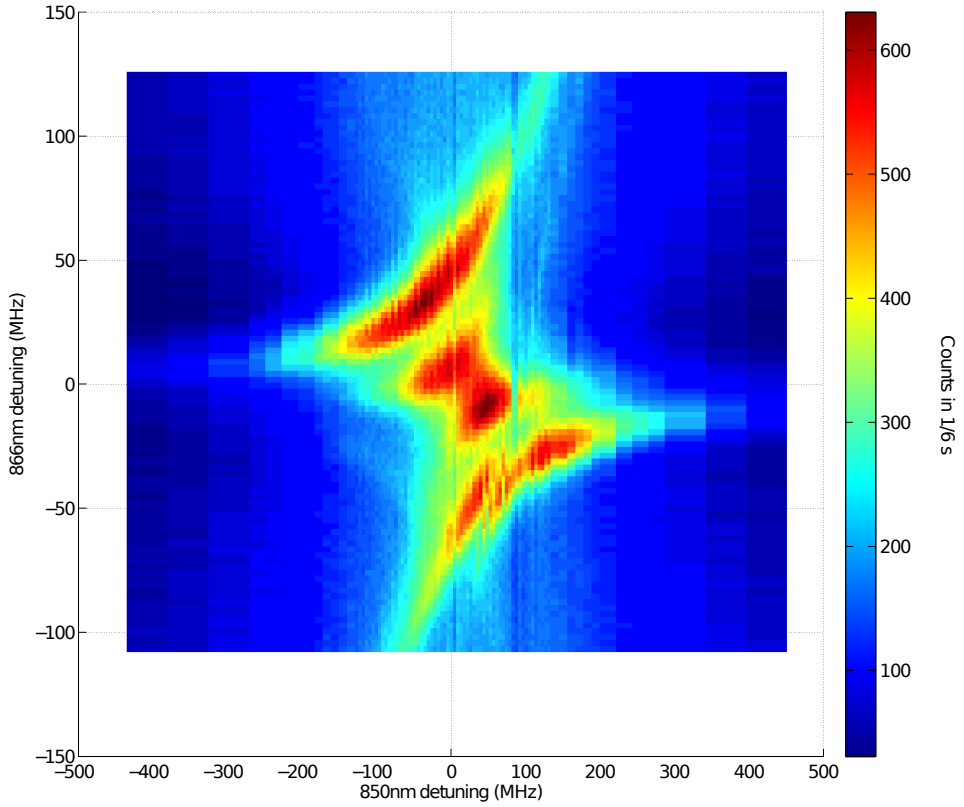

Figure 5.2: The fluorescence of a $^{40}Ca^{+}$ ion as a function of 866nm and 850nm detuning for the background-free detection scheme shown in figure 5.1. The structure of the output is complex, involving several arcs that are modified by features that follow vertical and diagonal paths across the surface. The techniques introduced in this chapter were developed in order to fully explain the origin of these features.

maximise the fluorescence, the intensity of the 850nm beam was set to a value much larger than the intensities of all other beams.

This 'two-dimensional' fluorescence scan shows a series of complex features, including various arcs and linear behaviour. These features are the result of various five-level coherent effects. By inputting the nominal experimental parameters into the Bloch equations simulation, we can roughly approximate the gross behaviour of the scan. However, the location and strength of the resonant features are strongly dependent on the laser parameters; matching the detailed resonance pattern manually is a difficult task.

This is compounded by the large number of variables that describe this

atomic and laser system. For the precise experiments to establish optimum conditions for quantum information processing, many variables were determined in separate experiments, but that was not the case in the work described here. We can describe each beam by five parameters. These are the intensity, the detuning, the linewidth and two parameters describing the polarisation state. In addition, each scanned beam has a frequency calibration variable associated with it describing the conversion factor between the DAC voltage recorded during the scan and the frequency scale. There are also three variables associated with the detected count rate; one variable gives the background count rate[1], and two variables each give the uncertainty in the collection efficiencies of 397nm and 393nm light. Finally, there is a variable associated with the magnetic field. In total, then, this system is described by 26 variables. Manually searching through parameter space to reproduce the resonant features seen in the data with no guide as to their origin would be very time-consuming.

One consequence of the combination of the large number of parameters that describe this system and the several different resonant effects involved is that a rich variety of behaviour can be seen in scans of this type. Figure 5.3 shows four different two-dimensional scans taken with parameter sets very different from the nominal experimental values associated with the data.

In order to simulate the resonant features seen in the data, we use a fitting protocol. Because the structure of the fluorescence surface is complicated, and because the detailed structure of the resonances is strongly dependent on the atomic parameters, a fit based on a single initial set of parameters will probably converge on a local minimum that is not the best match between the simulation and the data. We therefore make use of a global fitting routine to sample over the parameter space defined by the uncertainty in the measured parameters. This routine is similar to the routine described in section 3.3.6, but two modifications

[1]The nature of the scheme means that this is a very small value

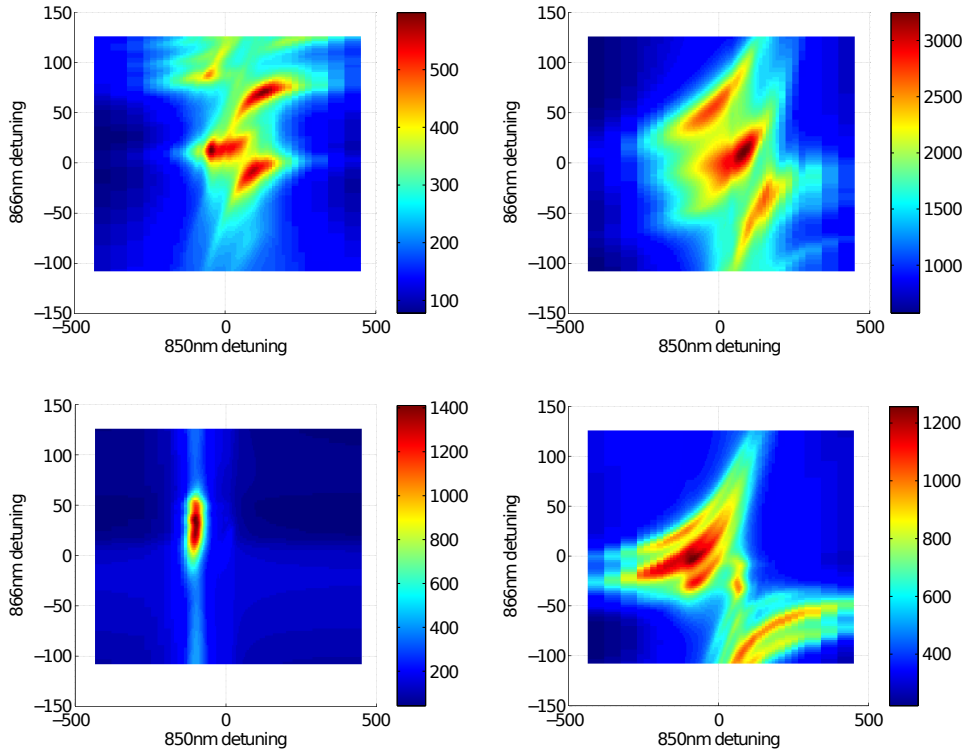

Figure 5.3: A series of four two-dimensional scans taken over the same frequency range as the original data. All 26 parameters describing each scan were randomly generated. These scans demonstrate some of the rich behaviour that can be seen in scans of this type in a five-level system.

are made to take account of the size of the data set. These modifications are discussed in the following section.

5.2 Fitting large, complex data sets

One difficulty in fitting this data is its size. The fluorescence scans we examined in the previous chapter typically involved no more than 500 points. This scan is formed of 86 individual 866nm scans, each comprised of 100 points. In addition, the five-level system contains more states and coherences than a three-level system, so each individual data point for this system will take much longer to calculate. Overall, this two-dimensional scan takes between one and two orders of magnitude longer to calculate than the three-level 866nm scans we have considered previously. The calculation time is $t_s \approx 100\text{s}$. Each iteration of a fitting

protocol will require of order the same number of function evaluations as the number of floated parameters. The total number of parameters in this system is 26. Even if we did not fit all of these, we would still expect one fitting iteration to take of order $10t_s$, which is approximately half an hour. A full fitting process will require of order 20 iterations with a calculation time of $200t_s$, or approximately five hours. Finally, the global fitting process makes use of order 10 full fitting processes. Overall, this process would involved a calculation time of order $2000t_s$, which is approximately two days.

While a calculation time of several days is not prohibitive for a single scan, it becomes unwieldy when we need to fit a series of scans. The global fitting process is therefore modified to calculate more efficiently when applied to large data sets. Two modifications are made. The primary bottleneck in the calculation time is the amount of time a local fit takes to compute. The first modification reduces the number of local fits that need to be fully computed, and the second modification speeds up the calculation time of the remaining fits. These modifications are discussed below.

5.2.1 Double vetoing

In general, the first step of any fitting process reduces the residuals of the output by the largest amount. In essence, this first step broadly identifies the local optimum nearest the initial set of parameters and moves into that region. All subsequent iterations then home in on the local optimum. We can make use of this process to significantly reduce the number of local fits we need to calculate in a global fitting process.

The structure of this global fitting routine is initially similar to that described in section 3.3.6. As before, we use a quasi-Monte-Carlo approach to sample over the parameter space defined by the relative uncertainly of the experimental values. An initial large number M of parameters is chosen at random. A

two-dimensional scan is calculated for each parameter set and the sum of the absolute value of the differences between points on this scan and the data points is calculated. The parameter sets are sorted according the value of this sum, and a small number N of the sets with the smallest sum are chosen.

However, at this point, the global routines differ. These N parameter sets are each put through a local fitting routine, but the fitting process is halted after the first iteration. The norm of the residuals of these outputs is calculated. Another small number L of these scans with the smallest residuals is chosen and then subjected to a second, modified fitting process. In practice, I tend to use values of M, N and L similar to:

$$M = 100 \tag{5.3}$$

$$N = 20 \tag{5.4}$$

$$L = 4 \tag{5.5}$$

The idea of this approach is to effectively sample over the parameter space described by N local fits by only actually using L fits. The local fits are sorted compared to their relative convergence, and fits that converge slowly are discarded. So, whilst we still sample over the space described by the N fits, we only expend most of our resources on the L fits that look most promising. Using the parameter set above, the calculation time of the global fitting routine can be reduced by a factor of five.

5.2.2 Jitter

It is also possible in some cases to speed up each of the L full local fits calculated at the end of the global fitting routine. The routine described in this section involves using fast random sampling of the data. Due to the nature of random sampling, the routine speeds the fitting in some but not all cases. In general, the likelihood of the routine failing increases the more complex the series of features seen in the data. The routine performs well on the original data, but not on

all of the more complicated data sets examined later in this chapter. When the routine performs successfully, calculation times for the full fitting routine can be speeded up by a factor of order two. In essence, the fast sampling is used to shunt the initial set of parameters applied to the full fitting routine closer to the local optimum before the fitting process begins, reducing the number of iterations required to reach convergence and therefore speeding up the calculation.

The fast fitting routine has calculation times two orders of magnitude faster than the slow routine. A small number m of points is chosen at random from the data. These points are then fitted using a local fitting routine that is terminated after a small number of iterations. This is a single step in the process. A second number m of points is then chosen, again at random, and the process repeated for n steps. I typically use values around:

$$m = 100 \tag{5.6}$$

$$n = 10 \tag{5.7}$$

Sampling only a few data points significantly speeds up the calculation time of the resulting fit. However, a lot of the detailed structure in the data is lost by only sampling a small portion of it. A fitting process based only on one small subset of points from the original data will converge on the local optimum of the full set of data in some dimensions relating to parameters that are important for the subset, but will quickly diverge in others. This routine deals with the problem of divergence by only calculating a few iterations of a fitting process for each subset of points and then making use of another, completely different set of points which has a different set of dimensions relating to convergence and divergence. The result is that, in parameter space, over the course of the routine the parameter set orbits around the local optimum of the full set of data and, in some cases, spirals inwards towards it. Of course, such a fitting process could never fully converge onto the local optimum due to the divergence introduced at every step. However, enough convergence can be achieved by using this ap-

proach to move the initial parameter set closer to the local optimum with only a small increase in calculation time. This significantly reduces the number of iterations required by the local fitting routine making use of the full set of points, resulting in a net decrease in calculation time of a factor of order two.

This 'orbiting' process is also beneficial if there are better local optima in the vicinity of the local optimum identified by the full local fitting routine. As the parameter set moves, it explores a larger region of parameter space than would be the case for the full routine. If the output detects a better local optimum during this process, the focus of the full fitting routine can be shunted from its original optimum onto this new optimum. The randomisation process ensures that this will only occur if the new local optimum is an improvement over the old; otherwise, the parameter set would latch back onto the original optimum when a new set of parameters was chosen.

5.2.3 The global optimum

The result of the global fitting process applied to this set of data is shown in figure 5.4. The set of features seen in the data is very well replicated by the output of the fit. The residuals are on the 20% level at worst and significantly smaller than this for most of the scan. We can be confident that the features seen in the data are due to atomic phenomena, and we can analyse the causes of these using the outputs of the Bloch equations simulation.

Before analysing this scan, we need to ensure that all the resonant features are fully resolved by the simulation. Otherwise, key features of the data might be missed that would aid in the interpretation. We therefore introduce a routine that actively detects resonant effects and changes the step-size of the calculation in their vicinity to ensure that they are resolved. This method is described in the following section.

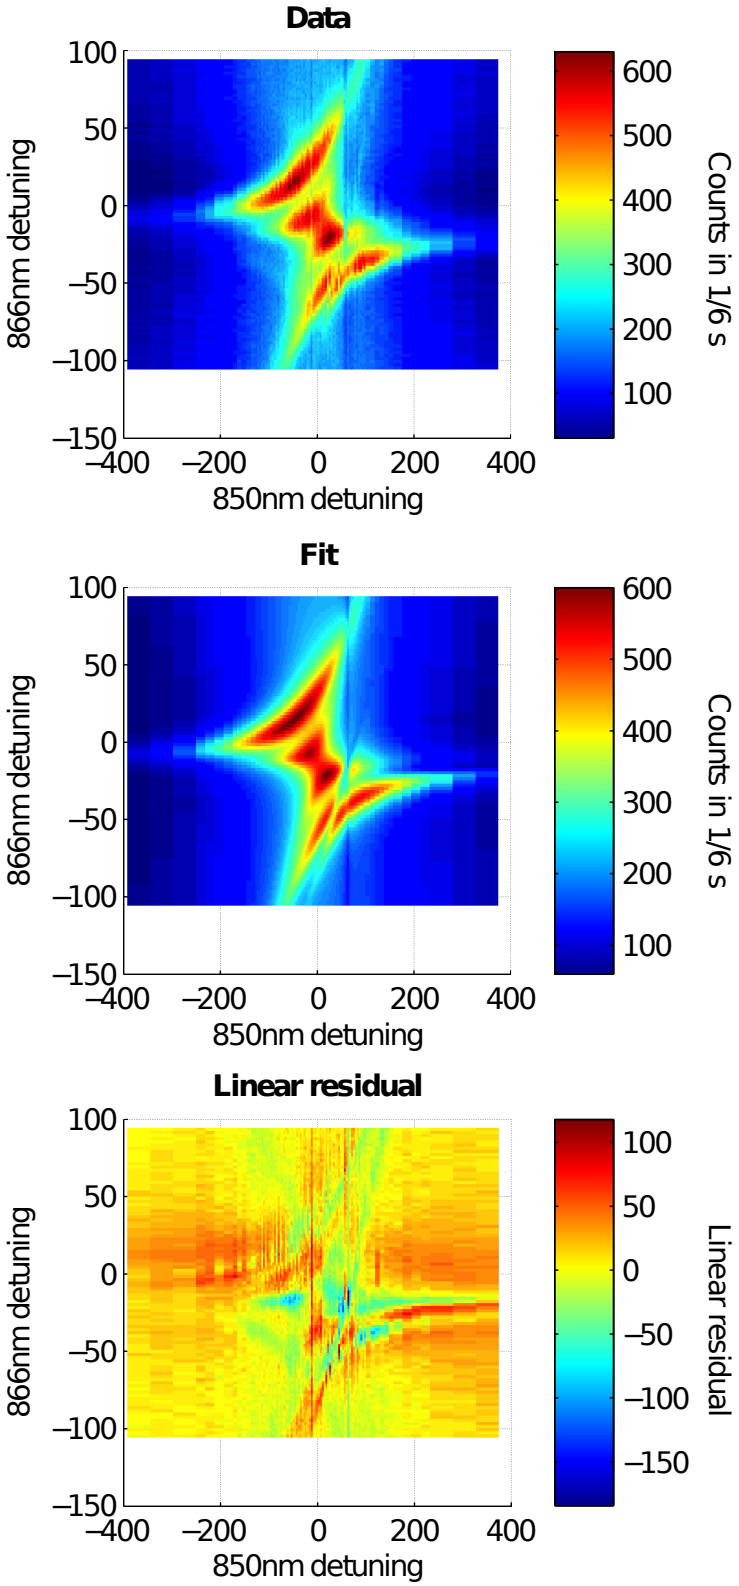

Figure 5.4: The results of the global fitting process applied to the original data. The features seen in the data are very well matched by the output. The errors are at worst on the 20% level and are much less than this for the vast majority of the scan.

5.3 Two-dimensional resonance detection

A resonance-detection scheme for data taken over a single dimension is described in section 3.3.5. In constructing a similar routine for two-dimensional data, there is an additional factor that we need to be aware of. The results of this routine are presented as a surface, so we need to ensure that the series of points chosen by the routine will be plotted well. If the surface is plotted poorly, than resonant effects may not be identified even if they are present in the output!

There are several possible extensions that could be made to add an additional dimension to the one-dimensional routine. However, many of these extensions will produce a set of points that are scattered over the surface following the contours of the data rather than using a fixed grid system. In order to plot this series of points, a triangulation routine must be employed. This poses a problem for this set of data. In general, scans of the type we are analysing in this chapter consist of sharp features with large gradients. Triangulation routines produce their results by looking at the projection of the surface onto the xy plane. In other words, it is assumed that the variations of the data in the z direction are small compared to the variations in the x and y directions. This assumption breaks down when there are large gradients in the surface being triangulated. In general, then, triangulation routines will produce poor results for the surfaces we are analysing, and the routines will produce poor output in exactly those portions of the scans we are interested in analysing! Triangulation routines will produce long, thin triangles in the vicinity of large gradients that obscure the features being studied. A more general triangulation routine would require an analytic expression for the height of the surface in order to function, and producing such an expression is difficult for curves such as the ones examined here with a large number of features.

Instead, we make use of a routine that ensures that all the points are laid out over a grid. Such a surface does not require triangulation to plot. The surface

is initially divided up into an $a \times b$ grid where a and b are both small. Each square of the grid is examined in turn. The four lines of points on the sides of each square are used to estimate the density of features within the square. The one-dimensional routine is applied to each side. The greatest point-density on each of the parallel sides is chosen and used to form a grid that extends over the square. Each square is then plotted as a separate surface, resulting in a connected, composite surface based on a grid layout where the density of points varies dependent on the level of detail present in the scan. a and b must be small enough that the level of detail in each square can be well approximated by the detail on its sides but large enough that the areas of the scan in which there are no features do not have too high a resolution associated with them. For the plots shown in this chapter, I generally choose $a = b = 100$. Figure 5.5 shows the output of this routine for the original data. This figure clearly shows all of the features in the scan. We are now in a position to interpret the resonant features seen in the original data.

5.4 Resonances in the original data

A general treatment of resonant behaviour in three, four and five level systems is too involved for the main body of this thesis and is given in appendix E. Let us consider what behaviour we would expect from this system.

For a five-level system with no Zeeman structure, we expect a bright resonance describing resonant effects in different subsets of levels within the system. This bright resonance takes the form of a 'modified hyperbola' with three asymptotes describing three-level, four-level and five-level type behaviour. We also expect a dark resonance connecting the three upper levels: the $D_{3/2}$, the $P_{3/2}$ and the $D_{5/2}$. These asymptotes and the dark resonance respectively lie along

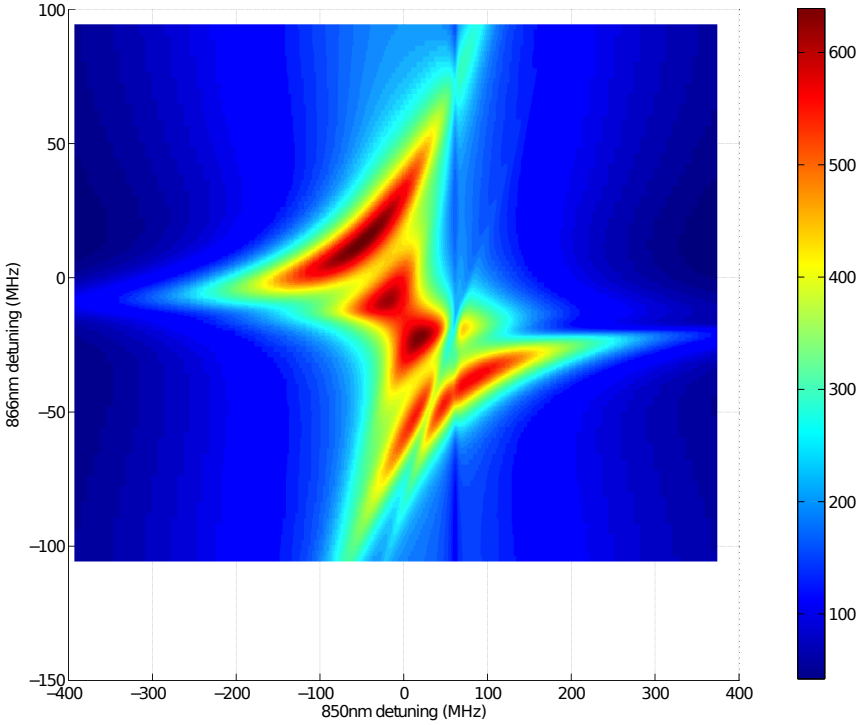

Figure 5.5: The output of the two-dimensional resonance-detection routine for the original data. All of the features of the data are clearly displayed.

the lines:

$$\alpha \quad : \quad \Delta_{866} = \Delta_{397} \tag{5.8}$$

$$\zeta \quad : \quad \Delta_{866} = \Delta_{850} + \Delta_{397} \tag{5.9}$$

$$\epsilon \quad : \quad \Delta_{866} = \Delta_{850} + \Delta_{397} - \Delta_{854} \tag{5.10}$$

$$\gamma \quad : \quad \Delta_{850} = \Delta_{854} \tag{5.11}$$

The equations for each line refer to the resonance conditions for coherences within the system. These coherent effects are shown on figure 5.6. The presence of Zeeman states splits the bright resonance into multiple arcs. The levels connected by the pump transition each consist of four Zeeman states, so we expect the resonance to split into at most four modified hyperbolae depending on the polarisation of the pump beam.

Figure 5.7 shows a high resolution two-dimensional scan of the original data

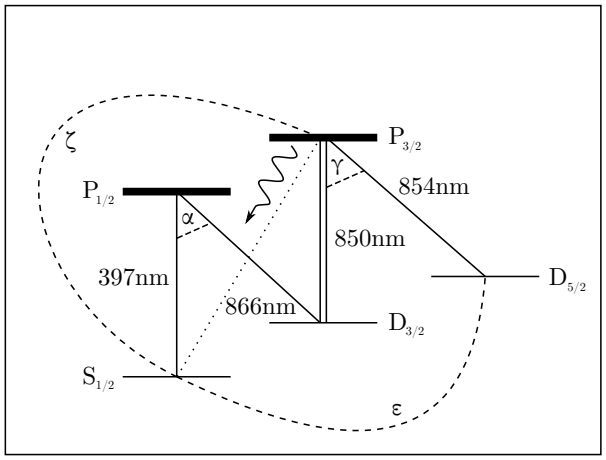

Figure 5.6: The important coherent effects in the original data (see text). States with large rates of spontaneous decay are shown by thick lines. The 850nm beam is much more intense than the other beams and is shown by a double line.

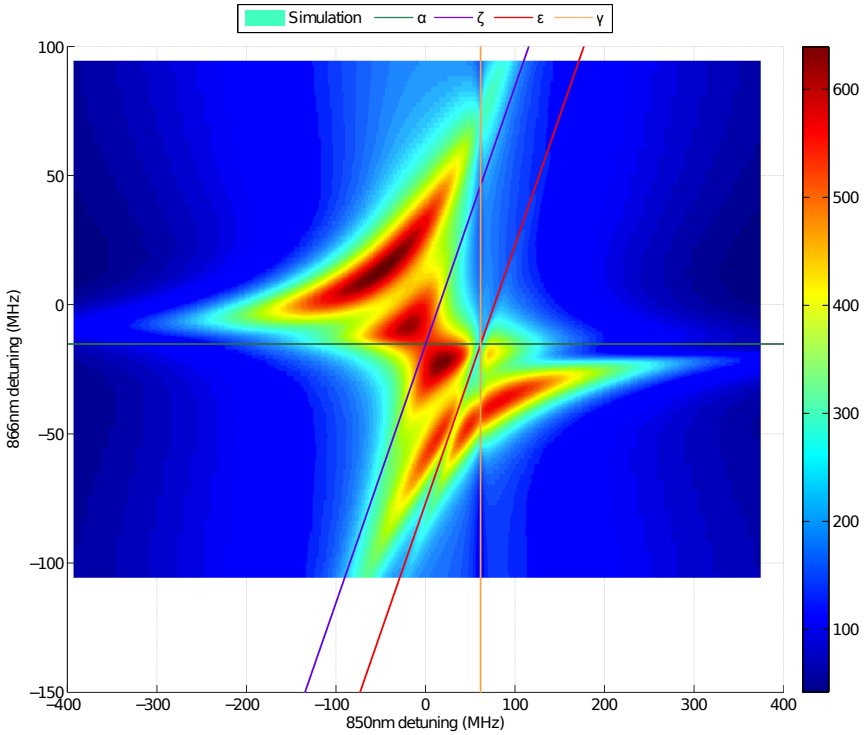

Figure 5.7: Applying the split five-level model to the original data. This plot shows the three asymptotes and dark resonance feature described in equations 5.8 to 5.11 overlaid over a high-resolution scan of the original data. The resonant features seen in this scan converge along lines that are described well by these asymptotes. In addition, a vertical line showing a drop in fluorescence over the entire scan is matched well by the dark resonance feature γ. 'Simulation' in the legend refers to the surface that uses the colour bar.

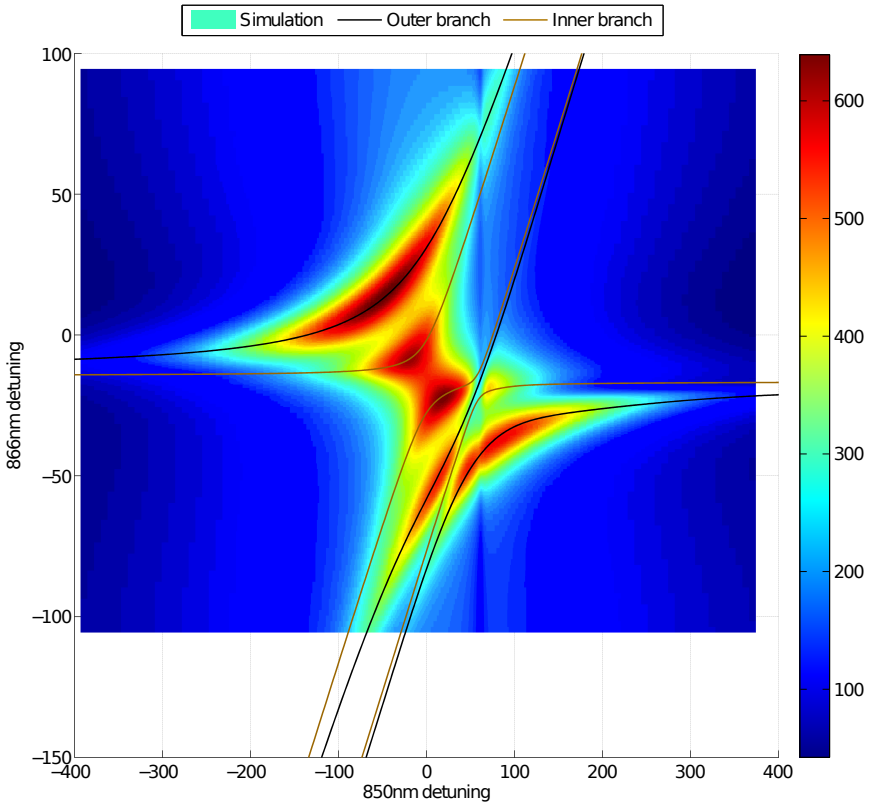

Figure 5.8: Applying the split five-level model to the original data. This plot shows two of the five-level modified hyperbolae described in section E.3 fitted to the fluorescence features in a high-resolution scan of the original data. These curves describe the position of the features very well, further supporting the claim that the split five-level model is a good model for the data.

overlaid by the three asymptotes and the dark resonance. The features of this scan converge to these asymptotes. Therefore, we can model the features in this scan using a pair of the modified hyperbolae described in section E.3. Figure 5.8 shows two such hyperbolae fitted to the two-dimensional scan. The curves describe the structure of the bright resonance extremely well. In addition, the fluorescence seen in the scan drops along a vertical line that is matched well by the position of the line γ. In summary, Zeeman-split five-level behaviour is a very good model for all the features seen in this scan.

We can interpret the behaviour as follows. The areas of high fluorescence in this scan are due to a bright resonance feature that coherently links all five levels in the system. This bright resonance converges asymptotically to three lines

which describe coherences driven in subsets of the system. When the 850nm beam is very far detuned from all the other beams, coherences are most strongly driven among the $S_{1/2}$, $P_{1/2}$ and $D_{3/2}$ levels and the bright resonance converges on the three-level resonance condition α for these three levels; the resonances converge to a horizontal asymptote. When the 850nm beam is close to resonance with the resonance condition ζ for the four lowest-energy levels of the system, coherences are strongly driven among these four levels. In this case, the resonance converges to the diagonal line ζ. Likewise, coherences are strongly driven among all five levels when the 850nm beam is close to resonance with coherent effects over the entire system. This is the cause of the third asymptote, described by the diagonal line ϵ. The bright resonance stretches between these four asymptotes. At any point, it describes how strongly coherences are driven in each of these three subsets.

The bright resonance is comprised of an inner and an outer pair of arcs. Each pair is described by one natural frequency at which the system is driven. Each resonant frequency produces a pair of arcs because the Autler-Townes splitting caused by the 850nm beam depends on the sign of the 850nm detuning relative to all the other beams.

Finally, the reduction in fluorescence in a vertical line around $75\mathrm{MHz}$ is due to a dark resonance being driven between the $D_{3/2}$, $P_{3/2}$ and $D_{5/2}$ levels. This resonance is strong as it coherently drives population that spontaneously decays from the $P_{1/2}$ on the 866nm transition.

5.5 Experimental analysis of five-level resonances

In the previous section, we fully described the features seen in the original data using a model of split five-level behaviour. To further test the validity of this model, a set of six additional two-dimensional scans was proposed. The parameters for these scans were chosen to demonstrate the contributions of the various

resonant effects leading to the full fluorescence structure. A series of scans was calculated at these predicted parameter sets, and a set of data was taken at nominally the same sets of parameters. In this section, we analyse the scans and compare them to the predictions of the model.

The split five-level model describes the shape of the feature seen in a two-dimensional scan using five parameters. The three asymptotes of the bright resonance feature are described by two parameters, Δ_{397} and Δ_{854}. The position of the vertical dark resonance is described only by Δ_{854}. The separation between the arcs of the bright resonance is dependent on the intensity of the 850nm beam. Finally, the splitting of the bright resonance depends on the polarisation of the 850nm beam, which is described by two parameters. To test the predictions of this model, we first take a standard base scan. Then, we take a series of subsequent scans varying each of the important parameters in turn, observing the effect of changing each parameter on the structure of the fluorescence scans. In this set of scans, we focused on the parameters describing the five-state structure; only the beam detunings and the 850nm intensity were altered. A further new set of over 20 scans has been taken recently which also examines the effect of the 850nm polarisation. See chapter 7.

Figure 5.9 shows a set of six scans taken to display variations in the five-state resonance structure. Scan 1 is the base scan. Each scan shares this set of parameters and varies one of the three important parameters. For these predictions, the base scan is the original set of data. Scan 2 alters the 397nm detuning. This detuning affects the position of all three asymptotes of the bright resonance. As a result, changing this parameter shifts the curve downwards along the Δ_{866} axis. In addition, the fluorescence seen in the scan is reduced. This is because population is pumped less efficiently out of the ground state, so less population reaches the $P_{3/2}$ level. Scans 3 and 4 increase and decrease the 850nm power respectively. This causes the bright resonance feature to expand and contract

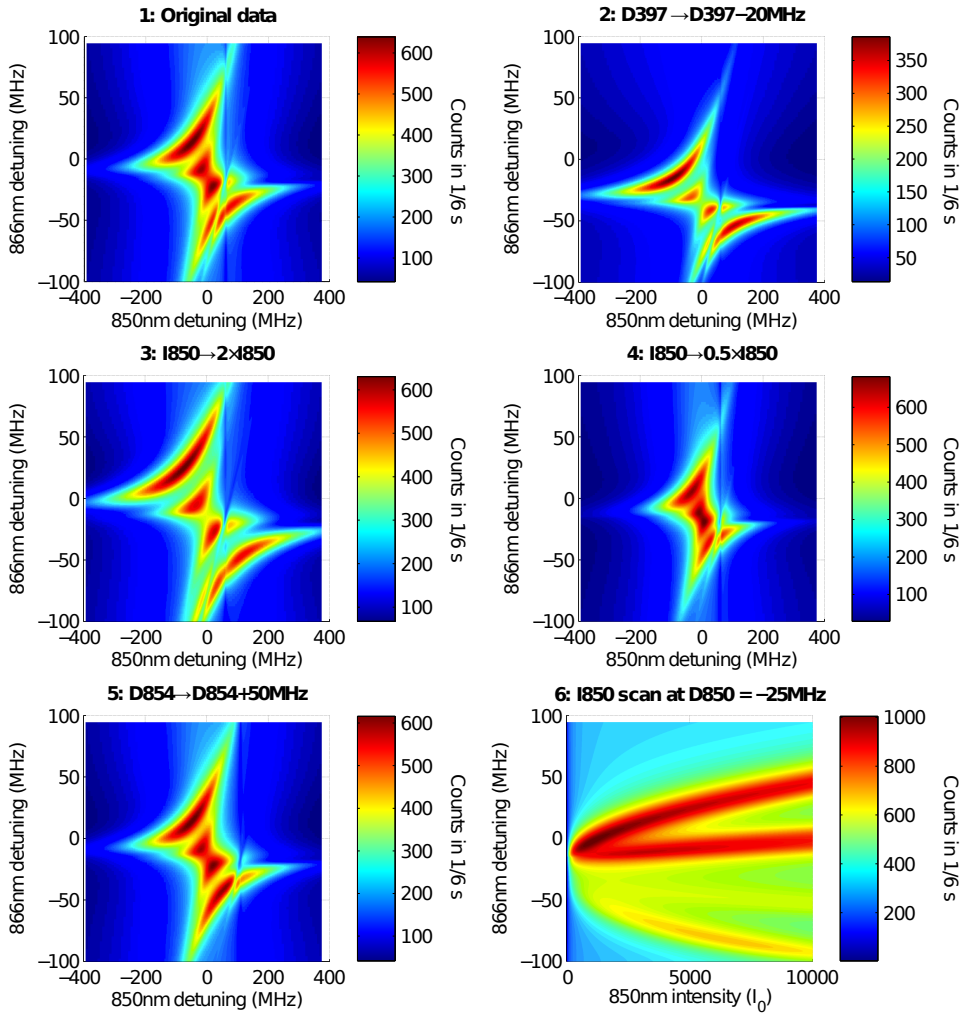

Figure 5.9: A series of scans designed to test the predictions of the split five-level model. Scan 1 is the original data. Each subsequent scan uses this parameter set, but varies one of the three parameters that characterises the five-state fluorescence structure. Scan 2 changes the 397nm detuning, translating the fluorescence structure in the vertical direction and reducing the fluorescence seen over the entire scan. Scans 3 and 4 change the 850nm intensity, increasing and decreasing the separation between the arcs of the modified hyperbolae respectively. Scan 5 changes the 854nm detuning, shifting the five-level diagonal asymptote and the vertical dark resonance. Scan 6 is a two-dimensional scan of the 866nm detuning and 850nm intensity, showing quantitatively how the separation of the arcs varies as a function of the pump intensity.

along its axis of symmetry. Scan 5 changes the 854nm detuning. This changes the position of both the five-level diagonal asymptote of the bright resonance and the vertical dark resonance.

Scan 6 is a two-dimensional scan taken with different axes. It shows the dependence of the fluorescence as a function of the 866nm detuning and the 850nm intensity. The aim of this scan is to examine the dependence of the arc separation on the 850nm intensity in a more quantitative manner than in scans 3 and 4. If the behaviour of all four arcs is well described by modified hyperbolae, then this separation should scale as the square root of the intensity. Note that, unlike the scans of intensity in appendix E, we are examining the position of the arcs along a vertical line rather than along the axis of symmetry; it is much simpler to change only one beam detuning in a scan than to try and change two simultaneously. The position measured in this manner will still show a square root dependence, but there will be an additional offset at low intensities due to the non-zero 850nm detuning.

Figure 5.10 shows the data taken at nominally this set of parameter values and figure 5.11 shows high-intensity scans taken from fits to the data. There are two immediate differences between the data and the predictions to mention. Firstly, the base scan 1 was taken at parameters only approximately describing the original data. The parameters were different enough from the original data to produce some slightly different features. These features are most clearly seen when the separation of the arcs is large compared to their width as in scans 2 and 3. In this case, it appears that the bright resonance is split not into two hyperbolae but instead into four or even five. However, this in as illusion; both levels directly connected by the pump transition involve four Zeeman states, so splitting arising from the transition only can only produce four separate hyperbolae. The splitting seen in this data is in fact due to a combination of factors. The outer and inner portions of the bright resonance are both slightly split by

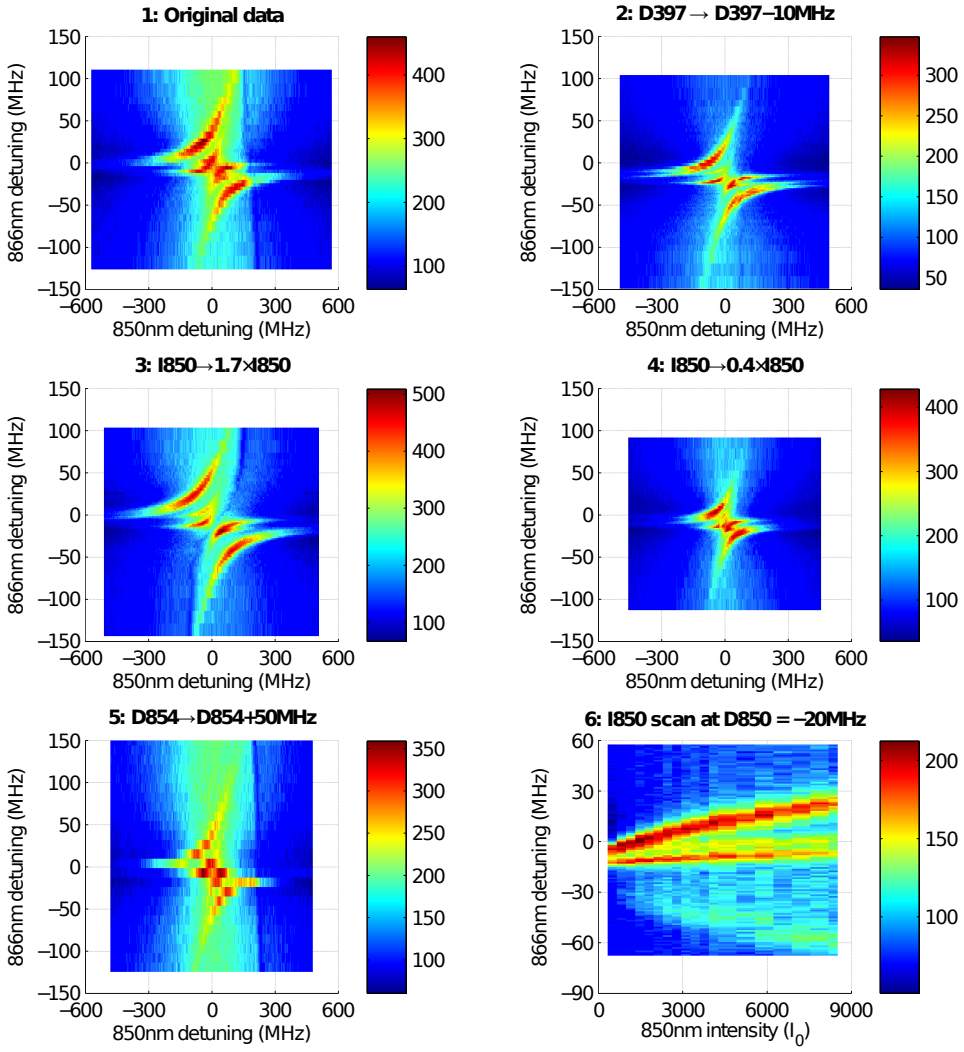

Figure 5.10: A set of data taken to test the predictions of the split five-level model. These six scans were taken at nominally the same sets of parameters as the equivalent predictions in figure 5.9. The parameter variations written above each scan are taken from fits to this data. All these scans reproduce the behaviour demonstrated by the set of predictions. For details of the differences between scan 1 in this figure and the scan 1 in figure 5.9, including the additional splitting and the shape of the dark resonance, see text. High-resolution scans of the data taken using these fits are shown in figure 5.11.

approximately the same amount. This splitting is due to the 397nm power. For this set of data, this power was slightly higher than for the original set, causing an additional splitting in the hyperbola due to the Autler-Townes effect. This results in a mirroring of the hyperbola in the vertical direction: see figures E.8 and E.14. This effect produces the small splitting in each arm of the bright resonance. However, this does not account for all of the splitting. The polarisation of the 850nm beam in this instance is also slightly different from the original data. Instead of splitting into two hyperbolae, the polarisation actually splits the curve into three in a similar manner to that seen in figure E.19. The split arcs of each modified hyperbola interfere, so that in general five split arcs will be seen rather than six. These five arcs are most clearly seen in high-resolution scan 2.

One further feature seen in this data arises from the beam setup. These sets of data are generally higher resolution than the original data. As a result, the scans generally took several times longer to run. Both the high resolution and the increased running time led to one feature of the optical setup being emphasised that was not visible in the original data. When this set of data was taken, all lasers except for the 854nm beam were frequency-locked. The drift of the 854nm frequency over time was small enough not be observed for the first set of data. However, for these scans, the 854nm frequency drift is quite marked. The effect of this drift is most clearly seen in alterations to the vertical dark resonance. The drift changes the resonance from following a vertical line to following a curve. These scans were taken in a similar manner to the original data by stitching together a series of 866nm scans taken at different 850nm detunings. The difference in 850nm detuning between each scan was modified at several points in the process to ensure that all the detail in the scan was recorded at a high resolution. The drift of the 854nm occurs at approximately the same rate throughout the process. As a result, the dark resonance will follow a piecewise line. This effect is most pronounced in scan 3, where the dark resonance

moves over a large portion of the scan. The newer set of scans was taken with the 854nm beam frequency locked, and these are currently being analysed.

This drift is in fact beneficial for our analysis. Our aim is to demonstrate that the original model can account for this data. The drift in the position of the dark resonance is a direct demonstration that the position of the resonance feature depends on the 854nm detuning. The fitting process for this set of scans was modified to take account of the 854nm drift. The drift between each 866nm scan was assumed to be constant. In the high-resolution scans, the piecewise line followed by the dark resonance in the data was fitted and used to adjust the 854nm detuning as a function of the 850nm intensity. This model of the drift matches the position of the dark resonance well for the majority of these scans. However, there are a few places where the 854nm detuning changes sharply that cannot be modelled using this approach. One example is in scan 2, where the positions of the five-level diagonal resonance and the vertical resonance cannot be reconciled unless the 854nm detuning changes sharply; the vertical resonance intercepts with the horizontal asymptote in a different place from the diagonal resonance, something which cannot occur unless there is a large step-change in the 854nm detuning between 866nm scans. However, this is a slight discrepancy; for the majority of the scans, the original model can still describe this data well even when the 854nm is allowed to drift. The scans fit the data very well, with residuals on the 10% level. The residuals for each of these six scans are shown in appendix D.

The relative changes in each of these scans relative to the base scan follow the differences seen in the original predictions (shown in figure 5.9). Scan 2 translates the base scan in the vertical direction and lowers the overall fluorescence. Scans 3 and 4 expand and contract the fluorescence arcs along their axis of symmetry. Scan 5 shifts the position of the five-level diagonal asymptote and the vertical resonance. The numerical values of the parameter shifts shown in the

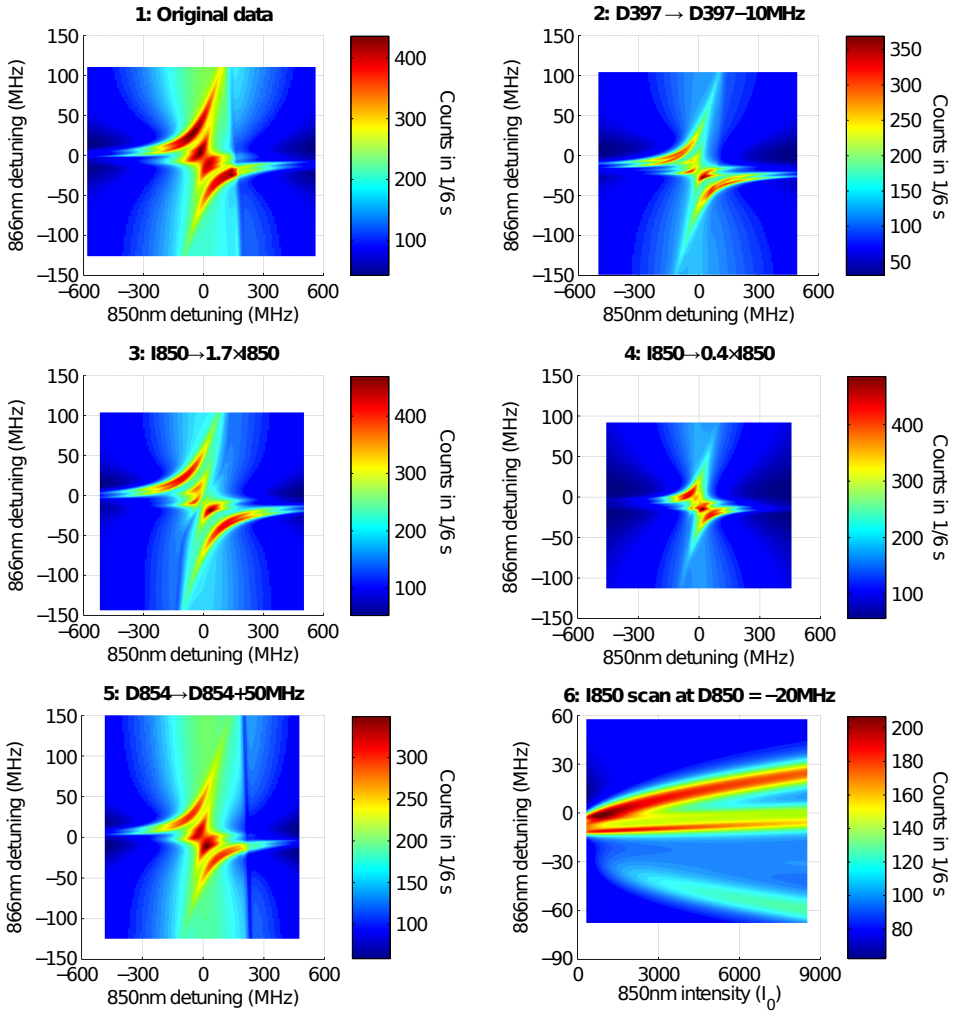

Figure 5.11: High-resolution scans of a set of data taken to test the predictions of the split five-level model. These scans are fits to the data shown in figure 5.10, and are nominally taken at the same set of parameters for the equivalent set shown in figure 5.9. The value of the 854nm detuning was fitted to a piecewise linear function assuming a constant drift between each of the 866nm scans that make up a two-dimensional scan. The changes in the structure in each feature closely match those of the predictions.

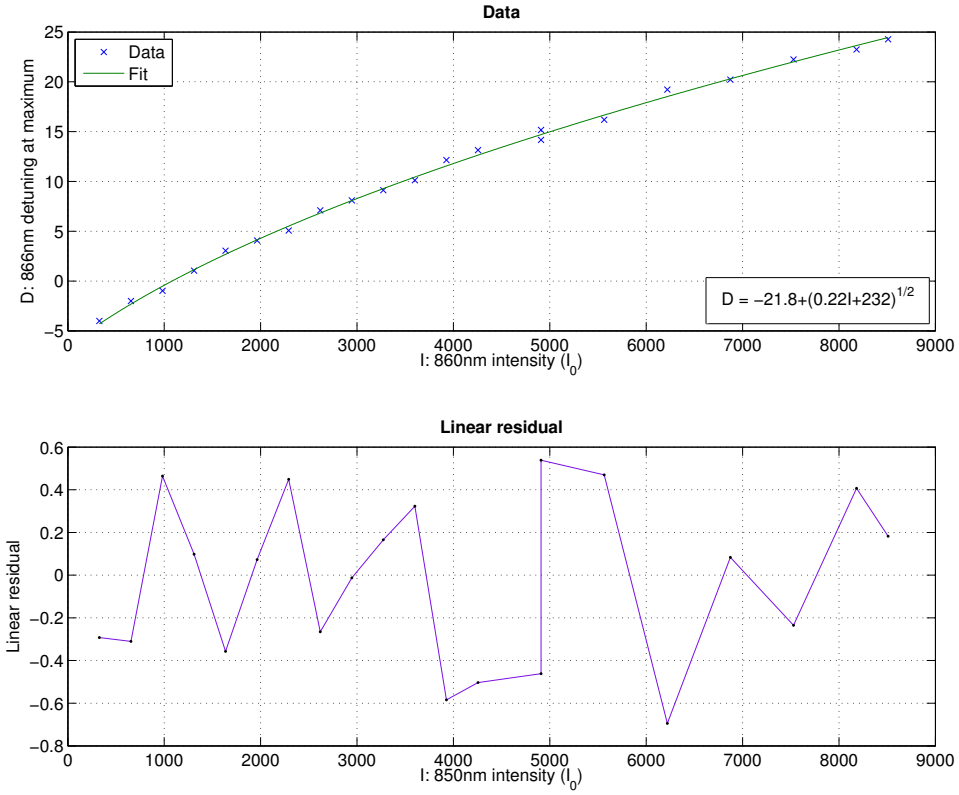

Figure 5.12: The position of the arc of the bright resonance involving the largest positive detuning of the 866nm beam as a function of the 850nm power. The position of the arc was calculated by taking the point of highest fluorescence for each 866nm scan comprising scan 6 of the data. This scan is shown in figure 5.10. Two scans were taken at $I \approx 5000I_0$. The position of the curve is fitted well by a square root dependence. We expect this dependence from the modified hyperbolae that describe bright resonances in the split five-level model. The residuals of the fitting process are of order 10% largely due to noise. A fit to the high resolution simulation has residuals an order of magnitude better than this. The curve has a non-zero offset as this scan is taken only varying the 866nm detuning. Scans of this type in appendix E have varied the frequencies of both beams simultaneously and looked at the position of the arcs along their axis of symmetry. That approach is impractical in experiment, so a scan of only a single detuning was used.

plots are taken from the output of the fitting routine.

Finally, we examine scan 6, which shows how the separation between the arcs varies as a function of the 850nm intensity. This structure seen in the scan is very similar to that in the prediction. The model predicts that the arcs in this plot will show a square-root dependence on the intensity. Figure 5.12 shows the position of the maximum fluorescence of the strong upper arc of this plot. This position is fitted well by a square-root dependence, with residuals again on the 10% level. This agrees with a hyperbolic description of this region of the arcs.

The predictions of the split five-level model are matched very well by the data taken in this set of experimental scans. We can therefore be confident that the model is a very good description of the behaviour of a five-level system.

6

A qubit in $^{43}Ca^{+}$

> Tchicaya nodded. She took him by the hands and lifted him up. He constructed a safe route through the processor and she ejected the tainted qubits, forming a tiny bubble of classical physics in the vendeks' quantum sea.
>
> *Schild's Ladder* by Greg Egan

In the previous two chapters, we modelled the temperature of a $^{40}Ca^{+}$ ion and analysed the resonant features present in three, four and five-level atomic systems. In this chapter, we analyse the temperature and resonant behaviour of $^{43}Ca^{+}$. This isotope that is a lot more complicated to work with than $^{40}Ca^{+}$, both theoretically and experimentally. However, this isotope posses qubit states that are first-order insensitive to magnetic field nose. By using such qubits, it is possible to increase the fidelity of all qubit operations by significantly reducing one source of decoherence. In order to perform gates on them, we need to cool the ion to the ground state of the trap so that we have complete control over the motional degrees of freedom. The technique used to achieve this is resolved sideband cooling (see [Har13] and [IBBW95] for a review). However, in order to use that method, we must first Doppler cool the ion to of order the Doppler limit.

With this aim, we studied Doppler cooling in $^{43}Ca^{+}$ and whether this level of cooling was possible. We set up the Bloch equations for this isotope and

solved both the time-dependent and time-independent problem. Different cooling schemes were proposed and analysed to examine the lowest feasible temperatures that can be achieved using Doppler cooling. We fit fluorescence scans to determine experimental parameters. For different experimental conditions, we were able to predict the ion's temperature and compare our results with data in the light of experimental limitations. In this chapter, we almost exclusively make use of the dynamic method to model the ion's temperature. This is the most detailed routine available to us that can compute on realistic timescales. While the motional method is more advanced, it calculates orders of magnitude more slowly than the dynamic method and is infeasible to use for this isotope.

We now describe $^{43}Ca^{+}$ and the field-insensitive qubits in more detail. This isotope has a non-zero nuclear spin $I = 7/2$ and therefore has a significantly more complex Zeeman structure than $^{40}Ca^{+}$. For example, in the three-level atomic system analysed in chapter 4, the number of states in $^{40}Ca^{+}$ is 8, whilst the number of states in $^{43}Ca^{+}$ is 64. The difference is even more pronounced for the five-level systems examined in chapter 5; the number of states involved for $^{40}Ca^{+}$ 18, whilst for $^{43}Ca^{+}$ the problem involves 144 states and over 20000 coherences.

The complexity of $^{43}Ca^{+}$ allows the ion to exhibit a richer variety of behaviour than the simpler isotope. If offers two very significant advantages for quantum information processing. First, a magnetic field causes mixing of the hyperfine states, leading to non-linear dependence of the state energies on the magnitude of $\mathbf{B}$. As a result, the energy E of some transitions at certain fields B will have a zero gradient dE/dB. At these points, the ion is first-order insensitive to magnitude field fluctuations. These 'atomic clock' transitions [HAB+14] significantly reduce the ion's sensitivity to one of the largest sources of decoherence for an ion qubit. A qubit on an atomic clock transition was first demonstrated in $^{9}Be^{+}$, a two level system [LOJ+05]. The work in this chapter describes

the first realisation of an atomic clock qubit in a three level system.

The second advantage arises from the large transition frequencies that are possible between states in the same energy level. In $^{43}Ca^{+}$, states in the same level can be separated by frequencies of over 3.2GHz. Such large separations allow for certain states to be selectively excited with a very high fidelity without any of the technical challenges arising from working with transitions that are not dipole-allowed. As a result, readout of a $^{43}Ca^{+}$ at a low magnetic field of order 2G has been observed with a fidelity of 99.9(1)% [BHL^{+}15, Bal14].

From a simulation point of view, the disadvantages of using a system with complex Zeeman structure are twofold. Firstly, a larger system is a more demanding computational problem. The number of elements in the matrix we use to encode the Bloch equations scales as the fourth power of the number of states. Many of the memory and time-saving approaches discussed in chapter 3 were implemented to allow calculations involving this isotope to be performed on standard hardware and on practical timescales.

Secondly, the size of the problem means that the atomic dynamics are very complicated. The fluorescence and temperature of an ion with a single dark resonance were discussed in chapter 4. The dynamics of $^{43}Ca^{+}$ involve thousands of dark resonances, and a single fluorescence scan from the ion takes the form of a forest of peaks and troughs. Problems such as maximising the ion's fluorescence or cooling to the ion to around the Doppler limit are highly non-intuitive and are generally approached using fitting and optimisation procedures. These can give significant insight into the behaviour of this isotope. For example, we found that if one exploits an appropriately chosen dark resonance, $^{43}Ca^{+}$ behaves like an ideal three-state system and can be cooled to temperatures below the Doppler limit.

The analysis in this chapter led to the first experimental realisation of an atomic clock transition in $^{43}Ca^{+}$. This qubit is the first trapped-ion qubit with

errors in initialisation, readout and gates that are all below the fault-tolerant threshold [Ste03] in the same set of apparatus [HAB+14]. This qubit has a combined state-preparation and readout error of 99.93%, a memory coherence time of $T_2^* = 50$s and an average single-qubit gate fidelity of 99.9999%. The qubit was Doppler cooled to a temperature 30% lower than the Doppler limit for the relevant single-photon transition [AHS+16]. Sideband cooling to the ground state of the trap from this starting point is straightforward.

Section 6.1 introduces the labelling conventions used to distinguish between intermediate field states in $^{43}Ca^{+}$. Section 6.2 describes the method used to search for field-insensitive qubits. The ground-state hyperfine microwave field-insensitive qubits in $^{43}Ca^{+}$ are discussed, and typical values of fluorescence and ion temperature that can be obtained with these are demonstrated. In section 6.3, methods used for optimising the temperature of an ion are discussed, and results of brute-force optimisation for the 146.09 gauss qubit are presented. In section 6.4, I describe the cooling scheme that allows for cooling below the Doppler limit to be observed in this isotope. This makes use of one of two coherences in the three-level system. Calculations are discussed that indicate that one coherence is more preferable to use. In addition, the variation of the optimal temperature with beam parameters is demonstrated. In section 6.5, a different repumping scheme is presented that allows for similar temperatures to the previous scheme to be obtained while increasing the fluorescence rate by over a factor of two. Section 6.6 describes a cooling scheme that makes use of five levels in the isotope and limits coherent behaviour. Section 6.7 shows fits to series of 866nm scans, including temperature predictions, and discusses the evidence for heating present in such scans. Section 6.8 shows the results of a series of scans in which Doppler cooling below the Doppler limit was both predicted by simulation and explicitly measured using sideband thermometry. Section 6.9 demonstrates that cooling below the Doppler limit can also be achieved for a

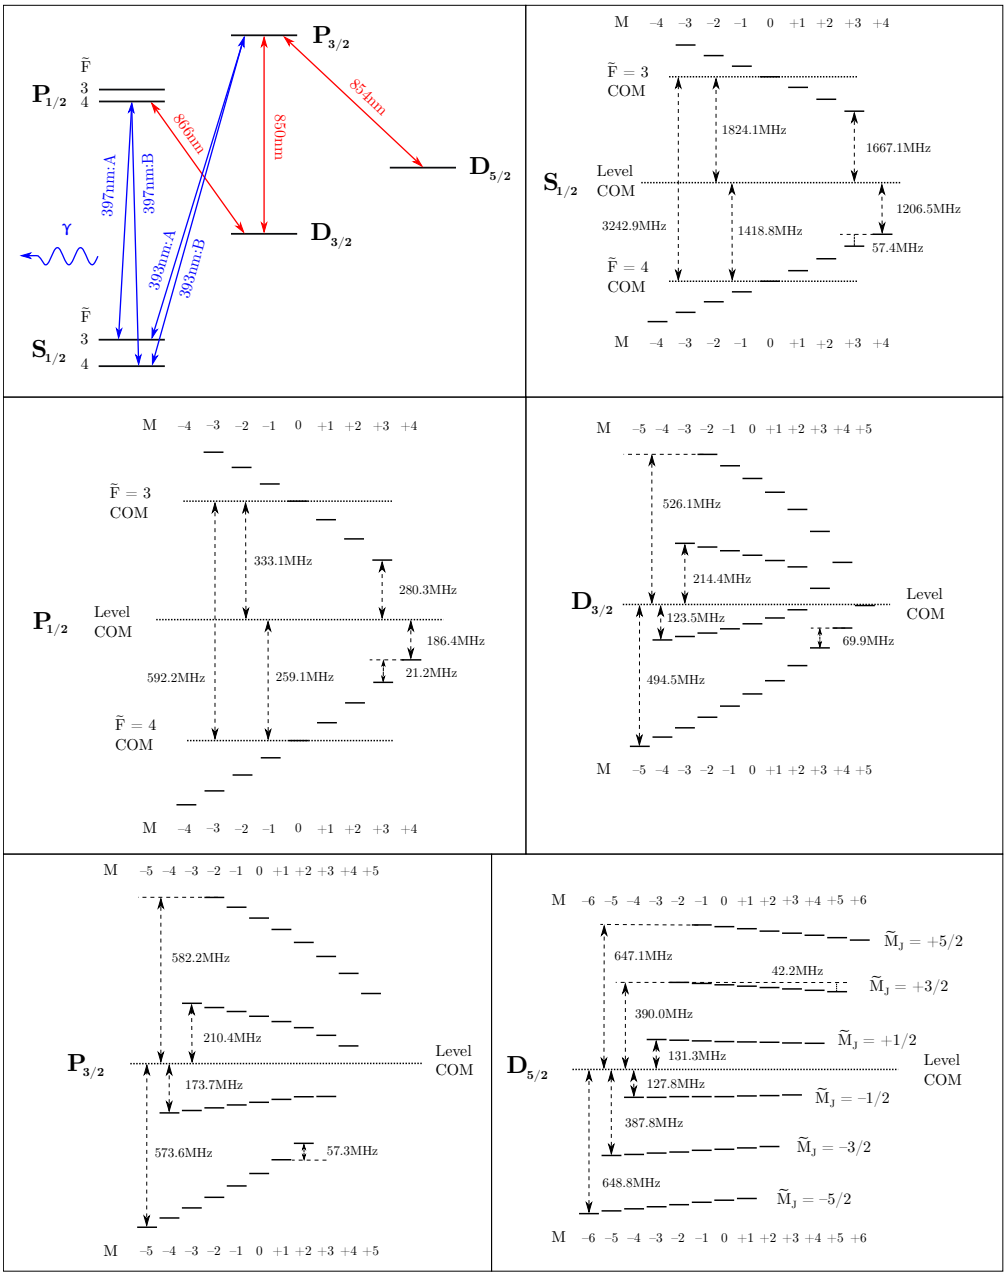

Figure 6.1: A schematic of the five lowest-energy levels of $^{43}Ca^{+}$ at a field of 146.09G and the Zeeman states in each level. The energies of the levels are not to scale. The states are ordered by the value of the angular momentum projection M. The state energies are shown to scale with the exception of the spacing between the two groups of states in the $S_{1/2}$ and $P_{1/2}$ levels; these spacings have been scaled down by a factor of five to ensure that the detail within each separate manifold can be seen clearly. The frequencies of characteristic transitions between states in each level are shown explicitly. For a discussion of the labelling conventions of the states and beams, see text.

field-insensitive qubit on a π transition at 287.78G. This qubit may be used in further work. Appendix G briefly discusses the use of the optical Bloch equations to simulate the initialisation and readout of a $^{43}Ca^{+}$ qubit. Appendix H lists the parameters used to produce the figures.

6.1 Labelling conventions

The Zeeman structure for the five lowest energy levels of $^{43}Ca^{+}$ is shown in figure 6.1. This structure is shown at a **B** field with a magnitude $|\mathbf{B}| = 146.09$ gauss, the field of the particular atomic clock transition we will discuss for the majority of this chapter. This section describes the notation used for labelling the states and beams.

Neither the low-field nor the high-field hyperfine bases are good descriptions of the Zeeman structure; within each level the states can only be described by a single good quantum number M. This quantum number alone is not enough to uniquely identify each state. However, certain groups of states can still approximately be described by the low or high-field quantum numbers. Figure 6.2 shows the energies of the Zeeman states in these five levels as a function of $|\mathbf{B}|$. At a field of 146.09 gauss, the states in the $S_{1/2}$ and $P_{1/2}$ levels form two distinct groups of states. The ordering of the states within these groups is the same as for the limit of low field. We describe each of these groups of states as a 'manifold' and label them according to the low-field quantum number associated with the groups, F. To make it clear that this is a label, rather than a good quantum number, we write the label using a superscript tilde: $\widetilde{F}$. We therefore uniquely describe states in the $S_{1/2}$ and $P_{1/2}$ using the notation $|\widetilde{F}, M\rangle$. States in the $D_{5/2}$ also group into manifolds, although these are ordered by the high-field number $\widetilde{M}_J$. We describe states in this level using the notation $|\widetilde{M}_J, M\rangle$. The states in the $D_{3/2}$ and $P_{3/2}$ levels cannot unambiguously be identified by groupings of either low or high-field states. We distinguish the states by their

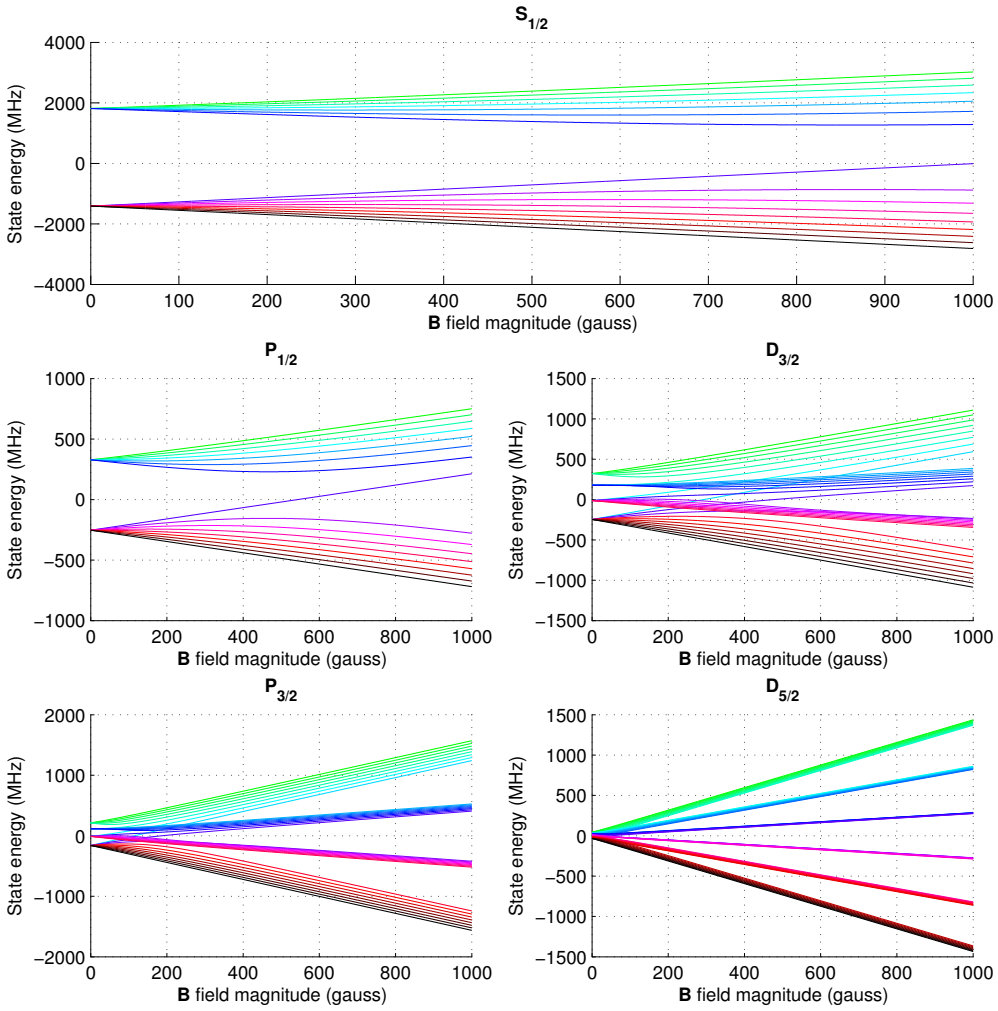

Figure 6.2: The energies of the Zeeman states of the five lowest energy levels of $^{43}Ca^{+}$ as a function of **B** field magnitude. The states are coloured from green to black according to the ordering of the states from high to low energy in the high-field regime. This is to demonstrate that the relative ordering of the states changes significantly when moving from the high-field to the low-field regime. This change in order is most easily seen in the plots for the $D_{3/2}$ and $P_{3/2}$ levels which show both extremes of behaviour.

energy E using the notation $|E, M\rangle$. One exception to this labelling scheme occurs for the two states in each level with the extreme values of M. These states are uniquely described by the value of M and are labelled only by this quantum number.

The energy difference between the two manifolds in the $S_{1/2}$ level is sufficiently large that a beam resonantly driving transitions from one manifold will

negligibly excite transitions from the other. To excite transitions from both manifolds simultaneously, we split each beam interacting with the level into multiple frequencies. The carrier and one sideband are each set to resonantly drive transitions from one of the manifolds. I label these beams alphabetically from the lowest to the highest frequency. I model each beam as a separate entity that only excites transitions from its target manifold. This avoids the complications of multiply-connected states which significantly complicates the structure of the Bloch equations; see section 2.1.1 and figure 2.1. The energy separation between the manifolds is so large compared to the energy difference between individual states that this introduces a negligible error.

6.2 Atomic clock transitions in ^{43}Ca^{+}

We define an 'atomic clock' transition as a transition between any two Zeeman states with a frequency $\nu(B)$ which, at a particular field B_0, has a gradient $(d\nu/dB)_{B=B_0} = 0$. At this 'field-insensitive point', the transition is first-order insensitive to magnetic field fluctuations. I will refer to B_0 as the 'critical field' of the transition and the frequency $\nu(B_0)$ as the 'critical frequency'.

A simple test can be performed to check whether any given transition has a field-insensitive point. The transition frequencies are the difference of two eigenvalues of the Hamiltonian. These eigenvalues vary smoothly as a function of B as the eigenstates rotate from the low-field basis states to the high-field states. As a consequence of this, the gradient of the transition frequency must vary monotonically. Any given transition can therefore have either no field-insensitive point or exactly one. A transition with a field-insensitive point can therefore be identified by calculating the gradient $d\nu/dB$ at zero field and at a very large field. If the signs of these two gradients differ, the transition has exactly one field-insensitive point that can be determined using a standard binary search.

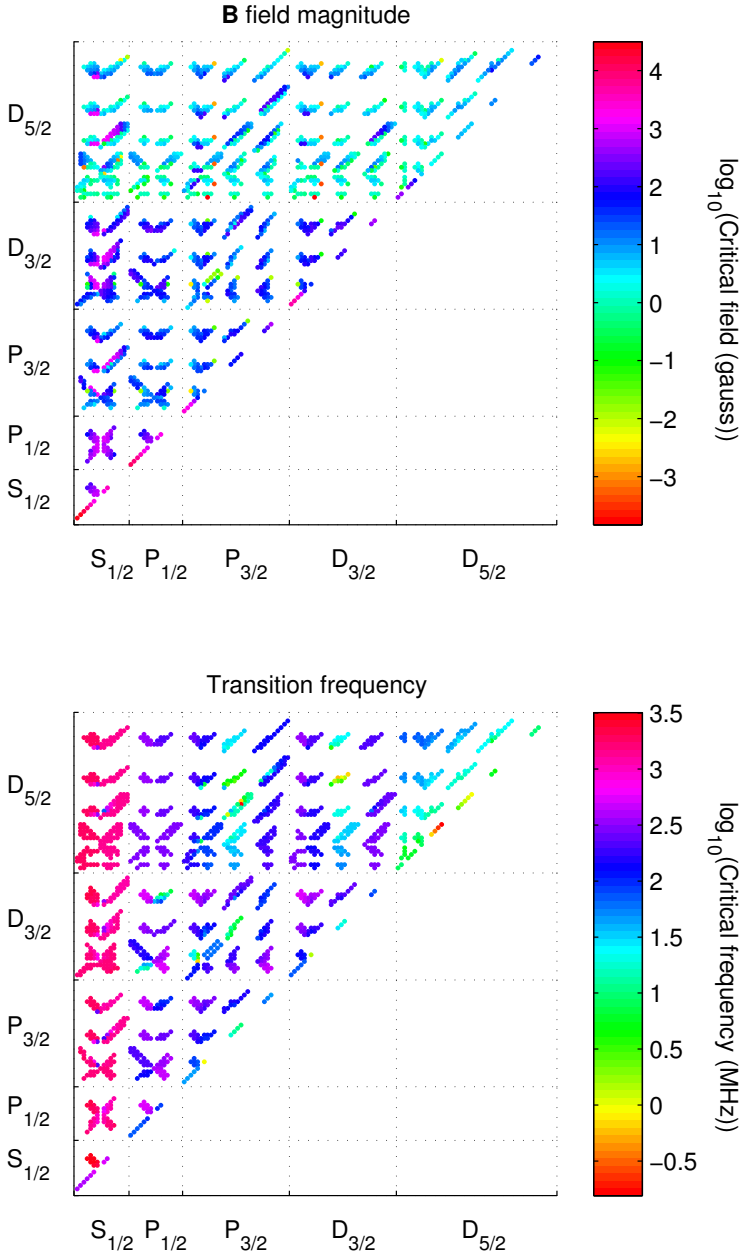

Figure 6.3: A plot of the field-insensitive transitions in the five lowest energy levels of $^{43}Ca^{+}$. This plot demonstrates that there are many transitions of this type between states in all five of the lowest energy levels with a very large range of both critical fields and frequencies. The upper plot shows the critical field and the lower plot shows the transition frequency. The transition frequency is shown without the fine-structure energy difference between levels in order to highlight small differences in frequency between different transitions. Only the upper-left triangular portion of the plots are shown for clarity: the lower-right portions are simply mirrored. Each point represents a transition between one state in each level shown on the axes. The states are ordered in each level according to their energy; the state energy decreases along the positive direction of each axis. The ordering is determined separately for each individual point at its critical field. The maximum change in M for a given transition is set by the change in L. These calculations assumed that LS coupling is still valid. The Paschen-Beck effect will only have a significant effect for $^{43}Ca^{+}$ for fields four orders of magnitude larger than those shown here.

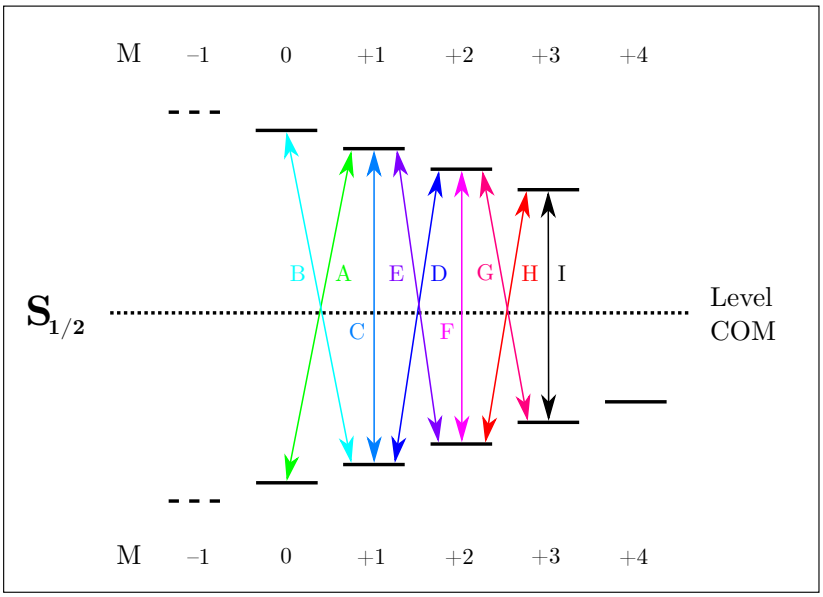

Figure 6.4: The nine hyperfine atomic clock transitions in $^{43}Ca^{+}$. The transitions are labelled alphabetically and by colour from green to black in order of ascending critical field. The transition $|\widetilde{F}=3, M=+3\rangle \leftrightarrow |M=+4\rangle$ does not have a field-insensitive point. For more information, see text.

Figure 6.3 shows critical fields and frequencies for all the field-insensitive transitions between the five lowest energy levels of $^{43}Ca^{+}$. Atomic clock transitions occur for a wide range of fields and frequencies, spanning from radio-wave transitions between adjacent M states through microwave transitions between hyperfine manifolds to optical transitions between different levels.

The work we discuss in this chapter involves using an atomic clock transition as a field-insensitive qubit. We restrict this qubit to the hyperfine transitions between the hyperfine manifolds of the ground state. This choice has the advantage of long coherence times [HAB+14] and high-fidelity readout due to the large energy separation between the manifolds [MSW+08, BHL+15] without the complications of narrow-linewidth or high power lasers required for qubit schemes on optical transitions (see, for example, [BHR+04]).

There are in total nine ground-state hyperfine atomic clock transitions in $^{43}Ca^{+}$. The states involved in these transitions are shown in figure 6.4 and the transition frequencies and critical fields are shown in figure 6.5. The atomic

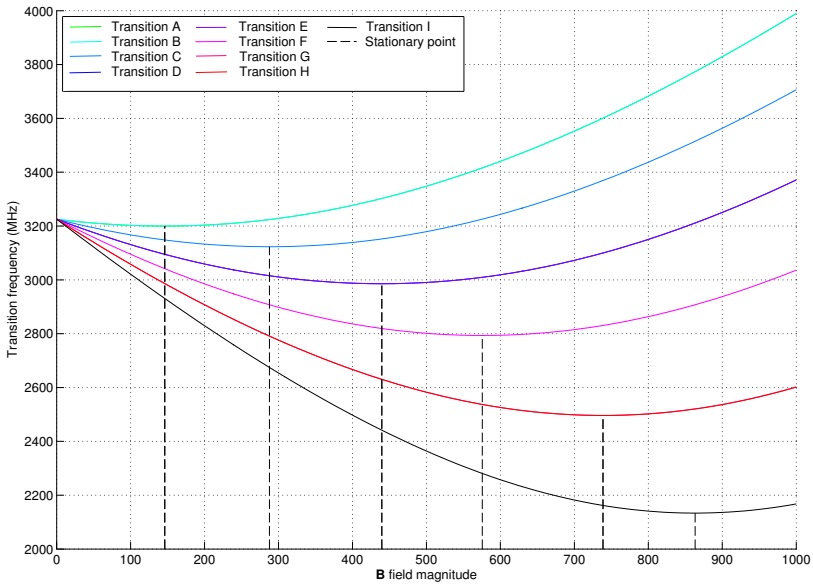

Figure 6.5: The transition frequencies and critical fields of the nine hyperfine atomic clock transitions in $^{43}Ca^{+}$. The field-insensitive point for each transition is marked by a vertical dashed black line. Clock transitions connecting states with the same total M have very similar frequencies and critical fields which cannot be resolved in this plot.

clock transitions connect the states with non-negative values of M. The one exception is the 'stretch state' $|\widetilde{F} = 3, M = +3\rangle \leftrightarrow |M = +4\rangle$. This state does not have a field-insensitive point, but instead tangentially approaches zero gradient as $|\mathbf{B}| \to \infty$. The behaviour of this state differs because it involves a state with a unique value of M. This state never mixes with any other state and has a transition frequency that varies linearly with $|\mathbf{B}|$. At high field, all other states in the same manifold tangentially approach the same gradient. As both $|\widetilde{F} = 3, M = +3\rangle$ and $|M = +4\rangle$ are in the same high-field manifold, the field-insensitive point for this transition must therefore occur at infinity.

The critical frequency, critical field and second-order sensitivity to the field for these transitions are listed in table 6.1. The fundamental, hyperfine and atomic constants used to calculate these values are listed in table 6.2.

Figure 6.6 shows characteristic fluorescence rates and temperatures that can be achieved from $^{43}Ca^{+}$ as a function of the **B** field magnitude using a three-level cooling scheme. This figure is the result of a simple optimisation process,

Qubit	$\widetilde{F}=3:M$	$\widetilde{F}:M=4$	$\lvert\mathbf{B}\rvert$ (G)	ν (MHz)	$\frac{d^2\nu}{d\mathbf{B}^2}(\mathrm{kHz}\cdot\mathrm{G}^{-2})$
A	+1	0	146.094151	3199.94108	-2.41560910
B	0	+1	146.331367	3199.85731	-2.41534531
C	+1	+1	287.782712	3123.18179	-2.51408134
D	+2	+1	439.582262	2985.79669	-2.58992317
E	+1	+2	439.803476	2985.54477	-2.58993661
F	+2	+2	575.565453	2793.45872	-2.81080616
G	+3	+2	738.637710	2496.48855	-3.11929489
H	+2	+3	738.821387	2496.06530	-3.11937664
I	+3	+3	863.348194	2133.53934	-3.68021026

Table 6.1: Parameters describing the nine microwave-frequency atomic clock transitions connecting the two hyperfine manifolds in the $S_{1/2}$. The transitions are labelled by the values of M for the states they connect. The critical field, critical frequency and second-order **B** field dependence for each qubit is shown.

Constant	Value
$A_J(S_{1/2})$	-806.40207160MHz [ABG+94]
$g_J(S_{1/2})$	2.00225664 [TPR+03]
μ_N	7.62259357MHz.T^{-1} [MTN12]
μ_I/μ_N	-1.31537 [Ols72]
μ_B	9.27400968$\times10^{-24}$J.T^{-1} [MTN12]
h	6.62606957$\times10^{-34}$Js [MTN12]
I	7/2

Table 6.2: The values of the fundamental, hyperfine and atomic constants used to calculate the fields and frequencies of the atomic clock transitions shown in table 6.1.

and cooling schemes will be examined in much greater detail in the following sections. The temperature was calculated using the dynamic method. In this and subsequent temperature calculations, the trap frequency was taken to be 3MHz unless otherwise stated. This is the value given in [Har13]. The highest fluorescence rates and lowest temperatures can be obtained using the lowest frequency atomic clock transitions. These two measures become less satisfactory as $|\mathbf{B}|$ increases; there is a corresponding increase in the detuning of some transitions from the laser frequencies, slowing down the internal dynamics of the ion. The values of the measures shown are presented as a rough guide to the values that can be expected; by using different beam arrangements or by adding side-

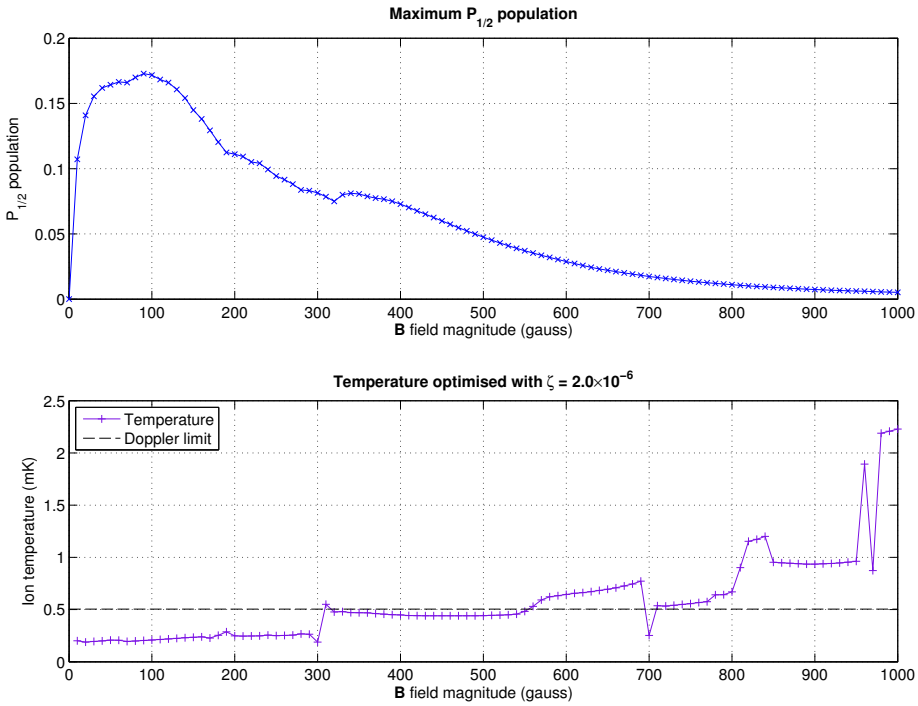

Figure 6.6: The maximum fluorescence and a characteristic value of the temperature that can be achieved in $^{43}Ca^{+}$ using a three-level cooling scheme as a function of $|\mathbf{B}|$. These results were obtained by optimising the intensity of the 397nm and 866nm beams and the detunings of the 397nm:A, 397nm:B and 866nm beams. The intensities of the two 397nm beams were kept equal. The polarisations of the beams were kept at $[f(\sigma^+), f(\sigma^-), f(\pi)] = [2/3, 0, 1/3]$ for the 397nm beams and $[f(\sigma^+), f(\sigma^-), f(\pi)] = [1/2, 1/2, 0]$ for the 866nm. The linewidths of the two beams were 500kHz for the 397nm beams and 100kHz for the 866nm beam. The temperature was calculated using a minimax process with $\zeta = 2.0 \times 10^{-6}$. ζ is defined in (4.115). The temperature varies in a discontinuous manner because the minimax process depends on both the population and temperature simultaneously. As $|\mathbf{B}|$ varies, the optimisation process may discretely jump between different optima that have almost identical values of the minimax parameter but different values of the temperature and fluorescence rate.

bands to more beams, it may be possible to achieve higher fluorescence rates or lower temperatures. However, for the beam arrangement considered here, it is possible to cool within a factor of two of the Doppler limit for all nine atomic clock transitions. For the five transitions with the lowest critical frequency, it is possible to cool the ion below the Doppler limit and observe a population in the $P_{1/2}$ of greater than 5%.

6.2.1 The 146.09 gauss qubit

We have seen that the fluorescence and temperature of atomic clock transitions become more favourable with lower critical field for the beam arrangement we have described. In addition, transitions with a lower critical field have a higher critical frequency, allowing higher possible readout fidelities. From a practical point of view, lower critical fields are easier to realise experimentally. The easiest and most practical qubits to make use of involve transitions A and B. The differences between these two qubits are slight.

The qubit chosen for experimental work was qubit A. The reason for this choice is based on the initialisation process. To initialise a qubit involving a transition between states in the centre of their respective manifolds, we use a two step process. The easiest states to initialise are those at the edges of the manifolds. These states can be initialised by applying a 397nm beam with a single sigma polarisation. This beam optically pumps population into the state with highest or lowest M depending on whether the polarisation is σ^+ or σ^- respectively. We then transfer population from this state to the qubit transition by using a sequence of microwave π pulses on the intervening hyperfine transitions. Transition A can be initialised from the state $|M = +4\rangle$ using three microwave pulses and transition B from the state $|M = -4\rangle$ using a similar three pulses. For more details on the initialisation process, see [All11, Har13]. Figure 6.7 shows the initialisation process for the chosen qubit, A.

Both qubits can be thus initialised with the same number of pulses from opposite ends of the $S_{1/2}$ manifolds. However, it is convenient to use the same σ polarisation on the 397nm beam as the dominant σ polarisation in the Doppler cooling scheme. As discussed in section 6.4.1, a Doppler cooling scheme using σ^+ light is more favourable than one using σ^- light. Qubit A was therefore chosen for experimental work.

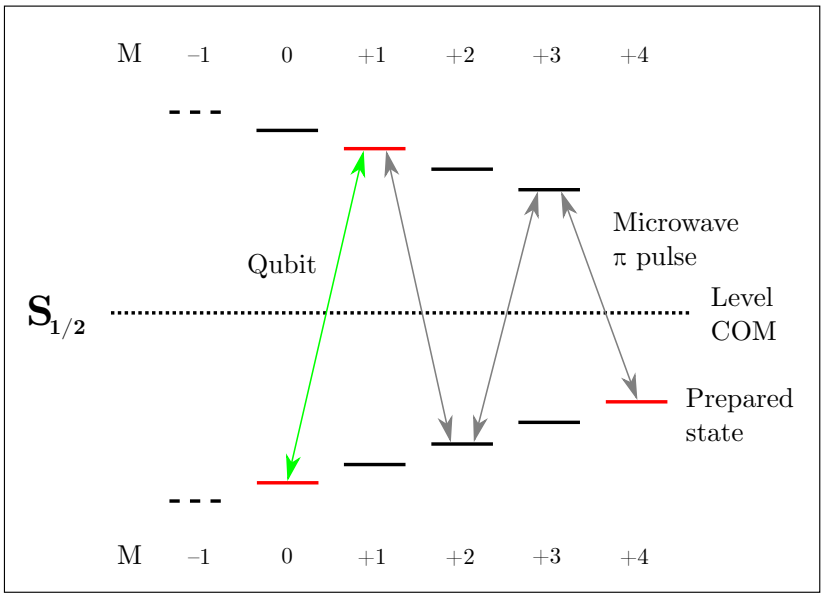

Figure 6.7: A demonstration of the initialisation process for qubit A. Population is first optically pumped to the 'prepared state' $|M = +4\rangle$ shown in red. Three microwave π pulses, shown in grey, are used to transfer population onto the qubit transition, shown in green.

6.3 Brute-force temperature optimisation

In section 4.8.3, we used a minimax optimisation process to examine the fluorescence rates and temperatures that can be achieved in $^{40}Ca^{+}$. It is instructive to run a similar optimisation on $^{43}Ca^{+}$ to compare how well the different isotopes can be cooled for a given fluorescence rate. In choosing an optimisation process, there are two features we need to consider.

Calculating temperatures by the dynamic method is slower for $^{43}Ca^{+}$ than for $^{40}Ca^{+}$ and is between two and three times slower than the steady-state approach. However, there are several reasons why it is necessary to use the slower method. It gives a better physical model of the behaviour of the ion. For $^{43}Ca^{+}$, the coldest temperatures predicted by the steady-state and dynamic methods are significantly different. This is demonstrated in figure 6.8, which shows the temperatures calculated using the two approaches over an 866nm scan. The minimum temperature predicted by the dynamic method is 2.3 times higher than that predicted by the steady-state method. We found that, in general, for tem-

peratures of the order the Doppler limit, the minimum temperature predicted by the dynamic method is of order twice that predicted by the simpler approach.

However, the dynamic method has a significant advantage. The first analysis of the temperature of $^{43}Ca^{+}$ was calculated using the steady-state method. As seen in figure 6.8, that method predicts low temperature spikes for a large number of the resonant effects in $^{43}Ca^{+}$, even if they are very weak. It can predict relatively strong cooling even on a tiny wiggle on the fluorescence curve, and so the temperature scan is a forest of a large number of spikes. Finding the lowest temperature in such a complicated function is a difficult task; it is easy for the optimisation process to stall on a local minimum, and there are a large number of these minima present in the function. The dynamic method, however, smooths over the vast majority of these weak resonant spikes. This method predicts only a handful of spikes, meaning that the optimal temperature point can be calculated much more easily. A global optimisation approach to calculate the ion's temperature using the dynamic method will run in less than an hour, whereas the same process using the steady-state method would calculate in of order a day.

One other consideration is the relative direction of propagation of the 397nm and 866nm beams. In the temperature models described in chapter 4, we allowed each beam to travel either parallel or antiparallel to a given propagation direction. The temperatures calculated in chapter 4 assumed that all beams were co-propagating. However, the temperature predictions for a system where the 397nm and 866nm beams are co-propagating and where they are counter-propagating will be different. Figure 6.9 shows the temperature predictions of the dynamic method for counter and co-propagating beams for a scan using the same parameters as the one shown in figure 6.8. The temperatures of the system with co-propagating beams are uniformly higher than that with counter-propagating beams; while the counter-propagating system cools to a tempera-

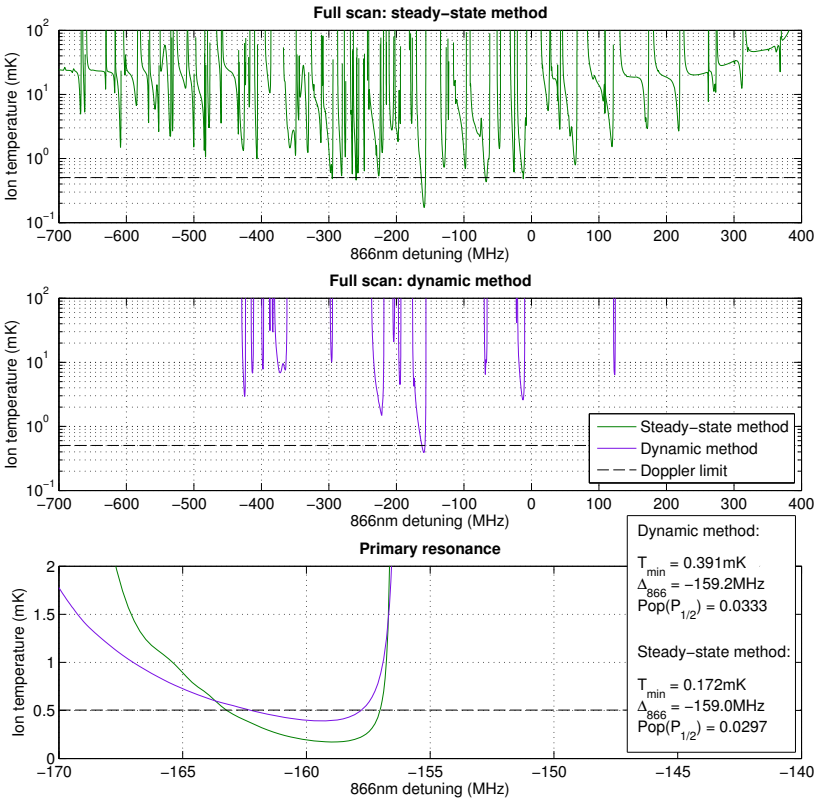

Figure 6.8: The temperature of an ion over an 866nm scan calculated using the steady-state and dynamic methods. The upper two plots show the temperature over the entire scan calculated using the steady-state and dynamic methods respectively. The lower plot shows the predictions of both methods over the resonance which most strongly cools the ion. The horizontal dashed line shows the Doppler limit. The steady-state method predicts low temperature spikes for every resonant feature seen in the scan. The output is a forest of many spikes, several of which extend to temperatures below the Doppler limit. By contrast, much of the behaviour is smoothed out in the predictions of the dynamic method. Only a few spikes occur over the scan for strong resonant features, and the resonance which most strongly cools the ion can easily be identified. However, the coldest temperature predicted by the dynamic method is approximately a factor of 2 higher than that predicted by the steady-state method.

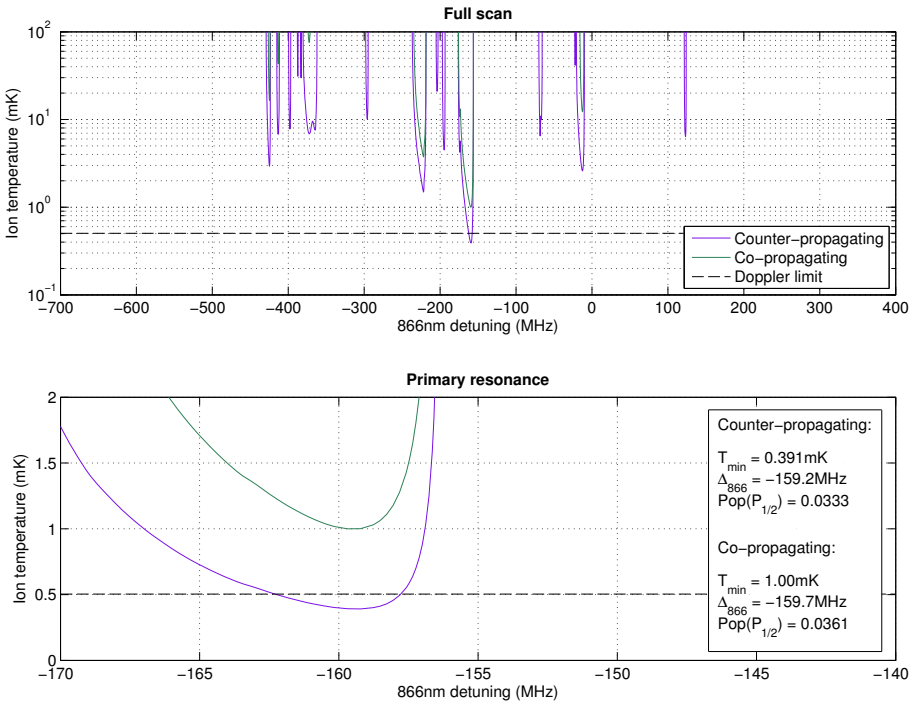

Figure 6.9: The temperature of an ion calculated using the dynamic method for a system with counter and co-propagating beams. The system shown uses the same parameters as figure 6.8. The structure of the spikes for both cases is very similar. However, the temperature predictions for the co-propagating case are uniformly hotter than for the counter-propagating case.

ture 30% lower than the Doppler limit, the co-propagating system only reaches a temperature of twice the Doppler limit.

We can understand this difference as follows. Consider an ion that sits on the optimum temperature point of a fluorescence curve. Cooling occurs because the laser frequencies in the frame of the ion vary with its velocity. Let us suppose that only one of the frequencies changes. Then the ion will climb a certain way up the wall of the dark resonance. We can take account of the frequency shift of the second beam by allowing the position of the dark resonance to shift as the ion moves up the wall of the resonance. If the beams are co-propagating, the centre of the dark resonance will move in the same sense as the ion. In this situation, the ion will move less far up the wall of the resonance in the same frequency interval than it would have done had the centre of the dark resonance remained stationary and the gradient the ion sees is reduced. By contrast, if the beams

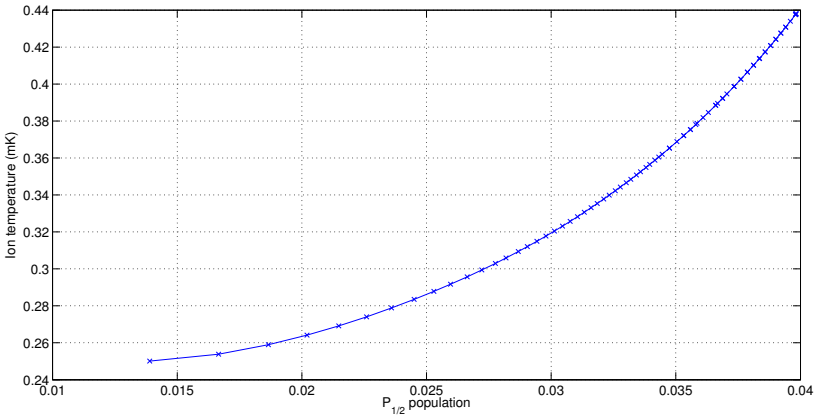

Figure 6.10: The optimal line produced by a minimax optimisation carried out for a cooling scheme using the three lowest energy levels of $^{43}Ca^{+}$. Two parameters were floated describing the intensities of the 397nm and 866nm beams. The two 397nm beams were fixed at equal intensity. Three parameters were floated describing the detuning of the 866nm beam and the two 397nm beams. Even for large populations, the ion can be cooled below the Doppler limit. Compare this figure to the similar scan taken for $^{40}Ca^{+}$ shown in figure 4.23.

are counter-propagating, the situation is reversed. The dark resonance moves in the opposite sense and the ion sees a higher gradient than it would have done otherwise. The temperature of the ion is inversely proportional to the gradient of the scattering rate and thus the fluorescence curve, so we could expect a system with counter-propagating beams to always produce colder temperatures than one with co-propagating beams, as observed. In all optimisation processes that follow, I make use of the dynamic method with counter-propagating beams.

Figure 6.10 shows the results of a minimax optimisation process carried out in $^{43}Ca^{+}$ at the critical field of qubit A. The process was chosen to be as similar to that for $^{40}Ca^{+}$ as possible. The temperature was optimised using the three lowest energy levels of the ion, and the same parameters were floated. Five parameters were floated describing the intensities and detunings of the three beams. The intensities of the two 397nm beams were kept equal. The beam polarisations were fixed at $[f(\sigma^{+}), f(\sigma^{-}), f(\pi)] = [2/3, 0, 1/3]$ for the 397nm beams and $[f(\sigma^{+}), f(\sigma^{-}), f(\pi)] = [1/2, 1/2, 0]$ for the 866nm beam. The linewidths of the 397nm and 866nm beams were set to $\Gamma_{L,397} = 0.5\mathrm{MHz}$ and $\Gamma_{L,866} = 0.1\mathrm{MHz}$ and

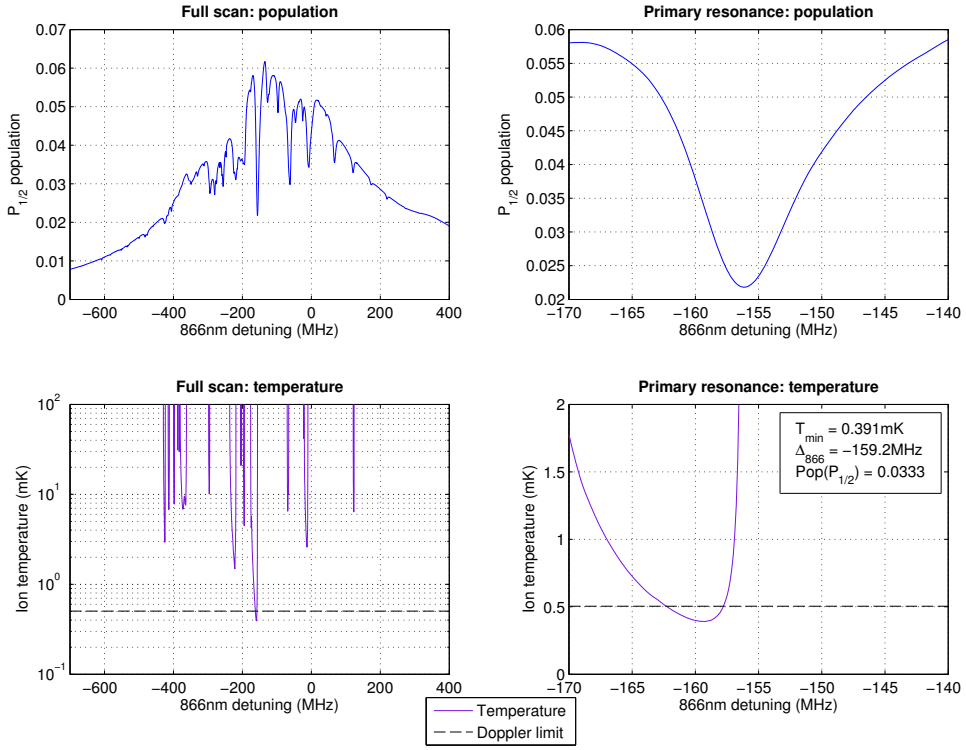

Figure 6.11: An 866nm scan showing the $P_{1/2}$ population and temperature for the minimax optimum with $\zeta = 2.46 \times 10^{-5}$ (see (4.115)). The lefthand plots show the population and temperature over the entire scan whilst the righthand plots focus on the behaviour on the resonance which cools the ion most strongly. The Doppler limit is shown on the temperature plots as a black horizontal dashed line. The minimum temperature of the ion is significantly below the Doppler limit. The $P_{1/2}$ population at this optimum point is 3.33%.

the magnitude of the B field was set to $|\mathbf{B}| = 146.09\text{G}$. The temperatures were calculated using the dynamic method. Temperatures lower than the Doppler limit can be achieved even for large populations.

Figure 6.11 shows the temperature and fluorescence for an 866nm scan taken around the minimax optimum. This optimum was taken from a point on the minimax line with large population. The value of the minimax parameter ζ, defined in (4.115), that produced this point is $\zeta = 2.46 \times 10^{-5}$. This value is used in all subsequent optimisations. The minimum temperature of the ion is significantly below the Doppler limit while still demonstrating a $P_{1/2}$ population on the side of the dark resonance of between 2% and 5%. Note that the complicated Zeeman structure means that there are many resonant effects spanning the en-

tire width of the scan. With all of these resonant features contributing to the ion's behaviour, it is perhaps surprising that temperatures that can be achieved in $^{43}Ca^{+}$ for this cooling scheme are only fractionally higher than for a similar system in $^{40}Ca^{+}$. In the following section, we will analyse the behaviour of $^{43}Ca^{+}$ for cold systems such as the one shown in this figure and determine the mechanism that allows this isotope to cool to such low temperatures.

6.4 Stretch state cooling

For a three-level cooling scheme, global temperature optimisation routines such as the one shown in the previous section strongly converge on two different resonances:

$$\alpha : |\mathrm{S}_{1/2}, M = +4\rangle \longleftrightarrow |\mathrm{P}_{1/2}, M = +4\rangle \longleftrightarrow |\mathrm{D}_{3/2}, M = +5\rangle \quad (6.1)$$

$$\beta : |\mathrm{S}_{1/2}, M = -4\rangle \longleftrightarrow |\mathrm{P}_{1/2}, M = -4\rangle \longleftrightarrow |\mathrm{D}_{3/2}, M = -5\rangle \quad (6.2)$$

The minimax optimum shown in figure 6.11 sits on resonance α. These resonances are unlike any others in the three-level system and cool the ion much more strongly than any others over a very large region of parameter space. Furthermore, optimisation routines that allow all the beam parameters to float will produce outputs that optically pump population onto these resonances. They have a special property which is demonstrated in figure 6.12 which shows the same system as shown in figure 6.11 but with the linewidths of the 397nm and 866nm beams set to zero. Note that when the system is exactly on resonance with α, the population in the $\mathrm{P}_{1/2}$ drops exactly to zero.

This behaviour is exactly what we would expect for a simple three-state system; see section E.1. For such a system, the population of the intermediate state of the resonance drops exactly to zero when the resonance condition is met when the two beams have zero linewidth. Thus, the two resonances α and β behave like perfect three-state resonances. In fact, when the resonance condition is met,

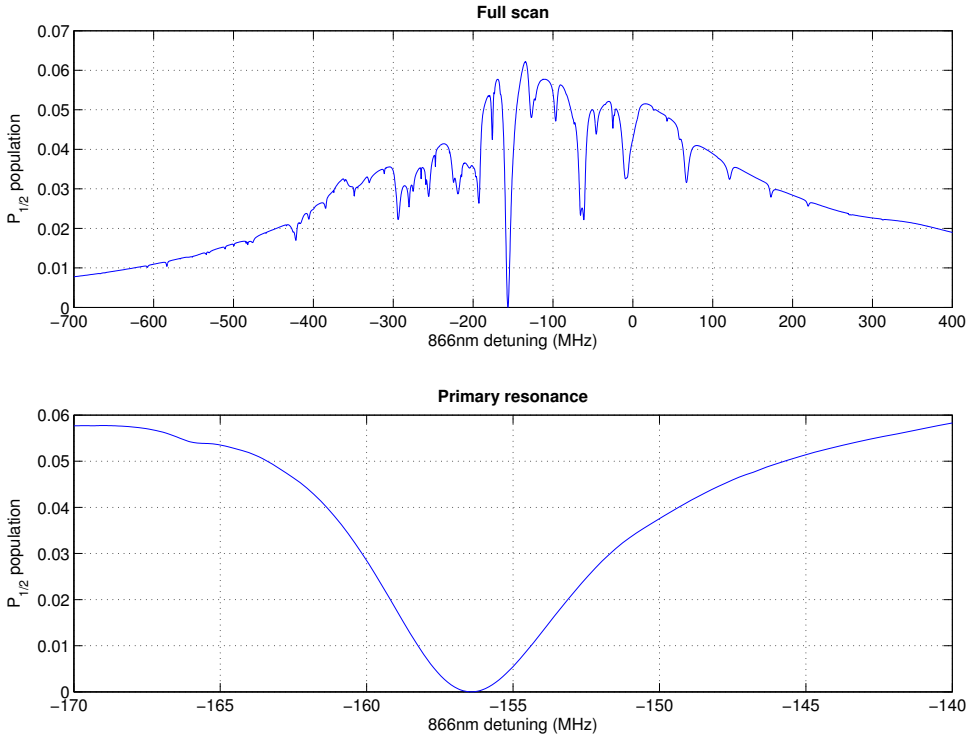

Figure 6.12: An 866nm scan taken using the same parameters as figure 6.11 but with the linewidths of both the 397nm and 866nm beams set to zero. The upper plot shows the $P_{1/2}$ population over the entire scan and the lower plot shows only the resonance that most strongly cools the ion, resonance α. The $P_{1/2}$ population drops exactly to zero when the resonance condition for α is met. This behaviour is only seen for resonance α; every other resonance seen on the scan, even those that cool the ion significantly, reduce the $P_{1/2}$ population to some non-zero value when the resonance condition is met.

population in not only the intermediate state but in *all other states not part of the resonance* drops exactly to zero; all the population is optically pumped onto the resonance.

To explain how this comes about for resonances involving states at the extremes of the level manifolds, we recall how the dark resonance occurs in a simple three-state system (see section E.2.1). With a certain distribution of population in the lower energy states $|a\rangle$ and $|c\rangle$, there is destructive interference preventing stimulated transitions from either to the common state $|b\rangle$. The population is thus trapped in $|a\rangle$ and $|c\rangle$.

Figure 6.13 shows schematically the behaviour of two sets of three states in a multi-state system. The lefthand plot shows a general resonance and the right-

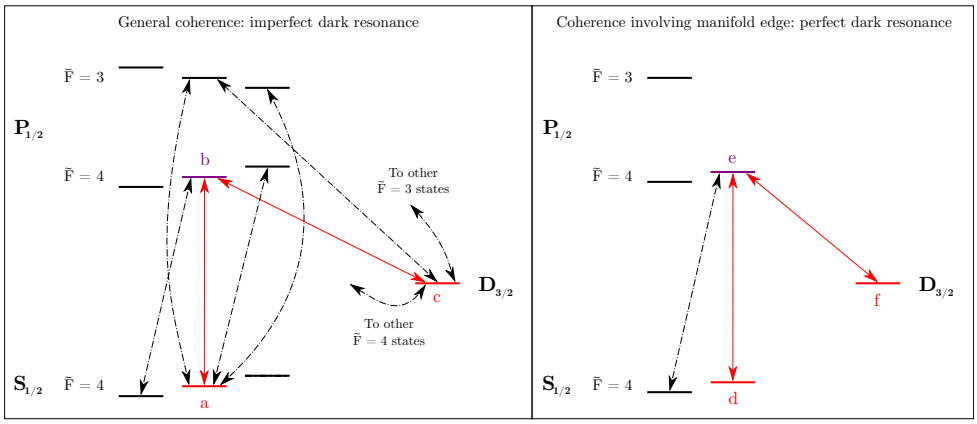

Figure 6.13: A schematic representation of 'imperfect' and 'perfect' dark resonances in three-level $^{43}Ca^{+}$. The energies of the states are not to scale. The lefthand plot shows a resonance involving states in the three level that do not have maximal magnitudes of M. The resonance is shown by the red arrows. The two states connected by the resonant effect are shown in red and the common state is shown in purple. Transitions from these three states to other states are shown by black dashed lines. The righthand plot shows one of the two resonances that makes use of states in the three levels with maximal magnitude of M. This resonance involves the states at the edge of the $S_{1/2}$ and $P_{1/2}$ manifold. Note that there are no additional transitions involving the red states d and f for the perfect resonance, but there additional transitions involving the states a and c for the imperfect resonance. For a discussion, see text.

hand plot resonance α involving the 'stretched states' at very edges of the level manifolds. For a general resonance it is usually possible for excitation to occur from either one or both of $|a\rangle$ and $|c\rangle$ to other states not involved in the resonance. Population can therefore generally reach $|b\rangle$ by an indirect route; from $|b\rangle$ it can decay to $|a\rangle$ and $|c\rangle$. In the steady state, therefore, there is population in $|b\rangle$.

In the case of α and β, however, the only stimulated transitions which can occur from the lower states (given our choice of laser polaristations) are to the common state. In steady state, therefore, when the populations of the lower states have values which close off this route by destructive interference, the populations of all other states must vanish. Spontaneous decay from the common state is the 'black hole' into which all population must eventually fall.

Figure 6.14 explicitly shows two properties of a perfect dark resonance. A dark resonance of this type relies on the 397nm beam having only one non-zero σ component. As soon as both σ components are present, a route out of the

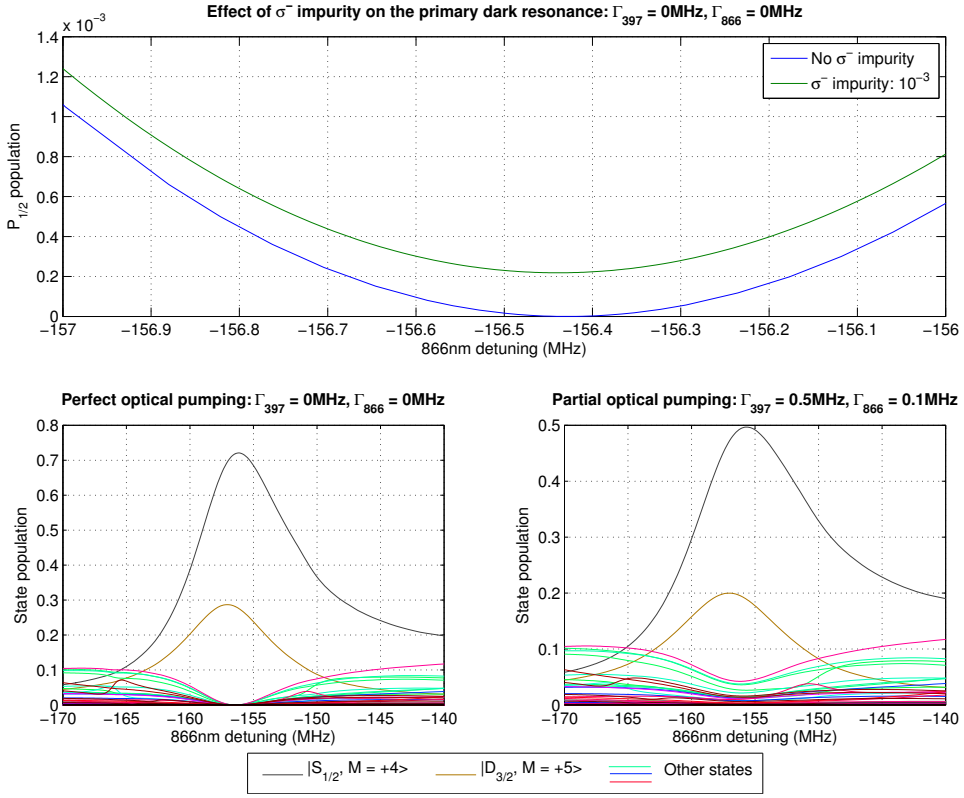

Figure 6.14: A demonstration of two properties of perfect dark resonance α. The upper plot shows the $P_{1/2}$ population as a function of 866nm detuning close to resonance. The linewidths of both beams are set to zero. The blue curve describes a system where the polarisation of the 397nm beam contains no σ^- fraction and the green curve a system with a small σ^- fraction. Only the blue curve drops exactly to zero on resonance. The σ^- polarisation disrupts the perfect dark resonance by adding an additional transition out of one of the coherence's connected states. The two lower plots show the population in every state in the system as a function of 866nm detuning close to resonance when there is no polarisation impurity. The connected state in the $S_{1/2}$ is coloured black and that in the $D_{3/2}$ coloured brown. Other states in the $S_{1/2}$ are coloured green, those in the $P_{1/2}$ are coloured blue and other states in the $D_{3/2}$ are coloured red. The lefthand plot shows the behaviour for a system where both beams have zero linewidth. The population of all states except for the connected states drop exactly to zero on resonance. The righthand plot shows the behaviour for a system where the linewidths of the beams are non-zero. The populations of the other states are still reduced on resonance, but they do not drop to zero.

connected states is provided and the dark resonance destabilises. The effect of adding a very small σ^- component to the 397nm beam for a system driving resonance α is shown in the upper plot of this figure. Even a very small component of the opposite σ polarisation destabilises the resonance and the population of the $P_{1/2}$ does not fall to zero. The two lower plots show the populations of all the states in the system as a function of 866nm detuning. The lefthand plot shows the results for a system where the linewidths of both beams are zero whilst in the righthand plot they are are nonzero. The populations of all states other than the connected states drops close to resonance. For beams with zero linewidths, the populations drop exactly to zero and population is shared between the two connected states. For beams with nonzero linewidths, the effect is less pronounced. Population is transferred into the connected states on resonance, but the populations of the other states no longer drop to zero.

6.4.1 Choosing positive or negative M

By focusing on one particular resonance, we can significantly decrease the amount of time a minimax routine takes to compute as we can ensure that every point is in the vicinity of a cooling optimum. We can use either resonance α or resonance β to optimally cool an ion using a three-state cooling scheme. The connected states in the $S_{1/2}$ for these two resonances lie at either end of the level's manifolds. Knowing this, we can optimise the ion's temperature without the need to optimise any polarisations. We know that population must be driven to the edges of the 397nm manifolds, so we simply set the polarisation of the 397nm beam to have as high a fraction of one σ polarisation and as small a fraction of the other σ polarisation as possible. The system is largely insensitive to the 866nm polarisation as this beam only repumps population onto the cooling transition.

The beam arrangement for the trap in which the experiments described later

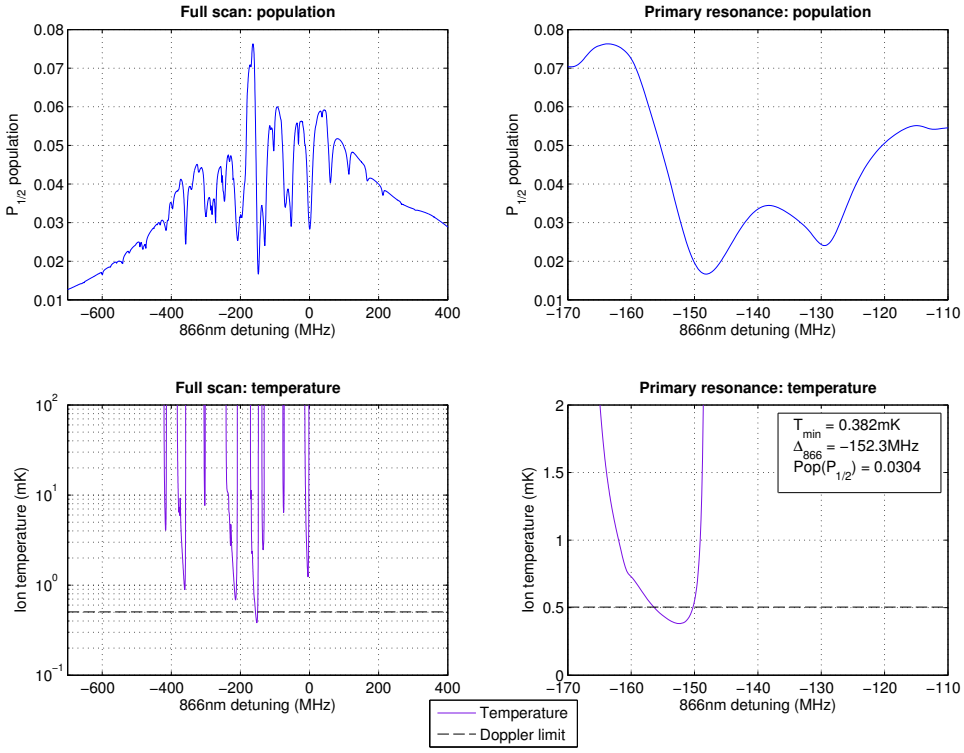

Figure 6.15: An 866nm scan around the minimax optimum calculated for polarisation α using $\zeta = 2.46 \times 10^{-5}$. The lefthand plots show the $P_{1/2}$ population and temperature over the entire scan. Resonance α can be clearly seen on the population curve as the tall slope with high negative gradient on the righhand edge of the high population spike. The righthand plots show the behaviour over the dark resonance α. On the temperature plots, the Doppler limit is shown as a horizontal dashed line. The optimum temperature is 0.382mK, well below the Doppler limit, with an associated $P_{1/2}$ population of 3.04%.

in this chapter were carried out has Doppler cooling beams at 45° to the **B** field. This limits the polarisation states we can use. The region in polarisation space that can be accessed traces out an ellipse, and this ellipse can be passed as a limit to an optimisation process. As can be understood from the foregoing discussion, all optimisations carried out using this ellipse converge to points where one σ polarisation is maximised at the expense of the other. The highest fraction of one σ polarisation that can be achieved is 2/3, with the remaining 1/3 comprising π light. See appendix B for a full derivation. We therefore have two choices of

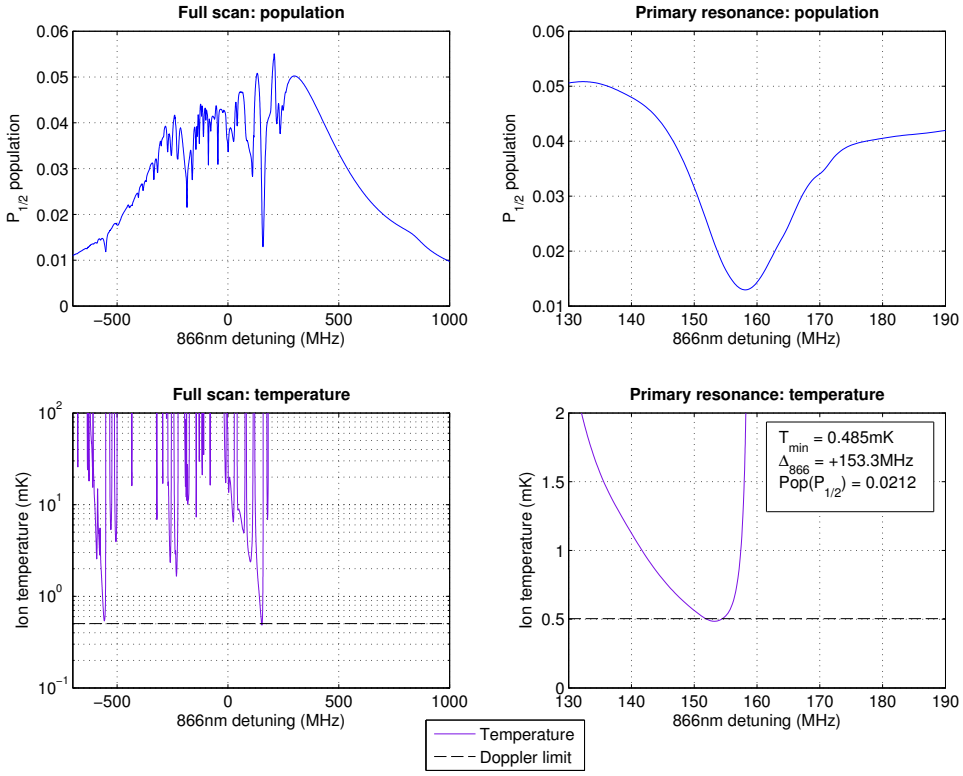

Figure 6.16: An 866nm scan around the minimax optimum for polarisation β using $\zeta = 2.46 \times 10^{-5}$. The lefthand plots show the $P_{1/2}$ population and temperature over the entire scan, while the righthand plots show the region around resonance β. The Doppler limit is shown on the temperature plots as a dashed horizontal line. The optimal temperature is 0.485mK, slightly below the Doppler limit, with an associated $P_{1/2}$ population of 2.12%.

397nm polarisation that will lead to low temperatures. These are:

$$\alpha : [f(\sigma^+), f(\sigma^-), f(\pi)] = [2/3, 0, 1/3] \tag{6.3}$$

$$\beta : [f(\sigma^+), f(\sigma^-), f(\pi)] = [0, 2/3, 1/3] \tag{6.4}$$

To determine which should be used, a minimax temperature optimisation and a fluorescence maximisation process were carried out for each of these polarisations.

Figures 6.15 and 6.16 respectively show the result of the minimax optimisation for polarisations α and β in the form of two 866nm scans. The value of the minimax parameter used in both cases was $\zeta = 2.46 \times 10^{-5}$. The minimal temperature for the scan around resonance α is well below the Doppler limit

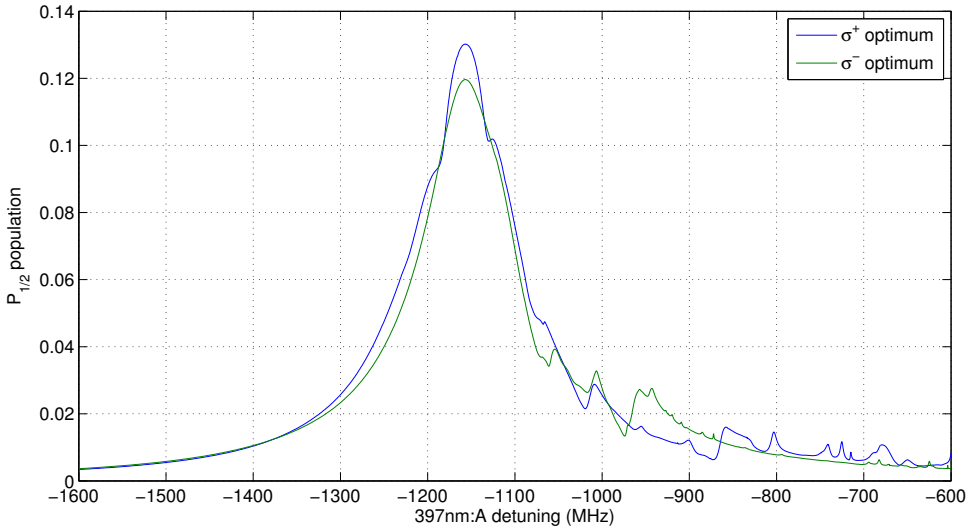

Figure 6.17: Two 397nm scans showing the results of a maximisation process on the $P_{1/2}$ population. The blue scan shows the optimum for resonance α, while the green scan shows that for β. The maximum populations respectively are 13.02% and 12.0%. It is possible to produce a higher fluorescence rate using the α optimum compared to the β optimum. The structure to the high frequency side of the maximum is due to other resonant effects in the system. The ion will in general be very hot in this region so these resonant effects will not be observed. Compare this plot to the similar 397nm scan in figure 6.9 of [All11]. The referenced plot shows a higher maximum fluorescence than this figure as the upper limit on the 397nm intensity was at a higher value.

for a $P_{1/2}$ population of several percent. By contrast, the minimal temperature for the scan around β is hotter than this by 0.1mK and only slightly below the Doppler limit. In addition, the associated $P_{1/2}$ population is lower than for the previous case. The minimax optimisation process using polarisation α found an optimum with a colder temperature and higher fluorescence than the optimisation process using polarisation β. This implies using resonance α to cool the ion is more favourable than using resonance β.

Figure 6.17 shows the maximum $P_{1/2}$ population that can be achieved using polarisation α and polarisation β. The maximum values of the population are 13.02% and 12.0% using α and β respectively. This also implies that using resonance α to cool is more favourable than using resonance β. In our experimental work, we use resonance α and σ^{+} light to cool the ion using a three-state cooling scheme.

6.4.2 Variation with beam parameters

We now examine how the temperature and fluorescence of the ion depend of the beam parameters. We use the set of parameters for the system shown in figure 6.15 and examine the effects of varying some of them.

Figure 6.18 shows the temperature and fluorescence as a function of the intensities of the 397nm and 866nm beams. The upper plot shows the variation of the two measures as function of the 397nm intensity. This plot was produced by stepping the 397nm intensity from the parameters describing the minimax optimum shown in figure 6.15 and subsequently optimising the 866nm detuning only. This second step ensures that the optimum refers to the lowest temperature on an 866nm scan, allowing a comparison with other 866nm scans shown in this chapter. We see similar behaviour to that mentioned in section 4.8.1, and this is one reason why we optimise temperature using a minimax process. The temperature of the ion will always become lower (and the cooling process slower) when the intensity of the 397nm beam is reduced until the point at which the heating rate of the trap prevents the ion from cooling any further. That scale is too small to be seen on this plot. As the temperature of the ion is reduced, the ion's fluorescence is also reduced. This is undesirable as the fidelity of readout drops with decreasing fluorescence [AHS+16]. In experiment, it is always possible to cool the ion further by reducing the 397nm intensity, but the drawback is a corresponding decrease in the fidelity of the diagnostics.

The variation with 866nm intensity is more complicated. As before, this plot was produced by stepping the 866nm intensity and then optimising the 866nm detuning only. The temperature has a minimum and the population a maximum at different values of the detuning. The position of the minimax optimum depends on the relative importance of the fluorescence and temperature set by the minimax parameter ζ. We see that the temperature of the ion is lower than the Doppler limit for a large range of 866nm intensities; the 866nm intensity

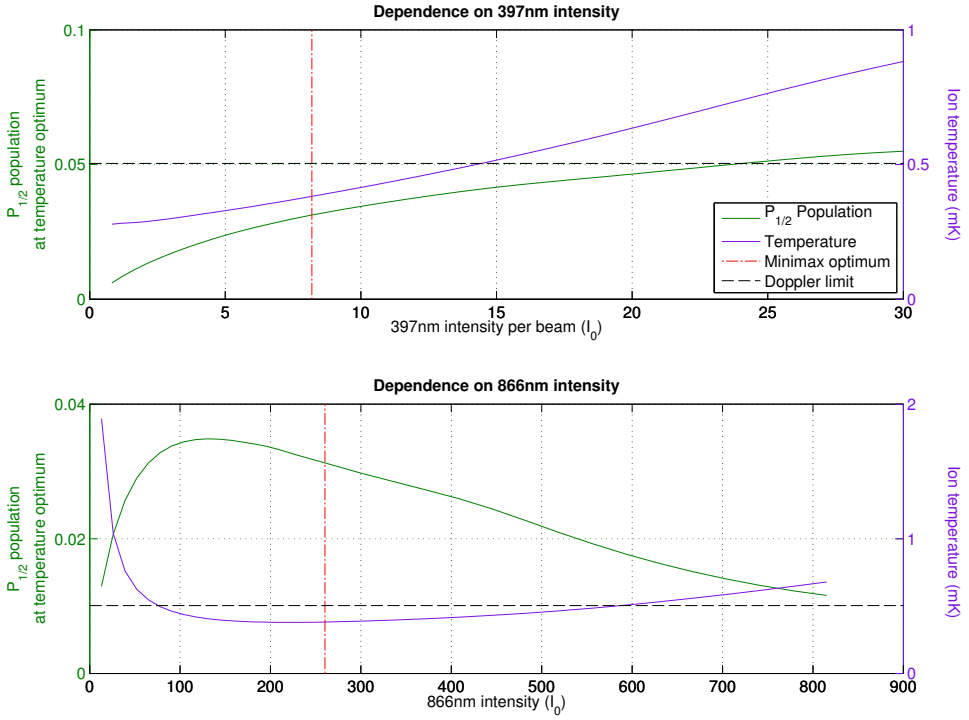

Figure 6.18: A demonstration of the dependence of the $P_{1/2}$ population and the temperature of the ion on the intensities of the 397nm and 866nm beams. The intensities of the two 397nm beams are fixed to be equal for these curves. The upper plot shows the $P_{1/2}$ population, in green, and the ion temperature, in purple, as a function of the 397nm intensity in each beam. The two curves use different vertical axes as shown on the left and righthand side of the plot. The lower plot shows the variation of the $P_{1/2}$ population and temperature as a function of the 866nm intensity. On both plots, the minimax optimum around which this scan was taken is shown as a vertical, red dot-dashed line. The optimum is at the same set of parameters as the scan shown in figure 6.15. The Doppler limit is shown as a horizontal dashed line. These plots were calculated starting at the minimax optimum and stepping one of the two intensities. The 866nm detuning was then optimised to ensure that the lowest temperature recorded refers to the lowest temperature on an 866nm scan. This allows an immediate comparison with other 866nm scans presented in this chapter.

can change by a factor two without significantly altering the ion's temperature. However, the ion's fluorescence will change by a larger amount over this range. The ion is less sensitive to changes in the 866nm intensity compared to changes in the 397nm intensity.

Figure 6.19 shows the variation of the fluorescence and temperature of the ion as a function of the 397nm detuning. Unlike the previous two plots, this plot was generated by stepping the 397nm detuning and subsequently minimax optimising all other parameters. The dependence of the two measures on this

variable is especially important from an experimental standpoint. The upper plot shows the variation of the two measures as a function of the 397nm detuning. The minimum of the ion's temperature coincides with the position of the minimax optimum. The ion's temperature and fluorescence change by a relatively small amount for a small region around this minimum. The temperature curve increases smoothly on either side of the minimum. The population curve, on the other hand, demonstrates several features (although the small features on the low-frequency side of the scan are due to computational noise). These features arise because the actual population at which the minimum occurs depends on several factors such as the relative height of the neighbouring resonances and the height of the Lorentzian envelope. We see, however, that we can alter the 397nm detuning by approximately ± 15MHz without significantly changing the ion's temperature or fluorescence.

Altering the 397nm detuning from the optimum position to a more negative value is often necessary in experiment. The reason for this is demonstrated in the two lower plots of this figure. These plots show the $P_{1/2}$ population over a 397nm scan for two sets of parameters. The scan on the left is taken using the parameters for the minimax optimum. The scan on the right is taken using the newly optimised parameters after the 397nm detuning has been altered by -19.6MHz. The centre of the Lorenztian envelope was calculated by fitting the wings of the profile and is shown as a pink dashed vertical line on both plots. The point at which the temperature is minimal is shown as a grey line. On the lefthand plot, the two vertical lines are coincident; the minimax optimum lies exactly on the centre of the Lorentzian profile. By contrast, the grey line lies -19.6MHz to the low frequency side of the pink line for the righthand plot. By changing the 397nm detuning, we have shifted the position of the low temperature point from the centre of the profile.

In section 4.1.1, we described how 397nm scans in $^{40}Ca^{+}$ show significant

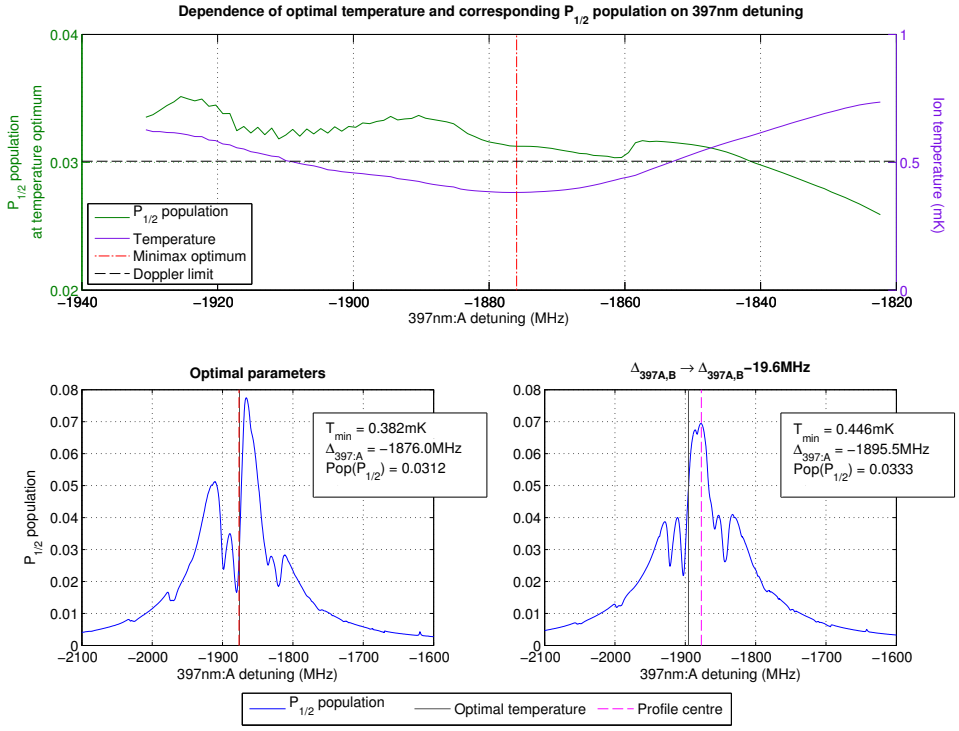

Figure 6.19: The dependence of the $P_{1/2}$ population and the ion's temperature on the 397nm detuning. The upper plot shows the variation of the two measures as a function of the 397nm detuning. The ion temperature is minimal at the same detuning as that of the minimax optimum. The $P_{1/2}$ population changes by a relatively small amount as the detuning varies. The population will change depending on the relative height of the principal resonant feature compared to neighbouring peaks along with the envelope of the 397nm profile. Because of this, the population traces out a curve with several features. The lower two plots show the population over a 397nm scan for two different sets of parameters. The lefthand plot shows the scan taken at the minimax optimum. The righthand plot shows a scan taken around the point on the upper plot with the 397nm detuning changed by -19.6MHz. The other parameters in the lefthand and righthand plots differ slightly due to the minimax optimisation process used to produce the upper plot. The position of the point with minimal temperature is shown by the grey vertical line. The approximate position of the profile centre is determined by a fit to the wings of the profile and is shown by the pink dashed vertical line. For a discussion of the importance of these two plots, see text.

heating to the high-frequency side of the centre of the profile. The same type of behaviour occurs in $^{43}Ca^{+}$, although the effect is less clearly defined. In general, there will be fewer heating regions on the low-frequency side of the profile and more on the high-frequency side. The minimax optimum lies at the centre of the envelope. It sits close to the high density of heating regions at higher frequency. If the 397nm detuning is increased by even a very small amount, the ion will become very hot. This has two consequences.

In order to cool a trapped ion, the cooling scheme must be able to deal with an ion that is at first very hot. A hot ion experiences very large Doppler shifts of both the 397nm and 866nm beams as it oscillates. The ion's final temperature will depend on the equilibrium between the heating and cooling effects seen over the extent of its external oscillations. We cool a hot ion by placing it in the vicinity of the strongest cooling resonance. However, for the system we are examining here, the heating region to the high-frequency side of the 397nm scan lies only a few MHz away from the coldest point on the scan. The cooling power of the resonance may not be able to compete with a large neighbouring region of the scan which strongly heats the ion. The ion will cool to some intermediate solution described by the excursion effect and not to the low temperatures we would predict using the dynamic method.

Secondly, diagnostics become more difficult. Consider an 866nm scan taken at parameters around the minimax optimum. At the point of optimum temperature, the ion experiences a small magnitude of its external oscillations and consequently small Doppler shifts of the 397nm and 866nm beams. We now move the ion away from the optimum up the wall of the dark resonance. The ion's temperature rises and the Doppler shifts the ion sees increases. The Doppler shift of the 397nm beam will eventually rise to of order a few MHz and the ion will see the effects of the very large heating region. The equilibrium temperature of the ion will therefore rise significantly. At all points of the scan at which the excur-

sion effect is important, the ion's temperature will be very large. An 866nm scan taken around the minimax optimum was shown in figure 6.15. The excursion effect is important at all points where the ion's temperature calculated using the dynamic effect is hotter than of order the Doppler limit. This is the case for the vast majority of the scan! If we were to take data for an ion under these conditions, those regions of the scan described by the excursion effect would show significant departures from the cold ion fluorescence model. In other words, most of the scan would consist of heating regions. This scan would therefore be very difficult to fit and examine using the fast cold ion model[1].

We can alleviate both of these problems by moving the optimal cooling point away from the large heating region. We do this by moving the 397nm frequency to the red of the minimax optimum as shown in the lower-right plot of figure 6.19. This ensures that larger Doppler shifts are required before the large heating region significantly increases the ion's equilibrium temperature. This has the effect of making it easier for a hot ion to cool to the optimal temperature predicted by the cooling scheme while also reducing the size of regions of fluorescence scans that cannot be fitted by the cold ion model, making them easier to fit. Of course, there is a price: the further the optimum is moved from the profile centre, the higher the temperature optimum becomes.

In practice, a compromise is best found experimentally. The intensities of the beams and the frequency splitting between the 397nm beams are set to the values of the chosen minimax optimum. This frequency splitting we use for the α optimum is 2.930GHz. This value was taken from a minimax optimisation process that produced sub-Doppler temperatures while still allowing populations of around 4%. There is some flexibility in the 397nm intensity depending on the level of cooling required. The 397nm and 866nm detunings are set to lower frequencies than optimal. These frequencies are then both slowly tuned as close

[1] By reducing the 397nm detuning, both α and β can be used to cool very hot ions

to the optimal values as possible before clear evidence of heating is observed such as, for example, sharp changes in the ion's fluorescence. This ensures that the ion's temperature is as low as can be achieved before the effects of the large heating region become apparent.

6.4.3 The excursion effect in $^{43}Ca^{+}$

Every resonant feature in a fluorescence scan will contribute to heating the ion on one side of the resonance. The number of resonance effects in $^{43}Ca^{+}$ means that there are many places over a fluorescence scan in which the ion can heat. If the parameters describing the scan are not chosen carefully, the heating effects of the resonances may be high enough that the excursion effect is important for large regions of the scan.

We can model this behaviour using the excursion and motional methods. However, the increased complexity of the structure of $^{43}Ca^{+}$ compared to the simpler isotope means that these approaches become unwieldy and time-consuming. Interestingly, the relative calculation times of the excursion and motional methods are reversed for $^{43}Ca^{+}$. The excursion method relies on repeated inversion of a matrix, while the motional method relies only on matrix multiplication. Matrix inversion scales strongly with matrix size; the number of operations[2] required to invert an $N \times N$ matrix scales as of order $N^{2.8}$. Even though the excursion method requires far fewer overall operations than the motional method, the increase in calculation time due to the matrix inversions means that the excursion method computes much more slowly than the motional method for $^{43}Ca^{+}$. As the motional method is more advanced than the excursion method, we only make use of the motional method for describing excursion effects for this isotope.

[2]One example is the Strassen algorithm which scales as $N^{2.807}$ [Str69]. Other algorithms have faster scalings, but these often have large constant overheads that make them slow to use for inverting most matrices in practice.

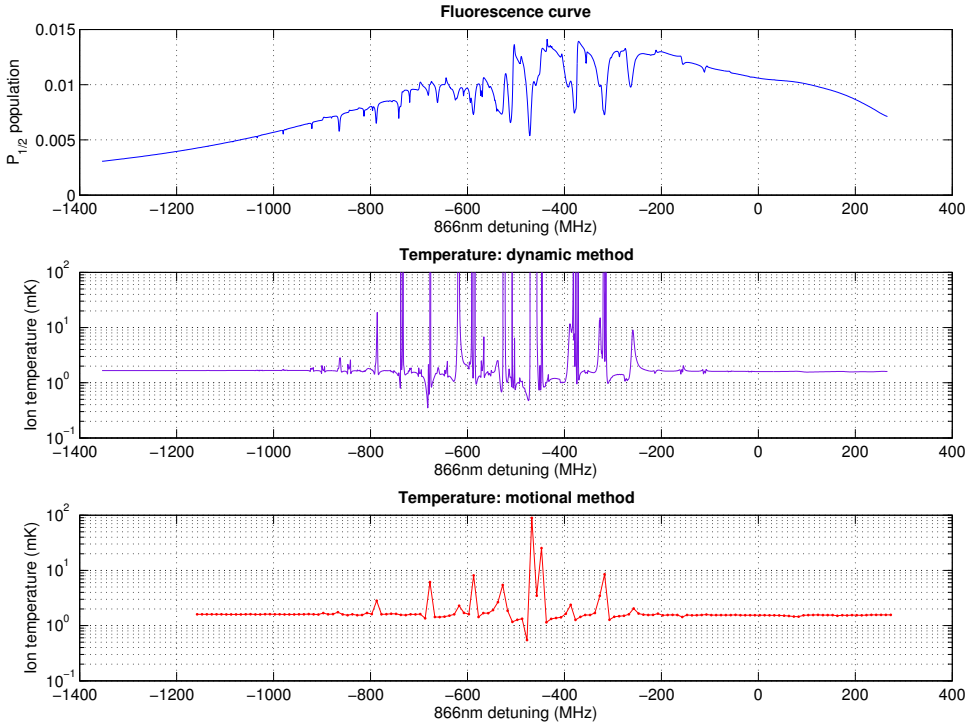

Figure 6.20: A demonstration of the temperature predictions of the dynamic method and the motional method taken over an 866nm scan. The parameters describing this scan were taken from the fit to an 866nm scan taken with a low 397nm intensity, resulting in low $P_{1/2}$ populations and correspondingly low temperatures. The upper plot shows the $P_{1/2}$ population. The middle plot shows the temperature calculated using the dynamic method and the lower plot shows the temperature calculated using the motional method. For more details, see text.

Unfortunately, this is still a slow process. By way of comparison, calculating the temperatures for a high resolution fluorescence scan in $^{40}Ca^{+}$ takes of order several hours. Calculating even a low resolution scan in $^{43}Ca^{+}$ takes a little over a week. However, we will rarely need to make use of the motional method for this isotope. We are interested primarily in those regions of parameter space where the ion is cold. When fitting a fluorescence scan, we can avoid dealing with the fluorescence modifications caused by excursion by not fitting those regions in which the excursion effect will be prominent.

Figure 6.20 shows a comparison of the temperature predictions of the dynamic and motional methods over an 866nm scan. This scan is a fit to set of data taken with very low 397nm intensity. As a result, both the $P_{1/2}$ population and

the ion's temperature are low over a large portion of the scan. However, the ion shows significant heating over several regions of the scan. As we would expect from our analysis in $^{40}Ca^{+}$, the motional effect smooths the behaviour of some of the smaller resonant effects in the scan and predicts large but finite temperatures in those areas of the scan where the dynamic method predicts the ion will heat indefinitely.

6.5 $\widetilde{\mathbf{F}}$ = 3 cooling

In our experiments, we make use of stretch state cooling. However, there is a more efficient beam arrangement that allows for higher fluorescence to be obtained for the same temperature. Unfortunately, experiments testing this arrangement were carried out before the development of the dynamic method and no improvement was observed. Whilst higher fluorescence is desirable for improved diagnostics, and this beam arrangement is recommended for future experiments, we have managed to Doppler cool $^{43}Ca^{+}$ without making use of this scheme.

Stretch state cooling makes use of a perfect dark resonance in the lowest three levels of $^{43}Ca^{+}$. The optimal temperatures were calculated using a global minimax process that sampled in a region of parameter space strongly centred around this resonance. In particular, the splitting between the 397nm beams was kept close to a value of 2.9GHz. The reason for this is to allow one of the two 397nm beams to pump population out of the $\widetilde{F} = 3$ manifold back onto the dark resonance. The large σ^{+} polarisation of the 397nm beams will push population into the high-M ends of both manifolds. In order to pump population efficiently out of the $\widetilde{F} = 3$ manifold, we set this beam to pump population mainly out of the $|\widetilde{F} = 3, M = +3\rangle$ state. The frequency difference between the two beams must therefore be approximately the frequency difference between the $|\widetilde{F} = 3, M = +3\rangle$ state and the $|M = +4\rangle$ state. This frequency difference is

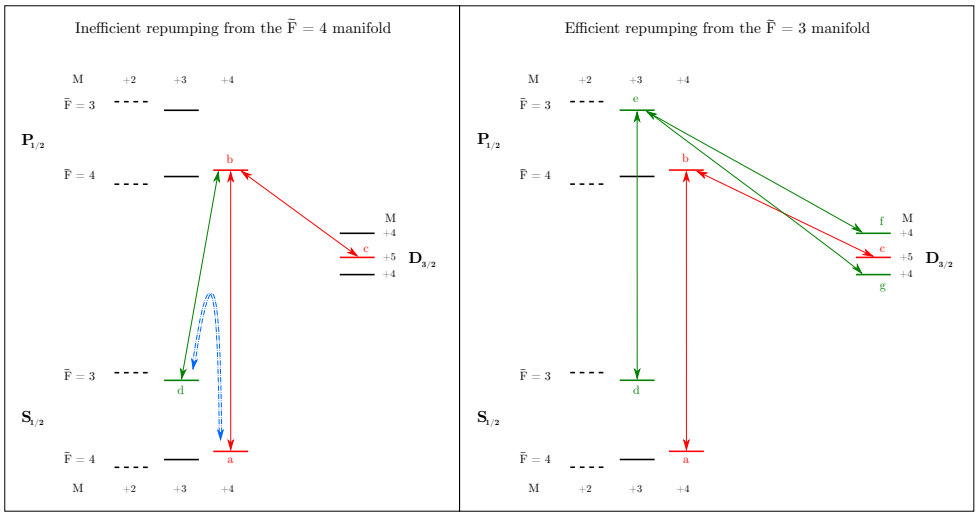

Figure 6.21: A demonstration of two different repumping methods for a cooling scheme making use of resonance α. The lefthand plot shows the repumping scheme used in section 6.4. Population is pumped out of the upper manifold of the $S_{1/2}$ using a beam that resonantly drives a transition from what would be expected to be the most strongly populated state in this manifold to the common state of the cooling coherence. The transitions and states involved with the coherence are shown in red. The $\widetilde{F} = 3$ state and transition involved in the repumping are shown in green. A strong resonant feature that drives population between the manifolds of the $S_{1/2}$ is shown in blue. The righthand plot shows the $\widetilde{F} = 3$ cooling scheme. Population is pumped out of the $\widetilde{F} = 3$ manifold of the $S_{1/2}$ into the $\widetilde{F} = 3$ manifold of the $P_{1/2}$. This removes the problem of the blue resonance feature. Population then decays back onto the cooling transition. For more information, see text.

2.9GHz.

However, this is not in fact the most efficient way to pump population out of the upper manifold. The problem with this approach is shown in figure 6.21. The lefthand plot shows the repumping scheme that has just been described and was used in the minimax optimisation processes discussed in section 6.4. The states and transitions involved in the cooling coherence are shown in red, and the $\widetilde{F} = 3$ state and transition involved in the repumping are shown in green. The problem that arises is that this cooling scheme excites another resonance effect. This resonance is shown in blue and connects the $S_{1/2}$ states $|\widetilde{F} = 3, M = +3\rangle$ and $|M = +4\rangle$. The two 397nm beams resonantly drive population between the two $S_{1/2}$ manifolds, actively counteracting the repumping process. This effect can be quite pronounced. Individual states in the $\widetilde{F} = 3$

manifold with $M < +3$ can contain over 10% of the ion's population, even for the minimax optima discussed in the previous section! This resonance effect cannot be avoided by detuning the repumping beam by a small amount as this will correspondingly reduce the efficiency of the repumping.

The righthand plot shows an alternative repumping scheme that avoids this resonance effect. I will refer to this scheme as the $\widetilde{F} = 3$ repumping scheme and the previous scheme as the $\widetilde{F} = 4$ repumping scheme. Population is instead repumped into the $\widetilde{F} = 3$ manifold of the $P_{1/2}$. It then decays back onto the cooling resonance. The two 397nm beams drive transitions between entirely different states and no longer resonantly drive population between the two $S_{1/2}$ manifolds. The effect is that population is driven out of the $\widetilde{F} = 3$ manifold far more efficiently using this scheme. For the same temperature, the $\widetilde{F} = 3$ scheme is associated with significantly higher fluorescence rates than the simpler scheme.

Figure 6.22 shows the results of a global minimax optimisation scheme calculated using the $\widetilde{F} = 3$ repumping scheme with minimax parameter $\zeta = 2.46 \times 10^{-5}$, the same as that used before (see figure 6.15). The value of the temperature in both cases is comparable, but the population at the coldest point is over a factor of two higher using the $\widetilde{F} = 3$ repumping scheme. This difference is summarised in table 6.3. Note also that the dynamic method predicts that the ion's temperature will be low for a large portion of the scan; the minimax optimum in this case does not lie on the centre of the 397nm Lorentzian envelope. As a result, this optimum will be able to cool even a very hot ion.

An experiment was carried out to try and cool using this repumping scheme. However, this repumping scheme was developed before the dynamic method. The original optimum produced using this scheme were calculated using the steady-state method. Unfortunately, this steady-state optimum was situated on the side of a thin resonance feature that is averaged out by the dynamic effect.

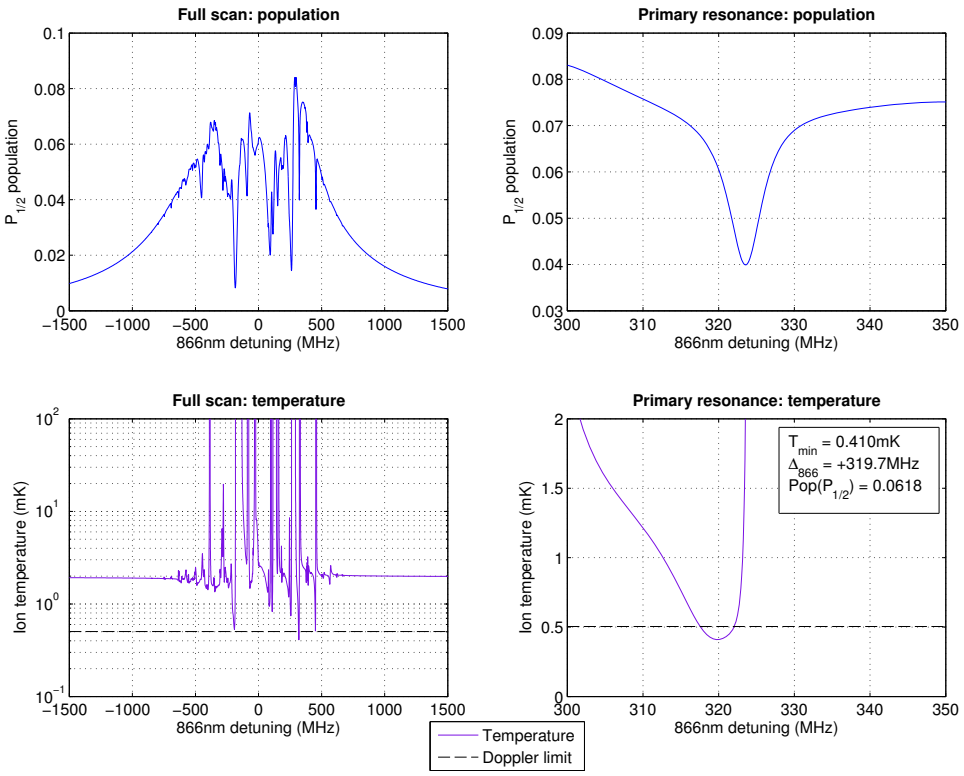

Figure 6.22: The $P_{1/2}$ population and temperature over an 866nm taken around the global minimax optimum parameters calculated using the $\widetilde{F} = 3$ repumping scheme with minimax parameter $\zeta = 2.46 \times 10^{-5}$. The lefthand plot shows the population and temperature taken over the entire scan. The righthand plot shows the region around the resonance that most strongly cools the ion. The Doppler limit is shown on the temperature plots as a horizontal dashed line. The value of the minimax parameter is the same as that used to produce the optimum using the simpler repumping scheme shown in figure 6.15. The temperature of the ion in both cases is comparable, but the population at the coldest point is over a factor of 2 higher for the optimum shown here. This population is 6.18%.

Scheme	$\widetilde{F} = 3$	$\widetilde{F} = 4$
Temperature (mK)	0.410	0.382
$P_{1/2}$ population	6.18%	3.04%

Table 6.3: A comparison of the temperature and population at the optimum point for two minimax processes with $\zeta = 2.4551 \times 10^{-5}$ for the $\widetilde{F} = 3$ and $\widetilde{F} = 4$ repumping schemes. The two temperatures are comparable, but the population is over twice as large for the $\widetilde{F} = 3$ repumping scheme.

The dynamic method predicts that an 866nm scan taken around those parameters will be very hot. This is exactly what was seen in the lab, and no cooling to of order the Doppler limit was observed. The optimal scan produced using the dynamic method shown in figure 6.22 will not suffer from this problem and can be used to cool the ion below the Doppler limit with a corresponding $P_{1/2}$ population of over 6%.

6.6 850nm/854nm cooling

It is interesting to compare the results of cooling schemes based on dark resonances to the predictions of a method that isolates the cooling transition from all other driven transitions, thereby removing all dark resonances and substantially simplifying the fluorescence profile. Figure 6.23 shows a five-level cooling scheme that avoids the need for a beam to repump population onto the cooling transition. Population escapes the cooling transition through spontaneous decay from the $P_{1/2}$ to the $D_{3/2}$. Population is transferred back to the cooling transition through another decay process. Population is excited from the $D_{3/2}$ to the $P_{3/2}$ using a beam of wavelength 850nm. Population in the $P_{3/2}$ then decays back to the $S_{1/2}$ and returns to the cooling transition. The $P_{3/2}$ also has a D level that it can decay to, however. Population must be repumped from this $D_{5/2}$ level using an additional beam with wavelength 854nm. This scheme makes use of five levels and three lasers and therefore involves the largest and most computationally challenging matrices examined in this chapter.

Figure 6.24 shows the two 397nm scans for optima that make use of this cooling scheme. The lefthand plots show the $P_{1/2}$ population and temperature for an optimum that maximises the $P_{1/2}$ population. The temperature of the ion is calculated using the steady-state method. Population is repumped using a spontaneous rather than a stimulated process; the cooling transition looks like a two state system. In general, this transition is not driven near its reso-

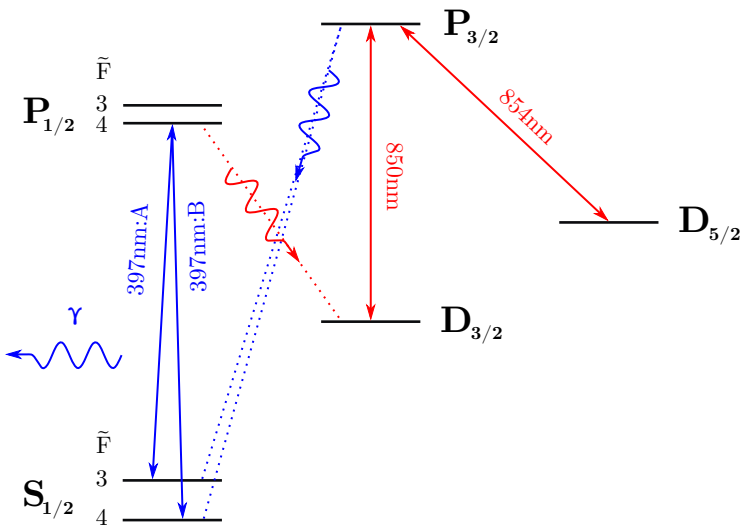

Figure 6.23: A cooling scheme for $^{43}Ca^{+}$ that makes use of the five lowest energy levels and avoids the need for a laser beam to repump population back onto the cooling transition. Population escapes the cooling transition by spontaneous decay from the $P_{1/2}$ to the $D_{3/2}$. This decay is shown by the red dotted line. The 850nm beam transfers population to the $P_{3/2}$, and population can decay from this state back to the $S_{1/2}$ and the cooling transition. This decay is shown by the blue dotted line. The $P_{3/2}$ can also decay to the $D_{5/2}$; the 854nm beam repumps population back from the $D_{5/2}$ to the $P_{3/2}$.

nant ion/photon frequency, so we do not need to use the dynamic method. The steady-state method calculates more quickly than the dynamic method, and this is an especially important consideration for a system with many states that computes slowly. The maximum population is significantly higher than can be observed using the three-level schemes discussed in sections 6.4 and 6.5. The maximum $P_{1/2}$ population that can be observed using this scheme is 21.79%. This is almost a factor of two larger than the populations that can be observed using a three-level scheme; see figure 6.17. The temperature of the ion is, however, correspondingly high.

The righthand plots of figure 6.23 show the two measures for a minimax optimum with $\zeta = 1 \times 10^{-5}$. Although this value favours low temperatures, the minimal temperature of the ion is four times higher than the Doppler limit and almost an order of magnitude higher than can be achieved using a three-level cooling scheme. Even though higher fluorescence rates can be achieved

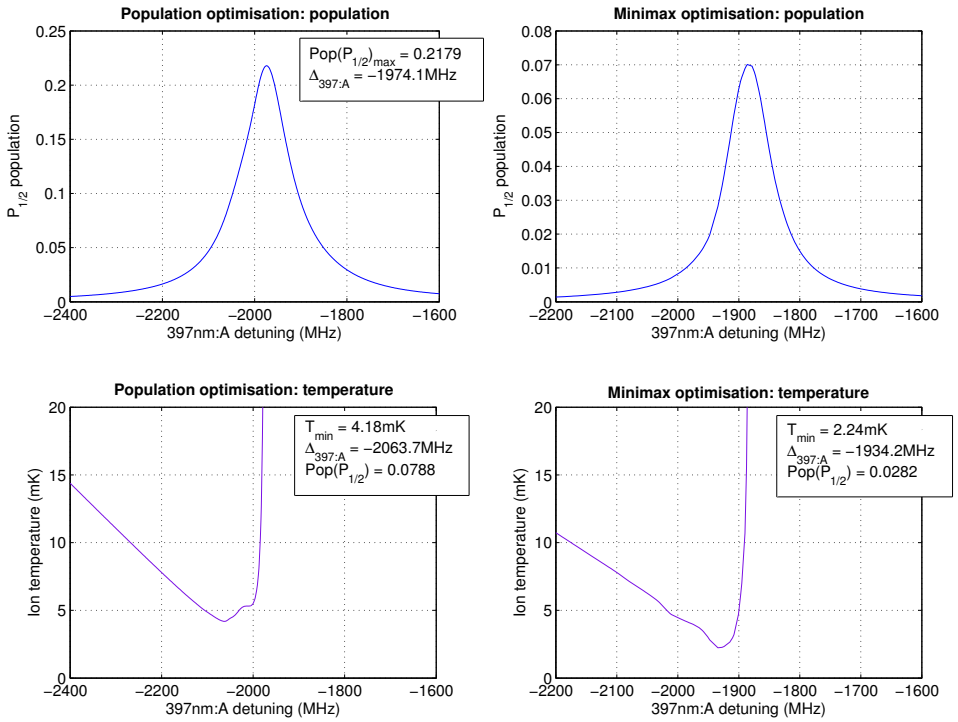

Figure 6.24: Two 397nm scans taken for different optima calculated using the 850nm/854nm cooling scheme. The lefthand plots show the $P_{1/2}$ population and temperature for a scan that maximises the $P_{1/2}$ population. The righthand plots show the same measures for a minimax optimum with minimax parameter $\zeta = 1 \times 10^{-5}$. The temperatures were calculated using the steady-state method; there is no repumper, so the predictions of the dynamic and steady-state method will be the same. The maximum $P_{1/2}$ population that can be observed in this system is 21.79%. The minimal temperatures seen in both cases are over four times higher than the Doppler limit.

using this scheme, it cannot be used to efficiently cool to temperatures below the Doppler limit. Dark resonances offer the best opportunities to cool an ion. They can provide high gradients of the fluorescence curve at points on the Lorentzian envelope that also produce high population. Dark resonances can produce far stronger cooling than a Lorentzian alone.

6.7 Fitting fluorescence scans with complicated structure

The cooling scheme that was used in all experimental work was a three-level cooling scheme using $\widetilde{F} = 4$ repumping and exciting the perfect dark resonance

α. This scheme was successfully used to cool an ion to the point at which fluorescence scans could be taken and analysed. This section examines a series of 866nm scans and demonstrates both that the ion behaves as predicted by the Bloch equations when it is cold and how the average temperature of the ion changes with beam parameters.

Figure 6.25 shows the very first 866nm scan that was fitted using the Bloch equations routines. The fitting process was global and constrained within a feasible range of the nominal experimental parameters. The floated parameters were the beam intensities, the beam detunings, the DAC calibration and the background fluorescence level. The frequency difference between the two 397nm beams was kept fixed at the known value of 2.930GHz and the intensities of these two beams were assumed to be equal. No fitted parameters took values at the extremes of their ranges; such behaviour implies that either the ranges or some assumptions about the data are incorrect, and this is checked in the analysis of any fit.

The agreement between the theory and data is impressive. The only major discrepancy occurs to the high-frequency side of one of the dark resonances. The is exactly the type of region where we would expect the ion to be hot and therefore to disagree with the cold ion model. In an 866nm scan, every resonance will contribute to heating the ion on the high-frequency edge. Those are the points of the scan where the ion is hottest, and thus where there is likely to be the greatest discrepancy between the ion's fluorescence and the predictions of the cold ion model. By the same logic, the significant agreement between the cold ion model and the observed fluorescence implies that, for the majority of the scan, the ion is cold to within a few multiples of the Doppler limit. The Doppler cooling scheme has successfully cooled the ion.

The ion was cooled on the perfect dark resonance α, so we expect this resonance to appear in the scan. We can easily identify it in the fitted 866nm scan

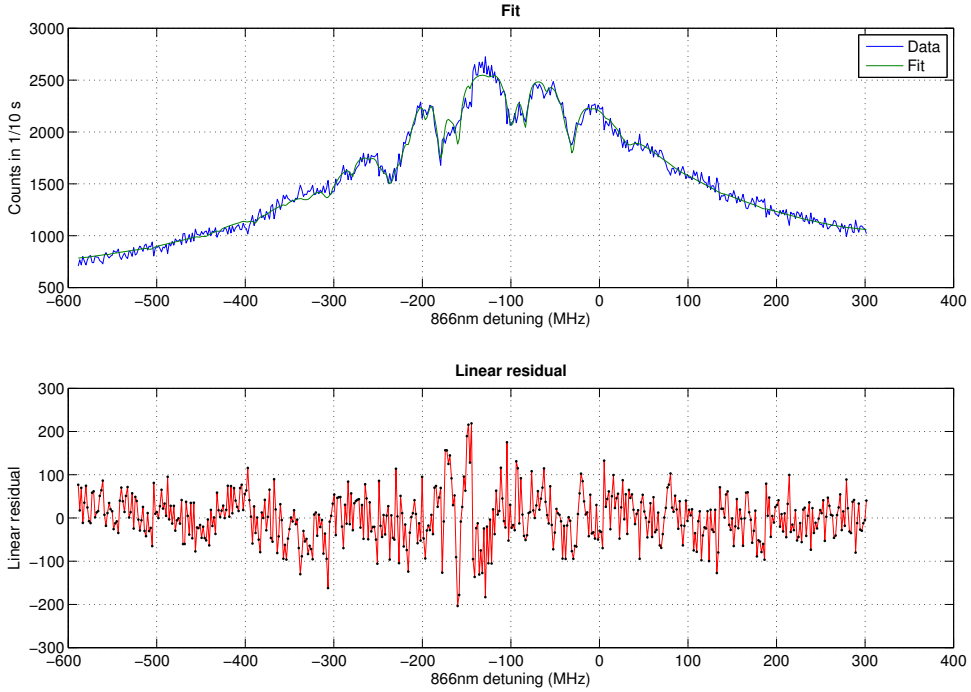

Figure 6.25: The very first fit to an 866nm scan that was calculated using the Bloch routines. The upper plot shows the data and fit. Each data point shows the number of counts detected in 1/10 s. The fitted curve was converted to this scale by adding on the fitted value of the background and using the conversion of equation 4.11. The lower curve shows the linear residuals. The fit is an impressive match to the data over a large portion of the scan. The only significant discrepancy is on the high-frequency side of the resonance centred at $\Delta_{866} = -180$MHz. We expect discrepancies to occur on the high-frequency side of resonances. In these regions, the resonance contributes to heating the ion. When the ion is hot enough, a discrepancy will be seen with respect to the cold ion model.

by rerunning the simulation and setting the linewidths of all beams to zero. The perfect dark resonance can then be identified as the single resonance where the $P_{1/2}$ population drops to zero. Figure 6.26 shows the result of such a test. There is only a single resonance where the population drops exactly to zero, that centred at $\Delta_{866} = -180$MHz. This is the same resonance that produces the only area of detectable heating of the scan; both the cooling and heating effects of this resonance are stronger than all other resonances, so we expect both the ion to be at its coldest to the low-frequency side of the resonance and at its hottest on the high frequency side. For any fluorescence scan taken of a cold $^{43}Ca^{+}$ cooled using this scheme, we expect that, on the heating side of the perfect dark

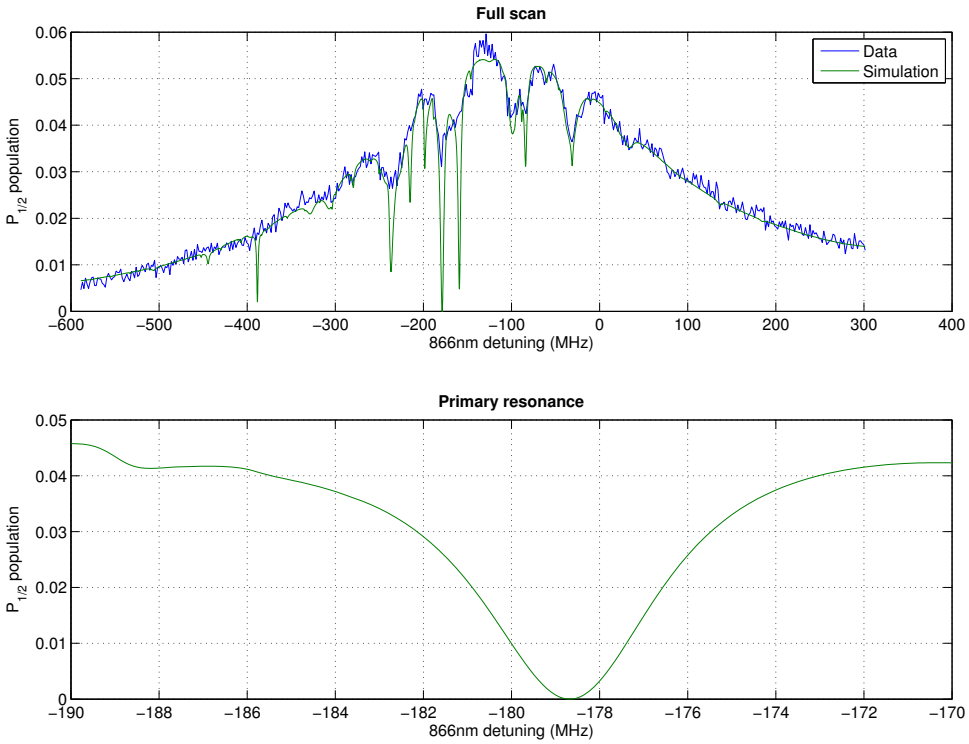

Figure 6.26: A demonstration of the perfect dark resonance in an experimental 866nm scan. The upper plot shows the measured fluorescence, adjusted to a scale showing the $P_{1/2}$ population using equation 4.11. Overlaid is a curve calculated using the fitted parameters but with the linewidths of the 397nm and 866nm beams set to zero. The population drops to zero on resonance α (see (6.1)). The lower plot shows the curve produced by the simulation centred around this resonance. This region contains the extremes of the ion's temperature, both in terms of cooling and heating the ion. The only region of the scan in which the ion is significantly hot is on the high-frequency side of this resonance.

resonance, the ion will be hot and discrepancies with the cold ion model will occur.

Figure 6.27 shows the temperature of the ion over this scan calculated using the dynamic method. This scan was taken with the beams co-propagating so the minimal temperature of the ion is not especially low; the relative direction of propagation of the beams was later changed to allow for lower temperatures to be obtained. However, even with a non-optimal beam arrangement, the ion was cooled to a minimum temperature approximately twice that of the Doppler limit. In addition, the temperature stays to within three or four times the Doppler limit for the majority of the scan. There are a few regions where the

Scan	397nm power (μW)	866nm power (μW)
a	20.0	30.0
b	20.0	60.0
c	20.0	120.0
d	40.0	60.0

Table 6.4: The measured powers of the 397nm and 866nm beams for the data shown in figure 6.28. This power refers to the total power of each beam. The 397nm power is shared between the carrier and two sidebands.

temperature becomes very large, but all but two of these occur over a very small range of frequencies. These narrow high temperature regions will be smoothed out in practice by the excursion effect. The only region of the scan where the ion is hot for a large frequency range is to the high frequency side of the perfect dark resonance. On the plot, this region is split in two by the cooling effect of the neighbouring strong dark resonance, but in practice the heating on either side will be strong enough that the excursion effect will ensure the ion remains hot over the entire region. This behaviour corresponds exactly to what we saw and inferred from the fluorescence scan itself. The ion is cold to within a few multiples of the Doppler limit for the majority of the scan and heats strongly to the high-frequency side of the perfect dark resonance.

We now examine a series of scans taken later. These scans used counter-propagating beams and consequently could reach significantly lower temperatures. They were taken to examine the behaviour of the temperature of the ion as the intensities of the 397nm and 866nm beams were varied. The first scan in the set was taken with parameters as close to those shown in figure 6.25 as possible. The subsequent scans were taken by altering the power in the various beams as shown in table 6.4 .

The four 866nm scans and the results of the fitting processes are shown in figure 6.28. To aid in the comparison between the plots, those regions of the theory curve for which the dynamic method predicts a finite temperature are shown in green while those regions for which the method predicts an infinite

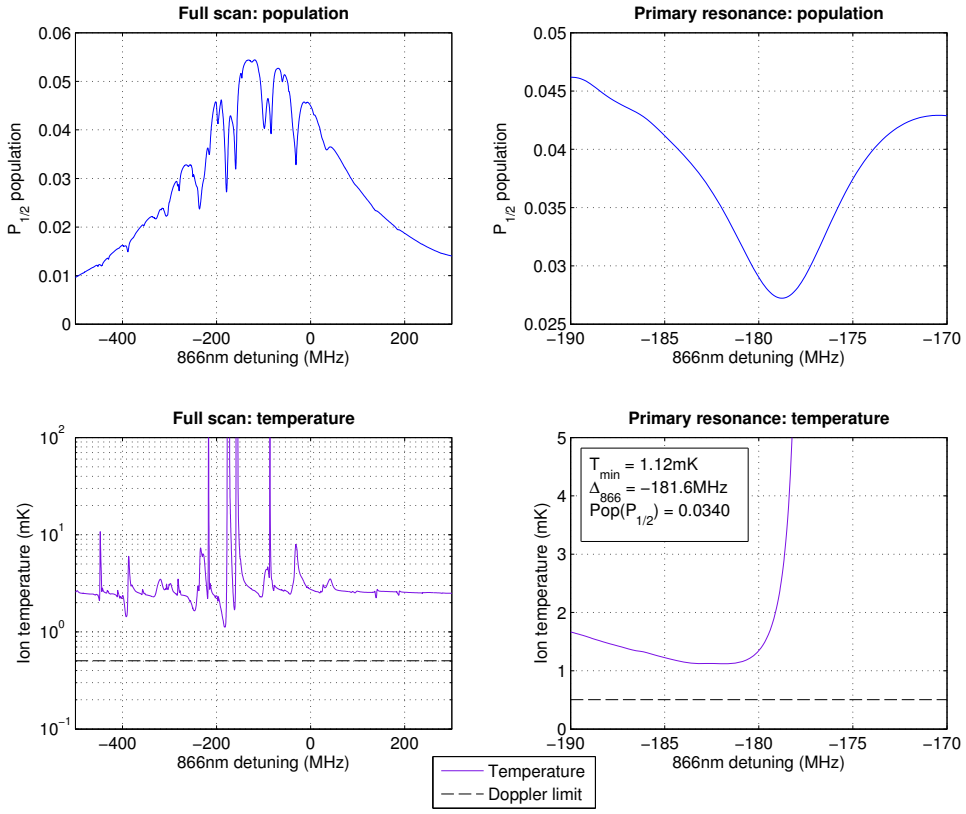

Figure 6.27: The temperature predictions of the dynamic method for the 866nm scan shown in figure 6.25. The two lefthand plots show the $P_{1/2}$ population and the ion's temperature over the entire scan. The righthand plots show the two measures over the primary cooling resonance. The beams in this arrangement were co-propagating, so the minimum temperature of the ion is not as low as the values shown in subsequent figures. However, the ion cools to within approximately a factor of two of the Doppler limit and is cold to within three or four times the Doppler limit for the majority of the scan.

temperature are shown in red. The behaviour of the ion is generally expected to match the cold ion behaviour for regions dominated by green and to show heating in regions dominated by red. The temperature predictions are shown separately in figure 6.29. Combined plots including the fit, temperature and linear residuals are shown in appendix F.

Scan a was taken with the lowest beam intensities and fits the data very well. The only significant discrepancy between the theory curve and the data is on the high-frequency side of the dark resonance, a region where we always expect the ion to be hot. As the beam powers increase in subsequent scans, the

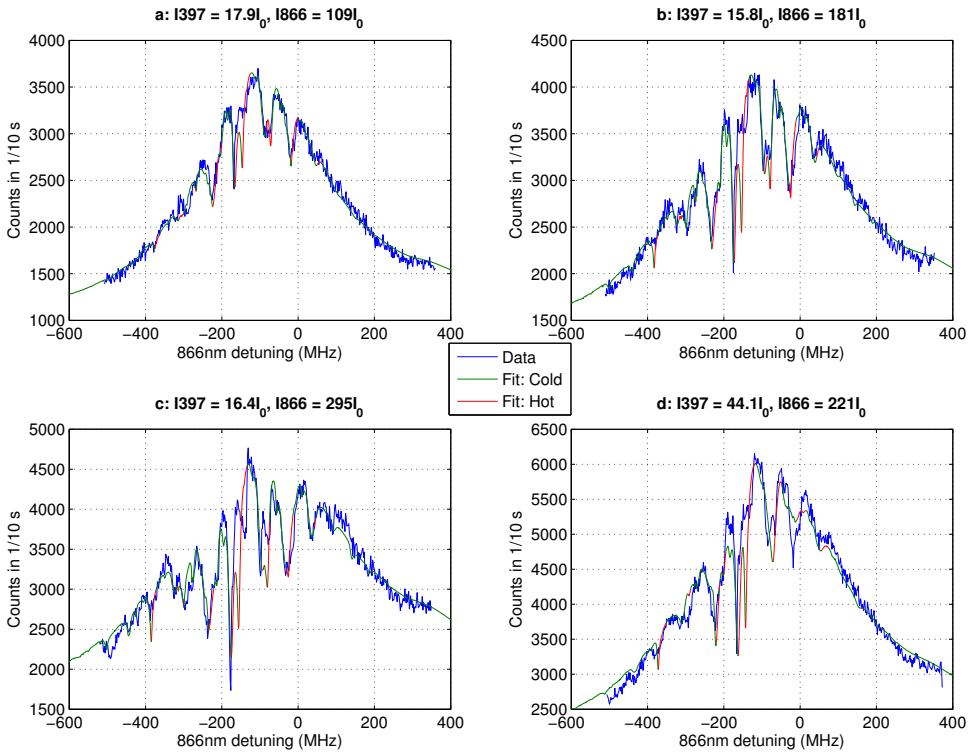

Figure 6.28: A series of 866nm scans taken by successively varying the intensities of the 866nm and 397nm beams. The nominal beam powers are listed in table 6.4. The fitted beam intensities are shown above each plot. Each plot shows the data and the results of a fitting process. The fitted curve is colour-coded. Green regions of the curve are those for which the dynamic method predicts a finite temperature and red regions are those for which the method predicts an infinite temperature. Generally, we expect green-dominated regions to be cold an red-dominated regions to be hot. The size and strength of the heating regions increases as the intensity of either beam increases. For more information, see text.

match between the fitted curve and the data becomes successively worse. This is not because the size of the heating regions predicted by the dynamic method increase; the size and positions of the red regions remain approximately the same for all four scans. However, the effect of heating in these regions becomes much more pronounced.

The fits were calculated by neglecting those regions of the scan in which heating was expected to occur. In scan a, only one of the resonances shows a discrepancy from the cold ion model; this scan is easy to fit. However, in subsequent scans, heating is observed on the high-frequency side of more and more resonances, and the discrepancy between the data and the cold ion model be-

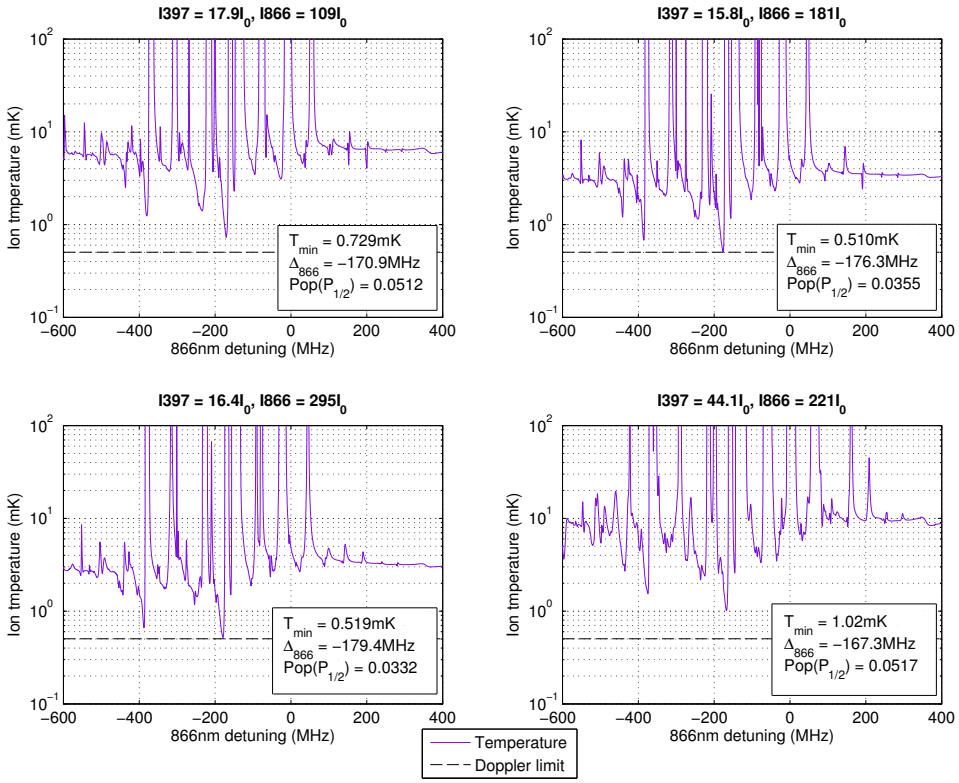

Figure 6.29: The temperature predictions of the dynamic method for the scans shown in figure 6.28. The Doppler limit is shown on each plot as a horizontal black dashed line. The increasing strength and size of the heating regions cannot be seen on the plot as this behaviour arises due to the excursion effect which is not modelled here. The minimal temperatures of the ion for scans a to c are of order the Doppler limit, although they are not lower than this limit. The ion is cold to within three or four times to Doppler limit for the majority of these scans. The temperatures for scan d are higher than the other three scans over its entire range. This is due to the larger 397nm power. The ion is especially hot for this scan, with temperatures typically an order of magnitude larger than the Doppler limit.

comes larger. Because the size of the regions in which there are large discrepancies increases from scan to scan, more and more regions needed to be removed from the data during the fitting process. Less information about the curve is being passed to the fitting routines and, consequently, the errors on the fitting parameters increase significantly. We can be confident that the values of the temperature predictions of scan a are are a good description of the ion, but we cannot say the same for scan d. A comparison of the measured and fitted values of the intensity is shown in table 6.5. The measured and fitted intensities change in the same sense between separate scans. Scans with higher beam powers show

Scan	Data: I397	Fit: I397	Data: I866	Fit: I866
a	20.0	17.9	30.0	109
b	20.0	15.8	60.0	181
c	20.0	16.4	120.0	295
d	40.0	44.1	60.0	221
Scan	Data: I397 ratio	Fit: I397 ratio	Data: I866 ratio	Fit: I866 ratio
a	1.0	1.0	1.0	1.0
b	1.0	0.88	2.0	1.7
c	1.0	0.92	4.0	2.7
d	2.0	2.5	2.0	2.0

Table 6.5: A comparison of the measured and fitted 397nm and 866nm intensities for the four 866nm scans shown in figure 6.28 and the ratio of the intensities compared to the values for scan *a*. The fitted 866nm intensities are uniformly larger than the measured powers. This is likely due to an error in the conversion between power and intensity such as, for example, an error in the spotsize. We remove the effects of such a conversion error by examining the ratios of the beam powers. The ratios for the data and the fits agree well, implying that the fits are a good model to the data.

less agreement between the measured and fitted intensities.

We explain the increase in the size and strength of the heating regions using the analysis of subsection 6.4.2. Whenever the intensity of either beam is increased, the minimum temperature that can be observed on any resonance rises[3]. In other words, the strength of the cooling effect of the resonance decreases. By reducing the cooling effect of all resonances, the heating effects in the scan become more pronounced. As a result, as the intensity of the beams increase, the size of all heating regions will increase and further heating regions will appear on the high-frequency side of dark resonances that were previously cold.

This increase in temperature is a result of the excursion effect, so cannot be directly observed in the dynamic temperature predictions; we infer the temperature increase due to the size of the regions in which there are large discrepancies between the data and the cold ion model. For all four scans, the average temperature disregarding the heating regions is predicted to be within a factor of

[3]We assume that the 866nm intensity is not very low here as reasonable repumping rates are required to achieve moderately high levels of fluorescence.

three or four of the Doppler limit. The minimum temperature of the scan is also of order the Doppler limit. The temperature predictions of scan d, however, are uniformly higher than the other three scans. Increasing the 397nm intensity has a greater effect on the ion's temperature than increasing the 866nm intensity. This is because the 397nm directly drives the cooling transition and therefore dominates the heating and cooling processes; the 866nm alters the ion's temperature only indirectly by controlling the rate at which population is repumped onto the cooling transition. The cold ion model is not a good one to describe a curve with a 397nm intensity that is this large, explaining why the agreement between the data and the model for scan d is significantly worse than the other three scans.

Figure 6.30 shows the results of a fit to a scan with very high beam powers. The fitted values of the 397nm and 866nm intensities are $I_{397} = 12I_0$ and $I_{866} = 518I_0$. Scans with intensities this large are desirable because they produce high fluorescence rates over a large portion of the scan. As with scan d in figure 6.28, the high intensities mean that the size of the heating regions are large and the scan is very difficult to fit. For this scan, the intensities are large enough that the size of the heating region on the high frequency side of the primary resonance extends over multiple neighbouring resonances. Fitting scans with very high intensities is a difficult task.

6.7.1 Temperature estimation using the fluorescence curve

For the scans discussed in this chapter so far, all the temperature predictions have been calculated by fitting a curve to the data and then using the fitted parameters as an input to the dynamic method. However, in some circumstances, it may not be possible to fit the data reliably without additional information. This may be because the intensities are too large to produce a good fit using the cold ion model. Alternatively, the range of frequencies over which the data is

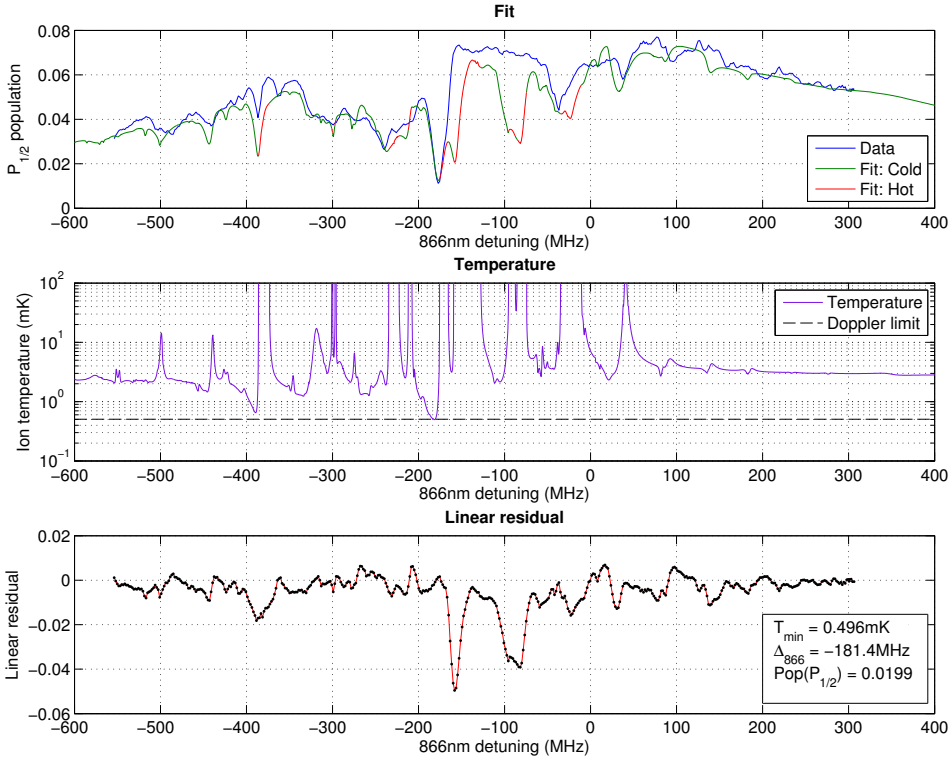

Figure 6.30: A fit and temperature predictions for an 866nm taken with very high beam intensities. The fitted values of the intensity are $I_{397} = 12I_0$ and $I_{866} = 518I_0$. The upper plot shows the data and the fit using the same colour convention as in figure 6.28. The middle plot shows the temperature predictions calculated using the dynamic method. The Doppler limit is shown as a horizontal black dashed line. The lower plot shows the linear residuals. There is a significant discrepancy between the fitted curve and the data. This is due to the large amounts of heating caused by the high beam intensities. The minimal temperature of the scan is approximately the Doppler limit.

taken may be too small to allow a fitting process to take place. For example, a scan may only include information around the coldest resonance. The set of parameters that can describe a single resonance is very broad indeed, so a scan like this cannot be fitted with a useful accuracy. In circumstances like this, we need another way to calculate the temperature.

It is possible to calculate the ion's temperature directly from the fluorescence curve, be that an 866nm or 397nm scan. We make use of the fact that the steady-state temperature is calculated using only the gradient and value of the scattering rate. We then approximate the scattering rate directly from the fluorescence curve.

Let us consider an ion sitting on the cooling side of a dark resonance on a fluorescence scan of the 866nm frequency ω_{866}. We change the frequency of this beam by a small amount $\delta\omega_{866}$. There will be a corresponding small change in the $P_{1/2}$ population. As a result, the scattering rate R will change by a small amount to $R + \delta R$ as the ion moves up the wall of the dark resonance. Now, we change the frequency of the 397nm beam ω_{397} by $\delta\omega_{866}$. The position of the dark resonance moves by $\delta\omega_{866}$ towards the ion's position. Assuming that R varies linearly with the two frequencies, the ion falls back down the wall of the resonance to its original position; the change in the population and scattering rate cancel out. We can therefore state:

$$\frac{\partial R}{\partial \omega_{866}} = -\frac{\partial R}{\partial \omega_{397}} \tag{6.5}$$

We can therefore measure the gradient on a scan of either beam and calculate the same temperature with no difference in scaling.

To calculate the steady-state temperature defined in (4.49), we need to calculate the gradient of the scattering rate with respect to the ion's velocity component v_x. We can express this derivative in terms of partial derivatives with respect to the two frequencies:

$$\frac{dR}{dv_x} = \frac{\partial R}{\partial \omega_{397}}\frac{d\omega_{397}}{dv_x} + \frac{\partial R}{\partial \omega_{866}}\frac{d\omega_{866}}{dv_x} \tag{6.6}$$

The frequency variations with velocity are given by the Doppler shift, so that:

$$\frac{dR}{dv_x} = \frac{\partial R}{\partial \omega_{397}}\frac{\omega_{397}}{c} \pm \frac{\partial R}{\partial \omega_{866}}\frac{\omega_{866}}{c} \tag{6.7}$$

where the positive sign describes co-propagating beams and the negative sign counter-propagating beams. Using (6.5), we can express dR/dv_x in terms of only a single partial derivative:

$$\frac{dR}{dv_x} = \frac{\partial R}{\partial \omega_{397}}\frac{1}{c}\left(\omega_{397} \mp \omega_{866}\right) \tag{6.8}$$

For a fluorescence scan, let us define the 'gradient temperature' as:

$$T_G = -\frac{\hbar}{k_B} R \left(\frac{\partial R}{\partial \omega}\right)^{-1} \tag{6.9}$$

where ω is the angular frequency of the scanned beam. Using (4.49) and (6.9), we relate the gradient temperature to the steady-state temperature using the factor θ_f:

$$T = \theta_f T_G \tag{6.10}$$

where:

$$\theta_f = \frac{\omega_{397}}{\omega_{397} \mp \omega_{866}} \quad . \tag{6.11}$$

This allows us to approximate the steady-state temperature directly from the gradient of the data without using any fitting or additional modelling. This factor is 1.9 for co-propagating beams and 0.69 for counter-propagating beams regardless of the frequency being scanned.

Figure 6.31 shows a comparison of the gradient temperature and the steady-state temperature for a scan taken using high beam powers. The ratio of the gradient temperature to the steady-state temperature at the coldest point on the scan is very similar to the approximate value in (6.11). The ratio calculated from the scan is $\theta'_f = 0.6515$ and the approximation is $\theta_f = 0.686$. The approximation is only valid for a small region centred on the coldest point of the scan. The approximation to the steady-state approach is very good in this region, predicting temperatures that differ from the state-method by only 5%. The ratio between the gradient and steady-state temperatures rapidly diverges from θ_f when the temperature of the ion increases; the assumption of a perfectly linear gradient of the scattering rate only holds when the ion is cold.

Figure 6.32 shows the predictions of the steady-state method, the gradient approximation and the dynamic method for a high resolution 866nm scan taken around the primary dark resonance. The steady-state and dynamic predictions were calculated directly from a fit to this scan. The approach used to fit this particular scan is discussed in section 6.8. The gradient approximation temperatures were calculated by fitting a cubic to the data on the low-frequency side of the dark resonance where the ion is cold. The steady-state and gradient ap-

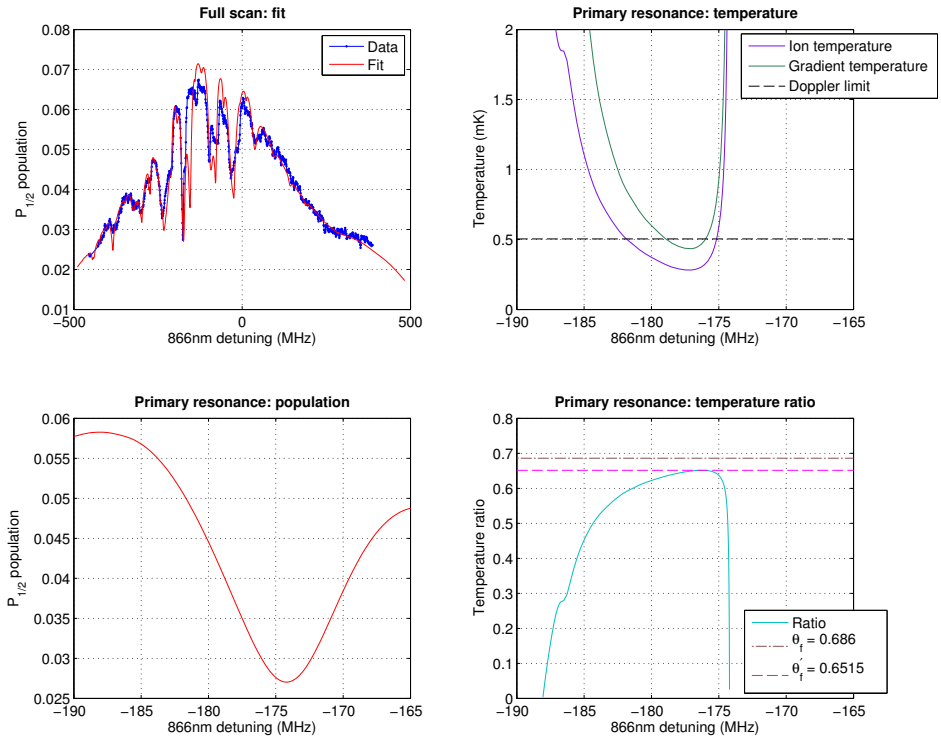

Figure 6.31: A demonstration of the difference between the steady-state temperature and the approximation to the steady-state calculated using the gradient of the fluorescence curve. The data analysed in this figure is an 866nm scan taken with high beam intensities. The plot on the upper-left shows the scan and fit. The lower-right plot shows the fitted curve in a region around the primary resonance. The upper-right plot shows the steady-state temperature and gradient temperature. The Doppler limit is shown as a horizontal black dashed line. The lower-right plot shows the ratio of the steady-state temperature to the gradient temperature. The gradient temperature is calculated using the fitted curve. The value of θ_f is shown as a horizontal brown dot-dashed line. The maximum value of the ratio calculated using the fit is shown as a horizontal pink dashed line. In a small region around the minimum temperature of the scan, the approximation to the steady-state temperature is very good, predicting a temperature that differs from the steady-state prediction by only 5%. The residuals for the fit and the temperature over the full scan are shown in appendix F.

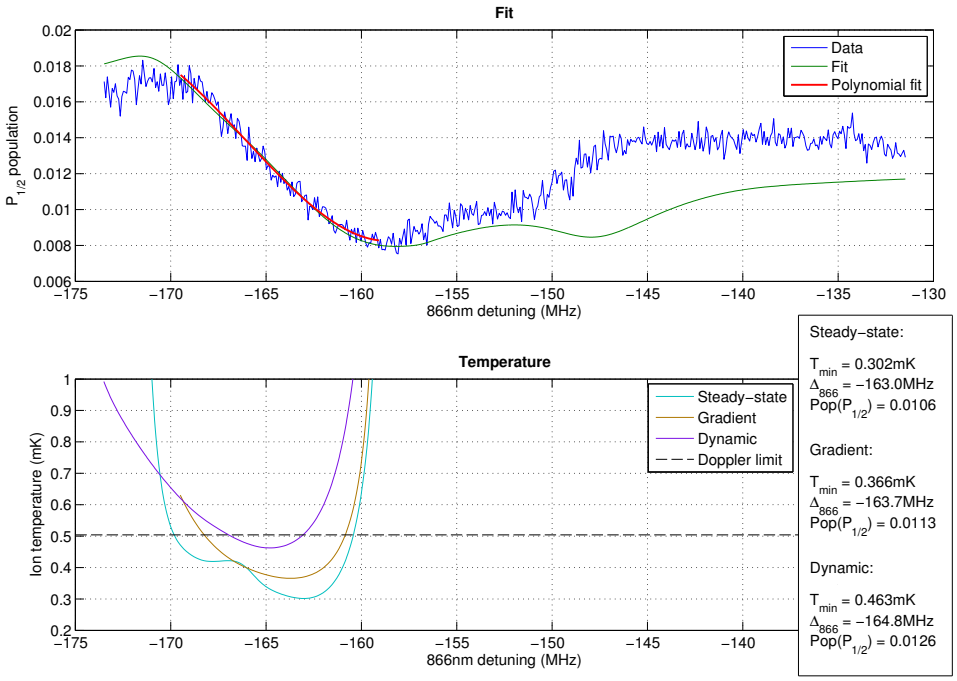

Figure 6.32: The predictions of the steady-state method, the gradient approximation method and the dynamic method for a high-resolution 866nm scan taken around the perfect dark resonance. The steady-state and dynamic predictions were calculated using a fit to this data. The method used to produce the fit is described in section 6.8. The upper plot shows the data, fit and a polynomial fit to low-frequency edge of the resonance. The lower plot shows the steady-state, gradient approximation and dynamic temperatures. The Doppler limit is shown as a horizontal black dashed line. The discrepancy between the steady-state and gradient approximation predictions is due to slight differences in the gradients of the two fits. The data has sufficient noise even at this resolution so that ambiguity in determining the gradient of the fluorescence curve can alter the predicted final temperature by approximately $\pm 20\%$. The dynamic method predicts a temperature that is significantly higher than both the steady-state method and its approximation. For more information about this scan, see section 6.8.

proximation methods disagree to a greater extent than the 5% value shown in figure 6.31. This is due to a slight difference in the gradients of the large range fit and the polynomial fit. Even for a high-resolution scan such as this, the noise on the data leads to a sizeable ambiguity in determining the gradient of the fluorescence curve. The steady-state fit takes more information into account than the polynomial fit, resulting in a reduced error when determining the gradient of the curve and a correspondingly reduced error in determining the ion's temperature. Both the steady-state and gradient approximation methods underestimate the temperature somewhat.

6.8 Cooling below the Doppler limit

In this section, we analyse a set of 866nm scans taken with the aim of reaching as low a temperature in $^{43}Ca^{+}$ as possible. These scans were taken around a series of parameter sets at which the temperature of the ion was directly measured using the technique of sideband thermometry which is discussed in more detail in [Har13]. The combination of the scans and the measured temperatures allows a direct test of the predictions of our temperature models.

The analysis presented in section 6.4.2 predicted that the ion's temperature always decreases with decreasing 397nm intensity. These scans were therefore taken with as low a 397nm intensity as could be achieved while still detecting sufficient fluorescence to produce an 866nm scan with acceptable noise. The ion was cooled using dark resonance α and the standard $\widetilde{F} = 4$ repumping scheme.

The output of the sideband thermometry measurements is a mean occupation number of an equivalent harmonic oscillator viewed along a particular axis direction in the trap. In order to compare the measurements to the temperature values produced by my simulations, we need to convert between a temperature expressed in terms of a mean occupation number $\bar{n}$ and a temperature expressed in units of Kelvin, T. The trap used to carry out these experiments is discussed in more detail in [AHB+13, All11, Har13]. The temperature of the ion is measured along one of the radial axes of the trap. For more information about the process of collecting the experimental data, see [Sep16].

The energy of a harmonic oscillator in one dimension takes the form:

$$E_m = \hbar\omega_m\left(\bar{n} + \frac{1}{2}\right) \tag{6.12}$$

where ω_m is the secular frequency along this dimension. We have defined the ion's temperature in terms of its energy:

$$E = 3k_B T \tag{6.13}$$

The cooling will only occur in one dimension along the line of the 397nm Doppler cooling beam. Let us assume that the energy of the ion associated with its motion along each cardinal direction is equal. In this case, the energy of the ion along the direction of the cooling beams is a third of E:

$$E_m = k_B T \tag{6.14}$$

Equating (6.12) and (6.14), we produce the following conversions between temperature and the mean occupation number:

$$T = \frac{\hbar\omega_m}{k_B}\left(\bar{n} + \frac{1}{2}\right) \tag{6.15}$$

$$\bar{n} = \frac{k_B}{\hbar\omega_m}T - \frac{1}{2} \tag{6.16}$$

For the purposes of this study, we first assume that the trap is isotropic with a secular frequency of 3MHz. This is the value recorded for the trap radial secular frequencies in [Har13]. We can express the conversions as:

$$T = 0.144\left(\bar{n} + \frac{1}{2}\right) \tag{6.17}$$

$$\bar{n} = 6.95T - \frac{1}{2} \tag{6.18}$$

where T is assumed to be in units of mK.

The value of the ion's temperature was optimised experimentally. In total, four parameters were adjusted. Two of these parameters were the total power in the 397nm and 866nm beams, and two described the detuning of the beams. The frequency splitting of the 397nm beam was fixed at 2.390GHz, the optimal value determined by my simulations. The beam polarisations were fixed at $[f(\sigma^+), f(\sigma^-), f(\pi)] = [2/3, 0, 1/3]$ for the 397nm beam and $[f(\sigma^+), f(\sigma^-), f(\pi)] = [1/2, 1/2, 0]$ for the 866nm beam. These choices ensured that the ion is cooled on the perfect dark resonance α. The ratio of the powers in the two 397nm beams was fixed at P(B)/P(A) = 0.85.

Each of the parameters was adjusted in turn. The temperature was determined approximately when adjusting the first three parameters by observing

the height of the red sideband. As the height of the sideband decreases, the temperature of the ion also decreases. The final parameter to be adjusted was the 866nm detuning. Sideband scans were recorded as this detuning was adjusted. The optimal temperature of the ion was calculated from these scans.

The standard method of measuring the temperature of an ion using sideband scans is to make use of the asymmetry between the red and blue sidebands. However, that approach cannot be used for this system as the magnitudes of the temperatures are too large to produce a sufficiently large sideband asymmetry. Instead, the temperatures are determined by fitting the sideband scans directly. For more information on this process, see [Sep16]. This fitting took place during the experimental run.

The lowest temperature reached by the ion was:

$$\bar{n}_{min} = 1.97(5) \tag{6.19}$$

Using conversion (6.17), we can express this in units of mK:

$$T_{min} = 0.356(7)\text{mK} \tag{6.20}$$

This temperature is below the Doppler limit:

$$\frac{T_{min}}{T_D} = 70\% \tag{6.21}$$

Once the lowest temperature of the ion had been determined, a high resolution 866nm scan was immediately taken around the dark resonance in order to provide data that could be analysed by my routines. The range of this 'narrow scan' was small, extending over only approximately 40MHz, and covered only the dark resonance. This scan alone does not contain enough information to allow fitting. A wide range of parameters could produce a scan that describes a resonance such as this. In addition, the lowest temperature of the scan determined using the dynamic method is a function not only of the slope of the

resonance but also of its height; we cannot be certain of the height of the resonance on the low-frequency side due to potential heating near the peak caused by other resonances.

To provide sufficient data to analyse the narrow scan, a second series of scans was taken later. An 866nm frequency scan was taken over the entire profile, including a significant portion of its wings. This was achieved by using three different 866nm ranges (the 866nm profile was too large to be covered with a single scan). The three different ranges extended over approximately 600MHz each. For each range, three scans were taken scanning the 866nm beam in a negative sense (from high to low frequency) and three taken scanning the beam in a positive sense. Each of the negative and positive scans were then knitted together by matching resonant features to produce a scan over a 1500MHz range. The result is two 866nm scans over a large range of frequency values taken in a positive and negative sense. I refer to these scans as 'wide scans'. By fitting them, an approximate series of parameters describing the narrow scan can be produced. The narrow scan can then be fitted using these parameters as a starting point and heavily constraining the optimisation process.

One advantage of producing a scan in this way is that some heating regions can easily be identified. By comparing scans taken in a negative and positive sense, regions of the curve that demonstrate hysteresis effects can clearly be seen. For at least one of the two curves, the ion must be hot in this region; the cold ion model cannot account for hysteresis. To ensure that no potential heating regions are passed to the fitting routine, these regions are not fitted for either curve.

Figure 6.33 shows the output of the fitting process for the wide scan taken in a positive sense. There are large regions in which the cold ion model shows significant discrepancies with the data; the ion is hot for a large part of this scan. This heating arises due to the choice of an optimal 397nm detuning discussed

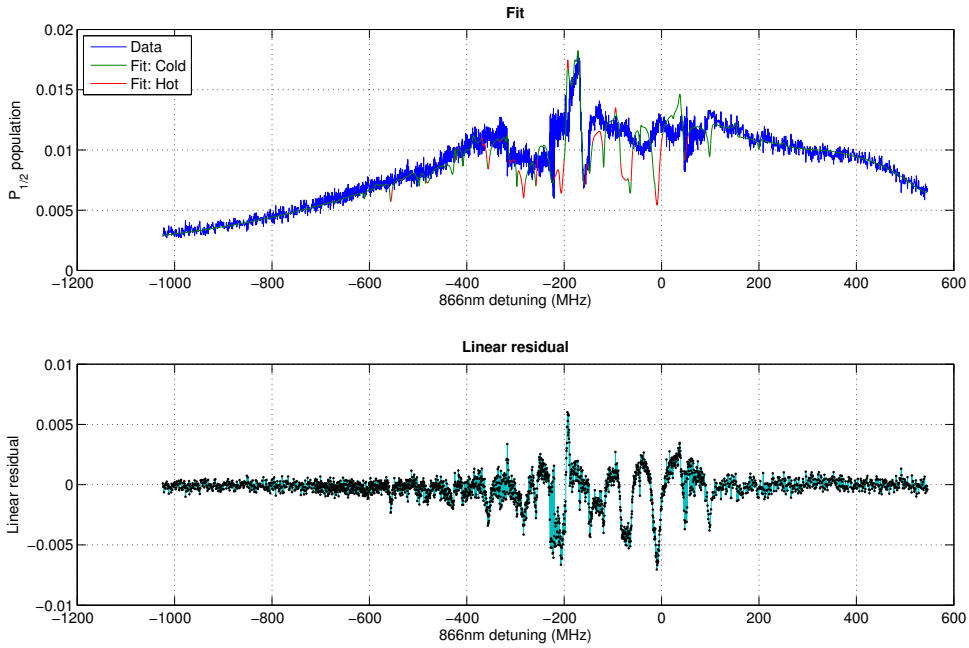

Figure 6.33: A fit to the 866nm wide scan taken in the positive sense. The upper plot shows the data and the fit. The fit is colour-coded using the same convention as in figure 6.28. Any parts of the curve dominated by red will show significant heating. There are a large number of red regions on the fit, and large regions of the scan in which there sizeable discrepancies between the data and the cold ion model. The ion is hot for a significant portion of the scan, meaning that this data is difficult to fit. The lower plot shows the linear residuals. For more information, see text.

in section 6.4.2. The temperature optimisation process started with an ion that was already cold and pushed the 397nm detuning as far to the blue as could be achieved. The 397nm detuning was close to the optimal position predicted by my simulations, that is, the centre of the envelope of a 397nm scan. Any regions of the curve in which excursion becomes important will see the large heating region on the high-frequency side of the profile and will significantly heat the ion. As a result, the ion is hot for large parts of the scan. This is an unfortunate side-effect of reducing the temperature of the ion as much as possible.

We fit the ion by using those portions of the curve in which we are confident that the ion is cold. This consists of the low-frequency side of the principal dark resonance, which can easily be identified by the large slope of negative gradient on the the high-frequency side of the large spike in the profile, and those parts of the wings in which there are no visible features. This scan was fitted using these

three portions of the curve. We can be confident that the fit is a good description of the cold portions of the scan by noting several points. The height and width of the primary dark resonance are matched well. On the low-frequency side of the fluorescence spike, the horizontal region of fluorescence sharply rises at approximately the point at which it intersects the cold portion of the fitted curve. In addition, the data approximately follows the dip in the fluorescence over the region $-400\text{MHz} < \Delta_{866} < -200\text{MHz}$. With the exception of the low-frequency side of the dark resonance, none of these regions was fitted.

The temperature of the ion calculated using the dynamic method is shown in figure 6.34. The large size of the heating regions is indicated by the large number of places in which the temperature is predicted to be very high or infinite. Even in the wings, the predicted temperature of the ion is over an order of magnitude larger than the Doppler limit. However, the minimal temperature of the scan is still cold:

$$T_{min} = 0.416\text{mK} \tag{6.22}$$

Unfortunately, this temperature cannot be directly compared to the measured temperature (6.20) as this scan was taken significantly after that measurement was taken.

The location of the point of lowest temperature is also shown in the figure. This location was an important consideration when fitting this scan. Several different fits were made to this scan using different initial conditions corresponding to different heights of the central spike and different widths of the dark resonance. The narrow scan contains information not only about the lowest temperature of the ion, but also at what point on the scan at which this occurs. Any output to the fitting process that placed the optimal temperature at a radically different point on the slope of the dark resonance compared to the temperature measurements was assumed to be a poor description of the data. The fit displayed here is the output of this process that best matched the position of the

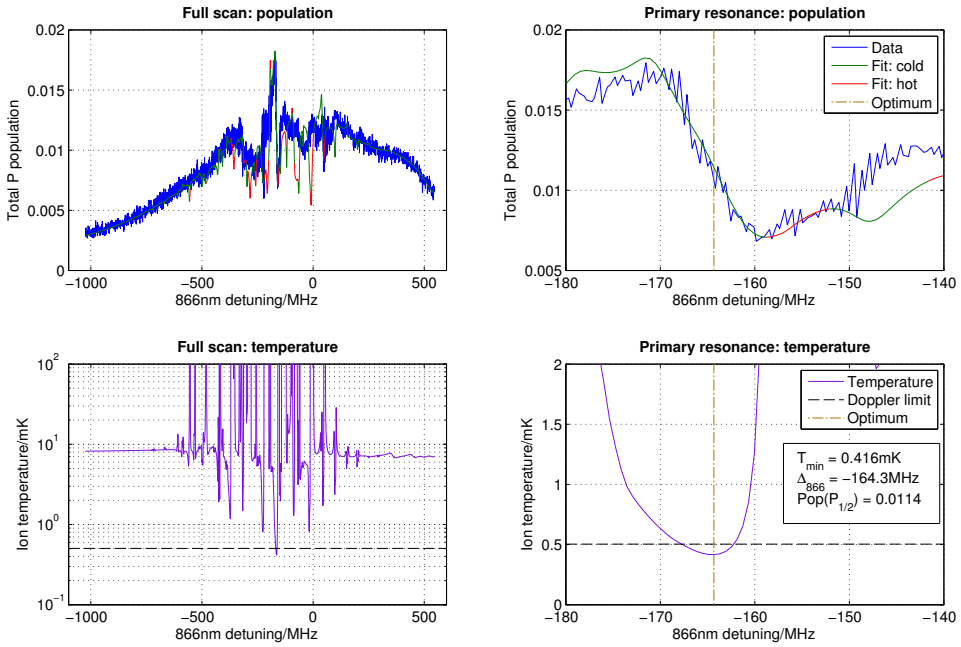

Figure 6.34: A fit to the 866nm wide scan taken in the positive sense. The lefthand plots show the $P_{1/2}$ population and the temperature calculated using the dynamic effect taken over the entire scan. The righthand plots show the two measures taken over the primary resonance. On these two plots, the position of the optimal temperature is shown by the vertical organ dot-dashed line. On the two temperature plots, the Doppler limit is shown by the horizontal black dashed line. The lowest temperature reached by the ion is 0.416mK.

optimal temperature point. It is also the scan that best matches the height and width of the primary dark resonance.

The parameters describing this fit were used an an input to a fitting process for the narrow scan. The range over which the parameters could float was heavily constrained. The 866nm intensity was very heavily constrained. This parameter directly controls the width of the Lorenztian profile of an 866nm scan; it can only vary by a small amount before it becomes inconsistent with the width of the wide scan. The results of this fitting process and the temperature of the ion are shown in figure 6.35. The position at which minimal temperature occurs, measured by the population at the optimal point, are very similar for both the data and the fit. The minimal temperature predicted by the dynamic method,

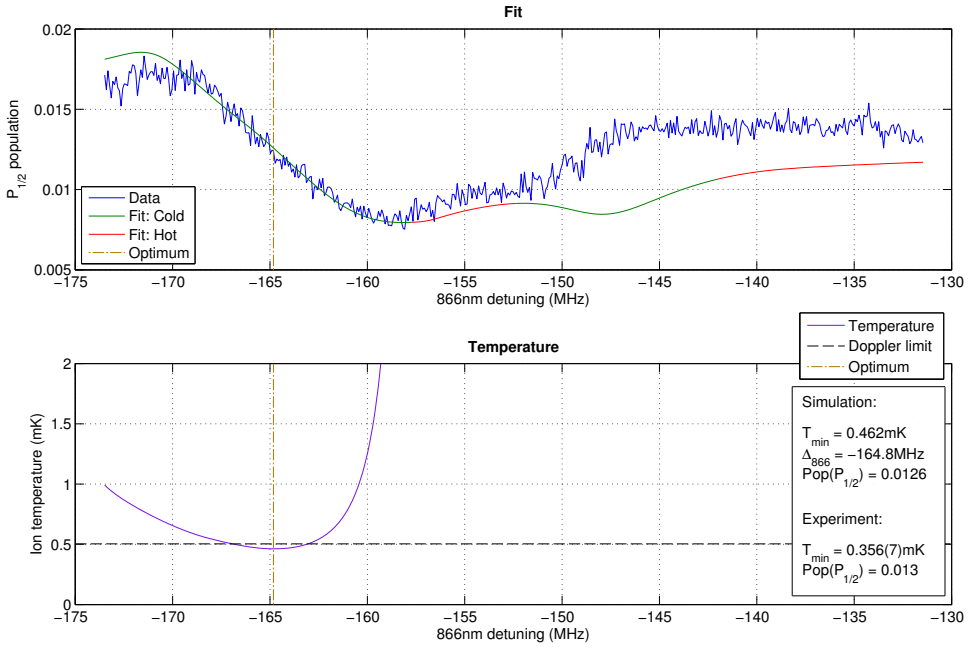

Figure 6.35: A fit to the narrow scan making use of the parameters describing the wide scan. The upper plot shows the fit and the data. The fit is colour-coded using the same scheme as shown in figure 6.28. The lower plot shows the temperature predicted by the dynamic method. On this plot, the Doppler limit is shown as a black horizontal dashed line. The point at which the temperature in optimal is shown on both plots as a vertical orange dot-dashed line. The minimal temperature and population of both the fit and the data, and the detuning of the optimal temperature point for the fit, are shown on the lower plot. This plot shows the same data and fit as that used to examine the gradient approximation method in figure 6.32.

along with the steady-state prediction and the measured value:

$$T_{min}(\text{Dynamic}) = 0.462\text{mK} \tag{6.23}$$

$$T_{min}(\text{Steady}-\text{state}) = 0.302\text{mK} \tag{6.24}$$

$$T_{min}(\text{Measurement}) = 0.356(7)\text{mK} \tag{6.25}$$

The prediction of the dynamic method is somewhat higher than the measured value, although the agreement between theory and experiment is still impressive. It is not possible to lower the dynamic method's prediction without disrupting the quality of the fit. The steady-state method also does not produce a temperature that is consistent with the measured value.

There are several reasons why we might expect a difference between the prediction of the dynamic method and the temperature of the ion. Primarily,

Axis direction	Secular frequency (MHz)
$\hat{x}$ (Radial)	3.270
$\hat{y}$ (Radial)	3.630
$\hat{z}$ (Axial)	0.477

Table 6.6: The secular frequencies along the three axis directions for the trap used to take this data. The temperature of the ion was measured along the $\hat{x}$ direction

the error listed with the measurement is only a statistical error produced by the routine used to fit the sideband scans. Systematic errors in the experiment are likely to produce much larger errors [AHS+16].

Errors might also arise from two assumptions of the routine used to model the dynamic effect. Firstly, the routine assumes that the ion behaves as a harmonic oscillator along the beam path. This trap is anisotropic, with differing secular frequencies along each of its three axis directions. These are listed in table 6.6. The temperature of the ion is measured along the $\hat{x}$ direction. The beams propagate along a direction that makes an equal angle to all three axes. The motion therefore traces out a path that is a function of all three secular frequencies and is therefore not harmonic in general. The dynamic method cannot be used rigorously to describe this system.

An extension to the motional method was constructed with the ability to describe the motion of an ion in three dimensions in a trap with three different secular frequencies. This extension used the one-dimensional cooling and heating terms used in the original model but, instead of reducing the size of the cooling term by a factor of a third, the energy of the ion in each dimension was equalised by adding an additional term to the equation of motion. The results of this approach show no significant temperature difference compared to the results of the dynamic effect in an isotropic trap. This implies that trap anisotropy is not the source of the discrepancy between theory and experiment.

There are two other possible sources for the disagreement. Firstly, the conversion from a temperature measured in $\bar{n}$ to a temperature in mK also assumed

that the ion behaved as a harmonic oscillator along the direction of propagation of the beams. This is not the case in the presence of the cooling beams. Secondly, a problem may arise due to the size of the temperatures being considered. The dynamic method is an inherently semiclassical model; the ion is assumed to oscillate sinusoidally in the trap. For a quantum harmonic oscillator with an $\bar{n}$ of order 2, this is not a good assumption. It may be that the quantum nature of these oscillations will reduce the strength of the dynamic effect. Because of the nature of these errors, we cannot claim that the predictions of the dynamic method is accurate to more than a few tens of percent.

A further set of data showed good agreement with our theoretical predictions. Figure 6.36 shows the results of a series of temperature measurements taken as each of the four optimised parameters are scanned from the point of optimal temperature. The trends seen as the parameters vary show impressive agreement between the model and the experiment.

6.9 A qubit at 287.78 gauss

In this chapter, we have primarily considered the behaviour of a qubit at 146.09 gauss. However, the cooling systems discussed here can also be applied to qubits at other fields. Figure 6.37 shows the results of a minimax optimisation process at a field of 287.78 gauss. This qubit makes use of the stretch state cooling scheme discussed in section 6.4; the scheme cools on perfect dark resonance α. The ion can be cooled to a temperature below the Doppler limit with a fluorescence rate that is only of order 10% less than the fluorescence rates that can be observed at 146.09 gauss. By making use of the $\widetilde{F} = 3$ repumping scheme discussed in section 6.5, the fluorescence of the ion could be increased by a factor of two from that shown here. This qubit will be the basis of further experimental work in ion traps that are currently being developed.

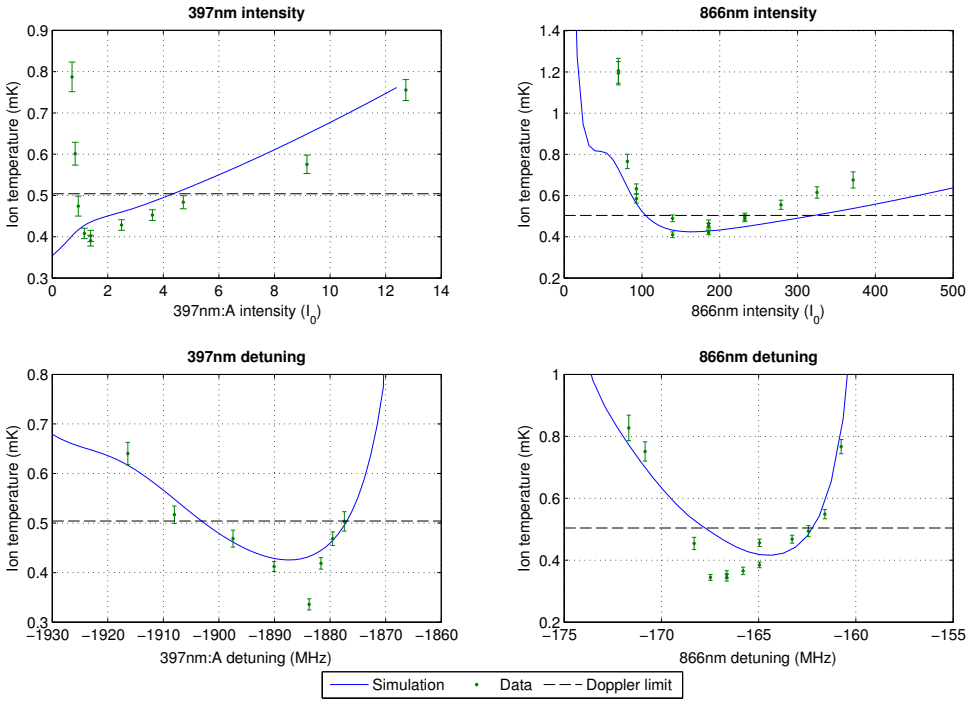

Figure 6.36: The variation in temperature of the data and simulation as each of the four optimised parameters are varied. The 397nm intensity, 866nm intensity and 866nm detuning curves were taken only stepping the relevant variable. The 397nm detuning curve was taken by stepping this detuning and then optimising all the other variables. The trend in the temperatures predicted by the theory show an impressive agreement with those in the data. The rise in temperature at low 397nm intensities is due to the ion failing to cool at all on the dark resonance. This is because a hot ion cannot resolve the dark resonance at these low powers; the size of the surrounding heating regions are too large and swamp the cooling power of the dark resonance.

6.10 Conclusion

In this chapter, we have analysed Doppler cooling in $^{43}Ca^{+}$ with the aim of developing a scheme to Doppler cool this isotope to a temperature of order the Doppler limit. We identified a field-insensitive hyperfine qubit in the ground state and, using the dynamic method, developed a dark resonance cooling scheme that was predicted to efficiently Doppler cool this isotope. This scheme was tested in experiment. Fluorescence scans taken from an ion cooled in this way were fitted and, taking account of regions in the scans where the ion is hot, a very good match between data and theory was observed. Finally, a comparison was made between the predictions of the simulations and measurements of

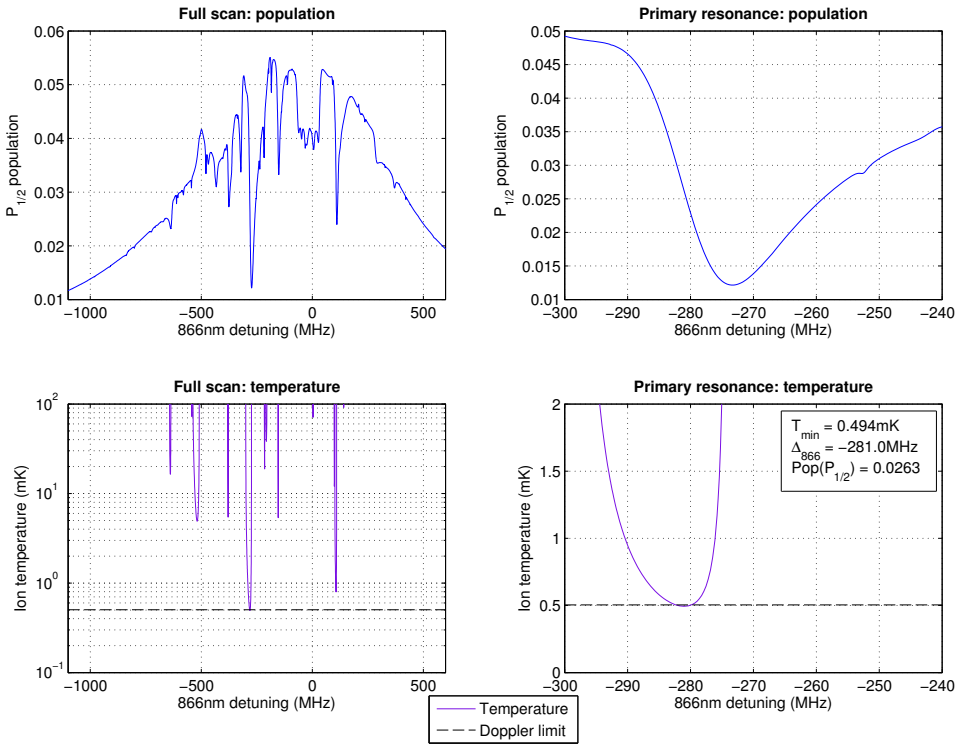

Figure 6.37: The results of a temperature minimax optimisation process at a field of 287.78 gauss using $\zeta = 2.46 \times 10^{-5}$. The lefthand plots show the $P_{1/2}$ population and temperature calculated using the dynamic effect over the entire scan. The righthand plots show the two measures calculated over the primary resonance. In the temperature plots, the Doppler limit is shown by the horizontal black dashed line. The minimum temperature of the ion is lower than the Doppler limit. The $P_{1/2}$ population is only of order 10% less than for optima calculated using this cooling scheme and this value of ζ at 146.09 gauss.

the ion's temperature for differing values of the beam parameters. The trends seen in data and simulation showed impressive agreement. The points of lowest temperature on these scans imply that the ion has been Doppler cooled to temperatures lower than the Doppler limit.

With the dark resonance cooling method as a starting point, experiments in the Oxford ion trap group have applied sideband methods to cool this ion to the ground state of motion to a good approximation [AHS+16]. This forms the starting point for ongoing quantum information experiments.

7

Summary and Outlook

That's the wonderful thing about crayons: they can take you to more places than a starship.

Guinan
Rascals, Star Trek: The Next Generation

This thesis described work on efficiently calculating the temperature and fluorescence of trapped ions using the optical Bloch equations, including identifying a cooling scheme which successfully Doppler cools $^{43}Ca^{+}$ at the critical magnetic field of a first-order magnetic field insensitive transition in the ground state to a temperature 30% lower than the Doppler limit.

The 'Global Linear Optical Bloch Equations Suite' (GLOBES) was developed in order to solve this problem. GLOBES is an efficient set of routines that solves the optical Bloch Equations for an arbitrary ion in the presence of an arbitrary set of laser beams. The suite can very efficiently solve the time independent problem, simulate fluorescence scans and solve the time dependent problem for small and large times. The 'Stepsize Iterations Locally Varied and Error Reduced IMPlicit EuleR' routine (SILVER IMPER) was developed to specifically solve the time dependent problem for very large times. This routine passes over the fast initial dynamics of such a problem to quickly and accurately model the approach to the steady state. In addition, a routine was developed to automatically detect

resonant behaviour in a fluorescence scan. GLOBES was designed to perform on a standard laptop without the need for advanced computational resources. All the calculations described in this thesis require no more than 4GB of RAM to run. The vast majority of the calculations were performed on a 2.53GHz Mac Book Pro (Late 2008) with 4GB RAM.

GLOBES was used to analyse Doppler cooling in $^{40}Ca^{+}$. Firstly, techniques to fit fluorescence scans were described, including how to deal with departures from a cold-ion model due to ion heating. Fluorescence scans were fitted and used to analyse the ion's temperature. The standard 'steady-state' approach was found to be insufficient to describe the behaviour of the ion; this is due to the presence of a low-lying D level. Two effects were identified that cause the ion's behaviour to depart from the steady-state approach. Firstly, the method assumes that, during the ion's oscillations, the scattering rate R varies linearly with the ion's velocity. The 'excursion effect' describes the modification to this behaviour when R does not vary linearly. Secondly, the steady-state method assumes that the ion's internal state is always in equilibrium over the course of its motion in the trap. The ion oscillates with a secular frequency of order 1MHz. This frequency is large enough that rate at which the beam frequencies change due to the ion's external motion is comparable to the characteristic timescale of the ion's internal dynamics; the ion is, in general, never in equilibrium. The 'dynamic effect' describes the modification to the steady-state theory when the ion's external motion is explicitly taken into account. The two theories referred to as the 'excursion method' and 'dynamic method' model the two respective effects. A further routine, the 'motional method', includes both effects; it is the most complete method of the four presented but also is the slowest to compute.

In addition, a 'minimax' optimisation method was developed to optimise the ion's temperature. This optimisation cannot be performed using a naïve minimisation routine: the temperature of the ion and the intensity of the 397nm cooling

beam are positively correlated, so that the temperature is minimal when the intensity of this beam is approximately zero (the minimum is not exactly at zero due to the heating rate of the trap, amongst other factors) and therefore the fluorescence is also approximately zero. A low intensity 397nm beam is undesirable as, the lower the intensity, the longer it takes a hot ion to reach equilibrium. The minimax optimisation process minimises the ion's temperature while simultaneously maximising the ion's fluorescence, ensuring that outputs of the routine show a cold ion with appreciable fluorescence and cooling rate. The tradeoff between the temperature and the fluorescence can be altered by adjusting an input parameter.

GLOBES was used to analyse a complicated set of data taken in a five-level, four-laser system in $^{40}Ca^{+}$ where one laser was significantly more powerful than the others. This data was a function of two laser detunings. An efficient method of fitting large data sets was described. This method makes use of two tiers of selection from a set of initial conditions sampled at random over the region of parameter space defined by the uncertainty in the experimental parameters: the first tier involves a single function evaluation and the second tier a single iteration of the local fitting process. This method is speeded up further by making use of 'jitter', a routine that continually randomly selects and fits small subsets of the whole dataset. In addition, an extension of the resonance detection scheme for data that is a function of two variables was detailed. The features seen in the data were explored using a general method for treating resonances in a pump/probe was developed for three, four and five state systems and extended to treat systems with Zeeman states. A new data run was then described. This run consisted of a series of six two-dimensional scans. The parameters describing these scans were chosen directly to test aspects of the five-level resonance model; for each scan, the position of selected resonant effects making up the entire structure are different. This data was analysed and the changes to the

resonance structure in each fit were a very good match to the predictions of the five-level model.

Finally, work relating to a magnetic-field-insensitive 'atomic clock' transition in $^{43}Ca^{+}$ was described. GLOBES was used to locate all transitions in the five lowest energy levels that each, at a particular critical field, had a zero value of the gradient of the transition frequency with respect to field strength. The critical field, transition frequency and second order dependence of the nine atomic clock transitions in the ground state were listed. Results of a minimax optimisation process for a three-level cooling scheme in this isotope were presented and, in a similar manner to the results from $^{40}Ca^{+}$ described previously, it was shown that it is possible to cool below the Doppler limit for a population in the upper state of up to $\sim 4\%$. These good results were explained in terms of a 'stretch state' cooling scheme that makes use of one of two dark resonances that involve states at the end of the M manifolds of the ground state. These dark resonances are 'perfect' in that there are no routes out of the resonance from the two lower energy states. As a result, population is trapped in these states and the resonance is especially strong. These resonances are very effective ones to use for dark resonance cooling. An improvement to this perfect dark resonance cooling was proposed that makes use of both of the $P_{1/2}$ manifolds instead of one. This scheme avoids the problem of the two 397nm sidebands pumping population out of the perfect dark resonance. This improved scheme was tested experimentally before the development of the dynamic method. The predictions of the steady-state method are not an adequate description of the ion's behaviour and, as a result, no improvement was observed. Experiments based on the predictions of the dynamic method would be expected to show improvements over the single-manifold approach. The results of these schemes were compared to a five-level scheme that removes multiple-level coherent effects from the cooling transition. This scheme produced temperatures an order of magnitude higher

than dark resonance cooling; the broad slopes of the fluorescence profile will not be a steep as the sides of a dark resonance.

A method of fitting very detailed fluorescence scans produced by an ion with a complicated internal structure was discussed. A series of scans taken with increasing beam power and increasing mean ion temperature over the scan were shown. Scans with moderate beams powers are the easiest to fit. The number of 'heating regions' which cannot be fitted increases with increasing beam power (although there are also large regions of heating for very low powers of the cooling beam). A technique for estimating the temperature of an ion from a fluorescence scan alone was described. This method is useful for scans which show a high density of heating regions. A high-resolution scan of a perfect dark resonance was then analysed and the results of the dynamic method compared favourably with a direct temperature measurement using sideband thermometry. Both the theory and measurement show that the temperature of the ion is below the Doppler limit. The variation in the measured temperature as a function of the beam powers and detuning was modelled using the dynamic method; the agreement between data and theory is impressive.

Optical initialisation and readout of an ion at the critical field was analysed. The time-dependent initialisation problem can be solved very efficiently by using SILVER IMPER. The results of a routine modelling a seven step readout process were compared to the results of a rate equations approach and found to be very similar. This is as expected due to the lack of any coherences connecting multiple beams during this process.

This thesis has demonstrated that it is possible to Doppler cool an atomic clock transition in $^{43}Ca^{+}$ at a critical field of 146.09 gauss to temperatures lower than the Doppler limit. This is sufficient for sideband cooling to operate, and this isotope can be cooled to the ground state using that technique. Very impressive results have been obtained from subsequent experiments on this qubit. All of

these results have been achieved in the same set of apparatus. The transition has demonstrated a coherence time of $50(10)$s, the largest ever recorded for a single qubit [HAB+14]. A combined error on state preparation and readout of $6.8(6) \times 10^{-4}$ has been measured. Gate errors for a single qubit of $1.0(3) \times 10^{-6}$ and for two qubits of $< 5.0 \times 10^{-2}$ have been observed [Sep16]. All of these results are at least an order of magnitude better than previous work. A high fluorescence rate of 25000 counts per second has also been observed. The optical Bloch equations routines used in the study of this isotope are entirely general and could be used to study a wide variety of problems. It is planned to release these routines as open source software to the community.

7.1 Further work

Further experimental work is under way to improve the recorded fidelities of the atomic clock transition described in this thesis. Firstly, a new iteration of a macroscopic ion trap is currently being constructed. The apparatus contains field coils that are able to produce a magnetic field of 287.78 gauss, and it is planned to use this trap to perform experiments on the 287.78 gauss atomic clock transition [Thi15]. This qubit will be cooled using the stretch state cooling technique; experimental parameters will be similar to those that produce the scan shown in figure 6.37. This qubit is on a π transition and will therefore couple slightly more strongly to the magnetic field gradients used in the experiment than the σ transition at 146.09 gauss [Har13]. It may therefore be possible to observe higher fidelity operations using this qubit. In addition, it is planned to use this trap for experiments in coupling photons with a trapped ion, a step towards realising a networked cell architecture for scalable ion trap quantum processing. This trap is currently being hard-baked, and it is expected that ions will be trapped in the system shortly.

Secondly, plans are being drafted for an ion trap that will by cryogenically

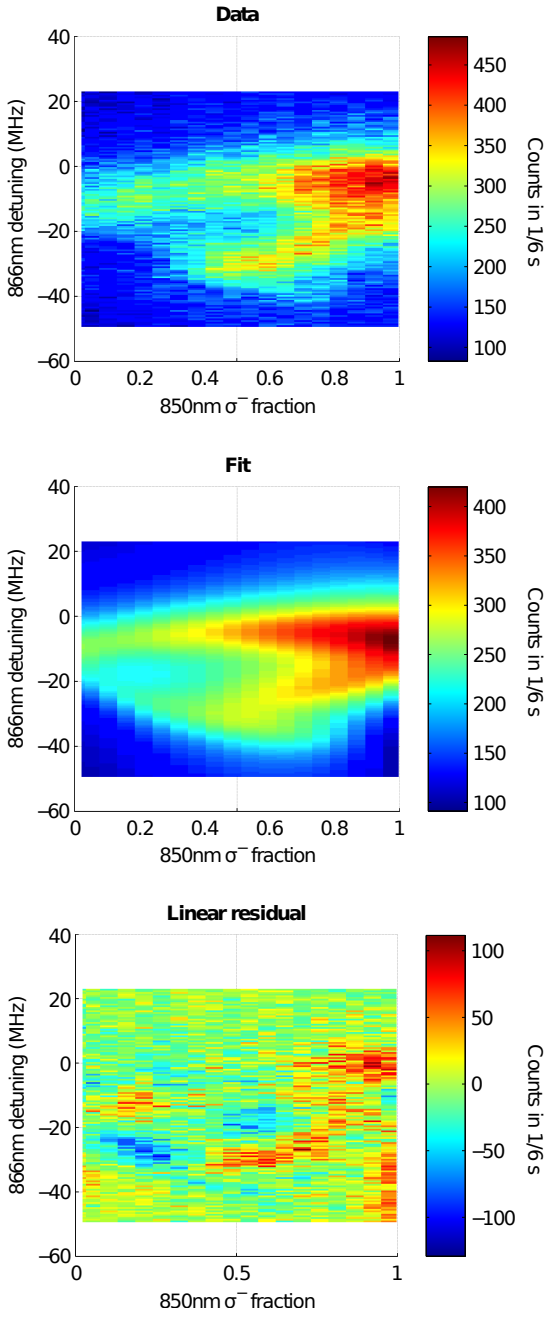

Figure 7.1: A preliminary analysis of the dependence of the fluorescence observed from the five-level system examined in chapter 5 as a function of the polarisation of the pump beam. The upper plot shows the data. The 850nm beam nominally contains no π polarisation component. The polarisation of the beam was altered by adjusting a waveplate in the beam path, changing the relative proportion of the two σ components. The middle plot shows the output of the first stage of the fitting process. The lower plot shows the residuals. Even at this early stage of fitting, the agreement between the data and theory is significant. Note that the hyperbola separation is at its greatest for equal sigma polarisations. This polarisation arrangement was used for the majority of the data taken in this system.

cooled. Cryogenic traps generally have significantly lower anomalous heating rates than uncooled traps (see [HCW+12] for a review of heating rates in cooled and uncooled traps, and [LGA+08] and [CS14] for examples of cryogenically cooled ion traps). A lower heating rate allows for greater control over the ion's motional state and thus higher gate fidelities. By combing a cryogenic trap with an atomic clock qubit, very low qubit operation fidelities could in principle be observed.

In addition, a new set of data has been taken studying further aspects of the four-laser, five-level pump/probe system in $^{40}Ca^{+}$ that was examined in this thesis. This data set consists of over 20 two-dimensional scans. In particular, a series of scans was taken to analyse the dependence of the splitting of the five-level hyperbola as a function of the polarisation of the pump beam. Preliminary results are very encouraging, and already show a good agreement between theory and experiment. Figure 7.1 shows an example of a preliminary fit of a two-dimensional scan that examines this behaviour. For this scan, the fraction of π polarisation in the pump beam was fixed at zero and the relative proportion of the two σ polarisations was varied. The x axis of the plot shows the fraction of the σ^- polarisation and the y axis 866nm detuning. The scan focuses on a small region around one arm of the hyperbola. This plot clearly shows that the hyperbola splits in two when there is a mixture of the two polarisations, but that the hyperbola is not split when the polarisation is pure σ^+ or σ^-. The maximal splitting occurs for an equal mix of the two σ polarisations. The agreement between the data and the fit is at the 20% level, and it is expected that further fitting will improve this agreement by at least a factor of two. The full set of scans is currently being fitted and analysed.

Bibliography

[ABG+94] F. Arbes, M. Benzing, T. Gudjons, F. Kurth, and G. Werth. Precise determination of the ground state hyperfine structure splitting of Ca-43-ii. *Zeitschrift Für Physik D*, 31(1–2):27–30, 1994.

[AGH+11] D. T. C. Allcock, L. Guidoni, T. P. Harty, C. J. Ballance, M. G. Blain, A. M. Steane, and D. M. Lucas. Reduction of heating rate in a microfabricated ion trap by pulsed-laser cleaning. *New Journal of Physics*, 13:123023, 2011.

[AHB+13] D. T. C. Allcock, T. P. Harty, C. J. Ballance, B. C. Keitch, N. M. Linke, D. N. Stacey, and D. M. Lucas. A microfabricated ion trap with integrated microwave circuitry. *Applied Physics Letters*, 102:044103, 2013.

[AHS+16] D. T. C Allcock, T. P. Harty, M. A. Sepiol, H. A. Janacek, C. J. Ballance, A. M. Steane, D. M. Lucas, and D. N. Stacey. Dark-resonance cooling and high fluorescence in Ca-43 ions at intermediate field. *New Journal of Physics*, 18:023043, 2016.

[All11] D. T. C. Allcock. *Surface-Electrode Ion Traps for Scalable Quantum Computing*. PhD thesis, University of Oxford, 2011.

[AT55] S. H. Autler and C. H. Townes. Stark effect in rapidly varying fields. *Physical Review*, 100(2):703–722, 1955.

[Bal14] C. J. Balance. *High-Fidelity Quantum Logic in Ca^+*. PhD thesis, University of Oxford, 2014.

[Ban06] Y. B. Band. *Light and Matter: Electromagnetism, Optics, Spectroscopy and Lasers*. Wiley, 2006.

[BCS+04] M. D. Barrett, J. Chiaverini, T. Schaetz, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. Deterministic quantum teleportation of atomic qubits. *Nature*, 429(6993):737–739, 2004.

[BDL+00] P. A. Barton, C. J. S. Donald, D. M. Lucas, D. A. Stevens, A. M. Steane, and D. N. Stacey. Measurement of the lifetime of the $3d^2$ $D_{5/2}$ state in $^{40}Ca^+$. *Physical Review A*, 62(3):032503, 2000.

[BHI+91] J. J. Bollinger, D. J. Heinzen, W. M. Itano, S. L. Gilbert, and D. J. Wineland. A 303-MHz frequency standard based on trapped Be^+ ions. *Institute of Electrical and Electronics Engineers: Transactions on Instrumentation and Measurement*, 40(2):126–128, 1991.

[BHL+15] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. Laser-driven quantum logic gates with precision beyond the fault-tolerant threshold. arXiv: 1512.04600, 2015.

[BHR+04] R. Blatt, H. Häffner, F. Roos, C. Becher, and F. Schmidt-Kaler. Ion trap quantum computing with Ca^+ ions. *Quantum Information Processing*, 3(1):61 – 73, 2004.

[BKR+07a] J. Benhelm, G. Kirchmair, U. Rapol, T. Körber, C. F. Roos, and R. Blatt. Erratum: Measurement of the hyperfine structure of the $S_{1/2}$-$D_{5/2}$ transition in $^{43}Ca^{+}$. *Physical Review A*, 75(4):049901, 2007.

[BKR+07b] J. Benhelm, G. Kirchmair, U. Rapol, T. Körber, C. F. Roos, and R. Blatt. Measurement of the hyperfine structure of the $S_{1/2}$-$D_{5/2}$ transition in $^{43}Ca^{+}$. *Physical Review A*, 75(3):032506, 2007.

[BS96] P. Bogacki and L.F. Shampine. An efficient Runge-Kutta (4,5) pair. *Computers and Mathematics with Applications*, 32(7):15–28, 1996.

[BW08] R. Blatt and D. J. Wineland. Entangled states of trapped atomic ions. *Nature*, 453(7198):1008–1015, 2008.

[CBK+09] M. Chwalla, J. Benhelm, K. Kim, G. Kirchmair, T. Monz, M. Riebe, P. Schindler, A. S. Villar, W. Hänsel, C. F. Roos, R. Blatt, M. Abgrall, G. Santarelli, G. D. Rovera, , and Ph. Laurent. Absolute frequency measurement of the $^{40}Ca^{+}$ 4s $^{2}S_{1/2}$-3d $^{2}D_{5/2}$ clock transition. *Physical Review Letters*, 102(2):023002, 2009.

[CGC+12] J. M. Chow, J. M. Gambetta, A. D. Córcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. R. Rozen, M. B. Ketchen, and M. Steffen. Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits. *Physical Review Letters*, 109(6):060501, 2012.

[Chi96] R. Y. Chiao, editor. *Amazing Light: A Volume Dedicated To Charles Hard Townes On His 80th Birthday*. Springer, 1996.

[CK90] J. R. Cash and A. H. Karp. A variable order Runge-Kutta method for initial value problems with rapidly varying right-hand sides. *Association for Computing Machinery: Transactions on Mathematical Software*, 16(3):201–222, 1990.

[CLH+13] D. P. L. Aude Craik, N. M. Linke, T. P. Harty, C. J. Ballance, D. M. Lucas, A. M. Steane, and D. T. C. Allcock. Microwave control electrodes for scalable, parallel, single-qubit operations in a surface-electrode ion trap. *Applied Physics B*, 114(1):3–10, 2013.

[CM69] E. Cuthill and J. McKee. Reducing the bandwidth of sparse symmetric matrices. *Proceedings of the 24th National Conference of the Association for Computing Machinery*, pages 157–172, 1969.

[CM11] E. Cuthill and J. McKee. Computing the action of the matrix exponential, with an application to exponential integrators. *Society for Industrial and Applied Mathematics Journal on Scientific Computing*, 33(2):488–511, 2011.

[Cor05] A. Corney. *Atomic And Laser Spectroscopy*. Oxford University Press, 2nd edition, 2005.

[Cra16] D. P. L. Aude Craik. *Near field microwave addressing of trapped-ion qubits*. PhD thesis, University of Oxford, 2016.

[CS14] J. Chiaverini and J. M. Sage. Insensitivity of the rate of ion motional heating to trap-electrode material over a large temperature range. *Physical Review A*, 89(1):012318, 2014.

[CZ95] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74(20):4091, 1995.

[Deo06] N. Deo. *Graph Theory with Applications to Engineering and Computer Science*. Prentice-Hall of India, reprint of 1974 edition, 2006.

[Deu85] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London A*, 400:97–117, 1985.

[Dit52] R. W. Ditchburn. *Light*. Blackie and Son, 1952. Available as part of the Internet Archive at `https://archive.org/details/Light_824`.

[DiV00] D. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9–11):771–783, 2000.

[DJ92] D. Deutsch and R. Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London A*, 439:553–558, 1992.

[Don00] C. Donald. *Development of an ion trap quantum information processor*. PhD thesis, University of Oxford, 2000.

[DP80] J. R. Dormand and P. J. Prince. A family of embedded Runge-Kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19–26, 1980.

[DS13] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: An outlook. *Science*, 339(6124):1169–1174, 2013.

[ER56] B. Edlén and P. Risberg. The spectrum of singly-ionized calcium, Ca-II. *Arkiv För Fysik*, 10(6):553–566, 1956.

[Fey82] R. P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6–7):467–488, 1982.

[FIM05] M. Fleischhauer, A. Imamoglu, and J. P. Marangos. Electromagnetically induced transparency: Optics in coherent media. *Proceedings of the International School of Physics "Enrico Fermi"*, 77(2):633–673, 2005.

[FN91] R. W. Freund and N. M. Nachtigal. QMR: a quasi-minimal residual method for non-Hermitian linear systems. *Numerische Mathematik*, 60(1):315–339, 1991.

[Foo05] C. J. Foot. *Atomic Physics*. Oxford University Press, 2005.

[Geo71] J. A. George. *Computer Implementation of the Finite Element Method*. PhD thesis, Stanford University, 1971.

[GKZ+08] R. Gerritsma, G. Kirchmair, F. Zähringer, J. Benhelm, R. Blatt, and C. F. Roos. Precision measurement of the branching fractions of the 4p $^2P_{3/2}$ decay of Ca II. *European Physical Journal D*, 50(1):13–19, 2008.

[Gou12] R. Gould. *Graph Theory*. Dover Publications, 2012.

[Gro97] L. K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical Review Letters*, 79(2):325–328, 1997.

[GWW+06] Y. Gu, L. Wang, K. Wang, Cheng Yang, and Qihuang Gong. Coherent population trapping and electromagnetically induced transparency in a five-level M-type atom. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 39(3):463–471, 2006.

[HAB+14] 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. *Physical Review Letters*, 113(22):220501, 2014.

[Har12] M. Harlander. *Architecture for a scalable ion-trap quantum computer*. PhD thesis, Universität Innsbruck, 2012.

[Har13] T. P. Harty. *High-fidelity microwave-driven quantum logic in intermediate-field* $^{43}Ca^{+}$. PhD thesis, University of Oxford, 2013.

[HCW+12] D. A. Hite, Y. Colombe, A. C. Wilson, K. R. Brown, U. Warring, R. Jördens, J. D. Jost, K. S. McKay, D. P. Pappas, D. Leibfried, and D. J. Wineland. 100-fold reduction of electric-field noise in an ion trap cleaned with *In Situ* Argon-ion-beam bombardment. *Physical Review Letters*, 109(10):103001, 2012.

[HHB08] C. F. Roos H. Häffner and R. Blatt. Quantum computing with trapped ions. *Physics Reports*, 469(4):155–203, 2008.

[HML+06] J. P. Home, M. J. McDonnell, D. M. Lucas, G. Imreh, B. C. Keitch, D. J. Szwer, N. R. Thomas, S. C. Webster, D. N. Stacey, and A. M. Steane. Deterministic entanglement and tomography of ion-spin qubits. *New Journal of Physics*, 8(9):188, 2006.

[Hom06] J. Home. *Entanglement of Two Trapped-Ion Spin Qubits*. PhD thesis, University of Oxford, 2006.

[HS00] L. Hales and S.Hallgren. An improved quantum Fourier transform algorithm and applications. *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 515–525, 2000.

[HW05] H. Haken and H. C. Wolf. *The Physics of Atoms and Quanta: Introduction to Experiments and Theory*. Springer, 7th edition, 2005.

[IBBW95] W. M. Itano, J. C. Bergquist, J. J. Bollinger, and D. J. Wineland. Cooling methods in ion traps. *Physica Scripta*, T59:106–120, 1995.

[IW82] W. M. Itano and D. J. Wineland. Laser cooling of ions stored in harmonic and penning traps. *Physical Review A*, 25(1):35–54, 1982.

[Jor15] S. P. Jordan. Quantum algorithm zoo (adapted from http://arxiv.org/abs/0809.2307). `http://math.nist.gov/quantum/zoo/`, 2015. Accessed: 28-09-2015.

[JSD78] R. S Van Dyck Jr., P. B Schwinberg, and H. G. Dehmelt. Electron magnetic moment from geonium spectra. *New Frontiers in High Energy Physics*, 1978.

[KBL+05] 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 ^{2}D-level lifetimes of $^{40}Ca^{+}$. *Physical Review A*, 71(3):032504, 2005.

[LAS+12] N. M. Linke, D. T. C. Allcock, D. J. Szwer, C. J. Ballance, T. P. Harty, H. A. Janacek, D. N. Stacey, A. M. Steane, and D. M. Lucas. Background-free detection of trapped ions. *Applied Physics B*, 107(4):1175–1180, 2012.

[LBMW03] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75(1):281–324, 2003.

[LDM+03] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422(1492):412–415, 2003.

[LGA+08] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibrandt, K. R. Brown, and I. L. Chuang. Suppression of heating rates in cryogenic surface-electrode ion traps. *Physical Review Letters*, 100(1):013001, 2008.

[Lin12] N. M. Linke. *Background-free detection and mixed-species crystals in micro- and macroscopic ion-traps for scalable QIP*. PhD thesis, University of Oxford, 2012.

[LJL+10] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. OBrien. Quantum computers. *Nature*, 464(8812):45–53, 2010.

[LL15] D. Leibfried and D. Lucas. Trapping and cooling of atomic ions. *Proceedings of the International School of Physics "Enrico Fermi"*, 189(Ion Traps for Tomorrow's Applications (Trappole ioniche per le applicazioni di domani)):1–64, 2015.

[LOJ+05] C. Langer, R. Ozeri, J. D. Jost, J. Chiaverini, B. DeMarco, A. Ben-Kish, R. B. Blakestad, J. Britton, D. B. Hume, W. M. Itano, D. Leibfried, R. Reichle, T. Rosenband, T. Schaetz, P. O. Schmidt, and D. J. Wineland. Long-lived qubit memory using atomic ions. *Physical Review Letters*, 95:060502, 2005.

[Lou97] R. Loudon. *The Quantum Theory of Light*. Oxford University Press, 3rd edition, 1997.

[LQNG15] B. Li, Y. Qi, Y. Niu, and S. Gong. Superluminal optical vector solitons in a five-level M-type atomic system. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 48(6):065501, 2015.

[Mar98] J. P. Marangos. Electromagnetically induced transparency. *Journal of Modern Optics*, 45(3):471–503, 1998.

[MK85] J. J. McClelland and M. H. Kelley. Detailed look at aspects of optical pumping in sodium. *Physical Review A*, 31(6):3704, 1985.

[MK13] C. Monroe and J. Kim. Scaling the ion trap quantum processor. *Science*, 339(6124):1164–1169, 2013.

[Mos13] H. G. J. Moseley. The high-frequency spectra of the elements. *Phil. Mag.*, 27:703–324, 1913.

[MSW+08] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Physical Review Letters*, 100(20):200502, 2008.

[MTN12] P. J. Mohr, B. N. Taylor, and D. B. Newell. CODATA recommended values of the fundamental physical constants: 2010. *Reviews of Modern Physics*, 84(4):1527–1605, 2012.

[MvdS99] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer, 1999.

[NBI+98] W. Nörtershäuser, K. Blaum, K. Icker, P. Müller, A. Schmitt, K. Wendt, and B. Wiche. Isotope shifts and hyperfine structure in the 3d $^2D_j \rightarrow 4p\ ^2P_j$ transitions in calcium II. *European Physical Journal D*, 2(1):33–39, 1998.

[NFB14] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin. Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links. *Physical Review X*, 4(4):041041, 2014.

[Nov01] E. Novak. Quantum complexity of integration. *Journal of Complexity*, 17(1):2–16, 2001.

[Ols72] L. Olschewski. Determination of the nuclear magnetic moments on free ^{43}Ca-, ^{87}Sr-, ^{135}Ba-,^{137}Ba-,^{171}Yb- and ^{173}Yb-atoms by means of optical pumping. *Zeitschrift für Physik*, 249(3):205–227, 1972.

[Orf14] S. J. Orfanidis. Electromagnetic waves and antennas. `http://www.ece.rutgers.edu/~orfanidi/ewa/`, 2014.

[ORVT06] C. Ottaviani, S. Rebić, D. Vitali, and P. Tombesi. Quantum phase-gate operation based on nonlinear optics: Full quantum analysis. *Physical Review A*, 73(1):010301, 2006.

[OWC+11] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland. Microwave quantum logic gates for trapped ions. *Nature*, 476:181–184, 2011.

[Pen36] F. M. Penning. Die glimmentladung bei niedrigem druck zwischen koaxialen zylindern in einem axialen magnetfeld. *Physica*, 3(9):873–894, 1936.

[PFTV92] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1992.

[PS53] W. Paul and H. Steinwedel. Ein neues massenspektrometer ohne magnetfeld. *Zeitschrift für Naturforschung A*, 8(7):448–450, 1953.

[QZH+12] Y. Qia, F. Zhoua, T. Huanga, Y. Niua, and S. Gonga. Three-dimensional atom localization in a five-level M-type atomic system. *Journal of Modern Optics*, 59(2):1092–1099, 2012.

[RAN+02] D. Reiß, K. Abich, W. Neuhauser, Ch. Wunderlich, and P. E. Toschek. Raman cooling and heating of two trapped Ba^+ ions. *Physical Review A*, 65(5):053401, 2002.

[RHR+04] M. Riebe, H. Häffner, C. F. Roos, W. Hänsel, J. Benhelm, G. P. T. Lancaster, T. W. Körber, C. Becher, F. Schmidt-Kaler, D. F. V. James, and R. Blatt. Deterministic quantum teleportation with atoms. *Nature*, 429(6993):734–737, 2004.

[RLB96] D. Reiß, A. Lindner, and R. Blatt. Cooling of trapped multilevel ions: A numerical analysis. *Physical Review A*, 54(6):5133–5140, 1996.

[Sep16] M. Sepiol. *Microwave driven quantum logic gates with calcium ions*. PhD thesis, University of Oxford, 2016.

[SF93] G. L. G Sleijpen and D. R. Fokkema. BiCGstab(l) for linear equations involving unsymmetric matrices with complex spectrum. *Electronic Transactions on Numerical Analysis*, 1(1):11–32, 1993.

[Sho97] P.W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

[Sho11] B. W. Shore. *Manipulating Quantum Structures Using Laser Pulses*. Cambridge University Press, 2011.

[SJS16] D. N. Stacey, H. A. Janacek, and A. M. Steane. Refinements to the theory of Doppler cooling. In preparation, 2016.

[SKDB10] G. Shu, N. Kurz, M. R. Dietrich, and B. B. Blinov. Efficient fluorescence collection from trapped ions with an integrated spherical mirror. *Physical Review A*, 81(4):042321, 2010.

[SKHR+03] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt. Realization of the Cirac-Zoller controlled-NOT quantum gate. *Nature*, 422(1494):408–411, 2003.

[SKK+00] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe. Experimental entanglement of four particles. *Nature*, 404:256–259, 2000.

[SLA+04] D. N. Stacey, D. M. Lucas, D. T. C. Allcock, D. J. Szwer, and S. C. Webster. Optical Bloch equations with multiply connected states. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 41(8):085502, 2004.

[SM99] A. Sørensen and K. Mølmer. Quantum computation with ions in thermal motion. *Physical Review Letters*, 8(9):1971–1974, 1999.

[SM00] A. Sørensen and K. Mølmer. Entanglement and quantum computation with ions in thermal motion. *Physical Review A*, 62(2):022311, 2000.

[Sob96] I. I. Sobel'man. *Atomic Spectra and Radiative Transitions*. Springer, 2nd edition, 1996.

[Sor91] D. C. Sorensen. Implicit application of polynomial filters in a k-step Arnoldi method. *Society for Industrial and Applied Mathematics Journal on Matrix Analysis and Applications*, 13(1):357–385, 1991.

[SSAT98] Y. Stalgies, I. Siemers, B. Appasamy, and P. E. Toschek. Light shift and Fano resonances in a single cold ion. *Journal of the Optical Society of America B*, 15(10):2505–2514, 1998.

[Sta03] J-P. Stacey. *Stabilization and Control In a Linear Ion Trap*. PhD thesis, University of Oxford, 2003.

[Sta04] P. Staanum. *Quantum Optics with Trapped Calcium Ions*. PhD thesis, University of Aarhus, 2004.

[Ste03] A. M. Steane. Overhead and noise threshold of fault-tolerant quantum error correction. *Physical Review A*, 68(4):042322, 2003.

[Str69] V. Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13(4):354–356, 1969.

[Szw09] D. J. Szwer. *High Fidelity Readout and Protection of a $^{43}Ca^{+}$ Trapped Ion Qubit*. PhD thesis, University of Oxford, 2009.

[Thi15] K. Thirumalai. Design and construction of a photonically entangled trapped ion computing experiment. University of Oxford Technical Report, 2015.

[TI97] L. N. Trefethen and D. Bau III. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 2nd edition, 1997.

[TPR+03] G. Tommaseo, T. Pfeil, G. Revalde, G. Werth, P. Indelicato, and J. Desclaux. The g_J-factor in the ground state of Ca^{+}. *European Physical Journal D*, 25(2):113–121, 2003.

[Web05] S. C. Webster. *Raman Sideband Cooling and Coherent Manipulation of Trapped Ions*. PhD thesis, University of Oxford, 2005.

[WI79] D. J. Wineland and W. M. Itano. Laser cooling of atoms. *Physical Review A*, 20(4):1521–1520, 1979.

[WMI+98] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259–328, 1998.

[Wol96] E. Wolf, editor. *Progress in Optics: Volume 35*, chapter V: Coherent population trapping in laser spectroscopy. Elsevier: North Holland, 1996.

[Woo83] G. K. Woodgate. *Elementary Atomic Structure*. Oxford University Press, 2nd edition, 1983.

[WvdBG+08] A. L. Wolf, S. A. van den Berg, C. Gohle, E. J. Salumbides, W. Ubachs, and K. S. E. Eikema. Frequency metrology on the 4s $^2S_{1/2}$-4p $^2P_{1/2}$ transition in $^{40}Ca^{+}$ for a comparison with quasar data. *Physical Review A*, 78(3):032511, 2008.

[WvdBUE09] A. L. Wolf, S. A. van den Berg, W. Ubachs, and K. S. E. Eikema. Direct frequency comb spectroscopy of trapped ions. *Physical Review Letters*, 102(22):223901, 2009.

[YRtNAT08] J. Reader Y. Ralchenko, A. Kramida and the NIST ASD Team. NIST Atomic Spectra Database (version 3.1.5). `http://physics.nist.gov/asd`, 2008. Accessed: 03-04-2016.

Appendices

A

Connectivity in an arbitrary system

> You must be the most soft-hearted griffin in the world! Teddy bears and moons! Come on Elda.
>
> Lukin of Luteria
> *Year of the Griffin* by Diana Wynne Jones

In this appendix, we describe methods implemented in our simulations to deal with arbitrary systems of states and lasers. In particular, we mention how to take account of elements in the Bloch equations that depend on the entire chain of stimulated transitions that connect any two states, along with with methods to detect multiply-connected states caused by loops of stimulated processes.

A.1 A universal input system

These routines have been designed to automatically solve the Bloch equations for any atomic system. As such, an input system was designed to describe an arbitrary atomic system and an arbitrary system of lasers in a compact and simple way. Information about the levels themselves is given in a series of vectors. An initial ordering of the levels is defined by the user, after which the properties of the levels are listed in the same ordering in their associated vector. This information will always be fixed for a given atomic system, and a simple list is

the most efficient way of storing the information.

A different approach was used for the properties of the laser interactions. These will be adjusted on a very regular basis, so it important that the inputs are not prone to user error. Changing the wrong element in a long list of values is a distinct possibility. Instead, the laser information is contained in a series of square matrices. Each matrix describes one property of the lasers. Each element refers to an interaction between the levels defined by the horizontal and vertical coordinates. The laser interactions are symmetric between a pair of levels. In these routines, the input matrices are modified after input to be symmetric; any laser interaction need only be written in a single element. The advantage of using this approach is that the states described by any element can be immediately determined by eye, rather than in a list where this information is not immediately clear.

A.2 Analysing systems of states and beams

For the most part, dealing with the inputs to the Bloch equations is straightforward. The stimulated and spontaneous rates between pairs of states depend only on atomic parameters and properties of any laser directly connecting the states. However, the Bloch equations deal with coherences between states that can be indirectly connected by, in principle, any number of lasers. As a result, the inputs to the equations also depend on the graph of laser-driven connections. For instance, cycles in this graph need to be carefully assessed, so that the detuning and linewidth terms are expressed correctly in the equations. As an example, the states j and k may be connected by the chain of laser-driven processes:

$$|j\rangle \xrightarrow{\gamma a} |a\rangle \xrightarrow{\gamma b} |b\rangle \xrightarrow{\gamma c} \cdots \xrightarrow{\gamma z} |k\rangle \tag{A.1}$$

The value of the detuning Δ_{jk} depends in principle on the detuning param-

eter of every laser in the chain:

$$\Delta_{jk} = \Delta_{ja} + \Delta_{ab} + \Delta_{bc} + \cdots + \Delta_{zk} \tag{A.2}$$

The intermediate detunings all cancel, meaning that the overall detuning only depends of the relative energies of the states $|j\rangle$ and $|k\rangle$. This cancellation does not occur for the linewidths, however, which do depend on the linewidths of every laser in the chain:

$$\Gamma_{L,jk} = \Gamma_{L,ja} + \Gamma_{L,ab} + \Gamma_{L,bc} + \cdots + \Gamma_{L,zk} \tag{A.3}$$

This is important because any routine using the linewidths as an input needs to know which lasers contribute to the connection between *any* pair of states. We could of course input this manually, but that would quickly become time-consuming for large systems.

We therefore need a routine that can calculate which lasers are involved in the coherences between any pair of states in the system. In addition, this routine needs to deal with two situations. Firstly, it must be able to model situations where two states are not connected by any laser beams at all. Secondly, it must be able to identify loops in the system, warn the user of their presence and calculate assuming that they are consistent. The Bloch equations can only model self-consistent loops of stimulated transitions between pairs of states as discussed in section 2.1.1.

This calculation is performed by GROVYLE (General Route-finding Obtained from Vertex connectivitY incorporating Loop Examination)[1]. This routine essentially solves two standard problems in graph theory, that is, route finding and the calculation of a minimal spanning tree (see, for example, [Deo06, Gou12]). Standard routines could be used to solve both these problems, but the methods described below are simple to implement in our system. The

[1]GROVYLE is an evolution of the 'Tree Connectivity-Knowledgable Orientation' routine (TreeCKO) mentioned in earlier reports.

timescale needed to calculate the connections is so small (the longest calculation time to solve both problems and construct the linewidth and detuning matrices for 144 states is of order 1s, and this calculation only needs to be carried out once) that any improvements in the calculation time achieved by using standard routines would be negligible.

A.2.1 Detecting connected components

The first goal of the routine is to determine which states have coherent connections between them of any kind. This is analogous to a graph theory problem. We are interested in the structure of the system on the order of the levels only. Let us represent the system as a graph. Each level is a node and each laser-driven transition an edge. The problem therefore reduces to finding the connected components with a particular graph. Notably, we need to deal with such subgraphs when looking at certain cooling schemes!

We define the adjacency matrix $\mathbf{G}$ of the graph as a square matrix with a side length equal to the number of levels in the system. G_{nm} is zero if the nodes n and m are not connected by an edge, and one if the nodes are directly connected (in graph terminology, they are neighbours). In other words, the matrix maps each node onto the set of nodes connected to it by a single edge. One can show the the matrix $\mathbf{G}^d$ is a matrix whose elements give the number of paths of length d between each pair of nodes. Hence, this matrix maps each node onto the set of nodes connected to it by a path of length d edges. We consider, then, the matrix exponential of the adjacency matrix:

$$\mathbf{E} = \exp(\mathbf{G}) = \sum_{d=0}^{\infty} \left(\frac{1}{d!} \mathbf{G}^d \right) \tag{A.4}$$

A matrix element of the matrix exponential E_{nm} can only be zero if and only if every term in the series is zero. In other words, no combination of the lasers in the system can map $|n\rangle$ onto $|m\rangle$. Therefore, the two states are not connected by any coherent process. As a result, we can detect non-connected components of

the graph by calculating the matrix exponential of the adjacency matrix. Connected components will appear as blocks of nonzero elements. The matrix exponential may seem a slow method of calculation, but it is possible to calculate all the information we need efficiently by only calculating portions of the exponential, once again using the routines of Awad H. Al-Mohy and Nicholas J. Higham [CM11].

A.2.2 Detecting loops

Once the connected components of the system have been found, we look at each in turn. Before we calculate routes between states, we need to deal with the problem of cycles in these subgraphs. A cycle means that there are in principle several different routes by which two nodes are connected. The presence of a cycle means that the Bloch equations cannot in general be reduced to a simple linear form, so the core routines will not be applicable. We therefore need to detect the presence of such cycles and either stop the calculation from running when they are present or take steps to calculate using the loop if the loop is consistent.

We detect whether loops are present in the system by using an iterative process. At each iteration, we remove all 'pendant' nodes (those that are connected to only a single edge). We detect these nodes simply from rows in the adjacency matrix that only have a single nonzero value. A spanning tree has exactly $N-1$ edges due to the requirement that every node be connected to an edge. As a result, if this process is applied to a spanning tree $N-1$ times, the result must be the null graph. If the process is applied to a more general graph, any nodes that remain must necessarily be part of a loop. This can be shown most easily by noting that any connected graph can be generated by adding edges to some spanning tree, and that any edges added to a spanning tree must necessarily form a loop. As a result, if the null graph is not returned after this process, the

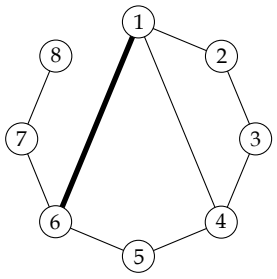

Figure A.1: A network to demonstrate the connectivity matrix $\mathbf{\Omega_d^{ij}}$. Let us consider the highlighted edge and paths of length d = 3. The element $\Omega_{3,18}^{16}$ is straightforward; there is only a single path of length three connecting 1 to 8 that passes through 16. Thus $\Omega_{3,18}^{16}$ = 1. The element $\Omega_{3,16}^{16}$ involves several routes. These are $1 \to 6 \to 1 \to 6$, which passes the edge three times, and $1 \to 2 \to 1 \to 6$, $1 \to 4 \to 1 \to 6$, $1 \to 6 \to 5 \to 6$ and $1 \to 6 \to 7 \to 6$ which all pass the edge once. Thus, $\Omega_{3,16}^{16}$ = 3+1+1+1+1 = 7.

system will contain at least one loop.

The only situation in which a loop is viable is when the detunings around the loop are self-consistent; the net detuning around the loop must be zero. In this case, there is no ambiguity in the detuning when multiple paths are taken between a pair of states. When a loop is detected, it is simple to determine whether the loop is consistent. For consistent loops, no information is lost if a single edge is removed from the loop for the purposes of this calculation only, provided that the removal of the edge does not disconnect the graph. This can be checked simply by calculating the matrix exponential. Each edge is removed in turn until an edge is found that can be removed without disconnecting the graph. We can then run the loop detection process again until either an inconsistent loop is detected and the calculation is stopped or the null graph is returned. The resulting spanning tree can be used to calculate the connections between the states.

A.2.3 Connecting the states

We can use powers of the adjacency matrix to also find the paths between pairs of states. Let us consider the matrix $\mathbf{D_{ij}}$ that is an adjacency matrix for a graph on the same nodes but containing only the edge ij. In addition, let us define the connectivity matrix $\mathbf{\Omega_d^{ij}}$ where each element $\Omega_{d,nm}^{ij}$ gives the total number of

times the edge ij is passed, summed over every path of length d connecting the two nodes n and m. So, for example:

$$\mathbf{\Omega_1^{ij}} = \mathbf{D_{ij}} \tag{A.5}$$

See figure A.1 for some other examples. The adjacency matrix maps nodes onto nodes that are directly connected to them by a single edge. The matrix $\mathbf{G\Omega_d^{ij}}$ therefore maps the connections described in the dth connectivity matrix onto a set of paths of length $d+1$. However, this doesn't include all the information about the connectivity; in fact, the matrix is not even necessarily symmetric. Not all new paths allowed by the larger path length can be generated directly from the smaller paths. Other paths exist only for connections of length $d+1$ or larger. They can be generated by mapping the edge in $\mathbf{D_{ij}}$ directly onto paths of length $d+1$ by means of the matrix $\mathbf{D_{ij}G}^d$. Note that the power of the adjacency matrix is d rather than $d+1$ as $\mathbf{D_{ij}}$ already includes a first order mapping. The $(d+1)$th connectivity matrix can therefore be generated iteratively from the dth order matrix:

$$\mathbf{\Omega_{d+1}^{ij}} = \mathbf{G\Omega_d^{ij}} + \mathbf{D_{ij}G}^d \tag{A.6}$$

As the 1st order matrix is simply $\mathbf{D_{ij}}$, we can therefore produce a direct expression for the dth order matrix:

$$\mathbf{\Omega_d^{ij}} = \sum_{n=0}^{d-1} \left(\mathbf{G}^n \mathbf{D_{ij}} \mathbf{G}^{(d-1)-n} \right) \tag{A.7}$$

For our purposes, we are only interested in paths of $d \leqslant N-1$ where N is the number of levels in the system. This is because a spanning tree must contain $N-1$ edges. The smallest value of $\Omega_{d,nm}^{ij}$ for $1 < d < N-1$ will give the smallest number of passes of the edge $i \leftrightarrow j$ in any path connecting n and m. This will either be 2 for a node connecting to itself or 1 for all other nodes. This method allows a second check that no loops are present as values larger than 2 can only be present for a network containing loops.

By summing over all the edges of the graph, we can determine which edges are passed on a path connecting any two pairs of nodes. It is then a simple matter of ordering these edges to find the complete path, if any, connecting all the nodes in the system. This information can be used to calculate the linewidths involved with any coherent process between any pair of states.

B

Polarisations in an angled beam

> Inspiration is needed in geometry, just as much as in poetry.
>
> Aleksandr Sergeyevich Pushkin

One issue we must consider prior to optimization using the Bloch equations are the possible polarisations we have access to. For reasons of practicality, it is not possible to illuminate the ion with all possible polarisations of all lasers all of the time. Therefore, we need to restrict the optimization process to those polarisations which we can implement. A limiting factor is that the angle of the beam relative to the coordinate axis defined by the **B** field must remain fixed for the lasers throughout the experiment. This angle produces a region in parameter space of possible polarisations that can be accessed, and this region must be passed as a constraint to the optimisation routine.

B.1 The polarisation ellipse

We wish to find what polarisation components can be produced by a beam propagating in a given direction. Let us consider a beam that is initially propagating along the $\hat{\mathbf{e}}_3$ direction. A general polarisation viewed in the plane perpendicular to the beam can be produced by using a combination of a quarter-wave plate

and a half-wave plate at an angle θ. The light is first prepared so that it is linearly polarised along the $\hat{\mathbf{e}}_1$ axis. The quarter wave plate converts this linear polarisation to a circular polarisation and the half-wave plate converts this to an elliptical polarisation. The **E** field of this beam after the two wave-plates is:

$$\begin{aligned} E_1 &= \cos(\theta)\cos(\omega t) \\ E_2 &= \sin(\theta)\sin(\omega t) \\ E_3 &= 0 \end{aligned} \tag{B.1}$$

The total magnitude of the **E** field has no effect on the polarisation states and is therefore not included in these expressions for simplicity. By varying θ, it is possible to reproduce any ratio of the amplitudes of the two components of the **E** field, and therefore produce any ellipse with axes aligned with the $\hat{\mathbf{e}}_1\hat{\mathbf{e}}_2$ axes.

We now introduce the beam angle γ by rotating the beam about the $\hat{\mathbf{e}}_3$ axis:

$$\begin{aligned} E_1 &= \cos(\theta)\cos(\omega t) \\ E_2 &= \sin(\theta)\cos(\gamma)\sin(\omega t) \\ E_3 &= \sin(\theta)\sin(\gamma)\sin(\omega t) \end{aligned} \tag{B.2}$$

The ellipse described by the $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_2$ components of the field has been scaled in the $\hat{\mathbf{e}}_2$ direction, producing a new ellipse.

We deduce the fraction of σ^- and σ^+ light by decomposing the elliptical polarisation into two counter-rotating circular polarisations. For the purposes of this derivation, I refer to the fraction of the intensity that is σ^- polarised and rotates anticlockwise as x and the fraction that is σ^+ polarised as y. I refer to the fraction of π light as z. We write the **E** field in terms of these fractions:

$$\begin{aligned} E_1 &= \tfrac{\sqrt{2}}{2}\left(\sqrt{x}+\sqrt{y}\right)\cos(\omega t) \\ E_2 &= \tfrac{\sqrt{2}}{2}\left(\sqrt{x}-\sqrt{y}\right)\sin(\omega t) \\ E_3 &= \sqrt{z}\sin(\omega t) \end{aligned} \tag{B.3}$$

The polarisation fractions appear as square roots in these expressions as the fractions refer to the beam intensity rather than the electric field. Using (B.2) and

(B.3), we write x, y and z in terms of θ and γ:

$$x = \frac{1}{2}\left(\cos\left(\theta\right) + \sin\left(\theta\right)\sin\left(\gamma\right)\right)^2 \tag{B.4}$$

$$y = \frac{1}{2}\left(\cos\left(\theta\right) - \sin\left(\theta\right)\sin\left(\gamma\right)\right)^2 \tag{B.5}$$

$$z = \sin^2(\theta)\sin^2(\gamma) \tag{B.6}$$

Equations (B.4) and (B.5) are parametric equations for another ellipse. However, this 'polarisation' ellipse is not in canonical form: the two axes are at an angle $\pi/4$ to the $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_2$ axes. The possible polarisations that can be achieved with this arrangement of wave-plates lie on the boundary of this ellipse.

The ellipse expressed in implicit form is:

$$(x-y)^2\sec^2\left(\gamma\right) + \left\{1 + 2(x+y-1)\mathrm{cosec}^2(\gamma)\right\}^2 = 1 \tag{B.7}$$

Either of the parametric of implicit equations may be used as constraints in any optimisation routines that float the polarisations of beams kept at a fixed angle. In my routines, I make use of the parametric equations. These equations are more efficient at tackling the more general problem for which the possible polarisations that can be achieved are bounded by the ellipse rather than lying on its edge. The parametric equations can more easily be adapted to force the polarisation state to lie within the boundary of the ellipse. See section B.2.

The polarisation ellipse is fully defined by three variables that describe the length of the two axes and the position of the centre. I define a and b as the length of the axes with gradient -1 and $+1$ respectively. The ellipse is symmetric about the line $x = y$ and therefore its centre coordinate (X_0, X_0) is described by only a single variable. These three variables scale with γ as:

$$a = \frac{\sqrt{2}}{2}\cos\left(\gamma\right) \tag{B.8}$$

$$b = \frac{\sqrt{2}}{4}\sin^2(\gamma) \tag{B.9}$$

$$X_0 = \frac{1}{4}(1 + \cos^2(\gamma) \tag{B.10}$$

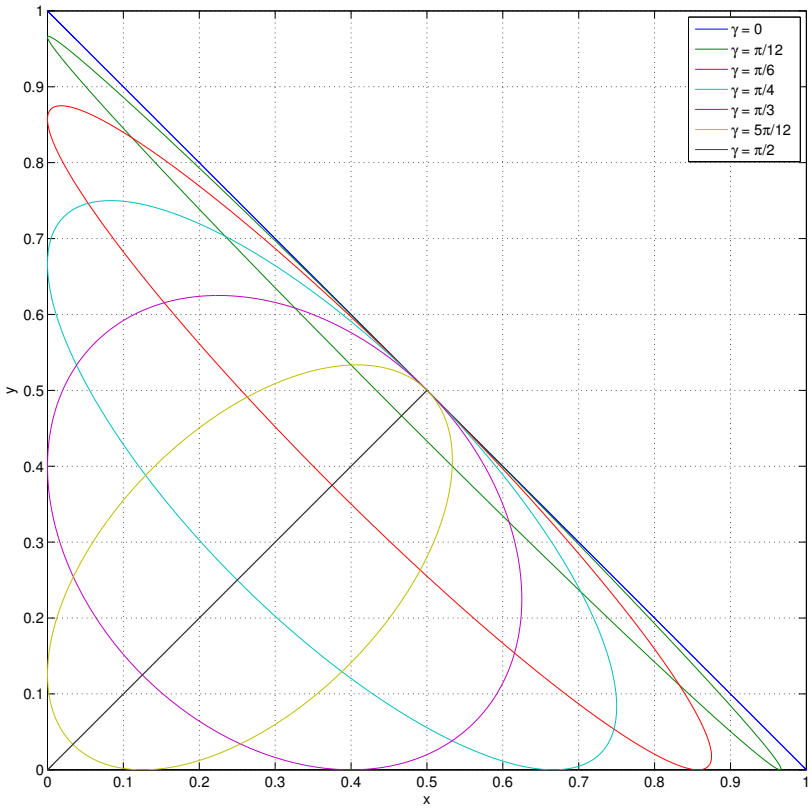

Figure B.1: A demonstration of the shape and size of the polarisation ellipse for a selection of values of the beam angle γ. The ellipses vary between the extremes of a line of gradient -1 for $\gamma = 0$ to a line of gradient $+1$ for $\gamma = \pi/2$. Note that every ellipse shares the common point [0.5, 0.5]; regardless of the beam angle, it is always possible to produce a polarisation state that comprises only of equal σ^- and σ^+ fractions.

Figure B.1 shows the polarisation ellipse for several values of γ between 0 and $\pi/2$.

The value of the wave plate angle required to produce a given polarisation state can be found by solving (B.4) and (B.5) for θ. If a given state $[x_0, y_0]$ is known to be on the circumference of the ellipse, the angle can be determined simply from:

$$\theta = \cos^{-1}\left(\frac{\sqrt{2}}{2}(\sqrt{x_0} + \sqrt{y_0})\right) \tag{B.11}$$

Otherwise, θ can be determined from either x_0 or y_0:

$$\theta = \sin^{-1}\left(\pm \sqrt{\frac{2x_0}{1+\cos^2(\gamma)}} \right) - \mathrm{atan2}(1, \cos(\gamma)) \tag{B.12}$$

$$\theta = \sin^{-1}\left(\pm \sqrt{\frac{2y_0}{1+\cos^2(\gamma)}} \right) - \mathrm{atan2}(1, -\cos(\gamma)) \tag{B.13}$$

This ensures that the polarisation state produced lies on the circumference of the ellipse. There are two values of θ that produce a given value of x_0 or y_0. These two solutions have different fractions of the undetermined sigma polarisation and z unless x_0 or y_0 are zero. In that situation, the two solutions are coincident.

B.1.1 Optimal polarisations for cooling at $\gamma = \pi/4$

In our experiments, our beams are at angle $\gamma = \pi/4$ with the **B** field. For our cooling schemes, we require the largest possible fraction of one sigma polarisation with the minimum possible fraction of the other. The optimal position for σ^- cooling requires $\theta = 0.30\pi$:

$$[x, y] = [2/3, 0] \tag{B.14}$$

The optimal position for σ^+ cooling requires $\theta = 0.70\pi$:

$$[x, y] = [0, 2/3] \tag{B.15}$$

This allows optimisation of the temperature of $^{43}Ca^+$ to occur without the need to explicitly optimise the polarisations. Optimising using a constraining ellipse confirms that optimal cooling occurs for these polarisation fractions.

B.2 Polarisations inside the ellipse

Some applications may require polarisations that cannot be realised using the scheme described above. For example, the beam angle might be fixed at a particular value due to other experimental constraints, limiting the range of polarisation states that can be implemented to the boundary of a single ellipse.

However, for a fixed beam angle, it is possible to make a modification to the scheme that increases the range of potential polarisation states. We do this by rotating the angle of the *quarter-wave* plate. We rotate the wave plate by angle ψ at the very first step of the process, so that **E** field of the beam with respect to an axis with $\hat{\mathbf{e}}_3$ parallel to the beam direction is:

$$\begin{aligned} E_1 &= \cos(\theta)\cos(\omega t)\cos(\psi) - \sin(\theta)\sin(\omega t)\sin(\psi) \\ E_2 &= \cos(\theta)\cos(\omega t)\cos(\psi) + \sin(\theta)\sin(\omega t)\cos(\psi) \\ E_3 &= 0 \end{aligned} \tag{B.16}$$

This is equivalent to expression B.1 for the original scheme. We now rotate by angle γ about the $\hat{\mathbf{e}}_3$ axis:

$$E_1 = \cos(\theta)\cos(\omega t)\cos(\psi) - \sin(\theta)\sin(\omega t)\sin(\psi) \tag{B.17}$$

$$E_2 = (\cos(\theta)\cos(\omega t)\cos(\psi) + \sin(\theta)\sin(\omega t)\cos(\psi))\cos(\gamma) \tag{B.18}$$

$$E_3 = (\cos(\theta)\cos(\omega t)\cos(\psi) + \sin(\theta)\sin(\omega t)\cos(\psi))\sin(\gamma) \tag{B.19}$$

Equations (B.17) and (B.18) describe an ellipse in the $\hat{\mathbf{e}}_1\hat{\mathbf{e}}_2$ plane. However, the ellipse is at an angle to these axes. In order to decompose the behaviour into two counter-rotating components, we need to rotate the ellipse so that it is coincident with the coordinate axes.

The angle ϕ needed to rotate the ellipse is complicated to determine analytically. For example, consider a straight line that is coincident with one axis of a rotated ellipse. When the line and ellipse are scaled in one coordinate direction (this is the effect of the rotation by γ), this line will *no longer be coincident* with the axis of the ellipse.

My routines solve this problem numerically. ϕ is determined by calculating the point on the ellipse that is furthest from its centre at the origin. The line between the origin and the farthest point defines the rotation angle. The ellipse is then rotated and the counter-rotating components determined using the magni-

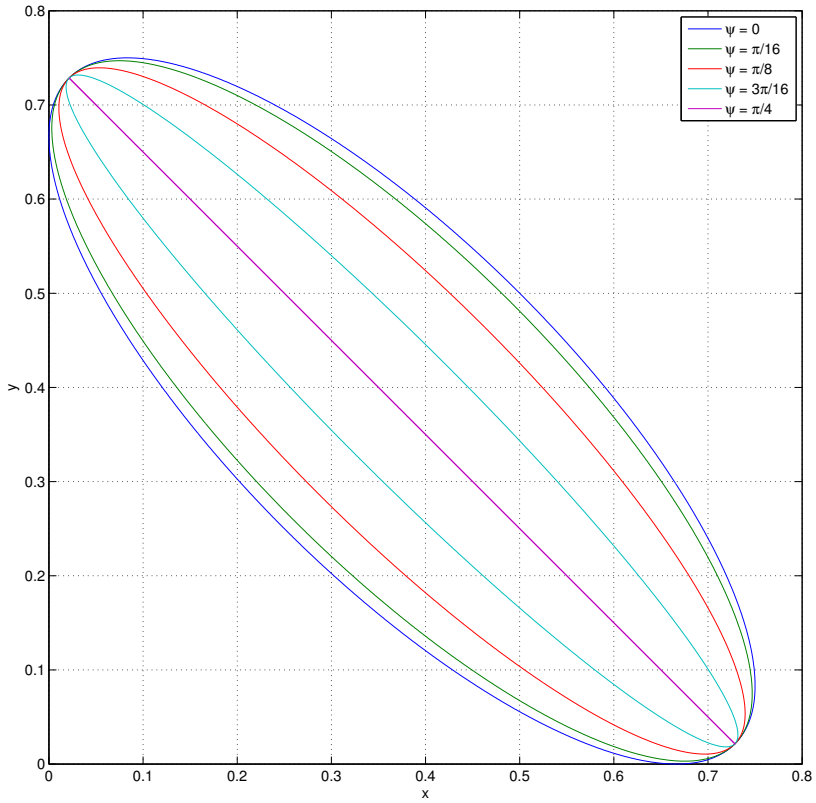

Figure B.2: A demonstration of the modification to the polarisation ellipse when the quarter-wave plate is rotated by an angle ψ. The 'standard' ellipse with $\psi = 0$ is shown in blue. All the other ellipses lie within the circumference of this standard ellipse. As the angle increases from 0 to $\pi/4$, the axis of the polarisation ellipse with gradient $+1$ decreases in length. By altering ψ, any polarisation state inside the standard ellipse can be realised.

tudes of the E_1 and E_2 components:

$$\begin{aligned} E_1 &= |E_1|\sin(\omega t + \Phi) \\ E_2 &= |E_2|\cos(\omega t + \Phi) \end{aligned} \tag{B.20}$$

We define:

$$\begin{aligned} \frac{\sqrt{2}}{2}\left(\sqrt{x} + \sqrt{y}\right) &= |E_1| \\ \frac{\sqrt{2}}{2}\left(\sqrt{x} + \sqrt{y}\right) &= |E_2| \end{aligned} \tag{B.21}$$

so that:

$$\begin{aligned} x &= \tfrac{1}{2}\left(|E_1| + |E_2|\right)^2 \\ y &= \tfrac{1}{2}\left(|E_1| - |E_2|\right)^2 \end{aligned} \tag{B.22}$$

The polarisations trace out an ellipse that is at angle $\pi/4$ to the coordinate axes.

Figure B.2 shows the modified polarisation ellipse for various values of ψ for a fixed beam angle $\gamma = \pi/4$. The modified ellipses all fall inside the boundary of the 'standard' ellipse described by $\psi = 0$ (i.e. the polarisation ellipse where the quarter-wave plate is not rotated). Rotating the quarter-wave plate allows access to polarisation states inside the boundary of the standard ellipse.

Each point on the circumference of the standard ellipse is defined by a value of θ. Rotating the wave plate moves each point on the circumference inside the ellipse along lines of constant θ with gradient $+1$. By choosing a suitable value of ψ and θ, any polarisation state $[x_0, y_0]$ bounded by the standard ellipse can be realised. The value of θ is calculated by determining the intersection points of the standard ellipse defined by (B.4) and (B.5) with the line of constant θ crossing the required polarisation state. This line is:

$$y = x + (y_0 - x_0) \tag{B.23}$$

The value of θ for a given polarisation state inside the ellipse is:

$$\theta = \frac{1}{2}\sin^{-1}\left(\frac{1}{2}(\pi \mp \pi) \pm \frac{x_0 - y_0}{\sin(\gamma)}\right) \tag{B.24}$$

The upper and lower signs respectively describe states $[x_0, y_0]$ that lie above and below the axis of the ellipse with gradient -1. The value of ψ is determined most easily by numerical means.

C

The temperature and natural frequency of a two-state system

> An atom blaster is a good weapon, but it can point both ways.
>
> Salvor Hardin
> *Foundation* by Isaac Asimov

This appendix contains analytic formulae for various properties of a two-state system calculated using the dynamic method. The amplitude of the in-phase oscillations of the higher-energy state is given by:

$$U_2 = -(8A\delta\Delta(2A^2+4A\Gamma+2\Gamma^2+4\omega^2)\Omega^4+(8A\delta(4A^2+4A\Gamma+4\omega^2)\Delta^3+8A\delta(A^4+3A^3\Gamma+3A^2\Gamma^2-3A^2\omega^2+A\Gamma^3-6A\Gamma\omega^2-3\Gamma^2\omega^2-4\omega^4)\Delta)\Omega^2)/((A^3+2A^2\Gamma+4A\Delta^2+A\Gamma^2+2A\Omega^2+2\Gamma\Omega^2)(\Gamma^4\omega^2+A^2\Gamma^4+16A\Delta^2\Gamma\omega^2+4A^5\Gamma+8A^4\Delta^2+6A^4\Gamma^2+9A^4\omega^2+4A^3\Gamma^3+16A^2\Delta^4+24A^2\omega^4+16\Delta^4\omega^2-32\Delta^2\omega^4+8\Gamma^2\omega^4+A^6+16\omega^6+16A^3\Delta^2\Gamma+20A^3\Gamma\omega^2+8A^2\Delta^2\Gamma^2-24A^2\Delta^2\omega^2+14A^2\Gamma^2\omega^2+4A\Gamma^3\omega^2+16A\Gamma\omega^4+8\Delta^2\Gamma^2\omega^2+12A^2\Gamma^2\Omega^2+12A^3\Gamma\Omega^2+8A^2\Omega^2\omega^2+32\Delta^2\Omega^2\omega^2+16A^2\Delta^2\Omega^2-8\Gamma^2\Omega^2\omega^2+8A\Gamma\Omega^4+4A\Gamma^3\Omega^2+16A\Delta^2\Gamma\Omega^2+4A^4\Omega^2+16\Omega^4\omega^2-32\Omega^2\omega^4+4A^2\Omega^4+4\Gamma^2\Omega^4)) \quad \text{(C.1)}$$

where A is the Einstein A-coefficient, δ is the assumed amplitude of the Doppler shift of the beam, Γ is the linewidth, ω is the trap frequency, Ω is the Rabi fre-

quency and Δ is the detuning. All quantities are in angular units. The amplitude of the in-quadrature oscillations of the higher-energy state is:

$$V_2 = (16A\delta\Delta\omega(A+\Gamma)\Omega^4+16A\delta\Delta\omega(2A^3+9A^2\Gamma/2+3A\Gamma^2+2A\omega^2+2\Delta^2\Gamma+\Gamma^3/2)\Omega^2)/((A^3+2A^2\Gamma+4A\Delta^2+A\Gamma^2+2A\Omega^2+2\Gamma\Omega^2)(\Gamma^4\omega^2+A^2\Gamma^4+16A\Delta^2\Gamma\omega^2+4A^5\Gamma+8A^4\Delta^2+6A^4\Gamma^2+9A^4\omega^2+4A^3\Gamma^3+16A^2\Delta^4+24A^2\omega^4+16\Delta^4\omega^2-32\Delta^2\omega^4+8\Gamma^2\omega^4+A^6+16\omega^6+16A^3\Delta^2\Gamma+20A^3\Gamma\omega^2+8A^2\Delta^2\Gamma^2-24A^2\Delta^2\omega^2+14A^2\Gamma^2\omega^2+4A\Gamma^3\omega^2+16A\Gamma\omega^4+8\Delta^2\Gamma^2\omega^2+12A^2\Gamma^2\Omega^2+12A^3\Gamma\Omega^2+8A^2\Omega^2\omega^2+32\Delta^2\Omega^2\omega^2+16A^2\Delta^2\Omega^2-8\Gamma^2\Omega^2\omega^2+8A\Gamma\Omega^4+4A\Gamma^3\Omega^2+16A\Delta^2\Gamma\Omega^2+4A^4\Omega^2+16\Omega^4\omega^2-32\Omega^2\omega^4+4A^2\Omega^4+4\Gamma^2\Omega^4)) \quad \text{(C.2)}$$

The fluorescence profile is described by:

$$R = \eta A\frac{A(A+\Gamma)\Omega^2}{A(A+\Gamma)^2+(A+\Gamma)\Omega^2+4A\Delta^2} = \eta A P_2 \quad \text{(C.3)}$$

where η is the collection efficiency and P_2 is the population of the upper state. The temperature of the ion is:

$$T = -\frac{\hbar}{k_B A}(A(A+\Gamma)\Omega^2)((A^3+2A^2\Gamma+4A\Delta^2+A\Gamma^2+2A\Omega^2+2\Gamma\Omega^2)(\Gamma^4\omega^2+A^2\Gamma^4+16A\Delta^2\Gamma\omega^2+4A^5\Gamma+8A^4\Delta^2+6A^4\Gamma^2+9A^4\omega^2+4A^3\Gamma^3+16A^2\Delta^4+24A^2\omega^4+16\Delta^4\omega^2-32\Delta^2\omega^4+8\Gamma^2\omega^4+A^6+16\omega^6+16A^3\Delta^2\Gamma+20A^3\Gamma\omega^2+8A^2\Delta^2\Gamma^2-24A^2\Delta^2\omega^2+14A^2\Gamma^2\omega^2+4A\Gamma^3\omega^2+16A\Gamma\omega^4+8\Delta^2\Gamma^2\omega^2+12A^2\Gamma^2\Omega^2+12A^3\Gamma\Omega^2+8A^2\Omega^2\omega^2+32\Delta^2\Omega^2\omega^2+16A^2\Delta^2\Omega^2-8\Gamma^2\Omega^2\omega^2+8A\Gamma\Omega^4+4A\Gamma^3\Omega^2+16A\Delta^2\Gamma\Omega^2+4A^4\Omega^2+16\Omega^4\omega^2-32\Omega^2\omega^4+4A^2\Omega^4+4\Gamma^2\Omega^4))/((A(A+\Gamma)^2+(A+\Gamma)\Omega^2+4A\Delta^2)(8A\Delta(2A^2+4A\Gamma+2\Gamma^2+4\omega^2)\Omega^4+(8A(4A^2+4A\Gamma+4\omega^2)\Delta^3+8A(A^4+3A^3\Gamma+3A^2\Gamma^2-3A^2\omega^2+A\Gamma^3-6A\Gamma\omega^2-3\Gamma^2\omega^2-4\omega^4)\Delta)\Omega^2) \quad \text{(C.4)}$$

The resonant frequency of the ion/photon system including Γ is:

$$\widetilde{\Omega} = \frac{1}{4}\sqrt{2\sqrt{25A^4 + 16\Delta^4 + 16\Omega^2 + 84A^3\Gamma + 40A^2\Delta^2 + 102A^2\Gamma^2 + 8\Gamma^2\Omega^2 + \sqrt{16A\Gamma\Omega^2 + 16A\Delta^2\Gamma + 9\Gamma^4 + 8A^2\Omega^2 + 52A\Gamma^3 - 24\Delta^2\Gamma^2 + 32\Delta^2\Omega^2} - \sqrt{6\Gamma^2 - 6A^2 - 12A\Gamma + 8\Delta^2 + 8\Omega^2}} \tag{C.5}$$

The limit of $\widetilde{\Omega}$ for $A \ll \Omega, \Gamma$ is:

$$\widetilde{\Omega} = \frac{1}{4}\sqrt{2\sqrt{\left(4\Delta^2 + 4\Omega^2 - 3\Gamma^2\right)^2 + 32\Gamma^2\Omega^2} + 8\Delta^2 + 8\Omega^2 - 6\Gamma^2} \tag{C.6}$$

The limit for $\Omega \ll A, \Gamma$ is:

$$\widetilde{\Omega} = \frac{1}{4}\sqrt{2\sqrt{25A^4 + 16\Delta^4 + 84A^3\Gamma + 40A^2\Delta^2 + 102A^2\Gamma^2 + 16A\Delta^2\Gamma + 9\Gamma^4 + \sqrt{52A\Gamma^3 - 24\Delta^2\Gamma^2} - 6\Gamma^2 - 6A^2 - 12A\Gamma + 8\Delta^2}} \tag{C.7}$$

D

Residuals for five-level scans

My attitude was always to keep open, to keep scanning.

Jonas Salk

This appendix shows the data, high-resolution scan and residuals for the fits discussed in section 5.5. The data and high-resolution scans are shown in figures 5.10 and 5.11.

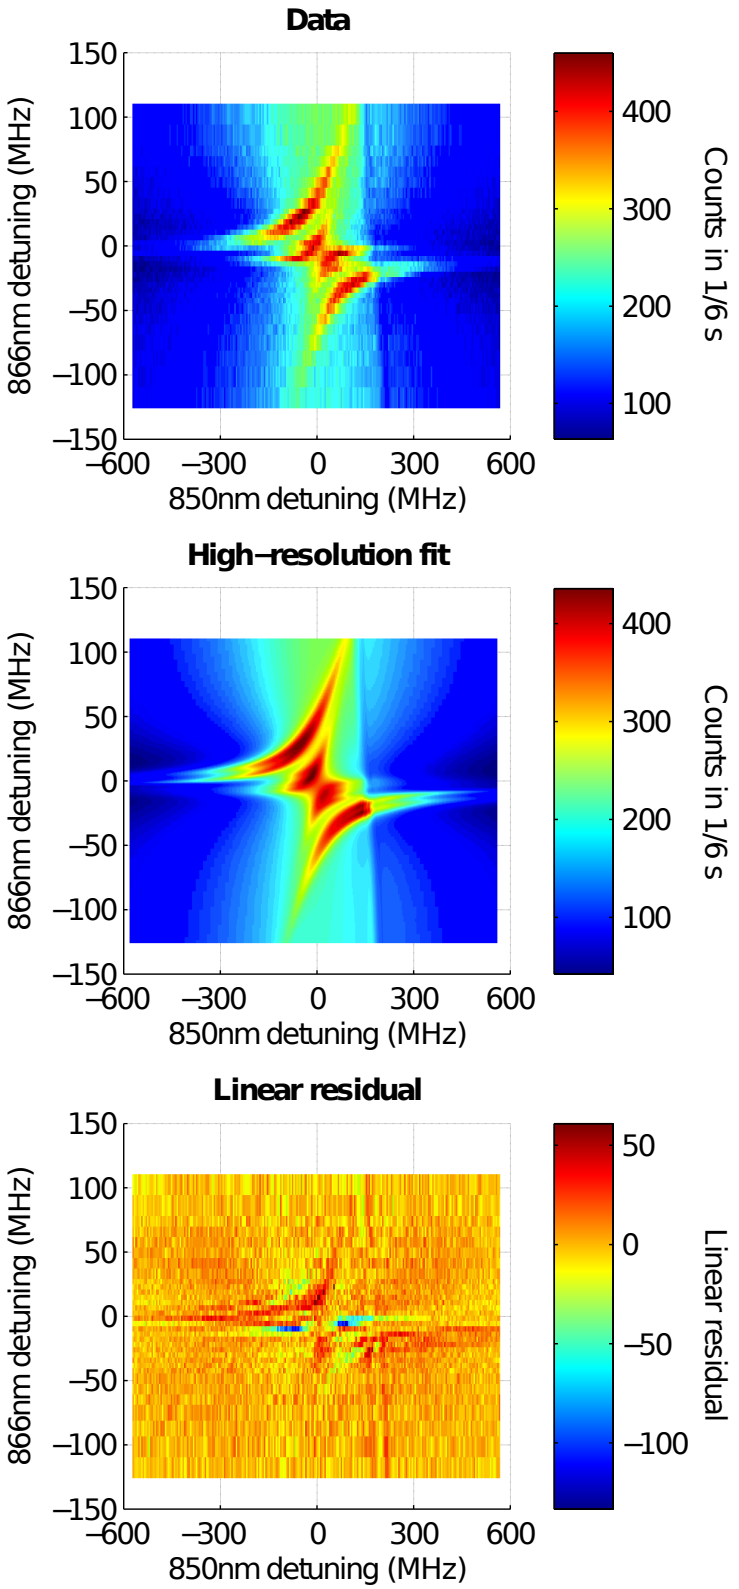

Figure D.1: The data, high-resolution scan and linear residuals of curve 1.

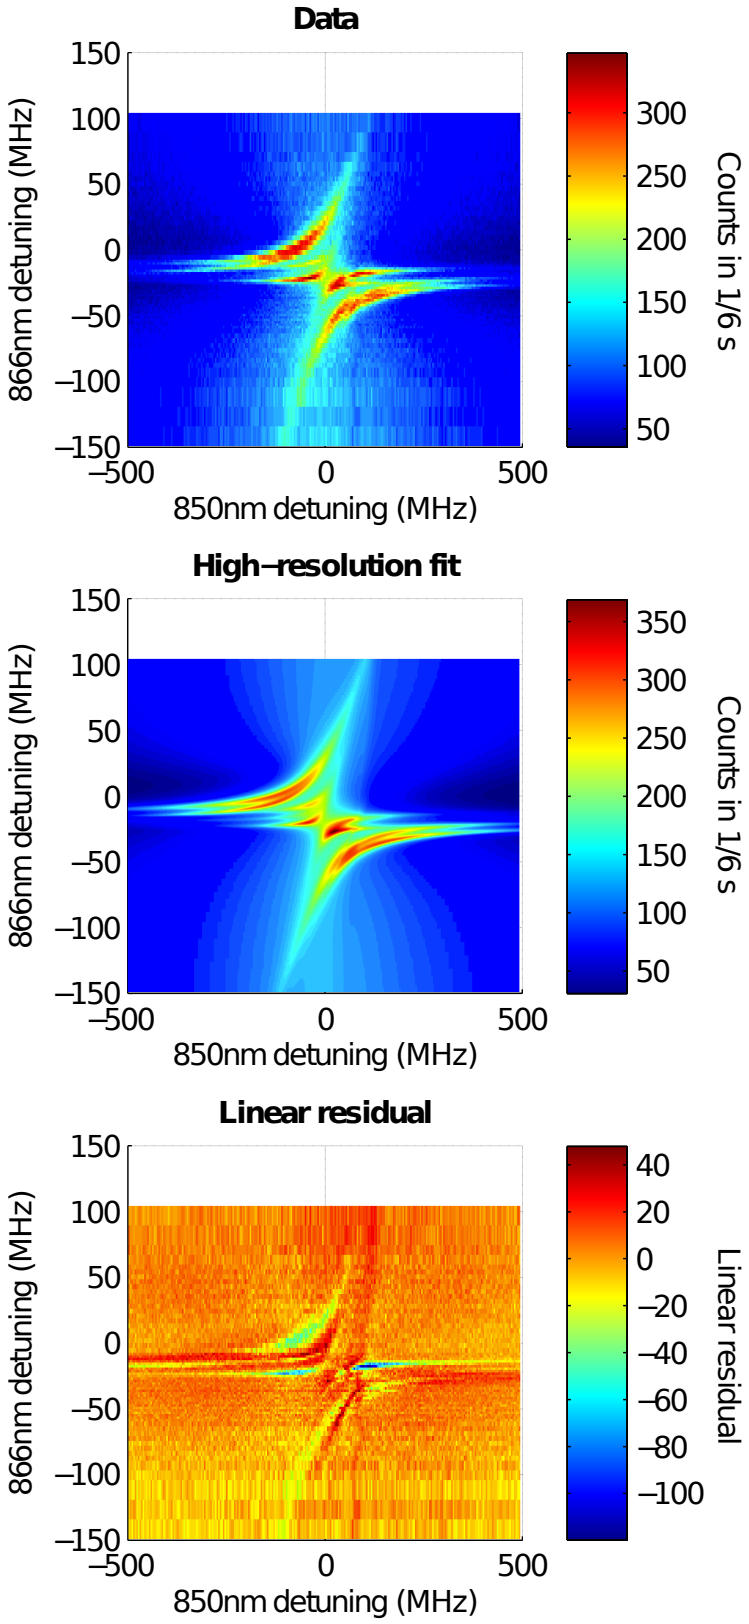

Figure D.2: The data, high-resolution scan and linear residuals of curve 2.

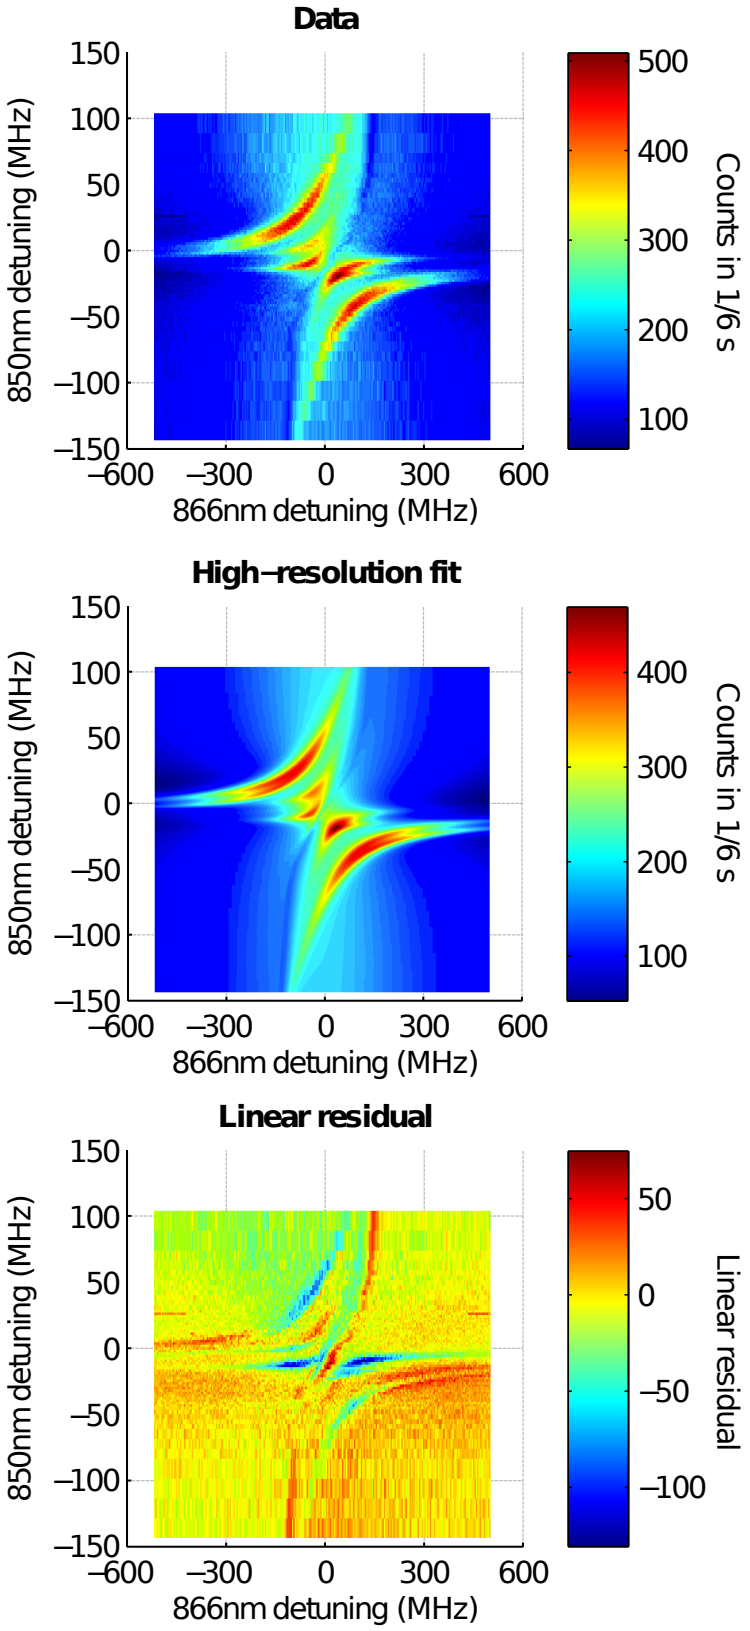

Figure D.3: The data, high-resolution scan and linear residuals of curve 3.

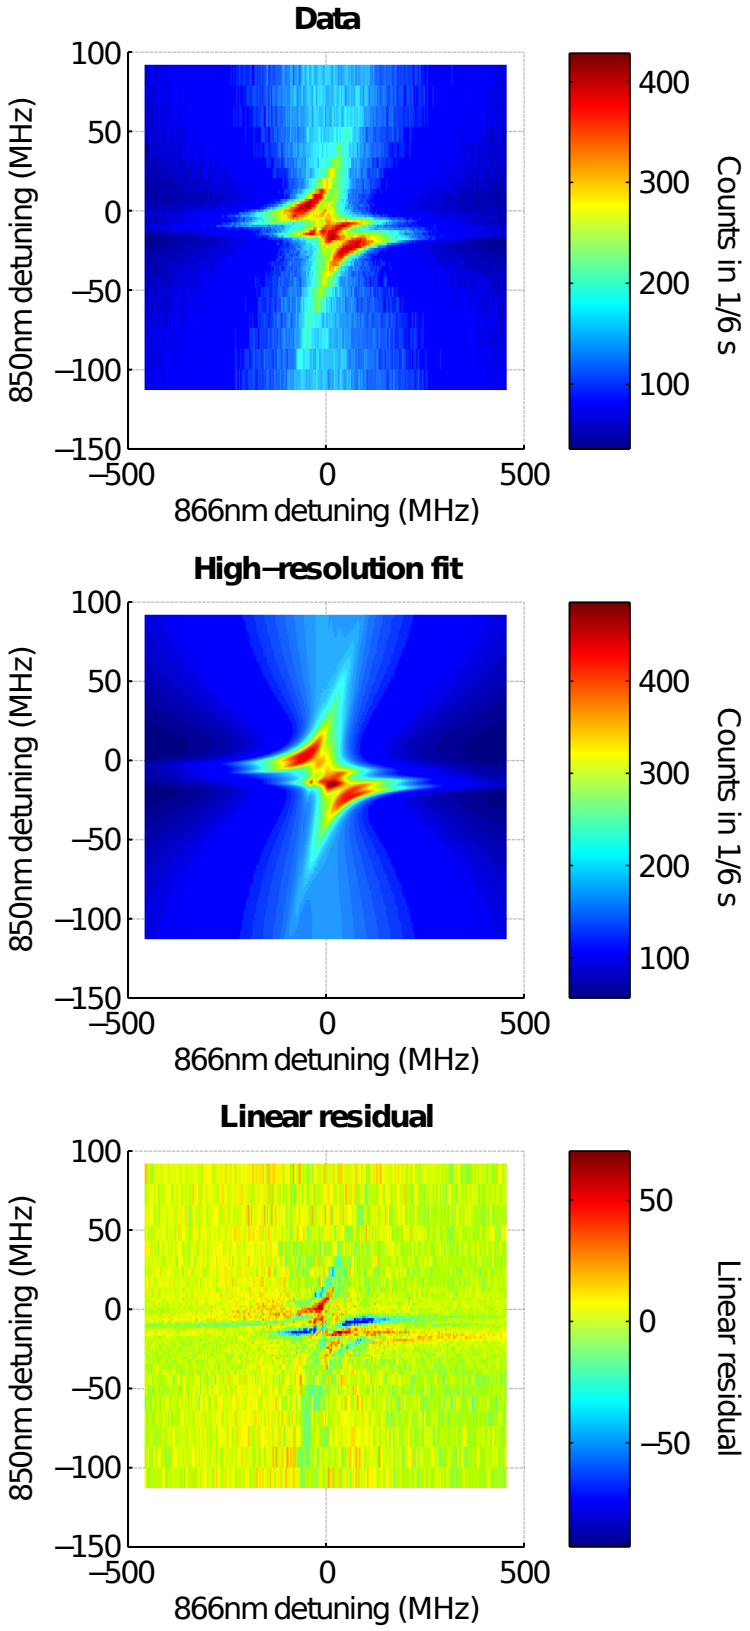

Figure D.4: The data, high-resolution scan and linear residuals of curve 4.

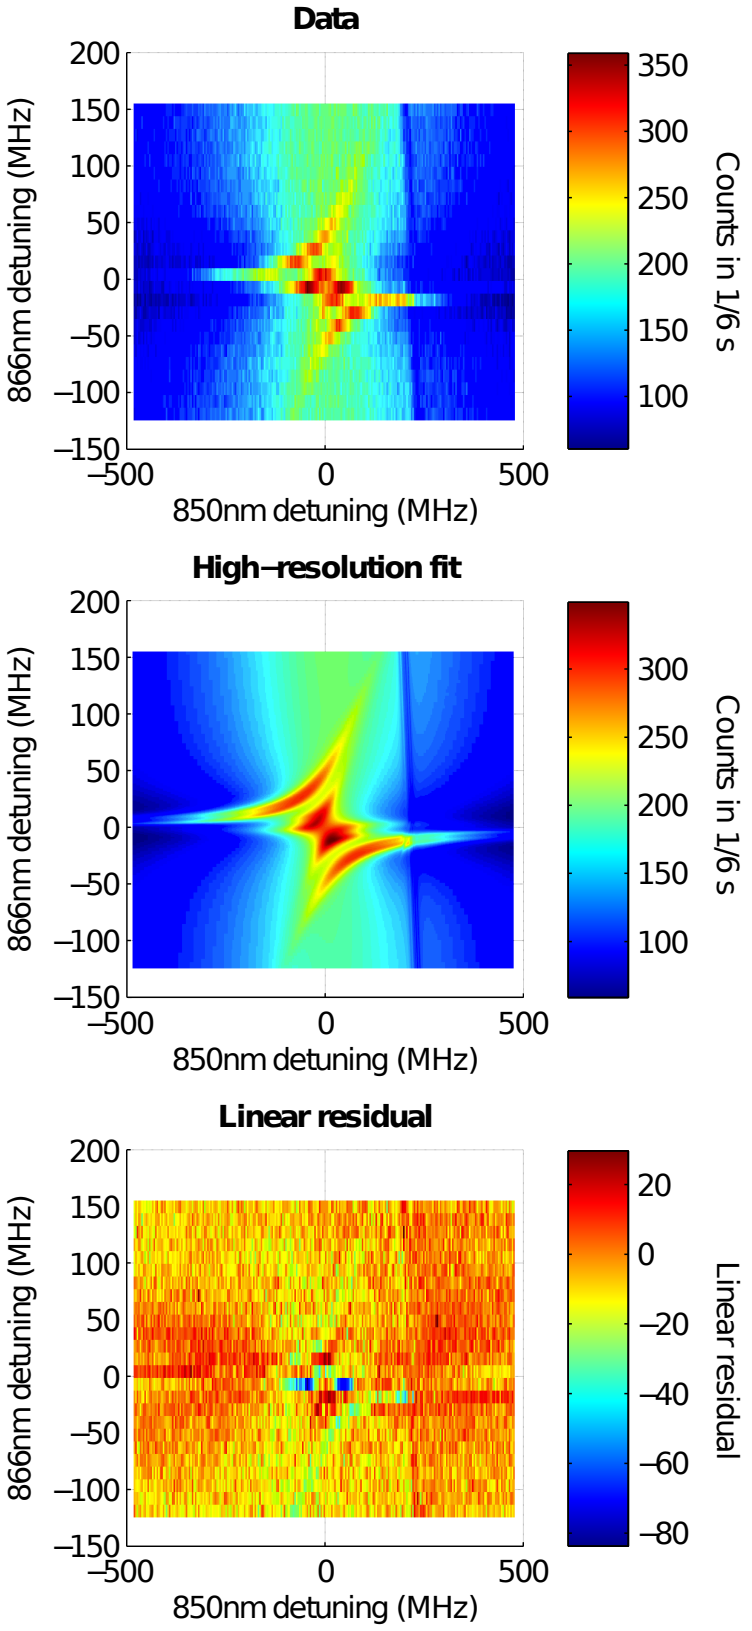

Figure D.5: The data, high-resolution scan and linear residuals of curve 5.

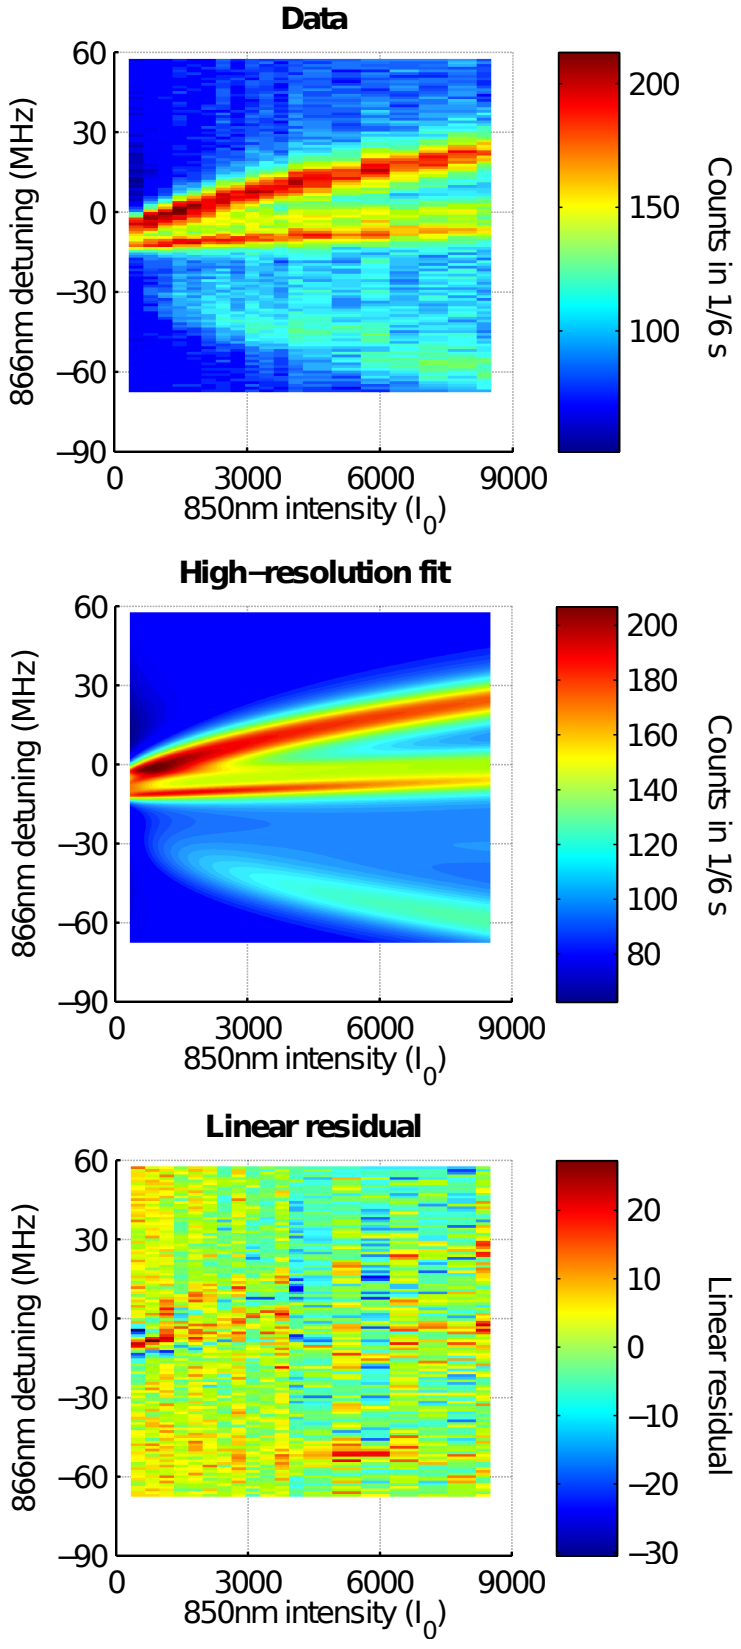

 Figure D.6: The data, high-resolution scan and linear residuals of curve 6.

E

Resonant effects in three, four and five-level systems

> It would be a long and hazardous climb to the stars. But surely it was in the very nature of man to make such journeys.
>
> *Sea-Horse in the Sky* by Edmund Cooper

The data described in chapter 5 shows resonant effects due to both its gross level structure and the presence of non-degenerate Zeeman states. In order to explain the behaviour, it is informative to first examine a set of much simpler systems. First, we look at the contributions arising from the level structure by examining effects that can be observed in systems with no Zeeman structure. We will first look at the simplest possible system that can produce a two-dimensional scan: the three-state system. We will then look at four-state and five-state systems in turn to demonstrate characteristic features that arise from resonant effects across three, four and five levels. The conventions I use to label the states, driven transitions and coherent effects are summarised in figure E.1. Appendix H lists the parameters used to produce the figures.

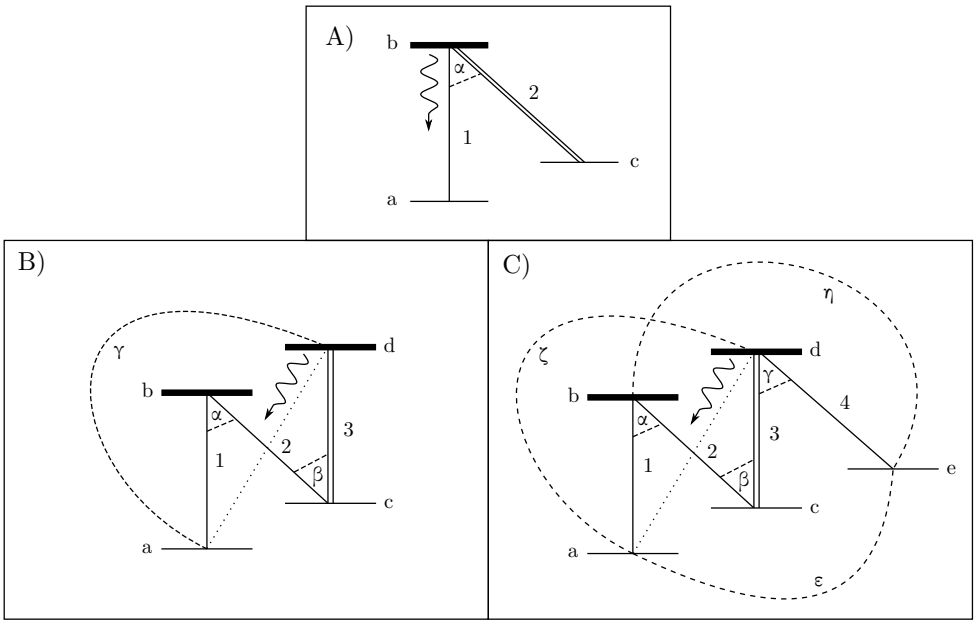

Figure E.1: A summary of the notation used to describe the states, driven transitions and coherent effects for the three, four and five-state systems discussed in sections E.1, E.2 and E.3. Transitions are labelled by numerals, states by roman letters and coherent effects by greek letters. States which show significant spontaneous decay are shown by a thick line. Driven transitions are shown as full lines; transitions that are not driven but from which fluorescence is collected are shown as dotted lines; and coherent connections between states are shown as dashed lines. Each system employs one beam that is much more intense than the others; this beam is shown by a double line. The transition from which fluorescence is collected is marked by the curved arrow. The population of the upper-energy state of this transition is used as a measure of this fluorescence.

E.1 The three-state system

The three-state system involving two-beams has been extensively studied. For a review of dark resonance behaviour demonstrated in these systems, see [Wol96]. Systems in which one beam is much more intense than the other are of particular interest due to Electromagnetically Induced Transparency (see [Mar98] for a review). Here, we look at the sort of features that can be seen in two-dimensional scans taken from such systems. To keep the analysis general, the spontaneous decay rates on each transition have not been set to equal values.

Figure E.2 shows a series of two-dimensional scans taken by successively increasing the intensity of beam 2. In this and all that follows, I will refer to the beam in the system with the largest intensity as the 'pump' beam and weaker

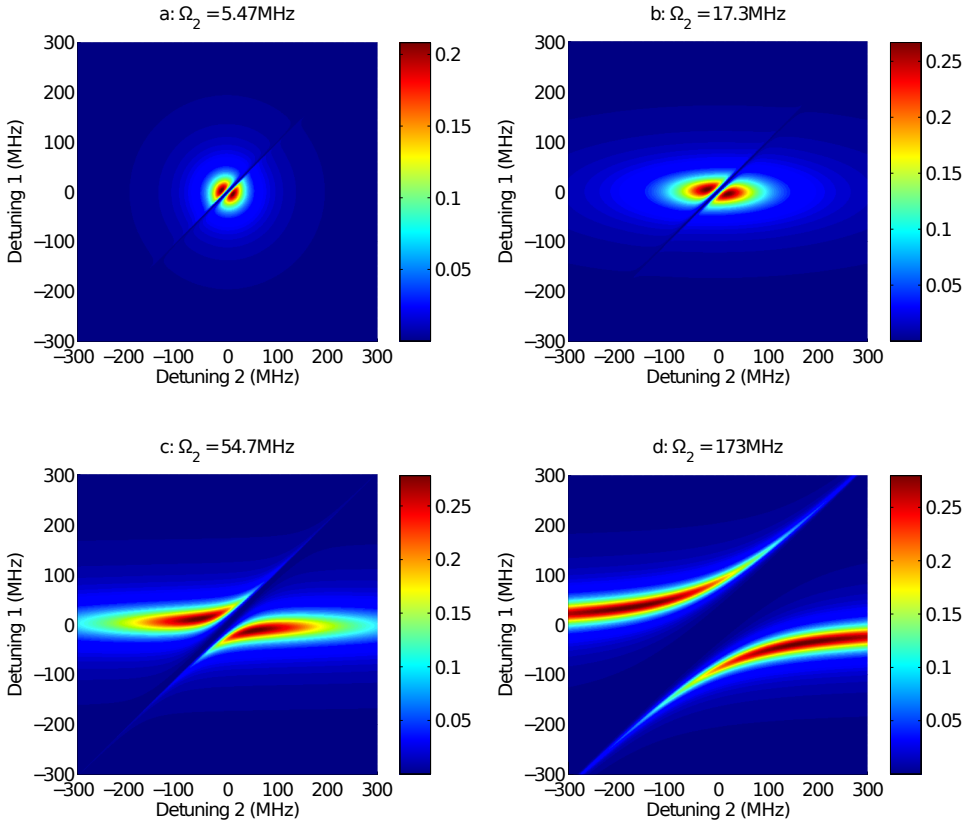

Figure E.2: The emergence of a hyperbola in the three-state system. This figure shows a series of two-dimensional scans varying the Rabi frequency of beam 2. The upper-left scan shows the output of a system for which the intensities of both beams are equal. The Rabi frequency of beam 2 is increased successively in the upper-right, lower-left and lower-right scans. The lower-right scan shows a characteristic hyperbolic effect. For more details, see text.

beams as 'probe' beams, in analogy to this system.

The plot on the upper-left shows a two-dimensional scan where beams 1 and 2 have the same small intensity. In this case, the output is straightforward. Each beam is weak; the intensity of both beams of order the saturation intensity. As a result, the fluorescence profile as a function of each separate detuning will be a Lorentzian with a dark resonance, the position of which will depend on the frequency not being scanned. The width of the Lorenztian will be equal for both beams as the intensities for each beam are equal and above saturation. When viewed in two-dimensions, the result is a two-dimensional Lorentzian with equal FWHM along both axis directions. The dark resonance appears as a

diagonal line cutting across this Lorenztian. As can be seen from the resonance condition:

$$\Delta_1 = \Delta_2 \tag{E.1}$$

this line has a gradient of +1 and passes through the origin.

The plot on the upper-right shows the result of slightly increasing the intensity, and thus also the Rabi frequency, of beam 2. The feature still has the shape of a two-dimensional Lorentzian, but the width of the Lorenztian along the horizontal axis has increased due to power broadening. The width of the dark resonance has also slightly increased, but it remains in the same position as the previous plot.

The lower two plots increase the Rabi frequency further and show new behaviour. The scans now no longer show a Lorenztian-type feature, but instead two arcs. These arcs tend to two asymptotes, one defined by the horizontal axis and the other defined by the diagonal dark resonance. In both plots, the position of the dark resonance is the same. Its width is also only slightly increased. Note that, while the separation between the arcs can become arbitrarily large, this separation is not due to the size of the dark resonance. The dark resonance still takes the form a thin diagonal line passing through the origin, but it is harder to see as it does not cut through a region with high population as in the previous plots. The splitting of the original resonant feature into two arcs due to the high intensity of the probe beam is a type of Autler-Townes splitting [AT55, Chi96].

The arcs in these systems where $\Omega_2 \gg \Omega_1$ are readily explained in the dressed-state approach. The intensity of light on the pump transition is so strong that it interacts with the atom and perturbs the resonant frequency of the atomic transition to a new resonant frequency in a combined atom/photon system. The frequencies at which the three-state system is driven are shifted from the dark resonance condition and a second type of resonance is formed: a bright resonance. The bright resonance occurs when the frequencies of the beams excite

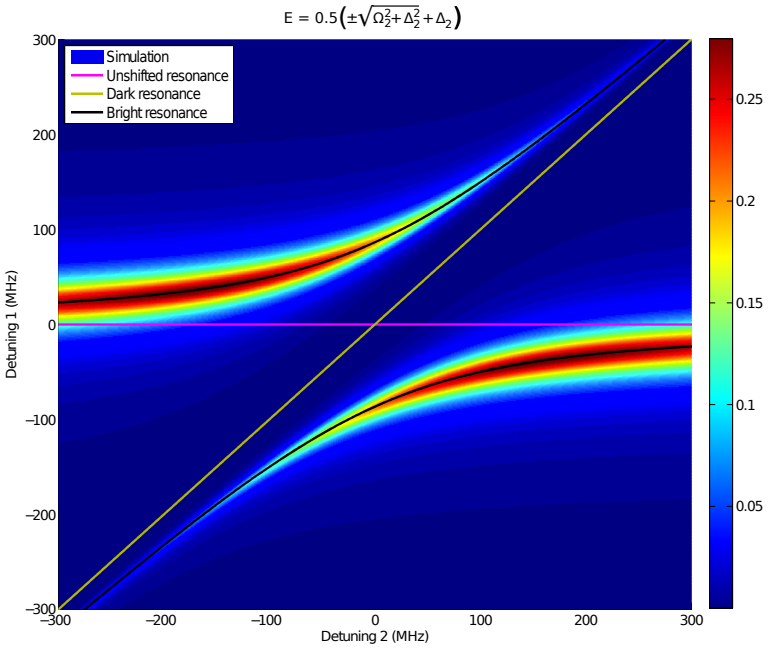

Figure E.3: The resonant effects in the three-state system. The resonant frequency of transition 1 is shown in pink. The dark resonance is shown in yellow. The bright resonance is shown in black. This takes the form of a hyperbola with asymptotes described the pink and yellow lines. 'Simulation' in the legend refers to the surface described by the colour bar.

the natural frequencies of the three-state system.

Let us simply derive the frequencies that drive this resonance. We first consider the two-state system where $A, \Gamma \ll \Omega$. The resonant frequency of this system is well known[Foo05, Ban06, Sho11], and is given by:

$$\widetilde{\Omega}^2 = \Omega^2 + \Delta^2 \quad . \tag{E.2}$$

Let us consider the situation for a system which has a negligible interaction with photons, so that $\Omega, A \to 0$. In this case:

$$\widetilde{\Omega}^2 = \Delta \quad . \tag{E.3}$$

Let us consider two systems. In system A, the laser light on the transition has a high intensity and a detuning Δ and the state energies have a bare frequency separation ω. In system B, the laser light has the same frequency, but the states are separated by a frequency $\omega + \Delta_s$. In addition, the system interacts negligibly

with photons. This ensures that the difference between the resonant frequencies of the atom and the atom/photon system is negligible. Because of the shift in the state energies in system B, the beam detuning is $\Delta_B = \Delta - \Delta_s$. We require that the resonant frequencies of these two systems are the same. Equating E.2 and E.3:

$$\Omega^2 + \Delta^2 = (\Delta - \Delta_s)^2 \quad . \tag{E.4}$$

The shift in the transition frequency for system 2 is:

$$\Delta_s = \pm\sqrt{\Omega^2 + \Delta^2} - \Delta \quad . \tag{E.5}$$

Assuming that both energy states are shifted by the same frequency difference κ in a symmetric manner, we determine the amount by which the laser shifts the energy of each state for system B:

$$\kappa = \frac{1}{2}\left(\pm\sqrt{\Omega^2 + \Delta^2} - \Delta\right) \quad . \tag{E.6}$$

κ is often termed as the 'light shift' of the states. See, for example, [SSAT98, MvdS99].

Let us now consider the three-state system. The effect of beam 1 on the state energies is negligible. The effective energy of state c will be shifted by κ due to the high intensity of beam 2. The net detuning between all three states will be zero for the bright resonance:

$$-\Delta_1 + \Delta_2 + \kappa = 0 \quad . \tag{E.7}$$

Writing κ in terms of Ω_2 and Δ_2 and solving for Δ_1:

$$\Delta_1 = \frac{1}{2}\left(\pm\sqrt{\Omega_2^2 + \Delta_2^2} + \Delta_2\right) \quad . \tag{E.8}$$

This is the bright resonance condition for the three-state system.

We can write this expression in a more convenient form for describing behaviour in a two-dimensional plot:

$$4\Delta_1^2 - 4\Delta_1\Delta_2 = \Omega_2^2 \quad . \tag{E.9}$$

This curve is a polynomial of degree two in Δ_1 and Δ_2 containing a cross-term and no terms of degree one. It therefore describes a hyperbola. Figure E.3 shows this hyperbola overlaid over the simulation and demonstrates that the hyperbola describes the location of the arcs extremely well. We can show more explicitly that (E.9) describes a hyperbola by actively constructing this hyperbola from its two asymptotes, $\Delta_1 = 0$ and $\Delta_1 = \Delta_2$. We take an equilateral North-South opening hyperbola:

$$y^2 - x^2 = a^2 \quad , \tag{E.10}$$

then stretch the x axis by $\cot(\pi/8)$ and rotate the curve through an angle of $\pi/8$ radians:

$$\left(-x\sin\left(\frac{\pi}{8}\right) + y\cos\left(\frac{\pi}{8}\right)\right)^2 - \left(\frac{x\cos\left(\frac{\pi}{8}\right) + y\sin\left(\frac{\pi}{8}\right)}{\cot(\pi/8)}\right)^2 = a^2 \quad . \tag{E.11}$$

This simplifies to:

$$4y^2 - 4xy = (2 + 2\sqrt{2})a^2 \tag{E.12}$$

which is exactly the form seen in expression (E.9) for the bright resonance condition. The bright resonance condition therefore describes a hyperbola. By comparing (E.9) and (E.12), we now know the minimum separation of the arcs of the hyperbola:

$$s = 2a = \frac{2}{\sqrt{2 + 2\sqrt{2}}}\Omega_2 \approx \Omega_2 \quad . \tag{E.13}$$

The separation between the two arcs therefore scales as the Rabi frequency, and thus as the square-root of the pump intensity, so long as $\Omega_2 \gg \Omega_1$. Figure E.4 shows the scaling of the separation between the arcs as a function of the pump intensity and demonstrates that it follows a square-root dependence.

In summary, there are two sorts of coherent effects demonstrated in two-dimensional scans taken using a three-state pump/probe system. These are a dark resonance, which always takes the form of a diagonal line, and a bright resonance which traces out a hyperbola.

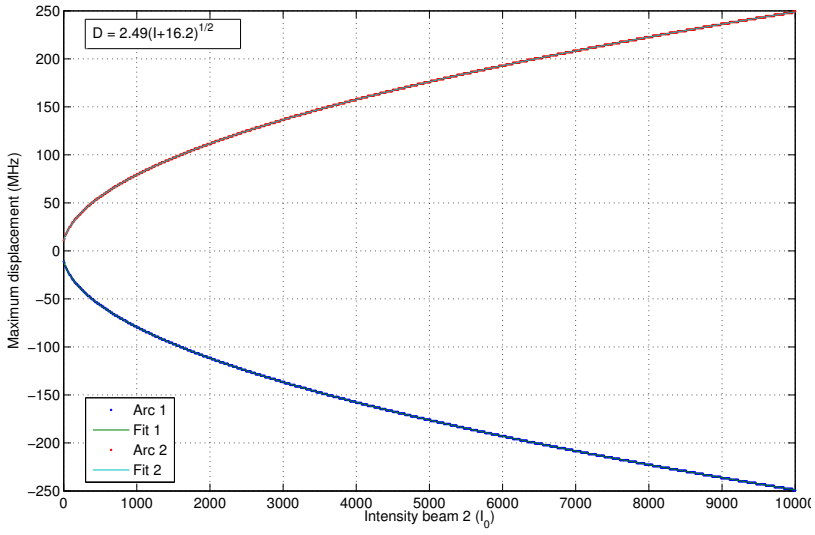

Figure E.4: The position of the two arcs of the bright resonance measured along the axis of symmetry as a function of the intensity of the pump beam. The position of both arcs, along with their separation, scales as the square-root of the intensity. The small offset at the origin arises from the finite lifetime of the two transitions. These limit the minimum size of the two-dimensional Lorentzian at small powers to the natural width of the respective transition along along each axis. The curves fit to a square-root dependence to machine precision.

E.2 The four-state system

We now examine the behaviour of the four-state system shown in figure E.5. This system shows an extension of the bright and dark resonance features seen in a three-state system.

Figure E.5 shows a series of two-dimensional scans at successively increasing values of the intensity of beam 3. This figure shows the emergence of two arcs in a similar manner to that seen in figure E.2 for the three-state system. However, there are several differences. These two arcs can be seen even at relatively low pump intensities, and the fluorescence is arranged more symmetrically between the two asymptotes. There is also a faint wide vertical line of increased fluorescence when the pump is on resonance. In addition, there is no visible dark resonance. Finally, and most importantly, the horizontal axis of the hyperbola is no longer aligned with a zero-detuning of beam 2; it has been displaced to a negative detuning.

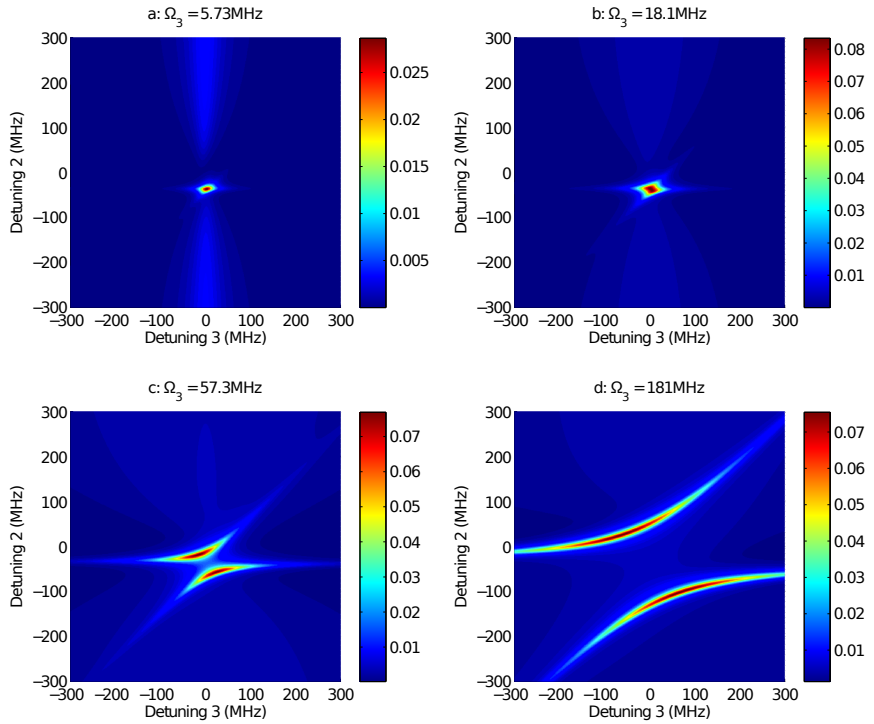

Figure E.5: The emergence of a hyperbola in a four-state system. This figure shows a series of two-dimensional plots increasing the intensity of the pump beam successively through the top-left, top-right, bottom-left and bottom-right scans. A hyperbola is seen even at relatively low intensities, and the fluorescence is distributed more symmetrically between the two asymptotes compared to the similar plots shown for the three-state system in figure E.2.

To explain these features, we examine the coherent effects present in the system. We split the system up into coherences connecting every state that is linked by one or more spontaneous transition and analyse these in turn. For the purposes of this study, we consider only coherences involving two or more spontaneous transitions. There are three different coherences of this type present in the system. Two involve connections between states separated by only one intermediate state, α and β. An additional resonance involves connections between a state separated by two intermediate states, γ. I will refer to the first sort of coherence as a 'three-level' coherence[1] and the second sort as a 'four-level' coherence, referring to the total number of states linking the two states involved

[1] I refer to these coherences as linking 'levels' rather than 'states' as these effects generalise to situations involving levels with Zeeman substates.

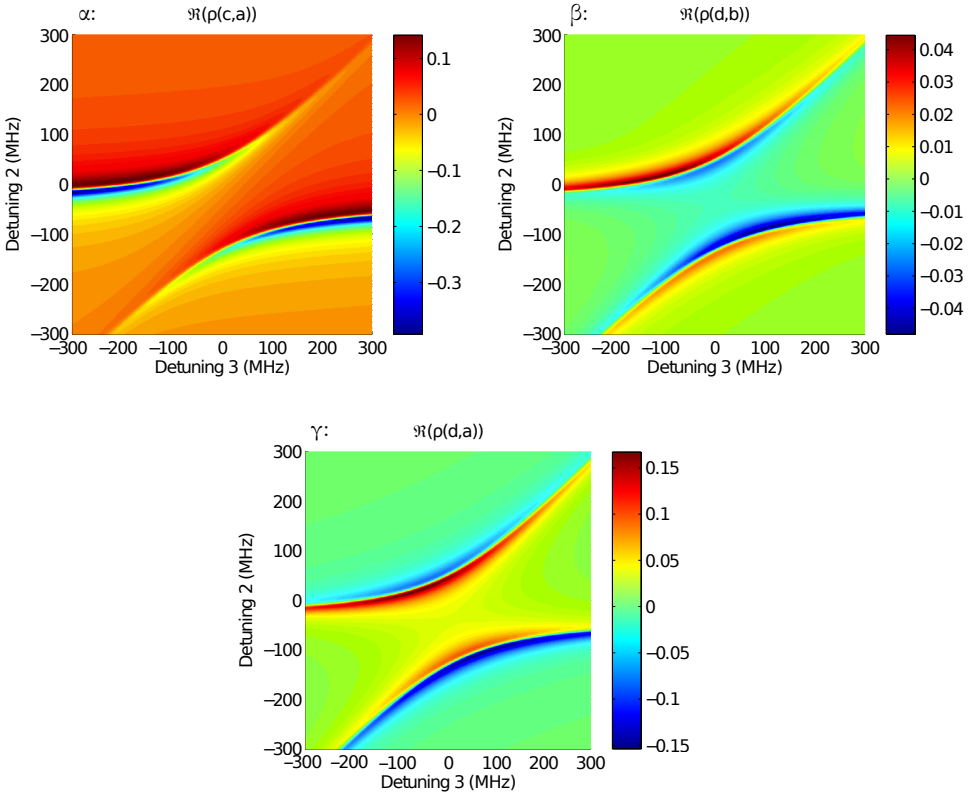

Figure E.6: The importance of the three coherences α, β and γ to the formation of the four-state hyperbola. Coherence α dominates close to the horizontal asymptote of the arcs and becomes negligible close to the diagonal asymptote. Coherences β and γ both follow the entire arc and are at their strongest when the fluorescence of the ion is strongest. β is weaker than γ by a factor of approximately three. For more details, see text.

in the coherence term[2].

Figure E.6 shows the relative strength of these coherent effects for the plot shown in the bottom-right of figure E.5 where the hyperbolic feature is pronounced. The coherences are shown on two-dimensional plots, and the strength of each coherence is measured using its real part. The advantage of plotting this measure, rather than the magnitude, is that the point at which the coherent effect is resonant is very clearly visible: over a resonance, the real part changes sign between two values with a large magnitude, so the resonance occurs at a zero-crossing. Resonance α is by far the strongest closest to the horizontal asymptote

[2]Such coherences could also be labelled as 'two-photon' and 'three-photon' coherences respectively. I have not used this terminology as referring to a multiple-photon effect often implies weak behaviour. These coherent effects, by contrast, dominate the behaviour of the system.

of the two arcs and disappears almost completely near the diagonal asymptote. Resonances β and γ both follow the entirety of the arcs and are strongest in those areas of the arcs when the fluorescence is strongest. However, resonance β has a weaker contribution to the dynamics than resonance γ. The three-level coherence is weaker by about a factor of three.

In the light of our analysis of the three-state system, we can interpret these plots as follows. The two arcs arise from a bright resonance driving the entire four-state system at its natural frequency. This bright resonance drives population between all four states of the system, and in general the coherence that best describes this effect is γ. When beam 3 is far detuned, however, it can only weakly drive coherent effects over all four states. Even though the whole system is coherently driven, the coherent effects are now by far the strongest between states a to c. In this regime, the coherence best describing the resonance effect is α. We can therefore describe the bright resonance as linking two regimes: a three-state regime described by α and a four-state regime described by γ. We describe these two regimes by using the resonance conditions for these coherences:

$$\alpha : \qquad \Delta_2 = \Delta_1 \tag{E.14}$$

$$\gamma : \quad \Delta_2 = \Delta_3 + \Delta_1 \quad . \tag{E.15}$$

The condition for α appears on a two-dimensional scan of Δ_2 and Δ_3 as a horizontal line which cuts the Δ_2 axis at Δ_1. The condition for γ appears as a diagonal line of gradient $+1$ passing through the point $(0, \Delta_1)$. The arcs approach these two lines asymptotically. Figure E.7 shows the position of these lines, and the line describing resonance β, overlaid over a two-dimensional scan.

In fact, the arcs seen in these scans are actually the two arms of a hyperbola, exactly as the arcs seen in the three-state system. When beams 1 and 2 have really very low intensities, the arcs can even be modelled by using an equivalent expression for the bright resonance condition in a three-state system defined

in (E.8), replacing variables originally describing beam 2 with those describing beam 3 and replacing the original Δ_1 with $\Delta_1 - \Delta_2$. This expression is:

$$\Delta_2 = \frac{1}{2}\left(\pm\sqrt{\Omega_3^2 + \Delta_3^2} + \Delta_3\right) + \Delta_1 \tag{E.16}$$

where the horizontal axis of the hyperbola has been shifted from the three-state expression so that it coincides with the resonance condition for α. More generally, the interactions of the two probe beams perturb the hyperbola slightly from this expression. The hyperbola is then better modelled using a more general expression:

$$\Delta_2 = \frac{1}{2}\left(\pm\sqrt{\overline{\Omega}_3^2 + (\Delta_3 + \delta_x)^2} + (\Delta_3 + \delta_x)\right) + \delta_y \quad . \tag{E.17}$$

This is a general expression for a hyperbola with one horizontal asymptote and one diagonal asymptote with a gradient of +1. δ_x and δ_y control the position of the centre of the hyperbola along the Δ_3 and Δ_2 axes respectively. $\overline{\Omega}_3$ controls the separation of the two arcs. The values of these variables are approximately $\delta_x \approx 0$, $\delta_y \approx \Delta_1$ and $\overline{\Omega}_3 \approx \Omega_3$. The bright resonance in a four-state system will *always* trace out a hyperbola that can be described by this expression.

Expression (E.17) can fit the calculated arcs to machine precision. Figure E.7 shows a hyperbola fitted to the fluorescence scan and the three coherent effects. The hyperbola clearly stretches between two lines approximately described by α and γ. The line describing β is given by the expression $\Delta_2 = \Delta_3$ and clearly does not relate to the hyperbola; the line intersects the curve with no visible interaction. This gives further support to the argument that the bright resonance extends between one regime best described by the three-level resonance α and another regime best described by the four-level resonance γ. Tracing along the hyperbola from a point close to the horizontal asymptote to one close to the horizontal asymptote, the bright resonance slowly shifts from driving population more strongly among the lower three states to driving population between all the states.

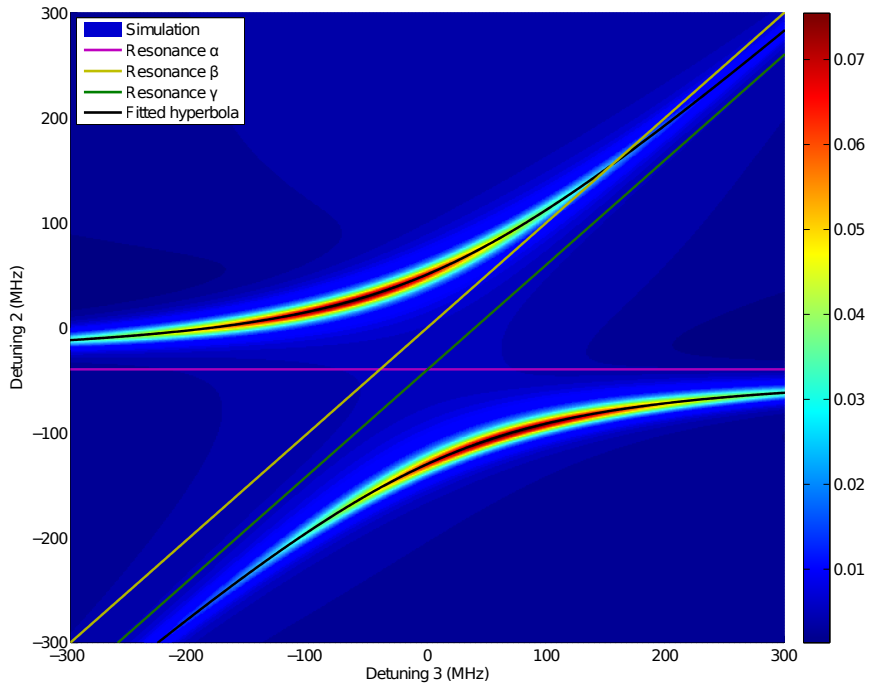

Figure E.7: A hyperbola in the four-state system. The black line shows a hyperbola fitted to the bright resonance in this system. This hyperbola fits to machine precision. Coherences α, β and γ are shown as lines overlaid on the plot. The lines describing α and γ clearly lie approximately on the asymptotes of the hyperbola. β is not related to the hyperbola. The line crosses the hyperbola with no discernible effect.

E.2.1 Differences between three-state and four-state hyperbolae

We can use this model of the behaviour to explain the observed differences between three-state and four-state hyperbolae.

The four-state hyperbola is displaced relative to the three-state hyperbola. A three-state hyperbola connects two regimes, one in which only the probe transition is strongly driven and the other in which all three states are driven. The four-state hyperbola connects a regime best described by a three-level coherence to one best described by a four-level coherence. The horizontal axis of a four-state hyperbola is defined by the detuning Δ_1, whereas the horizontal axis of a three-state hyperbola must always lie on the line $\Delta_1 = 0$. The centre of a four-state hyperbola will then in general be shifted in the vertical direction compared to a three-state hyperbola.

The fluorescence in the four-state system in the examples shown in this section is maximal closer to the centre of the hyperbola than for the equivalent examples in the three-state system. The three-state system will always have fluorescence biased towards the horizontal axis of the hyperbola. This is because a strong pump beam drives population out of the state that gives rise to the fluorescence signal. Strong fluorescence will only be seen when the pump beam is off-resonant so that population can build up in this state. However, the strongest fluorescence in a four-state system can occur at any point on the hyperbola. The hyperbola describes the relative importance of three-level and four-level coherences, and this relative importance is detained by the strength of the two probe beams. In the examples considered in this section, the intensities of the two beams were such that the fluorescence was maximal approximately around the semi-major axis of the hyperbola.

A characteristic feature of the three-state system is a dark resonance, but no such resonance is seen in the four-state system. This effect is a little more subtle to explain. A dark resonance can be seen in the three-state system because there is only a single coherent effect, meaning that this coherence will demonstrate very strong resonant behaviour. However, this is not the case for the four-state system. There are two three-level coherences that could show a dark-resonance. Resonance α drives population between states a and c through state b, while β drives population between b and d through state c. We do not see the effect of either dark resonance. To explain this, let us consider a situation in which α is perfectly on resonance and population is equally distributed between states a and c. The dark resonance prevents any population being transferred to state b from these states. However, a route out of the dark resonance exists: population can be excited to state d. From this state, it can be excited into the 'forbidden' state b. In essence, the additional transition out of the two populated states of the dark resonance destroys the coherence. In general, a dark resonance will be

disrupted if there is some additional route to its forbidden state involving states outside of the resonance. This is the case for both α and β, and therefore neither resonance is observed.

The four-state hyperbola can be seen even at low powers of the pump beam, whereas a Lorentzian behaviour is instead seen for the equivalent case in the three-state system. This is because of the different nature of the coherent effects linked by the two hyperbolae. A three-state system links two regimes, one in which only the probe is driven and the second in which the whole system is driven. However, this is only the case when one beam is stronger than the other. When both beams are of nearly equal intensity, there is no way to distinguish between the two transitions and no way to define a 'weaker beam'. In this situation, population is only driven coherently between the states when both beams are on resonance, and the two-dimensional plot simply looks like a two-dimensional Lorentzian. A hyperbola will therefore only be visible when one beam is significantly more intense than the other. By contrast, there is no ambiguity in a four-state system. The four-state hyperbola describes how strongly coherences are driven in a three-state subset of the system compared to the coherences driven over the entire system. When one beam is detuned, coherent effects involving this transition will always be reduced regardless of the intensity of this beam. The bright resonance will therefore always connect two asymptotes, one of which describes strong coherences in the three-state subset and one describing strong coherent effects over all four states. The intensity of the beam simply affects the separation of the two arcs. In figure E.5, the top-left plot shows a two-dimensional scan where all beams have approximately the same intensity. While the finite width of the arms of the hyperbola has caused the two arcs to apparently merge, the convergence to the horizontal and diagonal asymptotes is still observed as a horizontal and diagonal line of increased fluorescence passing through the central peak. A hyperbola will always be visible in a four-state

system, regardless of the intensity of the beams involved.

Finally, a broad vertical line of fluorescence centred on the zero of the x axis can be seen in the four-state system but not in the three-state system. This arises from an effect that can only be observed in the four-state system. This vertical line does not follow any of the coherence lines and therefore is not a coherent effect. In fact, it relies instead on spontaneous emission from state b to state c. Population that decays into state c can be driven into state d by beam 3, and will be driven most strongly when the beam is exactly on resonance. This is the cause of the vertical line. The width of this line depends on the intensity of beam 3 due to power broadening. At large intensities of beam 3, the line is spread across such a large width that it can be hardly be seen.

In conclusion, by treating the four-state system and three-state system in terms of a set of coherent effects linking subsets of their states, we can explain the difference in the features seen in two-dimensional scans taken in the two systems.

E.2.2 Behaviour with multiple pump beams

So far, we have considered the behaviour of a four-state system with one pump beam and two probe beams. In a three-state system, we can only ever probe a single pump beam. However, in the four-state system, it is possible to probe two pump beams. It will be instructive to look at what happens to the behaviour of the system we studied in the previous section when we increase the powers of each of the probe beams in turn so that they are no longer small compared to the power of the pump beam.

Such systems show interesting behaviour. Figure E.8 shows a series of plots in which the intensities of beam 1 and beam 2 are increased from their values in the set of parameters that describe the lower-right plot in figure E.5. The upper two plots show the effect of increasing the intensity of beam 2. The lefthand plot

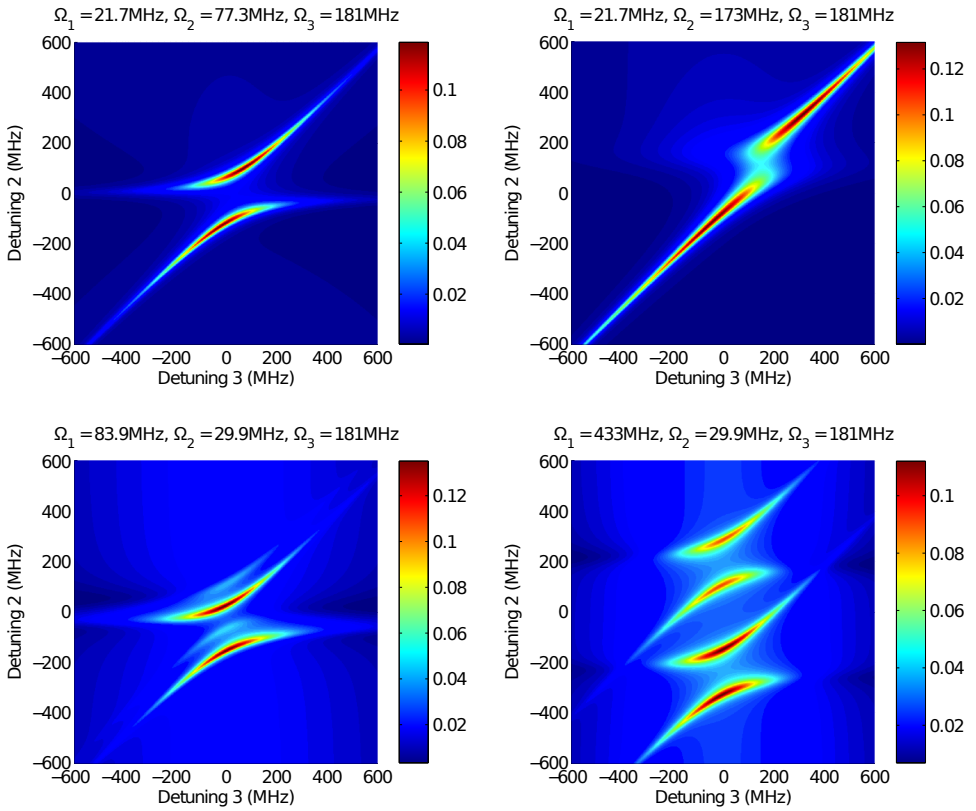

Figure E.8: The behaviour of the four-state system when two beams have a large intensity. The base set of parameters is that describing the lower-right plot in figure E.5. The two upper plots increase the intensity of beam 2. The lower two plots increase the intensity of beam 1. The lefthand plots show a small increase to the altered beam power and the righthand plots show a large increase. Increasing the intensity of beam 2 shifts the hyperbola upwards and moves the region of the hyperbola with the strongest fluorescence towards the diagonal asymptote. This effect is very pronounced for a large beam power, where the fluorescence lies very close to the asymptote. Increasing the intensity of beam 1 causes the hyperbola to split into two images that are displaced upwards and downwards from the unshifted hyperbola by an equal amount. For more details, see text.

shows a small increase and righthand plot a large increase. As the intensity of the beam is increased, the hyperbola is distorted. The regions of the hyperbola in which the fluorescence is largest are shifted towards the diagonal asymptote. In addition, the centre of the hyperbola is displaced. The effect is very pronounced in the righthand plot, where the fluorescence is aligned almost entirely along the diagonal asymptote. However, the behaviour still traces out a hyperbola.

We can explain this behaviour by noting that beam 2 controls the relative amount of population in state b compared to states a and c. When beam 2 is

very intense, it strongly drives population out of c into b; state c is sparsely populated. The fluorescence signal of the horizontal resonance α relies on stimulated transitions from state c to state d; these are strongly reduced due to the small population in c. By contrast, the diagonal resonance γ relies on strong coherent effects over the entire system. Increasing the intensity of beam 2 increases the strength of this coherence and the fluorescence signal associated with it. As a result, the regions of the hyperbola with the highest fluorescence will shift towards the diagonal asymptote as the power of beam 2 is increased.

The centre of the hyperbola is displaced because the natural frequency at which α is driven is a function of the intensity of beam 2 through the a.c. Stark shift. At small beam powers, the dependence is negligible. However, when it becomes prominent, the frequency will be shifted, changing the position of the horizontal asymptote of the hyperbola.

The lower two plots in figure E.8 show the effect of instead increasing the intensity of beam 1. The lefthand plot shows a small increase and the righthand plot a large increase. The effect is striking. The hyperbola seen in a system where beam 1 has a low intensity splits into two duplicates as the power of beam 1 is increased. These duplicates are shifted upwards and downwards by the same frequency compared to the low-intensity hyperbola. The relative strength of the two images depends on the polarisation of beam 1.

This effect can be explained by noting that beams 1 and 3 share a particular characteristic: both beams connect to a state at one end of the chain of transitions. Increasing the powers of either beam has a similar effect on the system. When all the beams are of similar intensity, the primary feature of the two-dimensional scans we have been considering is a small point of fluorescence (although other features are visible, such as the faint outlines of the hyperbola as it tends to the two asymptotes). As we have seen, increasing the intensity of beam 3 causes this point to split into two distinct arcs. The same effect occurs when

we instead increase the intensity of beam 1 compared to all other beams. However, on our two-dimensional scans, we now observe the bright central point splitting into two *points* translated in the vertical direction[3]. The critical beam determining how strongly each state interacts with the bright resonance is beam 1, a beam which is not scanned in this plot; there are only two values of beam 2 that can be resonant with a fixed frequency of beam 1. By making the intensities of both beams 1 and 3 large, splittings due to both beams will be observed. The hyperbola produced by the large intensity of beam 3 will be further split into two images along the vertical axis due to the splitting of beam 1.

Note that no splitting is observed when beam 2 has a large intensity. This is because this particular transition cannot participate in the mechanism that causes splitting in the other two beams. Beams 1 and 2 produce hyperbolic effects due to the interaction between a regime in which coherences are driven in a subset of three states and a regime in which coherences are driven in the entire system. However, there is no subset of three states connected coherently when beam 2 is far off resonant. In fact, there can be no coherent effects in the system at all! As a result, splitting can only be observed from beams 1 and 3.

In summary, there is only a single feature that can be seen in a four-state system. This is a bright resonance that takes the form of a hyperbola. When the power of the two 'pendant' beams involving the outer energy states are both large compared to the inner beam, the bright resonance traces out two duplicate hyperbolae translated in the vertical direction.

E.3 The five-state system

We now examine the behaviour of the five-state system shown in figure E.1. This system shows additional structure in the bright resonance features. This

[3]The splitting of the bright resonance into two is a manifestation of the Autler-Townes effect. This is the case whether the system splits into arcs or points, although most common examples of the effect involve a resonant point (such a spectral line) splitting into two duplicate points.

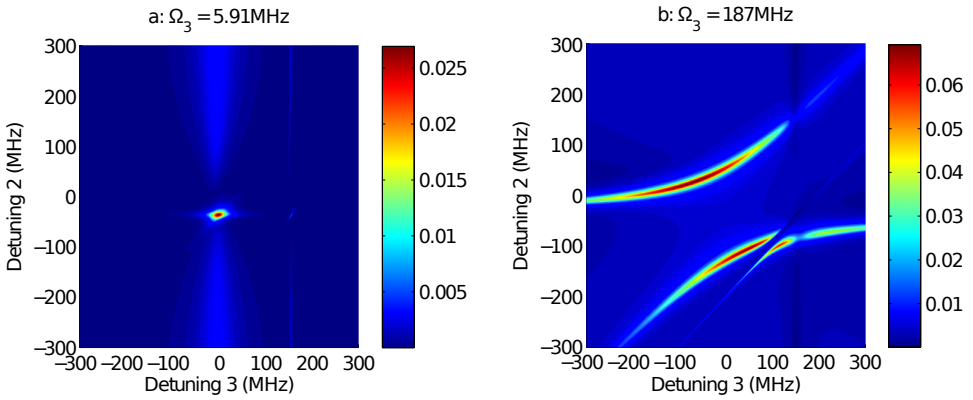

Figure E.9: The structure of two-dimensional fluorescence scans in a five-state system. The left-hand plot shows the behaviour when beam 3 has of order the same intensity as the other beams while the righthand plot shows the behaviour when it is significantly higher. These two plots share the parameters of beams 1 to 3 with two of the similar plots produced for the four-state system in figure E.5. The lefthand plot shares the parameters with the top-left plot in the four-state system and the righthand plot with the bottom-right plot. The gross structure of the five-state and four-state scans are very similar. However, two additional features can be seen in the five-state system. These are a vertical line centred on $\Delta_3 \approx 150\text{MHz}$ and a diagonal line that interferes with the hyperbolic feature. For more details, see text.

additional structure is the first stage in understanding the resonance structure seen in the experimental data examined in chapter 5.

Figure E.9 shows a pair of two-dimensional scans, one in which the intensity of beam 3 is of the order of all the other beams and one in which it is significantly higher. The parameters describing beams 1 to 3 in these two scans are the same as two of the plots shown in figure E.5 for the four-state system. The values of these parameters for the lefthand plot in figure E.9 are the same as those in the top-left plot of the figure E.5, and the values for the righthand plot are the same as those in the bottom-right plot[4] of figure E.5.

Most of the structure seen in the five-state plots is similar to that seen in the four-state system. The position, width and peak fluorescence associated with the hyperbolic feature are almost identical in the two cases. However, there are two further features present in both of the five-state plots. The first is a vertical line at

[4]The Rabi frequencies for beam 3 differ slightly between these two figures even though the intensities of beam 3 are the same; the additional transition from state d in the five-state system changes the decay constant on transition 3 and therefore slightly alters the Rabi frequency.

$\Delta_3 \approx 150\text{MHz}$. In the righthand plot, this feature cuts the hyperbola and reduces the fluorescence of the feature as it does so. In the lefthand plot, the feature takes the form of a vertical line of increased fluorescence in a similar manner to the much wider line at $\Delta_3 \approx 0\text{MHz}$ which we associated with driven transitions on the pump beam. The second feature is a diagonal line. This line has a more complicated effect on the fluorescence structure. In the low-intensity plot, it forms a narrow line of slightly increased fluorescence, although it is difficult to resolve on the scale of this figure. On the righthand plot, however, the line interferes with the hyperbolic feature. Close to the intersection of the line and the hyperbola, the bright resonance bends away from the base hyperbola and converges on the line asymptotically; there is an avoided crossing. In addition, there is a large drop in fluorescence on the intersection of the diagonal and the vertical line.

We explain these features using the same approach we took with the four-state system; we examine coherences between every pair of states linked by a stimulated process. The total number of these coherences in an $n-$state system is given by the number pairs of states in the system. After excluding the $n-1$ transitions between neighbouring states, the number of coherences is:

$$C_n = \binom{n}{2} - n + 1 \tag{E.18}$$

For the five-state system, this gives $C_n = 6$. For this system there are three three-level coherences α, β and γ and two four-level coherences ζ and η. In addition, there is a coherence connecting two states separated by three intermediate states, ϵ. I will refer to this coherence as a 'five-level' coherence. I will distinguish between coherent effects in the five-state and four-state systems by using a subscript 5 and 4 to avoid confusion.

A series of two-dimensional plots of the strength of these coherences is shown in figure E.10. We can be guided in our interpretation of these plots by our analysis of the four-state system. The behaviour of the resonances α_5, β_5

and γ_5 are very similar to their equivalents α_4, β_4 and γ_4 respectively. The shape and magnitudes of the features seen are almost identical except in those regions where the two additional five-state effects are prominent. The hyperbolic portion of the dark resonance is dominated by coherence α_5 closest to the horizontal asymptote and ζ in all other regions. By far the strongest coherent effect on the five-state vertical feature is γ_5. The new five-state diagonal feature is dominated by coherence ϵ. γ_5 and ϵ are equally strong about the point of intersection of the new five-state vertical and diagonal features. Finally, coherences β and η are dominated everywhere by other coherent effects.

We interpret these results as follows. The similar nature of the five and four-level coherences over the hyperbolic part of the fluorescence feature suggests that the mechanism giving rise to this portion of the feature is the same in both systems. For practicality, we consider the system as a chain of states stretching from the 'lowest' state a to the 'highest' state e. The bright resonance traces out a hyperbola that links together a region in which coherences are strongest in the three lowest states in the chain (the α asymptote involving a, b, c) to one in which coherences are strongest in the four lowest states (the ζ asymptote involving a, b, c, d). This change in the 'emphasis' of the five-level coherence is due to the detuning of beam 3. When it is far from resonance, coherent effects involving the fourth state in the chain d are suppressed.

This approach immediately lends itself to an extension. The resonance linking all five states in the system is ϵ. The resonance condition for this coherence is:

$$\Delta_2 = \Delta_3 + \Delta_1 - \Delta_4 \quad . \tag{E.19}$$

This traces out the diagonal line that cuts the hyperbola. When Δ_3 is close to resonance with this condition, coherences will be strongly driven between all the states in the system. The bright resonance in this system therefore connects *three* asymptotes, rather than two. Near this third asymptote, the bright reso-

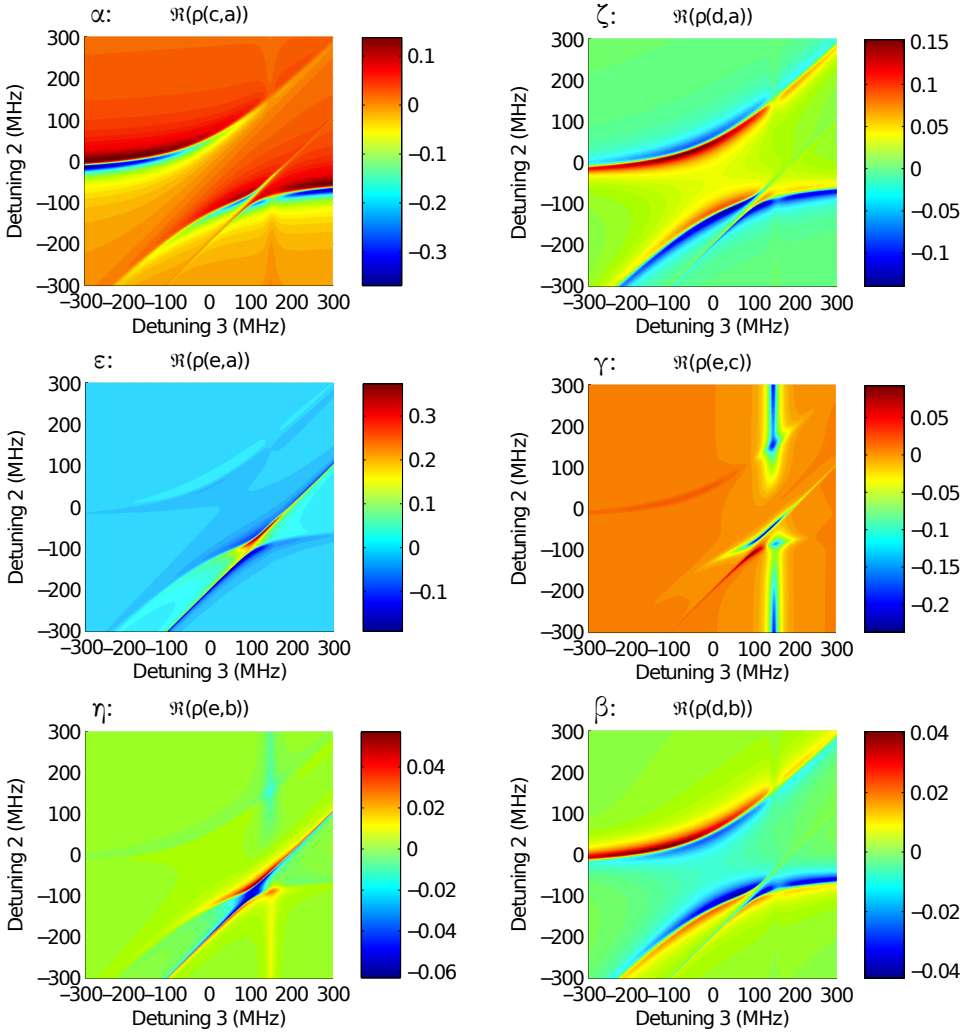

Figure E.10: The importance of the six coherence effects to the features seen in a two-dimensional scan in the five-state system. Resonance α is strongest closest to the horizontal asymptote. Resonance ζ is strongest over the hyperbolic portion. Resonance γ is strongest for the new vertical feature. Resonances γ and ϵ are strong for the new diagonal asymptote of the bright resonance. Resonances η and β are much weaker than than the other coherent effects. β is weaker than ζ on the hyperbolic portion by approximately a factor of three, whilst η is weaker than γ and ϵ on the new diagonal asymptote by a similar factor. For more details, see text.

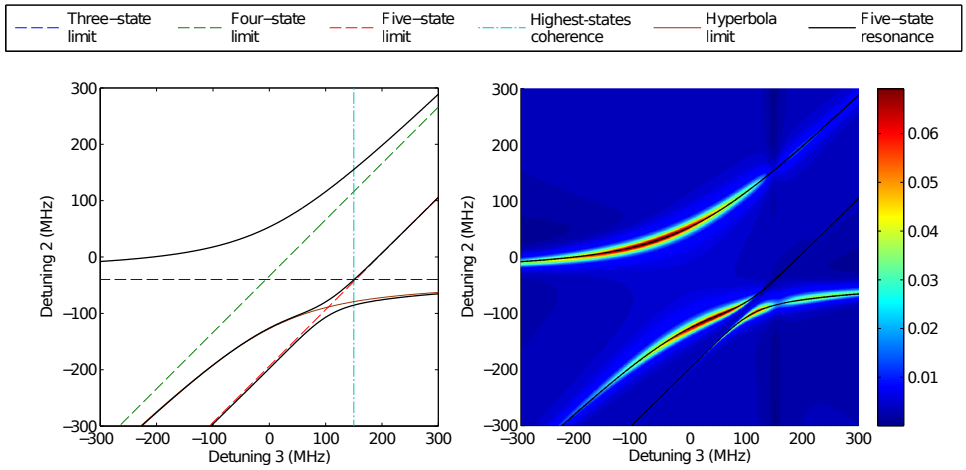

Figure E.11: A demonstration of the coherent effects giving rise to the modified hyperbolic feature. This feature connects three regions and is asymptotic to three lines. When beam 3 is very far detuned from resonance, coherences are most strongly driven in the three lowest states in the chain. When beam 3 is close to the resonance condition of ζ, coherences are strongly driven in the four lowest states. Finally, when beam 3 is close to the resonance condition of ϵ, coherences are strongly driven across all five states. We can consider the bright resonance as connecting three regions which primarily involve three-state, four-state and five-state behaviour. Each of these regions is described by an asymptote. The bright resonance therefore describes how strongly coherences are driven in the three, four and five-state subsets of this system. To complete the picture, the resonance causing the vertical feature is shown, and a fit is made to the bright resonance using a general expression for a curve with multiple asymptotes.

nance curves away from the hyperbolic behaviour described by the other two asymptotes to converge on this diagonal line.

The modified hyperbolic feature therefore links three regimes. The first regime one in which coherences are most strongly driven in the three lowest states of the chain (a, b, c), one in which they are driven in the four lowest states (a, b, c, d) and one in which they are driven over the entire chain (a, b, c, d, e). This behaviour is summarised in figure E.11.

E.3.1 Differences between four-state and five-state hyperbolae

We have two more features to explain. We treat the additional vertical feature first. The behaviour of this feature is similar to that of the much wider vertical feature that is most easily seen in the lowest-intensity plots in figures E.5 and E.9. The origin of that feature is due to population that spontaneously decays

from state b to state c being driven most strongly into state d when beam 3 is on resonance. However, with the five-state system, there is an additional feature that can arise because of the spontaneous decay on transition 2. Population driven into state c can now be driven coherently in the three highest states in the chain (c, d, e). The coherence that best describes this effect is γ. Close to resonance with γ, population in state c is coherently driven into state e, reducing the population in state d and thus the fluorescence. The horizontal resonance is independent of the coherences forming the modified hyperbolic feature and thus does not alter the shape of that feature; it merely reduces the fluorescence of the feature when it intersects it.

The final feature we examine is the large drop in fluorescence when the additional five-state horizontal and diagonal features interact. The feature shows the properties of a dark resonance; when the linewidths of all the beams are set to zero, the fluorescence at the intersection of γ and ϵ drops exactly to zero and the populations of states b and d also drop to zero. We observe dark resonance behaviour in the five-state system even though we could not in the four-state system.

Dark resonance behaviour is not observed in the four-state system because all the three-level coherences that could cause dark states interfere; additional paths to the forbidden state of one three-level coherence are provided by another. However, in the five-state system, there are two three-level coherences that do not interfere. α and γ only share a single state. When both resonance conditions are met, neither coherence provides additional routes to the forbidden state of its neighbour. Therefore, the two dark resonances are not destabilised. This drop in fluorescence is due to the combination of *two simultaneous* dark resonances. It occurs when both α and γ are on resonance. The resonance

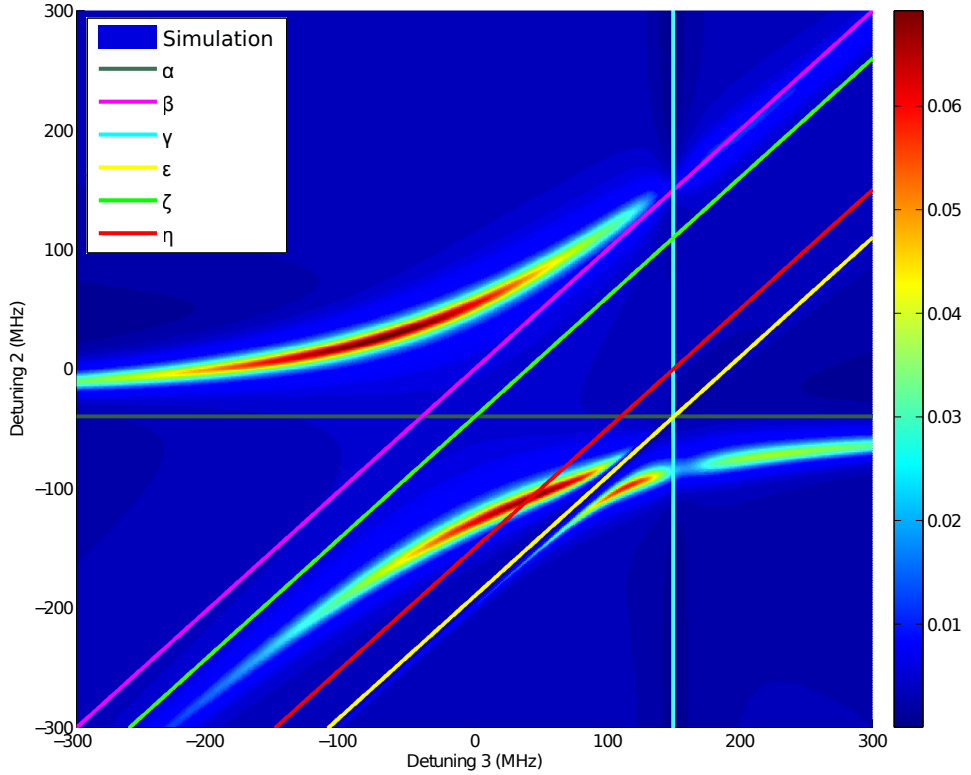

Figure E.12: The coherent effects giving rise to the modified hyperbola. This plot shows every coherent effect in the five-state system overlaid over the righthand two-dimensional scan in figure E.9. The asymptotes of the modified hyperbola approximately lie along the lines describing α, ζ and ϵ. The vertical feature lies approximately along the line describing γ. Resonances β and η cut the hyperbola with no visible effect. These coherences have a negligible effect on the dynamics of the feature.

conditions for these coherences:

$$\alpha: \quad \Delta_2 = \Delta_1 \tag{E.20}$$

$$\gamma: \quad \Delta_3 = \Delta_4 \quad . \tag{E.21}$$

This condition holds for exactly one point on the two-dimensional scan, (Δ_4, Δ_1). As expected, this is the intersection of the horizontal asymptote of the modified hyperbola and the five-level asymptote. At this point, the fluorescence of the scan is strongly reduced.

Figure E.12 shows a summary of the contributions of the resonance effects to the features seen in a two-dimensional scan, overlaying all six of the possible coherent effects on the righthand plot in figure E.9. This figure shows that three

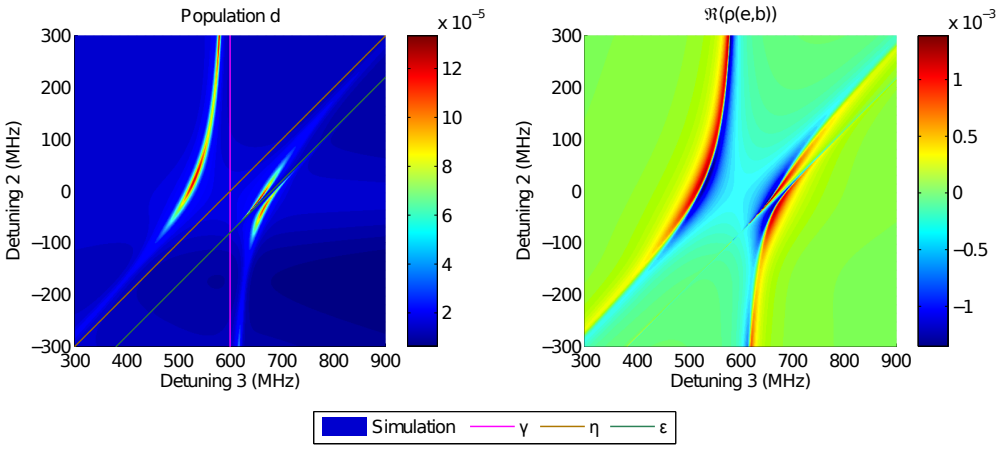

Figure E.13: An inverted modified hyperbola arising from subsets of three, four and five states ordered from the highest states in the chain. The lefthand plot shows a two-dimensional scan of the fluorescence and the righthand plot a scan of the strength of resonance η. This system was tailored to be biased towards the higher states in the chain by setting the intensity of beam 4 to a small value. The intensity of beam 2 was also raised slightly to increase the fluorescence. Most of the population is pumped into state e. The modified hyperbola has asymptotes described by γ, η and ϵ, which involve subsets in the higher states in the chain. The hyperbolic portion of the feature is described most strongly by resonance η. The low fluorescence rate is due to the low power of beam 4; only small amounts of population is excited from state e and fluoresceces.

asymptotes of the modified hyperbola are approximately given by α, ζ and ϵ. In addition, the position of the five-state vertical feature is approximately given by γ. Resonances β and η do not contribute to the dynamics of the modified hyperbola.

E.3.2 A bright resonance structure with asymptote η

The primary feature present in these scans is a bright resonance with three asymptotes describing three, four and five-level coherent behaviour. Each asymptote describes limiting coherent behaviour in a subset of states in the full system. The three, four and five-level asymptotes and subsets are α (a, b, c), ζ (a, b, c, d) and ϵ (a, b, c, d, e). As the number of states involved in the asymptotic behaviour increases, new states are added to the subset in alphabetical order. For example, moving from three-state behaviour to four-state behaviour adds state d to the subset (a, b, c). We have described the states as a chain stretching

from state a to state e, with the states involving in the asymptotic subsets being built up from the base of the chain.

However, this is not the only way to order the coherent subsets. The subsets could instead be ordered from the top of the chain, with states added to each larger subset in *reverse* alphabetical order. This type of ordering occurs only in special cases where the coherences in the higher states in the chain are stronger than the lower states. Such behaviour is unusual; it is much more common for coherences involving the ground state to dominate. This 'reversed' ordering of subsets can most easily be arranged by setting the intensity of beam 4 to a low value. This pumps most of the population into state e. If the beam is weak enough, coherences involving that state e are stronger than those involving the ground state.

Figure E.13 shows a pair of two-dimensional scans showing the fluorescence structure and the strength of the coherence η in such a situation. The result is a modified hyperbola with three asymptotes linking three-state, four-state and five-state behaviour. The states associated with these asymptotes use the reverse ordering system: the three, four and five-level coherent effects are γ (e, d, c), η (e, d, c, b) and ϵ (e, d, c, b, a). Unlike the modified hyperbolae of the more common ordering system, this hyperbola is asymptotic to the coherence η.

E.3.3 Behaviour with multiple pump beams

When analysing the four-state system, we observed the behaviour of increasing the intensities of the two probe beams so that they they were no longer small compared to the pump intensity. In particular, increasing the intensity of beam 1 caused the hyperbola seen for a small intensity to split into two duplicates separated in the vertical direction. This effect can also be seen in a five-state system. In fact, the presence of the additional state e allows another splitting effect to be observed. Figure E.14 shows a pair of scans which demonstrate these

two different types of splitting. The base set of parameters for both of these plots is the set describing the righthand plot in figure E.9.

The lefthand plot in this figure increases the intensity of beam 1 so that it is significantly larger than that of beam 3. The result is very similar to that seen in the four-state system; compare this plot to that shown in the bottom-right of figure E.8. The hyperbola splits into two duplicates translated in the vertical direction due to the same described in section E.2.2. Note that *all* the coherent effects are mirrored. This is because the position of every asymptote of the hyperbola is dependent on Δ_1, and this detuning is split in two.

A similar splitting effect occurs if the intensity of beam 4 is increased instead. The righthand plot shows the behaviour of the system when the intensity of beam 4 is significantly large compared to that of beam 3. In this plot, the original hyperbola is split into two duplicates, although this time the duplicates are separated in the horizontal direction. The mechanism is the same as the previous splitting. Coherent effects in the system see two different values of the detuning of beam 3 depending on whether this beam is red or blue-detuned compared to beam 4. This splits the hyperbola into two duplicates translated along the Δ_3 axis.

The five-state vertical and diagonal resonances γ and ϵ, however, do not move along with the duplicate hyperbolae this time; only the hyperbolic behaviour shifts. The horizontal resonance does not move because it is a dark resonance; its position is independent of the beam powers. The reason the diagonal resonance does not shift is more subtle. In the dressed-state picture, the two frequencies at which a transition is resonantly driven[5] are dominated by the Rabi frequency when the intensity of the beam is high enough; the contribution of the physical detuning of the beam tends to zero. The practical result of this is that the four-level and five-level asymptotes for both duplicate hyperbolae actu-

[5]These frequencies are determined by the values of the detuning in the dressed-state model $\widetilde{\Delta}$. See section E.1.

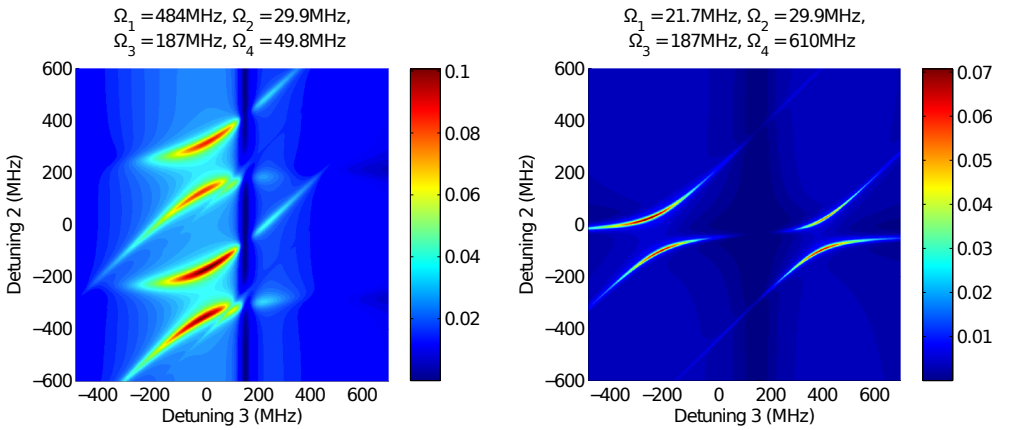

Figure E.14: The appearance of duplicate modified hyperbolae when two beams in the five-state system have a large intensity. The base set of parameters for both of these plots describes the righthand scan in figure E.9. The lefthand plot increases the intensity of beam 1 to a value large compared to beam 3. The original modified hyperbola is split into two duplicates translated in the vertical direction. All three asymptotes of the hyperbola are duplicated. This effect is analogous to the similar situation examined in the four-state system in the bottom-right plot of figure E.8. The righthand plot in this figure shows a similar splitting produced instead by increasing the intensity of beam 4 to a large value. This produces two duplicates which are translated in the horizontal direction. The difference between this case and the previous case is that the features of the modified hyperbola are not all preserved. The horizontal resonance does not move with the duplicates, and the five-level diagonal asymptote also does not apparently move with them. For more details, see text.

ally merge together, so that only a hyperbolic effect is seen. This behaviour does not change any of the resonant features seen in the lefthand scan as the physical detuning of beam 1 only affects the horizontal offset of the bright resonance. Removing this has very little effect on the plot.

In summary, several different resonance effects can be seen in a five-state system. A bright resonance takes the form of a modified hyperbola with three asymptotes. One asymptote is a horizontal line describing a strong three-level coherence, the second is a diagonal line describing a strong four-level coherence and the third a diagonal line describing a strong five-level coherence. The feature looks like a standard hyperbola which is distorted in the region of the third asymptote as the resonance converges onto a new line. In addition, a vertical dark resonance is observed in the three highest states in the state chain due to spontaneous decay on transition 2. When this dark resonance crosses the hor-

izontal asymptote of the modified hyperbola, a pair of dark resonances is observed in the system simultaneously that drive population out of both states b and d.

When the intensity of beam 1 is comparable to that of beam 3, the modified hyperbola describing the bright resonance splits into two duplicates translated in the vertical direction. By contrast, when the intensity of beam 4 is comparable to that of beam 3, the bright resonance splits into two similar curves described by the four-state hyperbola.

E.4 Splitting due to Zeeman states

The previous three sections analysed a series of systems with no Zeeman structure and demonstrated how the number of levels involved in an atom's dynamics contributes to the structure of two-dimensional fluorescence scans. We now examine the modifications that arise when each level is split into multiple Zeeman states.

We return to the simplest system of levels we analysed previously, the three-state system examined in section E.1 and shown in figure E.1. To start off as simply as possible, we split each level into only two states. I will refer to the system with no Zeeman structure as the 'three-state' system and the system with Zeeman structure as the 'three-level' system.

Figure E.15 shows a series of two-dimensional fluorescence scans. The intensities of the two beams are the same for all four plots. The intensity of the pump beam is slightly higher than the scans examined in section E.1 to ensure that all the detail in the three-level plots is clearly visible. The scan in the upper-left is taken in the three-state system. It shows a hyperbolic bright resonance with one asymptote lying on the line $\Delta_1 = 0$ and the other on the line $\Delta_1 = \Delta_2$. The remaining three plots in the figure show two-dimensional scans taken in the three-level system at various polarisations of the pump beam.

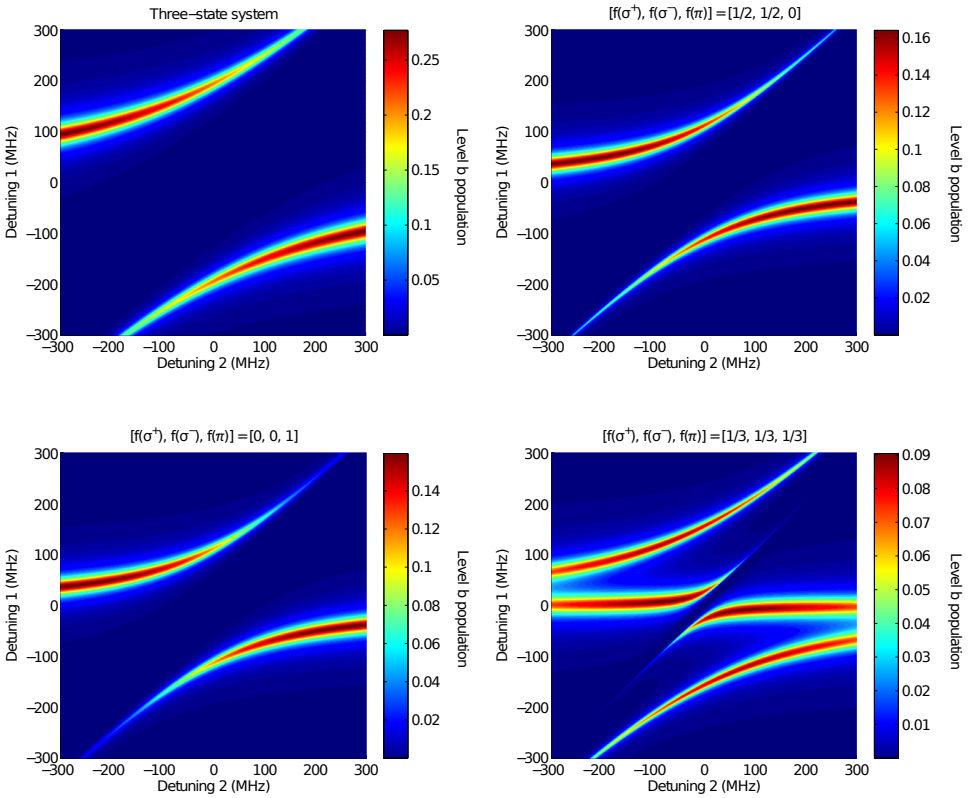

Figure E.15: The splitting of a hyperbola into multiple arcs in the three-level system. The intensities of both beams are the same for all four plots. The upper-left plot shows the behaviour of the three-state system with no Zeeman structure. This shows a hyperbolic bright resonance. The upper-right and lower-left plots show the behaviour of the three-level system subjected to only σ and only π light respectively, and shows two arcs which converge to the same asymptotes as the upper-left plot, but which are separated by a smaller distance. The behaviour in these two plots is almost identical. The lower-right plot shows the behaviour of the three-level system subjected to equal fractions of all three polarisation components. This plot shows *four* arcs, all of which converge to the same asymptotes as the other arcs in the figure. For more details, see text.

The scans in the upper-right and lower-left involve only σ and only π light respectively. The bright resonances seen in these scans are almost identical. They take the form of two arcs which share the same asymptotes as the hyperbola in the three-state system, but the separation between these arcs is significantly less than the three-state hyperbola. We infer that the arcs in the system with Zeeman structure are also hyperbolae. The small separation between the arcs is explained by the difference in Rabi frequencies; these are lower in the system with Zeeman structure because of the angular momentum coupling between the

Zeeman states. This coupling also reduces the efficiency with which population can be transferred to the upper state, therefore reducing the fluorescence of the system with Zeeman structure compared to the three-state system. Compare these two plots to the lower-right plot in figure E.2, which shows a vary similar structure for a plot in the three-state system where the intensity of the pump beam is less by a factor of five.

We confirm that the arcs are hyperbolic by examining the separation between them as a function of the intensity of the pump beam. This is shown in figure E.16. The extent of each arc along the axis of symmetry, and thus the splitting between the arcs, scales as the square root of the intensity. This is exactly what we would expect from a hyperbolic curve and another demonstration of the Autler-Townes effect.

The lower-right plot shows a significant departure from the three-state behaviour. This plot is taken with all three polarisation variables set to the same value. There are four arcs in this scan. All four tend to the same asymptotes as the arcs seen in the other three plots. The arcs fall into two pairs that are symmetric about the line $\Delta_1 = -\tan(5\pi/8)\Delta_2$. There is an inner pair of arcs that lie close to the asymptotes and an outer pair that has a smaller curvature.

It is possible to observe more arcs in the three-level system than in the three-state system. The width of the three-state hyperbola is determined to a very good approximation by the Rabi frequency on the pump transition. In the three-level system, however, there is no single Rabi frequency on this transition; there are in fact up to three different Rabi frequencies on transitions linking the states on the pump transition. These states and the transitions between them are shown in figure E.17. One Rabi frequency is associated with the π transitions A and D, and two separate Rabi frequencies are associated with the two σ transitions B and C.

Each hyperbola in the three-level system can be described by (E.9), which

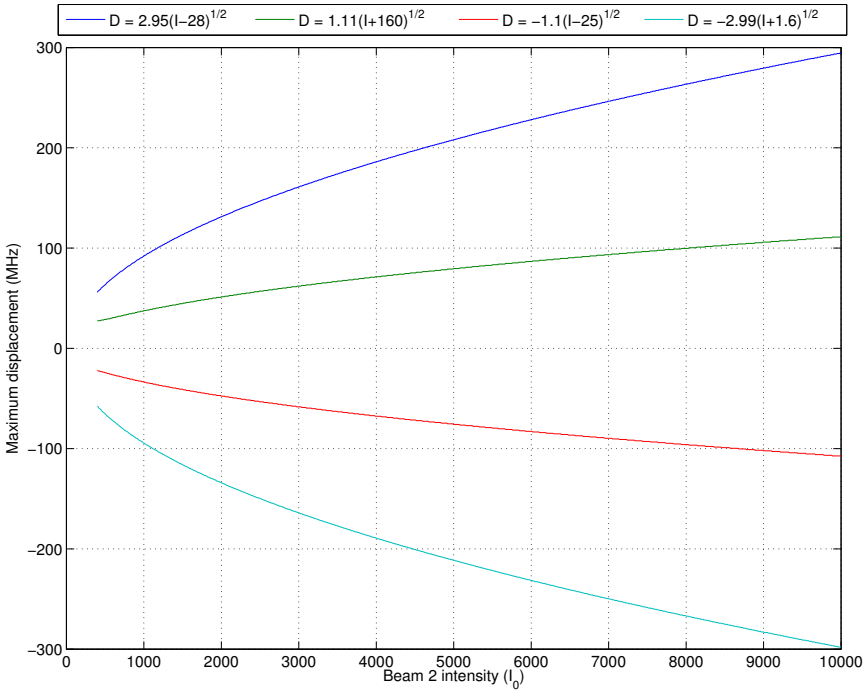

Figure E.16: The extent of each arc along its axis of symmetry as a function of the pump intensity. The polarisation of the pump beam in this case is $[f(\sigma^+), f(\sigma_-), f(\pi)] = [1/2, 1/3, 1/6]$. Each curve is fitted to a square root dependence. Due to ambiguity near to the point where the curves converge, the residuals of each fit do not reach machine precision but instead are of order 1MHz. However, this is still small compared to the extent of each arc. The distances scale as the square-root of the pump intensity, as we would expect from hyperbolic behaviour and an Autler-Townes splitting.

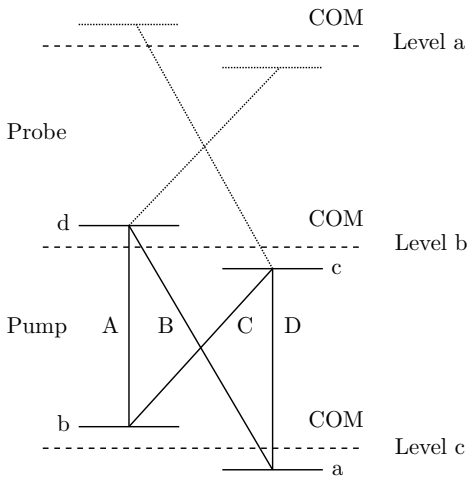

Figure E.17: A schematic diagram of the stimulated transitions in this three-level system. Transitions excited by the pump beam will dominate the structure of the bright resonance, so only these four transitions are labelled. The levels are arranged vertically to show the transitions most clearly. The relative energies of these three levels are shown in figure E.1. COM refers to the centre of mass of each level.

we used previously to describe the three-state hyperbola, with some Rabi frequency Ω. The four states on the pump transition are not independent, and the strengths of all four transitions are directly dependent on one another. This system has a set of frequencies at which it can be resonantly driven, and these frequencies are a function of all the Rabi frequencies connecting the states. The number of possible resonant frequencies is simply given by the total number of states in the levels linked together by the pump beam. In this case, there are four. Each hyperbola is defined by two frequencies; one arc is described by a positive frequency and the other by a negative one. Therefore, the total number of hyperbolae it is possible to see in this system is two.

It is possible to determine these resonant frequencies by calculating the eigenvalues of the Liouvillian, and then choosing those frequencies that correspond to strongly driven population and coherences related to the pump transition. However, we can gain a little more insight into the problem by instead directly looking at the Rabi frequencies on the pump transition. So long as a hyperbola is present in the three-level system, we know that the intensity on the pump transition must be high. As a result, spontaneous behaviour on the transition will strongly dominate over stimulated behaviour. We can therefore describe the time-dependence of the population in levels b and c entirely in terms of the Rabi frequency. If $\mathbf{P}$ is a column vector describing the population of states a to d and $\mathbf{M}$ a vector such that $M_{x,y} = \Omega(x, y)$, then:

$$\dot{\mathbf{P}} = \mathbf{M}\mathbf{P} \tag{E.22}$$

The eigenvalues of $\mathbf{M}$ therefore given the resonant frequencies at which the pump transition can be driven. For this system, $\mathbf{M}$ has the form:

$$\mathbf{M} = \begin{pmatrix} 0 & 0 & \Omega_A & \Omega_B \\ 0 & 0 & -\Omega_C & -\Omega_A \\ \Omega_A & -\Omega_C & 0 & 0 \\ \Omega_B & -\Omega_A & 0 & 0 \end{pmatrix}. \tag{E.23}$$

The two positive eigenvalues of this matrix are:

$$\widetilde{\Omega}_1 = \frac{1}{2}\left(\sqrt{4\Omega_A^2 + (\Omega_B - \Omega_C)^2} + \Omega_B + \Omega_C\right) \quad \text{(E.24)}$$

$$\widetilde{\Omega}_2 = \frac{1}{2}\left|\sqrt{4\Omega_A^2 + (\Omega_B - \Omega_C)^2} - \Omega_B - \Omega_C\right| . \quad \text{(E.25)}$$

When the fractions of the two sigma polarisations are equal, so that $\Omega_B = \Omega_C$, these two expressions simplify to:

$$\widetilde{\Omega}_1 = \Omega_A + \Omega_B \quad \text{(E.26)}$$

$$\widetilde{\Omega}_2 = |\Omega_A - \Omega_B| \quad \text{(E.27)}$$

This has a very simple interpretation. The larger frequency describes the situation where the π and σ transitions drive population through the states in the same sense so that the speed at which population is driven through the system is the sum of the separate speeds of the two types of transition. The smaller frequency meanwhile describes the situation where the two types of transitions drive population in the opposite sense and therefore compete. These expressions also show why the scans with pure σ and pure π light in figure E.15 have identical structure. The resonant frequency of the pump transition is the same for both these cases, producing an identical hyperbola.

Figure E.18 shows a two-dimensional scan taken in the three-level system where the polarisation of the probe beam is formed of unequal fractions of all three polarisation types. The bright resonance takes the form of two pairs of arcs. Two hyperbolae are overlaid on this figure; these hyperbolae are described by the equation (E.9) and using the two Rabi frequencies given in (E.26) and (E.27). They are an extremely good match to the bright resonance feature.

Let us briefly consider a three-level system in which each level directly connected by the pump beam is instead split into four Zeeman states. In this case, we could expect a maximum of four possible hyperbolae to be observed in two-dimensional scans. Figure E.19 shows a system in which the hyperbola splits

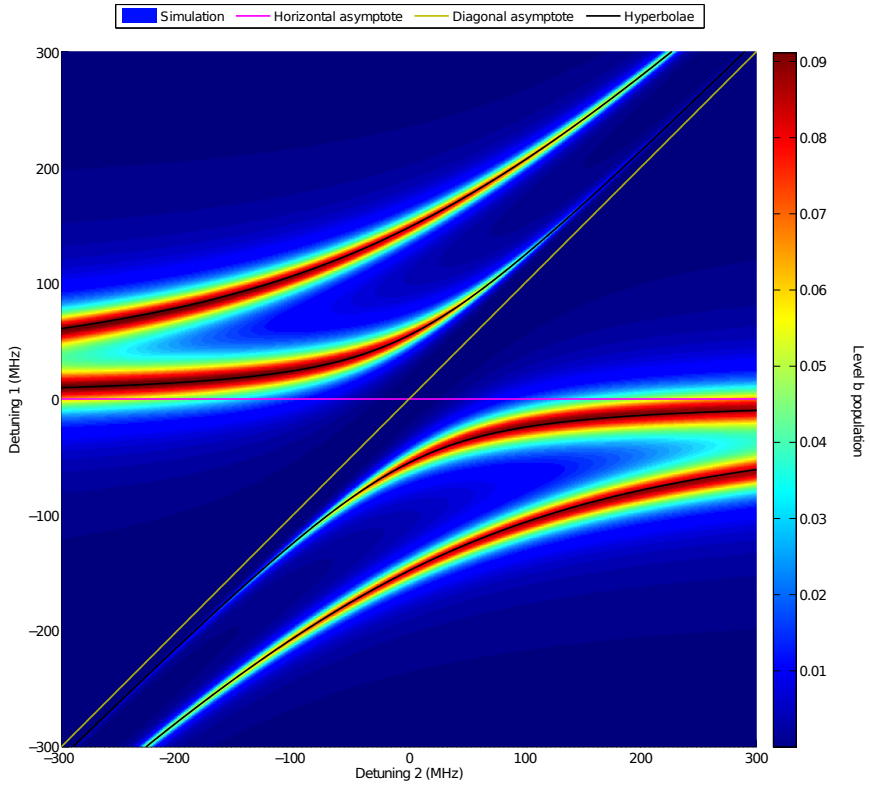

Figure E.18: The hyperbolic behaviour of the three-level split arcs. This plot shows a two-dimensional scan where the polarisation of the pump beam is given by $[f(\sigma^+), f(\sigma_-), f(\pi)] = [1/2, 1/3, 1/6]$. The four arcs are overlaid by hyperbolae. The positions of the hyperbolae match the arcs very well indeed, but these hyperbolae are not fits. The Rabi frequencies of each hyperbola were calculated directly from E.24 and E.25.

into three. This type of behaviour forms part of the structure seen in the series of experimental two-dimensional scans examined in section 5.5.

Finally, let us consider a system in which the pump beam connects two levels involving many Zeeman states. For the equivalent three-level system in $^{43}Ca^+$, the total number of states in the two levels connected to the probe transition is 48. As a result, we could expect there to be 24 hyperbolae visible in a two-dimensional plot in this system. This would produce a very complicated structure! A two-dimensional scan taken in $^{43}Ca^+$ is shown in figure E.20 along with the temperature of the ion. Many of the hyperbolae merge together due to their finite width, so only a handful of hyperbolae are visible in the scan. However, this gives a glimpse of the sort of structure that could be seen from this ion. The

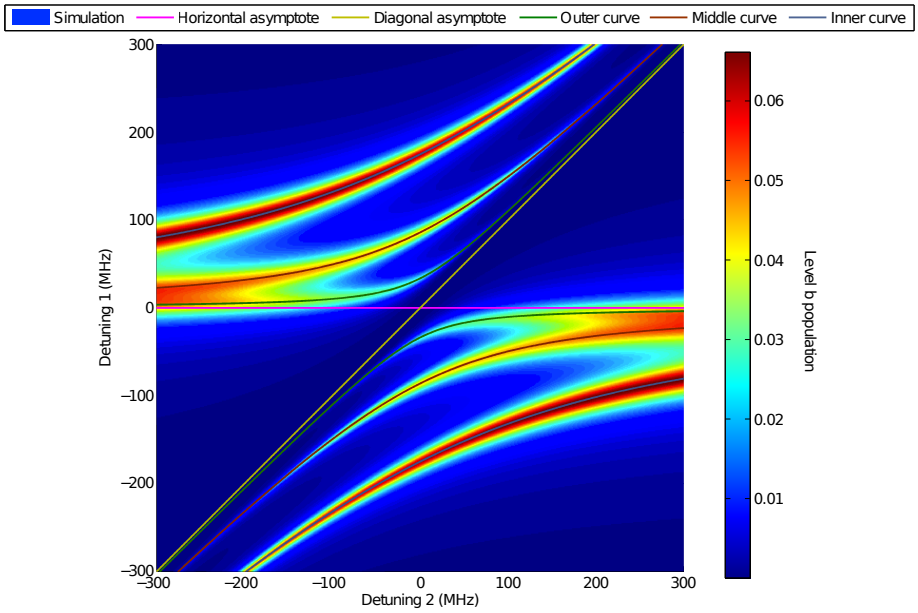

Figure E.19: Three hyperbolae seen in a three-level system where each level is split into four Zeeman states. The pump beam has a polarisation given by $[f(\sigma^+), f(\sigma^-), f(\pi)] = [1/3, 1/3, 1/3]$ in order to emphasise all the arcs. The bright resonances are overlaid by hyperbolae with resonant frequencies calculated from the eigenvalues of $\mathbf{M}$. Two eigenvalues are very close in value, producing two curves that merge to form the 'outer' curve. The outer hyperbola shown was produced using the mean value of the resonant frequency of these two eigenvalues. This type of splitting is a factor contributing to the structure of some of the experimental scans examined in chapter 5.

number of hyperbolae is seen much more clearly in the temperature plot. The ion is actually cold for a large portion of the scan. The ion heats in the region of the hyperbolae and around the dark resonances. In addition, the ion heats when the 397nm beam is at a higher detuning than approximately the position of the centre of mass. In these portions of the scan, the beam will cause heating that cannot be countered by the dynamic effect in the response on the 866nm beam. This is similar to the 397nm scan shown in $^{40}Ca^+$ in figure 4.4. However, the ion is cold for over half the scan, so portions of all the arcs could still be observed. It may therefore be possible to take data of this sort for this ion.

In general, we expect the hyperbolae seen in systems with no Zeeman structure to split into a series of hyperbolae when each level is split into multiple Zeeman states. The maximum number of split arcs is equal to half of total num-

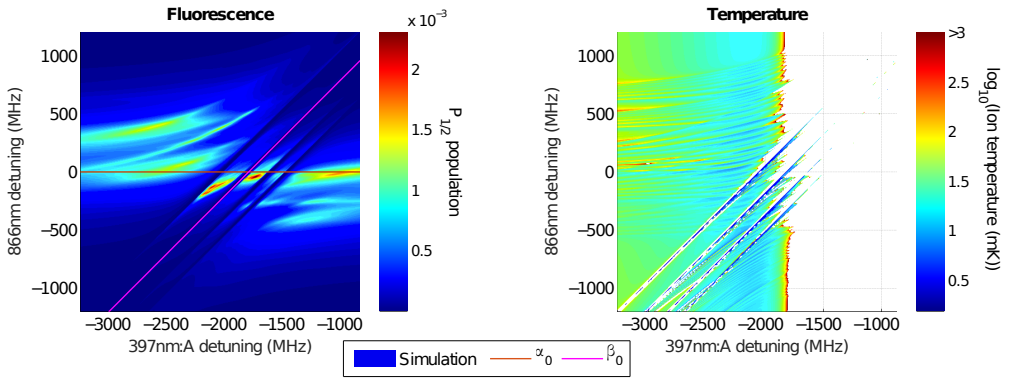

Figure E.20: The appearance of very many hyperbolae in two-dimensional scans in $^{43}Ca^{+}$. These two plots were taken in a three-level system in $^{43}Ca^{+}$ where the intensity of beam 1, the 397nm beam, is much greater than the intensity of beam 2, the 866nm beam. The lefthand plot shows the total population in the highest-energy level, the $P_{1/2}$ (equivalent to level b in figure E.17), and the righthand plot shows the temperature of the ion calculated using the dynamic method with a trap frequency of 816kHz. The horizontal axis shows the detuning of the lower-frequency 397nm beam. See figure 6.1. A large number of arcs can be seen, especially in the temperature plot. We would expect a maximum number of 24 hyperbolae to be seen in this system. The lines α_0 and β_0 show the positions of the base asymptotic lines that pass through the centres of mass of the $P_{1/2}$ level and the $\widetilde{F} = 3$ manifold of the $S_{1/2}$ level. Each hyperbola will converge to slightly different asymptotes due to the detunings of the states most strongly involved in that particular part of the bright resonance. When the temperature of the ion is not shown, the dynamic method predicts that the ion will heat. When the 397nm beam has a detuning approximately greater than the position of the centre of mass, the large intensity means that dynamic effect cannot counteract the 397nm heating with the response of the 866nm beam. Note that the ion is cold for over half of the scan, implying that it would be possible in principle of see half of the hyperbolic behaviour in an experimental scan.

ber of states in the levels directly connected by the pump transition. The Rabi frequencies associated with the separation of these arcs are given by the eigenvalues of the Rabi-frequency matrix $\mathbf{M}$.

F

Residuals for $^{43}Ca^{+}$ scans

> I particularly liked how they made Davros feel slightly like an Oxford academic.
>
> James Cave

This appendix contains plots of the data, high-resolution fit and residuals for the four 866nm scans shown in figure 6.28 and the 866nm scan shown in figure 6.31.

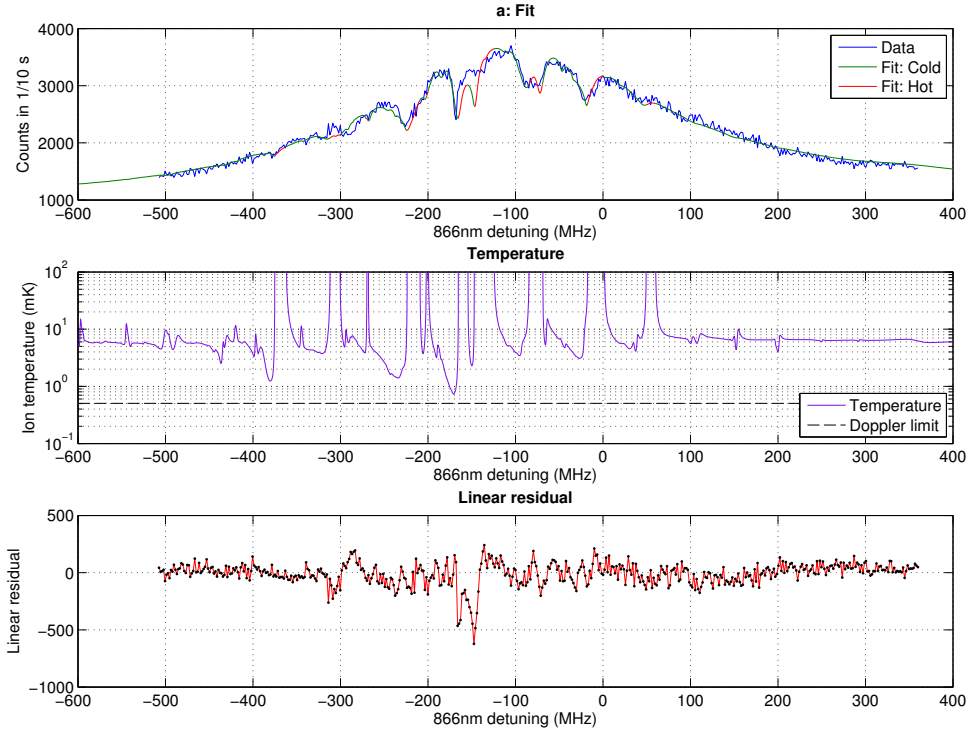

Figure F.1: The data, high-resolution fit and linear residuals of curve a in figure 6.28.

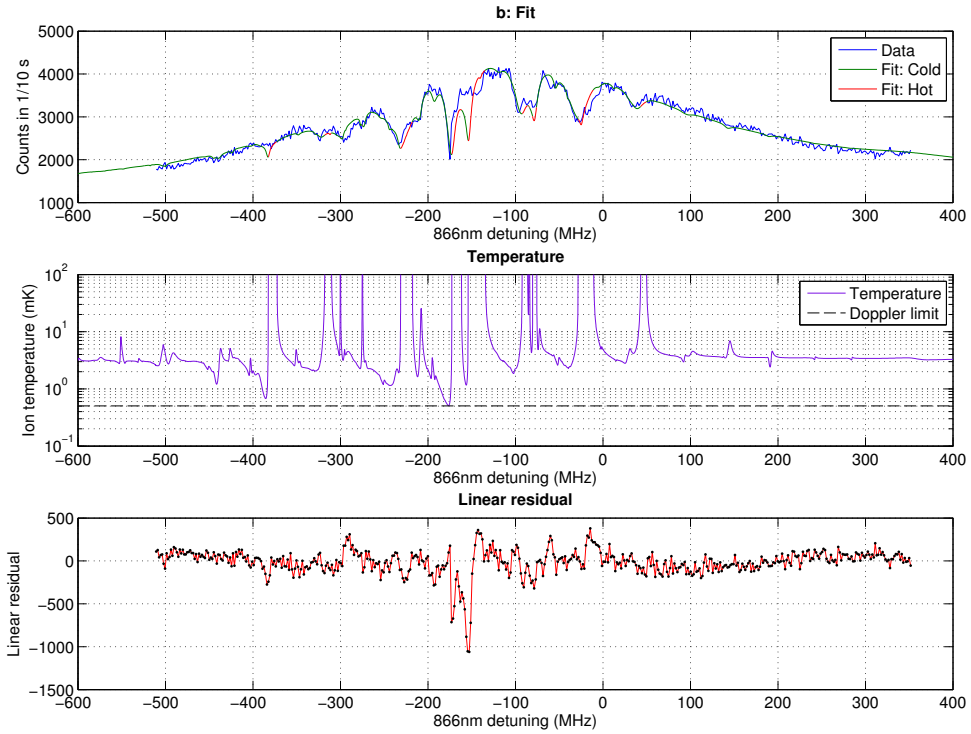

Figure F.2: The data, high-resolution fit and linear residuals of curve b in figure 6.28.

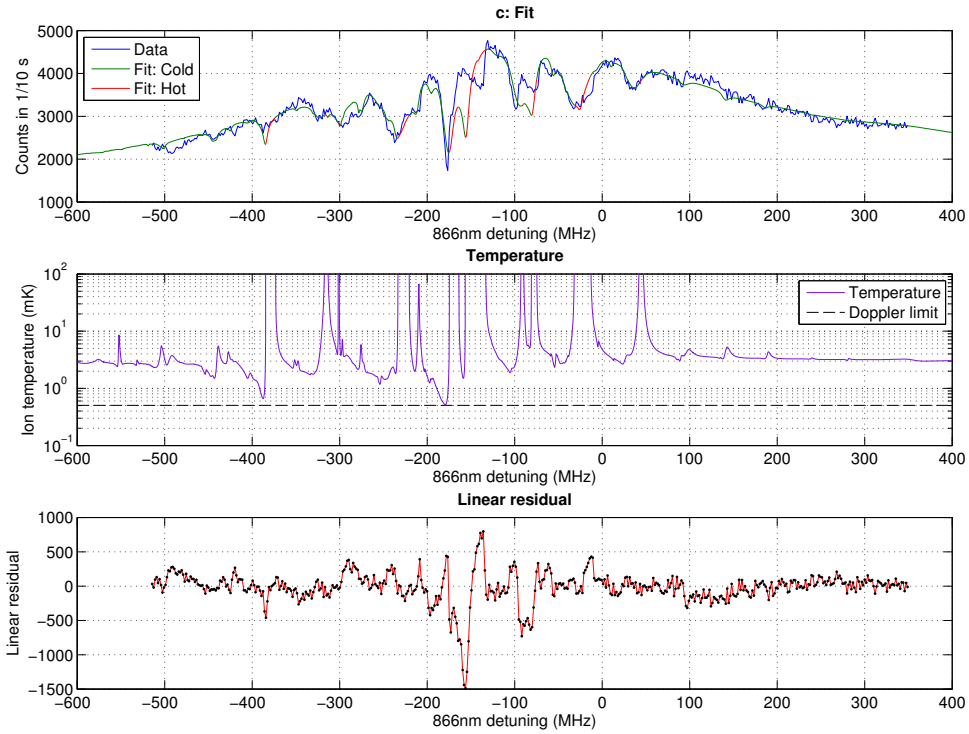

Figure F.3: The data, high-resolution fit and linear residuals of curve c in figure 6.28.

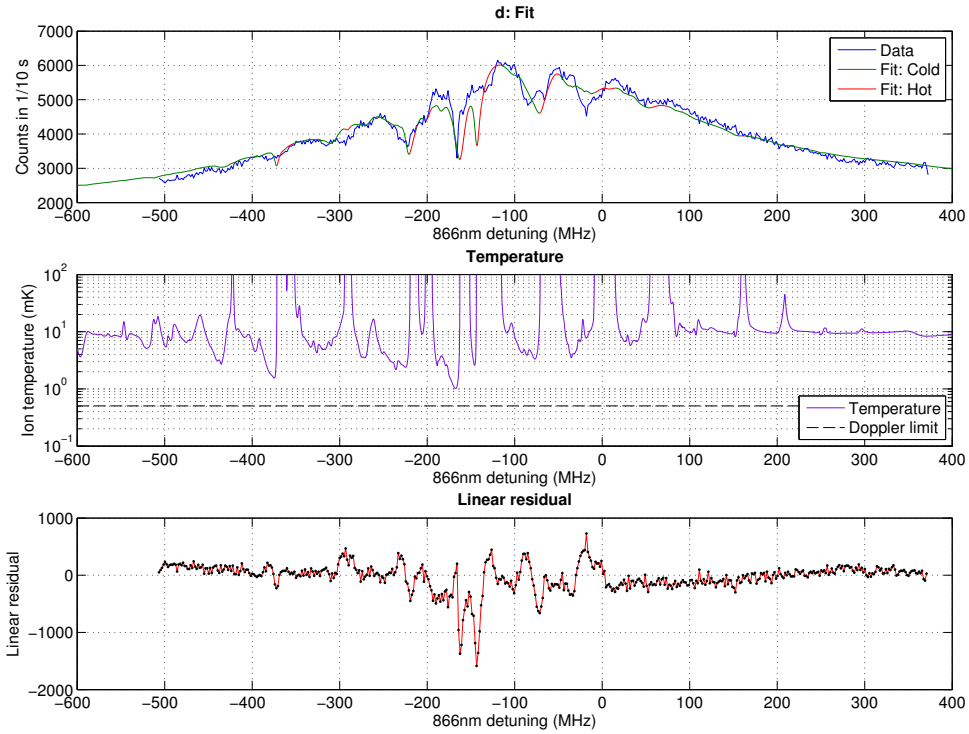

Figure F.4: The data, high-resolution fit and linear residuals of curve d in figure 6.28.

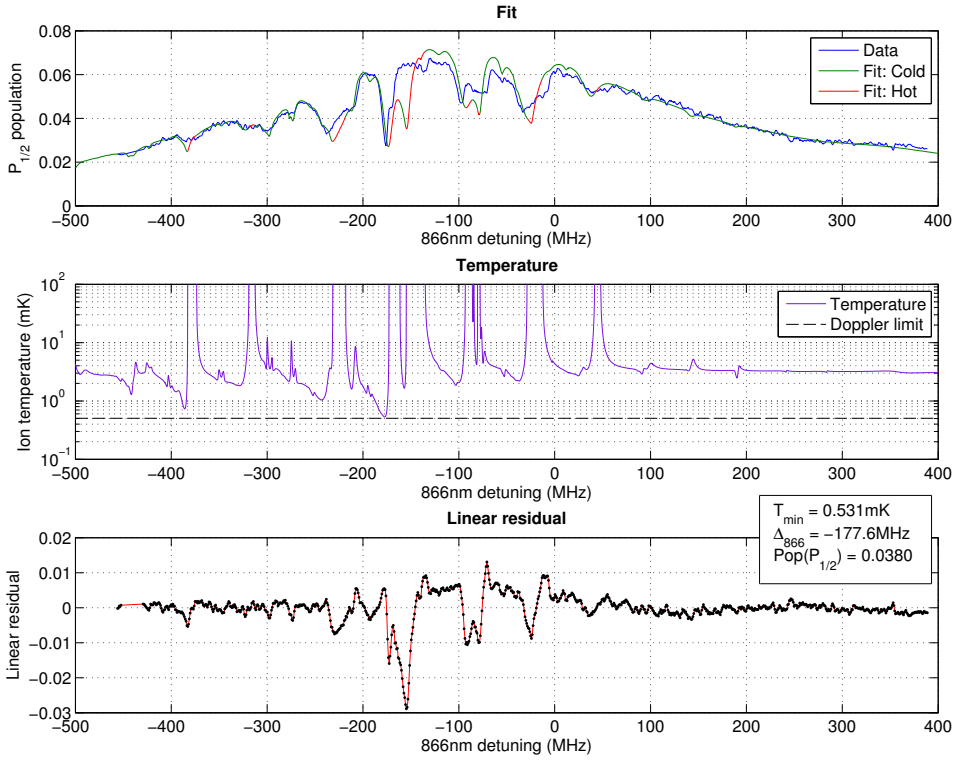

Figure F.5: The data, high-resolution fit and linear residuals of the 866nm scan shown in figure 6.31.

G

Initialisation and Readout

> It looks like a pair of scrolls, amazingly long, spiraling into darkness. The scrolls are twisted around each other. But instead of writing on paper, the scrolls are the message themselves, a long message in an unfamiliar language, in an alphabet of only four letters.
>
> *Geneforge* by Spiderweb Software

In chapter 6, we discussed the choice of a qubit in $^{43}Ca^{+}$ and techniques used to cool the ion. There are two further important aspects of such a qubit that we have analysed using the Bloch equations routines: qubit initialisation and readout. First, we discuss initialisation.

G.1 Initialisation

The 146.09 gauss qubit is initialised using a two-step process. The states of this qubit are hard to initialise directly with a high fidelity using laser light. Instead, we optically pump population into one end of the $S_{1/2}$ manifold using a combination of 397nm light with a polarisation that is nominally comprised of a pure σ component and a series of corrective microwave pulses. We then transfer population from this state to the qubit states using a series of microwave π pulses. The cooling scheme based on the dark resonance α involves 397nm light with a strong σ^{+} component and nominally no σ^{-} component. Under this scheme,

population is concentrated at the positive-M end of the $S_{1/2}$ manifolds. To ensure that initialisation is as fast as possible, we therefore optically pump the ion to the same end of the manifolds using pure σ^+ light. The prepared state for the optical pumping process is $|S_{1/2}, M = +4\rangle$. The microwave pulses used to transfer population to the qubit states are shown in figure 6.7. In this analysis, we consider only the optical pumping caused by the laser beams; the microwave pulses used to enhance the fidelity of this process are discussed in [Har13].

Figure G.1 shows an analysis of the fidelity of the optical pumping process as a function of time and as a function of the polarisation error in the 397nm beam. We are interested in the timescale in which a random input state approaches the steady-state solution. This is a problem that is very well suited to the routine SILVER IMPER described in section 3.4.4. In addition, the timescales of this integration problem are of the order of hundreds of microseconds. SILVER IMPER can calculate to these evaluation times very efficiently and significantly faster that the other time-dependent routines described in section 3.4 (e.g. see figure 3.19). The evolution of a random input state as function of time is shown in the upper plot of figure G.1. A state can be optically pumped to the prepared state with fidelities $\mathcal{F} = 99\%$ in 180μs and $\mathcal{F} = 99.9\%$ in 330μs. The evolution quickly converges to the steady state for pulse lengths greater than approximately 600μs.

The dependence of the fidelity of the process on the polarisation error can be determined very efficiently by directly calculating the steady state solution using matrix inversion. This calculation describes the limiting fidelity that can be reached for an optical pumping pulse with a duration $t \to \infty$. Matrix inversion is a much faster process than numerical integration. The dependence of the state preparation error on the fraction of σ^- light in the 397nm beam is shown in the lower plot. The behaviour is linear to a very good approximation for σ^- fractions less than approximately 0.02. At low polarisation errors, the σ^- fraction drives a small amount of population to the neighbouring states in the

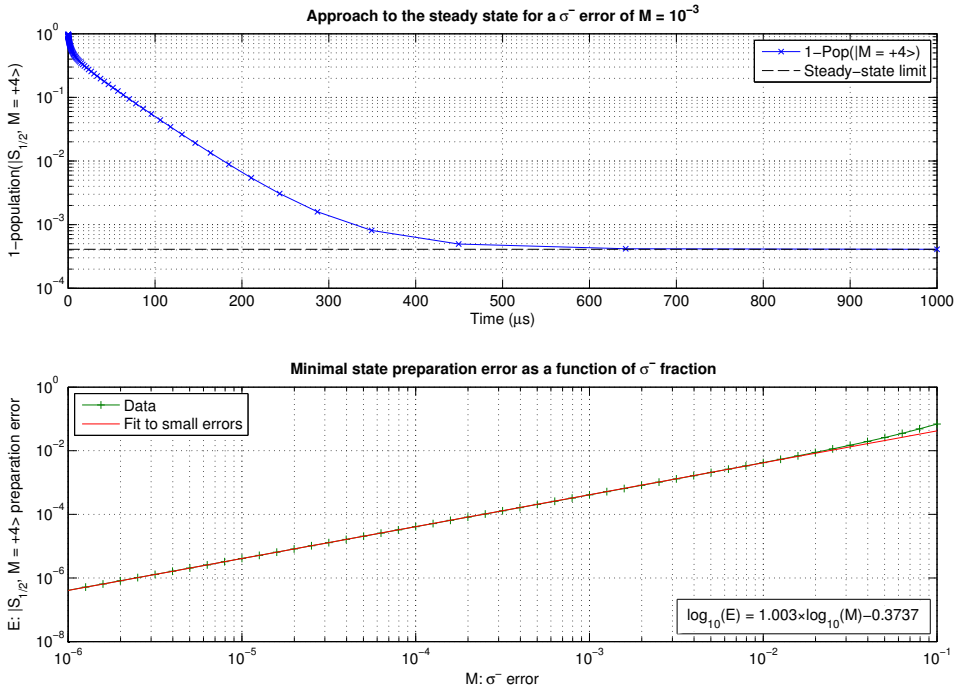

Figure G.1: The errors of a state preparation process that optically pumps population to the prepared state $|\mathrm{S}_{1/2}, M = +4\rangle$ as a function of time and polarisation error. The upper plot shows the preparation error as a function of time for a random initial state. The time dependence was calculated using SILVER IMPER. The steady state error is shown by the horizontal black dashed line. The initial state is pumped to the prepared state with fidelities $\mathcal{F} = 99\%$ in 180μs and $\mathcal{F} = 99.9\%$ in 330μs. There is asymptotic convergence to the steady state, the difference being negligible for $t \gtrapprox 600$μs. The lower plot shows the state preparation as a function of the polarisation error of the 397nm beam. The x axis shows the fraction of σ^- light present in the beam. The behaviour is linear for polarisation errors of less than approximately 0.02. For a pessimistic value of the polarisation error of 10^{-3}, the limiting error of the optical pumping process is 4×10^{-4}.

$\mathrm{S}_{1/2}$. At high errors, the fraction drives population into more states in the $\mathrm{S}_{1/2}$ and the behaviour deviates from a linear trend. Experimentally, with the technology available, we can certainly achieve a polarisation impurity of 10^{-3}. For this impurity, the limiting optical pumping error is $\epsilon = 4 \times 10^{-4}$.

This error is already very low. It can be decreased by using an iterative corrective process. At the end of the optical pumping pulse, the population of the states in both $\mathrm{S}_{1/2}$ manifolds decreases exponentially with decreasing M. The most heavily populated states besides the prepared state are those with $M = +2, +3$. Population in the $\widetilde{F} = 3$ manifold is transferred to the $\widetilde{F} = 4$ manifold by applying a short pulse that consists of only the 397nm:A and 866nm

beams. Population in the $|\mathrm{S}_{1/2}, \widetilde{F} = 4, M = +2\rangle$ and $|\mathrm{S}_{1/2}, \widetilde{F} = 4, M = +3\rangle$ is then transferred to the equivalent M states in the $\widetilde{F} = 3$ manifold by using two microwave π pulses. By applying another pulse of the 397nm:A and 866nm beams, this population is transferred back to the $\widetilde{F} = 4$ manifold, with states of higher M being preferentially populated. By repeating this series of pulses, population is transferred from the states with $M = +2, +3$ to the prepared state. This reduces the error in the optical pumping by approximately an order of magnitude. For more information, see [Har13].

G.2 Readout

The process used to detect the qubit state is an iterative process involving three separate pulses. This scheme is shown schematically in figure G.2. By making use of microwaves and an iterative series of three laser pulses, population in the lower qubit state $|\mathrm{S}_{1/2}, \widetilde{F} = 4, M = 0\rangle$ is selectively shelved in the $\mathrm{D}_{5/2}$ level. By then turning on the 397nm and 866nm Doppler beams, the state of the ion can be detected; only an ion that has not been shelved in the $\mathrm{D}_{5/2}$ will fluoresce, and so therefore only population in the lower-energy qubit state will produce a signal.

In more detail, population in the lower-energy qubit state is transferred to the state $|\mathrm{S}_{1/2}, M = +4\rangle$ using a series of microwave π pulses. A series of optical pulses is then used to transfer population in this state to the shelf level. This pulse sequence is:

$$\left(393\mathrm{nm}\sigma^+, 850\mathrm{nm}\sigma^+, 850\mathrm{nm}\pi\right)^n, 393\mathrm{nm}\sigma^+ \tag{G.1}$$

where n is an integer. The first pulse in the series is a 393nm:B σ^+ pulse. This beam is defined in figure 6.1. In this sequence, it transfers population from the $|\mathrm{S}_{1/2}, M = +4\rangle$ to the $|\mathrm{P}_{3/2}, M = +5\rangle$ state. From this state, population can decay either to the shelf, back to the $|\mathrm{S}_{1/2}, M = +4\rangle$, or to the three states in

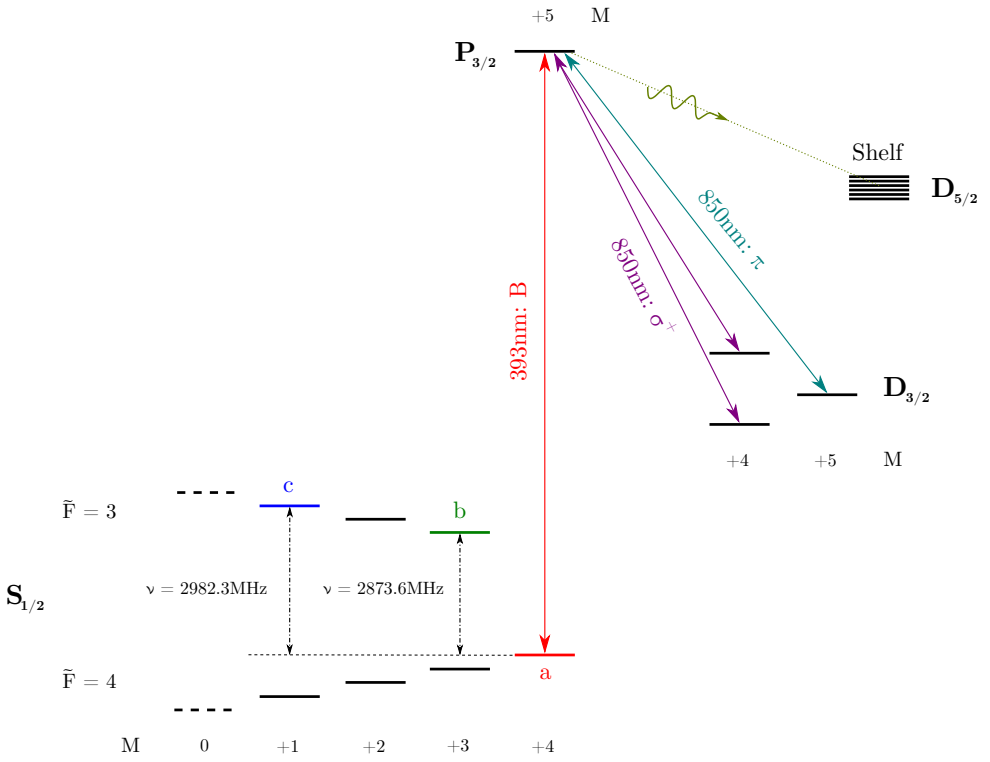

Figure G.2: A schematic diagram of the readout process used to detect the state of the hyperfine qubit. Population in the lower state of the qubit $|\mathrm{S}_{1/2}, \widetilde{F} = 4, M = 0\rangle$ is transferred to state a, shown in red, by using a series of microwave π pulses. The readout scheme selectively pumps population in this state into the shelf level $\mathrm{D}_{5/2}$. The state of the ion can be determined after the series of readout pulses by applying the 397nm and 866nm Doppler cooling beams. Fluorescence will only be observed if the ion is not shelved, and therefore a signal will be detected only if the initial state of the qubit contains population in the higher-energy state c. For more information about the pulse sequence, see text.

the $\mathrm{D}_{3/2}$ with $M = +4, +5$. Population is repumped from these latter states by using an 850nm σ^+ pulse to excite population from the $M = +4$ states followed by an 850nm π pulse to excite population from the $M = +5$ state. These three pulses are repeated n times to increase the probability that the ion will be shelved. Finally, a 393nm:B σ^+ pulse transfers any remaining population from the $|\mathrm{S}_{1/2}, M = +4\rangle$ to the shelf. These pulses are applied separately to avoid a similar problem to that described in section 6.5; if all three beams were applied simultaneously, resonant behaviour between the beams would reduce the amount of population that decays to the shelf. For more information about this scheme, see [Szw09, All11, Har13].

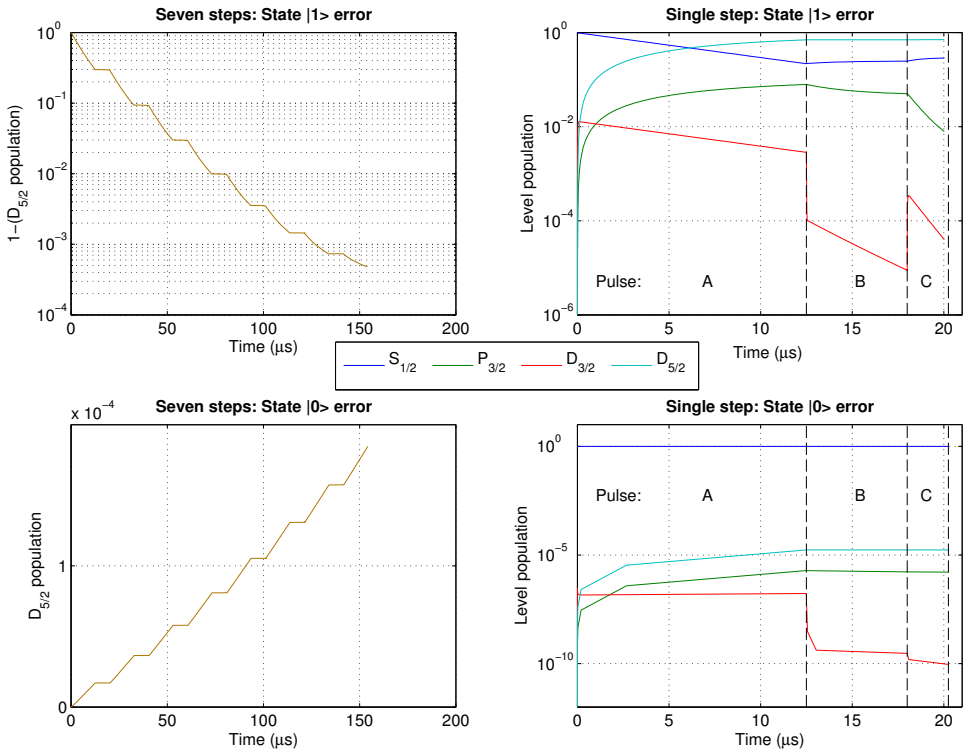

Figure G.3: A investigation of the readout process for the qubit pair of states. The upper plots show the behaviour of a system prepared in state $|1\rangle$ as a function of time and the lower plots show the behaviour for a system prepared in state $|0\rangle$. For both pairs of states, $|1\rangle$ is taken as $|\mathrm{S}_{1/2}, M = +4\rangle$. The lefthand plots show the readout errors for both states. This error is defined as the probability of incorrectly determining the state of the qubit. For a system in state $|1\rangle$, this is the probability of not shelving the qubit. For a system in state $|0\rangle$, the error is the probability of shelving the qubit. The righthand plots show the populations in the four levels as a function of time for the first set of three pulses. During pulse A, population is pumped out of the $\mathrm{S}_{1/2}$ into all the other levels. During pulses B and C, population is repumped from the $\mathrm{D}_{3/2}$ to the $\mathrm{D}_{5/2}$ and the $\mathrm{S}_{1/2}$. The population in the $\mathrm{D}_{5/2}$ shelf increases throughout all three pulses.

The behaviour of the readout process was analysed for two different pairs of states. The first is a pair of states that is simple to prepare and was used as a diagnostic of the readout process: $|\mathrm{S}_{1/2}, \widetilde{F} = 3, M = +3\rangle$ and $|\mathrm{S}_{1/2}, M = +4\rangle$. I refer to these states as the 'stretch' pair. The second consists of states $|\mathrm{S}_{1/2}, \widetilde{F} = 3, M = +1\rangle$ and $|\mathrm{S}_{1/2}, M = +4\rangle$. I refer to this pair of states as the 'qubit' pair. It is formed after population from the lower-energy state of the computational qubit is transferred to the $|\mathrm{S}_{1/2}, M = +4\rangle$ using a series of microwave pulses. The behaviour of the readout process for pulses with perfect polarisations and with polarisation errors was examined.

The readout error for any state is defined as the probability of receiving an incorrect signal if the system is prepared initially in that state. Figure G.3 shows the error of the readout process for both qubit states using the qubit pair. The set of three pulses was repeated $n = 7$ times. This number was chosen as a compromise between the time of the readout sequence and the fidelity of the operation. The behaviour of the ion was calculated using SILVER IMPER. State $|1\rangle$ is defined as the state which is to be shelved, and state $|0\rangle$ as the state which should not be shelved. For both the stretch and qubit pairs, state $|1\rangle$ is $|\mathrm{S}_{1/2}, M = +4\rangle$. The upper plots show the behaviour of a system prepared in state $|1\rangle$ as a function of time. The lower plots show the behaviour of a system prepared in state $|0\rangle$. For a system in state $|1\rangle$, this is the probability of not shelving the qubit, while for a system in the other state, the error is the probability of shelving the qubit. The lefthand plots show the readout error as a function of time. The righthand plots show the populations of the four levels for the three pulses of the first iteration. The populations of the levels change abruptly between pulses. This is because the initial behaviour of an ion when encountering a new pulse occurs on the timescale of spontaneous effects; this timescale is much shorter than the long duration of all three pulses. The 393nm:A pulse transfers population from the prepared state to the $\mathrm{D}_{5/2}$ shelf. Population that decays to the $\mathrm{D}_{3/2}$ is repumped during the two 850nm pulses. The $\mathrm{D}_{5/2}$ population increases during all three pulses.

There are two errors that are important in defining the readout fidelity: the errors in reading out states $|0\rangle$ and $|1\rangle$. We define these errors as $\epsilon(0)$ and $\epsilon(1)$ respectively. If we assume that it is equally likely that the qubit will be prepared in either state, then the mean error of the readout process is:

$$\bar{\epsilon} = \frac{\epsilon(0) + \epsilon(1)}{2} \tag{G.2}$$

The fidelity is simply calculated from this mean error:

$$\mathcal{F} = 1 - \bar{\epsilon} \tag{G.3}$$

Trial	Method	$\epsilon(0)$	$\epsilon(1)$	$\bar{\epsilon}$	$\mathcal{F}$
Stretch, pure	Bloch	3.75×10^{-4}	2.31×10^{-4}	3.03×10^{-4}	99.970%
	Rate	4.24×10^{-4}	1.22×10^{-4}	2.73×10^{-4}	99.973%
Stretch, impure	Bloch	3.73×10^{-4}	5.40×10^{-4}	4.57×10^{-4}	99.954%
	Rate	4.22×10^{-4}	5.14×10^{-4}	4.68×10^{-4}	99.953%
Qubit, pure	Bloch	1.85×10^{-4}	1.30×10^{-4}	1.57×10^{-4}	99.984%
	Rate	2.07×10^{-4}	0.926×10^{-4}	1.50×10^{-4}	99.985%
Qubit, impure	Bloch	1.84×10^{-4}	4.23×10^{-4}	3.03×10^{-4}	99.970%
	Rate	2.06×10^{-4}	4.66×10^{-4}	3.36×10^{-4}	99.970%

Table G.1: A comparison of the state errors, mean error and fidelity of the readout process defined in (G.1) for $n = 7$ pulses using the Bloch equations and a rate equations approach. The fidelities for the stretch and qubit pairs of states are shown. For each pair of states, the fidelity is calculated for pure polarisations and for impure polarisations. The polarisation impurities used to calculate these values are shown in table G.2. The predictions of the Bloch equations and rate equations are very similar, differing only by of order 10%.

Table G.1 shows a comparison of the errors and fidelities calculated by my Bloch equations routines and using a rate equations approach. Four situations are modelled. The stretch pair of states and the qubit pair of states are both modelled using pure polarisations and impure polarisations. The polarisation impurities used to calculate these values are shown in table G.2. Note that the fidelities calculated using the Bloch equations are always lower than using the rate equations for pure polarisations and higher for impure polarisations. The coherences that are present act to slightly reduce the amount of population each beam will excite from each level. This reduces the effectiveness of the readout process for pure polarisations. However, it also reduces the errors polarisation impurities introduce into the system. The combined effect is that singe-level coherences improve the fidelity of a readout process with polarisation impurities. The predicted errors in the readout process are a good match to the inferred errors in experiment.

The predictions of the Bloch equations and the rate equations agree to within of order 10% for all four cases. This is as we would expect. This readout process only ever has one beam shining on the ion at a time. Coherent effects therefore

Beam	σ^- error	π error	σ^+ error
393nm σ^+	0.5%	0.5%	-
850nm σ^+	0.5%	0.5%	-
850nm π	3%	-	-

Table G.2: The polarisation errors for the 'impure' trials used to compare the predictions of the readout fidelity using the optical Bloch equations and a rate equations approach. The results of these trials are shown in table G.1.

only occur between states in the same level and will be very weak. Calculating the fidelities using a rate equations approach takes at least an order of magnitude less time than with the Bloch equations so is to be preferred if a 10% error is acceptable. Readout from the qubit state is predicted by the Bloch equations to be 1.6×10^{-4}. The combined readout error in experiment, including the microwave pulse sequence, is inferred to be $\approx 5 \times 10^{-4}$ [HAB+14].

H

Simulation Parameters

> I wondered if I would ever be able to say with Mr. Valiant-for-Truth, out of Bunyan's *Pilgrims Progress*: "With great difficulty I am got hither, yet now do I not repent me of all the trouble I have been at to arrive at where I am".
>
> Dray Prescot
> *The Suns of Scorpio* edited by Alan Burt Akers

This appendix lists the parameters used to produce the figures presented in this thesis. Table H.1 shows variables describing the five energy levels in $^{40}Ca^{+}$ and $^{43}Ca^{+}$ and table H.2 shows variables describing the transitions between them. Table H.3 details the constants used in calculating the energy states and simulating ion's motion.

The long table H.4 lists all other parameters specifically relating to each figure. The figures are listed in the order of their appearance. Due to the length of this table and the different number of parameters related to different figures, the table is organised in order to preserve space. The figures are grouped into sets of four, and the portion of the table describing each set will not necessarily begin at the top of each page. The figure numbers are highlighted in bold for clarity. Some figures involve large numbers of parameters. These are organised into sets that in turn contain the beam parameters, the calibration parameters, the fitting variables and information about the data files. The parameter sets are separated from each other by double horizontal lines.

Level	A_J (MHz)	B_J (MHz)
$S_{1/2}$	-806.40207160[ABG+94]	0
$P_{1/2}$	-145.4[NBI+98]	0
$P_{3/2}$	-31.0[NBI+98]	-6.9[NBI+98]
$D_{3/2}$	-47.3[NBI+98]	-3.7[NBI+98]
$D_{5/2}$	-3.8931[BKR+07a, BKR+07b]	-4.241[BKR+07a, BKR+07b]

Table H.1: The hyperfine coefficients used to calculate the energy states of the five lowest-energy levels in Ca^+.

Sets of parameters that describe each laser beam appear first. The header for these is the wavelength of the appropriate beam. There are 9 possible beam parameters. I gives the intensity of the beam in saturation intensities (defined in (2.67)). For beams which make use of a carrier and a sideband, I_A and I_B refer to the intensity of each component. For the 397nm and 393nm beams, A and B are defined in figure 6.1. For the 850nm beam, the two components are shown in figure G.2. In both cases, beam A is taken to have a lower frequency than beam B. $f(\sigma^+)$, $f(\pi)$ and $f(\sigma^-)$ give the fraction of the intensity of the beam in the σ^+, π and σ^- polarisation components respectively. Δ gives the detuning of the beam. Beams using a carrier and sideband have detunings described by Δ_A and Δ_B and the frequency difference between the two beams is given by 'EOM'. Γ_L gives the linewidth.

The header 'Calibrations' is used to describe parameters that relate experiment to simulations. There are 36 possible parameters in this category. $|\mathbf{B}|$ gives the magnitude of the applied $\mathbf{B}$ field. The DAC calibration is described by three variables. C_{Err} gives the relative difference between the fitted calibration, C, and the nominal experimental calibration, C_0. The error is defined as:

$$C = C_0(1 + C_{\mathrm{Err}}) \tag{H.1}$$

To distinguish between different calibrations, these variables are listed along with the wavelength of the relevant beam.

Three parameters describe the conversion of 850nm intensity from beam

Transition	A ($\times 10^6 s^{-1}$)	λ (nm)
397nm	132[Szw09]	396.9589788[WvdBG$^+$08]
866nm	8.4[Szw09]	866.452[ER56, YRtNAT08]
393nm	135.0[GKZ$^+$08]	393.4774716[WvdBUE09]
850nm	0.955[GKZ$^+$08]	850.036[ER56, YRtNAT08]
854nm	8.48[GKZ$^+$08]	854.444[ER56, YRtNAT08]
732nm	0.850[KBL$^+$05]	732.591[ER56, YRtNAT08]
729nm	0.856[BDL$^+$00]	729.347276793955[CBK$^+$09]

Table H.2: The Einstein A-coefficient and wavelength used in calculations involving transitions between the five lowest-energy levels in Ca^+. The wavelengths have been converted from frequency units.

powers. The fitted relative difference, fitted value and nominal experimental conversion are given by $C_{I,\mathrm{Err}}$, C_I and $C_{I,0}$ respectively. These are defined using:

$$C_I = C_{I,0}\,(1 + C_{I,\mathrm{Err}}) \tag{H.2}$$

which is a similar expression to the DAC conversion.

The beam polarisation is altered by varying the angle of a quarter-wave plate in the beam. The calibration of the relationship between this angle and the beam polarisation is described by three variables in a similar manner to the DAC conversion. $\mathrm{Pol}_{\mathrm{Err}}$ gives the relative difference between the nominal experimental calibration, Pol_0, and the fitted calibration, Pol. These are defined as:

$$\mathrm{Pol} = \mathrm{Pol}_0\,(1 + \mathrm{Pol}_{\mathrm{Err}}) \tag{H.3}$$

Similarly, the angle where $f(\sigma^-) = 0$ is described by three variables using:

$$\mathrm{Pol}_Z = \mathrm{Pol}_{Z,0}\,(1 + \mathrm{Pol}_{Z,\mathrm{Err}}) \tag{H.4}$$

where $\mathrm{Pol}_{Z,\mathrm{Err}}$ is the relative difference between the nominal experimental calibration, $\mathrm{Pol}_{Z,0}$, and the fitted calibration, Pol_Z.

The variable Z gives the DAC voltage that corresponds to zero frequency. 'Back' gives the background counts of the scan in counts. This will have been determined either from the scan itself, or using a fitting process. t is the time interval over which fluorescence is collected.

The difference between the fitted collection efficiency and the nominal experimental collection efficiency is given by η_{Err}. This is defined in the same way as the DAC calibration:

$$\eta = \eta_0 \left(1 + \eta_{\mathrm{Err}}\right) \tag{H.5}$$

where η and η_0 are the fitted and experimental collection efficiencies respectively.

For the five-level, two-dimensional experimental fluorescence scans described in chapter 5 and appendix E, the drift of the 854nm beam over the scans was modelled. This drift was simulated by assuming that the 854nm beam shifted in frequency discretely by the same amount between each 866nm scan. As the frequency difference between the scans was usually altered at several points during the process, the result is that the 854nm drift as a function of 866nm detuning generally follows a piecewise curve made up of line segments. Several parameters describe the drift. $\bar{\Delta}_{854}$ describes the mean value of the 854nm drift over the scan and HW gives half the total drift. This is the half-width of the frequency distribution. The drift is modelled by a series of line segments in frequency space. These line segments are described by the vectors 'Int', 'Grad' and 'Pivot'. Int and Grad define the intercept and gradient of each line segment respectively. Pivot defines the region over which the line segment extends and is one element shorter than the other two vectors. Labelling the length of 'Int' and 'Grad' as m, and defining $1 < p < m$, the line segments are described using:

$$\Delta_{854} = \begin{cases} \mathrm{Grad}(1)\Delta_{866} + \mathrm{Int}(1), & \Delta_{866} < \mathrm{Pivot}(1) \\ \mathrm{Grad}(p)\Delta_{866} + \mathrm{Int}(p), & \mathrm{Pivot}(p-1) < \Delta_{866} < \mathrm{Pivot}(p) \\ \mathrm{Grad}(m)\Delta_{866} + \mathrm{Int}(m), & \mathrm{Otherwise} \end{cases} \tag{H.6}$$

Grad, Int and Pivot are calculated using the frequency difference between each 866nm scan and the total drift over the scan. If $m = 1$, the drift follows a single

Constant	Value
$g_J(\mathrm{S}_{1/2})$	2.00225664[ABG$^+$94]
μ_N	7.62259357MHz.T^{-1}[MTN12]
μ_I/μ_N	-1.31537[Ols72]
μ_B	9.27400968$\times10^{-24}$J.T^{-1}[MTN12]
u	1.660538921$\times10^{-27}$kg[MTN12]
c	299792458m.s^{-1}[MTN12]
k_B	1.3806488$\times10^{-23}$J.K^{-1}[MTN12]
h	6.62606957$\times10^{-34}$Js[MTN12]
I	7/2

Table H.3: Values of constants used in these simulations. This is a more complete version of table 6.2. u refers to the unified atomic mass unit and is used to calculate the mass of $^{40}Ca^+$ and $^{43}Ca^+$.

line segment for the entire scan with gradient and intercept described by Grad and Int respectively.

The final parameter in this header involves calculations of the ion's temperature. ν_T gives the trap frequency used by the excursion, dynamic and motional methods. When a figure is made up of several plots that use different trap frequencies, these are listed from top to bottom in the order of the temperature plots arranged from the top to the bottom of the figure.

There are five parameters grouped together under the header 'Fits'. 'Poly' is a vector describing the parameters of a polynomial fit. The elements of the vector from top to bottom give the coefficient of the polynomial from high to low order. The polynomial is defined as:

$$y = \sum_{n=1}^{m} \left(\mathrm{Poly}(n) x^{m-n+1}\right) \tag{H.7}$$

where m is the length of 'Poly'.

The other parameters under this header describe fitted hyperbolae and modified hyperbolae. A standard hyperbolae is described using three parameters a, b, c and the expression:

$$\Delta_y = \frac{1}{2}\left(\pm\sqrt{a^2 + (\Delta_x + b)^2} + (\Delta_x + b)\right) + c \tag{H.8}$$

where Δ_y and Δ_x are the frequencies of the beams shown on the vertical and horizontal axes of the two-dimensional plot respectively. For the four-state system, and using the terminology of section E.2, parameters a, b, c nominally correspond to Ω_3, 0 and Δ_1 respectively. The vector 'Hyp$_0$' gives the nominal values of these three parameters listed from top to bottom. 'Hyp' gives the fitted values.

The modified hyperbola seen in the five-state system is described by the eight parameters $a - h$ and the curves:

$$\begin{aligned} \Delta_y &= \frac{1}{2}\left(\sqrt{(f\Omega)^2 + (\Delta_x + 100d)^2} + (\Delta_x + 100d)\right) + \\ &\quad \Delta_1 + 100e \end{aligned} \tag{H.9}$$

$$\Delta_y = P\Delta_{hd} + (1-P)\,\Delta_\epsilon + 100cQ + 100e + 100g \tag{H.10}$$

$$\Delta_y = (1-P)\,\Delta_{hd} + P\Delta_\epsilon - 100cQ + 100e + 100g \tag{H.11}$$

where:

$$\begin{aligned} \Delta_{hd} &= \frac{1}{2}\left(-\sqrt{(f\Omega)^2 + (\Delta_x + 100d)^2} + (\Delta_x + 100d)\right) + \\ &\quad \Delta_1 + 100e \end{aligned} \tag{H.12}$$

$$\Delta_\epsilon = (\Delta_x + d) + \Delta_1 - \Delta_4 \tag{H.13}$$

$$P = 0.5\left(\tanh\left\{\frac{a}{100}\left((\Delta_x + 100d) + 100b\right)\right\}\right) \tag{H.14}$$

$$Q = \frac{h}{\cosh\left\{\frac{a}{100}\left((\Delta_x + 100d) + 100b\right)\right\}} \tag{H.15}$$

Δ_y and Δ_x are the detunings of the beams associated with the vertical and horizontal axes of the two-dimensional plot respectively. Ω refers to the Rabi frequency of the pump beam. For the five-state scans examined in section E.3, this is Ω_3. Δ_1 and Δ_4 are the detunings defined in figure E.1. These detunings refer to the transitions that excite population away from the two transitions being scanned. The parameters $a - h$ are given from the top to the bottom of the vector 'HypFull'. The final parameter describing hyperbolic systems in table H.4 is $\widetilde{\Omega}$. This gives the modified Rabi frequency associated with the different arcs of a hyperbola which is split due to the presence of Zeeman states.

ν_T gives the trap frequency

'Files' lists three final parameters. These provide information about the experimental data used in these simulations. This is provided to enable any further analysis if it is required. The data is stored on shared file storage area of the Oxford Ion Trap Group. The parameter 'Data' gives the file names of the set of data files. For file names which are long, only the end of the file name is used. This is still unique, and can be used to identify the file without the common file prefix. The date on which the data were taken is listed under 'Date'. Finally, 'Lab' gives the location of the data files. Data taken from the 'Old' lab are stored in folders named 'OldLab_x', where x is the year, and data taken from the 'New' lab are stored in folders named 'QC_x'. Data which required processing are labelled as 'Analysis'. These can be found in the path listed underneath and are labelled by the date of the analysis rather than the date the data were taken.

Figure	**3.4**	**3.6 Left**	**3.6 Right**	**3.9**
397nm:				
I_A	30			10.91
I_B	30			10.91
$f(\sigma^+)$	2/3			2/3
$f(\pi)$	1/3			1/3
$f(\sigma^-)$	0			0
Δ_A (MHz)	-1814.3			-1898.8
Δ_B (MHz)	1115.4			1030.9
EOM (MHz)	2929.7			2929.7
Γ_L (MHz)	0.5			0.5
866nm:				
I (I_0)	1000		1000	170.15
$f(\sigma^+)$	1/2		1/2	1/3
$f(\pi)$	0		0	1/3
$f(\sigma^-)$	1/2		1/2	1/3
Δ (MHz)			0	
Γ_L (MHz)	0.5		0.5	0.5
393nm:				
I_A		20.026	20.026	
I_B		20.026	20.026	
$f(\sigma^+)$		2/3	2/3	
$f(\pi)$		1/3	1/3	
$f(\sigma^-)$		0	0	
Δ_A (MHz)		-1822.4	-1822.4	
Δ_B (MHz)		1107.3	1107.3	
EOM (MHz)		2929.7	2929.7	
Γ_L (MHz)		0.5	0.5	
850nm:				
I (I_0)		1000	1000	
$f(\sigma^+)$		1/2	1/2	
$f(\pi)$		0	0	
$f(\sigma^-)$		1/2	1/2	
$f(\sigma^+)$		-68.4	-68.4	
Γ_L (MHz)		0.5	0.5	
854nm:				
I (I_0)		1000	1000	
$f(\sigma^+)$		1/2	1/2	
$f(\pi)$		0	0	
$f(\sigma^-)$		1/2	1/2	

Δ (MHz)		-14.0	-14.0	
Γ_L (MHz)		0.5	0.5	
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.1	146.1	146.1	147.33
t (s)		0.09984	0.09984	0.09984
$\eta_0(397)$ (%)				0.278
$\eta_0(393)$ (%)		0.278	0.278	
Figure	**3.11**	**3.12 Upper**	**3.12 Lower**	**3.14**
397nm:				
I (I_0)		18.74		
I_A	30		10.91	20
I_B	30		10.91	20
$f(\sigma^+)$	1/2	1/2	2/3	0.0005
$f(\pi)$	0	0	1/3	0.1442
$f(\sigma^-)$	1/2	1/2	0	0.8553
Δ (MHz)		-67.1		
Δ_A (MHz)	-1808		-1898.8	-2246.9
Δ_B (MHz)	1121.7		1030.9	1291.6
EOM (MHz)	2929.7		2929.7	3538.5
Γ_L (MHz)	0.5	0.5	0.5	0.5
866nm:				
I (I_0)	1000	165.6	170.15	1000
$f(\sigma^+)$	1/2	1/2	1/3	0.2759
$f(\pi)$	0	0	1/3	0.6704
$f(\sigma^-)$	1/2	1/2	1/3	0.0537
Γ_L (MHz)	0.5	0.5	0.5	0.5
Calibrations:				
$\|\mathbf{B}\|$ (G)	20	20	147.33	146.4
t (s)	0.09984	0.09984	0.09984	
$\eta_0(397)$ (%)	0.278	0.278	0.278	
Figure	**3.15**	**3.16**	**3.17**	**3.18**
397nm:				
I_A	20	20		20
I_B	20	20		20
$f(\sigma^+)$	0.0005	0.0005		0.0005
$f(\pi)$	0.1442	0.1442		0.1442
$f(\sigma^-)$	0.8553	0.8553		0.8553
Δ_A (MHz)	-2246.9	-2246.9		-2246.9
Δ_B (MHz)	1291.6	1291.6		1291.6
EOM (MHz)	3538.5	3538.5		3538.5

Γ_L (MHz)	0.5	0.5		0.5
866nm:				
I (I_0)	1000	1000		1000
$f(\sigma^+)$	0.2759	0.2759		0.2759
$f(\pi)$	0.6704	0.6704		0.6704
$f(\sigma^-)$	0.0537	0.0537		0.0537
Γ_L (MHz)	0.5	0.5		0.5
393nm:				
I (I_0)			1	
$f(\sigma^+)$			1	
$f(\pi)$			0	
$f(\sigma^-)$			0	
Δ (MHz)			1451.6	
Γ_L (MHz)			0.5	
850nm:				
I (I_0)			1000	
$f(\sigma^+)$			1	
$f(\pi)$			0	
$f(\sigma^-)$			0	
$f(\sigma^+)$			2300	
Γ_L (MHz)			0.5	
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.4	146.4	146.4	146.4
Figure	**3.19**	**4.2**	**4.3 Left**	**4.3 Lower**
397nm:				
I (I_0)		12.737	3.638	10.484
I_A	20			
I_B	20			
$f(\sigma^+)$	0.0005	0.7308	0.5336	0.7281
$f(\pi)$	0.1442	0	0	0
$f(\sigma^-)$	0.8553	0.2692	0.4664	0.2719
Δ (MHz)		-57.17	-1.55	-47.69
Δ_A (MHz)	-2246.9			
Δ_B (MHz)	1291.6			
EOM (MHz)	3538.5			
Γ_L (MHz)	0.5	0.5	1.5794	0.5
866nm:				
I (I_0)	1000	110.09	9.73	96.84
$f(\sigma^+)$	0.2759	0.3261	0.7120	0.3559
$f(\pi)$	0.6704	0	0.0167	0

$f(\sigma^-)$	0.0537	0.6739	0.2713	0.6441
Γ_L (MHz)	0.5	0.5	0.8911	0.5
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.4	1.9887	0.8849	1.8340
$C_{\mathrm{Err}}(866)$ (%)		-10.85	-26.61	-18.65
$C_0(866)$ (MHz.V^{-1})		0.094	0.094	0.094
$C(866)$ (MHz.V^{-1})		0.0838	0.0690	0.0765
$Z(866)$ (mV)		323.0	-113.7	272.3
Back (Counts)		0.37	26.8	0.37
t (s)		0.09984	0.09984	0.09984
$\eta_{\mathrm{Err}}(397)$ (%)		-2	-3.39	-10.71
$\eta_0(397)$ (%)		0.278	0.278	0.278
$\eta(397)$ (%)		0.2724	0.2686	0.2482
ν_T (MHz)		0.816		
Files:				
Data		009	002	009
Date		22/06/06	22/06/06	22/06/06
Lab		Old	Old	Old
Figure	**4.3 Right**	**4.4 Upper**	**4.4 Lower**	**4.5**
397nm:				
I (I_0)	6.681	3.317	3.317	10.484
$f(\sigma^+)$	0.4271	0.4999	0.4999	0.7281
$f(\pi)$	0	0.0001	0	0
$f(\sigma^-)$	0.5729	0.5000	0.5000	0.2719
Δ (MHz)	-33.6	-24.6		-47.69
Γ_L (MHz)	0.0463	0	0	0.5
866nm:				
I (I_0)	14.79	29.81	29.81	96.84
$f(\sigma^+)$	0.4406	0.5000	0.5000	0.3559
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	0.5594	0.5000	0.5000	0.6441
Δ (MHz)			62.41	
Γ_L (MHz)	2.2128	3.5837	3.5837	0.5
Calibrations:				
$\|\mathbf{B}\|$ (G)	1.5253	1.2079	1.2079	1.8340
$C_{\mathrm{Err}}(397)$ (%)			-31.58	
$C_0(397)$ (MHz.V^{-1})			1.00	
$C(397)$ (MHz.V^{-1})			0.6830	
$C_{\mathrm{Err}}(866)$ (%)	-6.72	50.00		-18.65
$C_0(866)$ (MHz.V^{-1})	0.094	1.00		0.094

$C(866)$ (MHz.V^{-1})	0.0877	1.50		0.0765
$Z(397)$ (mV)			2040.7	
$Z(866)$ (mV)	13.12	350.9		272.3
Back (Counts)	18.3	26.9	26.9	0.37
t (s)	0.16664	0.16664	0.16664	0.09984
$\eta_{\mathrm{Err}}(397)$ (%)	-2.80	-12.40	-12.40	-10.71
$\eta_0(397)$ (%)	0.11	0.11	0.11	0.278
$\eta(397)$ (%)	0.1069	0.0964	0.0964	0.2482
Files:				
Data	008	006	007	009
Date	22/06/06	22/08/07	22/08/07	22/06/06
Lab	Old	Old	Old	Old
Figure	**4.6**	**4.7**	**4.9**	**4.15**
397nm:				
I (I_0)	10.484	10.484	3.638	3.638
$f(\sigma^+)$	0.7281	0.7281	0.5336	0.5336
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	0.2719	0.2719	0.4664	0.4664
Δ (MHz)	-47.69	-47.69	-1.55	-1.55
Γ_L (MHz)	0.5	0.5	1.5794	1.5794
866nm:				
I (I_0)	96.84	96.84	9.73	9.73
$f(\sigma^+)$	0.3559	0.3559	0.7120	0.7120
$f(\pi)$	0	0	0.0167	0.0167
$f(\sigma^-)$	0.6441	0.6441	0.2713	0.2713
Γ_L (MHz)	0.5	0.5	0.8911	0.8911
Calibrations:				
$\lvert\mathbf{B}\rvert$ (G)	1.8340	1.8340	0.8849	0.8849
$C_{\mathrm{Err}}(866)$ (%)	-18.65	-18.65	-26.61	-26.61
$C_0(866)$ (MHz.V^{-1})	0.094	0.094	0.094	0.094
$C(866)$ (MHz.V^{-1})	0.0765	0.0765	0.0690	0.0690
$Z(866)$ (mV)	272.3	272.3	-113.7	-113.7
Back (Counts)	0.37	0.37	26.8	26.8
t (s)	0.09984	0.09984	0.09984	0.09984
$\eta_{\mathrm{Err}}(397)$ (%)	-10.71	-10.71	-3.39	-3.39
$\eta_0(397)$ (%)	0.278	0.278	0.278	0.278
$\eta(397)$ (%)	0.2482	0.2482	0.2686	0.2686
ν_T (MHz)	0.816	0.816	0.816	0.816
Files:				
Data	009	009	002	002
Date	22/06/06	22/06/09	22/06/06	22/06/06

Lab	Old	Old	Old	Old
Figure	**4.16**	**4.17**	**4.20**	**4.21**
397nm:				
I (I_0)	3.638	3.638	3.6	13.711
$f(\sigma^+)$	0.5336	0.5336	$1/2$	$1/2$
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	0.4664	0.4664	$1/2$	$1/2$
Δ (MHz)	-1.55	-1.55	-30.0	-49.48
Γ_L (MHz)	1.5794	1.5794	0.1	0.4276
866nm:				
I (I_0)	9.73	9.73	21.0	154.44
$f(\sigma^+)$	0.7120	0.7120	$1/2$	$1/2$
$f(\pi)$	0.0167	0.0167	0	0
$f(\sigma^-)$	0.2713	0.2713	$1/2$	$1/2$
Γ_L (MHz)	0.8911	0.8911	0.1	0.5
Calibrations:				
$\lvert\mathbf{B}\rvert$ (G)	0.8849	0.8849	5	1.48
$C_{\mathrm{Err}}(866)$ (%)	-26.61	-26.61		13.70
$C_0(866)$ (MHz.V^{-1})	0.094	0.094		0.094
$C(866)$ (MHz.V^{-1})	0.0690	0.0690		0.1069
$Z(866)$ (mV)	-113.7	-113.7		44.9
Back (Counts)	26.8	26.8		0
t (s)	0.09984	0.09984		0.09984
$\eta_{\mathrm{Err}}(397)$ (%)	-3.39	-3.39		0
$\eta_0(397)$ (%)	0.278	0.278		0.278
$\eta(397)$ (%)	0.2686	0.2686		0.2780
ν_T (MHz)	0.816		0.750	0.750
Files:				
Data	002			009
Date	22/06/06			22/06/06
Lab	Old			Old
Figure	**4.22**	**4.24 Blue**	**4.24 Green**	**4.24 Purple**
397nm:				
I (I_0)	1.559	0.290	0.720	2.6627
$f(\sigma^+)$	$1/2$	$1/2$	$1/2$	$1/2$
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	$1/2$	$1/2$	$1/2$	$1/2$
Δ (MHz)	-4.06	0.14712	-1.9661	-1.464
Γ_L (MHz)	1.5195	0.5	0.5	0.5

866nm:				
I (I_0)	76.82	5.4	9.31	17.251
$f(\sigma^+)$	1/2	1/2	1/2	1/2
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	1/2	1/2	1/2	1/2
Γ_L (MHz)	0.5	0.1	0.1	0.1
Calibrations:				
$\|\mathbf{B}\|$ (G)	1.48	0.1207	0.7822	2.4669
$C_{\mathrm{Err}}(866)$ (%)	-19.48			
$C_0(866)$ (MHz.V^{-1})	0.094			
$C(866)$ (MHz.V^{-1})	0.757			
$Z(397)$ (mV)	-42.6			
Back (Counts)	0			
t (s)	0.09984			
$\eta_{\mathrm{Err}}(397)$ (%)	0			
$\eta_0(397)$ (%)	0.278			
$\eta(397)$ (%)	0.2780			
ν_T (MHz)	0.816	0.816	0.816	0.816
Files:				
Data	003			
Date	22/06/06			
Lab	Old			
Figure	**4.25**	**5.2**	**5.4**	**5.5**
397nm:				
I (I_0)	1.0669		1.521	1.521
$f(\sigma^+)$	1/2		0.2854	0.2854
$f(\pi)$	0		0	0
$f(\sigma^-)$	1/2		0.7146	0.7146
Δ (MHz)	-2.489		-15.25	-15.25
Γ_L (MHz)			0.885	0.885
866nm:				
I (I_0)	10.4916		30.03	30.03
$f(\sigma^+)$	1/2		0.1894	0.1894
$f(\pi)$	0		0.7406	0.7406
$f(\sigma^-)$	1/2		0.0700	0.0700
Γ_L (MHz)	0.1		7.205	7.205
850nm:				
I (I_0)			4373.0	4373.0
$f(\sigma^+)$			0.4257	0.4257
$f(\pi)$			0.0014	0.0014
$f(\sigma^-)$			0.5729	0.5729

Γ_L (MHz)			0	0
854nm:				
I (I_0)			64.82	64.82
$f(\sigma^+)$			0.6441	0.6441
$f(\pi)$			0.0368	0.0368
$f(\sigma^-)$			0.3191	0.3191
Δ (MHz)			61.50	61.50
Γ_L (MHz)			1.087	1.087
Calibrations:				
$\|\mathbf{B}\|$ (G)	1.4678		0.230	0.230
$C_{\mathrm{Err}}(866)$ (%)			-14.42	-14.42
$C_0(866)$ (MHz.V^{-1})		0.093	0.093	0.093
$C(866)$ (MHz.V^{-1})			0.0804	0.0804
$C_{\mathrm{Err}}(850)$ (%)			-13.40	-13.40
$C_0(850)$ (MHz.V^{-1})		1.1	1.1	1.1
$C(850)$ (MHz.V^{-1})			0.9526	0.9526
$Z(866)$ (mV)		2940	3107.1	3107.1
$Z(850)$ (mV)		72	88.9	88.9
Back (Counts)			0	0
t (s)			$1/6$	$1/6$
$\eta_{\mathrm{Err}}(397)$ (%)			-69.97	-69.97
$\eta_0(397)$ (%)			8.7×10^{-4}	8.7×10^{-4}
$\eta(397)$ (%)			2.6×10^{-4}	2.6×10^{-4}
$\eta_{\mathrm{Err}}(393)$ (%)			-39.95	-39.95
$\eta_0(393)$ (%)			0.23	0.23
$\eta(393)$ (%)			0.1381	0.1381
ν_T (MHz)	0.816			
Files:				
Data		[011 - 096]	[011 - 096]	
Date		$18/11/10$	$18/11/10$	
Lab		Old	Old	
Figure	**5.7**	**5.8**	**5.9 1**	**5.9 2**
397nm:				
I (I_0)	1.521	1.521	1.521	1.521
$f(\sigma^+)$	0.2854	0.2854	0.2854	0.2854
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	0.7146	0.7146	0.7146	0.7146
Δ (MHz)	-15.25	-15.25	-15.25	-35.25
Γ_L (MHz)	0.885	0.885	0.885	0.885

866nm:				
I (I_0)	30.03	30.03	30.03	30.03
$f(\sigma^+)$	0.1894	0.1894	0.1894	0.1894
$f(\pi)$	0.7406	0.7406	0.7406	0.7406
$f(\sigma^-)$	0.0700	0.0700	0.0700	0.0700
Γ_L (MHz)	7.205	7.205	7.205	7.205
850nm:				
I (I_0)	4373.0	4373.0	4373.0	4373.0
$f(\sigma^+)$	0.4257	0.4257	0.4257	0.4257
$f(\pi)$	0.0014	0.0014	0.0014	0.0014
$f(\sigma^-)$	0.5729	0.5729	0.5729	0.5729
Γ_L (MHz)	0	0	0	0
854nm:				
I (I_0)	64.82	64.82	64.82	64.82
$f(\sigma^+)$	0.6441	0.6441	0.6441	0.6441
$f(\pi)$	0.0368	0.0368	0.0368	0.0368
$f(\sigma^-)$	0.3191	0.3191	0.3191	0.3191
Δ (MHz)	61.50	61.50	61.50	61.50
Γ_L (MHz)	1.087	1.087	1.087	1.087
Calibrations:				
$\|\mathbf{B}\|$ (G)	0.230	0.230	0.230	0.230
$C_{\mathrm{Err}}(866)$ (%)	-14.42	-14.42	-14.42	-14.42
$C_0(866)$ (MHz.V^{-1})	0.093	0.093	0.093	0.093
$C(866)$ (MHz.V^{-1})	0.0804	0.0804	0.0804	0.0804
$C_{\mathrm{Err}}(850)$ (%)	-13.40	-13.40	-13.40	-13.40
$C_0(850)$ (MHz.V^{-1})	1.1	1.1	1.1	1.1
$C(850)$ (MHz.V^{-1})	0.9526	0.9526	0.9526	0.9526
$Z(866)$ (mV)	3107.1	3107.1	3107.1	3107.1
$Z(850)$ (mV)	88.9	88.9	88.9	88.9
Back (Counts)	0	0	0	0
t (s)	$1/6$	$1/6$	$1/6$	$1/6$
$\eta_{\mathrm{Err}}(397)$ (%)	-69.97	-69.97	-69.97	-69.97
$\eta_0(397)$ (%)	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}
$\eta(397)$ (%)	2.6×10^{-4}	2.6×10^{-4}	2.6×10^{-4}	2.6×10^{-4}
$\eta_{\mathrm{Err}}(393)$ (%)	-39.95	-39.95	-39.95	-39.95
$\eta_0(393)$ (%)	0.23	0.23	0.23	0.23
$\eta(393)$ (%)	0.1381	0.1381	0.1381	0.1381
Figure	**5.9 3**	**5.9 4**	**5.9 5**	**5.9 6**
397nm:				
I (I_0)	1.521	1.521	1.521	1.521

$f(\sigma^+)$	0.2854	0.2854	0.2854	0.2854
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	0.7146	0.7146	0.7146	0.7146
Δ (MHz)	-15.25	-15.25	-15.25	-15.25
Γ_L (MHz)	0.885	0.885	0.885	0.885
866nm:				
I (I_0)	30.03	30.03	30.03	30.03
$f(\sigma^+)$	0.1894	0.1894	0.1894	0.1894
$f(\pi)$	0.7406	0.7406	0.7406	0.7406
$f(\sigma^-)$	0.0700	0.0700	0.0700	0.0700
Γ_L (MHz)	7.205	7.205	7.205	7.205
850nm:				
I (I_0)	8746	2186.5	4373.0	4373.0
$f(\sigma^+)$	0.4257	0.4257	0.4257	0.4257
$f(\pi)$	0.0014	0.0014	0.0014	0.0014
$f(\sigma^-)$	0.5729	0.5729	0.5729	0.5729
$f(\sigma^+)$				-25.00
Γ_L (MHz)	0	0	0	0
854nm:				
I (I_0)	64.82	64.82	64.82	64.82
$f(\sigma^+)$	0.6441	0.6441	0.6441	0.6441
$f(\pi)$	0.0368	0.0368	0.0368	0.0368
$f(\sigma^-)$	0.3191	0.3191	0.3191	0.3191
Δ (MHz)	61.50	61.50	111.50	61.50
Γ_L (MHz)	1.087	1.087	1.087	1.087
Calibrations:				
$\lvert\mathbf{B}\rvert$ (G)	0.230	0.230	0.230	0.230
$C_{\mathrm{Err}}(866)$ (%)	-14.42	-14.42	-14.42	-14.42
$C_0(866)$ (MHz.V^{-1})	0.093	0.093	0.093	0.093
$C(866)$ (MHz.V^{-1})	0.0804	0.0804	0.0804	0.0804
$C_{\mathrm{Err}}(850)$ (%)	-13.40	-13.40	-13.40	-13.40
$C_0(850)$ (MHz.V^{-1})	1.1	1.1	1.1	
$C(850)$ (MHz.V^{-1})	0.9526	0.9526	0.9526	
$Z(866)$ (mV)	3107.1	3107.1	3107.1	3107.1
$Z(850)$ (mV)	88.9	88.9	88.9	
Back (Counts)	0	0	0	0
t (s)	$1/6$	$1/6$	$1/6$	$1/6$
$\eta_{\mathrm{Err}}(397)$ (%)	-69.97	-69.97	-69.97	-69.97
$\eta_0(397)$ (%)	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}
$\eta(397)$ (%)	2.6×10^{-4}	2.6×10^{-4}	2.6×10^{-4}	2.6×10^{-4}
$\eta_{\mathrm{Err}}(393)$ (%)	-39.95	-39.95	-39.95	-39.95

$\eta_0(393)$ (%)	0.23	0.23	0.23	0.23
$\eta(393)$ (%)	0.1381	0.1381	0.1381	0.1381
Figure	**5.10 1**	**5.10 2**	**5.10 3**	**5.10 4**
Calibrations:				
$C_{\mathrm{Err}}(866)$ (%)	20.00	9.57	10.22	-4.37
$C_0(866)$ (MHz.V^{-1})	-0.034	-0.034	-0.034	-0.034
$C(866)$ (MHz.V^{-1})	-0.0408	-0.0373	-0.0375	-0.0325
$C_{\mathrm{Err}}(850)$ (%)	16.05	0.93	3.65	-7.06
$C_0(850)$ (MHz.V^{-1})	0.1976	0.1976	0.1976	0.1976
$C(850)$ (MHz.V^{-1})	0.2293	0.1994	0.2048	0.1836
$Z(866)$ (mV)	710.2	798.4	767.4	828.8
$Z(850)$ (mV)	538.5	510.7	531.7	496.6
Files:				
Data	[1.000 - 1.033]	[2.000 - 2.063]	[3.000 - 3.033 3.038 - 3.067]	[3b.002 - 3b.046]
Date	14/06/13	17/06/13	18/06/13	18/06/13
Lab	Old	Old	Old	Old
Figure	**5.10 5**	**5.10 6**	**5.11 1**	**5.11 2**
397nm:				
I (I_0)			0.3980	0.6140
$f(\sigma^+)$			0.5320	0.4244
$f(\pi)$			0	0.0463
$f(\sigma^-)$			0.4680	0.5293
Δ (MHz)			-4.75	-18.06
Γ_L (MHz)			0.321	1.026
866nm:				
I (I_0)			12.53	10.14
$f(\sigma^+)$			0.1086	0.1618
$f(\pi)$			0.6310	0.5501
$f(\sigma^-)$			0.2604	0.2881
Γ_L (MHz)			3.003	0
850nm:				
I (I_0)			4994	3735
$f(\sigma^+)$			0.4842	0.6071
$f(\pi)$			0	0
$f(\sigma^-)$			0.5158	0.3929
Γ_L (MHz)			0	0

854nm:				
I (I_0)			19.58	24.04
$f(\sigma^+)$			0.3507	0.9415
$f(\pi)$			0	0
$f(\sigma^-)$			0.6493	0.0585
Γ_L (MHz)			2.291	8.275
Calibrations:				
$\|\mathbf{B}\|$ (G)			2.505	2.09
$C_{\mathrm{Err}}(866)$ (%)	9.70	-29.95	20.00	9.57
$C_0(866)$ (MHz.V^{-1})	-0.034	-0.034	-0.034	-0.034
$C(866)$ (MHz.V^{-1})	-0.0373	-0.0238	-0.0408	-0.0373
$C_{\mathrm{Err}}(850)$ (%)	-2.32		16.05	0.93
$C_0(850)$ (MHz.V^{-1})	0.1976		0.1976	0.1976
$C(850)$ (MHz.V^{-1})	0.1930		0.2293	0.1994
$C_{I,\mathrm{Err}}(850)$ (%)		-17.78		
$C_{I,0}(850)$ (I_0.μW^{-1})		0.7962		
$C_I(850)$ (I_0.μW^{-1})		0.6546		
$Z(866)$ (mV)	859.2	1117.8	710.2	798.4
$Z(850)$ (mV)	516.8		538.5	510.7
Back (Counts)			0	12.1
t (s)			$1/6$	$1/6$
$\eta_{\mathrm{Err}}(397)$ (%)			-45.48	-52.27
$\eta_0(397)$ (%)			8.7×10^{-4}	8.7×10^{-4}
$\eta(397)$ (%)			4.8×10^{-4}	4.2×10^{-4}
$\eta_{\mathrm{Err}}(393)$ (%)			1.46	-17.39
$\eta_0(393)$ (%)			0.23	0.23
$\bar{\Delta}_{854}$ (MHz)			160.1	83.59
$HW(\Delta_{854})$ (MHz)			30.42	-38.81
Int(Δ_{854})			[16.1788 12.6608 7.4685 -22.5775 -9.1242 -17.9285]	[-26.4843 -17.8548 -3.0600 16.6732 14.1889 24.0353 30.1906]

Grad(Δ_{854})			[-0.1130 -0.1506 -0.2259 -0.4519 -0.2259 -0.1130]	[0.0827 0.1654 0.3307 0.6614 0.3307 0.1654 0.0827]
Pivot(Δ_{854}) (MHz)			[-93.4240 -68.9440 -44.4640 28.9760 77.9360]	[-104.3717 -89.4702 -59.6671 -7.5118 59.5450 74.4465]
Files:				
Data	[4.003 - 4.0028]	[008 - 0027]		
Date	14/06/13	19/06/13		
Lab	Old	Old		
Figure	**5.11 3**	**5.11 4**	**5.11 5**	**5.11 6**
397nm:				
I (I_0)	0.503	0.3320	0.3780	0.1450
$f(\sigma^+)$	0.4244	0.2088	0.4488	0.5257
$f(\pi)$	0.0463	0.0899	0	0.1317
$f(\sigma^-)$	0.5293	0.7013	0.5512	0.3426
Δ (MHz)	-8.31	-10.05	-0.3300	-8.20
Γ_L (MHz)	2.326	0	1.917	1.039
866nm:				
I (I_0)	9.02	13.17	10.13	9.65
$f(\sigma^+)$	0.1797	0.2710	0.2104	0.1600
$f(\pi)$	0.2854	0.3982	0.4683	0.5189
$f(\sigma^-)$	0.5349	0.3308	0.3213	0.3211
Γ_L (MHz)	0.13	1.906	2.214	0.0020
850nm:				
I (I_0)	7464	1567	4103	
$f(\sigma^+)$	0.6162	0.6162	0.5045	0.5363
$f(\pi)$	0	0.0264	0	0
$f(\sigma^-)$	0.3838	0.3574	0.4955	0.4637
$f(\sigma^+)$				-19.72
Γ_L (MHz)	0	0	0	0.681

854nm:				
I (I_0)	45.48	32.11	52.74	4.26
$f(\sigma^+)$	0.8507	0.7941	0.8962	0.3262
$f(\pi)$	0	0	0	0.2885
$f(\sigma^-)$	0.1493	0.2059	0.1038	0.3853
Γ_L (MHz)	2.474	10.000	0.173	4.008
Calibrations:				
$\|\mathbf{B}\|$ (G)	2.258	2.0310	2.032	1.47
$C_{\mathrm{Err}}(866)$ (%)	10.22	-4.37	9.70	-29.95
$C_0(866)$ (MHz.V^{-1})	-0.034	-0.034	-0.034	-0.034
$C(866)$ (MHz.V^{-1})	-0.0375	-0.0325	-0.0373	-0.0238
$C_{\mathrm{Err}}(850)$ (%)	3.65	-7.06	-2.32	
$C_0(850)$ (MHz.V^{-1})	0.1976	0.1976	0.1976	
$C(850)$ (MHz.V^{-1})	0.2048	0.1836	0.1930	
$C_{I,\mathrm{Err}}(850)$ (%)				-17.78
$C_{I,0}(850)$ (I_0.μW^{-1})				0.7962
$C_I(850)$ (I_0.μW^{-1})				0.6546
$Z(866)$ (mV)	767.4	828.8	859.2	1117.8
$Z(850)$ (mV)	531.7	496.6	516.8	
Back (Counts)	5.3	40	14.4	0.0009
t (s)	$1/6$	$1/6$	$1/6$	$1/6$
$\eta_{\mathrm{Err}}(397)$ (%)	-49.50	-1.43	-54.09	-79.23
$\eta_0(397)$ (%)	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}
$\eta(397)$ (%)	4.1×10^{-4}	8.6×10^{-4}	4.0×10^{-4}	1.8×10^{-4}
$\eta_{\mathrm{Err}}(393)$ (%)	-18.67	10.38	-14.96	3.85
$\eta_0(393)$ (%)	0.23	0.23	0.23	0.23
$\bar{\Delta}_{854}$ (MHz)	-1.86	-20.11	212.39	
$HW(\Delta_{854})$ (MHz)	-126.03	-39.94	15.51	
Int(Δ_{854})	[-87.6945 -65.3629 -32.7024 16.6147 47.3166 78.6714 98.3497]	[-24.1854 -20.7494 -15.6927 -4.1535 13.4787 17.6319 21.5241 24.6370 27.1012]	1.6923	

Grad(Δ_{854})	[0.2669 0.5338 1.0676 2.1353 1.0676 0.5338 0.2669]	[0.1396 0.1861 0.2792 0.5584 1.1167 0.5584 0.2792 0.1861 0.1396]	-0.1109	
Pivot(Δ_{854}) (MHz)	[-83.6675 -61.1826 -46.1927 28.7569 58.7367 73.7267]	[-73.8466 -54.3381 -41.3324 -31.5782 7.4389 13.9417 33.4502 52.9587]		
Figure	**5.12**	**6.6**	**6.8, 6.9**	**6.11**
397nm:				
I (I_0)	0.1450			
I_A		8.413	8.413	8.413
I_B		8.413	8.413	8.413
$f(\sigma^+)$	0.5257	2/3	2/3	2/3
$f(\pi)$	0.1317	1/3	1/3	1/3
$f(\sigma^-)$	0.3426	0	0	0
Δ (MHz)	-8.20			
Δ_A (MHz)		-1869.2	-1869.2	-1869.2
Δ_B (MHz)		1046.0	1046.0	1046.0
EOM (MHz)		2915.2	2915.2	2915.2
Γ_L (MHz)	1.039	0.5	0.5	0.5
866nm:				
I (I_0)	9.65	150	150	150
$f(\sigma^+)$	0.1600	1/2	1/2	1/2
$f(\pi)$	0.5189	0	0	0
$f(\sigma^-)$	0.3211	1/2	1/2	1/2
Γ_L (MHz)	0.0020	0.1	0.1	0.1
850nm:				
$f(\sigma^+)$	0.5363			
$f(\pi)$	0			
$f(\sigma^-)$	0.4637			
$f(\sigma^+)$	-19.72			

Γ_L (MHz)	0.681			
854nm:				
I (I_0)	4.26			
$f(\sigma^+)$	0.3262			
$f(\pi)$	0.2885			
$f(\sigma^-)$	0.3853			
Γ_L (MHz)	4.008			
Calibrations:				
$\|\mathbf{B}\|$ (G)	1.47	146.09		146.09
$C_{\mathrm{Err}}(866)$ (%)	-29.95			
$C_0(866)$ (MHz.V^{-1})	-0.034			
$C(866)$ (MHz.V^{-1})	-0.0238			
$C_{I,\mathrm{Err}}(850)$ (%)	-17.78			
$C_{I,0}(850)$ (I_0.µW^{-1})	0.7962			
$C_I(850)$ (I_0.µW^{-1})	0.6546			
$Z(866)$ (mV)	1117.8			
Back (Counts)	0.0009			
t (s)	$1/6$			
$\eta_{\mathrm{Err}}(397)$ (%)	-79.23			
$\eta_0(397)$ (%)	8.7×10^{-4}			
$\eta(397)$ (%)	1.8×10^{-4}			
$\eta_{\mathrm{Err}}(393)$ (%)	3.85			
$\eta_0(393)$ (%)	0.23			
ν_T (MHz)		3.00	3.00	3.00
Files:				
Data	[008 - 0027]			
Date	19/06/13			
Lab	Old			
Figure	**6.12**	**6.14**	**6.15**	**6.16**
397nm:				
I_A	8.413	8.413	8.1940	9.298
I_B	8.413	8.413	8.1940	9.298
$f(\sigma^+)$	$2/3$		$2/3$	0
$f(\pi)$	$1/3$		$1/3$	$1/3$
$f(\sigma^-)$	0		0	$2/3$
Δ_A (MHz)	-1869.2	-1869.2	-1876.0	-2277.8
Δ_B (MHz)	1046.0	1046.0	1053.7	1279.0
EOM (MHz)	2915.2	2915.2	2929.7	3556.8
Γ_L (MHz)	0		0.5	0.5
866nm:				
I (I_0)	150	150	260.2	416.4

$f(\sigma^+)$	1/2		1/2	1/2
$f(\pi)$	0		0	0
$f(\sigma^-)$	1/2		1/2	1/2
Γ_L (MHz)	0		0.1	0.1
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.09	146.09	146.09	146.09
ν_T (MHz)			3.00	3.00
Figure	**6.17 Sigma Plus**	**6.17 Lower Left**	**6.18 Upper**	**6.18 Lower**
397nm:				
I_A	30	30		8.1940
I_B	30	30		8.1940
$f(\sigma^+)$	2/3	0	2/3	2/3
$f(\pi)$	1/3	1/3	1/3	1/3
$f(\sigma^-)$	0	2/3	0	0
Δ_A (MHz)			-1876.0	-1876.0
Δ_B (MHz)			1053.7	1053.7
EOM (MHz)	2929.7	3487.6	2929.7	2929.7
Γ_L (MHz)	0.5	0.5	0.5	0.5
866nm:				
I (I_0)	1000	1000	260.2	
$f(\sigma^+)$	1/2	1/2	1/2	1/2
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	1/2	1/2	1/2	1/2
Δ (MHz)	131.07	284.03		
Γ_L (MHz)	0.1	0.1	0.1	0.1
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.09	146.09	146.09	146.09
ν_T (MHz)			3.00	3.00
Figure	**6.19 Upper**	**6.19 Lower Left**	**6.19 Lower Right**	**6.20**
397nm:				
I_A		8.1940	9.8990	2.32
I_B		8.1940	9.8990	2.32
$f(\sigma^+)$	2/3	2/3	2/3	2/3
$f(\pi)$	1/3	1/3	1/3	1/3
$f(\sigma^-)$	0	0	0	0
Δ_A (MHz)				-1930.4
Δ_B (MHz)				999.3

EOM (MHz)	2929.7	2929.7	2929.7	2929.7
Γ_L (MHz)	0.5	0.5	0.5	0.5
866nm:				
I (I_0)		260.2	171.3	192.95
$f(\sigma^+)$	$1/2$	$1/2$	$1/2$	$1/2$
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	$1/2$	$1/2$	$1/2$	$1/2$
Δ (MHz)		-152.39	-175.56	
Γ_L (MHz)	0.1	0.1	0.1	0.1
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.09	146.09	146.09	146.09
ν_T (MHz)	3.00			[3.00 0.816]
Figure	**6.22**	**6.24 Left**	**6.24 Right**	**6.25**
397nm:				
I_A	16.822	50	6.132	13.57
I_B	16.822	50	6.132	13.57
$f(\sigma^+)$	$2/3$	0.7791	$2/3$	$2/3$
$f(\pi)$	$1/3$	0.2209	$1/3$	$1/3$
$f(\sigma^-)$	0	0	0	0
Δ_A (MHz)	-1423.5			-1905.9
Δ_B (MHz)	1018.6			1023.8
EOM (MHz)	2442.1	2949.4	3382.2	2929.7
Γ_L (MHz)	0.5	1	0.5	0.5
866nm:				
I (I_0)	1000			79.65
$f(\sigma^+)$	$1/2$			$1/2$
$f(\pi)$	0			0
$f(\sigma^-)$	$1/2$			$1/2$
Γ_L (MHz)	0.1			0.5
850nm:				
I (I_0)		5000	1865.5	
$f(\sigma^+)$		$1/2$	0.5351	
$f(\pi)$		0	0.4649	
$f(\sigma^-)$		$1/2$	0	
$f(\sigma^+)$		41.97	218.48	
Γ_L (MHz)		0	0.5	
854nm:				
I (I_0)		2415.8	740.3	
$f(\sigma^+)$		$1/2$	$1/4$	

$f(\pi)$		0	1/2	
$f(\sigma^-)$		1/2	1/4	
Δ (MHz)		-73.28	-393.36	
Γ_L (MHz)		0	0.5	
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.09	146.4	146.4	146.09
$C_{\mathrm{Err}}(866)$ (%)				-5.10
$C_0(866)$ (MHz.V^{-1})				0.094
$C(866)$ (MHz.V^{-1})				0.0892
$Z(866)$ (mV)				1599.8
Back (Counts)				537.4
t (s)				0.09984
$\eta_{\mathrm{Err}}(397)$ (%)				0
$\eta_0(397)$ (%)				0.278
$\eta(397)$ (%)				0.278
ν_T (MHz)	3.00			
Files:				
Data				002
Date				23/01/12
Lab				New
Figure	**6.26**	**6.27**	**6.28 Upper Left**	**6.28 Upper Right**
397nm:				
I_A	13.57	13.57	17.8897	15.7629
I_B	13.57	13.57	17.8897	15.7629
$f(\sigma^+)$	2/3	2/3	2/3	2/3
$f(\pi)$	1/3	1/3	1/3	1/3
$f(\sigma^-)$	0	0	0	0
Δ_A (MHz)	-1905.9	-1905.9	-1893.1	-1900.3
Δ_B (MHz)	1023.8	1023.8	1036.6	1029.4
EOM (MHz)	2929.7	2929.7	2929.7	2929.7
Γ_L (MHz)	0	0.5	0.5	0.4
866nm:				
I (I_0)	79.65	79.65	108.5959	180.5751
$f(\sigma^+)$	1/2	1/2	1/2	1/2
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	1/2	1/2	1/2	1/2
Γ_L (MHz)	0	0.1	0.5	0.1
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.09	146.09	146.09	146.09

$C_{\mathrm{Err}}(866)$ (%)	-5.10		2.02	1.42
$C_0(866)$ (MHz.V^{-1})	0.094		0.08518	0.08518
$C(866)$ (MHz.V^{-1})	0.0892		0.0869	0.0864
$Z(866)$ (mV)	1599.8		837.5	909.9
Back (Counts)	537.4		926.1	1191.5
t (s)	0.099984		0.099984	0.099984
$\eta_{\mathrm{Err}}(397)$ (%)	0		0	0
$\eta_0(397)$ (%)	0.278		0.278	0.278
$\eta(397)$ (%)	0.278		0.278	0.278
ν_T (MHz)		3.00		
Files:				
Data	002		003	002
Date	23/01/12		02/11/12	02/11/12
Lab	New		New	New
Figure	**6.28 Lower Left**	**6.28 Lower Right**	**6.29 Upper Left**	**6.29 Upper Right**
397nm:				
I_A	16.4102	44.1287	17.8897	15.7629
I_B	16.4102	44.1287	17.8897	15.7629
$f(\sigma^+)$	2/3	2/3	2/3	2/3
$f(\pi)$	1/3	1/3	1/3	1/3
$f(\sigma^-)$	0	0	0	0
Δ_A (MHz)	-1902.4	-1889.8	-1893.1	-1900.3
Δ_B (MHz)	1027.3	1039.9	1036.6	1029.4
EOM (MHz)	2929.7	2929.7	2929.7	2929.7
Γ_L (MHz)	0.4	0.4	0.5	0.4
866nm:				
I (I_0)	295.0652	220.6287	108.5959	180.5751
$f(\sigma^+)$	1/2	1/2	1/2	1/2
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	1/2	1/2	1/2	1/2
Γ_L (MHz)	0.1	0.1	0.5	0.1
Calibrations:				
$\lvert\mathbf{B}\rvert$ (G)	146.09	146.09	146.09	146.09
$C_{\mathrm{Err}}(866)$ (%)	1.56	3.38		
$C_0(866)$ (MHz.V^{-1})	0.08518	0.08518		
$C(866)$ (MHz.V^{-1})	0.0865	0.0881		
$Z(866)$ (mV)	945.1	750.8		
Back (Counts)	1353.6	1891.1		
t (s)	0.099984	0.099984		

$\eta_{\mathrm{Err}}(397)$ (%)	0	0		
$\eta_0(397)$ (%)	0.278	0.278		
$\eta(397)$ (%)	0.278	0.278		
ν_T (MHz)			3.00	3.00
Files:				
Data	004	001		
Date	02/11/12	02/11/12		
Lab	New	New		
Figure	**6.29 Lower Left**	**6.29 Lower Right**	**6.30**	**6.31**
397nm:				
I_A	16.4102	44.1287	11.9627	13.5282
I_B	16.4102	44.1287	11.9627	11.5956
$f(\sigma^+)$	2/3	2/3	2/3	2/3
$f(\pi)$	1/3	1/3	1/3	1/3
$f(\sigma^-)$	0	0	0	0
Δ_A (MHz)	-1902.4	-1889.8	-1903.9	-1900.7
Δ_B (MHz)	1027.3	1039.9	1025.8	1029.0
EOM (MHz)	2929.7	2929.7	2929.7	2929.7
Γ_L (MHz)	0.4	0.4	0.4	0.65
866nm:				
I (I_0)	295.0652	220.6287	518.0733	191.0863
$f(\sigma^+)$	1/2	1/2	1/2	1/2
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	1/2	1/2	1/2	1/2
Γ_L (MHz)	0.1	0.1	0.1	0.1
Calibrations:				
$\lvert\mathbf{B}\rvert$ (G)	146.09	146.09	146.09	146.09
$C_{\mathrm{Err}}(866)$ (%)			-11.16	-4.7
$C_0(866)$ (MHz.V^{-1})			0.0976	0.0887
$C(866)$ (MHz.V^{-1})			0.0867	0.0845
$Z(866)$ (mV)			1376.8	361.7
Back (Counts)			455.2	989.9
t (s)			0.199968	0.199968
$\eta_{\mathrm{Err}}(397)$ (%)			0	0
$\eta_0(397)$ (%)			0.1645	0.233
$\eta(397)$ (%)			0.1645	0.233
ν_T (MHz)	3.00	3.00	3.00	
Files:				
Data			034	024

Date			10/10/13	18/10/13
Lab			New	New
Figure	**6.32**	**6.33**	**6.34**	**6.35**
397nm:				
I_A	1.2512	1.2408	1.2408	1.2512
I_B	1.0635	1.0547	1.0547	1.0635
$f(\sigma^+)$	2/3	2/3	2/3	2/3
$f(\pi)$	1/3	1/3	1/3	1/3
$f(\sigma^-)$	0	0	0	0
Δ_A (MHz)	-1887.6	-1887.6	-1887.6	-1887.6
Δ_B (MHz)	1042.1	1042.1	1042.1	1042.1
EOM (MHz)	2929.7	2929.7	2929.7	2929.7
Γ_L (MHz)	1.9927	1.5257	1.5257	1.9927
866nm:				
I (I_0)	162.51	160.9634	160.9634	162.51
$f(\sigma^+)$	1/2	1/2	1/2	1/2
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	1/2	1/2	1/2	1/2
Γ_L (MHz)	0.1	0.1	0.1	0.1
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.09	146.09	146.09	146.09
$C_{\mathrm{Err}}(866)$ (%)	0.01	-15.27	-15.27	0.01
$C_0(866)$ (MHz.V^{-1})	-0.008419	0.0964	0.0964	-0.008419
$C(866)$ (MHz.V^{-1})	-0.0084198	0.0817	0.0817	-0.0084198
$Z(866)$ (mV)	-15604	[3949.7 7538.9 -1693.4]	[3949.7 7538.9 -1693.4]	-15604
Back (Counts)	411.3	[535.8 527.8 538.7]	[535.8 527.8 538.7]	411.3
t (s)	0.3999		0.49992	0.49992
$\eta_{\mathrm{Err}}(397)$ (%)	0	0	0	0
$\eta_0(397)$ (%)	0.17	0.17	0.17	0.17
$\eta(397)$ (%)	0.17	0.17	0.17	0.17
ν_T (MHz)	3.00		3.00	3.00
Fits:				
Poly	[6.87×10^{-6} 3.43×10^{-3} 0.568 31.3]			

Files:				
Data	023	[001 - 003]	[001 - 003]	023
Date	05/02/15	06/02/15	06/02/15	05/02/15
Lab	New	New	New	New
Figure	**6.36**	**6.37**	**7.1**	**D.1**
397nm:				
I (I_0)			3.4111	0.3980
I_A		18.946		
I_B		18.946		
$f(\sigma^+)$	2/3	2/3	0.2557	0.5320
$f(\pi)$	1/3	1/3	0	0
$f(\sigma^-)$	0	0	0.7443	0.4680
Δ (MHz)			-29.9472	-4.75
Δ_A (MHz)		-1685.3		
Δ_B (MHz)		967.0		
EOM (MHz)	2929.7	2652.3		
Γ_L (MHz)	1.5257	0.5	7.4×10^{-4}	0.321
866nm:				
I (I_0)		755.3	16.9044	12.53
$f(\sigma^+)$	1/2	1/2	0.7234	0.1086
$f(\pi)$	0	0	0.2078	0.6310
$f(\sigma^-)$	1/2	1/2	0.0688	0.2604
Γ_L (MHz)	0.1	0.1	8.7896	3.003
850nm:				
I (I_0)			7086.4	4994
$f(\sigma^+)$				0.4842
$f(\pi)$				0
$f(\sigma^-)$				0.5158
$f(\sigma^+)$			-94.9373	
Γ_L (MHz)			7.0571	0
854nm:				
I (I_0)			16.5474	19.58
$f(\sigma^+)$			0.2074	0.3507
$f(\pi)$			0.0524	0
$f(\sigma^-)$			0.7402	0.6493
Δ (MHz)			-34.6040	
Γ_L (MHz)			6.8829	2.291
Calibrations:				
$\lvert\mathbf{B}\rvert$ (G)	146.09		2.6532	2.505
$C_{\mathrm{Err}}(866)$ (%)	-15.27		12.09	20.00
$C_0(866)$ (MHz.V^{-1})	0.094		-0.4617	-0.034

$C(866)$ (MHz.V^{-1})	0.0796		-0.5175	-0.0408
$C_{\mathrm{Err}}(850)$ (%)				16.05
$C_0(850)$ (MHz.V^{-1})				0.1976
$C(850)$ (MHz.V^{-1})				0.2293
$\mathrm{Pol}_{\mathrm{Err}}(850)$ (%)			0.3031	
$\mathrm{Pol}_0(850)$ ($^{\circ -1}$)			0.01087	
$\mathrm{Pol}(850)$ ($^{\circ -1}$)			0.01093	
$\mathrm{Pol}_{Z.\mathrm{Err}}(850)$ (%)			0.3031	
$\mathrm{Pol}_Z, 0(850)$ ($^\circ$)			8	
$\mathrm{Pol}_Z(850)$ ($^\circ$)			8.0	
$Z(866)$ (mV)			24.5	710.2
$Z(850)$ (mV)				538.5
Back (Counts)		287.8	0.02	0
t (s)	0.3999			1/6
$\eta_{\mathrm{Err}}(397)$ (%)	0		0.4196	-45.48
$\eta_0(397)$ (%)	0.17		0	8.7×10^{-4}
$\eta(397)$ (%)			9.7×10^{-4}	4.8×10^{-4}
$\eta_{\mathrm{Err}}(393)$ (%)			-53.26	1.46
$\eta_0(393)$ (%)			0.23	0.23
$\bar{\Delta}_{854}$ (MHz)				160.1
$HW(\Delta_{854})$ (MHz)				30.42
Int(Δ_{854})				[16.1788 12.6608 7.4685 -22.5775 -9.1242 -17.9285]
Grad(Δ_{854})				[-0.1130 -0.1506 -0.2259 -0.4519 -0.2259 -0.1130]
Pivot(Δ_{854}) (MHz)				[-93.4240 -68.9440 -44.4640 28.9760 77.9360]
ν_T (MHz)	3.00	3.00		

Files:				
Data	[power397 power866 detuning397 detuning866]		[155 - 176]	[1.000 - 1.033]
Date	08/06/15		07/08/15	14/06/13
Lab	[Analysis Tom]		Old	Old
Figure	**D.2**	**D.3**	**D.4**	**D.5**
397nm:				
I (I_0)	0.6140	0.503	0.3320	0.3780
$f(\sigma^+)$	0.4244	0.4244	0.2088	0.4488
$f(\pi)$	0.0463	0.0463	0.0899	0
$f(\sigma^-)$	0.5293	0.5293	0.7013	0.5512
Δ (MHz)	-18.06	-8.31	-10.05	-0.3300
Γ_L (MHz)	1.026	2.326	0	1.917
866nm:				
I (I_0)	10.14	9.02	13.17	10.13
$f(\sigma^+)$	0.1618	0.1797	0.2710	0.2104
$f(\pi)$	0.5501	0.2854	0.3982	0.4683
$f(\sigma^-)$	0.2881	0.5349	0.3308	0.3213
Γ_L (MHz)	0	0.13	1.906	2.214
850nm:				
I (I_0)	3735	7464	1567	4103
$f(\sigma^+)$	0.6071	0.6162	0.6162	0.5045
$f(\pi)$	0	0	0.0264	0
$f(\sigma^-)$	0.3929	0.3838	0.3574	0.4955
Γ_L (MHz)	0	0	0	0
854nm:				
I (I_0)	24.04	45.48	32.11	52.74
$f(\sigma^+)$	0.9415	0.8507	0.7941	0.8962
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	0.0585	0.1493	0.2059	0.1038
Γ_L (MHz)	8.275	2.474	10.000	0.173
Calibrations:				
$\lvert\mathbf{B}\rvert$ (G)	2.09	2.258	2.0310	2.032
$C_{\mathrm{Err}}(866)$ (%)	9.57	10.22	-4.37	9.70
$C_0(866)$ (MHz.V^{-1})	-0.034	-0.034	-0.034	-0.034
$C(866)$ (MHz.V^{-1})	-0.0373	-0.0375	-0.0325	-0.0373
$C_{\mathrm{Err}}(850)$ (%)	0.93	3.65	-7.06	-2.32
$C_0(850)$ (MHz.V^{-1})	0.1976	0.1976	0.1976	0.1976

$C(850)$ (MHz.V^{-1})	0.1994	0.2048	0.1836	0.1930
$Z(866)$ (mV)	798.4	767.4	828.8	859.2
$Z(850)$ (mV)	510.7	531.7	496.6	516.8
Back (Counts)	12.1	5.3	40	14.4
t (s)	$1/6$	$1/6$	$1/6$	$1/6$
$\eta_{\mathrm{Err}}(397)$ (%)	-52.27	-49.50	-1.43	-54.09
$\eta_0(397)$ (%)	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}	8.7×10^{-4}
$\eta(397)$ (%)	4.2×10^{-4}	4.1×10^{-4}	8.6×10^{-4}	4.0×10^{-4}
$\eta_{\mathrm{Err}}(393)$ (%)	-17.39	-18.67	10.38	-14.96
$\eta_0(393)$ (%)	0.23	0.23	0.23	0.23
$\bar{\Delta}_{854}$ (MHz)	83.59	-1.86	-20.11	212.39
$HW(\Delta_{854})$ (MHz)	-38.81	-126.03	-39.94	15.51
Int(Δ_{854})	[-26.4843 -17.8548 -3.0600 16.6732 14.1889 24.0353 30.1906]	[-87.6945 -65.3629 -32.7024 16.6147 47.3166 78.6714 98.3497]	[-24.1854 -20.7494 -15.6927 -4.1535 13.4787 17.6319 21.5241 24.6370 27.1012]	1.6923
Grad(Δ_{854})	[0.0827 0.1654 0.3307 0.6614 0.3307 0.1654 0.0827]	[0.2669 0.5338 1.0676 2.1353 1.0676 0.5338 0.2669]	[0.1396 0.1861 0.2792 0.5584 1.1167 0.5584 0.2792 0.1861 0.1396]	-0.1109
Pivot(Δ_{854}) (MHz)	[-104.3717 -89.4702 -59.6671 -7.5118 59.5450 74.4465]	[-83.6675 -61.1826 -46.1927 28.7569 58.7367 73.7267]	[-73.8466 -54.3381 -41.3324 -31.5782 7.4389 13.9417 33.4502 52.9587]	

Files:				
Data	[2.000 - 2.063]	[3.000 - 3.033 3.038 - 3.067]	[3b.002 - 3b.046]	[4.003 - 4.0028]
Date	17/06/13	18/06/13	18/06/13	14/06/13
Lab	Old	Old	Old	Old
Figure	**D.6**	**E.2 a**	**E.2 b**	**E.2 c**
397nm:				
I (I_0)	0.1450	1	1	1
$f(\sigma^+)$	0.5257			
$f(\pi)$	0.1317			
$f(\sigma^-)$	0.3426			
Δ (MHz)	-8.20			
Γ_L (MHz)	1.039	0	0	0
866nm:				
I (I_0)	9.65	1	10	100
$f(\sigma^+)$	0.1600			
$f(\pi)$	0.5189			
$f(\sigma^-)$	0.3211			
Γ_L (MHz)	0.0020	0	0	0
850nm:				
$f(\sigma^+)$	0.5363			
$f(\pi)$	0			
$f(\sigma^-)$	0.4637			
$f(\sigma^+)$	-19.72			
Γ_L (MHz)	0.681			
854nm:				
I (I_0)	4.26			
$f(\sigma^+)$	0.3262			
$f(\pi)$	0.2885			
$f(\sigma^-)$	0.3853			
Γ_L (MHz)	4.008			
Calibrations:				
$\|\mathbf{B}\|$ (G)	1.47			
$C_{\mathrm{Err}}(866)$ (%)	-29.95			
$C_0(866)$ (MHz.V^{-1})	-0.034			
$C(866)$ (MHz.V^{-1})	-0.0238			
$C_{I,\mathrm{Err}}(850)$ (%)	-17.78			
$C_{I,0}(850)$ (I_0.µW^{-1})	0.7962			

$C_I(850)$ (I_0.μW^{-1})	0.6546			
$Z(866)$ (mV)	1117.8			
Back (Counts)	0.0009			
t (s)	$1/6$	$1/6$		
$\eta_{\mathrm{Err}}(397)$ (%)	-79.23			
$\eta_0(397)$ (%)	8.7×10^{-4}			
$\eta(397)$ (%)	1.8×10^{-4}			
$\eta_{\mathrm{Err}}(393)$ (%)	3.85			
$\eta_0(393)$ (%)	0.23			
Files:				
Data	[008 - 0027]			
Date	19/06/13			
Lab	Old			
Figure	**E.2 d**	**E.3**	**E.4**	**E.5 a**
397nm:				
I (I_0)	1	1	1	1
Δ (MHz)				-40
Γ_L (MHz)	0	0	0	0
866nm:				
I (I_0)	1000	1000		30
Γ_L (MHz)	0	0	0	0
850nm:				
I (I_0)				10
Γ_L (MHz)				0
Figure	**E.5 b**	**E.5 c**	**E.5 d**	**E.6**
397nm:				
I (I_0)	1	1	1	1
Δ (MHz)	-40	-40	-40	-40
Γ_L (MHz)	0	0	0	0
866nm:				
I (I_0)	30	30	30	30
Γ_L (MHz)	0	0	0	0
850nm:				
I (I_0)	100	1000	10000	10000
Γ_L (MHz)	0	0	0	0

Figure	**E.7**	**E.8 Upper Left**	**E.8 Upper Right**	**E.8 Lower Left**
397nm:				
I (I_0)	1	1	1	15
Δ (MHz)	-40	-40	-40	-40
Γ_L (MHz)	0	0	0	0
866nm:				
I (I_0)	30	200	1000	30
Γ_L (MHz)	0	0	0	0
850nm:				
I (I_0)	10000	10000	10000	10000
Γ_L (MHz)	0	0	0	0
Fits:				
Hyp_0	[181.3508 -40 0]			
Hyp	[180.4148 -36.6587 -5.8778]			

Figure	**E.8 Lower Right**	**E.9 Left**	**E.9 Right**	**E.10**
397nm:				
I (I_0)	400	1	1	1
Δ (MHz)	-40	-40	-40	-40
Γ_L (MHz)	0	0	0	0
866nm:				
I (I_0)	30	30	30	30
Γ_L (MHz)	0	0	0	0
850nm:				
I (I_0)	10000	10	10000	10000
Γ_L (MHz)	0	0	0	0
854nm:				
I (I_0)		80	80	80
Δ (MHz)		150	150	150
Γ_L (MHz)		0	0	0

Figure	**E.11 Left**	**E.11 Right**	**E.12**	**E.13**
397nm:				
I (I_0)		1	1	1.675

Δ (MHz)		-40	-40	-80
Γ_L (MHz)		0	0	0
866nm:				
I (I_0)		30	30	800
Γ_L (MHz)		0	0	0
850nm:				
I (I_0)		10000	10000	1000
Γ_L (MHz)		0	0	0
854nm:				
I (I_0)		80	80	0.0426
Δ (MHz)		150	150	600
Γ_L (MHz)		0	0	0
Fits:				
HypFull	[1.5042 -0.8886 0.1128 -0.0364 0.0976 0.9137 -0.2113 1.2403]			
Figure	**E.14 Left**	**E.14 Right**	**E.15 Upper Left**	**E.15 Upper Right**
397nm:				
I (I_0)	500	1	1	1
$f(\sigma^+)$				1/2
$f(\pi)$				0
$f(\sigma^-)$				1/2
Δ (MHz)	-40	-40		
Γ_L (MHz)	0	0	0	0
866nm:				
I (I_0)	30	30	5000	5000
$f(\sigma^+)$				1/2
$f(\pi)$				0
$f(\sigma^-)$				1/2
Γ_L (MHz)	0	0	0	0
850nm:				
I (I_0)	10000	10000		
Γ_L (MHz)	0	0		

854nm:				
I (I_0)	80	12000		
Δ (MHz)	150	150		
Γ_L (MHz)	0	0		
Figure	**E.15 Lower Left**	**E.15 Lower Right**	**E.16**	**E.18**
397nm:				
I (I_0)	1	1	1	1
$f(\sigma^+)$	1/2	1/2	1/2	1/2
$f(\pi)$	0	0	0	0
$f(\sigma^-)$	1/2	1/2	1/2	1/2
Γ_L (MHz)	0	0	0	0
866nm:				
I (I_0)	5000	5000		5000
$f(\sigma^+)$	0	1/3	1/3	1/3
$f(\pi)$	1	1/3	1/6	1/6
$f(\sigma^-)$	0	1/3	1/2	1/2
Γ_L (MHz)	0	0	0	0
Fits:				
$\widetilde{\Omega}$ (MHz)				[110.0088 296.2929]
Figure	**E.19**	**E.20**	**F.1**	**F.2**
397nm:				
I (I_0)	1			
I_A		5000	17.8897	15.7629
I_B		5000	17.8897	15.7629
$f(\sigma^+)$	1/3	1/3	2/3	2/3
$f(\pi)$	1/3	1/3	1/3	1/3
$f(\sigma^-)$	1/3	1/3	0	0
Δ_A (MHz)			-1893.1	-1900.3
Δ_B (MHz)			1036.6	1029.4
EOM (MHz)		3225.7	2929.7	2929.7
Γ_L (MHz)	0	0	0.5	0.4
866nm:				
I (I_0)	7000	1	108.5959	180.5751
$f(\sigma^+)$	0.025	1/2	1/2	1/2
$f(\pi)$	0.95	0	0	0
$f(\sigma^-)$	0.025	1/2	1/2	1/2
Γ_L (MHz)	0	0	0.5	0.1

Calibrations:				
$\|\mathbf{B}\|$ (G)			146.09	146.09
$C_{\mathrm{Err}}(866)$ (%)			2.02	1.42
$C_0(866)$ (MHz.V^{-1})			0.08518	0.08518
$C(866)$ (MHz.V^{-1})			0.0869	0.0864
$Z(866)$ (mV)			837.5	909.9
Back (Counts)			926.1	1191.5
t (s)				0.099984
$\eta_{\mathrm{Err}}(397)$ (%)			0	0
$\eta_0(397)$ (%)			0.278	0.278
$\eta(397)$ (%)			0.278	0.278
ν_T (MHz)		0.816	3.00	3.00
Fits:				
$\widetilde{\Omega}$ (MHz)	[67.3472 171.9291 349.8027]			
Files:				
Data			003	002
Date			02/11/12	02/11/12
Lab			New	New
Figure	**F.3**	**F.4**	**F.5**	**G.1 Upper**
397nm:				
I_A	16.4102	44.1287	13.5282	1.373
I_B	16.4102	44.1287	11.5956	1.373
$f(\sigma^+)$	2/3	2/3	2/3	0.9990
$f(\pi)$	1/3	1/3	1/3	0
$f(\sigma^-)$	0	0	0	0.0001
Δ_A (MHz)	-1902.4	-1889.8	-1900.7	-1825.0
Δ_B (MHz)	1027.3	1039.9	1029.0	1100.4
EOM (MHz)	2929.7	2929.7	2929.7	2925.4
Γ_L (MHz)	0.4	0.4	0.65	0.5
866nm:				
I (I_0)	295.0652	220.6287	191.0863	1000
$f(\sigma^+)$	1/2	1/2	1/2	0.0530
$f(\pi)$	0	0	0	0.4424
$f(\sigma^-)$	1/2	1/2	1/2	0.5046
Δ (MHz)				64.433
Γ_L (MHz)	0.1	0.1	0.1	0.5
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.09	146.09	146.09	146.4

$C_{\mathrm{Err}}(866)$ (%)	1.56	3.38	-4.7	
$C_0(866)$ (MHz.V^{-1})	0.08518	0.08518	0.0887	
$C(866)$ (MHz.V^{-1})	0.0865	0.0881	0.0845	
$Z(866)$ (mV)	945.1	750.8	361.7	
Back (Counts)	1353.6	1891.1	989.9	
t (s)	0.099984	0.099984	0.099984	0.199968
$\eta_{\mathrm{Err}}(397)$ (%)	0	0	0	
$\eta_0(397)$ (%)	0.278	0.278	0.233	
$\eta(397)$ (%)	0.278	0.278	0.233	
ν_T (MHz)	3.00	3.00	3.00	
Files:				
Data	004	001	024	
Date	02/11/12	02/11/12	18/10/13	
Lab	New	New	New	
Figure	**G.1 Lower**	**G.3 Pulse A**	**G.3 Pulse B**	**G.3 Pulse C**
397nm:				
I_A	1.373			
I_B	1.373			
$f(\pi)$	0			
Δ_A (MHz)	-1825.0			
Δ_B (MHz)	1100.4			
EOM (MHz)	2925.4			
Γ_L (MHz)	0.5			
866nm:				
I (I_0)	1000			
$f(\sigma^+)$	0.0530			
$f(\pi)$	0.4424			
$f(\sigma^-)$	0.5046			
Δ (MHz)	64.433			
Γ_L (MHz)	0.5			
393nm:				
I (I_0)		0.0146		
$f(\sigma^+)$		0.990		
$f(\pi)$		0.005		
$f(\sigma^-)$		0.005		
Δ (MHz)		1451.6		
Γ_L (MHz)		0.5		
850nm:				
I_A (I_0)			4.49	2.05
I_B (I_0)			4.49	2.05

$f(\sigma^+)$			0.990	0
$f(\pi)$			0.005	0.97
$f(\sigma^-)$			0.005	0.03
$f(\sigma^+)$				248.699
Δ_A (MHz)			96.3079	
Δ_B (MHz)			340.8714	
EOM (MHz)			244.5635	
Γ_L (MHz)			0.5	0.5
Calibrations:				
$\|\mathbf{B}\|$ (G)	146.4	146.4	146.4	146.4

Table H.4: Parameters describing the figures presented in this thesis. For a description of the parameters, see text.