Durham E-Theses

First observations of Rydberg blockade in a frozen gas of divalent atoms

BODDY, DANIELLE

How to cite:

BODDY, DANIELLE (2014) *First observations of Rydberg blockade in a frozen gas of divalent atoms*, Durham theses, Durham University. Available at Durham E-Theses Online: http://etheses.dur.ac.uk/10740/

Use policy

First observations of Rydberg blockade in a frozen gas of divalent atoms

Danielle Boddy

This thesis details the first measurements of Rydberg dipole blockade in a cold ensemble of divalent atoms. Strontium atoms are cooled and trapped in a magneto–optical trap and coherently excited to Rydberg states in a two–photon, three–level ladder scheme. Owing to the divalent nature of strontium, one electron can be excited to the Rydberg state, whilst the other lower–lying electron is available to undergo resonant optical excitation to autoionising states, which ionise in sub–nanosecond timescales. The remaining ions that are recorded on a micro–channel plate are proportional to the number of Rydberg atoms.

The development of a narrow linewidth laser system necessary for an additional stage of cooling is explained and characterised. Two frequency stabilisation schemes are discussed: one to address the short–term laser frequency instabilities based on the Pound–Drever–Hall technique; the other to address the long–term laser frequency instabilities based on Lamb–dip spectroscopy in an atomic beam. The cooling dynamics on the narrow cooling transition is studied experimentally and modelled via theoretical simulations.

First observations of Rydberg blockade in a frozen gas of divalent atoms

Danielle Boddy

A thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

Department of Physics
Durham University

August 14, 2014

Contents

List of Figures

Declaration

I confirm that no part of the material offered has previously been submitted by myself for a degree in this or any other University. Where material has been generated through joint work, the work of others has been indicated.

Danielle Boddy
Durham, August 14, 2014

Acknowledgments

I would like to express my gratitude to the many people who helped me during my time as a Ph.D. student. I would like to thank them whole heartedly, and I apologise to anyone I have not specifically mentioned here. In particular, I must thank Matt Jones for his supervision, encouragement and guidance. I could not have asked for a more appreciative supervisor, especially one who understands the annoyance of electronics. I admire the love and enthusiasm he has for his work, and his ability to see the forest through the trees, or rather the sound of accomplishment through the moans and groans of his students.

My time in the lab would have been intolerable and extremely silent without the other team members. I am lucky to have worked with all team strontium members, both past and present. I must thank James Millen, the first student on this experiment, for forming the basis of the experiment and helping me through my first few months as a clueless first year student. I can not thank Graham Lochead enough for giving up his time to teach me the ropes of the experiment in my early Ph.D. years. In particular, I enjoyed his nonsensical chit–chat and receiving random facts during long days in the lab. My favourite being the Metro Trains video of 'Dumb Ways to Die'. I wish him all the best in his scientific endeavor, though I hope he no longer uses his awful attempt of a Geordie accent in public. The new lead scientist, Dan Sadler, is a credit to the team. His attention to detail and inquisitive questions, inspired me to follow suit, and for this I must also thank him immensely. His passionate ranting about something or other always entertained me. I look forward to following his work, be it in the scientific or business world. I must thank Liz Bridge for her patience and optimism. Having worked on a relatively identical Ph.D. at NPL, she understands the complexity and frustration of working with narrow linewidth lasers in strontium. Her caring nature shone from day one and since then she has adopted the motherly role in the group, which was needed during my last panic–stricken months in the lab when it seemed everything was going belly–up. The newest member of the team, Alistair Bounds, fits in perfectly with the group and I'm sure he will go on to do exciting things with the experiment. His witty sense of humour made my last few months writing a thesis a little more enjoyable, but just a little bit!

The research group AtMol are the friendliest and most helpful group of people I have come across. Of the people in this group, I would like to thank my second supervisor Charles Adams for his scientific guidance; Ifan Hughes for his *phenomenal* scientific explanations and humorous, mainly football–related, chat; and Kev Weatherill for inspiring me to do a Ph.D. in the first place. I must thank my other work colleagues David Paredes, David Szwer, Christophe Vaillant and Tommy Ogden for both scientific and running–related discussions.

To my mam and dad, for teaching me
the valuable lesson of never giving up.

For Jonathan and our journey together.

Chapter 1

Introduction

The ability to routinely create samples of ultracold atomic ensembles has opened new avenues to investigating quantum systems, providing one of the few systems where quantum states can be prepared and manipulated. Key interests of study are strongly correlated systems, whereby interactions play a strong role in many–body systems. Many–body systems are ubiquitous in nature and the description and understanding of their properties is an essential topic in many areas of physics [1], such as quantum information and high temperature superconductors.

A many–body system is generally very difficult to solve, owing to the vast quantity of information the system contains. The basic equations of quantum mechanics can not be solved directly and approximations are made to simplify calculations. The most common approach is to employ mean–field theory through which the effect of all other atoms on any given atom is approximated by a single averaged effect [2], thus reducing the many–body problem to a one–body problem. The pit–fall of the model, therefore, is the neglect of strong correlations, which are vital for study of phase transitions, for example, where correlations become important as the critical point is approached.

With the preparation of cold Rydberg atoms, there has been renewed interest in low temperature Rydberg systems as a solution to the mean–field problem. Rydberg atoms are excited–state atoms with one or more electrons having a very high principal quantum number n. Such atoms interact with each other over large distances, which lead to strongly correlated many–body dynamics,

providing a platform to study strong correlations.

Figure 1.1: Van der Waals (vdW) interaction strength between two ground state strontium atoms (green solid line) compared to the total dipole–dipole interaction strength between two strontium atoms in the 60s and 80s state (blue and red solid lines, respectively). At interatomic separations of $\approx$ 4 μm, the interaction strength U between atoms in the 80s Rydberg state is 19 orders of magnitude greater than between ground state atoms. The coulomb interaction between two doubly–charged ions is also shown (black, solid line).

The interactions between Rydberg atoms scale with their principal quantum number n [3], as shown in figure 1.1, and are nearly 19 orders of magnitude greater than the interactions between ground state atoms. Furthermore, Rydberg states can be easily prepared via laser excitation from the ground state, meaning the strong interactions can be 'switched on' and 'off' at the researcher's request. Ergo, Rydberg atoms have been proposed as a tool with which to carry out quantum information processing [4], as well as single–atom and single–photon sources [5].

The key ideas that lead to a real explosion of the field are the dipole blockade and collective Rabi oscillations. For two nearby atoms, the dipole–dipole interactions between Rydberg atoms result in an energy shift of the doubly–excited state, as illustrated in figure 1.2. If the atoms are resonantly driven from the ground state $|g\rangle$ to the Rydberg state $|r\rangle$ via laser excitation ($\hbar\omega$), the doubly–excited state becomes off–resonant due to the interaction shift. If the interaction shift is greater than the combined linewidth (γ_b) due to the Rabi frequency of the driving laser field, Ω, the linewidth of the state Γ, and the linewidth of the excitation lasers γ, then double excitation is forbidden [3, 6].

The distance over which only one excitation is permitted is called the "blockade radius" (R_B), which is given as

$$R_B = \left(\frac{D_\theta C_6}{\gamma_b} \right)^{1/6} \quad , \tag{1.1}$$

where D_θ is due to the angular dependence of the interaction and the C_6 coefficient, which is n dependent, governs the strength of the long–range interaction.

Figure 1.2: The dipole blockade between two Rydberg atoms. The black line shows the shift of the doubly–excited state, with the horizontal dashed line a non–interacting doubly–excited state $|rr\rangle$. The red, shaded region shows the width of the largest of the driving Rabi frequency (Ω), state linewidth (Γ), or the laser linewidth (γ). R_B is the blockade radius, whereby the interaction shift is larger than the broadening mechanisms, allowing only one excitation to be present.

Consequently, a signature of the dipole blockade is a suppression of further excitation in a volume around the first Rydberg atom. Figure 1.3 illustrates this phenomenon with increasing ρ/ρ_B, where ρ is the ground state atomic density and ρ_B denotes the critical density, where blockade phenomena start to play a role, and is given by

$$\rho_B = \frac{3}{\pi R_B^3} \frac{\Delta^2}{\Omega^2} \quad , \tag{1.2}$$

where Δ is the detuning.

Figure 1.3: Suppression of further Rydberg excitation due to the dipole blockade. At low densities, the dipole–dipole interactions are linear with a slope governed by R_B, Ω and Δ (blue, solid line) since the interatomic distances R are much larger than R_B. At high densities, the dipole–dipole interactions tend towards saturation (green, solid line). In this regime, all N atoms are within one blockade sphere of volume $\frac{4}{3}\pi R_B^3$. ρ_B denotes the critical density (red, dash–dotted line) where blockade phenomena start to play a role.

At low densities the interatomic distances R are much larger than R_B, therefore the atoms are independent and the interactions are linear with a slope governed by R_B, Ω and Δ. The number of atoms N within one blockade sphere of radius R_B is determined by ρ, such that

$$N = \frac{4}{3}\pi R_B^3 \rho = \frac{\rho}{f\rho_B} \ , \tag{1.3}$$

where f is the fraction of Rydberg atoms within the ensemble.

Above ρ_B no further Rydberg excitation is possible and the contribution for each individual atom is thus effectively vanishing as the density increases. At very high densities, all N atoms lie within one blockade sphere of volume $\frac{4}{3}\pi R_B^3$. Due to the blockade, all N atoms share *one* Rydberg excitation forming a collective state $\frac{1}{\sqrt{N}}\sum_{i=1}^{N}|g_1, g_2, ..., r_i, ..., g_N\rangle$. This state is the so–called *superatom*, which is coupled by the light field to the ground state with a collectively enhanced Rabi frequency $\sqrt{N}\Omega$ [3]. This mesoscopic entangled sample has unique applications in quantum information and entanglement enhanced measurement [6]. The suggestion to use Rydberg atoms and the blockade mechanism to produce a fast, neutral quantum gate was put forward by Jaksch *et al.* [7]. Using Rydberg atoms for quantum information has several advantages over using

quantum dots, superconductors or ions, such as the relatively long decoherence time. Dipole blockade has been experimentally observed in many systems [8–10]. In particular Schempp *et al.* [11] illustrate the excitation blockade effect by measuring a suppression in the Rydberg signal detected with changing the ground state density, which gives a similar result to that shown in figure 1.3.

Many of the Rydberg atom properties depend on n, not merely the dipole–dipole interaction strength, which scales as n^{11}. These include the radiative lifetime (n^3), orbital radius (n^2) and polarisability (n^7) [3] for low angular momentum states. On the last account, Rydberg atoms can be polarised very easily and show very large electric dipole moments when placed within small electric fields, making Rydberg atoms a useful tool in sensing devices, [12] for example.

To date, resonantly exciting to Rydberg states has proven to be an ideal platform to study strong interactions in many–body systems on short timescales, where atom manipulation occurs within a timescale set by the radiative lifetime of the excited states. During this time, the atoms appear to be 'frozen' in space. As a result, there have been experimental breakthroughs in investigations of long–range interacting quantum many–body physics [13], which includes dynamical crystallisation whereby the dipole blockade is manipulated to produce a well–controlled, self–ordered crystalline structure via tailored light pulses, and quantum non–linear optics [14], where the strong interactions are mapped into strong interactions between individual optical photons.

To model strongly correlated systems, powerful numerical tools beyond that of the mean–field theory are required, and, since the advent of the dipole blockade phenomenon, advances in theoretical models of strong interactions in many–body Rydberg gases have been developed. As a consequence, new intriguing many–body effects have arisen in such systems from theoretical work that will permit the study of interesting phases in ultracold gases, such as supersolids and solitons [15–17]. These phases require an interrogation time that extends beyond the frozen gas regime, which can be achieved by weakly dressing the atomic ground state with a small fraction of the Rydberg state, given by

$$f \approx \frac{\Omega^2}{\Delta^2} \ , \tag{1.4}$$

whereby the Rydberg fraction is greatest using small detunings and large Rabi frequencies to the Rydberg state.

The overall lifetime of the dressed ground state system is then enhanced by $1/f$. Theoretical proposals for a Rydberg–dressed ground state system in a Bose–Einstein Condensate (BEC) have recently increased interest in this excitation scheme [15, 16]. Extending this scheme to atomic lattices also promises a number of intriguing perspectives [18], such as spin–squeezing in optical lattice clocks [19], where strong correlations between the atoms can be used to beat the limit due to quantum projection noise and improve the signal–to–noise ratio.

However an experimental realisation of Rydberg–dressing has been elusive thus far. To dress an atomic ground state, the maximum effect is expected in the intermediate density regime where the energy difference per atom between the blockaded system and the non–blockaded system is highest, i.e. $\rho = \rho_B$, and for $f << 1$. Working in this density regime can be difficult to achieve and reducing the atomic density can be experimentally challenging, particularly in BECs where typical densities are 10^{13}–10^{14} atoms/cm^3. Instead of reducing the atomic density, other parameters can be varied to control f including: reducing R_B by exciting to a Rydberg state with lower principal quantum number, since this increases the atomic density where saturation occurs as a consequence of the $C_6^{1/6}$ dependence on R_B [3]; increase the excitation Rabi frequency Ω; or excite using very far–detuned radiation. Nonetheless, there are drawbacks to the first two suggestions. The radiative lifetime of the Rydberg state decreases (since it is $\propto n^3$) thereby reducing f, and there is a limit to Ω before scattering dominates. The final suggestion, however, should be possible through excitation to high–lying Rydberg states via very narrow linewidth transitions.

Figure 1.4: Two–photon, three–ladder Rydberg excitation scheme. Atoms in the ground state $|g\rangle$ are excited to the intermediate state $|i\rangle$ via a single–photon from the probe laser at λ_p. The atoms in state $|i\rangle$ are excited to the Rydberg state $|r\rangle$ via a single–photon at λ_c from the coupling laser. The probe (Δ_p) and coupling (Δ_c) laser detunings are also shown.

To excite to Rydberg states, a two–photon, three–level ladder scheme is often employed as the excitation wavelengths required are in the region of the electromagnetic spectrum that can be achieved via laser excitation, as shown in figure 1.4. Excitation from the ground state $|g\rangle$ to the intermediate state $|i\rangle$ requires a photon at λ_p from the probe laser, while excitation from the intermediate state to the Rydberg state $|r\rangle$ requires a photon at λ_c from the coupling laser. For small Δ_p, including on–resonance probe light, the linewidth of the entire excitation transition ($|g\rangle \to |r\rangle$) is determined by the linewidth of the $|g\rangle \to |i\rangle$ transition. As Δ_p increases, the entire excitation transition linewidth reduces as the $|i\rangle$ state is adiabatically eliminated [20]. Owing to the long–lived lifetimes of $|r\rangle$ states, $|i\rangle \to |r\rangle$ transition linewidths are typically tens of kilo–Hertz [21], therefore it is ideal to have an excitation linewidth of this magnitude from $|g\rangle \to |i\rangle$. Such narrow (i.e. kilo–Hertz linewidth) transitions are not common in all atoms, particularly in the alkali atoms, however they do exist in alkaline–earth atoms as a consequence of mixing of the singlet–triplet states. This places strontium, the atom of focus in this thesis, at the forefront of this research.

1.1   Motivation

Almost all quantum gas experiments work with alkali atoms, including dipole blockade experiments, which has been observed in extended ensembles [10, 11, 22, 23]. Such atoms have a comparatively simple electronic structure compared to alkaline–earth elements, which feature a handful of remarkable properties, such as long–lived metastable states, narrow singlet–triplet transitions, and a ground state with no hyperfine structure. Atoms with two electrons offer a new approach in Rydberg experiments as the inner valence electron provides an additional way to probe and manipulate Rydberg atoms. Recent experiments have shown that the inner electron can be used as a fast, state–selective probe of the interactions in a cold Rydberg gas [21, 24] and theory predicts that the polarisability of the additional electron also enables tight, magic–wavelength traps for Rydberg atoms [25].

Recent calculations of the interactions for strontium reveal interesting features that differ from the alkali metals. There exist both attractive and repulsive isotropic interactions depending on the Rydberg spin symmetry, i.e. $^1\mathrm{S}_0$ or $^3\mathrm{S}_1$ [26] (see figure 1.5 for the atomic energy level scheme in strontium). The ability to study attractive isotropic interactions is unique to strontium and is easily accessible via the 5s5p $^1\mathrm{P}_1$ state.

Figure 1.5: Atomic energy level scheme in strontium. Shown are the main cooling transitions at 461 nm and 689 nm, the clock transition at 698 nm, and the Rydberg excitation schemes via the 5s5p ^{1}P$_1$ state and 5s5p ^{3}P$_1$ state.

Owing to the narrow singlet–triplet transitions, the experimental limitations experienced in alkali atom Rydberg experiments, such as $f << 1$ where f is given by equation 1.4, are irrelevant to alkaline–earth atom experiments. Hence, Rydberg–dressing via these narrow transitions should be possible due to the experimental ease of achieving large detunings. Furthermore, unlike the alkali atoms, Rydberg and ground state strontium atoms can be trapped simultaneously in the same optical lattice [25], warranting the study of Rydberg–dressing [18] and spin–squeezing [19] in optical lattices.

Although strontium has many redeeming properties, to study strong–correlations in a many–body system a cold, dense atomic ensemble is required, such that experiments in the frozen gas regime can be performed. The focus of this thesis is to produce an atomic ensemble of strontium atoms that are: (1) ultra–cold such that the atoms appear frozen in space during Rydberg excitation; (2) have sufficiently high density such that dipole blockade phenomena can be studied.

1.2   Context of this thesis

Prior to the work in this thesis, the experimental apparatus was designed and constructed, a Rydberg excitation scheme, via the $5s^2\,^1S_0 \rightarrow 5s5p\,^1P_1$ transition, and detection scheme were developed [21], and spatial and temporal [24] techniques were implemented during the course of J. Millen's [27] and G. Lochead's [28] Ph. D. theses. Rydberg experiments were performed in a magneto–optical trap (MOT) on the main cooling transition at 461 nm, where typical cloud densities are 10^9 atoms/cm^3. Working at a Rydberg state of 5s75d $^1D_2$ via the 461 nm line, where the blockade density is $\approx$ 6 x 10^9 atoms/cm^3 (this density has been calculated using the C_6 coefficients stated in [26]), we were unable to observe blockade (see figure 1.6). The Rydberg signal clearly follows a linear trend with ground state density, whereas in an interacting Rydberg ensemble a saturation of the Rydberg signal with increasing density would be observed [11]. In addition, the atoms in this MOT are typically at a temperature of 5 mK translating to a mean propagation distance of 1 μm in 1 μs. The laser excitation pulses used are typically 1 μs, hence we are not working in the frozen gas regime in this MOT.
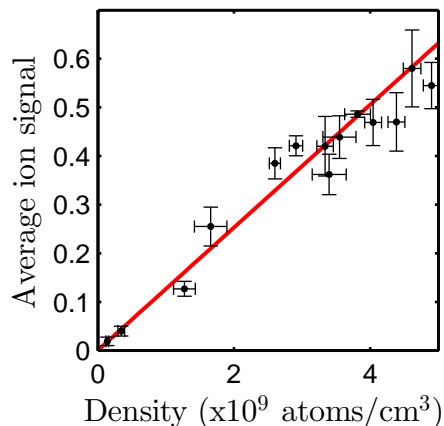
Figure 1.6: The average ion signal of the 5s75d $^1D_2$ state as a function of ground state density. The red, solid line is a linear fit to the data. Figure has been taken and modified from reference [28].

There does exist another cooling transition in strontium that will not only reduce the cloud temperature, but also increase the density by three orders of magnitude. This is achieved via cooling on the narrow singlet–triplet line at 689 nm (see figure 1.5), which is a standard protocol in laser–cooled strontium

experiments. The work in this thesis focuses on designing and building the narrow linewidth 689 nm laser system, as well as the cooling procedure performed to reach the long hoped–for density. The laser system will also form the basis of narrow–line Rydberg excitation, which will be used in future Rydberg–dressing experiments. The thesis will conclude with the first 'known' observation of blockade in a two–electron atom ensemble.

1.3   Thesis outline

The work presented in this thesis aims to give a review of the steps taken to successfully achieve an ultra–cold, ultra–dense sample of strontium atoms along with the first observation of the dipole blockade in a two–electron gas.

- Following an introduction to Rydberg physics, there will be an introduction to the alkaline–earth element, strontium. In this chapter, current strontium research will be discussed, including a discussion on narrow–line laser cooling dynamics and laser frequency stability requirements.

- In chapters 3 and 4 we introduce the reader to the concept of frequency stabilisation techniques for reducing the laser frequency instability on short– and long–interrogation times. In addition to the technical details of laser frequency stabilisation, a discussion regarding the laser linewidth evaluation is also included.

- Chapter 5 includes a discussion of the experimental apparatus and other laser systems used for producing the MOT on the 461 nm cooling line, the experimental sequence necessary for introducing the second cooling stage on the 689 nm transition, and finally the experimental evaluation of the MOT characteristics, including a comparison with theory.

- Chapter 6 contains the laser system and Rydberg experimental sequence, as well as the first 'known' observations of Rydberg blockade in a two–electron gas.

- Finally, in chapter 7 we conclude the findings in the thesis and give an outlook to future work.

1.4 Publications

Number–resolved imaging of excited–state atoms using a scanning autoionization microscope. Phys. Rev. A **87**, 053409 (2013). G. Lochead, D. Boddy, D. P. Sadler, C. S. Adams, and M. P. A. Jones

Chapter 2

Laser cooling strontium and its applications

Strontium can be cooled to a few hundreds of nK through optical methods alone owing to narrow linewidth optical transitions from the ground state as a consequence of two valence electrons. Cooling dynamics on these transitions differ greatly from conventional Doppler cooling, but can be modelled rather simply. A discussion on the electronic properties and the narrow intercombination transitions will be covered, in particular the cooling theory and experimental requirements for laser cooling on an intercombination line.

2.1   Strontium properties

Strontium ($Z = 38$) is an alkaline–earth element. It is a divalent atom (i.e. possesses two valence electrons) that has four naturally occuring isotopes, three of which are bosonic (all of which have I $= 0$), whilst the fourth is fermionic (I $= 9/2$) – ^{88}Sr, ^{86}Sr, ^{84}Sr and ^{87}Sr, listed in order of their natural abundance and statistic. The two valence electrons lie in the outer s–shell and the spins of these electrons can be aligned anti–parallel or parallel, permitting singlet (S $= 0$) or triplet (S $= 1$) states, respectively. The atomic ground state of ^{88}Sr, which is the focus of this thesis, is a spin singlet J $= 0$ state ($5s^2$ $^1S_0$), which removes the complications of ground state optical pumping, as well as making sub–Doppler cooling ineffective. Strontium ($Z = 38$) is an alkaline–earth element. It is a divalent atom (i.e. possesses two valence electrons) that has four

naturally occuring isotopes, three of which are bosonic (all of which have I = 0), whilst the fourth is fermionic (I = 9/2) – ^{88}Sr, ^{86}Sr, ^{84}Sr and ^{87}Sr, listed in order of their natural abundance and statistic. The two valence electrons lie in the outer s–shell and the spins of these electrons can be aligned anti–parallel or parallel, permitting singlet (S = 0) or triplet (S = 1) states, respectively. The atomic ground state of ^{88}Sr, which is the focus of this thesis, is a spin singlet J = 0 state ($5s^2$ $^1S_0$), which removes the complications of ground state optical pumping, as well as making sub–Doppler cooling ineffective.

Most atomic transitions are electric dipole allowed, which obey the ΔS = 0 dipole selection rule [29]. Dipole allowed transitions from the ground state to low–lying states, which are typically used in laser cooling, are characterised by broad linewidths (of the order of tens of MHz) and are in the visible range of the electromagnetic spectrum. Figure 2.1 illustrates the transitions between the low–lying states in strontium. Shown is a broad, singlet–singlet transition on the $5s^2$ $^1S_0 \rightarrow 5s5p$ $^1P_1$ line at 461 nm, which has a lifetime of (5.22 ± 0.03) ns [30] (translating to a linewidth of 2π x 30.2 MHz). This is the main cooling transition and is often referred to as a 'blue MOT' owing to the colour of the cooling light. Due to the simplicity of the ground state, the model of a two–level atom is nearly ideally realised.

Figure 2.1: Atomic level diagram depicting the two main cooling transitions (at 461 nm and 689 nm), the clock transition (at 698 nm) and the repump transitions (at 679 nm and 707 nm).

Although transitions between singlet and triplet states are 'forbidden', transitions are more weakly allowed. The electric dipole selection rules apply for lighter elements, such as helium, but as the atomic mass increases there is LS–coupling breakdown, and spin–orbit, and other relativistic effects, induce mixing of the states with the same total angular momentum, J, and different spin, S [29].

The most exploited triplet states in strontium are the $5s5p$ $^3P_{0,1}$ states, which are commonly referred to as the metastable states, as they are very long–lived. The singlet–triplet transitions, otherwise known as intercombination transitions, are on the $5s^2$ $^1S_0 \rightarrow 5s5p$ ^3P_i lines where $i = 0$, 1, 2 and are characterised by very narrow linewidths.

In the case of the $5s^2$ $^1S_0 \rightarrow 5s5p$ $^3P_1$ spin forbidden transition, the $5s5p$ $^3P_1$ state acquires a small admixture of the $5s5p$ $^1P_1$ wavefunction and with it a weak electric dipole moment. The $5s5p$ $^3P_1$ state has a lifetime of 21.4 μs with a corresponding linewidth of 2π x 7.4 kHz [31]. Since there is only one decay route back to the ground state and, as a result of its narrow linewidth, this line is well suited to narrow–line Doppler cooling, where temperatures as low as ≈ 200 nK have been achieved through optical cooling methods alone. Narrow–

line cooling of strontium was established by the Tokyo group and later studied extensively by the group at JILA [32, 33]. A detail discussion will be covered in section 2.2.

The 5s5p $^3P_0$ state is the clock state, where the $5s^2$ $^1S_0$ $\rightarrow$ 5s5p $^3P_0$ (i.e. $\Delta S = 1$, $J = 0 \rightarrow J = 0$) transition is doubly forbidden by both spin and angular momentum selection rules. For the bosons, the hyperfine interaction, which is responsible for breaking the angular momentum selection rule on the angular momentum, is absent and the lifetime of the state is calculated to be thousands of years. However, courtesy of the nuclear–spin–induced hyperfine interaction, the doubly–forbidden transition exists in the fermionic isotope with a linewidth $<< 1$ Hz, which is well suited for the clock transition. Using strontium in an optical lattice clock has produced the best agreed upon frequency standard, and has become a strong contender as the re–definition of the SI second [34]. Such clocks are currently in operation in laboratories worldwide, including JILA (Boulder, USA), LENS (Florence, Italy) and SYRTE (Paris, France), and further clocks still are being set–up (e.g. NPL (UK)). Improving the fractional stability of clocks is a key importance in the evolution of GPS technology, and, due to the sensitivity of quantum states to external perturbations, ultra–cold atoms make for incredibly precise sensors too. In addition, researchers are using the sensitive nature of cold atoms to approach new frontiers in gravitometry and magnetometry. The group at LENS are producing a transportable strontium lattice optical clock that can be used to measure the gravitational field in outer space [35].

A new addition to the quantum mechanical toolbox is quantum degenerate gases. By advancing cooling stages beyond the conventional laser cooling in a MOT, researchers are able to reach ultra–cold degenerate gas regime creating Bose–Einstein condensates (BECs) [36, 37] and degenerate Fermi gases. Recently (2009) two groups have created a strontium BEC [38] and later a Bose–Fermi mixture of strontium [39], with their goal being to create and investigate novel quantum systems that are beyond the reach of alkali quantum gases. The drive for their research is quantum computation and quantum simulation of many–body systems. On account of the nuclear spin in ^{87}Sr, information can be stored in a well–protected manner, while at the same time, owing to the complex electronic structure of strontium, the information can be manipulated

and extracted.

Regardless of the scientific quest, the strontium experiments discussed all utilise the same, two–stage laser cooling procedure (the second stage via the narrow $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition) in order to achieve typical atom temperatures of (1–2) μK. Owing to the very narrow linewidth, the cooling dynamics on this transition differ from that of standard Doppler cooling theory. An in–depth discussion is provided in the following section.

2.2   Laser cooling theory

Laser cooling is a well established technique, which can be used to create atomic samples at very low temperatures. Broad lines, historically used for typical laser cooling experiments with alkali atoms, are defined by $\Gamma/\omega_R >> 1$, where Γ is the atomic transition linewidth and ω_R ($= \frac{\hbar k^2}{2m}$, where $k = 2\pi/\lambda$ is the wavevector of the light field and λ the wavelength, m is the atomic mass, and $2\pi\hbar$ is Planck's constant) is the single–photon recoil frequency shift that originates as a consequence of conservation of momentum and energy as the atom absorbs the photon momentum [40]. When cooling on a broad line, the natural energy scale is governed by the power–broadened linewidth Γ_E, where

$$\Gamma_E = \Gamma\sqrt{1 + s}\ , \tag{2.1}$$

is defined by the saturation parameter $s = I/I_{\text{sat}}$ and $I(I_{\text{sat}})$ is the single–beam peak intensity (transition saturation intensity). Gravity is essentially negligible on these transitions since the ratio of the maximum radiative force ($\hbar k\Gamma_E/2$), which is dependent on s, to the gravitational force (mg), is typically on the order of 10^5, where g is the gravitational acceleration. Generally, the limit of the attainable temperature of the MOT is set by the Doppler limit of T_D when sub–Doppler cooling is absent, i.e.

$$k_B T_D = \frac{\hbar\Gamma}{2}\ , \tag{2.2}$$

where k_B is the Boltzmann constant. The cooling dynamics on these lines can be explained by standard Doppler cooling theory, whereby the scattering force provides a dispersion shaped feature [40].

In contrast to broad lines, narrow Doppler cooling lines are characterised by

$\Gamma/\omega_R \approx 1$. The minimal attainable temperature is determined by the photon recoil temperature, $T_R/2$ [41], where $T_R > T_D$, and the cooling forces are comparable to the local gravity (i.e. $a_{\max}/g \to 1$). The cooling dynamics on these lines deviate from that of semiclassical physics, as discussed in section 2.2.1.

Characteristics of the first stage MOT on the $5s^2\ ^1S_0 \to 5s5p\ ^1P_1$ transition, including the Doppler temperature given by equation 2.2, and the second stage MOT on the $5s^2\ ^1S_0 \to 5s5p\ ^3P_1$ transition are compared in table 2.1.

Parameter	First stage cooling	Second stage cooling
λ	461 nm	689 nm
$\Gamma/2\pi$	30.2 MHz	7.4 kHz
ω_R	$6.7 \times 10^4\ \mathrm{s}^{-1}$	$3.0 \times 10^4\ \mathrm{s}^{-1}$
I_{sat}	$42.7\ \mathrm{mW\,cm}^{-2}$	$3\ \mu\mathrm{W\,cm}^{-2}$
$a_{\max}$	$9.34 \times 10^5\ \mathrm{m\,s}^{-2}$	$157.2\ \mathrm{m\,s}^{-2}$
$T_D = \hbar\Gamma/(2k_B)$	725 μK	179 nK
Γ/ω_R	3×10^3	1.6
$T_R = \hbar^2 k^2/(k_B m)$	1.02 μK	458 nK

Table 2.1: Comparison between broad and narrow line MOTs in strontium, where Γ is the transition linewidth, T_D is the Doppler temperature, T_R is the single–photon recoil temperature, $\omega_R = \hbar k^2/2m$ is the single–photon recoil frequency shift ($k = 2\pi/\lambda$ is the wavevector of the light field and λ the wavelength, m is the atomic mass, and $2\pi\hbar$ is Planck's constant).

The minimum achievable temperature of the first stage of cooling is 725 μK, whereas the minimum temperature of the second stage of cooling is more than three orders of magnitude colder. Furthermore, there is also an increase in density from the first cooling stage to the second by a similar order of magnitude, which will be discussed in the next section.

2.2.1   Narrow line cooling dynamics

Since the Doppler and recoil temperature limits almost coincide when cooling on narrow lines, the kinetic energy of the atom depends on s, permitting an opportunity to investigate both semiclassical and quantum mechanical cooling processes. Hence the trap dynamics can be separated in three regimes:

(I) $s >> 1$ and $\delta > \Gamma_E$; (II) $s >> 1$ and $\delta < \Gamma_E$; and (III) s $\approx$ 1.

To grasp an understanding of these regimes, it is helpful to visualise the geometric region in which the atoms interact with the MOT light. In regime (I) we assume strong saturation and a large detuning, such that $\delta > \Gamma > \Gamma_E$. The light is only resonant with the atomic transition in regions where the Zeeman shift induced by the quadrupolar magnetic field ($m_J g_J \mu_B B$) balances the detuning, i.e.,

$$\delta = m_J g_J \mu_B B \ , \tag{2.3}$$

where $\mu_B = 1.4$ MHz/G is the Bohr magneton, and the Landé g–factor is $g_J = 1.5$ for the 5s5p ^{3}P$_1$ state. This forms a boundary for the atoms in the shape of an on–resonant ellipsoid shell with a vertical radius, ϵ, given by

$$\epsilon = \left[\frac{\delta}{m_J g_J \mu_B} \right] \left(\frac{\partial B}{\partial z} \right)^{-1} \ , \tag{2.4}$$

where $\frac{\partial B}{\partial z}$ is the quadrupolar magnetic field gradient along the vertical z direction. This regime is illustrated in figure 2.2 for various detunings and with $s = 250$. The atoms can be thought of as trapped within the shell and only interact with the shell boundary, as opposed to the traditional dispersion shaped force curve in strong cooling [40]. In (a) the scattering force is calculated as a function of position (assuming a velocity of 0 m s^{-1}) for cooling along the z–axis at different detunings. As the detuning increases, the ellipsoid shell volume increases with the thickness of the shell fixed by s, as we shall discuss later. Since the atoms interact with a single beam at each end of the trap, they experience no scattering in the central trapping region.
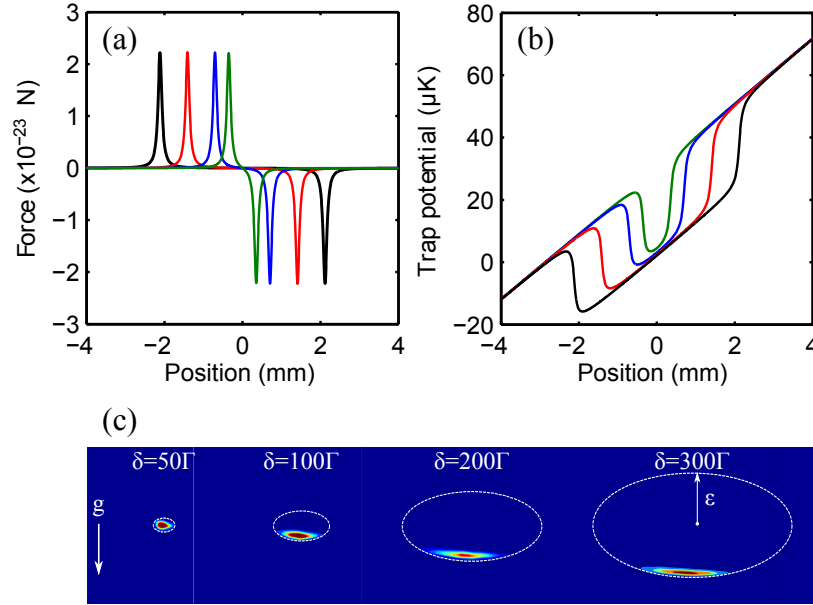
Figure 2.2: Semiclassical radiative force as a function of position of the 689 nm red MOT along the z–axis for (a) $s = 250$ and $\delta = -50\Gamma$ (green solid line), $\delta = -100\Gamma$ (blue solid line), $\delta = -200\Gamma$ (red solid line), and $\delta = -300\Gamma$ (black solid line). For small detunings the force adopts the typical dispersion shape, whereas at large detunings the MOT beams operate independently as boundaries. (b) Significance of gravity on the potential energy of the trap along the z–axis. (c) 'False' colour fluorescence images of the MOT cloud for the saturation parameter and detunings corresponding to those in (a) and (b). The resonant ellipsoid shell, which is centred at the magnetic field zero, with a vertical radius given by equation 2.4 is also displayed, as well as the direction of the gravity g. A velocity of 0 m/s and an axial magnetic field gradient of 5 G/cm is used (2.5 G/cm along z).

The effects of gravity play a significant role in this regime. Illustrated in figure 2.2 (b) is the significance of gravity on the potential energy of the trap along the z–axis. The inclusion of the gravitational potential results in a potential minimum in the trap near the resonant position of the lower beam. Therefore the atoms sag to vertical positions where the Zeeman shift balances δ.

'False' colour fluorescence images of the MOT cloud are shown in figure 2.2 (c) for the saturation parameter and detunings corresponding to those in (a) and (b). The resonant ellipsoid shell, which is centred at the magnetic field zero, with a vertical radius given by equation 2.4 is also displayed. The atoms sag

below the magnetic field zero due to the effects of gravity and form a lenticular cloud with increasing detuning. Gravity plays a large role in the temperature of the atoms in this regime. For standard cooling, we expect a strong detuning dependence on temperature. However, the detuning doesn't change the scattering rate, it only changes the shell width, ϵ as discussed. As a consequence, the MOT geometry can be manipulated without effecting the atom temperature.
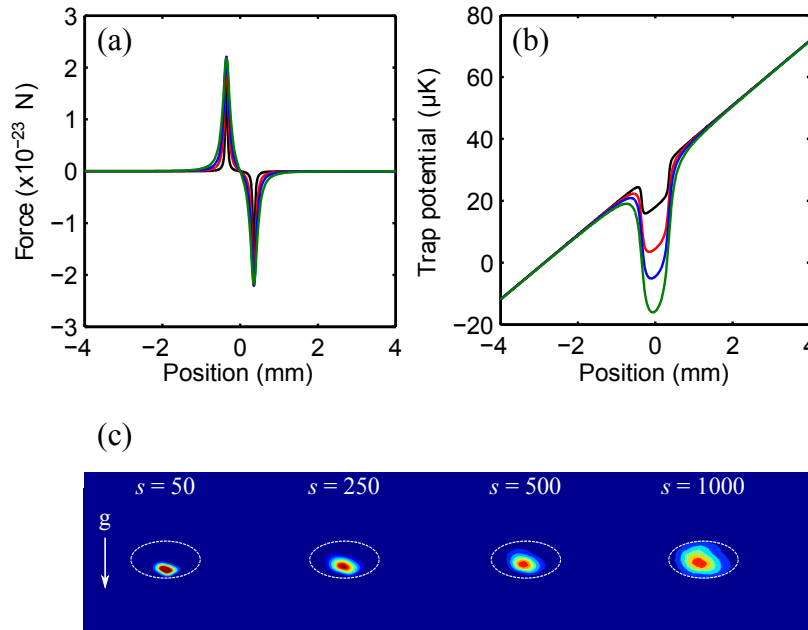
Figure 2.3: Semiclassical radiative force as a function of position of the 689 nm red MOT along the z–axis for $\delta = -50\Gamma$ and $s = 50$ (black solid line), $s = 250$ (red solid line), $s = 500$ (blue solid line), and $s = 1000$ (green solid line). For low intensities, the single beam hard wall properties of the MOT is apparent once more, indicative of regime (I). As the saturation parameter is increased, the cooling conditions evolve into the standard semiclassical cooling described by regime (II). (b) Significance of gravity on the potential energy of the trap along the z–axis. (c) 'False' colour fluorescence images of the MOT cloud for the detuning and saturation parameters corresponding to those in (a) and (b). The resonant ellipsoid shell, which is centred at the magnetic field zero, with a vertical radius given by equation 2.4 is also displayed, as well as the direction of the gravity g. A velocity of 0 m/s and an axial magnetic field gradient of 5 G/cm is used (2.5 G/cm along z).

If the laser detuning is reduced below Γ_E (or if the power broadened linewidth is made larger than the detuning), the cooling is described in regime (II), where $\delta < \Gamma_E \gg \Gamma$. Figure 2.3 illustrates this regime where the saturation param-

eter is varied over a large range with fixed detuning –50Γ. For low intensities, the single beam hard wall properties of the MOT is apparent once more as in regime (I). However as the saturation parameter increases, the cooling dynamics evolve into standard semiclassical physics, whereby cooling behaves as in a typical strong transition and the scattering force provides a dispersion shaped feature. Since the scattering force, which is determined by Γ_E, is large, the effect of gravity is negligible and we do not observe a change in MOT position.

Finally, the most exotic cooling regime is that of (III), where $\Gamma_E = \Gamma \approx \omega_R$. In this regime, the natural energy scale is governed by ω_R, since $\Gamma_E \approx \omega_R$, and trap dynamics are governed by single–photon recoils, i.e. quantum physics. When $s \approx 1$ and $\delta \to 0$, temperatures as low as $\Gamma_R/2$ [41] have been measured [42].

2.2.2 Laser cooling neutral strontium

Strontium has negligible vapour pressure at room temperature, therefore to obtain a sufficient number of atoms, we chose to heat the metal to ≈ 900 K to produce an atomic beam with a mean velocity of ≈ 500 m/s, further discussion follows in chapters 4 and 5. It is not possible to slow these atoms via the $5s^2\ ^1S_0 \to 5s5p\ ^3P_1$ transition due to the comparatively low photon scattering rate (refer to table 2.1) and capture velocity. Instead the atoms are slowed and cooled via the $5s^2\ ^1S_0 \to 5s5p\ ^1P_1$ transition by means of a Zeeman slower and a first stage MOT. The atoms are then transferred to a second stage of cooling on the $5s^2\ ^1S_0 \to 5s5p\ ^3P_1$ transition, which is referred to as the 'red' MOT owing to the colour of the cooling light.

However, when a laser interrogates an atomic transition, it is important that most of the power is concentrated around the carrier frequency. Since the $5s^2\ ^1S_0 \to 5s5p\ ^3P_1$ atomic transition linewidth is 2π x 7.4 kHz, the MOT cooling light must have a linewidth on the order of this. A discussion on the dominant noise sources in lasers, which lead to a finite laser linewidth, including methods of reducing the noise, follows in section 2.3.

Assuming one has reduced the laser noise, by virtue of the narrow laser linewidth the atoms are still too hot to load and trap a sufficient number into the second stage MOT, even though the first stage of slowing and cooling reduces the temperature of the atoms from ≈ 900 K down to several mK. To increase the

transfer efficiency, the laser light is artificially broadened to match the Doppler–broadened profile of the atoms emanating from the blue MOT, which is typically several MHz wide. This phase is often referred to throughout the text as the *broadband* red MOT. Once a sufficient number of atoms have been transferred from the blue MOT, the broadening is removed, leaving single–frequency red MOT light only. This phase is often referred to as the *narrowband* red MOT.

2.3 Laser frequency stabilisation

The natural frequency linewidth of the laser radiation is broadened by the several random physical phenomena that occur in the lasing process and in the laser oscillator. The presence of phase modulation due to noise processes leads to a spreading of the emission power among a continuum of frequencies, and thus to a finite emission linewidth of the laser oscillator. If the broadened laser linewidth is much greater than that of the atomic transition linewidth, spectroscopic resolution will be limited and narrow line cooling to the recoil temperature will not be possible. The remainder of this section discusses the dominant sources of noise in laser oscillators and methods of reducing and measuring noise.

2.3.1 Frequency noise

The electromagnetic field emitted by a laser can be viewed as an optical oscillator (i.e. characterised by an amplitude $A(t)$ and a time dependent phase factor ψ) such that,

$$E(t) = A(t) \exp i\psi(t) = A(t) \exp \left[i(2\pi\nu_0 t + \Phi(t)) \right] \ , \tag{2.5}$$

where ν_0 is the carrier frequency of the optical oscillator, $A(t)$ is the slowly (with respect to $1/\nu_0$) varying field amplitude, and $\Phi(t)$ is the instantaneous phase modulation of the light wave. The oscillating field frequency $\nu(t)$ is defined as the instantaneous variation of the phase i.e.,

$$\nu(t) \equiv \frac{1}{2\pi} \frac{d\psi}{dt} = \nu_0 + \frac{1}{2\pi} \frac{d\Phi}{dt} \ , \tag{2.6}$$

where the frequency noise of the field $\partial\nu(t)$ is represented by $\frac{1}{2\pi}\frac{d\Phi}{dt}$ [43]. Therefore, since the description of the laser field is related to the functions $\Phi(t)$, $\partial\nu(t)$ and $A(t)$, which are random functions of time, it is possible to perform phase sensitive detection to extract the individual noise components.

It is important, especially for short–interrogation times, to use laser frequency instability analysis in the frequency domain as it is directly related to the underlying noise processes affecting the oscillator frequency instability [43, 44]. Spectral analysis is particularly useful when a specific source of frequency noise is of interest and, in practice, it has been shown that the random fluctuations, which yield phase and frequency fluctuations, are the sum of five independent noise processes: *white phase*; *flicker phase*; *white frequency*; *flicker frequency*; and *random–walk frequency* noise (i.e. drift). All of these processes have different origins and effects in the diode laser linewidth. Such measurements in the frequency domain can be achieved via a beat measurement between identical laser oscillator systems.

For long interrogation times, it is common practice to implement measurements in the time–domain. This is performed by measuring the Allan variance $\sigma_y^2(\tau)$, which is defined as follows [43]

$$\sigma_y^2(\tau) = \left\langle \frac{(\bar{y}_{k+1} - \bar{y}_k)^2}{2} \right\rangle \; , \qquad (2.7)$$

where $\bar{y}_k$ is the fractional frequency fluctuation averaged over the time interval τ.

The rate at which an instability averages down depends on the source of the noise. Figure 2.4 shows a log–log plot of the Allan deviation as a function of the sampling time τ for a source that exhibits all five common noise types. The slopes of the white phase noise and the flicker phase noise are the same, therefore cannot be separated using this plot [45].
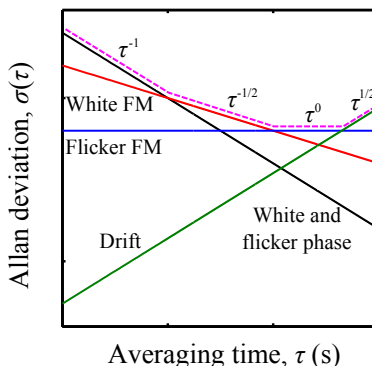
Figure 2.4: Log–log plot of $\sigma_y^2(\tau)$ showing the correspondence with the type of noise. The resulting slope including all five noise processes is also shown (pink, dashed line).

In the optical regime, such measurements can be achieved by measuring the beat between the laser light in question and a frequency comb stabilised to a hydrogen maser [33].

2.3.2   Linewidth and frequency noise in diode lasers

The dominant source of frequency fluctuations in diode lasers is due to the quantum mechanical nature of spontaneous emission of photons, i.e. white frequency noise. This source of noise results in a fundamental (quantum) limit for the linewidth of a laser known as the Schawlow–Townes linewidth [46], which is set by the length of the gain medium within the diode. For a typical diode, the Schawlow–Townes limit is in the region of a few tens of MHz. The next dominant contribution to the laser linewidth is flicker frequency noise, which is mainly due to technical effects/sources such as mechanical fluctuations, current supply noise, etc.

By increasing the length of the diode cavity, the Schawlow–Townes linewidth can be reduced. This is achieved by seating the diode within an external cavity, where a diffraction grating is one mirror and the diode is the other. A small fraction of light, which is diffracted off of the grating at an angle dependent on the laser frequency, is fed back into the laser diode and the frequency is tuned by varying the length of the cavity. A laser diode in this set–up is commonly referred to as an External Cavity Diode Laser (ECDL). The dominant noise

process is no longer due to the Schawlow–Townes limit, but rather due to flicker frequency noise, which leads to a finite laser linewidth $\delta\nu_I$ that is typically 500 kHz [47, 48].

Figure 2.5 aids the reader visually when discussing the different laser linewidths in the following few paragraphs. In the context of this thesis, $\delta\nu_I$ is commonly referred to as the instantaneous laser linewidth. The inclusion of the external cavity in the laser system introduces an additional linewidth constraint to that of $\delta\nu_I$ – a short–term laser linewidth $\delta\nu_S$ – that arises due to small cavity length changes as a result of temperature drifts over the period of several hours. Typically $\delta\nu_S \approx$ a few tens of MHz over several minutes.

Figure 2.5: The instantaneous laser linewidth $\delta\nu_I$ (black, solid line) arises due to white and flicker frequency noise in the laser diode. This limit can be reduced by seating the diode in an ECDL configuration. Due to small cavity length changes, there is a short–term instability, which occurs over several minutes (grey, solid lines), introducing a short–term frequency linewidth $\delta\nu_S$ (red, solid line). $\delta\nu_I$ and $\delta\nu_S$ can be stabilised to a high–finesse optical cavity (discussed further in chapter 3), which is subject to a long–term drift over several hours due to temperature and pressure variations, hence an additional long–term laser linewidth is introduced (blue, solid line). The linewidth is defined by $\delta\nu$, which corresponds to the full–width–at–half–maximum (FWHM) of the Lorentzian emission line.

To reduce both $\delta\nu_I$ and $\delta\nu_S$ to the order of the 5s^2 ^{1}S$_0$ $\rightarrow$ 5s5p ^{3}P$_1$ transition linewidth, a suitable stable frequency reference, which can measure the frequency fluctuations due to white, flicker and random walk noise, is required.

Active stabilisation, via electronic feedback, to the ECDL can compensate these fluctuations. Furthermore, since the noise processes resulting in $\delta\nu_I$ are linearly decreasing with averaging time, whereas those resulting in $\delta\nu_S$ are increasing with averaging time (refer to figure 2.4), the electronic feedback circuit can stabilise $\delta\nu_I$ and $\delta\nu_S$ independently. A convenient choice for a frequency diagnostic on such short–timescales is a high–finesse optical cavity, which consists of two highly reflective mirrors separated by a spacer. An in–depth discussion on the technique, including more technical details such as the design of the cavity and the feedback electronics, is provided in chapter 3.

Once $\delta\nu_I$ and $\delta\nu_S$ have been stabilised, the stability is now dependent on the drift of the high–finesse cavity length. The cavity length drift occurs over much longer timescales than the drift of the ECDL cavity length. Long–term stability is reduced by environmental temperature and pressure fluctuations, via thermal expansion properties of both the spacer and the mirrors, and typically occurs over many hours or days. To stabilise the cavity drift over the duration of a typical experimental cycle (which can be several hours), the cavity length is actively stabilised to the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition via spectroscopic techniques. A detailed discussion will follow in chapter 4.

Conclusion

The nuclear and electronic properties of strontium have been introduced, in particular the narrow intercombination lines that arise as a consequence of the spin–robit interaction. The significance of strontium in current research was touched upon, notably the $5s5p\ ^3P_0$ clock state in ^{87}Sr, which has become a strong contender as the re–definition of the SI second [34]. The main topic discussed was narrow–line cooling on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition, whereby temperatures as low as the photon recoil temperature have been achieved. The significance of the narrow linewidth transition on the cooling dynamics were discussed as well as the cooling and trapping procedure, which is a standard protocol in all cold strontium experiments. The stringent requirements of the linewidth of $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ MOT cooling light were introduced, including the different noise sources that contribute to the laser linewidth over numerous timescales.

Chapter 3

Short–term laser frequency stabilisation

A popular method to reduce the instantaneous and short–term laser linewidth is to use the phase locking technique developed by Drever and Hall, based on earlier work by Pound, which has been successfully demonstrated at various wavelengths [49, 50]. The technique requires a high–finesse optical cavity to provide a very narrow reference signal. The laser can be pinned to this relatively stable reference by means of high–bandwidth feedback to the laser diode current to reduce the instantaneous laser linewidth, and not–so–high–bandwidth feedback to the external cavity diode laser (ECDL) piezo to account for cavity length drifts, i.e. to reduce the short–term laser linewidth. When the laser is in this configuration, i.e. stabilised to the cavity, the notion of instantaneous and short–term laser linewidth is combined, and we say the laser has been narrowed.
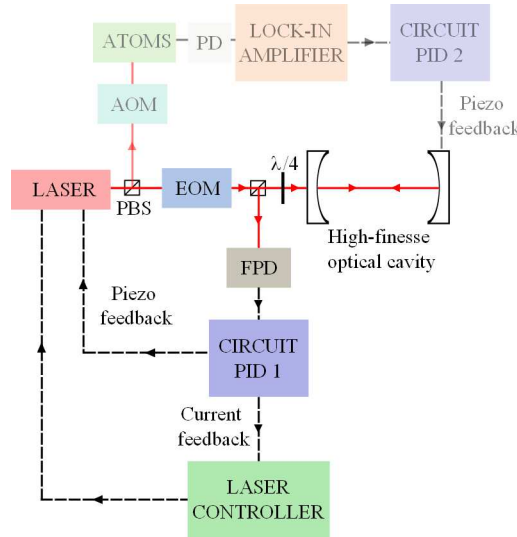
Figure 3.1: Simplified set–up of the 689 nm laser system. Shown are the optical and electrical components required for narrowing the laser linewidth, the laser beam path (red solid line) and the electrical feedback paths (black dashed line). The frequency of the light entering the optical cavity is modulated using an electro–optic modulator (EOM) and the reflected signal from the optical cavity is separated from the incoming light via a polarising beam splitter cube (PBS) and a quater waveplate, and is detected on a fast photodiode (FPD). The output of the FPD is sent to 'CIRCUIT PID 1'; this provides a correction signal to the laser piezo and laser injection current, via the laser controller. Also shown, but illustrated translucently, are the components required for stabilising the laser frequency on resonance with the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition. More is discussed about this set–up in chapter 4.

Figure 3.1 is a schematic of what is required to narrow a laser. The laser beam entering the high–finesse optical cavity is frequency modulated via an electro–optic modulator (EOM). The ingoing and outgoing beams to and from the cavity are distinguishable via a quarter waveplate and a polarising beam splitter cube (PBS). The reflected signal from the cavity is detected by a fast photodiode (FPD) and the signal is demodulated by a mixer that is referenced to the EOM driving frequency. Further discussion follows in section 3.3. The output of the mixer provides the error signal used for laser frequency stabilisation via feedback electronics ('CIRCUIT PID 1'), which provides corrections to the laser diode injection current, via the laser controller, and to the external cavity diode laser (ECDL) piezo. The components shown in figure 3.1 are discussed in great depth in the remainder of the chapter, in particular the cavity design constraints and circuit considerations when dealing with high–bandwidth feedback. In addition

the performance of each component in the feedback loop is discussed.

Also shown in the figure are the components to reduce the long–term laser linewidth (photodiode signal from the atoms, lock–in amplifier and a second PID circuit), however this method of stabilisation is discussed further in chapter 4.

3.1 689 nm laser system

The 689 nm laser system is a homebuilt ECDL based upon a modified version of the Arnold *et al.* [51] design that employs a Littrow ECDL configuration. Red laser diodes have been commercially available for many years and come at a reasonable cost. The laser diode is an Opnext HL6738MG with centre wavelength of 688 nm at laboratory temperature ($\approx$ 19 °C) and a maximal output power of 35 mW. The diode temperature is maintained at 25 °C to encourage the diode to lase at the desired frequency and the external cavity, from output facet to diffraction grating, is $\approx$ 15 mm long. This design has some drawbacks. For example, when tuning the laser frequency by moving the grating, the angle of the output beam changes and realignment further downstream the optical path is required. Alternatively, using a Littman–Metcalfe ECDL design would alleviate this irritation, hence this is a consideration for future designs, although this reduces the output power available from the laser.

The ECDL current is driven by a Vescent Photonics D2–105 Laser Controller, which is based on the Libbrecht–Hall circuit [52]. The controller has low current noise (100 pA/$\sqrt{\text{Hz}}$ [53]) and also has a front panel current servo BNC for high speed control of the laser's frequency. For lower bandwidth control, a piezo (Thorlabs AE0203D04F) is seated between the front and back mirror mount plate where the grating is mounted on. The addition of the piezo results in a mode–hop free tuning range of $\approx$ 3 GHz, without the aid of feed forward.

The laser diode temperature is stabilised using the D2–105 Laser Controller, which claims a long–term stability of 1 mK/day [53]. The controller provides two–stages of temperature control: one for the diode; the other for the diode housing. The diode housing is not stabilised, though it is something we might consider implementing as a possible improvement of the short–term laser instability. The temperature PID servo is manually tuned using the Ziegler–Nichols tuning method [54]. The laser frequency can drift by (0.4–4) GHz/hour, which

we believe is due to poor temperature insulation and a less–than–ideal thermistor position with respect to the laser diode.

3.2 High finesse optical cavity

An optical cavity consists of two highly reflective mirrors separated by a known distance. Depending on the choice of mirrors and cavity length, the amount and spectral bandwidth of the light transmitted varies. When designing an optical cavity, we must consider the performance characteristics of the cavity.

3.2.1 Cavity requirements for our system

We require the cavity to:

- **operate at 689 nm and 638 nm**: As well as narrowing the linewidth of the 689 nm laser, which is required for the second stage MOT, we will use the cavity as a transfer cavity to frequency stabilise the 638 nm laser system. We frequency–double 638 nm light to produce laser light at 319 nm. We will use 319 nm light to excite Rydberg states via the narrow $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition.

- **contain two piezos for cavity length stabilisation**: We will actively stabilise the length of the cavity using spectroscopy on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition and feedback to the piezos. We will use one piezo to provide an offset voltage to tune the laser onto resonance, and we will use the other piezo to stabilise the length of the cavity. Cavity length stabilisation is discussed in chapter 4.

- **have a finesse $> 10{,}000$**: Assuming we can lock to within 2–5 % of the top of the cavity peak (note that this is a rough estimate), and we have a cavity finesse of 10,000 and a cavity length of 100 mm, which is typically what other Strontium groups use for this laser system, we can obtain cavity resonances with a full–width–at–half–maximum of 150 kHz and obtain a locked laser linewidth of less than 7.5 kHz. This is narrower than the linewidth of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition so we should be able to cool on this transition using a cavity with those specifications. A finesse of much greater than this (i.e. of the order of hundreds of thousands) will increase design complications, such as requiring very good mechanical

isolation to prevent cavity vibrations, and we will not gain much with regards to linewidth. A discussion regarding basic cavity theory, including the definition of cavity finesse, follows in section 3.2.2.

- **have resolvable sidebands at 10 MHz**: A resonant EOM at a modulation frequency of 10 MHz is implemented in the Pound–Drever–Hall locking scheme. To resolve the 10 MHz sidebands, it is important for no higher–order TEM_{mn} modes to lie within, at least, 10 MHz of the carrier fringe. A discussion on higher–order modes follows in section 3.2.2.

In designing a high–finesse optical cavity, we have control over the mirror reflectivity, cavity length, and mirror radii of curvature. Each parameter contributes to at least one of our requirements, so we can carefully chose these to fulfill all of our requirements.

3.2.2   Designing an optical cavity

To design a cavity to fulfill ones needs, it is important to understand cavity theory.

Basic cavity theory

The frequency separation between neighbouring longitudinal modes of the same transverse mode is known as the free spectral range ($\Delta\nu_{\text{FSR}}$), which is dependent upon the cavity length L. In general, the free spectral range of an optical cavity is [55, 56]

$$\Delta\nu_{\text{FSR}} = \frac{c}{2nL} \ ,\tag{3.1}$$

where c is the speed of light and n is the refractive index of the gas between the mirrors. The only exception to this rule is the confocal cavity, where in this configuration, all of the modes are degenerate causing the free spectral range to halve.

The cavity finesse, $\mathscr{F}$, is a property dependent solely on mirror quality (i.e. mirror reflectivity $\mathscr{R}$, and to a lesser degree mirror substrate properties) and is independent of the cavity length, i.e.

$$\mathscr{F} = \frac{\pi}{2\arcsin\left(\frac{1-\mathscr{R}}{2\sqrt{\mathscr{R}}}\right)} \ .\tag{3.2}$$

As $\mathscr{R} \to 1$, equation 3.2 can be approximated as,

$$\mathscr{F} = \frac{\pi\sqrt{\mathscr{R}}}{1 - \mathscr{R}} \ .$$ (3.3)

Another important aspect is the cavity linewidth $\delta\nu$, which is related to the cavity finesse and free spectral range via,

$$\delta\nu = \frac{\Delta\nu_{\text{FSR}}}{\mathscr{F}} \ .$$ (3.4)

The narrower the cavity fringe the narrower the laser linewidth (when stabilised to the cavity), up to a limit. Illustrated in figure 3.2 is the effect of the mirror reflectivity on the cavity linewidth, and hence finesse, of an optical cavity with a known fixed length. The greater the reflectively, the narrower the cavity resonances.
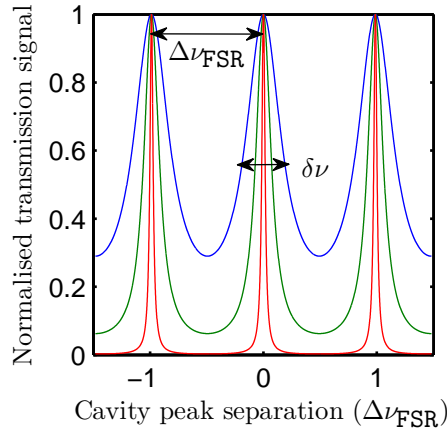
Figure 3.2: Theoretical cavity transmission for mirror reflectivities $\mathscr{R} = 30$ % (blue, solid line), $\mathscr{R} = 60$ % (green, solid line), and $\mathscr{R} = 90$ % (red, solid line). Also shown and labelled is the free spectral range $\Delta\nu_{\text{FSR}}$ and the cavity linewidth $\delta\nu$ for a mirror reflectivity of 30 %.

A cavity with a higher finesse will be more difficult to couple light into than one with a lower finesse. We have chosen two identical plane–concave fused–silica mirrors purchased from Layertec that are (25.0 ± 0.1) mm in diameter with a centre thickness of (6.35 ± 0.1) mm. The reflectivity of the mirrors is > 99.99 %, with losses from absorption or scattering ≈ 10 ppm, giving a theoretical finesse of $\approx 78,500$ (using equation 3.3) at a wavelength of 689 nm and 638 nm.

Using the theoretical finesse calculated from the mirror reflectivities and equation 3.2, the cavity length can be chosen such that the cavity linewidth is minimal (recall equations 3.1 and 3.4). In theory, very long cavity lengths (several metres or so) sound ideal, as the cavity fringes will be very narrow and the cavity itself will have a reduced frequency shift due to gravitational sag and mechanical vibrations. However, in practice, long cavities are not feasible in laboratory conditions due to limited bench space available. Having acknowledged the works of other groups, we decide to use a cavity length of 104 mm, once the piezo thicknesses have been taken into account. A cavity of this length has a $\Delta\nu_{\text{FSR}} \approx 1.5$ GHz, and cavity resonance widths of $\delta\nu \approx 35$ kHz.

Cavity stability criteria

Not all cavity configurations are stable, and it is important to acknowledge the beam propagation through the optical system. Propagation of a laser beam inside an optical cavity can be described by the paraxial Helmholtz equation [55] when cavity boundary conditions are taken into account. A solution to this equation is

$$A(\vec{r}) = \frac{A_1}{q(z)} \exp\left[-ik\frac{\rho^2}{2q(z)}\right] \quad , \tag{3.5}$$

where the wavenumber $k \equiv 2\pi/\lambda$ and A_1 is the amplitude of the wave. The radius ρ and complex radius of curvature $q(z)$ have been introduced such that $\rho^2 \equiv x^2 + y^2$, where x and y are the radial distances from the beam axis, and $q(z) \equiv z + iz_0$, where z is the axial position from the beam's waist, which is at z_0.

The complex nature of this solution gives rise to several features: at long distances from the focus the beam behaves like a spherical wave and close to the focus the beam diameter has a minimum ($\neq 0$). The intensity falls off in a Gaussian manner, therefore it is called a Gaussian beam. These features ensure that the beam can not carry infinite energy along the transverse direction and in the focus, which is not the case for the spherical wave.

We can define two real, and very useful, functions from the complex radius of curvature in the following way [57],

$$\frac{1}{q(z)} \equiv \frac{1}{r(z)} - i\frac{\lambda}{\pi w(z)^2} \quad , \tag{3.6}$$

and obtain the expressions for the radius of curvature of the beam inside the cavity $r(z)$ and the beam waist inside the cavity $w(z)$ by inserting $q(z) \equiv z + iz_0$:

$$r(z) = z \left[1 + \left(\frac{z_0}{z} \right)^2 \right] \tag{3.7}$$

and

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_0} \right)^2} \quad , \tag{3.8}$$

with $w_0^2 \equiv \frac{\lambda z_0}{\pi}$ that is defined in such a way that it corresponds to $w_0 \equiv w(z = 0)$, i.e. the value of the waist at the focus, and z_0 is the Rayleigh range. Furthermore, when a Gaussian beam passes through a focus, an additional phase shift of π is introduced, in addition to the usual $\exp[-ikz]$ phase shift that would be expected from a plane wave. The phase shift is known as the Gouy phase $\varsigma(z)$ [55, 56].

ABCD matrices are used to analyse the beam propagation through an optical system. For a beam reflected from a curved mirror of focal length f_1 onto a second curved mirror of focal length f_2, separated a distance L, with $f_{i=1,2} = R_{i=1,2}/2$ where R is the radius of curvature of the mirror in question, the corresponding ABCD matrix is [57]

$$\begin{pmatrix} h_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ -1/f_2 & 1 \end{pmatrix} \begin{pmatrix} 1 & L \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f_1 & 1 \end{pmatrix} \begin{pmatrix} h_1 \\ \theta_1 \end{pmatrix}$$

$$\begin{pmatrix} h_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 1 - L/f_1 & L \\ \frac{1}{f_1 f_2}(L - f_1 - f_2) & 1 - L/f_2 \end{pmatrix} \begin{pmatrix} h_1 \\ \theta_1 \end{pmatrix} \quad ,$$

where $h_{i=1,2}$ is the initial/final beam height and $\theta_{i=1,2}$ is the angle of the beam to the cavity axis.

If $h_2 > h_1$ and $\theta_2 > \theta_1$, the beam is on a diverging path that will lead to instability after many passes. Conversely, if $h_2 < h_1$ and $\theta_2 < \theta_1$, the beam will always tend to stability, since it will be converging towards the axis.

From this, the stability criterion can be inferred for a cavity of length L, having mirrors with radii of curvature R_1 and R_2 [57]

$$0 < (1 - L/R_1)(1 - L/R_2) < 1 \ , \tag{3.10}$$

which is shown in figure 3.3.

Figure 3.3: Cavity stability criterion for two mirrors with radii of curvature R_1 and R_2. For a stable cavity, the product of g_1 and g_2 must be less than one and lie between the blue, solid lines and the black, dash–dotted lines. Labelled are the g_1 and g_2 products for three cavity types: Concentric; Confocal; and Planar.

The cavity length is chosen subject to the free spectral range and cavity linewidth constraints. The remaining free parameter that influences the cavity stability is the radius of curvature of the mirrors. The tolerance for this is quite high permitting values of R within the range 100 mm $\leq R \leq$ 3000 mm. However, there is an additional requirement one must adhere to, which is stated at the beginning of this section: the 10 MHz sidebands must be resolvable. The value of R has a strong influence on this requirement due to higher–order TEM$_{mn}$ modes being present within the cavity, which is discussed in the following section.

Higher–order modes

In addition to the Gaussian beam, other solutions to the paraxial Helmholtz equation exist. Solving the equation in Cartesian coordinates leads to a family of solutions known as the Hermite–Gaussian modes, while solving the equation in cylindrical coordinates leads to the Laguerre–Gaussian modes. These solu-

tions are often abbreviated TEM$_{mn}$ modes. For both families, the lowest–order solution describes a Gaussian beam, i.e. TEM$_{00}$, while higher–order solutions describe higher–order transverse modes inside a cavity having spherical mirrors. In the remainder of the discussion, we shall only consider Hermite–Gauss modes. The lowest order modes in this family are shown in figure 3.4. The TEM$_{mn}$ mode of interest is the TEM$_{00}$ modes, which is displayed in the first panel in figure 3.4.

Figure 3.4: Theoretical calculations of the intensity distributions of the TEM$_{mn}$ Hermite–Gauss modes within an optical cavity. Starting from left to right, TEM$_{00}$, TEM$_{10}$, TEM$_{20}$, and TEM$_{30}$.

The higher–order modes have an additional phase factor when compared to the TEM$_{00}$ mode, which is related to the Gouy phase via [57]

$$\Phi(z) \equiv (m + n + 1)\varsigma(z) - k \left(\frac{\rho^2}{2r(z)} + z \right) \quad , \tag{3.11}$$

that lead to certain resonance frequencies inside the cavity. Using equation 3.11, as well as equations 3.7 and 3.10, the frequency separation Δf_{mode}, and mode number $(m + n)$, of the higher–order modes from the TEM$_{00}$ mode can be calculated via $\Delta f_{\text{mode}} = f_{\text{qmn}} - f_{\text{q00}}$, and

$$\Delta f_{\text{mode}} = \Delta \nu_{\text{FSR}} \left(\left(\frac{m + n + 1}{\pi} \right) \arccos \left[\sqrt{\left(\frac{1 - L}{R_1} \right) \left(\frac{1 - L}{R_2} \right)} \right] \right) \quad , \tag{3.12}$$

where q defines the longitudinal mode number, $R_{i\,=\,1,2}$ the radius of curvature of mirror 1 and mirror 2, L the length of the cavity and $\Delta \nu_{\text{FSR}}$ is the free spectral range of the cavity, which is defined in equation 3.1.

With the cavity length defined, specific mirror radii of curvature can be chosen such that Δf_{mode} is maximised using equation 3.12. Figure 3.5 (a) illustrates the mode separation when a length of 104 mm is chosen with mirrors having radii

Figure 3.5: (a) Theoretical calculations to determine the mode separation for differ-
ent radii of curvature R for a given longitudinal mode, which is defined by q. The
transverse mode number is defined by $(m + n) = $ x, and $(a + b) = $ x $+ 1$, where x
is an integer number. The cavity length L is fixed at 104 mm and we have assumed
the mirrors are identical. (b) Theoretical calculations to determine the mode number
$(m + n)$ of the TEM_{mn} that is nearest in frequency to the TEM_{00} mode within a
frequency range of ± 15 MHz for three radii of curvature: $R = 100$ mm (black circles);
$R = 400$ mm (blue diamonds); and $R = 1000$ mm (red squares).

of curvature R. The longitudinal mode number is defined by q, $(m + n) = $ x,
and $(a + b) = $ x $+ 1$, where x is an integer number. The stability condition
(equation 3.10) is valid for all values of R explored. We see that as $R \rightarrow 0$, the
mode separation increases then sharply drops off at $R = 100$ mm. The maxi-
mum mode separation we can achieve is ≈ 700 MHz, which is $\approx \Delta\nu_{FSR}/2$. This
makes sense as g_1 and g_2 are ≈ -0.06 making the cavity performance similar to
that of a confocal cavity, hence the higher–order modes are degenerate within
the cavity and causing $\Delta\nu_{FSR}$ to halve.

To reiterate the final cavity requirement, the 10 MHz sidebands must be re-
solvable. Therefore it is necessary to observe the mode number of TEM_{mn}
modes within a narrow frequency range centred at TEM_{00}. Figure 3.5 (b) illus-
trates the mode number of the TEM_{mn} within a frequency span of ± 15 MHz
for three values of R (100 mm, 400 mm, and 1000 mm). There appears to be
no apparent trend between the higher–order modes observed within the given
frequency range and the radii of curvature. Clearly, choosing $R = 400$ mm,
would not be beneficial since there are five higher–order modes present within
$\approx \pm 1$ MHz of the main carrier peak, which is centred at 0 MHz, and these modes

have comparatively low $(m + n)$ numbers. The lower the value of $(m + n)$ at a fixed separation from the main carrier resonance peak, the greater the influence it has on the system. This is a consequence of the mode beam radius matching more closely to that of the TEM_{00}, meaning it is easier for light to couple into the unwanted mode.

Alternatively, choosing $R = 100$ or $R = 1000$ mm, where the closest mode is ≈ 10 MHz $((m+n) = 56)$ and ≈ 14 MHz $((m+n) = 41)$ respectively, should make it possible to resolve the carrier resonance peak and the 10 MHz sidebands fairly easily. We choose to use a radius of curvature of 1000 mm, which results in a mode separation of ≈ 200 MHz, as these were readily and commercially available, in addition to fulfilling the final requirement.

Final cavity parameters

We have purchased commercially available mirrors that work at both 689 nm and 638 nm. We have chosen two identical plane–concave fused–silica mirrors that have a theoretical reflectivity of > 99.99 %, resulting in a theoretical finesse of $\approx 78,500$ at wavelengths of 689 nm and 638 nm. We have chosen a cavity length of 104 mm as this results in $\Delta\nu_{\text{FSR}} \approx 1.5$ GHz and, assuming $\mathscr{F} \approx 78,500$, $\delta\nu \approx 35$ kHz. Including piezos in the design permits active stabilisation of the cavity length (discussion to follow in chapter 4). Taking into account the thickness of each piezo, we have purchased a custom made, Zerodur$^{\circledR}$ spacer that is 100 mm in length. We have chosen Zerodur$^{\circledR}$ as it has a low thermal expansion coefficient [58]. The mirrors have a radius of curvature of 1000 mm, and, when separated by $L = 104$ mm, the cavity is stable and the higher–order modes are separated by 200 MHz from the TEM_{00} mode.

3.2.3 Cavity construction

An overview of the cavity and its enclosure is shown in figure 3.6. The custom made spacer from Hellma has a length of 100 mm and a square cross–sectional area of 30 mm. The spacer contains a longitudinal bore of 10 mm and an additional side bore, of 5 mm, to vent the central bore (see figure 3.6 (a)). The spacer is made of Zerodur$^{\circledR}$, since Zerodur$^{\circledR}$ has a low thermal expansion coefficient ($0.05 \times 10^{-6}/^{\circ}$C) [58]. In addition, when the temperature is constant and stable, but is warmer than room temperature, the thermal expansion coefficient

reduces.

(a)

(c)
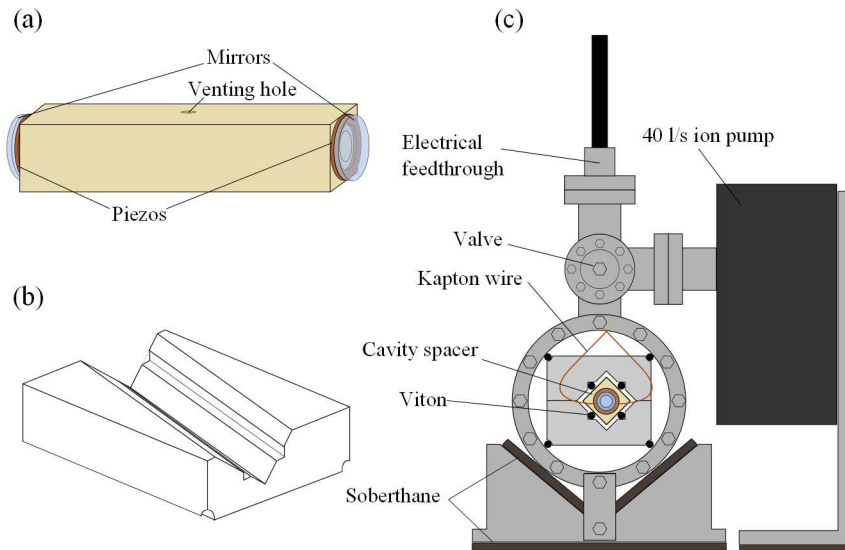
Figure 3.6: Schematic drawing of the high–finesse optical cavity. Note this is not drawn to scale. (a) Zerodur® cavity spacer with both piezos and mirrors glued at either end. Also shown is the side bore that is necessary for venting the central bore. (b) One half of the steel 'V' block mount that the cavity sits in. The indentation along the inside of the 'V' seats the Viton® rods. The axis of the cavity sits at an angle of 5° with respect to that of the steel block. (c) Front view of the cavity assembly. Labelled are the important aspects.

The piezos necessary for cavity length control are glued to either end of the cavity, between the spacer and the mirrors (see figure 3.6 (a)), using Norland NOA61 UV–curing epoxy. The use of two piezos is advantageous [59]: the large, static offset that is required to tune the cavity to resonance can be provided by one piezo, whilst the other piezo will be used to stabilise the cavity length to an atomic transition via spectroscopy techniques.

The piezo providing the offset (Noliac 2009223/GRC ring piezo) has an outer diameter of 30 mm, and an inner diameter of 10 mm. This has a large scan range of 3 μm and is able to withstand a maximum of 200 V, permitting a scan range in frequency of 12.5 GHz. To scan over a free spectral range of 1.5 GHz, a voltage of 24 V is required, which can easily be achieved using a series of batteries and a trimpot to provide tuning. It is ideal to use a battery powered reference voltage chip as the voltage source for the offset piezo, since this will have low inherent noise and a low offset drift. The stabilisation piezo (custom part Ring

RCM51) has an outer diameter of 25 mm, and an inner diameter of 15 mm. This piezo is more sensitive to small voltage changes than the aforementioned piezo. The stroke of the stabilisation piezo is ≈ 80 nm, and can also withstand a maximum of 200 V. Kapton® wire of thickness 150 μm is soldered onto the Ag–electrodes of each piezo using lead–free solder. Each wire is attached to an electrical feedthough for control.

Prior to gluing the mirrors onto the cavity spacer, a cavity is set up in free space using the purchased mirrors and the laser frequency is scanned using the ECDL piezo. The mirrors are mounted in a modified mirror mount, with each mirror mount supported on three translation stages (for x–, y–, and z–position control, with the laser beam propagating along the z–axis). Whilst maintaining a constant propagation direction of the laser beam, the mirror positions and angles are adjusted and the transmission is monitored and optimised using a photodiode. The spacer is placed between the mirrors and the mirror separation along the z–direction is reduced, such that the piezos on the spacer are in contact with the mirrors. The mirror angles and positions are adjusted to optimise the transmission once again.

Small quantities of Norland NOA61 UV–curing epoxy are dabbed on the piezo–mirror join. It is necessary to minimise, where possible, the quantity of glue used for two reasons: to limit the outgassing from the glue; and to prevent excessive glue smearing through the join onto the mirror surface. To cure, the glue is subjected to UV light for 20 s. After this initial curing time, the glue thickens and becomes more adhesive, but does not completely set. Fine mirror position adjustments are made to optimise the transmission, then the glue is cured for a further 30 minutes using the UV–lamp.

Once UV–curing is complete, the glue is 'aged' for the final curing stage. 'Aging' involves allowing the glue to set for a specific period of time before placing under vacuum. Neglecting this step risks smearing glue on the mirrors and/or mirror misalignment. 'Aging' requires 24 hours at room temperature, or fewer hours at a higher temperature. The spacer, piezos, and mirrors are heated to 50 ºC for several hours inside the vacuum can with the windows attached, but not vacuum sealed, to prevent contaminants defacing the mirror surfaces.

To mechanically isolate the cavity, the spacer is seated between two blocks of vacuum compatible stainless steel of 90 mm length, 69.7 mm width and 30 mm height, see figure 3.6 (c) for a schematic drawing of the block. The blocks are chosen to perform as a large thermal capacitor, as well as for vibrational dampening. The length is chosen to be slightly shorter than the total length of the spacer, piezos and mirrors, such that the mirrors hang out of the end facets of the block. The width, however, is maximised to fit into a DN100CF vacuum tube. The axis of the cavity sits in a 'V' shaped bore, which is at an angle of 5° with respect to that of the steel block. This is necessary to prevent reflections from the cavity coinciding with those from the vacuum window that may introduce unwanted parasitic etalons, see figure 3.6 (b).

The spacer is mounted inside the steel block on Viton® rods of ≈ 5.5 mm diameter and 20 mm in length. Viton® provides additional damping and, before use, is baked under vacuum up to a temperature of 200 °C for several days to free large amounts of water immersed in the material. The size of the material shrinks slightly and this needs to be taken into account when designing the inner bores of the 'V' block for the Viton® rods to sit in to ensure tight fitting.

The cavity is placed under vacuum such that any changes in air pressure and humidity will not lead to undesirable fluctuations of the effective cavity length and also convection currents will not hinder the temperature stability. The vacuum is maintained at a steady pressure of 2 x 10^{-9} Torr using a 40 l/s ion pump. Such a large ion pump is excessive for this use and a 2 l/s ion pump will suffice as it is only necessary to remove any products the Viton® might degas.

The DN100CF vacuum tube is seated on two 'V' blocks with 6.35 mm thick Sorbothane® sheeting sandwiched between. In an ideal scenario, the cavity vacuum tube is completely isolated from the table (i.e. no connection between the 'V' block and cavity vacuum tube). However, in this configuration and, due to the large vacuum pump mass, the cavity is susceptible to tilting if the optical bench is disrupted. Therefore, studding is used to connect one of the vacuum tube flanges to the 'V' block using a nut (see figure 3.6 (c)) and the entire vacuum system, including the ion pump, which is mounted and attached to an 'L' plate to hold the majority of the pump weight, is mounted on 6.35 mm thick Sorbothane® sheeting to isolate the cavity from mechanical vibrations from the

optical table. The height of the 'V' blocks are chosen such that the centre of the cavity mirrors are 85 mm above the optical bench, which is the typical working beam height. However, the inclusion of the Sorbothane® sheeting was an after thought, and, as a consequence, a periscope is required to couple light into the cavity. Periscopes can be unstable, however, to minimise this instability, a singular 25.4 mm thick post is used to mount both mirrors.

3.2.4 Cavity measurements

Cavity measurements are made in the following section to confirm whether the finesse, free spectral range, etc., are in agreement with those calculated and expected.

Cavity finesse

To measure the finesse of an optical cavity one would normally scan the laser frequency or the cavity length in order to obtain a transmitted spectrum like that shown in figure 3.2. One would then measure the frequency separation between two identical modes (i.e. between two TEM_{00} modes for example) and also the full–width–at–half–maximum (FWHM) of the cavity mode (i.e. the cavity linewidth $\delta\nu$). This is easy to do when measuring the finesse of a cavity with low reflectivity mirrors, but as the reflectivity, and hence the finesse, increase the cavity modes become much narrower than the linewidth of the laser light entering it, recall figure 3.2. The result is a transmission spectrum whose peak maximum fluctuates, making it difficult to extract a FWHM measurement. An alternative method to extract the finesse is to perform a cavity ring–down measurement, which is essentially measuring the average photon lifetime inside the cavity. In this measurement the amount of light that is 'trapped' within the cavity is constant when the cavity length is fixed. The lifetime of the 'trapped' light depends on the number of round trips the light undergoes, which is dependent on the reflectivity of the mirrors. To measure the lifetime of the 'trapped' light, the system must enter steady state, i.e. the amount of light entering the cavity must be equal to the amount of light escaping the cavity. This is achieved by stabilising the laser to the cavity using high–bandwidth feedback electronics, which are discussed in section 3.4. The circuit doesn't have to be fully–optimised, but must be sufficiently optimised to obtain detectable levels of transmitted light. Once in steady state, the light to the cavity is removed

using an acousto–optic modulator (AOM) and the decay time for the light to leak out of the cavity is measured.

The decay time, τ, is the time it takes for the intra–cavity power to decay to $1/e$ of its original value, P_0, i.e.

$$P(t) = P_0 \exp\left(-t/\tau\right) \ . \tag{3.13}$$

The number of round trips the light makes before having a $(1 - 1/e)$ probability of escaping is known as the bounce number, $\mathscr{B}$, which is related to the round trip time t_r and τ via,

$$\mathscr{B} = \frac{2\tau}{t_r} \ , \tag{3.14}$$

where $t_r = \frac{2L}{c}$, L is the cavity length, and c is the speed of light.

From the conservation of energy law, an assumption is made such that the light is either reflected (r), transmitted (t), absorbed (a) or scattered (s) off of the mirrors within the cavity, such that

$$
\begin{aligned}
r &= 1 - (t + a + s) \ , \\
&= 1 - \frac{1}{\mathscr{B}} \ , \\
&= 1 - \frac{L}{\tau c} \ .
\end{aligned}
\tag{3.15}
$$

Substituting equation 3.15 into equation 3.3 permits the cavity finesse to be calculated via the following equation,

$$\mathscr{F} = \left[\left(\frac{2\pi\tau c}{L} - \pi\right) + \sqrt{\left(\frac{2\pi\tau c}{L} - \pi\right)^2 - \frac{8\pi^2}{4}}\right]/4 \ . \tag{3.16}$$

In this measurement the transmission through the cavity is monitored using a fast photodiode (Thorlabs DET10A with 6 kΩ termination) whilst the laser is stabilised to the TEM_{00} mode of the cavity. An example of the decay of light out of the cavity is shown in figure 3.7, as well as the photodiode response time. From the average decay time of (4.39 ± 0.05) μs, we calculate a cavity finesse of $(41{,}300\pm600)$ using equation 3.16. The finesse is less than we were expecting ($\approx 78{,}000$), which we believe is a consequence of the mirrors having a lower reflectivity than that quoted on the manufacturers datasheet, which is

Figure 3.7: Measuring the decay time of light 'trapped' (red solid line) inside the high–finesse optical cavity. Also shown is a theoretical exponential fit (black solid line) with a decay time of 4.2 μs and the photodiode response time when terminated with a 6 kΩ resistor (blue solid line).

a theoretical estimate. Despite this, the finesse is still greater than the lower bound value of 10,000, the reason for which is stated in the requirements list at the beginning of this section.

Observing higher–order modes and measuring free spectral range

The frequency separation of the cavity modes from the TEM_{00} mode is measured using a CCD camera that monitors the transmission of the cavity, and determines the higher–order mode number, and a WS7 HighFinesse Wavemeter to monitor the frequency of the light entering the cavity. The laser frequency entering the cavity is varied using the laser piezo in the ECDL, whilst the cavity length is held fixed. A few of the lowest–order modes observed are shown in figure 3.8. The frequency separation between the higher–order modes and TEM_{00} mode is measured to be $\approx$ 200 MHz and $\Delta\nu_{FSR} = (1.43\pm0.04)$ GHz, which is expected from the theoretical calculations. Using equation 3.4, the cavity resonance linewidth is $\delta\nu = (35\pm1)$ kHz, which is slightly greater than expected owing to the reduced measured finesse.

Figure 3.8: TEM_{mn} modes that exist within the high–finesse cavity. Starting from left to right, TEM_{00} (i.e. Gaussian mode), TEM_{01} or TEM_{10}, and TEM_{03} or TEM_{30}.

Temperature stabilisation

Having tested the cavity piezos, only the small stroke piezo is connected to the feedthrough. During the cavity build, the Kapton® wire connected to the large stoke piezo disconnected from the electrical feedthrough, leaving only the small stroke piezo to provide a tuning offset voltage and a feedback voltage to stabilise the cavity length to an atomic resonance. For fear of contaminating the mirrors, we chose to leave the cavity under vacuum and use only one piezo.

There is another way to tune the TEM_{00} cavity mode onto resonance with the atoms and this is by changing the temperature of the cavity, since this changes the length of the cavity. Using a heater wire and a homebuilt PI feedback circuit, we stabilise the cavity temperature to $\approx \pm 250$ mK. This an upper bound limit since we are limited by the measurement technique. A detailed discussion follows in appendix B.

3.3   Pound–Drever–Hall frequency stabilisation

The idea behind the Pound–Drever–Hall (PDH) method is simple in principle: A laser's frequency is measured with an optical cavity, and using high–bandwidth electronics, this measurement is fed back to the laser to suppress frequency fluctuations and stabilise the laser to the top of a cavity fringe [49, 50]. The measurement is made using a form of nulled lock–in detection, which decouples the frequency measurement from the laser's intensity. An additional benefit of this method is that the system is not limited by the response time of the optical cavity. Frequency fluctuations that occur faster than the cavity ring–down time, can be measured and suppressed. The technique has been explained very well in other references, e.g. [49], [50] and [60], hence a detailed discussion will not

be carried out here, but rather a simplified one.

The PDH technique relies upon phase modulating the light that enters the optical cavity, such that the light consists of a carrier and two sidebands. This is achieved via an electro–optic modulator (EOM) driven at $\Omega_{\text{mod}} = 10$ MHz using an arbitrary function generator (arb). The internal oscillator clock of the arb is synchronised to the clock within an RF synthesizer, the output of which is connect to the 'LO' (i.e. local oscillator) input of the mixer. The phase between these signals is varied using the 'Phase' control on the arb front panel. Alternatively, it is possible to achieve a similar effect by modulating the laser diode current, though this is less desirable as the modulation will be present on all of the laser light, not just the light entering the cavity.

The optical set–up used for PDH frequency stabilisation is shown in figure 3.9. An additional polarising beam splitter (PBS) cube prior to the EOM is used, which is not shown in figure 3.9, to purify the polarisation of light entering the EOM and cavity. The lens is used for mode–matching the light entering the cavity for efficient coupling, see [57] for more information on mode–matching. Furthermore, an optical isolator is placed between the cavity and EOM to prevent back reflections feeding back to the master diode and the optics are mounted at an angle to the beam to prevent parasitic etalons.
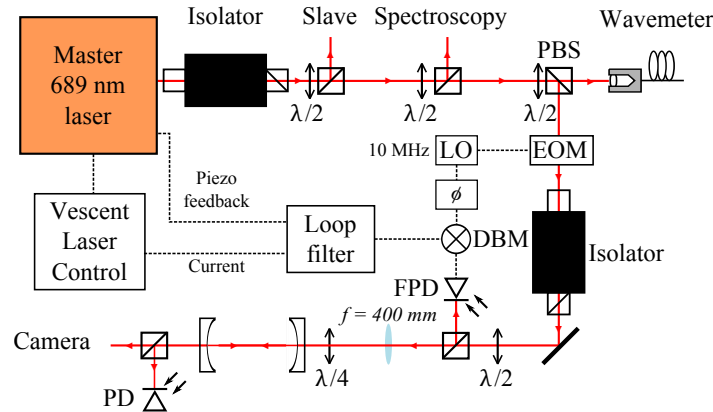
Figure 3.9: Optical set–up for Pound–Drever–Hall (PDH) laser frequency stabilisation. Shown are three optical paths: the first to the slave, further discussion is in section 5.2.1; the second for spectroscopy, which is discussed further in chapter 4; and the third for PDH locking. The beam in the PDH path is modulated using an electro–optic modulator (EOM) at 10 MHz and is coupled into the high–finesse cavity via an optical isolator, to prevent back reflections from the cavity interfering with the master laser, a polarising beam splitter (PBS) cube, mode–matching lens and quarter wave plate. The PBS cube and waveplate are required to extract the reflected signal. The reflected signal is detected using a homebuilt fast photodiode (FPD), the bandwidth of which is greater than 10 MHz. The local oscillator 'LO' 10 MHz signal is split into two: one drives the EOM; the other passes through a phase shifter 'ϕ', which is used for demodulating the FPD signal via a double–balanced mixer (DBM).

The beam that is sent to the optical cavity passes through a polarising beam splitter (PBS) cube and a quarter wave plate in order for the reflected light to be picked off from the incident beam. It is this signal that is detected and used for frequency stabilisation in the PDH scheme. It is necessary for the fast photodiode (FPD) to have a bandwidth that is greater than the modulation frequency, and should be incredibly low–noise as we do not wish for this noise to be put on the laser output. The signal from the homebuilt FPD is demodulated using a +7 dBm double–balanced mixer (DBM) and is processed by a homebuilt high–bandwidth PID circuit, referred to as the 'Loop Filter' in figure 3.9.

The transmitted light through the cavity is also monitored using a photodiode and a CCD camera. Observing the light on a photodiode aids cavity alignment and to check whether the laser is stabilised to the top of the cavity fringe. Whereas the CCD camera allows the user to optimise initial coupling for the

TEM$_{00}$ and to observe the TEM$_{mn}$ mode when the PDH frequency stabilisation scheme is engaged (i.e. locked).

3.3.1 Detecting the reflected signal

It is imperative for the fast photodiode (FPD) in the PDH frequency stabilisation scheme to be: (1) low–noise; and (2) operate at a speed that is high enough to detect the modulated signal. The photodiode circuit noise should not limit the linewidth of the laser, nor should it have a significant influence on the transfer function of the loop filter, which is discussed in section 3.4. Furthermore, the photodiode circuit should have sufficient gain at Ω_{mod} (i.e. 10 MHz in this set–up). A schematic of the circuit diagram for the fast photodiode is shown in figure 3.10. To reduce circuit noise, the power lines are capacitively decoupled and batteries are used for the supply voltages, which should alleviate 50 Hz mains noise on the signal. The photodiode is a Hamamatsu S5973. This has a maximum dark current of 100 pA, a typical photosensitivity of 0.51 A/W and a typical cut–off frequency of 1000 MHz [61]. A low–noise, very high speed operational amplifier (LT1226) is used, whose typical application is a photodiode preamplifier. In fact, the circuit diagram shown in figure 3.10 is the one recommended on the operational amplifier datasheet [62]. It is quoted as having an input voltage noise of 2.6 nV/$\sqrt{\mathrm{Hz}}$, a slew rate of 400 V/μs, and has a 1 GHz gain bandwidth.

Figure 3.10: Fast photodiode circuit diagram taken from [62]. The photodiode used is a Hamamatsu S5973 [61], and the operational amplifier is a low–noise, very high speed LT1226 [62].

The quoted bandwidth of the circuit in figure 3.10 is 15 MHz [62]. Although we have not measured the circuit bandwidth, we are confident it is > 15 MHz as we can easily resolve the 10 MHz modulation sidebands on the carrier peak. The demodulated and filtered PDH signal is shown in figure 3.11. Also shown is the transmitted light with sidebands centred at ±10 MHz from the main carrier peak. Both traces, the PDH and the transmission signal, are noisy owing to the fact the laser linewidth is greater than the cavity linewidth.

Figure 3.11: Transmitted cavity signal (red) with sidebands at ±10 MHz. This signal has been multiplied by two to make the sidebands visible. Also shown is the PDH error signal (black) that is used as a reference for stabilising the laser frequency. The PDH error signal is extracted from the monitor BNC output on the homebuilt high–bandwidth feedback electronics circuit, which is discussed in section 3.4. This is a filtered version of the error signal as it extracted after the first operational amplifier stage. The resistor and component values in the first operational amplifier stage are stated in [60].

3.4   High–bandwidth feedback loop

Feedback systems supply corrections to a device (e.g. a laser) with some electronic bandwidth. Such systems are typically referred to as feedback loops. A thorough explanation has been provided in reference [60], hence the remainder of this section will summarise the important points discussed in the reference.

3.4.1 Feedback loop criteria

When designing loop filters it is important for the system to have sufficient gain at low frequencies whilst also remaining stable, such that any small deviations of the error signal are greatly amplified and the laser is forced to move in the appropriate direction to suppress the deviation (i.e. providing negative feedback). Typically, the perturbations are reduced by the factor of $[1+A(f)]^{-1}$, where $A(f)$ refers to the system gain at Fourier frequencies f. To address the stability of the loop filter: the signal undergoes a phase shift as the feedback signal traverses the loop, which is not equal for all frequencies. Phase shifts increase with increasing frequency and, eventually, at very high frequencies the phase shift is 180º from the original correction signal (i.e. providing positive feedback). Positive feedback causes the laser to move more than the original perturbation, hence the loop becomes unstable and the laser oscillates. For the loop to remain stable, the loop gain must be less than 0 dB, or unity gain, in the frequency range where the loop phase shift is approaching 180º. To visualise these requirements, a Bode plot is often used to illustrate the frequency dependence of the system gain.

3.4.2 Bode plots

The Bode plot in figure 3.12 is an illustration of the gain as a function of frequency the laser experiences by means of feedback to the diode injection current, as well as the frequency response of other components involved in the feedback loop, such as the cavity, laser diode and high–bandwidth feedback electronics. This Bode plot is a theoretical calculation of the 689 nm laser–cavity system described in this chapter.

The bandwidth of the loop refers to the frequency at which the gain falls to 1 (or 0 dB), and this is commonly referred to as the unity gain point. The unity gain point is greater than the expected free running laser linewidth, and typically, it is in the range of several tens of MHz. In figure 3.12, the unity gain point is $\approx$ 20 MHz. There is a slope transition from –20 dB/decade to –40 dB/decade at $\approx$ 200 kHz, as it is required of the loop filter to have very high gain at low frequencies and to be able suppress higher frequencies as much as possible. The second slope transition from –40 dB/decade to –20 dB/decade at $\approx$ 2 MHz is required for loop filter stability. If the transfer function slope

is decreasing by more than 20 dB/decade of frequency at the unity gain point, the loop filter is rendered unstable [63].

Figure 3.12: Bode plot for stabilising a laser to a cavity using the laser diode's injection current. No feedback to the laser ECDL piezo is shown. The theoretical transfer functions of the high–bandwidth feedback electronics circuit (blue, solid line), the cavity (red, solid line), the laser diode (green, solid line), and the total response of the loop filter (black, solid line), which includes the circuit, cavity, laser diode and fast photodiode, is shown in the figure. The response of the fast photodiode is designed to have minimal effect on the total transfer function and so it is not shown in the figure. The black, dashed line indicates the unity gain point. The overall system response, in particular when the slope changes from a –20 dB/decade slope to a –40 dB/decade slope and vice versa, is shaped using capacitors, resistors and the switch S_1 in the high–bandwidth electronics circuit, which is shown in figure 3.13.

The elements that have the greatest effect on the loop filter are the cavity and diode laser, both of which are unique. It is for this reason why the high–bandwidth circuit needs to be optimised for each individual system. The fast photodiode, used for detecting the reflected signal from the cavity, is designed to have a minimal effect on the loop filter, which is why the response of this has not been included in figure 3.12. First to address the role of the cavity. The cavity acts like a low–pass filter having a f^{-1} slope for frequencies greater than the cavity linewidth $\delta\nu$. To obtain an overall transfer function like that in fig-

ure 3.12, whereby there is a slope change from –20 dB/decade to –40 dB/decade, the operation of the first operational amplifier in the high–bandwidth circuit is changed via switch S_1, which in shown in figure 3.13. With switch S_1 open, the first operational amplifier behaves like an integrator that has a gain governed by the open loop gain of the operational amplifier. The integrator contributes an additional –20 dB/decade slope to that already from the cavity. The frequency at which this slope change occurs is dependent, mainly, on the circuit components R_1, R_5, C_3 and C_4, which are shown in figure 3.13. Secondly, the transfer function of the laser diode is far more complicated than that of the cavity and it can vary among diode types and manufacturers. For most common low–cost diodes the response rolls off sharply at 1 MHz. For simplicity, and discussion purposes, it is shown in figure 3.12 to have a response like that of a low–pass filter with a cut off frequency of 1 MHz.

3.4.3 Circuit design

Oates *et al.* [60] discusses in detail the methodology of measuring the transfer function of the laser, the cavity and the detector. By measuring the transfer function of each individual component in the loop, it is possible to theoretically determine the circuit capacitor and resistor values necessary for producing the required loop filter transfer function. An alternative method is to optimise the circuit capacitor and resistor values in real time.

Figure 3.13: Simplified illustration of the high bandwidth loop filter circuit. Shown are the input and output ports, the important capacitors (C_n) and resistors (R_n), and finally the switches used for stabilisation. The switch configuration shows the laser locked via current and piezo feedback. The capacitors and resistors labels match with the values stated in appendix A.

A simplified version of the circuit to implement PDH frequency stabilisation is shown in figure 3.13. This is a modified circuit design to that discussed in [60]. Only the main contributors to the loop filter transfer function are shown, since the full circuit is much more complicated than this, which can be found in appendix A. The first stage operational amplifier filters and amplifies the error signal originating from the mixer. This signal is split into two feedback paths: a fast path that provides feedback to the laser's injection current, via another additional gain stage if more gain is required; and a slow path that provides feedback to the laser's piezo, via an integrator and offset stage, to control the length of the extended cavity.

The first stage consists of an operational amplifier with parallel RC feedback channels. When the integrator switch, S_1, is closed, the operational amplifier gain is $\frac{R_5}{R_1}$ at DC. However, when S_1 is open the gain is that of the open loop gain of the operational amplifier. Shown in figure 3.12 is the effect of the switch being closed and open. At a frequency $\frac{1}{2\pi R_5 C_4}$ the gain starts decreasing as f^{-1} – it is this stage that is responsible for one half of the –40 dB/decade slope shown in figure 3.12. The remaining half is attributable to the cavity since the cavity roll–off that occurs beyond the cavity linewidth contributes a f^{-1} slope. To shape the transfer function, such that there is a change in slope from a f^{-2} slope to f^{-1} slope prior to the unity gain frequency, the capacitor, C_3, in parallel with the input resistor, R_1, provides differential feedback that has a roll–off frequency of $\frac{1}{2\pi R_1 C_3}$. To centre the error signal around zero, a potentiometer, P_1, is used to provide an offset voltage before entering the first operational amplifier. Usually one would not need to zero the error signal when using an EOM as residual amplitude modulation is usually small. However, this is included as a precaution.

The resulting fast correction, after the second operational amplifier gain stage, is sent to the current modulation input of the Vescent laser current controller via a switch, S_3, and a resistor R_{26}. This resistance, as well as the inherent capacitance of the diode, forms a low pass filter, which might introduce a phase lag, limiting the servo bandwidth. It is for this reason why a capacitor, C_8, is placed in parallel with this resistor to act as a phase–lead compensator.

The transfer function in figure 3.12 does not include the low frequency gain

provided by the slow path. In this path, the gain for Fourier frequencies less than 1 kHz is significant and is dominant near DC. This is ideal as the laser current DC level can remain constant, preventing laser mode hops due to large changes in current. Including this stage means that under locked conditions, the DC level after the first operational amplifier should be equal to zero as the integrator in the slow path can apply a non zero correction. The frequency response of the slow path is given by $\frac{1}{2\pi P_6 C_5}$. To tune the laser onto a cavity resonance, an offset stage is included, which connects either the laser piezo to the loop filter or to a ramp signal via a ganged switch. This switch also controls whether the operational amplifier is integrating or not (S_{5-a} and S_{5-b} shown in figure 3.13). The offset stage can supply the piezo with a voltage between 0.5–11 V, permitting a laser scan range of approximately 1.3 GHz. An inverting amplifier is included before the integrator to change the polarity of the error signal. Furthermore, the inherent offset of the integrator operational amplifier is adjusted for more accurate locking.

3.4.4   Practical considerations

Oates *et al.* [60] discusses the importance of notch filters in the loop filter. Prior to the first operational amplifier in their circuit design, there are two notch filters to attenuate the unwanted signal at Ω_{mod} and $2\Omega_{\mathrm{mod}}$ that leaks through the mixer. In the discussion, notch filters are chosen as they introduce very little phase lag into the system and have sufficient attenuation. Having tested notch filters at 10 MHz and 20 MHz, they are not included in the final circuit design (shown in figure 3.13) owing to their susceptibility to stray capacitance pick–up. In fact, no filters are used between the mixer output and the circuit input.

Preliminary work trialled a strip board version of the circuit, which was very susceptible to pick–up and high frequency oscillations that could not be removed or suppressed. Instead, a printed circuit board design is used, which inherently reduces the signal path length (hence reduces phase lag) and warrants a grounding plane around the circuit (hence isolating the circuit from environmental factors). Also included is a monitor BNC output, where the output is extracted after the first operational amplifier, granting a monitor of the error signal. The current, or rather the fast feedback, output is connected to the

laser via the BNC modulation input on the Vescent photonics current controller.

One last consideration is the choice of operational amplifier. The operational amplifier must have a gain bandwidth product greater than 30 MHz such that there is sufficient gain at high frequencies, and have a noise level that is substantially less than that of the input signal, which typically depends on the amount of light incident on the photodiode. Additionally, the operational amplifiers must be unity gain stable. The operational amplifiers are surface mount OPA211a. They are quoted having a low voltage noise of 1.1 nV/$\sqrt{\text{Hz}}$ at 1 kHz and have a gain bandwidth product of 80 MHz (with a corresponding gain of 100). In the slow path, the operational amplifiers do not have to operate as fast but do have to be low–noise and have a low offset drift. The operational amplifiers are OPA227, which are quoted as having a low voltage noise of 3 nV/$\sqrt{\text{Hz}}$ (the frequency at which this was taken is not stated) and a low offset voltage of 75 μV maximum.

3.4.5   Initial locking

Prior to engaging the loop feedback, the following technical issues are addressed.

Employing an EOM to modulate the laser frequency entering the cavity, although advantageous to laser diode current modulation, can introduce technical problems, such as amplitude modulation of the light as a consequence of birefringence in the EOM crystal. As a result, a DC offset will be observed on the reflected signal. These offsets are removed by tuning the laser off–resonance with the cavity, monitoring the output of the FPD on a spectrum analyser centred at the modulation frequency and rotating the EOM axis until the observed peak at the modulation frequency is minimised. We do not, however, remove the peak completely.

It is necessary for the reflected light level to be above the noise floor of the FPD, which can be monitored via the spectrum analyser. For a suitable signal–to–noise ratio, a power of ≈ 13 μW on the FPD is used, with ≈ 24 μW entering the cavity resulting in a measured spectral power much greater than $+7$ dBm. On account of using a $+7$ dBm double–balanced mixer (Minicircuits ZLW–1–1+) to demodulate the FPD signal, the RF input requires an input of $+7$ dBm,

anything deviating from this and the mixer is under– or over–driven. Therefore a +10 dBm attenuator is placed between the FPD and the RF input of the mixer and the light levels are set accordingly. The 'LO' input is constrained to a maximum of 1 dBm, which is set using the amplitude of RF synthesizer output. When setting up the PDH frequency stabilisation scheme, one should look up their specific mixer datasheet and not rely on the values quoted here. The output of the mixer is attached directly to the input of the circuit using an SMA–SMA connector.

Early attempts of stabilising the laser frequency to the cavity use the capacitor and resistor values quoted in the Oates [60] paper. It is beneficial to optimise only one feedback path at a time, with the current (i.e. the fast path) feedback optimised first. While the laser ECDL piezo is scanning over a cavity resonance at a frequency of ≈ 30 Hz with the current feedback path engaged, the circuit attempts to stabilise the laser to the top of the cavity resonance, and in doing so, the resonance widens. This is the best time to optimise coupling into the cavity. If the phase is set incorrectly, the loop filter will try to lock to the sidebands instead, which can be observed via the transmission peaks and the monitor output of the loop filter. If this is the case, the phase of the signal from the local oscillator to the LO input of the mixer needs to be changed by 180°. If no locking is observed, then there is insufficient gain in the circuit (increasing the feedback resistor, R_{25}, in the second operational amplifier in the fast path will increase the gain, similarly reducing the resistance of R_1). On the other hand if there is too much gain, the servo will oscillate and the loop filter will suppress the transmission.

3.4.6 Optimising the loop filter

The optimum values of the resistors and capacitors in the circuit depend on the particular laser diode, and also on the cavity and detector. It is found that the capacitor and resistor values stated in figure 3.13 have the most significant effect on the transmission through the cavity, which is why these component values have been shown in the schematic circuit. Varying these components alone will not result in the optimium feedback circuit however, therefore, the remaining components shown in figure A.1 require changing also. The procedure for optimisation involves an iterative process of changing component values in

figure A.1 and monitoring the effect on the transmission. An increase in transmission infers an increase in circuit bandwidth. To determine whether more capacitance is required for the phase–lead capacitors, a simple test is to place a finger on top of the capacitor and observe whether the transmission increases or not.

Nearing ever closer to the optimum values, the circuit, via the current feedback, will stabilise the laser frequency tighter to the top of the cavity peak, resulting in an increase in transmission and reduction in transmission noise. When the loop filter is near optimum, the slow feedback path to the piezo is introduced. To optimise the piezo feedback path, the proportional gain for the piezo operational amplifier is reduced to a minimum. The integrator time constant is increased using P_6 until the locked DC error signal from the monitor BNC output is forced away from zero. The resistance of P_6 is set just below this value. Proportional gain is introduced and increased until oscillation is observed on the locked DC error signal.

For finer tuning of the loop parameters, it is best to monitor the in–loop locked error signal, which is possible using the coupled output of a directional coupler, whose input is attached to the FPD, and an RF spectrum analyser.

Fine tuning of loop parameters

In order to completely optimise the circuit, a better diagnostic is required. Using the coupled output of the directional coupler, whose input is attached to the output of the FPD, the in–loop locked error signal can be observed using an RF spectrum analyser. Figure 3.14 compares the system noise when the laser is 'free–running' and when it is locked to the cavity. For the 'free–running' trace, it is difficult to sit the laser on top of the cavity peak for long enough to extract a trace like that shown in figure 3.14. Instead the feedback loop is engaged and the proportional gain is reduced, using a variable resistor in place of R_{25} in the circuit, until the laser is on the cusp of unlocking. When the laser is 'free–running', we see a very broad noise spectrum centered around Ω_{mod}. This is not indicative of the laser linewidth however and is only used as a visual aid. When the laser is stabilised to the cavity, the noise observed on the 'free–running' laser is removed and deposited as far away as possible from

Ω_{mod}, leaving a very narrow central peak with sidebands at $\Omega_{\text{mod}} - f_{\text{band}}$ and $\Omega_{\text{mod}} + f_{\text{band}}$. At $\Omega_{\text{mod}} \pm f_{\text{band}}$ the signal has undergone a phase shift of 180º from the original correction signal. The loop is providing positive feedback, hence the noise is increased rather than suppressed. Increasing the loop gain forces these sidebands further away from the carrier peak, i.e. increasing the loop filter bandwidth, but if there is too much gain, the loop filter begins to oscillate (i.e. positive feedback causes the loop to become unstable). This becomes evident when the noise sidebands increase drastically in amplitude. A reduction in the transmission peak amplitude is also observed with an increase in noise on the error signal, which is observed via the monitor output BNC. The width of the noise spectrum, i.e. f_{band}, when the feedback loop begins to oscillate, is an indication of the bandwidth of the loop filter. In figure 3.14 the in–loop error signal has large servo bumps at 1.5 MHz away from the central peak at 0 MHz, which is centred at Ω_{mod}. Fine tuning of the feedback circuit components is achieved by, once more, changing individual circuit components such that f_{band} is as large as possible in the absence of oscillation. During optimisation, it is discovered that reducing the signal path length from the photodiode to the modulation input of the Vescent controller (i.e. reducing the total electrical path length from 3 m to 1 m) makes a significant improvement to the bandwidth of the loop filter. Before reducing the cable length, the best bandwidth obtained was 600 kHz.

An additional point to note. The bandwidth changes with the laser's operating parameters. For example, if the laser is near a mode hop, the bandwidth reduces and sometimes the laser will not lock. Similarly, if coupling into the cavity deviates from optimum, the amount of light incident on the FPD changes, which will inherently change the gain of the system. Since the circuit is optimised under very specific conditions, e.g. a given power incident on photodiode and for the TEM_{00}, the bandwidth of the loop filter reduces when stabilising to other TEM_{mn} modes.

Figure 3.14: In–loop locked error signal noise (red) observed on an RF spectrum analyser centred about the modulation frequency, which is 10 MHz, and the 'free–running' laser noise (black). The resolution bandwidth of the spectrum analyser during this measurement is 100 Hz. The centre peak at 0 MHz of the in–loop error signal is equivalent in power as the 'free–running' laser noise. The noise sidebands are peaking at a frequency of $\approx \pm 1.5$ MHz, indicating a servo bandwidth of this wide.
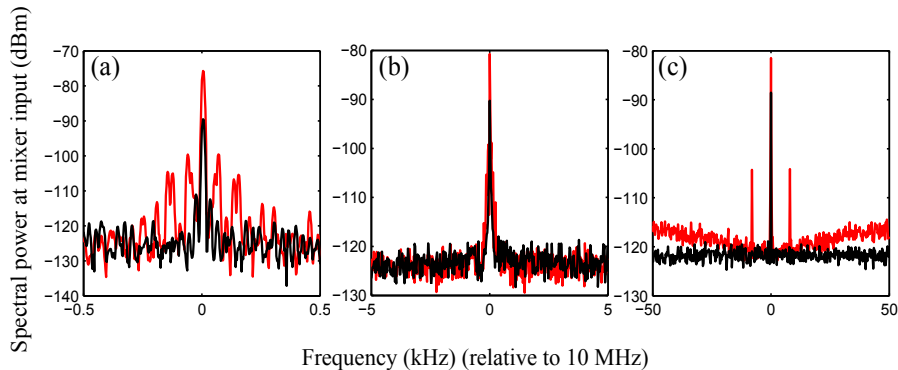
Figure 3.15: In–loop error signal noise (red) observed on a spectrum analyser centred about the modulation frequency, which is 10 MHz, and the laser AM noise (black), over three frequency spans: (a) a span of 1 kHz; (b) a span of 10 kHz; and (c) a span of 100 kHz. The resolution bandwidth of the spectrum analyser during this measurement is 10 Hz.

It is also advantageous to look at the laser amplitude modulation (AM) noise when fine tuning the feedback circuit, since this sets the minimum limit to which the low frequency noise can be reduced. The amplitude noise level of the laser

is measured by disengaging the loop filter, tuning the laser off–resonance with the cavity and measuring the power incident on the FPD using an RF spectrum analyser. It is ideal to reduce the frequency observation range considering this is where the gain of the transfer function is expected to be high. When optimising the circuit, the in–loop error signal noise should match that of AM noise over as large a frequency span as possible.

The laser AM noise and the locked in–loop error signal are shown in figure 3.15 for the optimised feedback circuit over three frequency spans: (a) 1 kHz; (b) 10 kHz; and (c) 100 kHz. The resolution bandwidth of the spectrum analyser is set to be minimum, which is 10 Hz. The laser AM noise and the locked in–loop error signal have the same baseline level, up until $\approx \pm 7$ kHz, as is evident in figure 3.15 (c). Only amplitude fluctuations are observed on the laser since laser frequency fluctuations are suppressed. There is no 50 Hz noise on the laser, which is apparent in figure 3.15 (a), but instead frequencies of ≈ 68 Hz and ≈ 153 Hz. The presence of 50 Hz noise is easily explained by mains noise. However, the origin of the noise at these frequencies is unknown. In view of the fact that the noise is not on the AM noise signal, the noise does not originate from the laser ECDL. A possible source might be the oscillation of the detached piezo, though it is difficult to test this unless the cavity vacuum is intentionally broken and the piezo is reattached. Alternatively, the noise might be a consequence of the Thorlabs piezo amplifier that is used to amplify the output of the low–bandwidth feedback circuit used in cavity length stabilisation. These amplifiers are renowned for having significant noise at 50–60 Hz. Further discussion regarding this follows in chapter 4.

In figure 3.15 (c) there are noise peaks at a frequency of $\approx \pm 8$ kHz. This frequency is too high to originate from the cavity piezos, which typically respond up to a maximum of 1 kHz. The noise might originate from the high–bandwidth electronics circuit, though, to date, there is no evidence to support this.

In all three subplots of figure 3.15, there is a carrier peak at Ω_{mod} for both traces, i.e. when the laser is locked to the cavity and when the laser is off resonance with the cavity. One would assume an absence of carrier peak when the light is off resonance. However, the presence suggests there is light incident on the FPD as a consequence of higher–order modes. If a small fraction of light

is coupled into these higher–order modes, the sidebands will undergo a slight phase shift, leading to a small frequency offset from the carrier peak, which will reduce the PDH error signal gradient. This does not seem to have a significant effect on the instantaneous or short–term linewidth, since we measure a temperature very close to that of the photon recoil limit, which is discussed in chapter 5.

It is not possible to extract a laser linewidth from the method used to optimise the feedback circuit. However, methods to evaluate the performance of the laser system are discussed in section 3.5.

3.5   Evaluating the loop filter

To obtain an absolute laser linewidth on short–interrogation times, a beat measurement in the frequency domain is required, then, by using spectral analysis (chapter 2), the dominant noise sources can be diagnosed and, therefore, reduced. Beat measurements can be made via a self–heterodyne method [64] or between two other identical laser–cavity systems. The self-heterodyne measurement requires a 10 km single–mode polarisation–maintaining fibre, which we do not have access to, whereas the latter involves building and optimising the entire 689 nm laser system twice more, which is unnecessarily time consuming. It is difficult to extract an absolute laser linewidth without significant effort. We have devised two alternative methods to determine the performance of the 689 nm laser system: measuring the loop filter performance; and using the cold atoms.

3.5.1   Feedback loop performance

Measuring the noise within the loop gives a good indication of the magnitude of the laser noise, but not of the source of the noise. For example, if the loop is trying to suppress noise at the integrator input of the high–bandwidth electronics circuit, and this noise is due to amplitude fluctuations, the loop filter will suppress this signal by increasing the frequency fluctuations [60].

The loop noise is measured by calibrating the in–loop locked RMS noise using the slope of the PDH error signal when the cavity length, or laser frequency, is in scanning mode. Since the modulation frequency ($\Omega_{\mathrm{mod}} = 10$ MHz) is much

greater than the linewidth of the cavity ($\delta\nu = 35$ kHz), the slope at the centre
of the error signal is:

$$D = \frac{2V_{\mathrm{pp}}}{\delta\nu} \quad , \tag{3.17}$$

where V_{pp} is the peak–to–peak amplitude of the PDH error signal extracted
before the high–bandwidth electronics circuit, commonly referred to as the AC
error signal, and $\delta\nu$ is the cavity linewidth.

Measuring V_{pp} via the method discussed is usually an underestimate of the
peak–to–peak amplitude, but will provide an upper bound for the system noise.
Assuming Gaussian noise on the laser, the noise (ν_n) is estimated using [49]

$$\nu_n = \frac{2.355 V_{\mathrm{rms}}}{D} \quad , \tag{3.18}$$

where V_{rms} is the root–mean–square of the locked error signal. The gradient of
the central PDH error signal feature, used throughout this thesis, is typically
$D = 47$ mV/kHz.

The noise of the AC error signal (V_{obrms}) is not only a consequence of the laser
noise, but also due noise on the electrical components, such as the FPD. By
blocking the light to the FPD, the electrical noise (V_{backrms}) can be measured
and removed via

$$V_{\mathrm{rms}}^2 = V_{\mathrm{obrms}}^2 - V_{\mathrm{backrms}}^2 \quad , \tag{3.19}$$

leaving the true RMS amplitude of the locked error signal. Using equation 3.18,
a 'linewidth', or rather the noise, can be extracted.

Figure 3.16 illustrates the noise for a given oscilloscope timebase setting. This is
very similar to an Allan deviation–type measurement, except that the averaging
time is set using the timebase, and hence bandwidth, of the oscilloscope. The
noise at very short times is (1.51 ± 0.06) kHz, which, we can assume, is the upper
bound to the laser linewidth. This seems a plausible result, since a tempera-
ture very close to that of the photon recoil limit is measured and discussed in
chapter 5. We will show in the following section, the temperature dependence
on the loop filter bandwidth, and hence the laser linewidth.

Figure 3.16: Laser noise from the RMS in–loop error signal for given oscilloscope timebase settings.

3.5.2 Laser linewidth performance on cold atoms

The temperature of the atoms in the MOT is governed solely on the cooling transition linewidth (recall equation 2.2). If the laser linewidth of the interrogating light is greater than that of the cooling transition, the atom temperature will depend on the laser linewidth rather than the transition linewidth. The performance of the single–frequency red MOT can be used as a diagnostic of the behaviour of the 689 nm laser system. By measuring the temperature of the single–frequency red MOT, the short–term frequency stability of the 689 nm laser can be deduced, however, an absolute linewidth can not be extracted. The temperature of the single–frequency red MOT is measured using the time–of–flight method. The cavity length has been stabilised to the $5s5p\,^1S_0 \rightarrow 5s5p\,^3P_1$ atomic line, which is discussed in chapter 4.

Figure 3.17: Atom temperature on laser linewidth. Shown in (a) and (c) are the in–loop locked error signal noise where S is the spectral power at the mixer input and the detuning is relative to 10 MHz, and in (b) and (d) the squared width versus the square of the expansion time is shown, from which the narrowband temperature can be deduced, for a total red MOT beam power of 200 μW. The 'poor' high–bandwidth laser stabilisation feedback loop to the high–finesse cavity is shown in (a) with a loop bandwidth $f_{\mathrm{band}} \approx 150$ kHz leads to a corresponding narrowband red MOT atom temperature of T = (14 ± 1) μK. The improved stabilisation feedback loop in (c) has a bandwidth of $f_{\mathrm{band}} \approx 1.5$ MHz leading to a narrowband red MOT atom temperature of T = (2.2 ± 0.1) μK

In this method the atoms are released from the trap for a free expansion time of Δt and a fluorescence image is taken using the blue MOT light. The power and detuning of the single–frequency red MOT is the same during each time–of–flight measurement. By measuring the $1/e^2$ radius of the MOT cloud (w), a temperature is inferred from the linear fit. Figure 3.17 compares the in–loop locked error signal noise with the corresponding temperature of the single–frequency

red MOT for high–bandwidth laser stabilisation feedback loops with different bandwidths. In figure 3.17 (a) the bandwidth of the loop filter is $\approx$ 150 kHz and the corresponding red MOT temperature is (14$\pm$1) μK (figure 3.17 (b)). By increasing the loop filter bandwidth by an order of magnitude (figure 3.17 (c)), the temperature is reduced by nearly an order of magnitude also (figure 3.17 (d)).

Conclusion

We have introduced the reader to the well–known Pound–Drever–Hall laser frequency stabilisation technique that is used to reduce the instantaneous and short–term laser linewidth of a homebuilt ECDL 689 nm laser. The equipment requirements of the technique were briefed, followed by a detailed discussion regarding the design, via a theoretical approach, and construction of a high–finesse optical reference cavity. The cavity properties, such as the finesse and free spectral range, were measured. The measured finesse is (41,300$\pm$600), which is not in agreement with the theoretical finesse of $\approx$ 78,500 and is likely due to the mirror reflectivity deviating from that stated on the manufacturers datasheet. In spite of this, the finesse is greater than the lower bound limit of 10,000, which was set during the design phase of the cavity.

A thorough discussion regarding high–bandwidth feedback electronics is given. In particular the importance of feedback loop response, as well as an introduction of how to shape the response via electronic circuitry. A basic step–by–step procedure on laser frequency stabilisation to a high–finesse optical cavity via high–bandwidth feedback electronics is covered, which includes a procedure for optimising the circuit components.

The absolute laser linewidth is not measured due to experimental limitations, however two approaches, which infer the performance of the laser system, were discussed. The first method involved measuring the noise of the feedback loop. When optimised, the upper bound measured noise results in a linewidth of (1.51$\pm$0.06) kHz on short–interrogation times of tens of μs, which is less than the 5s^2 ^{1}S$_0$ $\rightarrow$ 5s5p ^{3}P$_1$ transition linewidth. The second method involved measuring the single–frequency red MOT temperature for numerous loop bandwidths. By increasing the bandwidth by an order of magnitude, the temperature of the atoms reduces by an order of magnitude. Using these methods

and an iterative procedure of optimising the circuit components by measuring atom temperature, a feedback bandwidth of ≈ 1.5 MHz is obtained, resulting in a measured minimum single–frequency red MOT temperature of ≈ 460 nK (further discussion follows in chapter 5).

Chapter 4

Long–term laser frequency stabilisation

In chapter 3 we discuss how to narrow a relatively broad 1 MHz external cavity diode laser to several kHz using a high–finesse optical cavity and high–bandwidth feedback electronics. This, in itself, is a stabilisation technique, which can be used for both short–term frequency stabilisation (i.e. laser linewidth narrowing) and, should you so wish, long–term frequency stabilisation (usually over the period of hours or days). We have decided to actively stabilise the long–term frequency by providing low–bandwidth, electronic feedback to the cavity length, by means of a piezo seated between the cavity spacer and a cavity mirror, that is referenced to an atomic transition [59].

A schematic diagram of the long–term laser frequency stabilisation loop is shown in figure 4.1. The laser light, whose short–term frequency is stabilised to a high–finesse optical cavity, is frequency modulated via an acousto–optic modulator (AOM) and sent to a beam of strontium atoms. An error signal is generated via a lock–in amplifier for phase–sensitive detection (the atomic spectroscopic techniques available are discussed later). A low–bandwidth feedback circuit ('CIRCUIT PID 2') uses the error signal as a reference and provides corrections to the cavity length, which in turn changes the laser frequency.

Figure 4.1: Simplified stabilisation set–up of the 689 nm laser system. Shown are the optical and electrical components required for stabilising the high–finesse optical cavity length, laser beam path (red solid line) and electrical feedback paths (black dashed line). Also shown, but illustrated translucently, are the components required for narrowing the laser linewidth. The light required for this stabilisation scheme is extracted using a polarising beam splitter cube (PBS). More is discussed about this set–up in chapter 3.

There are many atomic spectroscopy methods we can employ for frequency stabilisation on the $5s^2$ $^1S_0$ $\rightarrow$ $5s5p$ $^3P_1$ atomic transition. These include polarisation–absorption spectroscopy, which is used by the LENS group [65, 66], and FM–absorption spectroscopy, which is used by the IQOQI group [67]. Both groups measure a minimum atomic spectroscopy linewidth of 50 kHz, and do not observe a considerable cavity length drift (or rather one that would change experimental parameters) over the course of the day. These techniques use either hollow cathode lamps or a heat pipe, where the absorption paths lengths are long.

We chose to employ saturated fluorescence spectroscopy, which was inspired by the SYRTE (Paris, France) group, since we have an atomic beam machine available. The absorption path length is comparatively short, leading to small absorptions of the incident laser beam, which is reduced by saturation and is also difficult to detect [68]. Although the fluorescence signal is reduced by saturation also, a sensitive modulation technique for isolating small changes in the fluorescence intensity have been developed [69] and implemented [70]. A discus-

sion of this technique will follow later in the chapter. A schematic diagram of the set–up is shown in figure 4.2. In this scheme a single, frequency modulated beam is directed through an atomic sample, and the fluorescence is detected.

Figure 4.2: Schematic diagram of the saturated fluorescence spectroscopy detection technique. A single, frequency modulated beam is retro–reflected orthogonally through an atomic beam. The fluorescence is collected via a 1:1 lens system, and detected via a low–noise photodiode. There are two photodiode outputs (one is the photodiode signal output, whilst the other is a reference ground) to prevent ground loops inducing a change in DC signal levels. The outputs of the photodiode are sent to a lock–in amplifier in differential configuration for phase sensitive detection.

The frequency modulated beam is retro–reflected back through the atomic beam to obtain a sub–Doppler feature. The SYRTE group measure an atomic spectroscopy linewidth of, at best, 200 kHz [71], which is broader than that obtained via polarisation– and FM–absorption spectroscopy. However, they are able to measure the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition frequency to within 20 kHz (i.e. 434 829 121 300 (20) kHz) [72].

4.1 Fluorescence imaging

If we replace the photodiode in figure 4.2 with a CCD camera we can measure both the velocity spread of the atoms in the atomic beam, and the lifetime of the 5s5p $^3P_1$ state. Figure 4.3 is a 'false' colour camera image taken when a single–frequency, single–pass, 689 nm laser beam is intersecting a hot strontium atomic beam orthogonally. This image is taken using a 'PixelFly qu' CCD camera with an exposure time of 65.5 ms. The laser beam is propagating along the y–axis, the atomic beam is propagating along the x–axis, and the CCD is along the z–axis. From this image, we observe the decay of atoms out of the 5s5p $^3P_1$ via a reduction in the fluorescence.
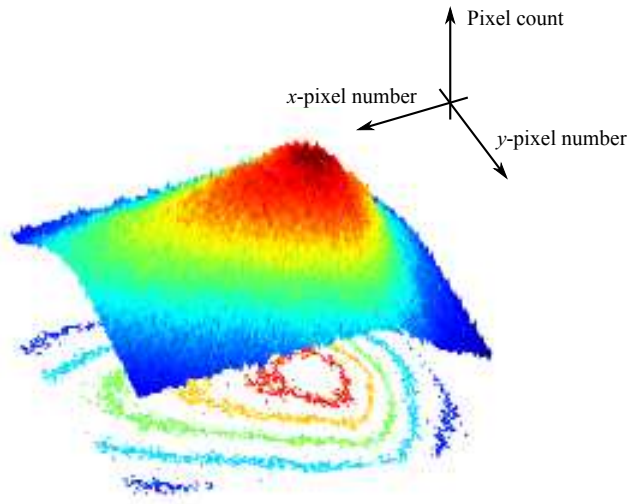
Figure 4.3: 689 nm fluorescence image on the $5s^2$ $^1S_0 \rightarrow 5s5p$ $^3P_1$ atomic line. The laser is propagating along the y–axis, the atomic beam is propagating along the x–axis, and the CCD is along the z–axis.

Taking slices out of the date shown in figure 4.3 along the direction of the laser (atomic) beam warrants the lifetime (velocity) information to be extracted. Inset figure 4.4 is a slice taken along the direction of the laser beam. Once the atoms have propagated through the laser beam, we can calculate the lifetime of the exited state, τ, via [40]

$$N(t) = N_0 \exp\left(-\frac{t}{\tau}\right) , \qquad (4.1)$$

where N_0 is the initial pixel count and $N(t)$ is the pixel count after a propagation time t.

The main body of figure 4.4 is the decay of atoms of the 5s5p ^{3}P$_1$ state. We assume all of the atoms in the atomic beam are propagating at the most probable speed of $\overline{v} = 500$ m/s with an uncertainty of $\pm$ 11 m/s due to the atomic beam divergence (43 mrad full width). The $1/e^2$ radius of the laser beam is ≈ 0.7 mm, both perpendicular to the optical bench and parallel to it. Using the gradient of the logarithmic least squares linear fit deduced for the data in the inset of figure 4.4 after a propagation distance of ≈ 2 mm ((–92$\pm$3) pixel count/m), which is more than twice the $1/e^2$ radius of the laser beam, we measure the lifetime of the 5s5p ^{3}P$_1$ state to be (21.8$\pm$0.9) μs. The linear fit, shown in the main body of figure 4.4, is calculated using the logarithmic least squares linear fit deduced for the data in the inset of figure 4.4 and $\overline{v} = 500$ m/s. This is in good agreement with the measured lifetime of the 5s5p ^{3}P$_1$ state is 21.4 μs [31], giving us confidence in the detection method.

Figure 4.4: 5s5p ^{3}P$_1$ lifetime measurement. Inset: a slice taken from the image in figure 4.3 along the direction of the atomic beam (black circles), where x defines the propagation distance and $N(x)$ the pixel count. Main: Decay time out of the 5s5p ^{3}P$_1$ state (black cirles) and a linear theoretical fit (red, solid line) calculated using the logarithmic least squares linear fit deduced for the data in the inset after a propagation distance of ≈ 2 mm and a most probable speed of $\overline{v} = 500$ m/s.

4.2 Saturation fluorescence spectroscopy

Saturation spectroscopy is based on the velocity–selective saturation of Doppler–broadened atomic transitions. In this scheme the spectral resolution is no longer limited by the Doppler width, which is present when a single beam passes through an atomic sample, but only by the much narrower width of the spectral

hole [68], which arises when two counter–propagating beams pass through an atomic sample.

Typically, in saturation fluorescence spectroscopy, the laser is split into a strong pump beam and a weak probe beam. For the conventional two–level atom depicted in figure 4.5 (a), the ground–state population as a function of laser frequency, ω_L, shows a Doppler–broadened profile (ω) with a spectral hole at a position depending on ω. The spectral hole arises due to stimulated absorption and emission. The conventional way to represent this is shown in figure 4.5 (b), where the spectral hole is burnt into the distribution at $v_z = +\Delta/k$, with $k = 2\pi/\lambda$ the wave vector of the pump beam and $\Delta = \omega_0 - \omega_L$ is the probe laser detuning from resonance, ω_0. The width and depth of the spectral hole depends on the saturation parameter s of the pump beam.

Figure 4.5: (a) A two–level atom with excited state $|e\rangle$ an energy $\hbar\omega_0$ above the ground state $|g\rangle$ is illuminated with laser light of angular frequency ω_L. The detuning from the resonance frequency is Δ. (b) The normalised ground–state population $\rho(v_z)$ of an ensemble of atoms with velocities v_z when the atoms are illuminated with a saturating pump beam detuned by Δ ($= \omega_0 - \omega_L$) from resonance, with an intensity equal to the saturation intensity, and a weak probe beam also at Δ. A spectral hole is burned into the distribution centred at $v_z = +\Delta/k$, with a maximum reduction in ground–state population (for this intensity) by a quarter. (c) The normalised ground–state population $\rho(v_z)$ in such a case of equal pump and probe beam intensities for $\omega_L \neq \omega_0$. Due to the opposite Doppler shifts of the two beams with angular frequency ω_L, two spectral holes are burnt into the distribution at velocity components $v_z = \pm\Delta/k$.

When the beams are of equal intensities two spectral holes are burnt into the distribution at velocity components $v_z = \pm\Delta/k$, owing to the opposite Doppler shifts of the two beams [68] (assuming that $\Delta > \Gamma_E$, where Γ_E is the power–broadened linewidth). As shown in figure 4.5 (c), when $\omega \neq \omega_0$ the atoms

with the velocity components $v_z = \pm(\Delta \pm \Gamma_E/2)/k$ undergo stimulated absorption and emission. However, when $\omega = \omega_0$, both beams are absorbed by the same atoms, which are essentially moving perpendicular to the laser beams at $v_z = (0\pm\Gamma_E/2)/k$, and the absorbed intensity is now twice as large. Figure 4.6 compares the normalised Doppler–broadened fluorescence profile that originates from a weak probe beam to that of the normalised fluorescence profile when two saturating pump and probe beams are present at $\omega = \omega_0$ on the $5s^2\,^1S_0 \rightarrow 5s5p\,^1P_1$ (a) and the $5s^2\,^1S_0 \rightarrow 5s5p\,^3P_1$ transition (b). This method of spectroscopy was initially described by Lamb, and is commonly known as the standing wave saturation effect with the spectral holes referred to as *Lamb dips*. The standing wave nature of the incident radiation does not play a major role in the Doppler–broadened line profile but does influence the Lamb dip. The width of the dip is dependent on both the transition linewidth and s, whereas the depth is dependent only on s.

Figure 4.6: Theoretical normalised ground–state population $\rho(v_z)$ as a function of atomic velocity v_z for (a) the $5s^2\,^1S_0 \rightarrow 5s5p\,^1P_1$ and (b) the $5s^2\,^1S_0 \rightarrow 5s5p\,^3P_1$ transition for saturation parameters $s = 10$ (red, solid line), $s = 100$ (blue, solid line), and $s = 1000$ (green, solid line). The normalised Doppler–broadened ground–state population in the absence of a counter–propagating pump beam is also shown (black, solid line).

To use the feature in figure 4.6 (b) as a reference for laser frequency stabilisation, the interrogating laser beam must be frequency modulated to permit phase–sensitive detection by restricting the frequency response of the detection system to a narrow frequency interval. By monitoring the fluorescence through a lock–in amplifier tuned to the modulation frequency, the demodulated output

of the lock–in amplifier results in a dispersion signal at $v_z = 0$ m/s (an example of such a dispersive signal is shown in figure 4.8).

Considering all of the fluorescence is monitored, removing the Doppler–broadened background requires modulating the pump and probe beams at different frequencies and referencing the lock–in amplifier to the sum of these frequencies. This technique, which detects very weak absorption lines, was demonstrated by Sorem and Schawlow [69], for which they received the Nobel prize in 1981 for their contributions to the development of laser spectroscopy. However, in the context of this thesis, both pump and probe beams originate from the same frequency modulated source, hence the Doppler–broadened background will always be present (which is evident in figure 4.8).

4.3   Experimental considerations

The Lamb dip shown in figure 4.6 (b) will be broadened from its natural 2π x 7.4 kHz wide Lorentzian profile by a number of broadening mechanisms. Power broadening has already been discussed, which is due to the saturation of the transition by the incident laser power resulting in a linewidth Γ_E given by equation 2.1. Other broadening effects arise due to the laser–atomic beam geometry and the modulation parameters of the interrogating light. Aligning the laser beam orthogonally to the atomic beam results in the atoms interacting with the light for a finite time of $t_{\text{trans}} \approx d/v_{\text{beam}}$, where d is the laser beam diameter and v_{beam} is the velocity of the atoms [56]. This finite interaction time Fourier limits the frequency resolution of the spectroscopy and results in a linewidth of $\nu_{\text{trans}} \approx 1/(2\pi t_{\text{trans}})$. To increase the interaction time, we expand the probing laser beam to a $1/e^2$ diameter of 10 mm, which leads to $\nu_{\text{trans}} \approx 42$ kHz for atoms propagating at the most probable speed ($\overline{v} = \sqrt{3k_{\text{B}}T/M}$ for atoms in an effusive beam [40], where T is the temperature of the atoms in the beam, M is the atomic mass, and k_{B} is Boltzman's constant) [40].

We shall discuss the other contributors to the broadening of the Lamb dip width, including a thorough investigation of the parameters of the interrogating light, in the following sections and will conclude with the optimum operating parameters for our system.

4.3.1 Frequency modulating the interrogating light

Due to equipment constraints, we only have access to one AOM to modulate both the pump and probe beam. Therefore, to produce the narrow Lamb dip feature, we retro–reflect the incident beam (see figure 4.2). The frequency modulated waveform is depicted in figure 4.7. Light from the master laser is aligned through an AOM, with centre operating frequency 80 MHz, at an angle that has optimum coupling efficiency at 75 MHz. The main carrier waveform is sinusoidal with a frequency $f_c = 75$ MHz (see figure 4.7 (a)). We chose to frequency modulate the main carrier waveform with a ramp waveform, since this gave a uniform spread in RF power over the frequency range we were interested in. We can control the deviation (i.e. the dithering amplitude) and the modulation frequency f_m (i.e. dithering rate) of this waveform (see figure 4.7 (b)). The frequency modulated waveform applied to the light is shown in figure 4.7 (c). The AOM is controlled using an arbitrary function generator (AFG3020), whose f_c, modulation shape, deviation and f_m can be varied easily. The AOM is aligned in double–pass configuration, such that the light is shifted by 150 MHz (i.e $2f_c$) from the master laser frequency. As a consequence, the deviation enhances by a factor of two, however f_m does not.

Figure 4.7: Frequency modulating an AOM. (a) Laser light is aligned through an AOM at an angle that has optimum coupling efficiency at f_c and a carrier waveform that is sinusoidal with a frequency f_c. (b) Modulation waveform applied to the waveform in (a), whose dithering amplitude (deviation) and dithering rate (f_m) can be varied. (c) Resulting frequency modulated waveform

.

The signal from the fluorescence photodiode, the circuit and design of which is given in section 4.3.3, is demodulated using a Standford Research lock–in amplifier (model SR830 DSP), which is referenced to the 'Trigger Output' of the arbitrary function generator at f_m. Figure 4.8 shows the magnetically insensitive atomic error signal, when the polarisation of light is set to drive to a π transition, at the output of the lock–in amplifier when the 5s5p ^{3}P$_1$ state is Zeeman split via a magnetic field (a detailed discussion on Zeeman splitting is provided in section 4.3.4). To obtain the error signal, the laser frequency is stabilised to a high–finesse optical cavity, then, to change the laser detuning, the cavity length is scanned via a piezo incorporated in the cavity length.

The entire Doppler–broadened background, as well as the error signal that will be used in the laser frequency stabilisation feedback loop, is shown in figure 4.8. We do not remove the Doppler–broadened background signal as a result of modulating both the pump and probe beams at the same frequency, and referencing the lock–in amplifier to this frequency. We can reduce the height of the

Doppler–broadened background signal by appropriately choosing the modulation parameters, as we shall see in section 4.3.5.

Figure 4.8: Atomic spectroscopy error signal, after the lock–in amplifier, on the $5s^2 \ ^1S_0 \ (m_J = 0) \rightarrow 5s5p \ ^3P_1 \ (m_J = 0)$ transition when the cavity length is scanned at a frequency of 10 mHz over the entire Doppler width. The detuning axis has been calibrated using the cavity piezo response quoted in chapter 3, i.e. 1.67 kHz/mV.

The finite interaction time that arises as a consequence of the scanning the cavity length Fourier limits the frequency resolution of the spectroscopy, hence it is necessary, particularly during the optimisation phase, to scan the cavity length as slowly, and over as small a cavity length change, as is possible. In figure 4.8, the cavity length is scanned at a frequency of 10 mHz. Equivalently, modulating the laser frequency, whereby the deviation and f_m are controllable, also Fourier limits the frequency resolution thereby introducing a linewidth ν_{mod}. We will shown in section 4.3.5 that the Lamb dip width, as well as the height, is governed by a combination of Γ_E, ν_{trans}, and ν_{mod}.

4.3.2 Role of the lock–in amplifier

The lock–in amplifier is a phase–sensitive detector that provides an additional gain stage. The main features are the 'Sensitivity' control and the 'pre–' and 'post–time constant' filters, τ_{pre} and τ_{post} respectively. The 'Sensitivity' amplifies the entire signal and does not change the signal–to–noise ratio. This can be considered as an additional proportional gain stage in the low–bandwidth locking circuit. The filters, τ_{pre} and τ_{post}, are low pass and should be chosen such that they filter out f_m. For example, if $f_m = 291$ Hz, then $\tau_{\mathrm{pre}} \geq 1/f_m \approx 3.4$ ms.

The role of τ_{post} is to increase the 3 dB roll off. The maximum f_m that we can use, and filter out using solely the lock–in amplifier, is 1 kHz (i.e. minimum value for τ_{pre} is 1 ms). The greater τ_{pre} that is used, the greater the signal–to–noise ratio. There is a constraint when increasing τ_{pre}, however, and this is a reduction of the bandwidth of the loop filter. Since the feedback loop is stabilising a slowly varying cavity length (i.e. due to temperature variations) this should not be hindrance, but it is something to take note of when using lock–in amplifiers.

4.3.3 Detecting the signal

The fluorescence is detected using a homebuilt, high–transimpedance gain photodiode. The circuit diagram is shown in figure 4.9. The photodiode (Centronic OSD15–5T) is suited to detect low light levels (with a quoted minimum responsivity of 0.18 A/W at $\lambda = 436$ nm) and has an active area of 3.8 x 3.8 mm^2, with a maximum dark current of 5 nA [73]. It is important to choose a photodiode with low inherent dark current, since this is the major source of noise. The operational amplifier used is an AD795 [74] and is quoted as having low noise (low input noise of 3.3 μV$_{\text{pp}}$ from 0.1 Hz to 10 Hz and a maximum noise level of 0.6 fA/$\sqrt{\text{Hz}}$ at 1 kHz) and a high DC accuracy (2 pA maximum input bias current, 500 μV maximum offset voltage and low supply current of 1.5 mA maximum). A low offset voltage is necessary as we do not want the error signal zero–crossing to drift, in particular during the experimental cycle time. Furthermore, the 10^{10} Ω common–mode impedance ensures that input bias current is independent of common–mode voltage and supply voltage variations.

The AD795 is designed for mounting on printed circuit boards. Special precautions need to be taken when designing the circuit since voltage differences between the input pins and other pins cause parasitic currents larger than the AD795's input current, hindering the pA resolution. There are two guidelines that must be adhered to: the first of which is to include guarding of the input lines; and the second is to maintain adequate insulation resistance, i.e. the circuit layout should be compact and surrounded by a ground plane to reduce the length of input lines and to shield the high impedance precision circuitry from electrical noise and interference, respectively. We have taken these considerations into account and the circuit is shown in figure 4.9.

Figure 4.9: Low–noise photodiode circuit to detect 689 nm fluorescence in an atomic beam. C_1 and C_2 are decoupling capacitors to minimise noise on the power lines, the source of which are two 9 V batteries. A switch, S_1, controls the 'on/off' state of the photodiode circuit. The feedback resistor, $R_1 = 20$ MΩ, sets the gain and the signal–to–noise ratio of the photodiode circuit. There are two signal outputs (signal + and signal –) to prevent ground loops inducing a change in DC level. Also shown is the guard ring around the inverting input of the operational amplifier.

The ground of the photodiode is decoupled from the environment to reduce any pick–up and is battery–operated. The photodiode itself is sealed within a metal container that behaves like a Faraday cage and has a 1:1 imaging system, i.e. a $2f$ imaging system with a 40 mm diameter bi–convex lens with a focal length of $f = 40$ mm (Comar 40VB40), with the focus at the centre of the atomic beam machine vacuum chamber ($\approx$ 80 mm from the lens), see figure 4.2. The container is mounted onto the top viewport of the atomic beam machine. The two photodiode output ports (one of which is the photodiode circuit ground) are connected to a Stanford Research lock–in amplifier (model SR830 DSP) in a differential configuration, permitting phase–sensitive detection to improve signal–to–noise ratio.

We cannot observe the raw, unmodulated signal as the photodiode does not have sufficient gain. A possible modification to the circuit design, which might overcome this issue, is increasing the feedback resistor R_1 in the photodiode circuit given in figure 4.9. Alternatively, we might be able to increase the signal–to–noise ratio by placing a curved mirror on the underside of the vacuum chamber. It has been shown by the YbCs group in Durham, who have used our photodiode imaging design as a detector in their 556 nm laser frequency stabilisation feedback loop, that placing a curved mirror on the underside of

their vacuum chamber increases their signal–to–noise ratio by a factor of two.

4.3.4 Zeeman splitting the 5s5p ^{3}P$_1$ state

Stray magnetic fields inside the atomic beam machine vacuum chamber will Zeeman split the 5s5p ^{3}P$_1$ state into the three m_J sublevels, effectively broadening the spectroscopy measurement. This phenomenon is illustrated in figure 4.10. To prevent broadening, coils are wound in Helmholtz configuration around the vacuum chamber along the axis of the atomic beam. The 5s^2 ^{1}S$_0$ ($m_J = 0$) $\rightarrow$ 5s5p ^{3}P$_1$ ($m_J = 0$) transition is chosen as it is insensitive to the Zeeman shift ($g_J \mu_B B$, where $g_J = 1.5$, which is the Landé g–factor for the 5s5p ^{3}P$_1$ states, $\mu_B = 1.4$ MHz/G is the Bohr magneton, and B is the applied magnetic field). Therefore, the polarisation is chosen to drive the π transition (figure 4.10 (a)). Figure 4.10 (b) illustrates the Zeeman splitting into the three magnetic sublevels when a field of 14 G is applied. The laser detuning is varied by scanning the cavity length and the detuning axis is calibrated using a WS7 HighFinesse wavemeter. The central atomic error signal is due to the $m_J = 0$ and has an amplitude roughly three times as large as the error signals due to the $m_J = \pm 1$. We do not expect to see the error signals due to the $m_J = \pm 1$ if the polarisation of light is pure. We might be able to remove the signal due to the $m_J = \pm 1$ states by placing a polarising beam splitter cube in the beam path before the atomic beam machine vacuum chamber to purify the polarisation. We measure a splitting of ≈ 30 MHz, which is in agreement with the splitting we expect from $g_J \mu_B B$.
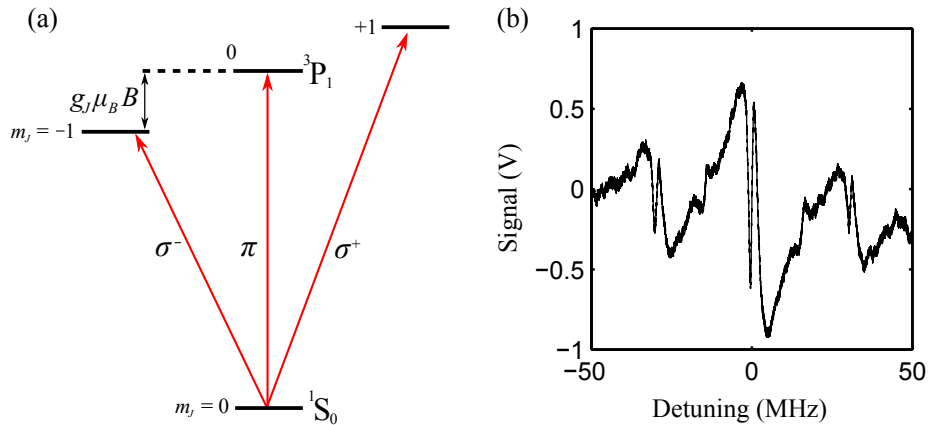
Figure 4.10: (a) Zeeman splitting of the 5s5p ^{3}P$_1$ state into the three magnetic sub-levels by an amount $g_J\mu_B B$ when a magnetic field B is present, where g_J is the Landé g–factor ($g_J = 1.5$ for the 5s5p ^{3}P$_1$ states) and $\mu_B = 1.4$ MHz/G is the Bohr magneton. The m_J states are selected via a choice of laser light polarisation. (b) Zeeman splitting of the 5s5p ^{3}P$_1$ state by applying a magnetic field using Helmholtz coils wound along the axis of the atomic beam when a magnetic field of 12 G is applied. The detuning axis is calibrated using a WS7 HighFinesse wavemeter.

We typically use a field of 7 G, which shifts the $m_J \pm 1$ states ≈ 15 MHz from the magnetic insensitive transition, thereby allowing us to probe the $\Delta\, m_F = 0$ transition only and effectively increasing the error signal gradient.

4.3.5   Optimising the error signal

Having already constrained ν_{trans} by fixing the laser beam diameter and scanning the cavity length as slow as we possibly can (10 mHz), we will investigate the effects of Γ_E and ν_{mod} to find the best combination of signal–to–noise ratio and error signal FWHM. The definitions of these are shown in figure 4.11.

Figure 4.11: Main: Error signal full–width–at–half–maximum (FWHM) and height definition. The gradient is defined as height/FWHM. Inset: Doppler–broadened background signal with the error signal due to the $5s^2\ {}^1S_0\ (m_J = 0) \to 5s5p\ {}^3P_1\ (m_J = 0)$ transition at 0 MHz detuning. The same definitions for height, FWHM and gradient apply for the Doppler–broadened background signal too. The signal–to–noise ratio is expressed as the ratio of the Error signal FWHM (height):Doppler–broadened background signal FWHM (height).

Modulation parameters

We use a value of f_m that is a prime number, such that 50 Hz mains noise will not be amplified by the lock–in amplifier, and the lock–in amplifier parameters, such as pre– and post–time constant, sensitivity, etc., remain constant for a controlled investigation of the deviation.

Figure 4.12: Error signal (a) FWHM (b) gradient for $f_m = 291$ Hz for different AOM deviations with a fixed single beam peak intensity of 800 μW/cm^2. The laser linewidth is narrowed via stabilisation to a high–finesse cavity. To obtain the error signal height and FWHM, the cavity length is varied by scanning the cavity piezo at a frequency of 10 mHz. Standard error bars are also shown.

The error signal FWHM and error signal gradient are shown in figure 4.12. At a single beam peak intensity of 800 μW/cm^2, the power–broadened linewidth is ≈ 162 kHz (using equation 2.1), therefore we will never measure an error signal width less than this. The minimum FWHM we measure is ≈ 400 kHz at a deviation of 100 kHz, as shown in figure 4.12 (a), and the FWHM increases linearly with increasing deviation. The broadening is due to the finite interaction time on resonance, which is approximately (Deviation x f_m)/Γ_E. For example, for a deviation of 1 MHz and $f_m = 291$ Hz, there is a broadening of ≈ 1.8 MHz, which is very close to what is measured in figure 4.12 (a). Figure 4.12 (b) illustrates the error signal gradient as a function of deviation. As the deviation increases, the gradient decreases, therefore it seems beneficial to use a deviation that is as small as possible.

Additionally, we find the deviation to have a profound effect on the ratio of the error signal to the Doppler–broadened background signal, as we observe in figure 4.13. As the deviation increases, the height ratio decreases fairly linearly (figure 4.13 (a)), meaning that we have restricted the observation range beyond the error signal feature, therefore we are sampling the fluorescence from the Doppler–broadened background. The maximum error signal height to Doppler– broadened background signal height we measure is ≈ 2 for a deviation between

100–200 kHz.

Furthermore, there appears to be a non–linear effect on the FWHM with deviation. From figure 4.13 (b) we find that with increasing deviation, the FWHM ratio follows a saturation–type trend, with a maximum error signal FWHM to Doppler–broadened FWHM ratio of ≈ 0.1 (i.e. the error signal FWHM saturates to a value that is one tenth of the Doppler–broadened background FWHM). It was discussed in section 4.3 how the saturation parameter has a non–linear effect on the Lamb dip width (i.e. the error signal FWHM) but it was not known that the deviation has a similar effect.

Figure 4.13: Error signal to Doppler–broadened background ratio of (a) height and (b) FWHM for $f_m = 291$ Hz and different AOM deviations with a fixed single beam peak intensity of 800 μW/cm^2. The laser linewidth is narrowed via stabilisation to a high–finesse cavity. To obtain the error signal height and FWHM, the cavity length is varied by scanning the cavity piezo at a frequency of 10 mHz. Standard error bars are also shown.

To summarise, using very low deviations will result in an error signal that has a height that is approximately double that of the Doppler–broadened background signal. Using large deviations result in a width that has a non–linear dependence on the deviation.

Power broadening

The effect of power broadening can only be observed if the linewidth of the Lamb dip, or rather the FWHM of the error signal, is not limited by ν_{trans},

ν_{mod} or the broadening due to scanning the cavity length. Assuming the latter is negligible, and setting the deviation to 50 kHz, which is on the order of ν_{trans}, we observe the power dependence on the error signal FWHM and height (figure 4.14). The FWHM shows a non–linear dependence on intensity, which is what we expect from the discussion in section 4.3. The minimum FWHM we measure is $\approx$ 50 kHz, which, again, is on the order of ν_{trans} and the deviation. With reducing intensity, the error signal height decreases (figure 4.14 (b)), as well as the signal–to–noise ratio, making the error signal difficult to find, and, subsequently, to lock to.

Figure 4.14: Variation of the (a) FWHM of the atomic error signal and (b) height with single laser beam peak intensity when a deviation and f_m of 50 kHz and 291 Hz are used, respectively. Standard error bars are also shown.

To summarise, using a deviation that is on the order of transit–time broadening, we observe the non–linear Lamb dip width dependence on the saturation parameter. Using low intensities hinders the signal–to–noise ratio, therefore when we come to use this error signal in frequency stabilisation, we will need to find a compromise between power and deviation.

4.3.6   Conclusion

We have investigated the parameters that influence both the error signal FWHM and height, such as laser light modulation parameters and intensity. We have discovered the deviation has a non–linear effect on the error signal FWHM, and hence the Lamb dip. We can reduce the error signal FWHM by reducing the deviation, which results in an error signal height that is double that of the

Doppler–broadened background. We cannot completely remove the Doppler–broadened signal, which is what we expect from section 4.2, but we can reduce it by using a small deviation.

We have also demonstrated the non–linear dependence of the FWHM on the saturation parameter when we use a deviation that is comparable to transit–time broadening effects. Using very low peak intensities ($\approx 100\ \mu\text{W/cm}^2$ in the single pass beam), we measure a FWHM of 50 kHz, however the signal–to–noise ratio is very small, making it difficult to visually extract and lock to the error signal. Furthermore, we have discussed the purpose of the lock–in amplifier sensitivity and time constants, and how to chose these appropriately.

As one final note, the SYRTE group [71], who also implement this fluorescence spectroscopy technique, claim to observe no better than 200 kHz, which we believe is accountable to power broadening, choice of deviation and, perhaps, cavity length scan frequency.

4.4   Low–bandwidth feedback electronics

An illustration of the low–bandwidth locking circuit used for laser frequency stabilisation is shown in figure 4.15. The circuit design is simple including a PI feedback stage and an offset stage for tuning the cavity, and hence the laser, onto resonance with the atoms. The circuit does not include a DC offset stage to zero the error signal as there is one inherent in the lock–in amplifier.
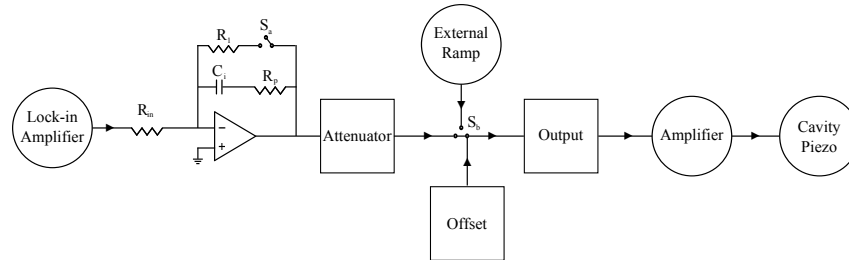
Figure 4.15: Low–bandwidth locking circuit (operational amplifier stages are shown in black boxes) including input and output stages (shown in black circles). The double-pole, double–throw switch (S_a and S_b) simultaneously connects the output to the error input and changes the operation of the operational amplifier to an integrator. The resistor R_1 prevents the capacitor C_i from charging and imposing a voltage spike onto the cavity piezo when the integrator is activated. The 'Attenuator' and 'Offset' stages are used to limit the integrator rail voltages, limit the minimum output voltage to ≈ 0.5 V and to manually tune the cavity length such that the laser light is resonant with the atoms. The 'Amplifier' stage is a commercial Thorlabs piezo amplifier.

The circuit contains a ganged switch, which, simultaneously connects the output to the error input and changes the operation of the operational amplifier to an integrator. The resistor R_1 prevents the capacitor C_i from charging and imposing a voltage spike onto the cavity piezo when the integrator is activated. The 'Attenuator' and 'Offset' stages are used to limit the integrator rail voltages, limit the minimum output voltage to ≈ 0.5 V and to manually tune the cavity onto resonance with the atoms.

The circuit does not have to be high–gain, high–bandwidth like the one discussed in chapter 3, but it does have to be low–noise to reduce any technical noise that might appear as flicker frequency noise on the laser. It is important when choosing components that the operational amplifier not only adheres to the low–noise criteria but that it also has a low voltage offset drift. If the operational amplifier is not a precision operational amplifier, any voltage drifts will lead to a DC offset on the error signal, hence changing the laser frequency as the zero–crossing moves around. We use an OP270 (low noise density of 5 nV/$\sqrt{\text{Hz}}$ at 1 kHz maximum, input offset voltage of 75 μV maximum and an offset drift of less than 1 μV/$^\circ$C according to the producers datasheet [75]). The OP270 is unity gain stable and capable of driving large capacitive loads, such as piezo attenuators, without oscillation. It is important when choosing input/feedback resistors for this choice of operational amplifier as some values

lead to lower inherent chip noise. To reduce total noise, which is dependent on operational amplifier voltage noise, operational amplifier current noise and resistor noise, the input resistance must be kept to a minimum. Additionally, to minimise output noise introduced via the amplifier supply pins, the power lines have been choked and capacitively decoupled. The full circuit diagram is shown in figure A.2, appendix A and the time constant of the PI feedback is on the order of a few seconds.

4.4.1 Circuit output amplification

In the cavity design, two piezos were included in the length of the cavity: one piezo would coarsely tune the cavity length such that the laser light is resonant with the atoms; the other piezo would be used in the stabilisation loop [59] to provide fine length corrections. Unfortunately, during the cavity build, the tuning piezo disconnected from the feedthrough, leaving only one piezo to provide a tuning offset voltage and a feedback voltage to stabilise the cavity length to an atomic resonance. Even though the cavity is temperature stabilised, if the cavity changes temperature by 100 mK (which is roughly what we believe the cavity temperature is stabilised to within, see appendix B), this would result in a frequency shift of 37 MHz. To bring the cavity back on resonance with the atoms, a voltage of 22 V needs to be applied to the piezo (i.e. 1.67 kHz/mV shift). The maximum output of the low–bandwidth locking circuit is 15 V. Thus we require a piezo amplifier on the output of the circuit. Amplifiers can add significant noise to the loop filter. Although there are low–noise amplifiers commercially available these are quite expensive. Throughout the context of this thesis, we implement a Thorlabs piezo amplifer (MDT693A) despite the fact that these have large inherent 50 Hz noise. The output noise is specified as 1.5 mV$_{\mathrm{RMS}}$ for an input signal of $\approx$ 9.9 V$_{\mathrm{pp}}$ when an external capacitive load is not connected (0.047 μF output impedance only). Adding a capacitive load, i.e. the cavity piezo, will decrease the noise by creating a low pass filter with the output resistance. We can assume a maximum of 1.5 mV$_{\mathrm{RMS}}$ of noise, which corresponds to a broadening of $\approx$ 2.5 kHz at 50–60 Hz [76].

The lock–in filters are set to $\tau_{\mathrm{pre}} = 100$ ms and $\tau_{\mathrm{post}} = 0.1$ s, such that the cut–off frequency of the filter is 10 Hz. Although we can not observe 50 Hz noise on the error signal, this does not mean the noise is not present on the

cavity piezo. The circuit is designed to remove slowly varying cavity length drifts, hence it is unable to remove 50 Hz noise. Possible methods for removing/suppressing the noise due to the piezo amplifier is discussed in section 4.5.

4.5 Evaluating the loop filter

To measure the long–term frequency stability, the laser frequency must be compared to a much more stable and reliable frequency standard. Frequency metrology groups typically measure the frequency stability of their laser system by monitoring a beat between the laser and a frequency comb that is locked to a hydrogen maser [33] over many hours or days. Although we do not have access to a frequency comb, we can use the cold atoms as a frequency standard. We can measure the Allan deviation (section 2.3) of the $1/e^2$ Gaussian width of the red MOT cloud in the horizontal direction to deduce whether there is a long–term frequency instability, i.e. a cavity length drift. Using equation 2.3, we can determine the detuning fluctuations that cause the MOT width variations.

4.5.1 Long–term frequency stability

Figure 4.16 is an Allan deviation measurement of the $1/e^2$ Gaussian width of the red MOT cloud for two AOM deviations, using equation 2.7. We expect the drift to be worse with increasing deviation used. The minimum absolute stability is ≈ 1 kHz and ≈ 2 kHz for a deviation of 100 kHz and 200 kHz, respectively, at an averaging time of $\approx 10^2$ s. This agrees with our earlier discussion that the modulation parameters in the saturated fluorescence spectroscopy technique contribute to the linewidth of the Lamb dip. At this time, the main contributor to noise is due to white frequency noise and we observe no drift in our system. To observe whether there is drift, we would need to measure the Allan deviation over much longer averaging times.

Figure 4.16: Cavity length stability for (a) a deviation of 100 kHz (black circles) and (b) a deviation of 200 kHz (red squares). All other parameters, such as f_m, saturation parameter s, τ_{pre} and τ_{post} etc. are those listed in 4.3.6 and they do not change.

On short timescales (on the order of several seconds), the absolute frequency stability is $\approx \pm 10$ kHz and the dominant noise is white phase and flicker phase. However, this is an upper bound for the absolute frequency stability since other effects, including an oscillating magnetic field gradient or polarisation fluctuations out of the MOT beam fibre, could be contributing to the changing MOT width. Unfortunately, during the work of this thesis, we were unable to find the cause. The AtMol group is, however, purchasing a frequency comb, which will make the process of finding the source of the noise much easier.

4.5.2   Remarks

We observe a DC offset drift in the fluorescence signal on the order of tens of kHz over the duration of the day, which should be apparent in the Allan deviation plot if we measured the fluctuating MOT cloud width for longer averaging times. We have yet to alleviate this problem, but we have considered implementing the intermodulated fluorescence spectroscopy technique discussed in section 4.2 should we have access to another AOM.

We also find the alignment of the incident beam and retro–reflected beam causes the zero–crossing of the error signal to change and also causes broadening. If the beams are poorly overlapped, atoms with a non–zero velocity class are addressed and we see the sharp peak move from the central position of ω_0, which agrees

with work in [77]. Courtillot's thesis [71] provides an in–depth discussion on a more precise beam alignment procedure than the one used in this thesis. We have not implemented the alignment technique in [71] due to time and equipment constraints. This is a possible improvement to the set–up, however, until we find a method of determining the absolute frequency of our laser system, this is a fruitless pursuit.

We have not undertaken a thorough optimisation of the low–bandwidth feedback circuit. We could do this by measuring the $1/e^2$ Gaussian width of the red MOT cloud over many experimental cycles for various gains and integrator time constants. However, this can not be fulfilled in real–time and is quite a tedious procedure. However we may be able to minimise the noise due to the piezo amplifier.

Conclusion

We have discussed and implemented a saturation spectroscopy technique to detect and stabilise a laser system to a narrow atomic transition. We use saturation fluorescence spectroscopy and find that the modulation parameters, as well as the laser beam power, contribute most to the linewidth of the error signal detected and also to the stability of the laser frequency when the feedback loop is engaged. Without the use of a stable reference as a frequency diagnostic, we have implemented an approach to measure the frequency stability of the laser system by measuring the $1/e^2$ Gaussian width of the red MOT cloud in the horizontal direction. Using homebuilt, low–bandwidth feedback electronics, we stabilise the short–term laser frequency to within, at worse, $\approx \pm 10$ kHz and the long–term frequency to within $\approx \pm 1$ kHz in 15 minutes. This is an upperbound for the long–term laser stability since other effects, including an oscillating magnetic field gradient or polarisation fluctuations out of the MOT beam fibre, could be contributing to the changing MOT width.

The long–term instability (on this interrogation time–scale) is less than the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition linewidth, hence we are pleased with this outcome. In particular, owing to the simplicity and ease of the set–up. A possible improvement to the technique includes placing a curved mirror on the underside of the vacuum chamber on the axis of the fluorescence photodiode to increase

the signal–to–noise ratio. In this configuration, we might be able to work at smaller deviations and a lower interrogating beam power, which will reduce the width of the Lamb dip feature, therefore improving the long–term frequency stability. One other possible improvement involves optimising the gains and integrator time constant of the feedback circuit via the Allan deviation measurement. This will be a tedious and laborious task, however we may be able to minimise the noise due to the piezo amplifier.

Chapter 5

Two stages of cooling

In this chapter the laser systems for cooling are discussed as well as the vacuum chamber set–up. Although one would like to cool the Sr vapour using only the narrow linewidth transition, the photon scattering rate is simply too low to capture the atoms. Thus, we use two stages of cooling where we first load the atoms into a standard MOT operating on the broad, strong transition at 461 nm, which is often referred to as the 'blue' MOT. The atoms are pre–cooled to a few mK before they are transferred to a second stage cooling MOT operating on the weak intercombination transition at 689 nm, hence we call this the 'red' MOT, whereby we achieve an atom temperature of a few μK at densities a thousand–times greater than those obtained in the blue MOT [33, 38, 78].

5.1 First stage: Blue MOT

The laser systems for the creation of the blue MOT, experimental vacuum chamber and imaging technique (including atom number and peak density extraction) are described in this section. The majority of the construction and characterisation of the blue MOT was carried out by J. Millen [27] and G. Lochead [28].

5.1.1 461 nm laser system

The laser radiation used to slow and capture the Sr atoms on the dipole allowed $5s^2$ $^1S_0 \rightarrow 5s5p$ $^1P_1$ transition is a frequency–doubled diode laser source at 922 nm. We use a Toptica TA–SHG 110, which has a typical output power of

250 mW. The laser is fibre coupled at its output as the beam is susceptible to walking when, for example, cavity realignment is carried out to optimise output power. We typically work with ≈ 150 mW of power out of the fibre. The beam at the output of the fibre is split for laser locking, the Zeeman beam (a discussion of which will follow), the MOT beams, and finally the probe beam, which is used in Rydberg experiments (further discussion follows in chapter 6), see figure 5.1. The frequency of each beam, as well as experimental on/off control, is set using acousto-optic modulators (AOMs).
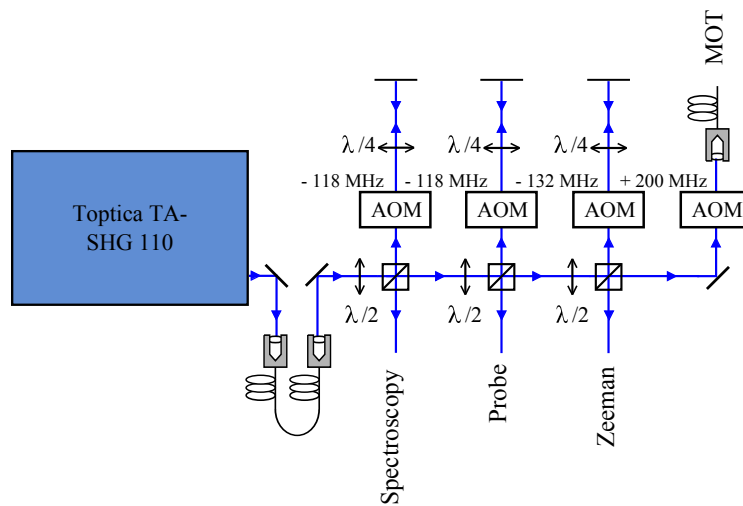
Figure 5.1: Schematic set–up of the 461 nm laser system. The light is coupled down a single–mode polarisation–maintaining fibre as the Toptica output beam is susceptible to walking. We typically work with ≈ 150 mW of power out of the fibre. The output of the fibre is split into four main paths for: spectroscopy; probe light, which is used in Rydberg experiments; Zeeman slowing light and; MOT light. The frequencies of these beams are set using acousto–optic modulators (AOMs) either aligned in double–pass or single–pass configuration.

5.1.2   Laser stabilisation

The frequency of the 461 nm laser requires stabilisation in order to produce a blue MOT. This is achived via stabilisation to an atomic signal, though there are many methods employed such as locking to a stable reference cavity (NPL [79]) or using hollow cathode lamps (JILA [80]). Since strontium has negligible vapour pressure at room temperature, and is highly corrosive with glass, the strontium source is produced by heating strontium dispensers contained within a vacuum cell [81].

Figure 5.2: 461 nm laser frequency stabilisation by means of modulation transfer spectroscopy (MTS) using a strontium vapour cell. This scheme employs a strong pump, which is modulated using an electro–optic modulator (EOM), and a counter–propagating weak probe. Modulation is transferred from the pump to the probe and this is detected, via the probe light, on a fast photodiode (FPD). A phase shifter is used to set the phase, ϕ, between the local oscillator input and the RF input of the double–balanced +7 dBm mixer. The output of the mixer is filtered and is fed back to the laser ECDL piezo. The electronic and optical paths are illustrated using black, dashed lines and blue, solid lines, respectively.

There are many ways to stabilise a laser to an atomic signal, however some methods are more suited depending on the type of atom. We employ a modulation transfer spectroscopy (MTS) laser frequency stabilisation scheme [82]. The MTS error signal is almost free of Doppler–broadened background signal (see figure 5.3 for an example error signal) and strong signals are produced for closed transitions. This method is ideal for use in strontium as there is only one decay channel out of the 5s5p $^1P_1$ to the 5s4d $^1D_2$ state; this is not problematic as this occurs once in every 50,000 cycles, hence atom losses from the stabilisation transition are small.

Figure 5.3: Example modulation transfer spectroscopy error signal used for stabilising the 461 nm laser frequency via electronic feedback to the ECDL piezo. The probe and pump intensities are 6 mW/cm^2 and 1 W/cm^2, respectively. The detuning axis is calibrated using a WS7 HighFinesse wavemeter. The figure has been taken from G. Lochead's thesis [28].

The MTS locking scheme is shown in figure 5.2. In the MTS scheme a pump beam is modulated by an electro–optic modulator (EOM) at 9.3 MHz, which passes through the vapour cell. The modulation is transferred to a probe beam that is counter–propagating to the pump beam via a four–wave mixing process in the atomic media. Prior to the EOM, the beam used for frequency stabilisation is aligned through an AOM in double–pass configuration at an angle that has optimum coupling efficiency at 118 MHz. This is split into the pump and probe beam via a half waveplate and polarising beam splitter cube. We use $\approx$ 2 mW of power in the pump beam and $\approx$ 120 μW of power in the probe and the probe is detected on a homebuilt fast photodiode (FPD). For more information on the photodiode design see G. Lochead's thesis [28]. The error signal is fed back to the Toptica piezo in the ECDL and is stabilised via a homebuilt PI feedback circuit, which is available in appendix A.

5.1.3 Vacuum chamber and MOT optics

An overview of the apparatus, which was developed during the thesis of J. Millen [27], is shown in figure 5.4. An atomic vapour is produced by heating 5 g of strontium to 350–400 °C. The atomic beam source is a conventional effusive oven design based on that of [83], whereby dendritic pieces of 99.9% pure strontium

(Sigma–Aldrich 460346–5G) are heated in the oven using resistive heater wire (Thermocoax SEI 20/150). The atoms traverse 169 stainless steel capillaries of 170 μm inner diameter to form an atomic beam of 43 mrad full width divergence at a mean velocity of $\bar{v} \approx 430$ ms^{-1}. The atom beam is directed towards an ultra–high vacuum (UHV) chamber via a spin–flip Zeeman slower [40] that is 25 cm long and is not water cooled. The magnetic field for the Zeeman slower was designed by M. P. A. Jones using algorithms from [84], and has been designed such that the stopping distance at $I = I_{\text{sat}}$ (i.e. $a = \frac{a_{\text{max}}}{2}$) is ≈ 20 cm. The Zeeman slower coils are housed in a mild steel yolk (red box shown in figure 5.4) to reduce the magnetic field at the chamber–side of the slower. Finally, the slowed atomic beam enters the UHV chamber, which is pumped with a 20 l/s ion pump (Gamma Vacuum TiTan 20S) and a non–evaporative getter pump (SAES Getters CapaciTorr B200) to remove hydrogen.

Figure 5.4: Vacuum chamber used in the cold atom experiment. The apparatus consists of an oven, whereby the atoms are heated to a temperature of ≈ 400 °C, to produce an atomic beam. The atoms are slowed via an on–resonant laser beam and a 'spin–flip' Zeeman slower, and cooled and trapped using three pairs of orthogonal retro–reflected MOT beams and a quadrupole magnetic field gradient [85].

To load the first stage MOT, Zeeman light that is red–detuned by $16.6\Gamma \approx 500$ MHz from the 5s^2 $^1\text{S}_0 \rightarrow 5\text{s}5\text{p}$ $^1\text{P}_1$ transition, which is counter–propagating to the atomic beam, is used. This Zeeman beam is focused onto the aperture of the atomic beam using a $f = +750$ mm lens, such that the Zeeman beam is mode–matched to the atomic beam. A typical Zeeman power of 35–40 mW is used, hence the blue MOT is not operating in the saturated atom number regime (refer to G. Lochead's thesis [28]). There is not enough 461 nm power available

to alleviate this.

The 461 nm MOT light is coupled into a single–mode, polarisation–maintaining fibre, which is mounted into a cage system to aid the 689 nm MOT light overlap. At the output of the fibre, the beam is expanded to a $1/e^2$ radius of 0.5 cm, by means of a telescope that is also mounted in the cage system, and is split into three beams. These beams intersect orthogonally and are retro–reflected to produce a six beam MOT [85]. Approximately 5 mW in each axis is used at a detuning of $-1.2\Gamma \approx -36$ MHz. The gradient of the quadrupole field, which is created by a pair of anti–Helmholtz coils inside the vacuum chamber, is 24 G/cm in the axial direction.

5.1.4   Fluorescence imaging

To obtain an atom number and cloud density, fluorescence imaging is implemented, whereby an image of the atom cloud is obtained using the blue MOT beams and a charge–coupled device (CCD) camera (pco. "PixelFly qe") with a 1:1 imaging system. The background of the fluorescence image is removed within the LabVIEW experimental control program, and slices, centred around the point of maximum intensity, are extracted along the z– and x–axis. Each slice has a width of a single pixel, which has a size of 6.45 μm x 6.45 μm. These slices permit the physical size of the cloud to be measured. By fitting a Gaussian to the MOT cross–sections, the $1/e^2$ Gaussian width of the cloud, w_x and w_z, is extracted along the x– and z–axis, respectively.

The number of atoms contained within the cloud is estimated using the fluorescence image and the detection efficiency of the CCD. The collection efficiency of the CCD, which was determined during the thesis of J. Millen, is measured to be $\mu = (0.132\pm0.006)$ counts per pixel per photon. Using this value of μ, the total number of photons detect in an exposure time t_{exp}, which is variable, is calculated via,

$$N_\gamma = \frac{\sum_{x,z=1}^{n} C}{\mu t_{\mathrm{exp}}} \, , \tag{5.1}$$

where C is the recorded number of counts per pixel.

Modelling the 5s^2 ^{1}S$_0$ $\rightarrow$ 5s5p ^{1}P$_1$ transition as a two–level atom, and neglecting small losses via decay to the 5s4d ^{1}D$_2$ state, the number of photons

scattered per atom is given by [40]

$$R = \frac{\frac{\Gamma}{2} s}{1 + s + \frac{4\Delta^2}{\Gamma^2}} \quad , \tag{5.2}$$

where Γ is the linewidth of the $5s^2\ ^1S_0 \to 5s5p\ ^1P_1$ transition, Δ is the detuning of light from resonance, and $s = I/I_{\text{sat}}$ is the saturation parameter, where I is the intensity of the MOT beams, and I_{sat} is the saturation intensity.

Using equations 5.1 and 5.2, an atom number N_a is extracted,

$$N_a = \frac{N_\gamma}{\Omega R (1 - A)} \quad , \tag{5.3}$$

where Ω is the solid imaging angle (which is 0.0037 steradian in this system) and A is the beam power losses due to absorption/scattering off of viewports.

Using equation 5.3, a peak density n_0 can be deduced from

$$n_0 = \frac{N_a}{(2\pi)^{3/2} w_x w_y w_z} \quad , \tag{5.4}$$

where w_y is the width of the cloud in the unseen direction, which is set to be the geometric mean of w_x and w_z.

5.1.5   Repump lasers

There is a small leakage of atoms out of the $5s5p\ ^1P_1$ state to the lower lying $5s4d\ ^1D_2$ state allowing atoms to escape the cooling cycle and be lost from the trap. In spite of the fact that the decay rate is sufficiently slow (branching ratio of approximately 50,000:1 [42]) there is an appreciable reduction in atom number, MOT density and trap lifetime observed. In order to recycle the atoms stored in the metastable $5s5p\ ^3P_2$ state, two repump lasers at 707 nm and 679 nm are used. The so–called 'repump' lasers pump the atoms into the $5s6s\ ^3S_1$ state, from which they quickly decay back to the ground state via the $5s5p\ ^3P_1$ state, see figure 5.5. There exist two other repumping schemes that require a single laser that excite the atoms via the 5s4d or $5s5d\ ^3D_2$ states [86, 87]. However, the laser system required is far more complex and costly than the approach adopted here.

Figure 5.5: Atomic level diagram depicting the two main cooling transitions (at 461 nm and 689 nm), the clock transition (at 698 nm) and the repump transitions (at 679 nm and 707 nm).

Repump stabilisation

Both repump lasers are homebuilt external cavity diode lasers (ECDLs) with piezos for wavelength control. The laser diodes are SDL–7311 (at 679 nm) and Toptica LD–0710–0010–AR–1 (at 707 nm). To stabilise the frequency, each repump laser is coupled into a multi–mode fibre leading to a HighFinesse WS7 wavemeter. A LabVIEW frequency stabilisation program is employed to set the desired laser frequency, read the laser frequency from the wavemeter then apply a voltage correction to the piezo via a DAQ and piezo amplifier, see figure 5.6. Although the wavemeter frequency reference tends to drift throughout the day (typically 30 MHz/day), this does not have an appreciable effect on the MOT performance.

Figure 5.6: Repump locking scheme using a wavemeter as a frequency diagnostic and a LabVIEW PID controller to feed corrections back to the ECDL piezo.

Repump light control

For repump light control, both lasers are aligned through a single–pass AOM that has been aligned for optimum coupling efficiency at 80 MHz, see figure 5.7. The output of the AOM is coupled into a multi–mode fibre leading to the main experiment table. The output beam from the fibre is superimposed with the Zeeman beam via a dichroic mirror. An atom number enhancement of a factor of ≈ 5 is observed when the repump lasers are included in the sequence, and 15×10^6 atoms are routinely trapped in the blue MOT, with a density of 5×10^9 atoms/cm^3. However, a slight pit–fall is the loading rate increases for a repumped blue MOT from ≈ 50 ms to ≈ 200 ms. The temperature of the blue MOT is extracted using the time–of–flight technique and is measured to be (5.5 ± 0.5) mK. A detailed study of the blue MOT can be found in J. Millen and G. Lochead's thesis [27, 28].

Figure 5.7: Repump lasers optical set–up. After the optical isolators, a small fraction of power is picked off from the main beam using a glass plate (GP), in each laser path, to use in the wavemeter laser frequency stabilisation scheme, see figure 5.6. The 707 nm and 679 nm beams are overlapped using a non–polarising 70:30 beam splitter, and aligned through an AOM aligned to have optimium coupling efficiency at 80 MHz. The first order of the AOM is aligned down a multi–mode fibre. The output of the fibre goes to the experimental vacuum chamber.

5.2   Second stage: Red MOT

Introducing a further stage of cooling increases the experimental complexity, particularly since it is not possible to optimise the system in real time. A typical example of a red MOT experimental sequence is shown in figure 5.8. This scheme has been implemented by strontium groups at LENS [88], in Tokoyo [59], and JILA [33, 42], to name just a few. The atoms from the hot atomic beam are slowed via the Zeeman effect whereby a changing magnetic field ensures the slowing atoms are on–resonance with a fixed detuning (with the $5s^2$ $^1S_0 \rightarrow 5s5p$ $^1P_1$ transition) of 461 nm light. The atoms are cooled and trapped using 461 nm MOT light and a quadrupole magnetic field gradient of 24 G/cm in the axial direction to a temperature of several mK. Further discussion on the blue MOT can be found in section 5.1.

Figure 5.8: Red MOT control sequence including both broadband (BB) and narrow-band (NB) phases.

For red MOT loading to commence, the blue MOT light is removed and the quadrupole magnetic field gradient is reduced to ≈ 3 G/cm and the red MOT light is turned on. Low magnetic field gradients, compared to those used for the blue MOT, are adequate as a consequence of the increased magnetic field sensitivity on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ atomic transition. For effective transfer from blue MOT to red MOT, the frequency of the red MOT light is artificially broadened to match the Doppler–broadened profile of the atoms emanating from the relatively hot blue MOT (a detailed discussion is given in chapter 2). This phase is often referred to as the broadband red MOT. Once a sufficient number of atoms are transferred to the broadband red MOT, the quadrupole magnetic field gradient is increased to ≈ 10 G/cm to compress the atom cloud and, therefore, to increase the density. After the compression phase, the frequency broadening is removed, leaving single–frequency MOT light. This phase is often referred to as the narrowband red MOT, particularly in the context of this thesis. The cooling dynamics of the narrowband red MOT are discussed thoroughly in chapter 2.

It is evident from figure 5.8 that the following criteria need to be fulfilled before attempts for a red MOT are made:

- **control the laser frequency and power of the red MOT beams**: Be able to broaden the 689 nm light and be able to turn this broadening

on and off as we so wish.

- **spatially overlap the blue and red MOT beams**: The beams must intersect at the MOT field zero and have the correct polarisation for producing a MOT.

- **control the quadrupole magnetic field gradient**: Be able to ramp the magnetic field gradient with a controllable duration and final field, such that there is an increase in the number of atoms transferred from blue MOT to red MOT, as well as an increase in cloud density, is observed.

Before attempting for a red MOT, the available power for the red MOT beams must be increased since the majority of the 689 nm laser power is used in two stabilisation techniques. There are two obvious means of doing so, either using a tapered amplifier (TA) or a slave laser. The latter option is chosen since red diodes are cheap and commercially available, and TA systems are quite costly. The 689 nm slave laser system is discussed in section 5.2.1. Additionally, we must have control of the red MOT light frequency and quadrupole magnetic field gradient, which is discussed in sections 5.2.3 and 5.2.7, respectively.

5.2.1   689 nm slave laser system

An additional 689 nm laser is required as a majority of the power from the homebuilt 689 nm ECDL is used in frequency stabilisation, leaving very little for use in the red MOT beams. To increase the 689 nm laser power, a laser diode (Roitner QL68J6SA, P = 50 mW) is injection locked using light from the main laser, which is narrowed and stabilised on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition. The light is aligned through the 'wrong' port of the slave optical isolator, which can then reach the slave diode and injection lock it, see figure 5.9. The injection locked laser is called the slave laser as it takes on the properties of the main, or rather master, ECDL. The $1/e^2$ radius of the injecting master laser beam is mode–matched to that of the slave output beam for efficient injection. Fortunately, the beam radii are within 10% of one another, both in the plane perpendicular to the optical bench and parallel, requiring no need for beam shaping optics, such as anamorphic prism pairs or cylindrical lenses.

In the initial stages of injection locking, owing to a small leakage of slave light out of the wrong port of the isolator travelling towards the master laser, these

beams are coarsely overlapped using two coupling mirrors. It is helpful having a diagnostic of the slave frequency, for example picking off some slave light and coupling into a wavemeter, or alternatively coupling the light into an optical cavity, whilst scanning the master laser's frequency or the cavity length. If the slave is injection locked, cavity peaks appear. For even tighter injection locking, the slave injection current is reduced to threshold and master alignment into the slave diode is optimised, such that the lasing threshold current is reduced. This is a similiar procedure to aligning the grating of an ECDL. Once the minimum threshold current is achieved using a specific master power, the master power is reduced and the process is repeated until the slave refuses to inject, due to too little master power.

To injection lock the slave laser, a master power of 200 μW and a slave diode current of (109$\pm$1) mA is required whilst the slave laser diode is stabilised to a temperature of 50 °C. A fraction of the slave laser light is coupled into a waveme-ter and a low–finesse optical cavity to monitor the injection. The slave laser typically stays locked for several days before the slave current needs changing slightly to re–engage the injection.

Unfortunately, due to the small leakage of slave light out of the isolator and down the master input path, the slave light injects into the master diode, caus-ing the frequency of the master laser to fluctuate (this frequency fluctuation is observed on the high–finesse optical cavity transmission peaks). The slave light is attenuated using an OD 1.3 neutral density filter, which means, in order to have 200 μW of power at the slave diode input, $\approx$ 5 mW of master power is required in the input path, which is incredibly wasteful, and undesirable. A possible solution is to place an optical isolator in the injection path, though this is increasing the cost of the relatively 'cheap' method of increasing the 689 nm MOT beam power.

Figure 5.9: Optical set–up of the 689 nm slave laser. The 'Master' 689 nm light is sent through the 'wrong' port of the optical isolator to injection lock the slave diode. After the optical isolator, the slave light is separated into two main paths: the light in the first path is used on the cold atoms and is sent through an AOM in double–pass configuration for 'on/off' and 'broadband/narrowband' light control; the light in the second path is used to monitor the performance of the injection lock via a wavemeter and a low–finesse optical cavity.

The quality of the injection locking is measured via a beat measurement between the master and slave light. The frequency of the slave light is shifted by 150 MHz, with respect to that of the master laser frequency, using an AOM aligned in double–pass configuration and optimised for a frequency of 75 MHz. A non–polarising 50:50 beam splitter is used to overlap the beams, as it is crucial for the beam polarisations to be the same. The beams are coupled down a 10 m single–mode polarisation–maintaining fibre and the beam powers at the end of the fibre are set to be equal. At the output of the fibre a Thorlabs DET10A photodiode, which is terminated using a 10 kΩ resistor, detects the beat. The terminating resistor is chosen such that the bandwidth, i.e. response time, of the photodiode is greater than the frequency of the beat, which is 150 MHz. The output of the photodiode is observed on an RF spectrum analyser. Figure 5.10 shows the beat with a FWHM of $\approx$ 20–30 Hz. The narrower the beat, the smaller the frequency deviations between the two lasers. The beat could be narrower than this, however it is not possible to confirm this, owing to minimum resolution bandwidth (10 Hz) of the RF spectrum analyser. It is important to note that the reference clock of the spectrum analyser and of the AOM frequency driver are synchronised. There are small peaks at $\approx$ 200 kHz from the carrier frequency. The origin of these peaks are unknown, but since there is very little power in these peaks, they should not be a cause for concern.

Figure 5.10: Beat measurement between the master and slave 689 nm lasers. The detuning is relative to 150 MHz, which is the frequency shift between the slave and master laser. The resolution bandwidth of the RF spectrum analyser is 10 Hz.

The output of the slave laser, after the optical isolator, is aligned through an AOM in double–pass configuration at an angle that has optimum coupling efficiency at 75 MHz. The output of the AOM is coupled into a single–mode polarisation–maintaining fibre and sent to the main chamber. The AOM controls the frequency and power of the red MOT beams. The maximum power available for the MOT beams is 15 mW.

5.2.2 Red MOT optics

At the output of the 689 nm fibre the beam is expanded to a $1/e^2$ radius of 0.5 cm by means of a telescope that is mounted in an adjustable cage mount. The red and blue beams are superimposed using a dichroic mirror, which is also held in the cage mount. Hereafter, the red and blue MOT beams use the same mirrors, cubes and waveplates. The superimposed red and blue beam is split into three independent MOT beams, for the x, y, and z direction, which are retro–reflected for three dimensional confinement. The power of the blue MOT beams in the x, y and z direction are balanced, whereas for the red beams, the beam power along z (aligned with gravity) is approximately double that of the beams along x and y.

5.2.3   Laser frequency and power control

Single–frequency 689 nm MOT beams do not provide sufficient capture efficiency, hence the MOT light is frequency–broadened to match the frequency spectrum to the velocity distribution of the atoms. This is achieved by frequency modulating the slave AOM with a sinusoidal waveform, such that a comb with discrete frequencies is produced, which is set by the deviation and FM frequency. The deviation sets the optical span and the FM frequency sets the spacing between each discrete frequency. A dual channel arbitrary function generator (AFG 3000 series) modulates the AOM. Channel 1 provides the modulation to the light used in the atomic spectroscopy feedback lock to stabilise the cavity length, i.e. stabilising the 689 nm laser frequency (chapter 4), and channel 2 provides the modulation for the broadband MOT.

Figure 5.11: Beat measurement between the master 689 nm laser light and the broadened slave laser light with a deviation of 1 MHz and an FM frequency of 10 kHz. The detuning is relative to 150 MHz, which is the frequency shift between the slave and master laser. (a) illustrates the optical width of the broadened light, whereas (b) shows the sideband spacing within this width. The resolution bandwidth of the RF spectrum analyser is 100 Hz.

Figure 5.11 illustrates a beat between the modulated broadened light with that of the unmodulated light. The light is modulated by 1 MHz with an FM frequency of 10 kHz. A deviation of 1 MHz results in an optical span of 4 MHz, as shown in figure 5.11 (a). There is an enhancement of four due to the light modulating by ± 1 MHz around the centre frequency (this accounts for an enhancement of two) and owing to the fact that the AOM is aligned in double–pass configuration (this accounts for the other enhancement of two). Figure 5.11 (b)

shows a smaller frequency range than that in (a), but clearly showing the sideband spacing, which is at 10 kHz.

To switch from broadband light to narrowband light, an additional arbitrary function generator is required for the narrowband parameters, as well as an RF switch. The two arbitrary function generators used have their internal clocks synchronised. It is not necessary to use such expensive equipment to operate an AOM, but it is useful during the optimisation phase.

5.2.4   Controlling the magnetic field gradient

When transferring atoms from the blue MOT to the red MOT the quadrupole magnetic field gradient is lowered to 3 G/cm in the axial direction to increase the capture volume (refer to figure 5.8), which is then increased to $\approx$ 10 G/cm in the final loading stage. The most common experimental method of doing this is to control the magnetic field gradient via an analogue ramp. Unfortunately in the current set–up this is difficult to implement due to a short in the MOT coil requiring the control electronics to be floating. However, it is possible to switch between two MOT field gradients (a 'high' state and a 'low' state) of our choosing using the current MOT coil circuit, see figure A.3 in appendix A, by applying a TTL pulse to either the 'high' or 'low' input of the switch. The 'high' field gradient is set to 24 G/cm for the blue MOT, and the 'low' field gradient is set to 3 G/cm for the red MOT.

Concerning the magnetic field gradient ramp (as shown in figure 5.8), a detailed discussion follows in section 5.2.7 and in appendix C.

5.2.5   First siting of the red MOT

Early attempts to load into a broadband red MOT included using maximum 689 nm MOT light available ($\approx$ 12–15 mW), an optical span set to match the Doppler–broadened profile of the blue MOT atoms (i.e. 4 MHz), an FM frequency of 50 kHz (i.e. results in the minimum number of sidebands), and a quadrupole magnetic field gradient of 3 G/cm. The detuning of the light closest to resonance was –200 kHz and the repump lasers were not included in the sequence to simplify the procedure.

The first image of the broadband red MOT is shown in figure 5.12. The $1/e^2$ Gaussian MOT width along the x– and z–axis are comparable to those obtained in the blue MOT ($\approx 500\ \mu$m). From a time–of–flight measurement, the temperature of the MOT is $(44.4 \pm 0.4)\ \mu$K. With no attempts of optimising the sequence or beam overlap, the temperature of the atoms is reduced by an order of magnitude.

Figure 5.12: 'False' colour image of the first red MOT obtained in our experiment. The direction of gravity, g, along the z axis is shown. The Gaussian $1/e^2$ width along x and z are $\approx$ 500–600 μm, which is comparable to the size of the blue MOT.

On first siting, the red MOT cloud centre was displaced vastly from that of the blue MOT (by several mm) and was found by defocusing the imaging system of the camera. The compensation of stray fields inside the vacuum chamber had not been attempted, and, since atoms in the red MOT are more sensitive to magnetic fields than those in the blue MOT, a small shift in the zero of the quadrupole field resulted in a displacement of the two MOTs. Using the red MOT atoms, the stray fields inside the chamber are measured and compensated for, which is discussed in section 5.2.6.

5.2.6   Compensating stray magnetic fields

For maximal transfer efficiency from blue to red MOT, it is necessary for the MOT cloud centres to be overlapped. Any weak, stray magnetic fields within the vacuum chamber shifts the zero of the quadrupole magnetic field gradient. Since the red MOT is more sensitive to magnetic fields, the centre position of the red MOT is displaced from that of the blue MOT if stray fields are present (as

discussed in section 5.2.5). Thus it is imperative to compensate stray magnetic fields within the MOT region, which is possible via three sets of coils wound in Helmholtz arrangement whose axes are orthogonal. Each coil pair can be independently controlled: a 1 G field is produced when 1 A is supplied to the coils on the x– and y–axis, and a 2 G field on the z–axis.

For each axis shown in figure 5.4, the stray field within the chamber can be determined by measuring the position of the broadband red MOT (r_i, where i denotes the x–, y– or z–axis) via fluorescence imaging. By varying the current through the compensation coils, which effectively changes the 689 nm light de-tuning (δ) for at least three MOT field gradients ($\partial B/\partial i$) [88], the stray field B_0 can be determined by solving the equations for B_0 simultaneously via,

$$r_i = \left[\frac{h}{g_J \mu_B} \delta - B_0 \right] \left(\frac{\partial B}{\partial i} \right)^{-1}, \tag{5.5}$$

where $g_J \mu_B B/h$ is the Zeeman shift the atoms experience when a magnetic field is present. For the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition, the atoms experience a shift of 2.1 MHz/G. Using equation 5.5 the compensation field B_0 required to produce a common red MOT position r_{i0} for all $\partial B/\partial r_i$ on any given axes, can be determined.

Figure 5.13 illustrates the MOT position shift along the x–axis for different $\partial B/\partial x$. Also shown are the reduced χ^2 linear fits. Using equation 5.5, and solving the three equations simultaneously, we find $r_{x0} \approx 3.6$ mm when $I_{\text{comp}} = (-201\pm5)$ mA, which, for this coil, gives a compensation field of (-201 ± 5) mG. The stray fields along the y– and z–axis are measured using the same method. Measuring the stray field on the y–axis requires a change of camera axis, as this is the blind axis in the current set–up. A stray field of (-510 ± 20) mG along z, and a stray field of (-166 ± 5) mG along y are measured. The stray field along z is most likely due to the Earth's magnetic field, and the negative signs are an indication of the direction of the stray field in the experimental arrangement.

Figure 5.13: Compensating stray magnetic fields along the x–axis. The MOT position r_x is changed by varying the current through a pair of compensation coils I_{comp}. Implementing this for three magnetic field gradients $\partial B/\partial x$ of 1.5 G/cm (black circles), 3.1 G/cm (red squares) and 5.3 G/cm (blue diamonds) permits the stay field to be deduced from the intersection point. Also shown are the reduced χ^2 linear fits for $\partial B/\partial x$ of 1.5 G/cm (black, solid line), 3.1 G/cm (red, solid line) and 5.3 G/cm (blue, solid line), as well as the standard error for r_x.

An important point to note when compensating stray fields using the red MOT position is that this is an iterative process, requiring many cycles of optimisation. The stray fields stated here are the result of three iterative cycles.

5.2.7   Compressing the cloud

It is possible to ramp the magnetic field gradient by implementing a technique known as pulse width modulation (PWM). In this technique, the average voltage (and current) fed to the load is controlled by rapidly switching between the 'on' and 'off' state. The longer the switch is 'on', the greater the power supplied to the load. A detailed discussion is available in appendix C.

Figure 5.14: A differential measurement across the power resistor in the MOT coil circuit, see appendix A for figure, in order to measure the current through the coils, which is used to calculate the magnetic field gradient along the z–axis ($\partial B/\partial z$). Three magnetic field gradients are shown: 3 G/cm (red solid line); 10 G/cm (green solid line); and the ramp sequence shown in figure C.1 (a) (blue solid line).

Figure 5.14 is a measure of the magnetic field gradient via the method discussed in appendix C. Shown in the figure are examples of the field gradient when the coil is in the 'high' state (i.e. a field gradient of 10 G/cm) and the 'low' state (i.e. a field gradient of 3 G/cm). The ramp sequence is very similar to that in figure 5.8, which is what we expect. A slight concern is that towards the end of the ramp and during the intermediate state the magnetic field gradient seems noisy.

The best measurement of the magnetic field gradient noise is to measure the detuning fluctuations of the narrowband MOT, which is discussed in section 5.3.3. On the whole, this has been a relatively quick and easy method to implement, and it is fairly easy to change the ramp parameters. It does, however, require an expensive arbitrary function generator. In the long term it would be best to either rebuild the MOT coil circuit such that it contains an analogue optocoupler, or, instead of ramping the magnetic field gradient, ramp the broadband red MOT parameters, a method implemented by Florian Schrek's group [67].

5.3   Red MOT optimisation

There are numerous parameters that can be varied to optimise the red MOT, most of which are dependent upon one another, hence for a fully optimised system many iterative cycles are involved. The optimisation procedure is separated into three main phases: the broadband red MOT; the compression phase; and the narrowband red MOT. Furthermore, to simplify optimisation, the repump lasers are not included in the sequence.

5.3.1   Loading the broadband MOT

The main parameters that can be varied in the broadband phase are: detuning (Δ); optical span of the broadband light; the number of sidebands (N_{bands}), which depends on the optical span and FM frequency (f_m); the power in the MOT beams; and loading time (τ_{BB}). The red MOT light parameters are illustrated in figure 5.15. In the following experimental tests of broadband parameters, the magnetic field gradient during the red MOT phase is held constant at 3 G/cm (i.e. no ramp) and only the frequency modulated red MOT light is used. Fluorescence on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition is used to image the red MOT.
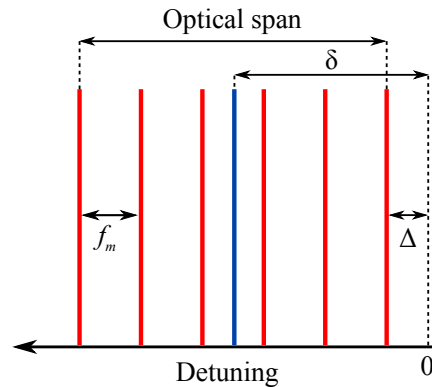
Figure 5.15: 689 nm light parameters for broadband MOT light (red lines) and narrowband MOT light (blue line). The red lines illustrate the sidebands within a given optical span separated by a modulation frequency of f_m. The detuning of the sideband closest to resonance is Δ. The detuning of the narrowband MOT light is δ.

Optical span

The 689 nm cooling light lineshape is broadened to cover the Doppler shift of the atoms from the blue MOT [89]. The best transfer efficiency, from blue MOT to broadband red MOT is obtained when an optical span of 4 MHz is used, which is expected. If the atoms in the blue MOT are at a temperature of 5 mK, this leads to a Doppler FWHM of ≈ 3.5 MHz (using equation 2.2).

Detuning

Figure 5.16 demonstrates the transfer efficiency as a function of Δ. There is a slight peak in the number of atoms transferred at ≈ 200 kHz.

Figure 5.16: Transfer efficiency as a function of broadband red MOT light detuning (Δ) when an optical span of 4 MHz and $f_m = 50$ kHz is used. The number of atoms in the blue MOT is measured to be $\approx 15 \times 10^6$. Standard error bars are included and the total power in the red MOT beams is $P_{\text{total}} = 12$ mW.

Modulation frequency

When an optical span of 4 MHz is used, there is no significant change in atom number observed when f_m is varied. However, using an optical span of 8 MHz, the effect of N_{bands} becomes more apparent, and figure 5.17 illustrates this. If there are too few sidebands, there is not enough cooling light sidebands for a considerable number of absorption/scattering events on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition to slow the relatively hot atoms emanating from the blue MOT. Conversely, if there are too many sidebands there is too little power distributed in the sidebands to make a significant effect. If an optical span of 8 MHz is

required due to a warmer initial cloud temperature, the value of N_{bands} that results in the greatest transfer efficiency is ≈ 270 (i.e. $f_m \approx 30$ kHz).

Figure 5.17: Sideband number influence on the number of atoms transferred to the broadband red MOT when the optical span is 8 MHz. Standard error bars are included and the total power in the red MOT beams is $P_{total} = 12$ mW.

MOT beam power

As the power in the MOT beams is reduced, the Gaussian $1/e^2$ MOT width (w) reduces in the z direction, figure 5.18 (a), irrespective of the detuning. However, in the x direction, there is no significant reduction in w with a decrease in MOT beam power, there is, on the other hand, a strong dependence on detuning, figure 5.18 (b). When the power–broadened linewidth (Γ_E) is greater than the detuning (i.e. $\Gamma_E > \Delta$, where Γ_E is given by equation 2.1), the radiative force produces damped harmonic motion, and the MOT thermodynamics is set by the velocity dependent force [32, 90] (i.e. standard Doppler cooling). The Gaussian $1/e^2$ MOT width in the x direction compared to that in the z direction approaches a ratio of 2:1, which is the ratio of the strength of the quadrupole magnetic field gradient. This is a result of gravitational perturbations to the potential energy becoming less pronounced [42].

Figure 5.18: Gaussian $1/e^2$ MOT width (w) dependence on total 689 nm power used in the MOT beams (P_{total}) in the (a) z and (b) x direction for five 689 nm MOT light detunings Δ: 0 kHz (black circles); –150 kHz (blue diamonds); –300 kHz (red squares); –450 kHz (green triangles); and –600 kHz (pink stars). The $1/e^2$ radius of the MOT beams is 0.5 cm. Standard error bars are also included. The optical span is 4 MHz with $f_m = 50$ kHz.

Furthermore, the transfer efficiency from blue MOT to broadband red MOT decreases linearly as the total power in the red MOT beams is reduced (figure 5.19). A global maximum is not observed, inferring an increase in power will result in an increase in transfer efficiency.

Figure 5.19: Transfer efficiency as a function of the total 689 nm power used in the MOT beams (P_{total}) when $\Delta = -150$ kHz. The optical span is 4 MHz with $f_m = 50$ kHz. The number of atoms in the blue MOT is measured to be $\approx 15 \times 10^6$. Standard error bars are also included.

To obtain a very dense cloud, very little power and a detuning very close to resonance is required. Unfortunately, using very little power results in a poor blue–to–red MOT atom transfer, and using a detuning very close to resonance limits the capture volume too (see equation 5.5). Ideally, two broadband phases would be used. The first phase would use maximum power available and a relatively large detuning, such that the capture volume matches the $1/e^2$ radius of the MOT beams, to provide the initial cooling and transfer of atoms from the blue MOT. The second phase would ramp the detuning closer to resonance, and also reduce P_{total}, in order to reduce w in both the x and z directions, respectively, see figure 5.20. A similar effect to ramping the detuning can be achieved by ramping the quadrupolar magnetic field gradient. However, the method used to broaden the 689 nm light does not permit control of the MOT beam intensity during an experimental sequence. Possible changes to the broadband MOT scheme is implementing a technique similar to that of Florian Schreck's group at IQOQI [67]. In their technique, a constant quadrupolar magnetic field gradient is maintained (at 1.15 G/cm) and, instead, the detuning and intensity of the broadband light is ramped. Although an atom number, temperature or Gaussian $1/e^2$ MOT width after the second broadband phase is not stated, a transfer of 10^7 atoms into the narrowband red MOT at a temperature of 800 nK is achieved.

Figure 5.20: Schematic diagram of a broadband sequence containing two phases, whereby the detuning (Δ), the optical span, and saturation intensity per sideband (s_{bands}) is swept. Florian Schreck's group at IQOQI implement this technique [67].

Broadband loading time

Figure 5.21 illustrates the atom number and density dependence on broadband red light loading time for $\Delta = -200$ kHz, an optical span of 4 MHz and $f_m = 50$ kHz. The final atom number and density are lower than those obtained in the optimised system, which is to be expected. The trend, however, is important irrespective of the absolute values. In figure 5.21 (a) the atoms, with a temperature of a few mK, expand outside of the observation region (because they are still comparatively hot) within a free expansion time of 2 ms, which is why there is an initial reduction in atom number. After ≈ 20 ms, the atom number starts to increase as the atoms begin to cool. After a broadband cooling time $\tau_{\mathrm{BB}} \approx 40$ ms the atom number saturates. Figure 5.21 (b) shows a steep increase in density at 20 ms, followed by a linear increase up until 80 ms when saturation is reached.

Figure 5.21: (a) Atom number dependence on broadband red light loading time, τ_{BB}. (b) Density dependence on τ_{BB}. Standard error bars are included and the total power in the red MOT beams is $P_{\mathrm{total}} = 12$ mW.

It is important to note that the greater Δ the longer it takes to enter the saturated atom number regime. This is accountable to an increasing capture volume.

Broadband phase conclusion

The best parameters to obtain the greatest transfer efficiency from blue MOT to broadband red MOT in our system, as well as typical broadband parameters used in other strontium groups, are shown in table 5.1. All groups modulate with

a sinusoidal waveform to broaden their 689 nm MOT light and the optical span varies slightly from group to group, which is most likely due to a variation of blue MOT temperature. Schreck's modulation parameters are vastly different as their aim is to load into a dipole trap and produce degenerate quantum gases of strontium (^{88}Sr, ^{87}Sr and ^{84}Sr [67]), rather than produce an ultra–cold, ultra–dense strontium gas. We use a broadband loading time of 100 ms such that we are working in the density and number saturated regime. Using these parameters, we transfer typically 50% of atoms from the blue MOT to the broadband red MOT. We could possibly improve upon the number of atoms transferred from the blue MOT to the broadband red MOT if we used more 689 nm MOT beam power.

Group	Span (MHz)	Δ (kHz)	f_m (kHz)	N_{bands}	s_{beam} (I/I_{sat})	s_{bands} (I/I_{sat})
Jones, Durham	4	–200	50	80	≈ 10000	≈ 125
Schreck, IQOQI [38, 67]	≈ 8	–250	35	≈ 229	≈ 4500–6000	20–26
Tino, LENS [88]	2	–	50	40	1600	40
Katori, Tokyo [91]	1.5	–100	50	30	900	≈ 3
Ye, JILA [33, 42, 80]	≈ 2	–	–	–	≈ 50000	–

Table 5.1: Broadband cooling parameters: Span is the optical width of the broadened light, which is set by the deviation of the modulated AOM; Δ is the detuning of the broadened light closest to resonance; N_{bands} is the number of sidebands within the optical width; s_{beam} is the saturation parameter per MOT beam; and s_{bands} is the saturation parameter per sideband.

5.3.2   Ramping the magnetic field gradient

To ramp the magnetic field gradient a technique known as pulse width modulation is employed (as discussed in section 5.2.7). The technique is tested by comparing the MOT characteristics, such as atom number and $1/e^2$ Gaussian

width w, in the presence and absence of the ramp. In the remainder of this section, the ramp duration (τ_{ramp}) and the intermediate field gradient ($\frac{\partial B_{\text{int}}}{\partial z}$ where the axial magnetic field gradient is $\frac{1}{2}\frac{\partial B_{\text{int}}}{\partial z}$), are investigated. The experimental sequence for the 689 nm MOT light and quadrupole magnetic field gradient is shown in figure 5.22. The final MOT cloud is imaged using fluorescence on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition.

Figure 5.22: Experimental sequences for investigating (a) the intermediate field gradient ($\frac{\partial B_{\text{int}}}{\partial z}$) and (b) the ramp duration (τ_{ramp}). The broadband (BB) MOT light duration is variable depending on τ_{ramp}, whereas the narrowband (NB) duration is always fixed at 50 ms. The final MOT cloud is imaged using fluorescence on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, which is not shown in the figure.

The MOT size depends on the magnetic field gradient, as represented in figure 5.23 (a) and (b). We expect the $1/e^2$ Gaussian MOT width in the z and x direction, w_z and w_x, respectively, to reduce with increasing field gradient. However, w_z increases linearly with increasing $\frac{\partial B_{\text{int}}}{\partial z}$ irrespective of detuning, suggesting there is heating due to an oscillating quadrupole magnetic field gradient as proposed (and feared) in section 5.2.7. On the other hand, w_x converges to a width of $\approx 150\ \mu$m for all detunings shown, implying the switching frequency of the pulse width modulation sets the final value of w_x. Similar ef-

fects are observed with the ramp duration, as shown in figure 5.23 (c) and (d). These findings agree with the measured magnetic field gradient in figure 5.14, though the severity of the oscillating field gradient could not be deduced from that figure alone.
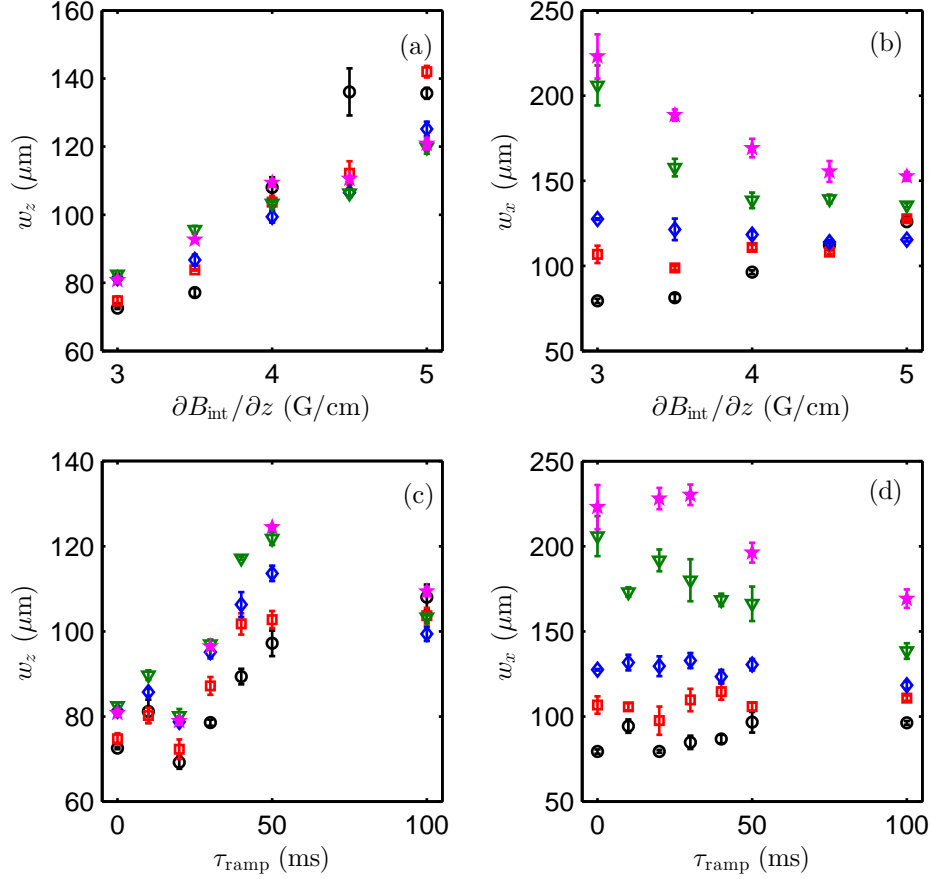
Figure 5.23: Gaussian MOT $1/e^2$ width, w, with magnetic field gradient ramp. Top: Fixed ramp duration with a variable magnetic field gradient ($\frac{\partial B_{\text{int}}}{\partial z}$) on the width in the (a) z direction and (b) the x direction for five narrowband detunings. Bottom: Fixed magnetic field gradient with a variable ramp duration τ_{ramp} on the width in (c) z direction and (d) the x direction for five narrowband detunings. The narrowband detunings used are $\delta = -100$ kHz (black circles); -150 kHz (red squares); -200 kHz (blue diamonds); -300 kHz (green triangles); -400 kHz (pink stars). The total power in the narrowband MOT beams is ≈ 180 μW. Standard error bars are also included.

Ramp parameters conclusion

A technique known as pulse width modulation (PWM) is implemented to ramp the quadrupole magnetic field gradient. Including the ramp, increases the $1/e^2$

Gaussian MOT width along the z–axis (w_z) with increasing magnetic field gradient and ramp duration, irrespective of laser detuning. Whereas, the $1/e^2$ Gaussian MOT width along the x–axis (w_x) converges to a fixed value of w_z. These features are characteristic of an oscillating magnetic field gradient, hence, unfortunately, this method of ramping the magnetic field gradient will not suffice. In the remainder of the thesis, the magnetic field gradient during the red MOT sequence will be held at 3 G/cm.

5.3.3 Loading the narrowband red MOT

The temperature and atom number of the final atom sample depend on the detuning and intensity of the single–frequency MOT light, which is evident in figure 5.24. For large detunings, the temperature is fairly independent of the detuning since gravity is comparable to the scattering force along the z direction, resulting in a constant photon scattering rate. We find the temperature be to lower than the semiclassical limit, which is in agreement with theoretical studies discussed in [41]. For detunings approaching resonance, the temperature begins to increase, which is a typical behaviour of standard Doppler cooling [33].

At large intensity, atoms are in resonance with the light on a single shell whose thickness is governed by the power–broadened linewidth Γ_E (equation 2.1). The scattering rate is high enough to keep the atoms in the MOT and the temperature is in agreement with Doppler theory [80]. As the intensity is reduced, the temperature approaches the recoil temperature T_R, and at low saturation, the recoil temperature is reached. Using even smaller saturation might permit the a temperature of half the recoil temperature [42].

Figure 5.24: Detuning (δ) dependence on (a) MOT temperature and (b) atom number after a free expansion time of 15 ms for $P_{\text{total}} = 18$ μW (red squares); 190 μW (black circles); 1.56 mW (blue diamonds). Also shown in (a) is the recoil temperature (black, dashed–dot line).

Reducing the intensity, not only reduces the temperature but also the atom number. The atoms lie of the bottom on an ellipsoid shell (chapter 2) and predominantly interact with the beam on the z–axis. At low intensity, the atoms are at risk of falling through the shell if an absorption/scattering event does not occur. Furthermore, reducing the detuning also reduces the atom number. In a simplified picture, decreasing the detuning compresses the cloud, whereas decreasing the intensity cools the cloud. A very dense sample of the atoms is required for Rydberg blockade experiments (as discussed in chapter 1), thus a compromise between atom number, temperature and MOT size is made. Working fairly close to resonance (i.e. 50–100 kHz) and using an intensity that results in a sufficient transfer of atoms from the blue MOT into the red MOT ($\approx 25\%$), densities in 10^{11} atoms/cm^3 are achieved. These are similar results achieved by other groups (refer to table 5.2).

5.3.4   Final MOT sequence and remarks

The experimental MOT sequence we use is shown in figure 5.25. We find we transfer more atoms from the blue MOT to the single–frequency red MOT when both blue and red broadband light MOT beams are on during blue MOT loading. We find the red MOT is not particularly sensitive to polarisation, but it is, however, incredibly sensitive to beam alignment in the single–frequency

Group	δ (–kHz)	s (I/I_{sat})	T (μK)	Density (atoms/cm^3)	Transfer (%)
Jones, Durham	50–100	16–34	≈ 2	$\approx 10^{11}$	≈ 25
Schreck, IQOQI [38, 67]	800	30	2.5	–	≈ 25
Tino, LENS [78, 88]	350	≈ 23	≈ 1	1.4×10^{11}	≈ 20
Katori, Tokyo [91]	100	10	≈ 1	–	–
Ye, JILA [33, 42, 80]	100	75	≈ 1	$\approx 5 \times 10^{11}$	≈ 30

Table 5.2: Narrowband cooling parameters comparison between the strontium groups in Durham, IQOQI, LENS, Tokyo and JILA. Parameters compared are: narrowband 689 nm MOT light detuning, δ; saturation parameter, s; narrowband temperature, T; density; and transfer efficiency from blue MOT to narrowband red MOT.

MOT, particularly at low saturation.
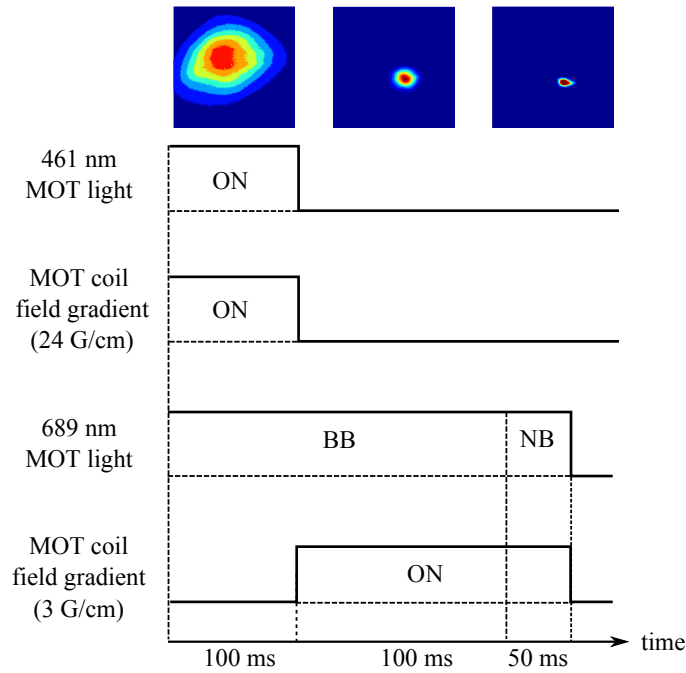
Figure 5.25: Red MOT control sequence including both broadband (BB) and narrowband (NB) phases, with their corresponding MOT fluorescence image.

During the optimisation of the red MOT, we did not include the repump lasers in the sequence. The final atom numbers and densities quoted in table 5.2 do not include the repump lasers hence we expect these to increase slightly, though we have yet to confirm this.

Conclusion

Starting from a 900 K atomic beam, we are able to cool $\approx$ 15 x 10^6 atoms on the strong $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition to 5 mK in a first stage of cooling, and transfer $\approx$ 25 % of atoms to a second stage of cooling on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition to a typical atom temperature of $\approx$ 1–2 μK, within 250 ms. The minimum atom temperature we have achieved ($\approx$ 460 nK) is very near to the photon recoil temperature limit and is indicative of a fairly narrow instantaneous 689 nm laser linewidth. With the parameters we typically use, we can obtain densities on the order of 10^{11} atoms/cm^3, which is two orders of magnitude greater than the density we typically obtain with the first cooling stage. We

have shown the cooling dynamics model, discussed in chapter 2, agrees with experimental findings. Although we were unable to successfully implement a magnetic field gradient ramp (via PWM), we are able to produce temperatures, densities and transfer efficiencies similar to those in other strontium groups. To conclude, this is one of, at most, ten strontium red MOTs world wide.

Chapter 6

Dipole blockade of divalent atoms

As outlined in chapter 1, the dipole blockade is a phenomenon that arises when the dipole–dipole interactions between Rydberg atoms induce an energy shift of the collective excited state. This causes the driving laser radiation to become off–resonant and access to the collective excited state is forbidden. A signature of dipole blockade is a suppression of further excitation in a volume around the first Rydberg atom. To test for dipole blockade the average Rydberg signal as a function of ground state density is measured (chapter 1). If there are interactions present, the Rydberg signal will saturate with ground state density [11] as more atoms are within the volume governed by the blockade radius, which cannot be excited. The first attempts of dipole blockade were performed in the blue MOT during the thesis of G. Lochead [28], where typical ground state densities are ≈ 5 x 10^9 atoms/cm^3. The Rydberg signal for 5s75d ^{1}D$_2$ state follows a linear trend with ground state density, indicating a non–interacting Rydberg cloud. The blockade density for this state, which is calculated from the angular average of C_6 and the excitation linewidth (≈ 8 MHz), is 6 x 10^9 atoms/cm^3. To observe blockade, a cold atomic ensemble of at least an order of magnitude greater in density is required. Such densities are achieved via cooling on the 5s^2 ^{1}S$_0$ $\rightarrow$ 5s5p ^{3}P$_1$ transition, hence it should be possible for us to observe blockade in a divalent atom ensemble for the first time.

6.1 Rydberg creation and detection

To access the Rydberg states, a two–photon, three–level ladder type scheme is used. The wavelengths of the lasers used and the states they access are shown in figure 6.1.

Figure 6.1: Two–photon, three–ladder Rydberg excitation scheme. The states and approximate wavelengths required for excitation in our experiment. Also shown are the probe detuning (Δ_p) and coupling detuning (Δ_c).

Here, the 461 nm laser beam is referred to as the probe beam, and the 413 nm laser beam is referred to as the coupling beam with detunings Δ_p and Δ_c, repectively.

6.1.1 Rydberg laser system

The Rydberg, or coupling, laser light is a frequency–doubled diode laser source (Toptica TA–SHG 110). The beam layout is very simple and is shown in figure 6.2. The output power at 413 nm is typically 60 mW, where $\approx$ 50 mW is used in spectroscopy for stabilising the laser frequency whilst the remaining $\approx$ 10 mW is available for use on the cold atoms in Rydberg experiments.
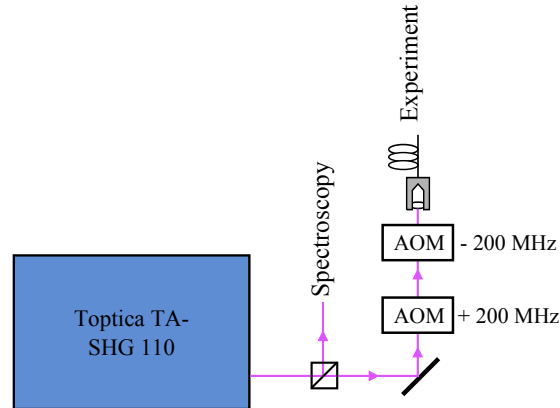
Figure 6.2: Schematic set–up of the 413 nm Rydberg laser system. Typical output power is $\approx$ 60 mW, with $\approx$ 50 mW used in spectroscopy for stabilising the laser frequency and $\approx$ 10 mW is available for use in the Rydberg experiments on the cold atoms. The frequencies of these beams are set using acousto–optic modulators (AOMs) either aligned in double–pass or single–pass configuration.

Electromagnetically–induced transparency (EIT) spectroscopy using a strontium vapour cell is employed to stabilise the frequency of the Rydberg coupling laser [92]. EIT is a process whereby the probe transition becomes transparent due to the coupling to another state [93]. The transparency arises in the weak probe regime and, if both probe and coupling laser beams are phase coherent, features much narrower than the intermediate state linewidth can be measured, with the lifetime of the Rydberg state being the limiting linewidth [94].
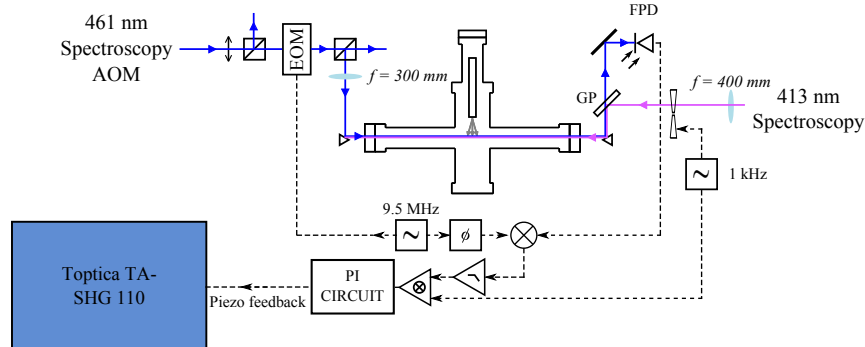
Figure 6.3: 413 nm laser frequency stabilisation by means of electromagnetically–induced transparency (EIT) spectroscopy using a strontium vapour cell. This scheme employs a weak 461 nm probe, which is modulated using an electro–optic modulator (EOM), and a counter–propagating strong coupling beam, which is amplitude modulated using an optical chopper at ≈ 1 kHz, that are aligned using a glass plate (GP). The EIT signal is detected, via the probe light, on a fast photodiode (FPD). A phase shifter sets the phase between the local oscillator input and the RF input of the double–balanced $+7$ dBm mixer. The output of the mixer is fed into a lock–in amplifier, which is referenced to the optical chopper driving frequency, and the demodulated signal is fed back to the laser piezo via the PI circuit. The electronic and optical paths are illustrated using black, dashed lines and blue/purple, solid lines respectively.

The beam layout is shown in figure 6.3. The 461 nm laser frequency is stabilised using a spectroscopy technique known as modulation transfer spectroscopy (MTS), which is discussed in chapter 5, in a strontium vapour cell [81]. Once the 461 nm laser frequency is stabilised, the 413 nm laser frequency is stabilised to the on–resonance EIT feature using the same strontium vapour cell, which is detected as a peak in transmission when monitoring the probe light using a fast photodiode. This stabilisation technique was developed during the thesis of G. Lochead [28] and it was found that the more 413 nm laser power used, the more robust the lock, which is why the majority of the laser power is directed down this path. To increase the coupling Rabi frequency further, the beam is focused to a $1/e^2$ beam radius of ≈ 180 μm, with the focus at the centre of the atomic beam. The probe beam is also focused and has a $1/e^2$ beam radius of ≈ 110 μm. Irrespective of the probe and coupling Rabi frequencies used in the EIT spectroscopy, the signal is still too small to obtain a robust feedback loop. To alleviate this problem, an optical chopper (Scitec Instruments 300H) amplitude modulates the signal at a frequency of ≈ 1 kHz and a lock–in

amplifier (Standford Research SR510) demodulates the signal to improve the signal–to–noise ratio. The lock–in amplifier filters are set to $\tau_{\mathrm{pre}} = 100$ ms and $\tau_{\mathrm{post}} = \mathrm{NONE}$, which limits the bandwidth of the loop filter to a maximum of 10 Hz. We can easily lock the 413 nm to any value of n up to 75, and the laser typically remains locked for several hours. The laser frequency drifts by no more than 1 MHz/hour, which is not crucial for the experiments in this thesis.

The remaining laser power is directed through two acousto–optic modulators (AOMs) aligned in single–pass configuration, such that there is a net frequency change of zero, for 'on/off' control. This permits on–resonance coherent population trapping (CPT) in the cold atoms. CPT is a similar phenomenon to EIT [95, 96], however in this scheme the probe Rabi frequency is no longer vanishingly weak, such that there is population in the Rydberg state but none in the intermediate state.

In the cold atom experiments, an ion signal proportional in magnitude to the population in the Rydberg state is created via a technique known as 'autoionisation'.

6.1.2 Autoionisation

To study Rydberg atom dynamics, most experiments using alkali metal atoms use an electric field pulse to field–ionise the Rydberg atoms [3], which are detected using a microchannel plate (MCP). Although field ionisation in strontium has also been implemented [97], in our experiment we chose to use autoionisation. Once one electron is excited to the Rydberg state, the other lower–lying valence electron is available to undergo optical excitation. The process of exciting the inner core electron is known as 'isolated core excitation' (ICE) [98]. The level scheme for ICE is shown in figure 6.4.

Figure 6.4: Isolated core excitation level scheme for autoionisation, including a visual representation of the atomic orbitals.

The autoionising laser (λ_3) is tuned near the Sr^+ $5s_{1/2} \rightarrow 5p_{3/2}$ transition at 408 nm, and is supplied by a homebuilt ECDL using a wavelength selected diode (Sanyo DL–7146–101S). The laser frequency is set to give the maximum probability of autoionisation, see G. Lochead's thesis for more information regarding this [28]. The advantage of using this scheme is the atom ionises very fast (≈ 10 ps) and with very high probability, permitting temporal information of the population dynamics to be obtained [21]. Furthermore, the autoionisation signal is much greater than that obtained from spontaneous ionisation alone [28].

6.1.3   Coherent excitation to Rydberg states

The beam geometry for coherent excitation to Rydberg states is shown in figure 6.5. All three excitation lasers, as shown in figure 6.5, are coupled into single–mode polarisation–maintaining fibres. The fibre outputs of λ_2 and λ_3 are attached to lens systems (designed during the thesis of G. Lochead [28]) that focus the beams to ≈ 14 μm and (6.3 ± 0.1) μm, respectively. The beams propagate orthogonally with respect to one another and the foci intersect at the centre of the single–frequency red MOT. The fibre output of λ_1 is collimated using a fibre collimator, with a $1/e^2$ beam radius of (1.07 ± 0.02) mm in both directions, and is aligned counter–propagating to λ_2 to reduce Doppler broadening of the CPT spectrum [68].

Figure 6.5: Beam geometry for measuring Rydberg blockade, (a) illustrates the top view and (b) the side view. The probe and coupling lasers are counter–propagating along the quantisation field axis (z–axis). The autoionisation laser enters the vacuum chamber from an orthogonal axis (y–axis) to the excitation beams. The figure is modified from that shown in [28].

6.1.4   Timing sequence

The timing sequence for the creation and detection of Rydberg atoms in a single–frequency red MOT is shown in figure 6.6. The single–frequency red MOT is loaded (see figure 5.25, chapter 5) and the quadrupole magnetic field gradient is removed, whilst the magnetic quantisation field is introduced, further discussion of the importance of this follows in section 6.1.5. The excitation beams are pulsed on simultaneously for 1 μs, and 100 ns after this pulse, the autoionising laser is pulsed for 2 μs to ionise the Rydberg atoms.

Figure 6.6: Dipole blockade experimental timing sequence that commences after the creation of a single–frequency red MOT (the MOT experimental sequence for this is shown in figure 5.25). λ_1, λ_2, and λ_3 are the probe, coupling, and autoionising laser pulses, respectively. The magnetic quantisation field Zeeman splits the magnetic sublevels, whilst the electric field pulse directs the ions to the MCP for detection.

Once the atoms have undergone ICE (section 6.1.2), the created ions are collected on an MCP (Hamamatsu Compact MCP Assembly F4655). The MCP cathode is held at –2 kV and the anode is held at 100 V. Electrons, or ions, impacting the MCP create an avalanche of electrons, resulting in a voltage pulse at the anode. The cathode is covered by a metal mesh, referred to as the 'grid' (Goodfellow FE228710), to shield the atoms from the cathode's large electric field. The grid is held at a constant, negative voltage to attract the ions to the MCP. To direct the ions from the MOT towards the MCP, a small electric field pulse is applied. The electric field pulse is created by a pair of split ring electrodes attached to the MOT coil formers, based on the design in [99]. These are shown in figure 6.5 (a). The electrodes closest to the MCP are grounded and the pair furthest away from the MCP have a positive voltage applied, creating an electric field towards the MCP. A voltage of –35 V is supplied to the grid and 10 V to the electrodes, which gives an electric field strength of ≈ 3 V/cm when a pulse of 5 μs is applied, as this is found to give the maximum ion signal [28]. The ion signal is amplified and signal ion events are counted using an oscilloscope, as shown in figure 6.7. In (a), the MCP signal is shown from the start of the excitation sequence, which corresponds to a time of –47 μs in the figure. Immediately after the excitation sequence commences, there are two pulses. These are due to the applied electric field pulse, the separation of which depends on the duration of the pulse. The effect of the ion signal in the presence

and absence of the autoionisation laser is also shown. The MCP detects a much greater and more frequent ion signal in the presence of the autoionisation laser. To count the number of ions incident on the MCP, an oscilloscope is used as a discriminator and counter with the number of pulses below a certain reference value recorded per experimental shot, which is illustrated in figure 6.7 (b). Out of the five troughs shown, two are counted by the oscilloscope, whilst three are missed. The reference level, also shown in the figure, is set such that electronic ringing, which we believe is due to impedance mismatching, is not counted. In this example, the reference level is set to –98 mV. To avoid missing counts, the time base setting of the oscilloscope is set to 1 μs.

Figure 6.7: Example of ion signals used for counting. (a) MCP signal from the start of the excitation sequence in the presence (red solid line) and absence (black, solid line) of the autoionisation laser. The two pulses immediately after the start of the sequence are a consequence of the electric field pulse, and are separated by 5 μs (i.e. the duration of the electric field pulse). As a consequence of time–of–flight, there is a time delay of the Rydberg ion signal. (b) Same as in (a) over a smaller time frame. The black, dash–dotted line shows the reference level used in experiments for counting ions.

To extract information regarding the MOT atom number and density, the MOT is imaged via fluorescence on the $5s^2\ ^1S_0 \to 5s5p\ ^1P_1$ transition for a variable exposure time.

6.1.5 State selection

To selectively excite specific m_J states a defined quantisation field is required [40]. To create a quantisation field, a magnetic field along the axis of the excitation beams is used. The position of the coil is shown in figure 6.5 (b). The magnetic field at the position of the atoms is (3.0 ± 0.2) G. As a consequence of the direction of the magnetic field with respect to the excitation beams, only σ^+ or σ^- transitions (or both) can be driven, hence only three of the five magnetic sublevels of the $^1D_2$ state can be accessed ($m_J = 0, \pm 2$) by choosing the circular polarisation states of the excitation lasers, as shown in figure 6.8 (a).

Figure 6.8: (a) Zeeman splitting ($g_J \mu_B m_J B$) of the m_J states for a D state when a magnetic field of B is applied. The excitation routes permitted by circularly polarised light illustrate the accessible m_J states in the D state. Shown bottom right are the excitation probabilities for the accessible routes. (b) The CPT spectra showing the average ion count with probe detuning Δ_p, as illustrated in figure 6.1, when probe and coupling light are driving the $\sigma^- \sigma^+$ (black circles) and $\sigma^- \sigma^-$ (red squares) transitions. Theoretical Gaussian fits are also shown for the $\sigma^- \sigma^+$ (black, solid line) and the $\sigma^- \sigma^-$ (red, solid line) transitions.

Figure 6.8 (b) compares the $\sigma^- \sigma^-$ spectra of the stretched state, with the spectra of the $\sigma^- \sigma^+$ state in a single–frequency red MOT. The fitted centre–to–centre splitting is ≈ 20 MHz, which is expected from the theoretical splitting calculated for the applied field. The $\sigma^- \sigma^-$ feature height is less than the $\sigma^- \sigma^+$ feature, as is expected since the stretched states $m_J = \pm 2$ have the strongest excitation probability. By applying a magnetic quantisation field, the excitation polarisation is chosen such that the maximum ion signal is obtained (i.e. either $\sigma^- \sigma^-$ or $\sigma^+ \sigma^+$).

6.2 Dipole blockade experiment

To reproduce figure 1.3, the average Rydberg signal is measured as a function of the ground state density in the single–frequency red MOT, which is controlled via the single–frequency red MOT loading time. The coupling laser (λ_2) is on–resonance with the 5s5p $^1P_1 \rightarrow$ 5snd $^1D_2$ transition and 30 averages are taken at each density. The probe beam (λ_1) frequency is fixed on–resonance with the $\sigma^- \sigma^-$ transition, as shown in figure 6.8 (b).

Figure 6.9: Average ion count with ground state density for Rydberg states 5s40d $^1D_2$ (black circles) and 5s56d $^1D_2$ (red squares). Also shown is a linear least squares fit to 5s40d $^1D_2$ (black, solid line) and a saturation function fit ($y(x)=a/(1+b/x)$, where a and b are free parameters and x is the ground state density) to 5s56d $^1D_2$ (red solid line). Standard error bars are included.

Figure 6.9 illustrates the dependence of the average ion count with ground state density for Rydberg states 5s40d $^1D_2$ and 5s56d $^1D_2$. The probe power during 5s56d $^1D_2$ state excitation is set to maximum ($\approx 500~\mu$W), such that the greatest possible population is excited to the intermediate state. The coupling power is set to give the maximum ion count before the oscilloscope saturates (the maximum number of ions the oscilloscope can count is 250), which is $\approx 410~\mu$W. The average ion count follows a saturation–type trend of the form $y(x)=a/(1+b/x)$ (where a and b are free parameters and x is the ground state density) indicative of dipole blockade as more atoms are within the blockade radius (defined by equation 1.1), which cannot be excited. This is a phenomenological fit that takes into account the Gaussian ground state density distribution in the MOT and the finite linewidth of the excitation lasers, hence we do not expect a hard–

edge turn on of the dipole blockade, but rather a smooth, gradual transition. The fitted saturation density is (3.9 ± 0.3) x 10^{10} atoms/cm^3 with a corresponding $\chi^2_{\min}$ fit of ≈ 12, which is in agreement with the calculated dipole blockade density of ≈ 4 x 10^{10} atoms/cm^3. Fitting a linear function yields $\chi^2_{\min} > 10^{10}$ further confirming dipole blockade has been observed at 5s56d ^{1}D$_2$.

Whereas for the 5s40d ^{1}D$_2$ state, the probe and coupling powers for the 5s40d ^{1}D$_2$ state are ≈ 500 μW and ≈ 320 μW, respectively. The average ion count for the 5s40d ^{1}D$_2$ state follows a linear trend with ground state density with a least squares line–of–best–fit of $y = (9\pm0.1)$x10^{-10} $x - (24\pm6)$ and $\chi^2_{\min}$ of ≈ 0.4, demonstrating a non–interacting Rydberg ensemble. This is to be expected as the dipole blockade density is ≈ 3 x 10^{11} atoms/cm^3. Furthermore, there are no ions present when there is an appreciable MOT density (i.e. 1 x 10^{10} atoms/cm^3) suggesting there is an experimental procedure error.

During this experiment, we attempted to scale the coupling power such that the Rabi frequencies for the transitions to the 5s40d ^{1}D$_2$ state and 5s56d ^{1}D$_2$ state were the same. Ideally, this yields the same slope at low density in figure 6.9. However, we suspect the reduced slope and the offset of the 5s40d ^{1}D$_2$ curve are due to a small detuning offset. When changing n, the shape of the EIT error signal used to stabilise the 413 nm laser frequency changes, leading to an n–dependent offset in the laser frequency. While we checked that we were on–resonance with the 5s56d ^{1}D$_2$ state transition in the cold atoms, we did not do this for the 5s40d ^{1}D$_2$ state transition.

Background contribution from spontaneous ions

The advantage of using autoionisation is that we can clearly separate the Rydberg ion signal from the background signal. Figure 6.10 illustrates the dependence of average ion count with density for the 5s56d ^{1}D$_2$. Shown in (a) is a comparison of the ion count due to autoionisation and those due to spontaneous ions. It was shown in reference [100] that resonant dipole–dipole interactions lead to significant motion of the atoms, which is responsible for initiating the spontaneous evolution of the cold ensemble into a plasma. For instance, a 5 MHz shift, which corresponds to an electric field of ≈ 0.1 V/cm (as measured in G. Lochead's thesis [28]), occurs at an inter–particle separation of

$r \approx 12 \ \mu$m (i.e. $E = \frac{e}{(4\pi\epsilon)r^2}$, where E is the electric field, e is the electron charge and ϵ is the permittivity of free space). This is equates to a density of $\approx 0.5 \times 10^9$ atoms/cm^3, which is much less than the blockade density for this state, therefore the saturating feature we observe at $(3.9\pm0.3) \times 10^{10}$ atoms/cm^3 is likely to be caused by dipole blockade. Furthermore, since there is minimal background contribution from the spontaneous ions, particularly at a density of $\approx 0.5 \times 10^9$ atoms/cm^3, the excitation saturation is not due to the Stark shift caused by the ions. At high densities, the data becomes more scattered, which is likely due to the inaccuracy of the density measurement.

Figure 6.10: Average ion count of the 5s56d ^{1}D$_2$ as a function of ground state density. In (a) the ion count due to autoionisation (black circles) and spontaneous ionisation (red squares) are shown. (b) The resulting ion count when spontaneous ions are subtracted. The red, solid line is a fitted saturation function. Standard error bars for ion count and density are included.

Shown in figure 6.10 (b) is the ion count difference between those due to autoionisation and spontaneous ionisation. The saturation feature is more pronounced than it is in (a).

Limitations on the quantitative analysis

Using the Beer–Lambert law [40] as a back–of–the–envelope calculation, when the single–frequency red MOT has a width of $\approx 50 \ \mu$m and a density of $\approx 10^{11}$ atoms/cm^3, there will be ≈ 10–20% of photons from the incident imaging beams (i.e. the MOT beams) scattered out of the cloud and, at most, detected. Since the number of photons (N_γ) is lower than expected, extracting an atom

number (via equations 5.1 and 5.2) results in an underestimated density when calculated using equation 5.4, i.e. the cloud is optically thick. This is confirmed by imaging the cloud after a free expansion time of ≈ 10 ms and comparing the atom numbers extracted from this measurement with those when the MOT is imaged immediately after release. Furthermore, the relation between measured atom number and actual atom number becomes density–dependent, which makes correcting the density axis of figure 6.9 very difficult. Plotting the average ion count against true density will stretch the saturation curve, making the dipole blockade phenomenon more pronounced in the data.

This inability to measure ground state MOT densities via fluorescence further supports a change in Rydberg excitation scheme to that via the 5s5p ^{3}P$_1$ line, whereby a probe 689 nm laser, including an AOM for frequency and 'on'/'off' control, will be implemented as well as a coupling laser beam at 319 nm (further discussion in chapter 1). Imaging the MOT via fluorescence using far–detuned 689 nm probe light (10Γ, where Γ is the atomic transition linewidth) will result in nearly all of the scattered photons escaping the dense atomic cloud, available for detection. Alternatively, the cloud can be imaged using absorption of the 461 nm probe beam.

Conclusion

We have introduced the reader to the experimental apparatus and timing sequence required for producing Rydberg atoms in a two–electron atom ensemble. We have shown that by including the second stage of cooling (i.e. the red MOT), which increases (decreases) the density (temperature) of the ensemble, dipole blockade in the 5s56d ^{1}D$_2$ state is observed via a saturation of the Rydberg signal with increasing density. In addition, the dipole blockade density extracted from the saturation fit to the data agrees with theoretical calculations. However, theoretical calculations and experimental findings reveal the atomic cloud is optically thick and using fluorescence imaging on the 5s^2 ^{1}S$_0$ $\rightarrow$ 5s5p ^{1}P$_1$ transition yields underestimated atom numbers and densities as we approach densities of 10^{11} atoms/cm^3. By changing the Rydberg excitation scheme to that via the 5s5p ^{3}P$_1$ line, whereby a probe 689 nm laser and a coupling 319 nm laser will be used, the cloud can be diagnosed via absorption imaging of the 461 nm probe beam.

Chapter 7

Discussion and Outlook

In this thesis we have demonstrated the first known measurements of Rydberg dipole blockade in a cold, dilute gas of strontium atoms. To facilitate these measurements a reliable, flexible source of strontium atom is used. A magento-optical trap (MOT) on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition loaded from a Zeeman slowed atomic beam creates a confined gas of 15 x 10^6 strontium atoms at a typical density of 5 x 10^9 atoms/cm^3 and a temperature of 5 mK, when repumped out of the metastable $5s5p\ ^3P_2$ state. To obtain sufficiently high densities to observe dipole blockade, a second stage of cooling on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition is implemented and creates a confined gas of 4 x 10^6 strontium atoms at typical densities on the order of 10^{11} atoms/cm^3 and a temperature of (1–2) μK. The $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ MOT dynamics are simulated theoretically and agree with experimental findings.

For creation of the second stage MOT (i.e. red MOT), a laser system on the narrow (2π x 7.4 kHz wide) line is required. The laser is a homebuilt external cavity diode laser (ECDL) operating at 689 nm. The short–term frequency fluctuations are stabilised to a high–finesse optical cavity. Using basic cavity theory, the high–finesse optical cavity is designed to operate at 689 nm and 638 nm, have a finesse > 10,000, and have resolvable carrier sidebands. Cavity performance measurements agree with the specified requirements.

Frequency stabilisation of the 689 nm laser to the cavity implements the Pound–Drever–Hall (PDH) phase sensitive detection technique [49, 50]. A high–bandwidth feedback electronic circuit, based on the design in [60], is constructed

and optimised to reduce laser frequency instabilities. In the absence of two identical laser–cavity systems to perform beat measurements necessary for evaluating laser frequency instabilities, a method for determining the laser performance on short–interrogation times is implemented by means of measuring the temperature of the cold atomic ensemble. The minimum single–frequency red MOT temperature achieved is ≈ 460 nK, indicating very small laser frequency instabilities on short–interrogation times.

The long–term frequency stability of the laser system in question is minimised by stabilising the length of the high–finesse optical cavity. This is achieved via active stabilisation to an atomic transition ($5s^2$ $^1S_0 \rightarrow 5s5p$ $^3P_1$) provided by low–bandwidth electronic feedback to a piezo, which is included in the cavity length during construction. Due to experimental constraints, saturated fluorescence spectroscopy [69] is used, whereby a single, frequency modulated beam is retro–reflected through an atomic beam. The fluorescence is detected and demodulated using a lock–in amplifier to provide the error signal for cavity length stabilisation. The width of the Lamb dip that originates from this technique is shown to have a non–linear dependence on the saturation and modulation parameters of the laser beam interrogating the atoms. The narrowest feature measured is ≈ 50 kHz. Again, without the aid of a stable frequency reference, such as a frequency comb stabilised to a hydrogen maser, a method for determining the long–term laser frequency instability is improvised. Using the $1/e^2$ Gaussian MOT widths of the single–frequency red MOT and implementing an Allan deviation measurement, the stability over a 15 minute period is measured to be ± 1 kHz with no evidence of laser frequency drift.

To excite ground state atoms to high lying Rydberg states a resonant, two–photon, three–level ladder scheme is used as this allows use of visible wavelength lasers. The first photon is resonant with the $5s^2$ $^1S_0 \rightarrow 5s5p$ $^1P_1$ transition at 461 nm, whilst the second photon is resonant with the $5s5p$ $^1P_1 \rightarrow 5snd$ $^1D_2$ transition at 413 nm. The 461 nm laser is frequency stabilised to an atomic source using modulation transfer spectroscopy [82], whereas the 413 nm laser is frequency stabilised to an atomic source using electromagnetically–induced transparency (EIT) [92]. Both frequency stabilisation schemes occur simultaneously in the same, novel dispenser–based vapour cell [81] and both lasers remain locked for several hours.

The dipole blockade phenomenon is demonstrated in the single–frequency red MOT via coherent Rydberg excitation in ensembles of differing ground state densities. Although the cloud is optically thick, and an accurate density measurement was not made, the excitation saturation is due to the dipole blockade effect, and not due to the Stark shift caused by the ions, as shown by the spontaneous ion background count. By changing the imaging scheme such that absolute density measurements can be extracted, the excitation saturation curve can be then modeled theoretically.

Outlook

Such a blockade result can be extended to perform dynamical crystallisation in the atomic sample, whereby the dipole blockade is manipulated to produce a well–controlled, self–ordered crystalline structure via tailored light pulses, facilitating controlled quantum state preparation in large ensembles [13, 101]. Additionally, the 689 nm laser system and second stage MOT developed in this thesis can be extended to look at weakly–dressed Rydberg states via the $5s^2 \, ^1S_0 \rightarrow 5s5p \, ^3P_1$ line. Owing to the very narrow transition (2π x 7.4 kHz), the experimental limitations experienced in the alkali atom Rydberg experiments (such as a Rydberg atom fraction in the cold atom ensemble that is much less than one) are irrelevant to alkaline–earth atom experiments, since large detunings from the intermediate state can easily be achieved. The 319 nm laser system, required for coupling the intermediate state to the Rydberg state, is currently being developed and stabilised to the high–finesse optical cavity developed in this thesis.

Appendix A

Circuit diagrams

A.1   High–bandwidth feedback circuit

The final component values that obtain a bandwidth of 1.5 MHz for the TEM_{00} with a power of ≈ 13 μW incident on the FPD are shown in table A.1. The roll–off frequency of the integrator and differential stage for the first op amp is calculated as $f_i = \frac{1}{2\pi R_5 C_4} \approx 2.4$ MHz and $f_d = \frac{1}{2\pi R_1 C_3} \approx 440$ kHz, respectively. This is a little confusing as it disagrees with the discussion in 3.4.3. However, calculating the transfer function using these parameters, see figure 3.12, gives the necessary shape with the –20 dB/decade to –40 dB/decade slope transition at ≈ 440 kHz and the –40 dB/decade to –20 dB/decade slope transition at ≈ 2.4 MHz.

Figure A.1: High–bandwidth feedback circuit used for stabilising the laser frequency to the high–finesse optical cavity.

Component	Value	Units
R_1	2.4	kΩ
R_5	8.0	kΩ
R_{23}	135	Ω
R_{25}	14.8	kΩ
R_{26}	2	kΩ
C_3	150	pF
C_4	8.2	pF
C_5	2.2	μF
C_7	15	pF
C_8	4.7	pF
P_1	Variable depending on DC offset	
P_6	Variable depending on TEM$_{mn}$ mode	

Table A.1: Final component values used in the optimised loop filter circuit.

A.2 Low–bandwidth feedback circuit

The component values are shown in table A.2. The resistors R_{in} and R_p are variable to enable real–time optimisation of the integrator time constant and proportional gain. Using the values stated, the integrator time constant can range from anything between $\approx$ 145 ms to $\approx$ 7 s, and the proportional gain can be anything up to 100.

Component	Value	Units	Comment
R_{in}	50	kΩ	Variable trimpot
R_1	10	kΩ	Fixed resistor
R_p	100	kΩ	Variable trimpot
C_i	22	μF	Fixed capacitor
S_a/S_b	N/A	N/A	Double–pole, double–throw switch

Table A.2: Final component values used in the optimised loop filter circuit.

Figure A.2: Low–bandwidth feedback circuit

A.3 MOT coil circuit

Figure A.3: MOT coil switch used for controlling the quadrupole magnetic field gradient.

Appendix B

Temperature stabilising a high-finesse optical cavity

To measure the temperature dependent frequency shift, the cavity temperature is changed using heater tape and monitored using a thermocouple. By monitoring the frequency of the light entering the cavity (using a WS7 High–Finesse Wavemeter) and observing the transmitted light on a CCD camera, such that the cavity TEM_{mn} mode is known, the approximate frequency change as a function of cavity temperature can be measured. The cavity resonance frequency varies by ≈ 370 MHz/°C, and it decreases with increasing temperature.

To provide temperature control of the cavity, 10 m insulated constantan cable of thickness 0.56 mm is wound in a field–free manner around the vacuum chamber, giving a total of 18 turns that are spaced by 10 mm. The wire has a resistance of 18 Ω. Two 10 kΩ Negative Temperature Coefficient (NTC) thermistors (Farnell 167–2402) are attached to the centre of the chamber on each side of the cavity: one is used to monitor the temperature; the other is used in the control of the cavity temperature. A homebuilt proportional–intergal (PI) circuit is used, see figure B.1 for a simplified diagram, to temperature stabilise the cavity to approximately 28.3 °C via control of the current through the wire. The entire cavity is wrapped in fibre glass insulation and a wooden box, with holes at either end to permit beam entry/escape, is placed over the vacuum chamber to provide additional thermal insulation.

	Ziegler–Nichols [54]			Tyreus–Luyben [102]		
Control	K_P	τ_I	τ_D	K_P	τ_I	τ_D
P	$K_u/2$					
PI	$K_u/2.2$	$P_u/1.2$		$K_u/3.2$	$2.2P_u$	
PID	$K_u/1.7$	$P_u/2$	$P_u/8$	$K_u/2.2$	$2.2P_u$	$P_u/6.3$

Table B.1: Optimum proportional ($K_P = R_p/R_{in}$), integral, and differential gains for the Ziegler–Nichols and Tyreus–Luyben method, where K_u is the ultimate gain, P_u the ultimate period of oscillation, K_P the proportional gain, $\tau_I = \frac{1}{(2\pi R_{in} C_i)}$ the integrator time constant, and τ_D the differential time constant.

The PI circuit is housed within the wooden box to provide a temperature stabilised environment for the temperature sensitive components in the Wheatstone bridge circuit. Figure B.1 illustrates the basic assembly of the PI temperature circuit. The thermistor attached to cavity vacuum chamber forms a section of the Wheatstone bridge, such that small resistance changes of the thermistor, due to changes in temperature, can be measured. The output of the bridge is amplified and converted into a voltage. To set the cavity to a desired temperature, an operational amplifier, whose input can be changed via a potentiometer, is placed between the path of the bridge amplifier and PI operational amplifier. The PI operational amplifier provides the corrections, which are sent via a FET amplifier output stage, to maintain the set temperature of the cavity.

There are numerous procedures one can carry out to optimise the PID gains of a circuit. The Ziegler–Nichols (ZN) method [54] is used to tune the temperature of the 689 nm laser. In this method, one drives the system into oscillation using only proportional feedback: this is the ultimate gain of the system (K_u). From the oscillation period (P_u), the integral and differential gains can be set, see table B.1 for optimum values. A modified version of this is the Tyreus–Luyben (TL) method [102]. The procedure is the same however the optimum values are set a little differently. Past experience of temperature tuning large thermal loads have taught us that the TL tuning method performs slightly better than the ZN method, in that the system reaches its set–point quicker with little ringing or over–damping. It is for this reason why we have coarsely tuned the PI gains of the cavity temperature controller using the TL method. For more fine tuning we slightly tweaked the gains for optimum performance.
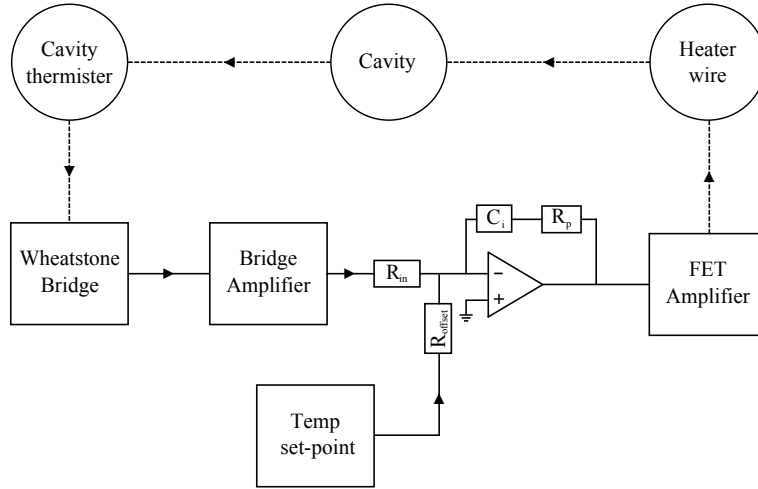
Figure B.1: Schematic cavity temperature stabilisation feedback loop. Stages in the PI circuit are shown in black, outlined boxes, with electrical paths shown using black, solid lines. Components external to the circuit are shown in black, outlined circles, with electrical paths shown using black, dashed lines. Note that R_{offset} should be equal to R_{in}.

To measure the temperature of the cavity we supply a constant current through the second thermistor. Using a National Instruments DAQ (NI USB–6008) we can measure, and monitor, the voltage across the thermistor over time in conjunction with a LabVIEW monitoring program. From the measured voltage, and knowing the current through the thermistor, we can infer the resistance of the thermistor using Ohm's law, which can then be converted to a temperature using

$$T = \frac{1}{\left[\ln\left(\frac{R}{R_{25}}\right)\frac{1}{B} + \frac{1}{T_{25}}\right]} - 273, \tag{B.1}$$

with $B = 4100$ K (± 1 %), $R_{25} = 10$ kΩ, which is the thermistor resistance at 25 °C, $T_{25} = 298$ K, and $R = R_{25}\exp\left[B\left(\frac{1}{T} - \frac{1}{T_{25}}\right)\right]$. The equation and constants can be found on the thermistor datasheet [103, 104].

Monitoring the cavity temperature over a period of several hours and calculating the Allan deviation using equation 2.7 gives an indication of the temperature stability and drift. Figure B.2 illustrates the absolute temperature stability of the cavity when the temperature is set to 28.3 °C and monitored for ≈ 28 hours. The temperature stability reaches a minimum of 250 mK, which we believe is the noise floor of the measurement apparatus. There is no obvious indication that this increasing (i.e. an indication of slow temperature drift) but

this might be apparent after an averaging time of 10^4 s. However, the DAQ used to convert the measured resistance from the thermistor into a voltage can only resolve 100 mK steps, therefore we would need to improve the precision of the temperature measurement before monitoring temperature for longer durations.


Figure B.2: Allan deviation of long–term cavity temperature.

Appendix C

Ramping the magnetic field

We are unable to supply an analogue ramp to our coils. During the vacuum chamber build, one of the MOT coils shorted to the inside of chamber. To prevent current flowing into the chamber walls, a digital optocoupler is used in the MOT coil circuit to isolate the TTL computer ground from the circuit ground. This essentially means the MOT coil circuit is 'floating'. It is possible to switch between high and low magnetic fields by implementing a DG413 analogue single–pole, single–throw switch in the MOT coil circuit. The use of the switch makes it possible to control whether the MOT field gradient is in the 'high' (24 G/cm), 'low' (3–5 G/cm) or 'no field' state. The magnetic field gradient of the 'high' and 'low' levels are set using potentiometers. The switch logic is shown in table C.1.

$\mathbf{TTL}_1$	$\mathbf{TTL}_2$	Coil	Field gradient
1	1	ON	LOW
1	0	ON	HIGH
0	1	OFF	N/A
0	0	OFF	N/A

Table C.1: MOT coil circuit logic table.

If an analogue ramp is required, the MOT coil circuit will need rebuilding to include an analogue optocoupler instead. However, an alternative, and simpler, method is to use pulse width modulation (PWM). In this technique, the average voltage (and current) fed to the load is controlled by rapidly switching between

the 'on' and 'off' state. The longer the switch is 'on', the greater the power supplied to the load. For example, shown in figure C.1 (a) is an example magnetic field gradient ramp sequence. It is necessary to reduce the magnetic field gradient down from 24 G/cm to 3 G/cm to increase the capture range for the initial broadband phase, then after, ≈ 10 ms, the MOT cloud is compressed by increasing the magnetic field gradient from 3 G/cm to 10 G/cm within 50 ms, which is held at this value for a further 20 ms. Figure C.1 (b) shows the PWM waveform corresponding to this ramp sequence. Recalling the logic of the circuit switch, given in table C.1, the pulse waveform begins in the 'high' state (10) as this state of the field gradient after the blue MOT. When the field gradient is 3 G/cm the switch is in the 'low' state (11) for the entire duration. During the ramp phase the MOT coil circuit switches between the high level (24 G/cm) and low level (3 G/cm), i.e rapidly switches between (10) $\leftrightarrow$ (11). The pulse durations decrease from τ_1 to from τ_2 during the 50 ms, then remains fixed at a duration of τ_2, resulting in a field gradient of 10 G/cm. Note that in the experiment the PWM waveform is of higher resolution than that shown in figure C.1 (b).

Figure C.1: (a) Example ramp of the magnetic field gradient (b) Corresponding pulse width modulation waveform for the ramp in (a). Note that in the experiment the pulse width modulation waveform is of higher resolution than that shown in (b). The duration of the TTL pulse at the beginning and end of the ramp are labelled τ_1 and τ_2 respectively.

A small piece of code has be written to generate the PWM waveform. The basic procedure is as follows. First the ramp is separated into five stages: the initial 'high' state; the 'low' state; the ramp; the intermediate state; and the final 'high' state. The duration of each stage can be inputted, as well as the magnetic field gradient intermediate state. The code generates a sawtooth wave that has a frequency that can be inputted and an amplitude that depends on the difference between the magnetic field gradients between the 'high' and 'low' states, e.g. in figure C.1 (a) the maximum amplitude is 21 G/cm. For each stage, the shape of the waveform is defined, e.g. the ramp has a linearly increasing waveform, whereas the intermediate stage has a fixed amplitude, then the amplitude of the waveform in each stage is compared to that of the carrier sawtooth wave. If the waveform amplitude is greater than that of the carrier sawtooth wave, the TTL is in the 'off' state, otherwise it is in the 'on' state.

An arbitrary function generator is used to output the PWM waveform to the

'low' state TTL input using a support software package (ArbExpress). It is important to check the memory size of the arbitrary function generator as this will set an upper limit on the carrier sawtooth wave frequency. Too little memory will result in a noisy, oscillating magnetic field.

Bibliography

[1] I. Bloch, J. Dalibard, and S. Nascimbene. Quantum simulations with ultracold quantum gases. *Nature Physics*, 8:267–276, 2012.

[2] M. Opper and D. Saad. *Advanced Mean Field Methods: Theory and Practice.* MIT Press, 2001.

[3] T. F. Gallagher. *Rydberg Atoms.* Cambridge University Press, first edition, 1995.

[4] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with Rydberg atoms. *Rev. Mod. Phys.*, 82:2313–2363, Aug 2010.

[5] M. Saffman and T. G. Walker. Creating single-atom and single-photon sources from entangled atomic ensembles. *Phys. Rev. A*, 66:065403, Dec 2002.

[6] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller. Dipole blockade and quantum information processing in mesoscopic atomic ensembles. *Phys. Rev. Lett.*, 87:037901, Jun 2001.

[7] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. Fast quantum gates for neutral atoms. *Phys. Rev. Lett.*, 85:2208–2211, Sep 2000.

[8] A. Gaëtan, Y. Miroshnychenko, T. Wilk, A. Chotia, M. Viteau, D. Comparat, P. Pillet, A. Browaeys, and P. Grangier. Observation of collective excitation of two individual atoms in the Rydberg blockade regime. *Nature Physics*, 5:115–118, 2009.

[9] D. Tong, S. M. Farooqi, J. Stanojevic, S. Krishnan, Y. P. Zhang, R. Côté, E. E. Eyler, and P. L. Gould. Local blockade of Rydberg excitation in an ultracold gas. *Phys. Rev. Lett.*, 93:063001, Aug 2004.

[10] U. Raitzsch, V. Bendkowsky, R. Heidemann, B. Butscher, R. Löw, and T. Pfau. Echo experiments in a strongly interacting Rydberg gas. *Phys. Rev. Lett.*, 100:013002, Jan 2008.

[11] H. Schempp, G. Günter, C. S. Hofmann, C. Giese, S. D. Saliba, B. D. DePaola, T. Amthor, M. Weidemüller, S. Sevinçli, and T. Pohl. Coherent population trapping with controlled interparticle interactions. *Phys. Rev. Lett.*, 104:173602, Apr 2010.

[12] A. Osterwalder and F. Merkt. Using high Rydberg states as electric field sensors. *Phys. Rev. Lett.*, 82:1831–1834, 1999.

[13] P. Schauß, J. Zeiher, T. Fukuhara, S. Hild, M. Cheneau, T. Macri, T. Pohl, I. Bloch, and C. Gross. Dynamical crystallization in a low-dimensional Rydberg gas. *arXiv*, 2014.

[14] Y.O. Dudin and A. Kuzmich. Strongly interacting Rydberg excitations of a cold atomic gas. *Science*, 336:887–889, 2012.

[15] N. Henkel, R. Nath, and T. Pohl. Three-dimensional roton excitations and supersolid formation in Rydberg-excited Bose-Einstein condensates. *Phys. Rev. Lett.*, 104:195302, 2010.

[16] G. Pupillo, A. Micheli, M. Boninsegni, I. Lesanovsky, and P. Zoller. Strongly correlated gases of Rydberg-dressed atoms: Quantum and classical dynamics. *Phys. Rev. Lett.*, 104:223002, 2010.

[17] F. Maucher, N. Henkel, M. Saffman, W. Królikowski, S. Skupin, and T. Pohl. Rydberg-induced solitons: Three-dimensional self-trapping of matter waves. *Phys. Rev. Lett.*, 106:170401, 2011.

[18] T. Macrìand T. Pohl. Rydberg dressing of atoms in optical lattices. *Phys. Rev. A*, 89:011402, 2014.

[19] L. I. R. Gil, R. Mukherjee, E. M. Bridge, M. P. A. Jones, and T. Pohl. Spin squeezing in a Rydberg lattice clock. *Phys. Rev. Lett.*, 112:103601, 2014.

[20] A. Grabowskir, R. Heidemann, R. Löw, J. Stuhler, and T. Pfau. High resolution Rydberg spectroscopy of ultracold rubidium atoms. *Fortschr. Phys.*, 54:8–10, 2006.

[21] J. Millen, G. Lochead, and M. P. A. Jones. Two-electron excitation of an interacting cold Rydberg gas. *Phys. Rev. Lett.*, 105:213004, Nov 2010.

[22] D. Tong, S. M. Farooqi, J. Stanojevic, S. Krishnan, Y. P. Zhang, R. Côté, E. E. Eyler, and P. L. Gould. Local blockade of Rydberg excitation in an ultracold gas. *Phys. Rev. Lett.*, 93:063001, Aug 2004.

[23] J. D. Pritchard, D. Maxwell, A. Gauguet, K. J. Weatherill, M. P. A. Jones, and C. S. Adams. Cooperative atom-light interaction in a blockaded Rydberg ensemble. *Phys. Rev. Lett.*, 105:193603, Nov 2010.

[24] J. Millen, G. Lochead, G. R. Corbett, R. M. Potvliege, and M. P. A. Jones. Spectroscopy of a cold strontium Rydberg gas. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(18):184001, 2011.

[25] R. Mukherjee, J. Millen, R. Nath, M. P. A. Jones, and T. Pohl. Many-body physics with alkaline-earth Rydberg lattices. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(18):184010, 2011.

[26] C. L. Vaillant, M. P. A. Jones, and R. M. Potvliege. Long-range Rydberg–Rydberg interactions in calcium, strontium and ytterbium. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 45(13):135004, 2012.

[27] J. Millen. *A cold strontium Rydberg gas*. PhD thesis, Durham University, 2011.

[28] G. Lochead. *Excited state spatial distributions in a cold strontium gas*. PhD thesis, Durham University, 2012.

[29] B. H. Bransden and C. J. Joachain. *Physics of Atoms and Molecules*. Prentice Hall, second edition, 2003.

[30] S. B. Nagel, P. G. Mickelson, A. D. Saenz, Y. N. Martinez, Y. C. Chen, T. C. Killian, P. Pellegrini, and R. Côté. Photoassociative spectroscopy at long range in ultracold strontium. *Phys. Rev. Lett.*, 94:083004, 2005.

[31] The National Institute of Standards and Technology (NIST) database, available online at http://www.nist.gov/pml/data.pdf.

[32] X. Xu, T. H. Loftus, M. J. Smith, J. L. Hall, A. Gallagher, and J. Ye. Dynamics in a two-level atom magneto-optical trap. *Phys. Rev. A*, 66:011401, 2002.

[33] Thomas H. Loftus, Tetsuya Ido, Martin M. Boyd, Andrew D. Ludlow, and Jun Ye. Narrow line cooling and momentum-space crystals. *Phys. Rev. A*, 70:063413, 2004.

[34] S. Blatt, A. D. Ludlow, G. K. Campbell, J. W. Thomsen, T. Zelevinsky, M. M. Boyd, J. Ye, X. Baillard, M. Fouché, R. Le Targat, A. Brusch, P. Lemonde, M. Takamoto, F.-L. Hong, H. Katori, and V. V. Flambaum. New limits on coupling of fundamental constants to gravity using ^{87}Sr optical lattice clocks. *Phys. Rev. Lett.*, 100:140801, 2008.

[35] N. Poli, F.-Y. Wang, M. G. Tarallo, A. Alberti, M. Prevedelli, and G. M. Tino. Precision measurement of gravity with cold atoms in an optical lattice and comparison with a classical gravimeter. *Phys. Rev. Lett.*, 106:038501, 2011.

[36] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269(5221):198–201, 1995.

[37] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969–3973, 1995.

[38] S. Stellmer, M. K. Tey, B. Huang, R. Grimm, and F. Schreck. Bose-Einstein condensation of strontium. *Phys. Rev. Lett.*, 103:200401, 2009.

[39] M. K. Tey, S. Stellmer, R. Grimm, and F. Schreck. Double-degenerate Bose-Fermi mixture of strontium. *Phys. Rev. A*, 82:011608, 2010.

[40] C. J. Foot. *Atomic Physics*. Oxford University Press, first edition, 2005.

[41] Y. Castin, H. Wallis, and J. Dalibard. Limit of Doppler cooling. *J. Opt. Soc. Am. B*, 6(11).

[42] M. Boyd. *High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time*. PhD thesis, Colorado University, 2007.

[43] D. W. Allan. Time and frequency (time-domain) characterization, estimation, and prediction of precision clocks and oscillators. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 34:647–654.

[44] L. S. Cutler and C. L Searle. Some aspects of the theory and measurements of frequency fluctuations in frequency standards. *Proceedings of IEEE*, 54:136–154, 1966.

[45] NIST frequency standards characterisations, available online at http://www.tf.nist.gov/timefreq/general/pdf/1286.pdf.

[46] A. L. Schawlow and C. H. Townes. Infrared and optical masers. *Phys. Rev.*, 112:1940–1949, 1958.

[47] C. Wiemann and L. Hollberg. Using diode lasers for atomic physics. *Rev. Sci. Inst. (invited review)*, 62, 1991.

[48] H. Loh, Y. Lin, I. Teper, M. Cetina, J. Simon, J. K. Thompson, and V. Vuletić. Influence of grating parameters on the linewidths of external-cavity diode lasers. *Appl. Opt.*, 45(36):9191–9197, 2006.

[49] E. D. Black. An introduction to Pound-Drever-Hall laser frequency stabilization. *American journal of physics*, 69:79, 2001.

[50] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied physics B: Lasers and optics*, 31:97, 1983.

[51] A. S. Arnold, J. S. Wilson, and M. G. Boshier. A simple extended-cavity diode laser. *Rev. Sci. Inst*, 69(3):1236–1239, 1998.

[52] K. G. Libbrecht and J. L. Hall. A low-noise high speed diode laser current controller. *Rev. Sci. Inst.*, 64:2133–2135, 1993.

[53] Vescent photonices d2-105 laser controller datasheet, available online at http://www.vescent.com/wp-content/uploads/2011/02/d2-105-laser_controller_flyer.pdf.

[54] J. G. Ziegler, N. B. Nichols, and N. Y. Rochester. Optimum settings for automatic controllers. *American Society of Mechanical Engineers*, 171, 1992.

[55] G. Brooker. *Modern Classical Optics*. Oxford University Press, 2003.

[56] A. E. Siegman. *Lasers*. University Science Books, 1986.

[57] H. Kogelnik and T. Li. Laser beams and resonators. *Preceedings of the IEEE*, 54(10):1312 – 1329, 1966.

[58] T. Doiron and J. Beers. The gauge block handbook. *Dimensional Metrology Group - The National Institute of Standards and Technology (NIST) database.*

[59] Y. Li, T. Ido, T. Eichler, and H. Katori. Narrow-line diode laser system for laser cooling of strontium atoms on the intercombination transition. *Applied Physics B*, 78(3-4):315–320, 2004.

[60] R. W. Fox, C. W. Oates, and L. W. Hollberg. Experimental methods in the physical sciences: Stabilising diode lasers to high finesse cavities. *Elsevier Science (USA)*, 40, 2003.

[61] Hamamatsu datasheet, available online at http://www.hamamatsu.com/resources/pdf/ssd/s5971etckpin1025e07.pdf.

[62] Linear technology datasheet, available online at http://cds.linear.com/docs/en/datasheet/lt1226.pdf.

[63] P. Horowitz and W. Hill. *The Art of Electronics - 2nd Edition.* Cambridge University Press, 1989.

[64] H. Ludvigsen, M. Tossavainen, and M.] Kaivola. Laser linewidth measurements using self-homodyne detection with short delay. *Optics Communications*, 155(1–3):180 – 186, 1998.

[65] G. M. Tino, M. Barsanti, M. de Angelis, L. Gianfrani, and M. Inguscio. Spectroscopy of the 689 nm intercombination line of strontium using an extended-cavity InGaP/InGaAIP diode laser. *App. Phys. B*, 55:397 – 400, 1992.

[66] M. L. Harris, C. S. Adams, S. L. Cornish, I. C. McLeod, E. Tarleton, and I. G. Hughes. Polarization spectroscopy in rubidium and cesium. *Phys. Rev. A*, 73:062509, 2006.

[67] S. Stellmer. *Degenerate quantum gases of strontium.* PhD thesis, Innsbruk University, 2013.

[68] W. Demtröder. *Laser Spectroscopy.* Springer, second edition, 1998.

[69] M. S. Sorem and A. L. Schawlow. Saturation spectroscopy in molecular iodine by intermodulated fluorescence. *Opt. Commun.*, 5:148, 1972.

[70] C. Freed and A. Javan. Standing-wave saturation resonances in the CO_2 10.6μm transitions observed in a low-pressure room-temperature absorber gas. *Appl. Phys. Lett.*, 17, 1970.

[71] I. Courtillot. *Premiere observation de la transition fortement $^1S_0 \rightarrow {}^3P_0$ du strontium, pour une horloge optique a atomes pieges*. PhD thesis, Paris University, 2003.

[72] A. Quessada. *Developpement d'une horlage optique a atomes de strontium pieges: Realisation d'un laser ultra-stable et stabilite de frequence*. PhD thesis, Paris University, 2005.

[73] Datasheet available online at http://www.farnell.com/datasheets/316996.pdf.

[74] Datasheet available online at http://www.analog.com/static/imported-files/data_sheets/ad795.pdf.

[75] Analogue devices datasheet, available online at http://www.analog.com/static/imported-files/data_sheets/op270.pdf.

[76] Thorlabs piezo amplifier datasheet, available online at http://sites.fas.harvard.edu/ phys191r/bench_notes/d4/thorlabs_mdt693a.pdf.

[77] A. K. Mollema, L. W. Wansbeek, L. Willmann, K. Jungmann, R. G. E. Timmermans, and R. Hoekstra. Laser-frequency locking using light-pressure-induced spectroscopy in a calcium beam. *Phys. Rev. A*, 77:043409, 2008.

[78] N. Poli, G. Ferrari, M. Prevedelli, F. Sorrentino, R.E. Drullinger, and G.M. Tino. Laser sources for precision spectroscopy on atomic strontium. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, 63(5):981 – 986, 2006.

[79] I. R. Hill. *Development of an Apparatus for a Strontium Optical Lattice Optical Frequency Standard*. PhD thesis, Imperial College London, UK, 2012.

[80] A. Ludlow. *The Strontium Optical Lattice Clock: Optical Spectroscopy with Sub-Hz Accuracy*. PhD thesis, Colorado University, 2008.

[81] E. M. Bridge, J. Millen, C. S. Adams, and M. P. A. Jones. A vapor cell based on dispensers for laser spectroscopy. *Review of Scientific Instruments*, 80(1), 2009.

[82] D. J. McCarron, S. A. King, and S. L. Cornish. Modulation transfer spectroscopy in atomic rubidium. *Measurement science and technology*, 19:105601, 2008.

[83] I. Courtillot, A. Quessada, R. P. Kovacich, J-J. Zondy, A. Landragin, A. Clairon, and P. Lemonde. Efficient cooling and trapping of strontium atoms. *Opt. Lett.*, 28(6):468–470, Mar 2003.

[84] C. J. Dedman, J. Nes, T. M. Hanna, R. G. Dall, K. G. H. Baldwin, and A. G. Truscott. Optimum design and construction of a Zeeman slower for use with a magneto-optic trap. *Rev. Sci. Instrum.*, 75:5136, 2004.

[85] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard. Trapping of neutral sodium atoms with radiation pressure. *Phys. Rev. Lett.*, 59:2631, 1987.

[86] P. G. Mickelson, Y. N. Martinez de Escobar, P. Anzel, B. J. DeSalvo, S. B. Nagel, A. J. Traverso, M. Yan, and T. C. Killian. Repumping and spectroscopy of laser-cooled Sr atoms using the $(5s5p)^3P_2$ –$(5s4d)^3D_2$ transition. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 42(23):235001, 2009.

[87] N. Poli, R. E. Drullinger, G. Ferrari, J. Léonard, F. Sorrentino, and G. M. Tino. Cooling and trapping of ultracold strontium isotopic mixtures. *Phys. Rev. A*, 71:061403, Jun 2005.

[88] N. Poli. *Raffreddamento ed intrappolamento di atomi di stronzio: verso un nuovo standard di frequenza nella regione ottica*. PhD thesis, Universita degli studi di Firenze, Florence, Italy, 2005.

[89] K. R. Vogel, T. P. Dinneen, A. Gallagher, and J. L. Hall. Narrow-line Doppler cooling of strontium to the recoil limit. *IEEE transactions on instrumentation and measurement*, 48:618, 1999.

[90] T. H. Loftus, T. Ido, A. D. Ludlow, M. M. Boyd, and J. Ye. Narrow line cooling: Finite photon recoil dynamics. *Phys. Rev. Lett.*, 93:073003, 2004.

[91] H. Katori, T. Ido, Y. Isoya, and M. Kuwata-Gonokami. Magneto-optical trapping and cooling of strontium atoms down to the photon recoil temperature. *Phys. Rev. Lett.*, 82:1116–1119, 1999.

[92] R. P. Abel, A. K. Mohapatra, M. G. Bason, J. D. Pritchard, K. J. Weatherill, U. Raitzsch, and C. S. Adams. Laser frequency stabilization to excited state transitions using electromagnetically induced transparency in a cascade system. *Appl. Phys. Lett.*, 94:071107, 2009.

[93] K.-J. Boller, A. Imamolu, and S. E. Harris. Observation of electromagnetically induced transparency. *Phys. Rev. Lett.*, 66:2593–2596, May 1991.

[94] M. Fleischhauer, A. Imamoglu, and J. P. Marangos. Electromagnetically induced transparency: Optics in coherent media. *Rev. Mod. Phys.*, 77:633–673, Jul 2005.

[95] E. Arimondo. Coherent population trapping in laser spectroscopy. *Progress in Optics*, 35:257, 1996.

[96] O. A. Kocharovskaya and Ya. I. Khanin. Population trapping and coherent bleaching of a three-level medium by a periodic train of ultrashort pulses. *Sov. Phys. J.E.T.P.*, 63:945, May 1986.

[97] A. Schwarzkopf, R. E. Sapiro, and G. Raithel. Imaging spatial correlations of Rydberg excitations in cold atom clouds. *Phys. Rev. Lett.*, 107:103001, Aug 2011.

[98] W. E. Cooke, T. F. Gallagher, S. A. Edelstein, and R. M. Hill. Doubly excited autoionizing Rydberg states of Sr. *Phys. Rev. Lett.*, 40:178–181, Jan 1978.

[99] R. Löw, U. Raitzsch, R. Heidemann, V. Bendkowsky, B. Butscher, A. Grabowski, and T. Pfau. Apparatus for excitation and detection of Rydberg atoms in quantum gases. arXiv:0706.2639v1, 2007.

[100] W. Li, P. J. Tanner, and T. F. Gallagher. Dipole-dipole excitation and ionization in an ultracold gas of Rydberg atoms. *Phys. Rev. Lett.*, 94:173001, 2005.

[101] T. Pohl, E. Demler, and M. D. Lukin. Dynamical crystallization in the dipole blockade of ultracold atoms. *Phys. Rev. Lett.*, 104:043002, Jan 2010.

[102] B. D. Tyreus and L. L. Luyben. Tuning PI controllers for integrated/dead time processes. *Industrial Engineering Chemistry Research*, 31(171):2625 – 2628, 1992.

[103] Datasheet, available online at http://www.farnell.com/datasheets/86141.pdf.

[104] Datasheet, available online at

http://assets.newport.com/webdocuments_en/images/an04_thermistor_calibration_ix.pdf.