

Atoms in Microcavities: Detection and Spectroscopy

Joanna Kenner

Thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy of the Imperial College London
and for the Diploma of Membership of Imperial College.

Department of Physics, Imperial College London

March 2010

Referees

Referee: Professor Wolfgang Lange
 Department of Physics
 University of Sussex

Co-referee: Professor John Tisch
 Department of Physics
 Imperial College London

Declaration

I declare that this thesis is my own work and has been completed without any aids, other than those declared.

Abstract

This thesis presents work undertaken with cold rubidium atoms interacting with an optical microcavity. The optical microcavity used is unique in its design, being formed between an optical fibre and silicon micromirror. This allows direct optical access to the cavity mode, whilst the use of microfabrication techniques in the design means that elements of the system are inherently scalable. In addition, the parameters of the system are such that a single atom has a substantial impact on the cavity field.

In this system, two types of signal arise from the atoms' interaction with the cavity field; a 'reflection' signal and a 'fluorescence' signal. A theoretical description for these signals is presented, followed by experiments which characterise the signals under a variety of experimental conditions. The thesis then explores two areas: the use of the microcavity signals for atom detection and the investigation of how higher atom numbers and, as a result, a larger cooperative interaction between the atoms and the cavity field, impacts the signals.

First, the use of these signals to detect an effective single atom and individual atoms whilst falling and trapped is explored. The effectiveness of detection is parameterised in terms of detection confidence and signal to noise ratio, detection fidelity and dynamic range.

In the second part of this thesis, the effect of higher atom numbers on the reflection and fluorescence signals is investigated. A method for increasing the atom number is presented, alongside experiments investigating the impact on the measured signals. This is followed by experiments which explore the dispersive nature of the atom-cavity interaction by measuring the excitation spectrum of the system in reflection and fluorescence. In doing so, it is shown that, for weak coupling, these two signals are manifestly different.

Acknowledgements

There are many people I want to thank for their help and support during the last four years. Firstly, thank you to my supervisor, Prof. Ed Hinds, for giving me the opportunity to join CCM as a postgraduate and believing that two years as a management consultant hadn't destroyed my aptitude for physics. CCM is a great place to work; its members being committed and enthusiastic about their projects, but also willing to help! Thank you to all I've worked with, past and present, with particular note to Ben, Jony, Jos, Rob, Benoit, Isabel, Sean, Joe, Florian and Sanja.

A cohort of five began their PhD in the group alongside me; Dhiren, Thom, Jake, Sam and Manuel. The lion's share of answering my questions has fallen to them, so thank you! I have also been fortunate to have three whole post docs working with me on the cavity experiment, all of whom deserve individual thanks. Adrian, thank you for patiently answering a million different hypothetical, physics-related questions. Micheal, thank you for your 'can-do' attitude, your great ideas and keeping me company during the nights and weekends of data-taking. Jon, thank you for teaching me how to build and use (literally) everything on the experiment and for being a wonderful touchstone for feedback whilst I was writing up. To all three of you I can only say: TIME, CHANGE, QUARTZ!

Out of the lab, I am lucky to have some excellent friends, who've provided great company during good times and suitable distraction when things have not been going so well. Particular thanks belong to Laura, Donna, Carol, Jo, Liam and Aodhnait and the Pembroke Wardening team.

My final thanks are reserved for my family; Mum, Dad, Jamo and Izzy, and Becca, without whom I'd be nothing; and for Steve, for putting up with me (and my nose). Thank you all.

Contents

1	Introduction	31
1.1	Cavity QED	31
1.2	Realisations of cavity QED	34
1.3	Experiments and applications of cavity QED	36
1.3.1	Single atom detection	36
1.3.2	Spectroscopy of the coupled atom-cavity system and non-linear optics .	38
1.3.3	Quantum information processing	39
1.4	Outline of this thesis	40
2	Atoms and cavities	42
2.1	Optical cavities	42
2.1.1	Reflection and transmission properties of an optical cavity	42
2.1.2	Mode structure of an optical cavity	46
2.2	The quantum electric field	48
2.2.1	Quantising the electric field	48
2.2.2	The driven, damped quantised cavity field	51
2.3	The interaction between atoms and a quantised cavity field	52
2.3.1	The two-level atom interacting with an electric field	53
2.3.2	The Jaynes-Cummings Hamiltonian	54
2.3.3	The Master equation for a driven atom-cavity system	57
2.3.4	The coherent cavity field in the limit of weak excitation	58
2.3.5	Enhanced spontaneous emission: the Purcell effect	60
2.3.6	Extending to multiple atoms	61
2.3.7	High intensity limit and optical bistability	62
2.4	Cavity signals	63
2.4.1	Cavity reflection signals	63

<i>CONTENTS</i>	6
2.4.2 Cavity fluorescence signals	66
3 Experiment set up	70
3.1 High level description of experiment procedure and layout	70
3.2 The cavities	71
3.2.1 Design and fabrication	71
3.2.2 Characterisation of properties	74
3.2.3 Cavity pumping and stabilisation	76
3.3 Atom preparation	80
3.3.1 Rubidium	80
3.3.2 The MOT	81
3.3.3 The molasses	85
3.3.4 Transport and driving beams	88
3.4 Computer control	89
3.5 Laser sources	90
3.5.1 Overview of laser sources	90
3.5.2 Reference laser	91
3.5.3 DL100	96
3.5.4 Repump laser	100
3.5.5 The Cavity laser	104
3.5.6 Tapered amplifier system (Toptica BoosTA)	106
4 Atom detection	108
4.1 Producing cavity signals	108
4.1.1 Cavity reflection signals	109
4.1.2 Cavity fluorescence signals	112
4.2 Characterising the cavity signals	116
4.2.1 Intensity	117
4.2.2 Detuning	122
4.2.3 Effective atom number	126
4.2.4 Polarisation	128
4.3 Single atom detection	132
4.3.1 Signal to noise ratio and Poissonian statistics	133
4.3.2 Detecting $N_{\text{eff}} = 1$	134

4.3.3	Detecting individual atoms falling through the cavity mode	143
4.3.4	Detecting individual atoms trapped inside the cavity	151
5	Increasing the atom number	158
5.1	Motivation for increasing N_{eff}	158
5.2	Increasing N_{eff}	159
5.2.1	Experimental method	159
5.2.2	Characterisation	164
5.3	Reflection and fluorescence signals at higher N_{eff}	168
5.3.1	Reflection signals and N_{eff}	168
5.3.2	Fluorescence signals and N_{eff}	171
5.4	The dispersive atom-cavity interaction: 2D scans of reflection and fluorescence signals	174
5.4.1	2D scan in reflection	177
5.4.2	2D scan in fluorescence	182
6	Future directions	187
6.1	Conveying and trapping atoms inside the cavity	187
6.1.1	An intra-cavity dipole trap	187
6.1.2	Conveying atoms to the trap	193
6.2	Coherent control of atoms inside the cavity	194
6.2.1	Incoherent control: state selectivity of atom detection in the cavity . . .	195
6.2.2	Coherent control: EIT	196
7	Conclusions and outlook	200
7.1	Conclusions	200
7.2	Outlook	202
A	The fluorescence signal and increasing drive beam	214

List of Figures

1.1	Schematic of a coupled atom-cavity system. The system is described by three rates; the rate of coherent coupling between the atom and the field, g ; the spontaneous emission rate of the atom, γ ; and the decay rate of the cavity field, κ .	33
2.1	A schematic of a Fabry-Perot optical cavity. The cavity is formed by two mirrors separated by a distance d . Each mirror has an associated reflection and transmission field-amplitude coefficients, r_1, r_2 and t_1, t_2 , respectively. A field of amplitude E_{in} impinges on mirror one, leading to a circulating field, E_{circ} inside the cavity and a transmitted field, $E_{out} = t_2 E_{circ}$.	43
2.2	The reflected intensity as a function of cavity length, d , shown for two values of the mirror reflectivities. The solid, light green line is calculated for $r_1 = r_2 = 0.99$ and the dashed, dark green line is calculated for $r_1 = r_2 = 0.8$. In both cases $\lambda = 780\text{nm}$. Each cavity fringe is caused by the destructive interference between the initially reflected light at mirror one and the backwards leakage through mirror one of the circulating cavity field. As the mirror reflectivities increase, the resonance fringes become narrower.	44
2.3	A schematic showing the ladder of states of doublet states formed as a result of the atom-cavity interaction. The LHS ladder is for $g = 0$, whilst the RHS ladder shows the same but now $g > 0$. On the LHS, the atom and cavity are detuned by δ_{AC} , which produces the small splitting to what would otherwise be degenerate states.	54

- 2.4 The system eigenvalues as a function of atom-cavity detuning, δ_{AC} , plotted for $n = 1$. The ‘bare’ eigenvalues for the composite atom-cavity system, for which there is no interaction between the atom and the cavity field, are represented and labelled by the dotted lines. The ‘dressed state’ eigenvalues include an interaction term and are represented by the solid line. In the latter case, the interaction causes the eigenvalues to repel one another close to $\delta_{AC} = 0$, causing an avoided crossing. 56
- 2.5 The cavity transmission as a function of detuning, Δ where $\Delta = \Delta_C - \Delta_A$. The solid line represents an empty cavity which shows one peak centred on $\Delta = 0$. The dashed line represents a cavity with one atom present, with $(g, \gamma, \kappa) = 2\pi \times (80, 3, 50)$ MHz, so that $C = 21.3$ and the system is in the strong coupling regime, to be defined later. The normal mode splitting of the cavity resonance symmetrically about $\Delta = 0$ is clearly visible. 58
- 2.6 The effect of increasing pumping powers on the intra-cavity field. $|\alpha_{res}|^2$ plotted as a function of cavity pumping rate, η for $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz. The lighter green, solid line is calculated for $N = 1$, corresponding to $C_N = 0.3$. In this case, $|\alpha_{res}|^2$ increases monotonically with increasing η . The darker green, dashed line is calculated for $N = 13$, corresponding to $C_N = 4$. For this case, $|\alpha^{res}|^2$ is multi-valued as the system becomes bistable. 62
- 2.7 The impact of different fringe visibilities, v , on the reflection signal. LHS: The normalised reflection signal, J_R/J_{in} as a function of cooperativity, plotted for three different values of v . From darker to lighter green, $v^2 = 0.01, 0.25, 0.5$. The reflection signal saturates with increasing C_N , saturating more quickly for smaller values of v . The grey dots indicate how an increase in cooperativity has an equivalent effect to a reduction in fringe contrast. RHS: The differential of J_R/J_{in} as a function of atom number, N for three different cavity contrasts. From darker to lighter green correspond to $v = 0.1, 0.2$ and 0.3 . As v increases (and thus the fringe contrast decreases), the change in signal caused by a single atom decreases, making it harder to discriminate between atom numbers. . . . 65

- 2.8 Schematic of fluorescence experiment. An atom inside the cavity is pumped into an excited state by transverse pump beams, passing vertically through the cavity from the top and/or the bottom. The atom's decay rate has two components: γ into free space and γ_{cav} into the cavity mode. Here, the cavity is formed between an optical fibre and a silicon wafer, as will be described in Chapter 3. 67
- 2.9 The fluorescence signal as a function of N_{eff} , in the 'ordered' and 'disordered' limits. This is plotted for the system parameters $(g, \kappa, \gamma, \Omega) = 2\pi \times (100, 5200, 3, 10)$ MHz. In the 'ordered limit', shown in dark red, J_{fl} initially increases quadratically with atom number, but for higher atom numbers the signal saturates. In contrast, the 'disordered limit', shown in light red, initially increases linearly with atom number and then rolls over, peaking when $N_{eff} = 1/2C$. As $N_{eff} \rightarrow \infty$, $J_{fl} \rightarrow 0$ 69
- 3.1 The optical microcavity experiment, showing all the major components. The typical position of the MOT is indicated by a red dot. 71
- 3.2 Schematic of the optical microcavity, showing the key features. The optical fibre is glued onto a piece of Macor, which itself is glued onto a glass block held underneath the gold mirror used to form the MOT, as discussed later in this section. A dielectric mirror is glued onto the front facet of the fibre to form one end of the cavity. The other end of the cavity is formed by one of an array of micromirrors etched onto the surface of a Silicon wafer. The wafer is glued onto a piece of Macor held on top of a piezo stack, which allows the length of the cavity to be adjusted. The white arrow indicates the direction travel of the piezo. The piezo stack is held in place using Macor pieces screwed together. . . 73
- 3.3 The reflected cavity fringes. LHS: The reflected cavity fringes for the 804 nm light used for locking and trapping (shown in light green) and the 780 nm light used for atom detection (shown in dark green), taken whilst the cavity length is being scanned. The relative amplitudes should not to be directly compared; the 780 nm is measured using a low pass filter to convert APD counts into an analogue voltage for the purposes of locking whilst the 804 nm is measured using a photodiode. RHS: A 'zoomed' scan of just one cavity fringe. The dashed line indicates an inverted Lorentzian lineshape, which is a good approximation for the cavity fringes. 74

3.4	Measuring cavity parameters. LHS: Three scans taken to measure the cavity linewidth, κ . The bottom scans show cavity fringes for only trapping light (RHS) and only repumping light (LHS) circulating inside the cavity. The top scan is for both trapping and repumping light, where the slight notch in the bottom of the fringe and increased width indicates that this fringe is the sum of two fringes of different frequency. RHS: The fringe contrast measured as a function of cavity length before gluing.	76
3.5	A schematic showing the arrangement of optics prior to the cavity fibre. This setup is designed to filter 804 nm light from the 780 nm which is detected at the APD. Four dichroic filters are used; three in reflection, labelled R_1 , R_2 and R_3 and one in transmission, T_1	78
3.6	The transmission (solid line) and reflection (dashed line) through the dichroic as a function of angle. The same units are used for transmission and reflection amplitudes, so that the scales are one-to-one. This indicates that a loss in transmission was accompanied by an equivalent gain in reflection.	79
3.7	A schematic of the D_2 manifold of ^{87}Rb . The upper and lower level hyperfine states are shown with the splitting between consecutive levels. The relative splitting is not to scale. To the right are arrows indicating the $ F = 2\rangle \rightarrow F' = 3\rangle$ cycling transition and the $ F = 1\rangle \rightarrow F' = 2\rangle$ repumping transition.	81
3.8	A schematic for the arrangement for the mirror MOT. A pair of beams diagonal to the mirror surface are reflected to form counter-propagating pairs. Two further beams directed in/out of the page, indicated by the dashed circle, are used to form a MOT. The MOT coils are aligned diagonally either side of the mirror.	82
3.9	CCD images of the MOT (LHS) after 3 s of loading and the atom cloud after 30 ms of molasses (RHS). The bright disk at the bottom of each image is scattered light from the 1 mm diameter hole in the mirror through which the atoms fall. The LHS and RHS pictures have spatial dimensions 6.5×6.2 mm and 5.5×5.2 mm, respectively, found by calibrating the distance per pixel to the known size of the hole.	84

- 3.10 Optimising the molasses stage. LHS: The ratio of cloud width with and without a molasses stage, measured as a function of molasses stage duration. The ratio appears to stabilise ~ 0.88 , with no apparent benefit from long durations. An exponential fit is shown alongside the data, with a $1/e$ lifetime of 1.7 ms. RHS: The normalised reflection signal amplitude, measured as a function of molasses beam power for two different frequencies; 47 MHz red-detuned (dark green); and 62 MHz red-detuned (light green). The maximum achievable power, set by the efficiency of the AOM at that frequency, is indicated by the dashed lines. The power is measured using a photodiode located inside the fibre port, which measures the pick-off from a 99:1 beam splitter, the output of which is measured with a voltmeter. In the case of the smaller detuning a peak of ~ 1.85 in the drop signal is visible for powers corresponding to ~ 10 mV. Higher drop signals can be achieved with the larger detuning, although no maximum was measured due to the inability to test higher powers. 86
- 3.11 The width of the atom cloud with time. Each point is found by fitting a Gaussian profile to the fluorescence image taken by the camera. Shown alongside the widths is a fit to Eq.(3.4), which estimates $T = 38.7 \pm 2.4 \mu\text{K}$ 88
- 3.12 Schematic of the setup for the beams used for pushing atoms and driving fluorescence. The beams are located vertically above and below the MOT and cavities and are delivered to the chamber using single mode, polarisation maintaining fibres. 89
- 3.13 TOP: A photograph of the Reference laser in situ. BOTTOM: A photograph of the same laser, but with the aluminium box open to show the internal parts. . . 92
- 3.14 Schematic of the optical layout for polarisation spectroscopy, with all the relevant parts labelled. $\frac{\lambda}{2}$ and $\frac{\lambda}{4}$ refer to a half-wave and quarter-wave plate respectively. ‘PBS’ and ‘NPBS’ refer to polarising beam splitters and non-polarising beam splitters. ‘PD 1’ and ‘PD 2’ are the two photodiodes. The difference signal from these is used as an error signal. 94

- 3.15 Polarisation spectroscopy, showing the dispersive features across the $|F = 2\rangle \rightarrow |F' = 1, 2, 3\rangle$ of ^{87}Rb and $|F = 3\rangle \rightarrow |F' = 2, 3, 4\rangle$ of ^{85}Rb . LHS: A typical polarisation spectrum, with the lock point is indicated. RHS: Polarisation spectroscopy for three different values of the current through the solenoid surrounding the cell; $I = 0\text{ mA}$ (TOP); $I = 10\text{ mA}$ (MIDDLE); and $I = 20\text{ mA}$ (BOTTOM). The top and bottom graphs have been shifted vertically on the graph for clarity. 95
- 3.16 LHS: Saturated absorption spectroscopy for the $|F = 2\rangle \rightarrow |F' = 2, 3\rangle$ transitions using the DL100. RHS: The error signal used to frequency lock the DL100 to the Reference laser. The lock point is indicated and corresponds to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition, 266 MHz red-detuned from the Reference laser. 96
- 3.17 A schematic of the locking system used to frequency lock the DL100 to the Reference laser. The Reference and DL100 beatnote at $\sim 260\text{ MHz}$ is detected at an amplified fast photodiode. This signal has an amplitude sufficient to saturate the limiter, which then provides a constant beatnote signal at 10dBm. The signal gets split, with one branch used as a monitor. The other branch is amplified and then mixed down to $\sim 80\text{ MHz}$ at a mixer by mixing with a $\sim 180\text{ MHz}$ signal from the VCO. The mixer output is amplified and then split into two branches, producing an error signal which is sent to the piezo driver for the DL100. 98
- 3.18 The DL100 error signal as a function of input signal power. LHS: Oscilloscope traces of the error signal shown for 8 different input powers. From dark to light green corresponds to $-9 \rightarrow -30\text{ dBm}$ in 3 dBm steps. As the power increases the shape of the error signal and position of the zero crossing change, with the gradient around the zero crossing increasing with increasing power. RHS: The frequency position of the error signal as a function of power, for the data shown on the LHS. 99
- 3.19 The modulation amplitude (dark green) and offset (light green) of the error signal as a function of input frequency. Whilst the zero crossing (i.e. position when the error signal equals zero) occurs at 80.6 MHz, the minimum in the modulation amplitude occurs at the higher frequency of 81.2 MHz, implying that this is where the error signal is least sensitive to power variations. 100

3.20	Schematic of the optical layout for the DAVLL, with all the relevant parts labelled. $\frac{\lambda}{2}$ and $\frac{\lambda}{4}$ refer to a half-wave and quarter-wave plate respectively. ‘PD 1’ and ‘PD 2’ are the two photodiodes. The difference signal from these is used as an error signal.	100
3.21	LHS: Saturated absorption spectroscopy for the $ F = 1\rangle \rightarrow F' = 0, 1, 2\rangle$ transitions. The dashed line indicates the repumping transition on $ F = 1\rangle \rightarrow F' = 2\rangle$. RHS: The DAVLL error signal, with the lock point indicated.	101
3.22	The gradient of the DAVLL error signal about the zero-crossing for different vapour cell temperatures.	102
3.23	Schematic of the components used in the phase lock. The top left hand side shows the set up for a phase-locked voltage controlled oscillator, whose signal, at 6.035 GHz, is used to mix down the optical beatnote signal to ~ 800 MHz. The top right hand side shows the set up for detecting the optical beatnote, at 6.835 GHz. The mixed down signal is the ‘VCO’ input to the Evaluation Board, where it is compared to a ‘Reference clock’ signal, which is provided by the Meguro Signal Generator. The Evaluation Board contains a phase-frequency detector and charge pump, which outputs current pulses at a frequency proportional to the phase difference between the two input signals. The loop filter integrates these currents pulses, providing a voltage signal used for the peizo and current servo loops.	103
3.24	LHS: Saturated absorption spectroscopy for the $ F = 1\rangle \rightarrow F' = 0, 1, 2\rangle$ transitions. The dashed line indicates the repumping transition on $ F = 1\rangle \rightarrow F' = 2\rangle$. RHS: The phase lock error signal.	105
3.25	The performance of the BoosTA. LHS: Measured output power versus current, for a constant injection power of 2 mW. At this injection power a maximum output power of 210 mW can be produced. RHS: Measured output power versus injection power, for a constant current of 1500 mA. This indicates that the saturated output power for this current is 375 mW.	106

- 4.1 The recorded counts from the APD for a single shot of a reflection experiment, taken with $\tau_{int} = 50 \mu s$. At time $t = 0$ on the graph the MOT has already loaded for 3 seconds and the dashed lines indicate other key points in the experiment. At $t = 10$ ms, the MOT field is switched off and the cooling light is red-detuned for the optical molasses. At $t = 40$ ms the trapping light is switched off, allowing the atoms to fall. The signal produced by the atom cloud as it falls through the cavity mode is centred at $t \sim 75$ ms, reaching a peak count level of 25 counts in $50 \mu s$ (or 540 kCounts/sec) compared to a background level of 12 counts in $50 \mu s$ (or 240 kCounts/sec). 109
- 4.2 Understanding the reflection signal. LHS: N_{eff} as a function of time, calculated using Eq.(4.1) and assuming $N_A = 19.7 \times 10^6$, $T = 38.7 \mu K$ and $y_0 = 5.1$ mm. This implies that at the peak in atom density, we would expect there to be ~ 1 atom averaged across the cavity mode. RHS: The APD counts averaged over ten shots of the experiment. The Gaussian-like signal caused by the falling atoms is visible, centred at ~ 30 ms. Shown alongside the data is a fit, described in the main text, which produces estimates of $N_A = 1.89 \pm 0.08 \times 10^7$, $y_0 = 5.12 \pm 0.02$ mm and $T = 29.7 \pm 0.8 \mu K$ 111
- 4.3 Understanding the fluorescence signal. LHS: The counts recorded by the APD for a single shot of the experiment. At $t = 0$ the MOT has been loading for 3 seconds. Dashed lines indicate key points in the experimental; at $t = 10$ ms the MOT field is turned off and the atoms are released into the optical molasses; at $t = 20$ ms the shutter for the cavity pump light is closed and, following a ~ 2 ms delay, the cavity count rate falls to a low background level of ~ 1 count per $50 \mu s$; at $t = 40$ ms the cooling light is switched off and the atoms fall towards the cavity. At $t \sim 73$ ms a narrow peak is visible, corresponding to the collected fluorescence produced by the atoms as they fall through the cavity. At $t = 125$ ms the pump light is turned on to check that the cavity remained on resonance. RHS: A section of a recorded trace averaged across 250 fluorescence experiments. This data was taken using $\tau_{int} = 2 \mu s$ (as opposed to $\tau_{int} = 50 \mu s$ shown on the LHS), allowing the structure of the fluorescence signal to be visible. Shown alongside the data is an exponential fit, which gives an estimate of the $1/e$ lifetime of $114.0 \pm 0.4 \mu s$. The background counts have also been reduced to 0.004 counts in $2 \mu s$, the result of improved beam alignment. 114

- 4.4 Fluorescence signal produced using one beam (shown in the darker red) and two counter-propagating beams (shown in the lighter red). The y-axis corresponds to the number of counts at the APD per $100\mu\text{s}$ time bin, summed for five experimental shots. The two-beam signal has an amplitude of nearly three times that of the one-beam case. In addition, the two-beam signal decays more slowly, having a FWHM of $500\mu\text{s}$ compared to $300\mu\text{s}$ in the one-beam case. 116
- 4.5 The effect of atomic saturation on the reflection signal. LHS: The normalised reflection signal measured for increasing cavity pump rates, taken with (in dark blue) and without (in light blue) repumping light circulating inside the cavity. There appears to be no significant difference between the two and fits to Eq.(4.5) confirm this. Shown alongside the data are fits to Eq.(4.5), where the atom-cavity coupling rate, g and a factor parameterising the pumping efficiency are free parameters. The fit to the case without repump estimates $g = 2\pi \times (82.4 \pm 1.6)\text{MHz}$. RHS: The normalised reflection signal calculated for pump rates between $0.1 - 500\text{MCounts/sec}$ using the fit parameters found from the LHS data. The effect of saturation is clear, with the reflection signal reducing substantially for pump rates $> 10\text{MCounts/sec}$. The vertical, dashed line indicates the typical upper pumping rate limit used for reflection experiments. 119
- 4.6 Peak fluorescence counts, N_{fl} , measured as a function of pump beam power, P . The measured data is shown alongside a fit to Eq.(4.7) which gives $N_{sat} = 34.8 \pm 1.9$ and $P_{sat} = 50.4 \pm 10.6\mu\text{W}$ 120

- 4.7 The reflection signal as a function of pump beam detuning. LHS: The reflection signal, calculated using Eq.(4.8), shown for three different coupling regimes. (g, γ) are kept constant at $2\pi \times (98, 3)$ MHz respectively, whilst the value of κ is changed. The dark blue shows $\kappa = 2\pi \times 5200$ MHz or $\kappa \gg g$, which is the regime we work in; the medium-toned blue shows $\kappa = 2\pi \times 141$ MHz, or $\kappa = \sqrt{2g^2 - \gamma^2}$ and is the threshold for the strong coupling regime; the light blue shows $\kappa = 2\pi \times 50$ MHz or $\kappa < g$, well into the strong coupling regime. The increased height and width of the lineshape as κ reduces (and thus C_N increases) is visible. RHS: The normalised reflection signal measured as a function of pump beam detuning (shown in units of γ). A fit to Eq.(4.8) is shown alongside the data, which estimates $C_N = 0.39 \pm 0.01$. The inner pair of dashed lines indicate the bare-atom linewidth of γ . The outer pair of dashed lines indicate the Purcell-enhanced linewidth of $2C_N\gamma$ 123
- 4.8 The fluorescence signal as a function of pump beam detuning. LHS: The calculated fluorescence signal plotted for three coupling regimes, using Eq.(4.9). (g, γ, Ω) are kept constant at $2\pi \times (98, 3, 10)$ MHz respectively, whilst the value of κ is changed. The dark red shows $\kappa = 5200$ MHz or $\kappa \gg g$, which is the regime we work in; the medium-toned red shows $\kappa = 141$ MHz, or $\kappa = \sqrt{2g^2 - \gamma^2}$ and is the threshold for the strong coupling regime; the light red shows $\kappa = 50$ MHz or $\kappa < g$, well into the strong coupling regime. The increased height and width of the lineshape as κ reduces (and thus C_N increases) is visible. RHS: The average peak fluorescence counts, N_{fl} , collected at the APD in a $50 \mu\text{s}$ time bin versus detuning, Δ_a of the pumping beam. The fit shown is to Eq.(4.9), which estimates $C_N = 1.01 \pm 0.16$ 125
- 4.9 The reflection and fluorescence signals as a function of N_{eff} . LHS: The reflection spectral width as a function of atom number density at the cavity. There appears to be a linear dependence and a linear fit shown alongside the data gives a y -intercept of 1.09 ± 0.04 and a gradient of $(0.06 \pm 0.01) \times 10^8 \text{ cm}^3$. RHS: The fluorescence counts measured in $5 \mu\text{s}$ as a function of N_{eff} , calibrated using reflection signals. A fit to the function $y = ax^b$, shown alongside the data, gives an estimate of $b = 1.05 \pm 0.06$, confirming the linear dependence of fluorescence signal with N_{eff} in the regime of low atom number. 127

- 4.10 The fluorescence counts as a function of pump beam polarisation. LHS: The peak fluorescence counts recorded at the APD in $50 \mu s$, measured as a function of waveplate angle. The data show a sinusoidal variation. Fitting to a sinusoidal function of fixed period 180° estimates a peak/trough contrast of 1.81 ± 0.14 , which compares well to master equation simulations discussed in the main text. It should be noted that the zero in degrees is arbitrary and just corresponds to $2\times$ the waveplate angle. RHS: Fluorescence lineshapes taken at polarisation angles corresponding to the maximum (lighter red) and minimum (darker red) in fluorescence counts (290° and 200° respectively). Shown alongside the data are the respective fits to Eq.(4.9). For the ‘max’ data, this fit estimates $C_N = 1.01 \pm 0.16$. For the ‘min data’ the fit estimates $C_N = 0.77 \pm 0.24$. Dashed lines indicate the FWHM for each dataset; the outer pair show the FWHM for the ‘max’ data, the middle pair the FWHM for the ‘min’ data and the inner pair for the atom in free space. 129
- 4.11 The excited state populations as a fraction of total population for the magnetic hyperfine levels $m_F = -3, -2, -1, 0, 1, 2, 3$, calculated from the master equation simulation discussed in the main text. LHS: Fractional populations in the case of pumping by $\sigma^\pm$ -polarised light. The atomic population builds up in the higher magnitude m_F levels as transitions with a net angular momentum change of ± 1 are favoured. RHS: Fractional populations in the case of pumping by π -polarised light. Population builds up in the lower magnitude m_F states as transitions with no net angular momentum change are favoured. 131
- 4.12 Comparison of the signal variance and mean. LHS: The variance/mean as a function of time for the reflection signal, calculated for 300 shots. The white line shows the variance/mean calculated for groups of integration time-bins, 200 long. RHS: The variance versus the mean for the fluorescence signal, calculated for 250 shots. The linear fit with zero intercept shown alongside the data produces a gradient of 1.04 ± 0.01 134

- 4.13 The average drop signal in reflection and calculated SNR. LHS: The average of 300 reflection experiments. The white line shows a Gaussian fit to the average, from which the peak normalised reflection signal is 2.00 ± 0.01 lab units, equivalent to an $N_{\text{eff}} = 1.1$, and the FWHM is 8.6 ms. RHS: The SNR as a function of integration time, calculated for the average drop signal shown on the LHS. The blue data points show calculations of the SNR from the data and using Eq.(4.11). The solid and dashed lines are theoretical calculations of the SNR as a function of τ_{int} . The solid line takes into account the Gaussian time dependence of the signal, whilst the dashed line does not. From this, the SNR rises quickly, equalling 1 after $37 \mu\text{s}$ and peaking at 10 after 6.65 ms. This latter roll-over occurs when τ_{int} equals $1.5 \times$ the FWHM of the Gaussian signal. . . . 136
- 4.14 Using Poisson statistics to calculate the detection fidelity. LHS: The distribution of counts per $2 \mu\text{s}$ integration time bin for the signal counts (lighter blue) and the background counts (darker blue). Fitting these distributions to a Poisson distribution gives estimates of the respective means as $\bar{x}_{\text{sig}}^{\text{fit}} = 0.3008 \pm 0.0002$ and $\bar{x}_{\text{bg}}^{\text{fit}} = 0.1621 \pm 0.0003$. RHS: The joint probability for detecting a signal count level $\geq u$ and a background count level $< u$ as a function of integration time. Three values of u are shown; $u = 3$ is shown in the darkest blue; $u = 2$ is shown in the medium-toned blue; and $u = 1$ in the lightest blue. As u increases the peak fidelity decreases and occurs for larger τ_{int} 138
- 4.15 The range of N_{eff} which can be discriminated by our atom detector. The solid lines show the differentiated signal with respect to N_{eff} as a function of N_{eff} , for two different pumping rates: 810 kCounts/sec (in lighter blue) and 81 kCounts/sec (in darker blue). The dashed lines show the signal fluctuations as a function of N_{eff} . Where the solid and dashed line of each colour meet, shown by the solid points, indicates the upper limit of the detection range. For the higher pumping rate this occurs at $N_{\text{eff}} = 26.2$ and for the lower pumping rate this occurs at $N_{\text{eff}} = 14.1$ 139

- 4.16 The SNR as a function of integration time, calculated for a fluorescence signal averaged over 250 drops. The period of τ_{int} always begins for the time bin after the pump beam is switched on. The dark red data points are discrete calculations of the SNR from the averaged data, calculated by summing over increasing numbers of time bins. The solid line is the SNR, calculated assuming an exponential signal, with parameters found by fitting to the average signal. It shows good agreement to the data points. The dashed line is the SNR calculated assuming a constant, peak signal — this diverges from the solid line and data points once $\tau_{int} \geq 20 \mu s$. SNRs > 1 are easily achievable, with the SNR equalling 1 after only $4 \mu s$ 139
- 4.17 Calculating the detection fidelity using the fluorescence signal. LHS: The distribution of signal counts, shown in the lighter red, and background counts, shown in the darker red. Unlike reflection, the count distributions are quite different. RHS: The joint probability to detect $N_{sig} \geq u$ AND $N_{bg} < u$ as a function of τ_{int} . Three threshold values are shown; $u = 1$ is shown in the darkest red; $u = 2$ is shown in the medium-toned red; and $u = 3$ in the lightest red. As u increases so too does the fidelity of the measurement, although a corollary of this is having to wait longer. A fidelity of 95 % can be achieved after $10.3 \mu s$ if we use $u = 1$ counts as the counting threshold. This can be increased to 99 % after $15.6 \mu s$ using $u = 2$ and even further increased to 99.7 % after $22.6 \mu s$ using $u = 3$ 141
- 4.18 Box-car analysis of simulations of the signal from a single atom passing through the cavity. LHS: An example input (bottom) and output (top) trace used to determine the optimal τ_{int} . The input trace is $250 \mu s$ long, in $50,000 \times 5$ ns time bins. The counts across the trace have a Poissonian distribution, with a mean of 0.0025 counts/bin. The output trace is calculated, per bin, using Eq.(4.13). RHS: The same output traces, but now re-binned with $\tau_{int} = 500$ ns (top) and $\tau_{int} = 10 \mu s$ (bottom). 145
- 4.19 The SNR of the reflection signal as a function of τ_{int} , produced using simulations of atom traces as detailed in the main text. A SNR optimum of ≈ 2.1 occurs at $\tau_{int} = 20 \mu s$, indicating that this is the best value of τ_{int} to use. 145

- 4.20 The probability of an atom having a given peak coupling rate, calculated by randomly assigning coordinates inside the cavity mode to one atom, one million times. The bounds for each coordinate are set by the cavity length and cavity waist. LHS: The probability of occurrence of a given peak coupling rate, g . The coupling rates are grouped into 1 MHz wide bins spanning a range $-98 \rightarrow 98$ MHz. RHS: The probability of occurrence of a given peak C . The peak cooperativities are grouped into 0.005 wide bins spanning a range $0 \rightarrow 0.31$. . . 146
- 4.21 The SNR of the reflection signal calculated as a function of peak atom cooperativity, assuming $\tau_{int} = 20 \mu s$. This indicates that a threshold of $C \geq 0.13$ produces an $SNR \geq 1$ 147
- 4.22 The normalised reflection signal calculated as a function of input counting rate, for three values of the cooperativity; $C = 0.1$ in dark blue, $C = 0.2$ in the medium-toned blue and $C = 0.3$ in light blue. For all three cases it appears as though saturation begins to take place at ~ 50 MCounts/sec, but the larger the cooperativity the larger the rate of decrease in the signal. 148
- 4.23 Detecting a single falling atom with the fluorescence signal (1). LHS: The signal to noise ratio as a function of integration time, calculated for a beam power of $170 \mu W$. Two peak cooperativities are shown; $C = 0.3$ in light red and $C = 0.1$ in dark red. The dashed line indicates the average time the atom spends in the cavity before being pushed out by the fluorescence beam. This is where the SNR peaks, at $SNR = 1.01$ and $SNR = 0.57$, respectively. RHS: The peak signal to noise ratio as a function beam power. The same two cases as for the LHS are shown here. For $C = 0.3$ shown in light red, a peak of $SNR = 1.01$ occurs for a beam power of $180 \mu W$ - although this value varies by only 1% between $80 - 350 \mu W$. For $C = 0.1$ shown in dark red, the peak is $SNR = 0.57$ and has a similarly large range over which it barely changes. 150
- 4.24 Peak fluorescence detection SNR as a function of cooperativity, calculated assuming a pump beam power of $200 \mu W$. This indicates that atoms with $C \geq 0.29$ will produce a signal with $SNR \geq 1$ 151

- 4.25 Detecting a single trapped atom in reflection (1). LHS: The mean lifetime, τ_{mean} for an atom inside a dipole trap in the cavity, calculated as a function of J_{in} . The primary cause of atom loss out of the dipole trap is assumed to be from scattering cavity pumping light. Its is evident that τ_{mean} is inversely proportional to J_{in} ; for $J_{in} = 1$ MCounts/sec, $\tau_{mean} = 6.7$ ms, whereas for $J_{in} = 0.5$ MCounts/sec, $\tau_{mean} = 13.3$ ms. RHS: The reflection measurement SNR as a function of τ_{int} for $J_{in} = 0.25$ MCounts/sec, so that $\tau_{mean} = 26.6$ ms. As expected, the peak of SNR = 64 occurs when $\tau_{int} = \tau_{mean}$ 152
- 4.26 Detecting a single trapped atom in reflection (2). LHS: The highest achievable SNR as a function of J_{in} . τ_{mean} is calculated for each value of J_{in} and it is assumed that the highest achievable SNR occurs when $\tau_{int} = \tau_{mean}$. Higher J_{in} results in higher values of the SNR, until $J_{in} \sim 1$ MCounts/sec, when the SNR begins to saturate at ~ 75 . This is because the constant dipole trap light contribution to N_{bg} becomes negligible in comparison to that from the cavity pumping light and, once in this regime, any increase in N_{sig} is counteracted by a corresponding increase in N_{bg} . RHS: The minimum value of τ_{int} required to achieve an SNR ≥ 1 , calculated as a function of J_{in} . As J_{in} increases, this minimum time reduces, enabling faster detection. A corollary of this is that the atom spends less time in the trap. 153
- 4.27 Detecting a single trapped atom using the fluorescence signal (1). LHS: The mean lifetime for an atom inside the dipole trap when being detected using the fluorescence signal. The atom's scattering rate increases for higher pump beam powers, so that the time taken to scatter sufficient photons to be lost from the trap reduces. For a beam power of $100 \mu\text{W}$, the mean lifetime is $56 \mu\text{s}$. The curvature is caused by saturation of the atomic scattering rate. RHS: The SNR as a function of τ_{int} , calculated for a beam power of $100 \mu\text{W}$. The dashed line indicates the average lifetime for an atom in the trap of $56 \mu\text{s}$. This is the point at which the SNR peaks, at a value of 2.2. 155

- 4.28 Detecting a trapped atom using the fluorescence signal (2). LHS: The SNR as a function of beam power, calculated assuming $\tau_{int} = \tau_{mean}$. As for the case in reflection, the SNR increases with increasing pump beam power, but when the power $\geq 100 \mu\text{W}$ the SNR saturates. The saturated SNR is ~ 2.3 . RHS: The minimum τ_{int} required to achieve an $\text{SNR} \geq 1$, calculated as a function of pump beam power. The minimum required τ_{int} decreases with increasing beam power. However, the rate at which it reduces saturates quickly; only $7 \mu\text{W}$ of power are needed for a minimum τ_{int} of $10 \mu\text{s}$, but reducing this further to $1 \mu\text{s}$ requires beam powers of several watts. 155
- 5.1 The counts measured at the APD averaged over ten shots of a push beam experiment. The reflection signal resulting from pushed atoms is clearly visible as the tall, narrow feature, centred at approximately 35 ms. The residual reflection signal from falling atoms is centred at approximately 75 ms. The relative signal heights are 2.70 ± 0.09 for the pushed atoms and 1.75 ± 0.02 for the dropped atoms (both in lab units), corresponding to collective cooperativities of 1.75 ± 0.39 and 0.31 ± 0.01 respectively. This implies that the pushed signal corresponds to $N_{\text{eff}} = 5.5$ atoms. 161
- 5.2 Example experiment traces using the multi-push method. TOP LHS: The reflection signal, showing 30 pushes out of the MOT at a repetition rate of 2 kHz. Each pulse contains two peaks; one is scattered light from the push beam itself and the other is the reflection signal caused by atoms. This latter signal decays exponentially with time, as indicated by the decreasing, dashed line. The horizontal, dashed line indicates the constant level of the scattered light pulses. TOP RIGHT: Zoomed-in version of the trace on the left, from 32 ms to 54 ms. At this scale it is possible to clearly distinguish between the scattered light pulses and the reflection signal pulses. BOTTOM: The normalised reflection signal in lab units for the data shown on the LHS as a function of pulse number. This demonstrates the utility of this method, which allows us to test amplitudes ranging from $\sim 2.5 - 1$ in one experimental run. 162

- 5.3 Example trace of a multi-push experiment with a sequence alternating between reflection and fluorescence measurements. The trace shows 40 pushes; 20 fluorescence pulses interspersed with 20 reflection pulses. To collect fluorescence the cavity drive beam is shuttered off from the cavity, explaining the alternating height of the background. 163
- 5.4 Characterising the effect of the push beam. LHS: The normalised reflection signal for different durations of the push beam. Shown alongside the data is a saturation fit, which estimates the fully saturated reflection signal to be 3.04 ± 0.33 . RHS: Fluorescence lineshapes taken for atoms pushed for $300 \mu\text{s}$ (in lighter red) and $100 \mu\text{s}$ (in darker red), alongside Lorentzian fits to the datasets. Both lineshapes are red-shifted with respect to the fluorescence beam (which is located below the MOT and cavity), with their lineshapes shifted by several MHz. As might be expected, the effect is more pronounced for longer durations. The fits produce central frequencies which are red-shifted from resonance by $-9.82 \pm 0.86 \text{ MHz}$ and $-30.79 \pm 1.06 \text{ MHz}$ for $100 \mu\text{s}$ and $300 \mu\text{s}$, respectively. These correspond to velocities of $7.67 \pm 0.67 \text{ m s}^{-1}$ and $24.0 \pm 0.83 \text{ m s}^{-1}$, respectively. 165
- 5.5 The normalised reflection signal with cavity pump beam detuning. The lighter blue dataset corresponds to a detuning scan across the first pulse in a multi-push experiment and the darker blue dataset corresponds to the same scan for the seventeenth pulse. Lorentzian fits to each dataset are shown alongside, producing estimates of $C_N = 3.55 \pm 0.11$ and $C_N = 0.40 \pm 0.04$, respectively. 167
- 5.6 Comparison of the different methods for estimating C_N . LHS: A graph of ‘ C_N from width’ versus ‘ C_N from height’, where the dropped atom data is shown in dark blue and the pushed atom data is shown in light blue. Ideally, these two methods should produce identical results and all the data points would lie on a straight line of unity gradient, as indicated by the dashed line. The data do broadly follow this pattern, with some clear exceptions caused by failings in the fitting procedure, discussed in the main text. RHS: A graph of ‘ C_N from width’ versus ‘ C_N from height’, this time calculated using ‘synthetic’ data produced by randomly distributing atoms inside the cavity mode. Some of the features of the real data are replicated here, such as the scatter about the unity gradient line for higher C_N and, to a lesser extent, the stratum of data at very low C_N 169

- 5.7 A comparison of the ‘known C_N ’ with the ‘measured C_N ’ for the synthetic data, calculated using the single-parameter fitting method. The stratum at low C_N has been eliminated and the scatter about the unity-gradient line is substantially reduced. 169
- 5.8 The normalised reflection signal as a function of collective cooperativity, C_N , calculated using the single parameter fitting method. The dropped atom data are shown in dark blue circles, whilst the pushed atom data are shown in light blue squares. Fits to Eq.(5.2) are shown alongside. The discrepancies between the data are caused by differences in fringe visibility. 171
- 5.9 Simulations of the behaviour of the fluorescence signal with increasing N_{eff} . LHS: Comparing the effects of having a uniform Ω for all atoms (in light red) to the real case, where Ω_j varies with the position of the j th atom (in dark red). The non-uniform case is simply reduced by $\sim 1/2$, equivalent to the ratio of mean-to-peak intensity, with the rollover position unchanged. RHS: The effect of detuning the driving beam, with a position-dependent Ω_j . Four detunings are shown, with darker to lighter red corresponding to $\Delta_a = 2\pi \times (0, 3, 6, 9)$ MHz, respectively. As the detuning increases, the rollover appears to shift to higher values of N_{eff} 172
- 5.10 The fluorescence signal with increasing N_{eff} . LHS: Experimental data showing the fluorescence signal measured for different N_{eff} (bottom horizontal axis) and C_N (top horizontal axis). The dashed line at $N_{\text{eff}} = 1/2C$ indicates where the rollover in fluorescence signal amplitude would be expected to occur. The data appear to increase linearly with N_{eff} at first, consistent with the ‘disordered scattering’ limit - but no rollover is apparent. RHS: Comparison of simulated data (the dark red line) and experimental data from the LHS (the light red points). The simulated data assumes a drive beam Rabi frequency of $2\pi \times 10$ MHz and a detuning of $2\pi \times 5.4$ MHz. There is a good match between them, indicating that detuning may well be the reason for the lack of rollover in the signal. . . . 173

- 5.11 The reflection (LHS) and fluorescence (RHS) signals, calculated as functions of detuning, Δ , where $\Delta = \Delta_c = \Delta_a$. These are shown for three different values of N_{eff} , with the cavity parameters kept constant at $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz respectively. TOP: $N_{\text{eff}} = 1$. For the reflection signal, the Lorentzian cavity profile is visible, with a small perturbing feature caused by the presence of an atom, centred at $\Delta = 0$. In contrast, the fluorescence signal is a Purcell-broadened Lorentzian, but shows no other features caused by being coupled to the cavity. MIDDLE: $N_{\text{eff}} = 1.6$, which corresponds to $N_{\text{eff}}g^2/\kappa\gamma = 1$. The perturbation at the centre of the reflection signal has increased, but now the fluorescence signal is qualitatively different, with the lineshape flattened in the region of $\Delta = 0$ and the second derivative of the signal equalling zero when $\Delta = 0$. BOTTOM: $N_{\text{eff}} = 80$. The coupling is now substantially larger, although it is not yet into the strong coupling regime, where $g\sqrt{N_{\text{eff}}} \gg (\kappa, \gamma)$. The reflection and fluorescence signals now both show a clearly resolved splitting. 176
- 5.12 The two-dimensional transmission spectrum of our cavity, with and without an atom. For each plot, Δ_a is given by the horizontal axis and Δ_c is given by the vertical axis. LHS: The separate transmission spectra for the cavity (top) and the atom (bottom). Each show a narrow band of frequencies when light is either transmitted (by the cavity) or absorbed (by the atom). MIDDLE: The transmission spectrum of light passing consecutively through an atom and a cavity. The cavity resonance is split where the atom absorbs strongly, but there is no dispersive effect as the atom and cavity do not interact. RHS: The transmission spectrum from a coupled atom-cavity system. The cavity resonance is still split, but now the dispersive effect of their interaction causes the resonance to shift in frequency. The white, dashed line indicates the value of Δ_c corresponding to the maximum in transmission as one sweeps Δ_a 178
- 5.13 Two dimensional scans of the cavity reflection signal with respect to the cavity and pump beam detuning. The LHS of each plot shows real data, the RHS shows theory plots to Eq.(5.5). Three different values of N_{eff} are shown. TOP: $N_{\text{eff}} = 0$. MIDDLE: $N_{\text{eff}} = 1.1$. BOTTOM: $N_{\text{eff}} = 7.1$ 180

5.14	The value of Δ_c at which the reflection signal is minimised, plotted as a function of Δ_a . Each plot corresponds to a contour plot shown above; darker to lighter blue gives of $N_{\text{eff}} = 0, 1$ and 7 . The dispersion clearly increases with larger coupling.	181
5.15	Two-dimensional scans of the fluorescence signal. TOP LEFT: The counts recorded at the APD as a function of pump beam detuning, Δ_a , and cavity detuning, Δ_c , integrated over $10 \mu\text{s}$, or one time bin. TOP RIGHT: The same as on the LHS, but now integrated over $100 \mu\text{s}$, or ten integration time bins. BOTTOM: A theory plot of the same two-dimensional scan, calculated for our cavity parameters and $N_{\text{eff}} = 1.1$ and using Eq.(5.5).	183
5.16	Theory plots showing two-dimensional scans of fluorescence signal, calculated for our cavity parameters and using Eq.(5.5). LHS: For $N_{\text{eff}} = 1.6$. The atom-cavity interaction has sufficiently distorted the Lorentzian peak along the line of $\Delta_a = \Delta_c$ to produce the flat-topped function shown earlier. RHS: For $N_{\text{eff}} = 7$. The atom-cavity interaction is now sufficient to split the resonance, becoming qualitatively similar to the reflection signal.	185
5.17	The recorded APD counts measured as a function of detuning. The counts are averaged over 50 integration bins, equal to $500 \mu\text{s}$. The shift of the atomic resonance with respect to the pump beam frequency is equal to $\sim 2.5\gamma$	186
5.18	The shift of the atomic resonance from $\Delta_a = 0$ as a function of averaging time. The value for each shift was produced by a making a Lorentzian fit along the line of $\Delta_c = 0$ for each averaging time. On timescales of $100 \mu\text{s}$ or less, the shift remains constant at $\gamma/2$, likely due to a mis-estimate of where zero detuning actually was. For timescales longer than $100 \mu\text{s}$, the shift increases and quickly tends towards a limiting value of 2.9γ	186
6.1	The dipole trap depth and atomic scattering rate as a function of trap wavelength, calculated using Eq.(6.2) and assuming an input power of $500 \mu\text{W}$. The trap depth, shown in the lighter green, is given by U_{dip}/k_B and for red-detuned light gives a negative potential, attracting the atoms to regions of high intensity. The scattering rate, Γ_{scatt} , shown in the darker green shows two poles in Γ_{scatt} , corresponding to the D_2 and D_1 transitions.	189

- 6.2 The Stark shift caused by increased trapping light circulating inside the cavity. LHS: Atomic lineshapes measured for different intra-cavity trap powers. The detuning of the pump beam is varied with respect to the cavity resonance. Lighter to darker blue correspond to lower to higher trap powers. The shift of the resonance is visible, but is accompanied by a reduction in the amplitude and an increase of the width as a result of different atoms being Stark-shifted by different amounts. RHS: A comparison of the increase in measured Stark shift as a function of the increase in power, taken from the lineshapes centres on the LHS. The increase in power in the dark blue dataset is estimated by measuring the reflected trap power and using this to estimate a cavity pumping rate. The increase in power in the light blue dataset is estimated from the counts at the APD when the cavity is only pumped with trap light. For both, the Stark shift increases with increasing power, as would be expected. 191
- 6.3 The three-level ‘lambda’ system, consisting of levels $|1\rangle$, $|2\rangle$ and $|3\rangle$. The transitions are addressed by two coherent laser fields; Ω_c and Ω_p , with respective detunings $\Delta_p = \omega_{31} - \omega_p$ and $\Delta_c = \omega_{32} - \omega_c$ 195
- 6.4 The normalised reflection signal in lab units, averaged for ten shots. Initially, atoms entering the cavity are in the $|F = 1\rangle$ state, caused by switching off the repumping light during the optical molasses phase of the experiment. In this state they will not interact with the light pumping the cavity. At $t = 69\text{ms}$, indicated by the vertical dashed line, the repumping light is switched back on. This pumps the atoms into the $|F = 2\rangle$ state, where they can be detected through their interaction with the pumping light. This is evidenced by the sudden retrieval of the drop signal. 196

- 6.5 Electromagnetically induced transparency in an Rb vapour cell. LHS: The transmission profile of the probe beam is shown as Δ_p is tuned across the ^{87}Rb $|F = 0\rangle \rightarrow |F = 0, 1, 2\rangle$ manifold. The darkest green shows the Doppler-broadened absorption profile. The medium-toned green shows the saturated absorption spectrum of the same transitions, where now the probe beam is retro-reflected through the cell. The lightest green shows the same saturated absorption spectrum, but now with the co-propagating coupling laser present. The narrow EIT feature is clearly visible at zero detuning. RHS: A ‘zoomed in’ version of the EIT feature. Shown alongside the scan is the fitted Lorentzian lineshape, which estimates the HWHM to be 357 ± 4 kHz. 198
- A.1 The fluorescence signal, J_{fl} as a function of transverse pumping rate Ω , where $J_{fl} = 2\kappa|\alpha_F|^2$ and α_F calculated using a numerical solution to Eq.(A.3). The cavity parameters used were $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz and it is assumed that for all atoms, $g_j = g$ and $\Omega_j = \Omega$. LHS: $N_{\text{eff}} = 1$. At this value, J_{fl} increases steadily and then smoothly rolls over when $\Omega = \sqrt{2}\gamma$, from whence it monotonically decreases - leading to a dark cavity when the atoms are fully saturated. RHS: The case of $N_{\text{eff}} = 13$, which corresponds to $C_N = 4$. This shows similar behaviour, with a smooth roll over in J_{fl} , then decreasing to a dark cavity for fully saturated atoms. However, when $\Omega = \sqrt{2N_{\text{eff}}}\gamma$, J_{fl} is multi-valued, showing optical bistability in the cavity. 216
- A.2 The results of 10 simulations calculating the intra-cavity photon number as a function of the saturation parameter, I/I_{sat} . For each simulation, three atoms are assigned random positions throughout the cavity volume, such that $N_{\text{eff}} = 0.2$. Eq.(A.2) is then solved, given each atom’s coupling rate, g_j and beam intensity, I_j . Whilst each ‘throw’ of atoms is different, corresponding to a different configuration of atoms, the general features for each simulation are the same; $|\alpha_F|^2$ initially increases with increasing I/I_{sat} ; then reaches a maximum, which covers several orders of magnitude in I/I_{sat} ; and then rolls over so that at as $I/I_{\text{sat}} \rightarrow \infty$, $|\alpha_F|^2 \rightarrow 0$. This is qualitatively similar to the case of purely coherent scattering. The dashed line represents the upper limit of $|\alpha_F|^2$, which occurs at $|\alpha_F|^2 = N_{\text{eff}}C\gamma/\kappa$ 217

List of Tables

3.1	A summary of the light needs on the experiment.	91
4.1	Summary of the key detection parameters for when $N_{\text{eff}} = 1$ for the reflection and fluorescence signals.	143
4.2	Summary of the results from analysing the prospect for detecting a single, trapped atom in our cavity using the reflection and fluorescence signals.	156

Chapter 1

Introduction

1.1 Cavity QED

The interaction between matter and light is of fundamental interest. One of its most basic manifestations occurs when a single, two-level atom interacts with a single mode, or photon, of the electromagnetic field. This can be described by the Jaynes-Cummings model [1]. In the absence of any mutual interaction, the atom is either in its ground state, $|g\rangle$ or excited state, $|e\rangle$ and the field can either have photon number $|0\rangle$ or $|1\rangle$. In the Jaynes-Cummings model, when these entities are coupled by an interaction of strength, or ‘Rabi’ frequency, g , the system can be described in terms of two ‘dressed states’; $|g, 1\rangle$ and $|e, 0\rangle$. In this way, a single excitation is shared between the atom and the field and, as the coupled system evolves coherently, this excitation oscillates between the dressed states at a rate, $2g$. A more formal discussion of this model is presented later, but for now it is important to note that the timescale on which this interaction occurs, and any effects arising from it, will be determined by the size of g .

This model has some limitations. For one, it treats both the atom and field as simple two-state objects. Firstly, an atom (which is not hydrogen) is really a complex, multi-electron system; even without external forces it has fine and hyperfine structure above and beyond its basic transitions between electronic energy states, see for example [2]. In reality, however, one can instead address a quasi two-level system found in, for example, alkali atoms such as cesium and rubidium. For such atoms, a single ‘outer’ electron enables an approximation to a simplified level scheme which can contain atomic ‘cycling transitions’ connecting two hyperfine states. On these transitions, an electron can be multiply excited and de-excited with only a small probability of coupling to other states. The two states of the cycling transition can thus be treated as the ground, $|g\rangle$ and excited, $|e\rangle$ states of the atomic two-level system.

Secondly, an atom in free space does not generally couple to a single mode of the field, but rather a continuous spectrum of electromagnetic modes, causing its excited state population to irreversibly decay at a rate 2γ , see for example [3]. By tailoring the atom's surrounding environment, however, one can create circumstances where the effects of a single quanta of excitation are observable. From Fermi's Golden Rule we know that the atom's decay rate is dependent on the density of electromagnetic states and this can be substantially altered by the presence of an optical cavity. The cavity, formed by two mirrors separated by a distance, d and with a mode volume, V , will only allow a discrete set of spatial, polarisation and frequency modes to resonate and one can hope to arrange it so that an atom placed inside it will only interact strongly with one such mode, with strength g . When $g \gg \gamma$, a photon in the cavity mode will have a larger influence over the atom than the other modes in the 4π steradians of free space. This modification of the atom's decay rate is known as the Purcell effect [4] and the regime within which it occurs is that of cavity quantum electrodynamics [5].

Critical then to achieving a system which accurately approximates the Jaynes-Cummings model is a large Rabi frequency between the cavity mode and the atom, given by:

$$g = \sqrt{\frac{\omega_{eg}|\mu_{eg}|^2}{2\epsilon_0\hbar V}} \quad (1.1)$$

In the above, ω_{eg} and μ_{eg} are the frequency and matrix element of the atomic transition, respectively; ϵ_0 is the permittivity of free space; and $\hbar$ is Planck's constant. From this, the rate is proportional to the energy density of the cavity mode, so that experiments looking to achieve a large g must work with small mode volumes. For the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition of the D₂ line of ⁸⁷Rb discussed in this thesis, $\gamma = 2\pi \times 3.03$ MHz, so that the critical mode volume, for which $g = \gamma$, occurs when $V = 2300 \times 10^{-15} \text{ m}^3$. Typically, experiments working in the optical regime have cavities with $V \sim (3 - 300) \times 10^{-15} \text{ m}^3$.

Besides its simplified treatment of the atom and the field, the Jaynes-Cummings model also ignores any losses which would lead to decoherence. One such decoherence mechanism already discussed is the spontaneous decay of the atom, at a rate 2γ . Loss of photons from the cavity mode caused by transmission, absorption or scattering at the cavity mirrors has a similar decohering effect. The rate at which this occurs is given by 2κ , motivating the use of high finesse cavities to increase the lifetime of photons inside the cavity. Improvements in the mirror fabrication process have enabled the development of smooth, high-reflectivity and low-loss mirrors, with cavity finesse typically scanning ranges $100 - 10^6$ in the optical regime and reaching as high as $\sim 10 \times 10^9$ for microwave cavities [6].




Figure 1.1: Schematic of a coupled atom-cavity system. The system is described by three rates; the rate of coherent coupling between the atom and the field, g ; the spontaneous emission rate of the atom, γ ; and the decay rate of the cavity field, κ .

Figure 1.1 shows a schematic of a cavity QED system, with the three main rates; g , γ and κ indicated. The figure of merit in this system is the size of g , compared to the damping rates; γ and κ . Ideally, one needs a g such that a single photon saturates the atomic response and, conversely, that a single atom has a profound effect on the field.

Considering the effect of a cavity photon on the atom first, an atom will be saturated when it is driven by a field with a Rabi frequency equal to $\sqrt{2}\gamma$. As will be shown later, in a cavity with n photons, the Rabi frequency is given by $2g\sqrt{n}$. Thus, the number of photons required to saturate the atom is known as the *saturation photon number*, n_{ph} , given by:

$$n_{ph} = \frac{\gamma^2}{2g^2}$$

The atomic influence on the cavity field can be quantified by considering a group of N atoms inside the cavity mode, each with a resonant absorption cross section given by σ_0 , see for example [7]. If the cavity supports m reflections of the photon from the internal mirrors faces, in the limit of low saturation the fraction of intra-cavity excitation which will be absorbed by the atoms is given by $Nm\sigma_0/A$, where A is the cross sectional area of the cavity mode. Demanding that this equals unity gives $N = A/m\sigma_0$, where now $N = N_{crit}$, the *critical atom number*. We can make the approximation that $A \sim V/d$ and $\sigma_0 \sim \lambda^2$. In addition, m will be given by the product of the photon round-trip time and the loss rate; $m = (c/2d)1/\kappa$. Substituting in these expressions gives $N = 2V/\lambda^2 c$. Using Eq.(1.1) to replace V gives:

$$N_{crit} \sim \frac{\kappa}{g_0^2} \left(\frac{|\mu_{eg}|^2}{\hbar \epsilon_0 \lambda^3} \right) = \frac{\kappa \gamma}{g^2}$$

The values of N_{crit} and n_{ph} give thresholds for when we might expect to see quantum phenomena arising from the atom-cavity interaction. For $n_{ph} < 1$, or, equivalently, $g \gg \gamma$, the interaction between a single atom and a single photon inside the cavity will cause nonlinear optical effects. When $N_{crit} < 1$, or $g^2/\kappa\gamma \gg 1$, a single atom will profoundly affect the cavity field, enabling its detection. These regimes have become increasingly accessible to experimentalists in the last two decades, with astounding results.

1.2 Realisations of cavity QED

There are many different physical implementations of cavity QED; including using Rydberg atoms traversing microwave cavities [8], trapped ions [9, 10] and neutral atoms in optical cavities [11]. Outside of these branches of physics, cavity QED has been explored in solid state environments, for example using quantum dots embedded within semiconductor micropillars [12] and photonic crystal nanocavities [13].

Cavity QED experiments using atoms fall into two distinct categories: those using Rydberg atoms interacting inside superconducting microwave cavities and those using low-lying atoms interacting inside optical cavities. Initially, due to technical limitations, many early cavity QED experiments were of the former type [14, 15], as the atomic and field lifetimes in the Rydberg systems are on the order of several microseconds, substantially lowering the threshold g required to reach strong coupling. The microwave cavities used are larger, on the order of several millimetres in length, although working with this type of radiation means that the cavity must be cryogenically cooled to limit the effect of thermal photons. Atoms are delivered into the cavity using an atomic beam directed through the mode centre, the timescales being such that the transit time for an atom is substantially smaller than the cavity field decay rate, so that consecutive atoms effectively sample the same field. In this way, effects resulting from the interaction of the cavity mode with, on average, small numbers of atoms could be measured [16, 17]. Whilst the measured effects are not the result of the interaction of one-and-the-same atom, a feature of these systems is that the atom-field interaction takes place slowly, making it possible to study entanglement and produce tailored quantum states [8].

Achieving strong coupling in the optical regime is more challenging, firstly because the decay rates are much faster (e.g. tens of nanoseconds), requiring a larger g and setting a limit on the dimensions of the cavity; optical cavities are typically only a few hundred microns long. A corollary of this is that the interaction between the atom and the cavity takes place on a faster timescale. With improvements in the production of cavity mirrors, the regime of

strong coupling became possible to reach in the optical domain [18, 19]. Like their Rydberg counterparts, these experiments used atomic beams, although here the transit time became a limiting factor as it necessarily resulted in interaction times smaller than the equivalent time in which the atom-cavity interaction took place. In addition, these experiments are stochastic, with fluctuations in the number of atoms in the cavity.

The introduction and development of laser cooling of neutral atoms, see for example [20], mitigated these factors, making it possible to trap and cool atoms to sub-milliKelvin temperatures in a magneto-optical trap (MOT) [21]. By doing so, experiments could be performed in both the strong coupling regime and a regime where the atomic kinetic energy is smaller than the atom-cavity interaction energy, enabling investigations of the motional dynamics of the atom inside the cavity. In addition, these experiments involved detecting the interaction between a single atom and the cavity field in real time. Typically, atoms could be loaded into a MOT and dropped through the cavity [22] or pushed into the cavity using an atomic fountain [23, 24].

These sources of colder, slower atoms have also made it possible to trap a single atom inside the cavity using optical dipole traps (see for example, [25]) implemented either as intra-cavity dipole traps [26, 27] or as transverse standing-wave optical dipole traps which are also used to controllably transport atoms into the cavity [28, 29, 30, 31]. These trapping schemes allow experiments to take place on one-and-the-same atom and provide a high degree of localisation and position control. It should be noted, however, that whilst these cold atom systems have developed to enable trapping of an atom inside a cavity for several seconds, experiments using one-and-the-same trapped ion inside a cavity have been performed for durations several orders of magnitude longer, for example in [32], exploiting the tight localisation produced by the use of electrostatic traps.

Optical trapping inside the cavity has also been implemented in tandem with cavity cooling [30, 33] to increase the storage time for atoms inside the cavity, in some cases to several seconds [30]. This cooling mechanism can be considered classically as being caused by the position-dependent change in refractive index [34] caused by the atom moving around the cavity mode. For a blue-detuned cavity this has a cooling effect, caused by the finite response time of the cavity to changes in the atom-cavity coupling and, as a consequence, the intracavity intensity. In addition, the atom-cavity interaction has been used to cool atoms to their lowest ground state of axial motion inside the trap [35].

Whilst the work described above focuses on the interaction between a single atom strongly

coupled to a cavity, there has been work at the other extreme, studying the interaction of high density atom clouds [36, 37] and Bose Einstein Condensates (BECs) [38, 39] inside a cavity. These experiments are able to exploit the square-root proportional increase in coupling with atom number [40] to increase the collective coupling of the atom cloud to the cavity. Such clouds of atoms have shown additional cooling forces caused by self-organisation inside the cavity mode [41, 42]. In the case of the BEC experiments, all the atoms are in an identical quantum state and can be treated as a single ‘super atom’ coupled to the cavity.

The majority of experiments discussed thus far have been ‘macro’ experiments, with standalone atom traps and cavities; but cavity QED has also been implemented using atom chips [38, 43, 44, 45]. Atom chips aim to miniaturise atomic physics experiments by integrating magnetic and electric-field producing elements on devices of typically a few cm^2 [46, 47]. In some experiments, current-carrying wires on atom chips have been used to trap and guide atoms through ‘macro’ cavities [44, 45]. Alternatively, Fabry-Perot type microcavities have been built by integrating optical elements, such as single-mode fibres and micromirrors, onto atom chips [38, 43, 48]. Technical advances are moving towards monolithic, fully integrated devices; a monolithic microcavity achieving high coupling rates with low losses has also been implemented with whispering gallery mode microtoroids [49] and a recent proposal forms a Fabry-Perot microcavity using a waveguide integrated on a silicon substrate [50].

1.3 Experiments and applications of cavity QED

The applications of cavity QED are varied and wide-ranging, but this section focuses on three areas relevant to the optical microcavity experiment discussed in this thesis; single atom detection, coupled atom-cavity spectroscopy and quantum information processing.

1.3.1 Single atom detection

Confidently detecting a single atom is of huge practical importance to many branches of physics. In quantum information processing, for example, fast detection of a single atom enables one to determine the presence of a qubit which could subsequently be used to perform quantum logic operations. In general, one can detect an atom by illuminating it with resonant or near-resonant light and measuring either a dispersion (i.e. a phase shift) or absorption of the light, or by collecting emitted fluorescence.

One difficulty associated with measuring either a dispersion or absorption are that the changes one measures as a result of a single pass of the light by a single atom are small

[51]. Single atom sensitivity is possible, but principally by integrating over long periods [52], requiring the atom to be localised in a trap. Small numbers of atoms have been detected in this way using atom chips; one demonstration measures the change in coupling between two optical fibres mounted on the chip, caused by the presence of ≤ 100 atoms [53] and a more recent experiment measures the change in coupling produced by one atom on average between two waveguides integrated onto the chip [54].

Alternatively, one can collect fluorescence emitted by the atom. This has the inherent difficulty of distinguishing a small signal count rate from a single atom from the background noise. Again, this can be overcome in trapped ion [55] and trapped atom [56, 57] experiments, where fluorescence detection is in standard use, because the tight localisation of the ion or atom combined with trapping times of several milliseconds make it possible to integrate the resulting fluorescence signal for long periods, thereby confidently distinguishing it from the background light collected by the detector.

Detecting non-localised, neutral atoms by collecting fluorescence is more difficult because of the finite duration over which fluorescence is emitted (i.e. the short time that the atom is in the detection region, typically only a few tens of microseconds), which limits the integration time available for detecting the signal. One method for overcoming this uses near- 4π collection mirrors [58] to achieve a high collection efficiency of signal photons, detecting an atom during its $60\mu\text{s}$ transit through the setup with virtually 100% certainty. Alternatively, one can use filtering optics to distinguish signal photons from background counts. One such example involves using two optical fibres mounted at 90 degrees to each other on an atom chip to detect atoms which are magnetically guided between them; one fibre emits light to excite atoms and the multimode fibre collects the fluorescence. This system detects single atoms with a high signal to noise ratio and with an efficiency of 66%. Another proposed implementation involves collecting fluorescence with a miniature aspheric lens and fibre optic mounted onto an atom chip surface [59].

Detection of a single atom by absorption, dispersion and collecting fluorescence can be significantly enhanced using an optical cavity. For absorption and dispersion measurements, the multi-pass effect of the cavity means that a single atom can produce an easily detectable change in the light. In the limit of strong coupling, the transits of single atoms through an optical cavity are clearly discernable as changes in the transmission of the cavity [22, 23, 60], which can be measured in real time. Aside from detecting the presence of a single atom, such experiments can also provide information about the atom's motion inside the cavity [61] and

changes in its internal hyperfine states [31]. In fluorescence, the Purcell effect can be exploited to enhance the collection of the atom's fluorescence. In the strong coupling regime streams of single photons produced by a trapped atom can be collected [62, 63].

Strong coupling isn't, however, a prerequisite for atom detection using optical cavities. An atom resolution of 0.1 has been achieved by measuring changes in cavity transmission with a $g^2/\kappa\gamma < 1$ [44]. Separately, an experiment using a medium finesse cavity surrounding a single atom magnetically trapped on an atomic chip achieved a 75% detection efficiency in 250 μs by collecting atomic fluorescence and single atom resolution by monitoring changes in transmission [45]. Alternatively, coincidence counting can be used to increase the detection fidelity. This has been implemented in an intermediate coupling regime, detecting an atom with more than 99.7% confidence after 1 μs [64].

Optical cavities integrated onto atom chips have also been proposed as a means to confidently detect single atoms [65]. A cavity formed by two fibre-tips mounted on an atom chip was used to detect < 50 rubidium atoms [66] and an etched micro-toroid formed from silicon has been used to detect single falling atoms [49]. The optical microcavity experiment at CCM, discussed throughout this thesis has also demonstrated single atom sensitivity and high detection fidelities [43, 67], using both changes in the reflected light from the cavity and collection of emitted fluorescence.

1.3.2 Spectroscopy of the coupled atom-cavity system and non-linear optics

The spectroscopy of a coupled atom-cavity system can be quite different to that of its constituent parts. When strongly coupled, an atom in the cavity can lead to a resolved splitting of the cavity resonance. This 'normal mode splitting' has been measured for a one atom on average in optical cavities [18, 19] and microwave cavities [17] and has been subsequently repeated for one-and-the-same trapped atoms. The splitting itself is a manifestation of the two dressed states of the system, which show an avoided crossing as the cavity, atom and pumping light are tuned into resonance. This avoided crossing can be measured directly because the excitation spectrum of the atom-cavity system is proportional to the transmission spectrum. It has been measured for a single, trapped, strongly-coupled atom [68] and a BEC [39]. The presence of the cavity also modifies the atom's spectroscopy, leading to a measurable enhancement [69] or inhibition [70] of the atom's spontaneous emission.

As mentioned, these experiments all measure the transmitted light through the cavity. In contrast, the spectroscopy presented here measures the reflected light from the cavity.

1.3.3 Quantum information processing

Cavity QED systems have been proposed as a building block for quantum networks [71]. In such a scheme, a single atom trapped in an optical cavity (a ‘quantum node’) is a means to store a ‘qubit’, with photons transferring qubits between nodes. This relies on being able to efficiently transfer information between the atom and photon and being able to generate single photons on demand.

Substantial progress has been made towards this latter criterion [62, 63, 72] using optical cavities. These implementations generally make use of a ‘dark state’ to transfer the atom from one hyperfine level to another, with the atom excited by a laser field and de-excited by the cavity vacuum field, thereby producing a single photon. Using such methods, streams of several thousand photons from a single, trapped atom have been produced [62]. These photon streams have a manifestly quantum nature, exhibiting sub-Poissonian fluctuations and photon anti-bunching [62, 63]. There has also been success in controlling the amplitude, frequency [73], polarisation [74] and, recently, the phase [75] characteristics of the emitted photons.

A second prerequisite for quantum information processing with cavity QED is the ability to efficiently transfer information between atoms and photons. This has been demonstrated in the optical regime, with the atom-cavity interaction used as a mechanism for entangling atoms and photons [76], for generating multi-photon entanglement [77] and for transferring a quantum state between an atom and photon [78]. Using Rydberg atoms, increasingly complex quantum states can be tailor-made by multiple interactions between a cavity field and consecutive atoms, for example [79] and [80].

In contrast to these experiments, which are deterministic, an alternative approach uses ensembles of atoms coupled to a cavity to generate photons and transfer quantum states. Whilst this is probabilistic — the generation of a photon or of state transfer is heralded [37, 81, 82, 83].

Cavity QED systems have also been proposed as a tool for building quantum gates using nonlinearities arising from an atom’s interaction with a single photon. The low n_{sat} required in cavity-QED systems means that non-linear optical effects, such as absorptive optical bistability [84, 28] and phase shifts [85] can be demonstrated for cavity fields with less than one photon on average.

One factor to be considered, however, when appraising any system’s potential for quantum information processing is its scalability. This underlines the motivation for cavity QED research using scalable devices such as atom chips. The microcavity discussed in this thesis, formed using scalable, microfabrication techniques, represents a step towards this.

1.4 Outline of this thesis

This thesis presents experimental research carried out with cold rubidium atoms interacting with an optical microcavity, at the Centre for Cold Matter (CCM), Imperial College London and under the supervision of Prof. E. A. Hinds. The optical microcavities discussed here have previously been developed and used at CCM [43, 86], demonstrating the ability to detect atoms with single atom sensitivity by measuring changes in the intensity of the light reflected back from the cavity and by collecting atomic fluorescence. This work forms part of a broader research effort throughout the group to trap, manipulate and detect atoms on atom chips, with recent successes including the successful implementation of etched pyramids for use in atom trapping [87] and a microfabricated photonic waveguide chip [54].

The optical microcavity discussed in this thesis is unique in its design, being formed between the tip of an optical fibre and silicon micromirror and with parameters such that a single atom has a substantial impact on the cavity field. This thesis presents experiments characterising the atom-cavity interaction in the microcavity and demonstrating the single atom sensitivity of the cavity. The results are used to quantify the system's functionality as an atom detector. The relation between the cavity signals and atom number inside the cavity is also explored and atom-cavity spectroscopy is performed for different atom numbers to demonstrate the dispersive nature of the interaction. This latter measurement highlights key differences between the two types signals we measure in our system.

In Chapter 2 the basic theoretical results describing the atom-cavity interaction are presented, using the Jaynes-Cummings Hamiltonian as a starting point. These results are used to provide a theoretical description for the signals arising from the cavity.

Chapter 3 provides an overview of the experimental apparatus used for the work presented in this thesis. This includes a description of the optical cavities, including calculations of its key parameters; the laser system used for preparing and loading atoms into the cavity; the computer control which automates data collection; and the optical setup used for detecting and filtering atom-cavity signals.

Chapter 4 focuses on the use of the microcavities for atom detection. The two types of signal resulting from the atom-cavity interaction are discussed and characterised, with particular consideration to the optimal conditions for detection. This is followed by a discussion of atom detection using these signals, including achievable signal-to-noise ratio, signal fluctuations, experiment sensitivity and detection fidelity.

Chapter 5 presents the method for and results from increasing the number of atoms inside

the cavity. In this way, a regime of higher cooperativity is accessed, allowing a series of experiments investigating the dispersive nature of the atom-cavity interaction and demonstrating the key differences between the signals arising from the cavity.

In Chapter 6 two goals for the immediate future of the experiment are discussed; the loading and trapping of a single atom inside the cavity, and the coherent control of the atom inside the cavity. Progress towards these goals is presented, including a discussion of dipole trapping and the necessary Stark shift of the atoms' internal states. The development of a moving lattice for controllably loading atoms into the cavity is discussed and the technology required to coherently control an atom, via electromagnetically induced transparency in a vapour cell, is demonstrated.

Chapter 7 summarises the key results presented in this thesis and proposes areas for future research, including the possibility for further integration of our microcavities with other atom-chip devices.

Chapter 2

Atoms and cavities

This chapter presents the underlying theory describing the experiments discussed in this thesis. The key concepts are optical cavities, the normal mode structure of their electric fields and how atoms interact with them. The latter topic is modeled by the famous Jaynes-Cummings Hamiltonian [1], which describes the interaction between a single two-level atom and a single mode of the electromagnetic field. Furthermore, it is shown that the atom-cavity interaction has two distinct limits of interest; that of so-called ‘strong coupling’ or of a ‘fast cavity’. Finally, a theoretical form for the signals we can expect to measure arising from the atom-cavity field interaction is presented.

2.1 Optical cavities

The focus of this section is optical cavities; including the properties of their internal, transmitted and reflected electric fields and their internal mode structure.

2.1.1 Reflection and transmission properties of an optical cavity

Optical cavities are formed between two reflecting surfaces and light entering the cavity will be reflected back and forth between them. One simple cavity geometry to consider is a ‘Fabry Perot’ cavity, shown schematically in Figure 2.1. This is formed between two parallel, plane mirrors separated by a distance d . With each mirror is associated a field amplitude reflection and transmission coefficient, quantified by r_1, r_2 and t_1, t_2 respectively.

When a plane wave enters the cavity it will be reflected back and forth between the mirrors, with each round trip accumulating a phase shift $\phi_c = 2nkd$, where k is the wavenumber of the light and n is the refractive index of the medium. Successive reflections will interfere




Figure 2.1: A schematic of a Fabry-Perot optical cavity. The cavity is formed by two mirrors separated by a distance d . Each mirror has an associated reflection and transmission field-amplitude coefficients, r_1, r_2 and t_1, t_2 , respectively. A field of amplitude E_{in} impinges on mirror one, leading to a circulating field, E_{circ} inside the cavity and a transmitted field, $E_{out} = t_2 E_{circ}$.

constructively when ϕ_c equals an integer number of wavelengths, leading to the cavity resonance condition:

$$2kd = 2\pi m \quad \text{or} \quad \lambda = \frac{2d}{m}, \quad \text{where } m \text{ is an integer.} \quad (2.1)$$

In addition to an accumulating phase shift, each round trip will also be associated with a reduction in field amplitude, given by $r_1 r_2$. For an initial transmitted field, $E_{in} t_1$, after n round trips, the circulating intra-cavity field will be given by $(r_1 r_2)^n E_0 t_1 e^{i2knd}$, leading to a total field (summing over $n \rightarrow \infty$) of:

$$E_{circ} = \sum_{n=0}^{\infty} (r_1 r_2 e^{i2kd})^n t_1 E_{in} = \frac{E_{in} t_1}{1 - (r_1 r_2) e^{i2kd}}$$

The circulating intensity, $I_{circ} \propto |E_{circ}|^2$:

$$I_{circ} = \frac{I_{max}}{1 + \frac{4r_1 r_2}{(1 - r_1 r_2)} \sin(kd)^2} \quad \text{where} \quad I_{max} = \frac{I_{in}}{(1 - r_1 r_2)^2} \quad (2.2)$$

This is the well-known Airy Function, see for example [88], which results in a series of equally-spaced resonance peaks in the reflected and transmitted intensity, as either the length of the cavity, or the frequency of the light, is changed. The frequency separation of the peaks is referred to as the free spectral range of the cavity, ν_{FSR} , where $\nu_{FSR} = c/2d$, in units of Hz.

In our experiment, we measure the light reflected back from the cavity. This can be thought of as having two contributions; the light initially reflected from the first mirror, $-E_0 r_1$; and the fraction of light circulating inside the cavity which is re-transmitted through the mirror, $E_{circ} t_1 r_2 e^{i2kd}$. The total reflected intensity then becomes:




Figure 2.2: The reflected intensity as a function of cavity length, d , shown for two values of the mirror reflectivities. The solid, light green line is calculated for $r_1 = r_2 = 0.99$ and the dashed, dark green line is calculated for $r_1 = r_2 = 0.8$. In both cases $\lambda = 780$ nm. Each cavity fringe is caused by the destructive interference between the initially reflected light at mirror one and the backwards leakage through mirror one of the circulating cavity field. As the mirror reflectivities increase, the resonance fringes become narrower.

$$I_r \propto |-E_{in}r_1 + E_{circ}t_1r_2e^{i2kd}|^2 = I_0r_1^2 \left| \frac{(\frac{t_1r_2}{r_1} + r_1)e^{i2kd} - 1}{1 - r_1r_2e^{i2kd}} \right|^2 \quad (2.3)$$

Figure 2.2 shows a plot of Eq.(2.3) as ϕ_C is scanned, in this case by changing the cavity length; the solid line is calculated for $r_1 = r_2 = 0.99$ whilst the dashed lines is for $r_1 = r_2 = 0.8$. In both cases, $\lambda = 780$ nm. The cavity fringes are caused by destructive interference between the two contributions of Eq.(2.3). The finite width of these fringes is caused by the cavity losses and, as is evident from the figure, as these losses increase so too does the width of the fringes.

The cavity losses can be quantified by the cavity finesse, $\mathcal{F}$, which is defined as:

$$\mathcal{F} = \frac{\pi\sqrt{r_1r_2}}{1 - r_1r_2}$$

The finesse directly relates the cavity losses to the full width half maximums (FWHM) of the resonance peaks. From Eq.(2.2), the circulating cavity intensity falls to half its maximum value when:

$$\frac{4r_1r_2}{(1 - r_1r_2)} = \left(\frac{2\mathcal{F}}{\pi}\right)^2 = \frac{1}{(\sin kd)^2}$$

This occurs at wavenumbers:

$$k = \pm \frac{1}{d} \sin^{-1} \left(\frac{\pi}{2\mathcal{F}} \right) \approx \pm \frac{1}{d} \frac{\pi}{2\mathcal{F}}$$

Re-writing this in terms of frequency units, this implies that the FWHM occurs at a distance, $\Delta\nu$, from the maxima of:

$$\Delta\nu = \frac{c}{2d} \frac{1}{\mathcal{F}} = \frac{\nu_{\text{FSR}}}{\mathcal{F}}$$

Thus, $\mathcal{F}$ is the ratio of the fringe width to the fringe separation. In the limit of high reflectivity one can isolate a single mode, which is well approximated by a Lorentzian of HWHM κ . In this limit, $\mathcal{F} \simeq \pi/(1 - r_1 r_2)$ and one finds that:

$$2\kappa \equiv \frac{\nu_{\text{FSR}}}{\mathcal{F}} = \frac{1 - r_1 r_2}{\pi} \frac{c}{2d} = \frac{\varrho}{\tau_{rt}}$$

The latter ratio expresses κ , in angular frequency units, in terms of the cavity field loss per round trip, $\varrho = 1 - r_1 r_2$ and the round trip duration, $\tau_{rt} = 2d/c$. This is an important experimental parameter, as 2κ is the lifetime of a photon inside the cavity and thus measures the rate at which photons are lost from the cavity field. In this limit we find that κ can be measured directly by measuring the width of the cavity fringe, enabling us to infer the value of $\mathcal{F}$, using $\kappa = \pi c/2d\mathcal{F}$.

The depth of the fringes, normalised to the maximum reflected intensity, is known as the ‘fringe contrast’, θ , and is set by the internal losses of the cavity:

$$\theta = 1 - \frac{I_{\min}}{I_{\max}} \quad (2.4)$$

As will shown, this parameter sets the dynamic range of the reflection signals we expect to measure. This expression also allows us to write the reflected intensity as a function of frequency detuning, $\Delta_C = \omega - \omega_C$ from a particular cavity resonance at ω_C , in a full, Lorentzian form:

$$\frac{I_r}{I_{\max}} = 1 - \frac{\theta}{1 + \left(\frac{\Delta_C}{\kappa}\right)^2}$$

From this, we know that an optical cavity is frequency selective, supporting a narrow band of longitudinal modes of width $\sim 2\kappa$ which satisfy the resonance condition of Eq.(2.1). At these allowed frequencies, the intra-cavity field ‘builds up’, with the input field amplified by a factor $\sim \mathcal{F}/\pi$.

The Fabry-Perot cavity discussed here is a useful pedagogical tool for understanding the properties of the internal and reflected fields from a cavity. In reality, however, such cavities are plagued by diffraction losses, to the extent that a stable internal mode cannot form. For this reason, the cavity used in this experiment has one spherical mirror, providing sufficient re-focusing to counteract the effect of diffraction. The stable mode inside the cavity has spatial structure set by the cavity dimensions, as discussed in the following section.

2.1.2 Mode structure of an optical cavity

The internal cavity field must satisfy Maxwell's equations at the mirror boundaries. From Maxwell's equations, one can derive the Helmholtz wave equation describing the propagation of the wave, $E(\mathbf{r}) = \psi(\mathbf{r})e^{ikz}$, in the forward z direction. This equation can be simplified to the paraxial Helmholtz equation by making the paraxial approximation, where one neglects the divergence of the light as it propagates, valid when the divergence angle of the wave is small¹. The equation is given by:

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + i2k \frac{\partial}{\partial z} \right) \psi(\mathbf{r}) = 0$$

One complete set of solutions to this equation are known as Hermite-Gaussian modes, given in cylindrical coordinates, (r, z) , where z is directed along the longitudinal cavity axis and r is the radial coordinate, $r^2 = x^2 + y^2$, as:

$$\begin{aligned} E_{l,m}(\mathbf{r}) = & E_0 \frac{w_0}{w(z)} \exp \left[\frac{-r^2}{w(z)^2} \right] \times \sqrt{\frac{1}{2^{l+m} l! m!}} H_l \left[\frac{\sqrt{2}x}{w(z)} \right] H_m \left[\frac{\sqrt{2}y}{w(z)} \right] \times \\ & \exp \left[ikz - i(1+l+m) \arctan \left(\frac{z}{z_0} \right) \right] \exp \left[ik \frac{r^2}{2R(z)} \right] \end{aligned} \quad (2.5)$$

The solution is formed of several factors, the first of which describes the electric field envelope, where E_0 is the peak field amplitude at origin:

$$E_0 \frac{w_0}{w(z)} \exp \left[\frac{-r^2}{w(z)^2} \right] \quad (2.6)$$

The beam amplitude is maximised on axis and reduces with increasing $|r|$, following a Gaussian profile. From this, the radius of the beam, $w(z)$, is defined as the radial distance over which the electric field amplitude falls to $1/e$ of its original value. $w(z)$ is specified by two parameters specific to the beam; the beam waist, w_0 and the Rayleigh length, z_0 :

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_0} \right)^2} \quad \text{where} \quad z_0 = \frac{\pi w_0^2}{\lambda} \quad (2.7)$$

This indicates that $w(z)$ is minimised when $w(z) = w_0$, which occurs when $z = 0$. $w(z)$ is a hyperbolic function of z with distinct regimes of behaviour demarcated by z_0 . When $|z| \ll z_0$, the beam is collimated. In the opposite limit, when $|z| \gg z_0$, the beam diverges from the axis at a divergence angle of w_0/z_0 . This sets the limits for when it is reasonable to neglect the

¹Here, 'small' means that it is valid to make the approximation $\sin \theta \simeq \tan \theta \simeq \theta$, where θ is the divergence angle of the beam. Roughly speaking, this is equivalent to $w_0 k \gg 1$.

increase in the transverse mode size with increasing z . For our cavity, which has a length, $d = 139 \mu\text{m}$, comparable to the Rayleigh range, with $z_0 = 85 \mu\text{m}$, this should be accounted for.

The second factor in Eq.(2.5) describes the radial mode structure that lies within the field envelope. The factor contains Hermite polynomials, $H_l(x)$ and $H_m(y)$, of order l, m . When the beam entering the cavity is mode-matched and carefully aligned to the $l = m = 0$ mode, as for our cavity, one can simplify to the case when $H_l(x) = H_m(y) = 1$, referred to as the fundamental mode, or TEM₀₀ mode.

The third term in Eq.(2.5) describes how the phase of the beam propagates:

$$\exp \left[ikz - i \arctan \left(\frac{z}{z_0} \right) \right] \times \exp \left[-ik \frac{k(r^2)}{2R(z)} \right]$$

The first exponential describes the longitudinal phase factor, for $(l, m = 0)$. This includes a sum of the phase for a plane wave propagating along the z axis and the Gouy phase shift, which is the phase retardation associated with the discrepancy from true planarity; over all z this will retard the phase by π . The second exponential describes how the phase varies transversely across the beam, resulting in a spherical wavefront with a radius of curvature, $R(z)$ given by:

$$R(z) = z + \frac{z_0^2}{z}$$

Again, the Rayleigh length defines two regimes for the wavefront curvature: for $|z| \ll z_0$, $R(z) \sim z_0^2/z$, whilst for $|z| \gg z_0$, $R(z) \sim z$, so that a Gaussian beam approaches a spherical beam as $|z| \rightarrow \infty$. The maximum in the curvature occurs when $dR(z)/dz = 0$, at $z = \pm z_0$.

These expressions can be used to find the spatial intensity distribution inside the cavity. For a Gaussian beam in free space, $I_{l,m}(\mathbf{r}) \propto |E_{l,m}(\mathbf{r})|^2$. In addition to this, a resonant mode in the cavity will form a standing wave with an intensity $\propto \sin^2(kz)$. Combining these factors gives the total cavity intensity distribution:

$$I_{00}(\mathbf{r}) = I_0 \left[\frac{w_0}{w(z)} \right]^2 e^{\frac{-2r^2}{w(z)^2}} \sin^2(kz) \quad (2.8)$$

One can integrate this distribution over the cavity length to find the cavity mode volume. When d equals an integer number of wavelengths, this yields:

$$V = \int I_{00}(\mathbf{r}) d^3\mathbf{r} = \pi w_0^2 \frac{d}{4} \quad (2.9)$$

As previously mentioned, the stability of the mode inside the cavity is determined by the need to limit diffraction losses. The incoming beam will make multiple round trips of the cavity and will be repeatedly re-focused at each cavity mirror. This re-focusing must be sufficient to

counteract the effect of beam divergence, which would otherwise result in increasing amounts of energy being lost. In order for there to be a stable beam mode inside the cavity, there must be sufficient re-focusing for the beam to be identical after one round trip. This condition and applying the ABCD ray transfer matrices to describe a round trip through the cavity, see for example [89], leads to the stability condition:

$$0 \leq \left(1 - \frac{d}{R_1}\right) \left(1 - \frac{d}{R_2}\right) \leq 1 \quad (2.10)$$

In the above, R_1 and R_2 are the radii of curvature for mirrors 1 and 2, respectively. The optical cavity used in our system is plano-concave, so that $R_1 = \infty$. This leads to the reduced stability condition that $0 \leq d \leq R_2$ and sets the range of stability for the cavity length.

These optical properties have been understood by treating the electromagnetic field as a continuous, classical variable. At very low intensities however, it is more appropriate to treat a quantised field.

2.2 The quantum electric field

The majority of the experiments discussed here will involve very low amplitude fields, corresponding to only a few photons. In this situation a quantum mechanical treatment is necessary. We begin by quantising the field.

2.2.1 Quantising the electric field

To quantise the electromagnetic field, one can first consider the field inside a box of volume, V . The electric and magnetic fields, $\mathbf{E}$ and $\mathbf{H}$, inside the box can be written in terms of an expansion of plane waves, oscillating at frequencies $\omega_{\mathbf{k}} = |\mathbf{k}|c$ and with polarisation vectors $\epsilon_{\mathbf{k}}, \epsilon'_{\mathbf{k}}$:

$$\begin{aligned} \mathbf{E} &= \sqrt{\frac{1}{\epsilon_0 V}} \sum_{\mathbf{k}} \epsilon_{\mathbf{k}} \omega_k E_{\mathbf{k}} e^{i(\mathbf{k} \cdot \mathbf{r} - \omega_k t)} + c.c \\ \mathbf{H} &= \sqrt{\frac{1}{\epsilon_0 V}} \sum_{\mathbf{k}} \epsilon'_{\mathbf{k}} \omega_k H_{\mathbf{k}} e^{i(\mathbf{k} \cdot \mathbf{r} - \omega_k t)} + c.c \end{aligned} \quad (2.11)$$

These fields must satisfy the Maxwell equations for all $\mathbf{r}$ and t and, in addition, each Fourier component must satisfy the reciprocal space Maxwell equations:

$$\begin{aligned} (\mathbf{k} \times \epsilon_{\mathbf{k}}) E_{\mathbf{k}} &= \epsilon'_{\mathbf{k}} \mu_0 \omega_k H_{\mathbf{k}} \\ (\mathbf{k} \times \epsilon'_{\mathbf{k}}) H_{\mathbf{k}} &= -\epsilon_{\mathbf{k}} \epsilon_0 \omega_k E_{\mathbf{k}} \end{aligned}$$

$$\mathbf{k} \cdot \boldsymbol{\epsilon}_{\mathbf{k}} = 0$$

$$\mathbf{k} \cdot \boldsymbol{\epsilon}_{\mathbf{k}}' = 0$$

The last two equalities show that $\boldsymbol{\epsilon}_{\mathbf{k}}, \boldsymbol{\epsilon}_{\mathbf{k}}'$ are perpendicular to $\mathbf{k}$. The first two equalities are satisfied when²:

$$\begin{aligned} H_{\mathbf{k}} &= \frac{E_{\mathbf{k}}}{c\mu_0}; \\ \boldsymbol{\epsilon}_{\mathbf{k}}' &= \hat{\mathbf{k}} \times \boldsymbol{\epsilon}_{\mathbf{k}}, \end{aligned}$$

In the above, $\hat{\mathbf{k}}$ is a unit vector. These expressions can be used to re-write Eq.(2.11):

$$\begin{aligned} \mathbf{E} &= \sqrt{\frac{1}{\varepsilon_0 V}} \sum_{\mathbf{k}s} \boldsymbol{\epsilon}_{\mathbf{k}s} \omega_k E_{\mathbf{k}s}(t) e^{i\mathbf{k} \cdot \mathbf{r}} + c.c \\ \mathbf{H} &= \sqrt{\frac{1}{\varepsilon_0 V}} \frac{1}{c\mu_0} \sum_{\mathbf{k}s} (\hat{\mathbf{k}} \times \boldsymbol{\epsilon}_{\mathbf{k}s}) \omega_k E_{\mathbf{k}s}(t) e^{i\mathbf{k} \cdot \mathbf{r}} + c.c \end{aligned}$$

The expansions now contain a sum over the two orthogonal polarisation modes, s , and the field amplitude has absorbed the explicit time dependence; $E_{\mathbf{k}s}(t) = E_{\mathbf{k}s} e^{-i\omega_k t}$. These equations can be used to find the energy of the field, U integrated over the volume of the box, written classically as:

$$U = \frac{1}{2} \int_V (\varepsilon_0 |\mathbf{E}|^2 + \mu_0 |\mathbf{H}|^2) dV = 2 \sum_{\mathbf{k}s} \omega_k^2 |E_{\mathbf{k}s}(t)|^2$$

The latter evaluation assumes that the Fourier modes are orthonormal³. One can now introduce the variables, $q_{\mathbf{k}s}$ and $p_{\mathbf{k}s}$, written in terms of the electric field amplitude:

$$q_{\mathbf{k}s}(t) = E_{\mathbf{k}s}(t) + E_{\mathbf{k}s}^*(t) \quad \text{and} \quad p_{\mathbf{k}s}(t) = -i\omega_k E_{\mathbf{k}s}(t) - E_{\mathbf{k}s}^*(t)$$

The field energy can be expressed in terms of $q_{\mathbf{k}s}$ and $p_{\mathbf{k}s}$ as:

$$H = U = 2 \sum_{\mathbf{k}s} [p_{\mathbf{k}s}^2(t) + \omega_k^2 q_{\mathbf{k}s}^2(t)]$$

This expression for the field energy has the same form as the Hamiltonian of a harmonic oscillator with unit mass, see for example [3], implying that the electromagnetic field can be treated as a sum over independent harmonic oscillators. The quantisation of the field now follows by switching to a new basis, formed by the non-Hermitian operators $\hat{a}_{\mathbf{k}s}$ and $\hat{a}_{\mathbf{k}s}^\dagger$:

$$\hat{a}_{\mathbf{k}s}(t) = \sqrt{\frac{1}{2\hbar\omega_k}} [\omega_k \hat{q}_{\mathbf{k}s}(t) + i\hat{p}_{\mathbf{k}s}(t)]$$

²Using the fact that $\mathbf{k} \times (\mathbf{k} \times \boldsymbol{\epsilon}_{\mathbf{k}}) = -k^2 \boldsymbol{\epsilon}_{\mathbf{k}}$.

³Using the fact that $\int_V e^{i(\mathbf{k}-\mathbf{k}') \cdot \mathbf{r}} dV = V \delta_{\mathbf{k}\mathbf{k}'}$ and that $(\hat{\mathbf{k}} \times \boldsymbol{\epsilon}_{\mathbf{k}s}) \cdot (\hat{\mathbf{k}} \times \boldsymbol{\epsilon}_{\mathbf{k}s}') = \delta_{ss'}$.

$$\begin{aligned}
&= \sqrt{\frac{2\omega_k}{\hbar}} \hat{E}_{\mathbf{k}s}(t) \\
\hat{a}_{\mathbf{k}s}^\dagger(t) &= \sqrt{\frac{1}{2\hbar\omega_k}} [\omega_k \hat{q}_{\mathbf{k}s}(t) - i\hat{p}_{\mathbf{k}s}(t)] \\
&= \sqrt{\frac{2\omega_k}{\hbar}} \hat{E}_{\mathbf{k}s}^\dagger(t)
\end{aligned}$$

These operators satisfy the commutation relations:

$$[\hat{a}_{\mathbf{k}s}, \hat{a}_{\mathbf{k}s}^\dagger] = \delta_{\mathbf{k}\mathbf{k}'} \delta_{ss'}$$

The Hamiltonian for the field is expressed in terms of these operators as:

$$\hat{H} = \sum_{\mathbf{k}s} \hat{H}_{\mathbf{k}s} = \frac{1}{2} \sum_{\mathbf{k}s} \hbar\omega_k \left(\hat{a}_{\mathbf{k}s} \hat{a}_{\mathbf{k}s}^\dagger + \frac{1}{2} \right)$$

$\hat{H}_{\mathbf{k}s}$ is the Hamiltonian for a single electromagnetic mode and the $\hbar\omega/2$ term represents the energy of the vacuum. A single electromagnetic mode has the energy eigenstate, $|n_{\mathbf{k}s}\rangle$, with energy E_n :

$$\hat{H}_{\mathbf{k}s} |n_{\mathbf{k}s}\rangle = \hbar\omega_k \left(n_{\mathbf{k}s} + \frac{1}{2} \right) |n_{\mathbf{k}s}\rangle = E_n$$

From this, $n_{\mathbf{k}s}$ can be treated as the number of excitations, or photons, in a given field mode, with each photon contributing energy $\hbar\omega_k$ to the field. States of different photon number are equally spaced in energy, forming a ladder of energy eigenstates for a given wavenumber, $\mathbf{k}$ and polarisation, s .

Concentrating now on a single eigenstate, $|n\rangle$, one can define a number operator, $\hat{N} = \hat{a}^\dagger \hat{a}$, so that the Hamiltonian is re-written as:

$$\hat{H} = \hbar\omega \left(\hat{N} + \frac{1}{2} \right)$$

This implies that:

$$\hat{N}|n\rangle = \left(\frac{E_n}{\hbar\omega} - \frac{1}{2} \right) |n\rangle = n|n\rangle$$

The eigenvalue of $\hat{N}$ is n , so that $|n\rangle$ is therefore referred to as a *number state* or *Fock state*. This operator also enables us to derive the following equations⁴:

$$\begin{aligned}
\hat{N}(\hat{a}|n'\rangle) &= (n' - 1)(\hat{a}|n'\rangle) \\
\hat{N}(\hat{a}^\dagger|n'\rangle) &= (n' + 1)(\hat{a}^\dagger|n'\rangle)
\end{aligned}$$

Thus, the eigenvalue of the state $(\hat{a}|n'\rangle)$ is $(n' - 1)$ and the eigenvalue of the state $(\hat{a}^\dagger|n'\rangle)$ is $(n' + 1)$, implying that applying $\hat{a}$ and $\hat{a}^\dagger$ to the state $|n'\rangle$ decreases and increases the excitation

⁴Using the fact that $\hat{N}\hat{a} = \hat{a}(\hat{N} - 1)$ and $\hat{N}\hat{a}^\dagger = \hat{a}^\dagger(\hat{N} + 1)$

number of the state by one, respectively. For this reason, $\hat{a}$ and $\hat{a}^\dagger$ are referred to as annihilation and creation operators of the field. More formally, the normalised action of $\hat{a}$ and $\hat{a}^\dagger$ is given by:

$$\hat{a}^\dagger|n\rangle = \sqrt{n+1}|n+1\rangle \quad \text{and} \quad \hat{a}|n\rangle = \sqrt{n}|n-1\rangle$$

The electric and magnetic field operators for a single electromagnetic mode can also be expressed in terms of $\hat{a}$ and $\hat{a}^\dagger$ as:

$$\begin{aligned} \hat{\mathbf{E}} &= \sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}} \epsilon_{\mathbf{k}s} (\hat{a} e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} + \hat{a}^\dagger e^{-i(\mathbf{k}\cdot\mathbf{r}-\omega t)}) \\ \hat{\mathbf{H}} &= \sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}} \frac{1}{c\mu_0} (\hat{\mathbf{k}} \times \epsilon_{\mathbf{k}s}) (\hat{a} e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} + \hat{a}^\dagger e^{-i(\mathbf{k}\cdot\mathbf{r}-\omega t)}) \end{aligned} \quad (2.12)$$

This gives the ‘electric field per photon’ as $[\hbar\omega/2\varepsilon_0 V]^{\frac{1}{2}}$. This quantised description of the field can now be used to understand the quantised intra-cavity field. Whilst this derivation assumes a box of volume V , in principle this box can be expanded to have an infinite volume, so as to describe the field in free space. As it happens, however, in our experiment the electric field is ‘confined’ inside our cavity.

2.2.2 The driven, damped quantised cavity field

For a cavity containing n photons at frequency ω_C , the Hamiltonian describing the quantised cavity field is given by:

$$H_C = \hbar\omega_C (\hat{a}^\dagger \hat{a} + \frac{1}{2})$$

This Hamiltonian describes an isolated cavity system, but in reality this is never the case. Typically, the cavity mode will be pumped by an external light source impinging on one of the cavity mirrors, with frequency ω_L . This can be represented in the Hamiltonian by a pumping term, $-i\hbar\eta(\hat{a} - \hat{a}^\dagger)$, where η is the pumping rate amplitude. The Hamiltonian in the interaction picture (i.e. a frame of reference rotating at frequency ω_C) then becomes:

$$H_C^{int} = -\hbar\Delta_C \hat{a}^\dagger \hat{a} - i\hbar\eta(\hat{a} - \hat{a}^\dagger) \quad (2.13)$$

Here, $\Delta_C = \omega_L - \omega_C$ and is the detuning between the cavity mode and the pumping beam. In addition, losses have thus far been ignored which, in reality, would cause the cavity field to decay until the vacuum state, $|0\rangle$, is reached. We know that the field amplitude decays at a rate κ and the effect can be taken into account using the master equation, which gives the

equation of motion of the reduced density operator of the cavity interacting with a reservoir of field modes in thermal equilibrium, see for example [90]. The master equation for the cavity field is given by:

$$\dot{\rho} = -i[H_C^{int}, \rho] + \kappa(2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) \quad (2.14)$$

A steady state solution to this equation is $\rho = |\alpha\rangle\langle\alpha|$, where $|\alpha\rangle$ is a coherent state. These states are defined by satisfying the factorisation condition $\langle\alpha|\hat{a}^\dagger|\alpha\rangle\langle\alpha|\hat{a}|\alpha\rangle = \langle\alpha|\hat{a}^\dagger\hat{a}|\alpha\rangle$, a corollary of which is that they are eigenstates of $\hat{a}^\dagger$. Given this, they can be written as a superposition of number states, see for example [3], written in a normalised form as:

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_n \frac{\alpha^n}{\sqrt{n!}} |n\rangle$$

As an aside, this description of the field state demonstrates one of its statistical features, which is that the probability, P_n for finding n photons is given by the Poisson distribution:

$$P_n = |\langle n|\alpha\rangle|^2 = e^{-|\alpha|^2} \frac{|\alpha|^{2n}}{n!}$$

From this, we would expect fluctuations of the number of photons in the field to be characterised by having a variance equal to the mean.

Given that $\alpha = \langle a \rangle$, we can find a steady-state solution for α by multiplying Eq.(2.14) by $\hat{a}$ and taking the trace, knowing that⁵:

$$\langle a \rangle = \text{Tr}[\rho\hat{a}] \quad \text{and} \quad \text{Tr}[\rho] = 1 \quad (2.15)$$

This gives:

$$\alpha = \frac{\eta}{\kappa - i\Delta_C}$$

From this, on resonance the cavity photon number, $\langle n \rangle = |\alpha|^2$, is given by η^2/κ^2 . The field will change when an atom is present inside the cavity, the effect of which is discussed in the following section.

2.3 The interaction between atoms and a quantised cavity field

This section outlines the key features of the interaction between an atom and a quantised field, with and without decoherence. We begin by considering the term in the Hamiltonian which results from this interaction.

⁵This procedure also makes use of the cyclical properties of the trace, i.e. $\text{Tr}[ABC] = \text{Tr}[CAB] = \text{Tr}[BCA]$.

2.3.1 The two-level atom interacting with an electric field

For the purposes of this section, the atom is treated as a simple, closed two-level system with ground state $|g\rangle$ and excited state $|e\rangle$. These states are separated in energy by $\hbar\omega_A$. The Hamiltonian describing this free atom is given by:

$$H_A = \hbar\omega_A|e\rangle\langle e| \quad \text{or} \quad \hbar\omega_A\sigma^+\sigma^- \quad (2.16)$$

The latter expression uses the atomic raising and lowering operators, $\sigma^+ = |e\rangle\langle g|$, $\sigma^- = |g\rangle\langle e|$.⁶

The free atom has an electric dipole moment, $\mathbf{d}$, although when the atom is in an energy eigenstate $\mathbf{d} = 0$, as it is non-polar. This means that $\mathbf{d}$ can be expressed as:

$$\mathbf{d} = \mathbf{d}_{eg}\sigma^+ + \mathbf{d}_{ge}\sigma^- \quad \text{where} \quad \mathbf{d}_{eg} = \langle e|\mathbf{d}|g\rangle \quad (2.17)$$

If this atom is now placed in an electric field, the interaction energy is described by the dipole interaction Hamiltonian:

$$H_I = -\mathbf{d}\cdot\mathbf{E}(t)$$

This expression makes the Dipole Approximation, see for example [3], for which it is assumed that the physical extent of the atom is sufficiently small compared to the wavelength of the field that the field phase can be treated as constant (i.e. $\mathbf{k}\cdot\mathbf{r} \ll 1$). Substituting in the expression for the electric field operator from Eq.(2.12)⁷ and the dipole operator from Eq.(2.17) gives:

$$H_I = -\sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}}(\mathbf{d}_{eg}\sigma^+ + \mathbf{d}_{ge}\sigma^-)\cdot\hat{\mathbf{e}}_{\mathbf{k}s}(\hat{a}e^{-i\omega t} + \hat{a}^\dagger e^{i\omega t})$$

Multiplying out gives:

$$H_I = -\sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}} \left[\langle e|\mathbf{d}\cdot\hat{\mathbf{e}}_{\mathbf{k}s}|g\rangle (\sigma^+ \hat{a}e^{-i\omega t} + \sigma^+ \hat{a}^\dagger e^{i\omega t}) + \langle g|\mathbf{d}\cdot\hat{\mathbf{e}}_{\mathbf{k}s}|e\rangle (\sigma^- \hat{a}e^{-i\omega t} + \sigma^- \hat{a}^\dagger e^{i\omega t}) \right] \quad (2.18)$$

If Eq.(2.18) is now transformed into the interaction picture using the unitary transformation $U^\dagger H_I U$, where $U = e^{iH_A t/\hbar}$, two pairs of terms are evident; the pair evolving at frequencies $\omega + \omega_A$ and another pair evolving at $\omega - \omega_A$. These pairs of terms are evolving on starkly different timescales, with the former taking place, in this case, at optical frequencies whilst the latter is evolving slowly near resonance. Given this, one can make the Rotating Wave Approximation (RWA), see for example [3], where it is assumed that on the time scale of the

⁶Here, $\sigma^\pm$ are defined in terms of the Pauli spin matrices as $\sigma^\pm = \frac{1}{2}(\sigma_x \pm i\sigma_y)$, where $\sigma_x = |e\rangle\langle g| + |g\rangle\langle e|$ and $\sigma_y = -i|e\rangle\langle g| + i|g\rangle\langle e|$.

⁷The Dipole Approximation is made again in neglecting the $\mathbf{k}\cdot\mathbf{r}$ phase terms in the field operator.




Figure 2.3: A schematic showing the ladder of states of doublet states formed as a result of the atom-cavity interaction. The LHS ladder is for $g = 0$, whilst the RHS ladder shows the same but now $g > 0$. On the LHS, the atom and cavity are detuned by δ_{AC} , which produces the small splitting to what would otherwise be degenerate states.

near-resonant interaction between the atom and the light, the fast-evolving pair average to zero and can be neglected. Having done this one finds a simplified interaction Hamiltonian:

$$H_I = -\sqrt{\frac{\hbar\omega}{2\varepsilon_0 V}} \left[\langle e | \mathbf{d} \cdot \hat{\mathbf{e}}_{\mathbf{k}s} | g \rangle \sigma^+ \hat{a} e^{-i(\omega - \omega_A)t} + \langle g | \mathbf{d} \cdot \hat{\mathbf{e}}_{\mathbf{k}s} | e \rangle \sigma^- \hat{a}^\dagger e^{i(\omega - \omega_A)t} \right] \quad (2.19)$$

From this, one can define the maximum atom-cavity coupling rate, g :

$$g = \sqrt{\frac{\omega |\mu_{eg}|^2}{2\varepsilon_0 V}} \quad \text{where} \quad \mu_{eg} = \langle e | \mathbf{d} \cdot \hat{\mathbf{e}}_{\mathbf{k}s} | g \rangle \quad (2.20)$$

Thus, g gives the interaction energy between the electric field of a single photon interacting with a single atom, in units of $\hbar$. Substituting this into Eq.(2.19) and making another unitary transformation to eliminate the time dependence gives the final Hamiltonian:

$$H_I = -\hbar \left(g \sigma^+ \hat{a} + g^* \hat{a}^\dagger \sigma^- \right) \quad (2.21)$$

This term is one of the constituents of the Jaynes-Cummings Hamiltonian.

2.3.2 The Jaynes-Cummings Hamiltonian

Having derived Hamiltonians describing the cavity field, the atom and the atom-cavity field interaction it is now appropriate to combine them in the famous Jaynes-Cummings Hamiltonian [1]. For now, we treat only the coherent interaction between the atom and cavity field, neglecting the effect of the finite cavity decay rate, κ .

For a stationary atom, using the Dipole and RW approximations and choosing an atomic quantisation axis so that g is real, one has:

$$H_{JC} = \hbar\omega_A\sigma^+\sigma^- + \hbar\omega_C a^\dagger a - \hbar g(\sigma^+ a + a^\dagger \sigma^-) \quad (2.22)$$

As mentioned, this Hamiltonian consists of three parts. The first term gives the internal energy of the free atom. The second term gives the energy of the cavity field excitation. The third term gives the energy associated with the interaction between the atom and the field and is proportional to the atom-field coupling rate, g . This rate is proportional to the electric field inside the cavity and, as such, it should be modified by a spatial factor to take into account the mode structure of the Gaussian standing wave formed inside the cavity. Using Eq.(2.8), one finds that:

$$g = g_0\psi(\mathbf{r}) \quad \text{where} \quad \psi(\mathbf{r}) = \left[\frac{w_0}{w(z)} \right] \sin(kz) e^{-\frac{r^2}{w_0^2}} \quad (2.23)$$

The above expression for g assumes the simplified case of a two-level atom with no internal structure and corresponds to the maximum achievable coupling rate. In reality the rubidium atoms used in this experiment have magnetic hyperfine structure, so that each possible transition has a g weighted by a Clebsch-Gordan coefficient. This also means that the atoms exhibit polarisation effects, a topic which is considered in more detail in Chapter 4.

Returning to the JC Hamiltonian, the basis vectors for the coupled atom-cavity states can be written as tensor products of the uncoupled atomic and Fock states, i.e. $\{|g, 0\rangle, |e, 0\rangle, |g, 1\rangle, |e, 1\rangle \dots\}$. In this closed atom-cavity system, energy conservation requires that an emission or absorption process by the atom must be reflected by a change in the cavity field state, which should increase or decrease respectively by one⁸. This creates a ladder of independent doublet states describing the coupled atom-cavity system, shown schematically in Figure 2.3. The LHS of this figure shows the situation where $g = 0$, with a finite atom-cavity detuning, $\delta_{AC} = \omega_C - \omega_A$, between the atom and the cavity which produces a splitting between otherwise degenerate eigenstates. The RHS shows the situation with $g \neq 0$, where now the eigenstates are split by a frequency $2\sqrt{ng}$.

The Hamiltonian for the doublet states, $|g, n\rangle, |e, n-1\rangle$, can be written in a general matrix format as:

⁸This restriction was actually made by the RWA, where terms which would have connected states such as $|g, 0\rangle$ and $|e, 1\rangle$ were neglected.




Figure 2.4: The system eigenvalues as a function of atom-cavity detuning, δ_{AC} , plotted for $n = 1$. The ‘bare’ eigenvalues for the composite atom-cavity system, for which there is no interaction between the atom and the cavity field, are represented and labelled by the dotted lines. The ‘dressed state’ eigenvalues include an interaction term and are represented by the solid line. In the latter case, the interaction causes the eigenvalues to repel one another close to $\delta_{AC} = 0$, causing an avoided crossing.

$$H_{JC} = \hbar \begin{bmatrix} \omega_C n + \delta_{AC} & g\sqrt{n} \\ -g\sqrt{n} & \omega_A + \omega_C(n-1) + \delta_{AC} \end{bmatrix}$$

Diagonalising this Hamiltonian gives two non-degenerate eigenvalues:

$$E_n^{\pm} = \frac{\hbar}{2} \left(2n\omega_C + \delta_{AC} \pm \sqrt{4g^2n + \delta_{AC}^2} \right)$$

The interaction between the atom and the cavity has split the cavity photon states into two non-degenerate ‘dressed states’. These eigenvalues are plotted in Figure 2.4. The dashed lines indicate the ‘bare’ atom-cavity eigenvalues, with no interaction, whilst the solid lines indicate the new dressed state eigenvalues. The interaction causes these eigenvalues to ‘repel’ one another, producing an ‘avoided crossing’ close to resonance. The energy difference between the states is minimised at $2g$ when $\delta_{AC} = 0$. The avoided crossing can be measured directly by measuring the excitation spectrum of the cavity — which is later shown to be proportional to the transmission spectrum.

In analogy to an atom interacting with a classical field, the atom undergoes coherent Rabi oscillations between the dressed states with an effective Rabi frequency, $\Omega_{\text{eff}} = \sqrt{4ng^2 + \delta_{AC}^2}$, with the excitation being transferred between the atom and the cavity field. In this ideal, closed quantum system, these oscillations would continue indefinitely. In reality, however, dissipation

due to atomic spontaneous emission and cavity decay, hitherto ignored in this model, will cause these oscillations to decay. As before, these effects can be accounted for using the master equation.

2.3.3 The Master equation for a driven atom-cavity system

There are two sources of dissipation in the coupled atom-cavity system: the atom causes loss through spontaneous emission at a rate 2γ and the cavity can cause loss through the mirrors at a rate 2κ . These can be included using the master equation, which describes the non-unitary evolution of the joint atom-cavity density matrix, ρ :

$$\frac{d}{dt}\rho = -i[H, \rho] + \gamma(2\sigma^-\rho\sigma^+ - \sigma^+\sigma^-\rho - \rho\sigma^+\sigma^-) + \kappa(2a\rho a^\dagger - a^\dagger a\rho - \rho a^\dagger a) \quad (2.24)$$

The Hamiltonian, H included here refers to the Jaynes-Cummings Hamiltonian of Eq.(2.22), with the inclusion of a cavity pumping term, $-i\hbar(\hat{a} - \hat{a}^\dagger)$. In the interaction picture this Hamiltonian is:

$$H = -\Delta_C \hat{a}^\dagger \hat{a} - \Delta_A \sigma^+ \sigma^- - g(\hat{a}\sigma^+ + \hat{a}^\dagger\sigma^-) - i\eta(\hat{a} - \hat{a}^\dagger) \quad (2.25)$$

The cavity and atomic detunings are defined with respect to the frequency of the pump light, so that:

$$\Delta_C = \omega_L - \omega_C \quad \text{and} \quad \Delta_A = \omega_L - \omega_A$$

The evolution of the system described by Eq.(2.24) cannot, in general, be solved analytically. One can reach an approximate solution, however, by treating the cavity field semiclassically and again assuming it is in a coherent state, α , where $\alpha = \langle \hat{a} \rangle$. In this approximation, one also assumes that the joint expectation values can be factored out, i.e. $\langle \hat{a}\hat{\sigma}^- \rangle \rightarrow \langle \hat{a} \rangle \langle \hat{\sigma}^- \rangle$ ⁹. Using this approach, the master equation can be used to derive the Heisenberg equations of motion for the expectation values of the cavity field operator, $\langle \hat{a} \rangle$ and the atomic raising and lowering operators, $\langle \hat{\sigma}^+ \rangle$ and $\langle \hat{\sigma}^- \rangle$, following Eq.(2.15):

$$\frac{d}{dt}\alpha = -(\kappa + i\Delta_C)\alpha + \eta + ig\langle \sigma^- \rangle \quad (2.26)$$

$$\frac{d}{dt}\langle \sigma^- \rangle = -(\gamma + i\Delta_A)\langle \sigma^- \rangle + ig\alpha[1 - 2\langle \sigma^+ \sigma^- \rangle] \quad (2.27)$$

$$\frac{d}{dt}\langle \sigma^- \sigma^+ \rangle = -2\gamma\langle \sigma^+ \sigma^- \rangle + ig[\alpha\langle \sigma^+ \rangle - \alpha^*\langle \sigma^- \rangle] \quad (2.28)$$

⁹This approximation is only valid when the cavity is weakly excited.




Figure 2.5: The cavity transmission as a function of detuning, Δ where $\Delta = \Delta_C - \Delta_A$. The solid line represents an empty cavity which shows one peak centred on $\Delta = 0$. The dashed line represents a cavity with one atom present, with $(g, \gamma, \kappa) = 2\pi \times (80, 3, 50)$ MHz, so that $C = 21.3$ and the system is in the strong coupling regime, to be defined later. The normal mode splitting of the cavity resonance symmetrically about $\Delta = 0$ is clearly visible.

These equations have the steady state solutions:

$$\alpha = \frac{1}{\tilde{\kappa}} [\eta + ig\langle\sigma^-\rangle] \quad (2.29)$$

$$\langle\sigma^-\rangle = \frac{ig\alpha}{\tilde{\gamma}} [1 - 2\langle\sigma^+\sigma^-\rangle] \quad (2.30)$$

$$\langle\sigma^+\sigma^-\rangle = \frac{g^2|\alpha|^2/|\tilde{\gamma}|^2}{1 + 2g^2|\alpha|^2/|\tilde{\gamma}|^2} \quad (2.31)$$

Here, the atom and cavity detunings have been absorbed into the new variables, $\tilde{\kappa} = \kappa + i\Delta_C$ and $\tilde{\gamma} = \gamma + i\Delta_A$. This set of equations does not form a closed system, so that further approximations must be made to find analytical solutions.

2.3.4 The coherent cavity field in the limit of weak excitation

A reasonable approximation for the majority of experiments discussed here is to assume that the cavity is only weakly excited, so that $\langle\sigma^+\sigma^-\rangle \ll 1$. In this case, the steady state solution for the intra-cavity field can be found from Eq.(2.29), giving:

$$\alpha = \frac{\eta}{\tilde{\kappa}} \frac{1}{1 + 2\tilde{C}} \quad \text{where} \quad \tilde{C} = \frac{g^2}{2\tilde{\kappa}\tilde{\gamma}} \quad (2.32)$$

On resonance, $\tilde{C}$ is known as the *single atom cooperativity* and parameterises how strongly the atom couples to the cavity. When $C \sim 1$, the coherent atom-cavity interaction is comparable to the dissipation rates of the system.

The steady-state cavity photon number, $\langle n \rangle$, is given by $\langle n \rangle = |\alpha|^2$, so that the cavity transmission becomes:

$$T(\Delta) \propto \langle n \rangle = |\alpha|^2 = \left(\frac{\eta^2}{|\tilde{\kappa}|^2} \right) \left| \frac{1}{1 + 2\tilde{C}} \right|^2$$

For suitable values of (g, γ, κ) , the cavity transmission (and reflection) show a ‘normal mode splitting’, where the cavity resonance is split by $\pm \hbar g$, arising from the splitting of the system eigenstates into the two dressed states, an example of which is shown in Figure 2.5.

The damped system can be described by the non-Hermitian hamiltonian [91]:

$$H = H_{JC} - i\hbar\gamma\sigma^+\sigma^- - i\hbar\kappa\hat{a}^\dagger\hat{a}$$

This Hamiltonian has the modified eigenvalues:

$$E^\pm = \omega_C(n + \frac{1}{2}) + \frac{1}{2}\sqrt{4ng^2 - (\gamma - \kappa)^2} - \frac{i(\gamma + \kappa(2n - 1))}{2}$$

From this, the relative sizes of κ and γ with respect to g define different regimes of cavity QED. For the normal mode splitting to be resolved, $\Im[\tilde{E}^\pm] < \Re[\tilde{E}^\pm]$, which occurs when $g^2 > (\gamma^2 + \kappa^2)/2$. This is satisfied in the ‘strong coupling’ regime, characterised by $g \gg \kappa, \gamma$. The strong coupling regime is not necessary however for quantum effects to be visible. The *saturation photon number*, n_{sat} and the *critical atom number*, N_{crit} , already introduced, are given by:

$$N_{crit} = \frac{\kappa\gamma}{g^2} = \frac{1}{2C} \quad \text{and} \quad n_{sat} = \frac{\gamma^2}{2g^2}$$

As discussed in Chapter 1, having $(N_{crit}, n_{sat}) < 1$ means that individual quanta, be it atoms or photons, will have a pronounced effect. With $n_{sat} < 1$ a single photon will saturate the atom, giving rise to nonlinear optical effects. With $N_{crit} \sim 1$ a single atom will have a pronounced effect on the cavity field, enabling its detection. Both of these conditions can be met with a ‘fast cavity’, for which $g \gg \gamma$ and $\kappa \gg g$. This is the case in our system, which has:

$$g = 2\pi \times 98 \text{ MHz}$$

$$\kappa = 2\pi \times 5200 \text{ MHz}$$

$$\gamma = 2\pi \times 3 \text{ MHz}$$

These quantities lead to an $n_{sat} = 4.5 \times 10^{-4}$ and $N_{crit} = 1.6$. In Chapter 4 it will be shown we can detect a single atom with a fidelity of 99% in 16 μs . In Chapter 5 it is shown that we are able to measure the dispersive effect of the atom-cavity interaction for one atom on average inside the cavity mode. In addition, the fast cavity limit also has an impact on the atom’s spectral qualities, as discussed in the following section.

2.3.5 Enhanced spontaneous emission: the Purcell effect

Up to now we have primarily been interested in the effect that the atom has on the cavity field (e.g. its impact on the transmission spectrum). Conversely, the atom's spectral qualities can be modified by the presence of the cavity. More specifically, the composite effect of the atom-cavity coupling and cavity loss rate can be such that cavity provides an additional loss channel for the atom, enhancing its spontaneous emission rate. This is known as the Purcell effect [4]. Following the method outlined in [91], the time evolution of the system's state vector can be used to find the emission properties of the atom when coupled to the cavity. Assuming no cavity pumping and, as a consequence, weak excitation, the Hilbert space is truncated to allow up to a maximum of one excitation between the cavity and the field. The state vector is then given by:

$$|\psi\rangle = c_g|g, 1\rangle + c_e|e, 0\rangle + c_{g0}|g, 0\rangle$$

Given no cavity pumping, there is a high probability for there to be no excitation and $c_{g0} \sim 1$. The Schrodinger equation for the state coefficients is then given by:

$$\frac{d}{dt} \begin{bmatrix} c_e \\ c_g \end{bmatrix} = \begin{bmatrix} -\gamma & -g \\ g & -\kappa \end{bmatrix} \begin{bmatrix} c_e \\ c_g \end{bmatrix}$$

In the 'fast cavity' limit, where $\kappa \gg g$, Rabi oscillations are inhibited by the quick loss of photons before they can re-excite the atom. In these conditions, the transient population in the $|g, 1\rangle$ state is negligible, allowing the approximation $\dot{c}_g \sim 0$, giving:

$$\dot{c}_e = -\left(\frac{g^2}{\kappa} + \gamma\right) c_e \quad \text{leading to} \quad c_e = e^{-\left(\gamma + \frac{g^2}{\kappa}\right)t} \quad (2.33)$$

The presence of the cavity has modified the atomic emission rate, which now consists of two parts — one part describing the 'bare atom' decay rate, γ and another describing the decay rate into the cavity, γ_{cav} :

$$\gamma_{tot} = \gamma + \gamma_{cav} = \gamma \left(1 + \frac{g^2}{\kappa\gamma}\right) = \gamma(1 + 2C)$$

The cavity dominates as a decay channel for the atom if $g^2/\kappa\gamma = 2C \gg 1$ and this effect has been demonstrated with atoms inside resonators at optical and microwave frequencies [69, 92]. In our system, the decay rate into the cavity mode is of a similar order of magnitude as the free space decay rate.

The effect of detuning the cavity from the atom (and therefore reducing the mode density available to it) can be taken into account by including a factor $1/(1 + \Delta_C^2)$, giving:

$$\gamma_{tot} = \gamma \left(1 + \frac{2C}{(1 + \Delta_C^2)} \right)$$

This indicates that the cavity decay channel can be ‘switched off’ by sufficiently increasing Δ_C . In addition, spontaneous emission can be inhibited by the cavity if the solid angle subtended by the cavity mode is non-negligible and the cavity dimensions are on the order of the emission wavelength of the atom [93, 94, 95]¹⁰. In this case, the cavity can be sufficiently detuned from atomic resonance to eliminate decay via the cavity channel, but the atomic decay rate is smaller than a free space atom because the cavity mode has reduced the available density of states.

2.3.6 Extending to multiple atoms

Thus far only a single atom interacting with the cavity field has been considered. The treatment can, however, be readily extended to N identical atoms, by re-writing the Hamiltonian of Eq.(2.25) as:

$$H = -\Delta_C \hat{a}^\dagger \hat{a} - \Delta_A \sum_{j=1}^N \sigma_j^+ \sigma_j^- - \sum_{j=1}^N g_j \left(\hat{a} \sigma_j^+ + \hat{a}^\dagger \sigma_j^- \right) - i\eta(\hat{a} - \hat{a}^\dagger) \quad (2.34)$$

The subscripts j for the raising and lowering operators indicate raising and lowering operations for the j th atom and g_j indicates the local coupling rate for the j th atom, arising from the spatial variation of the atom-cavity coupling $g_j = g_0 \psi(x_j, y_j, z_j)$, causing atoms at different positions in the cavity to ‘feel’ different couplings strengths. Solving as before we find a modified expression for α in the weak driving limit:

$$\alpha = \left(\frac{\eta}{\tilde{\kappa}} \right) \frac{1}{1 + 2C_N} \quad \text{where} \quad C_N = C \sum_{j=1}^N \left| \frac{g_j}{g_0} \right|^2 \quad (2.35)$$

Here, the effect of multiple atoms has now been absorbed into a *collective cooperativity*, C_N , which contains a sum over all the individual atomic contributions [40, 96]. This is possible because the atoms are driven in phase by the cavity pump beam, ensuring that their contributions add ‘coherently’. This leads to the concept of the *effective* number of atoms in the cavity, N_{eff} , which gives the total contribution from N atoms in terms of the equivalent number of maximally coupled atoms; i.e. $N_{\text{eff}} = C_N/C$. This is an important concept for our experiment because, as will be discussed later, our method for loading atoms into the cavity means that

¹⁰By using the free space atomic decay rate, γ throughout this section, it has implicitly been assumed that the solid angle subtended by the cavity is sufficiently small as to be negligible.




Figure 2.6: The effect of increasing pumping powers on the intra-cavity field. $|\alpha_{res}|^2$ plotted as a function of cavity pumping rate, η for $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz. The lighter green, solid line is calculated for $N = 1$, corresponding to $C_N = 0.3$. In this case, $|\alpha_{res}|^2$ increases monotonically with increasing η . The darker green, dashed line is calculated for $N = 13$, corresponding to $C_N = 4$. For this case, $|\alpha_{res}|^2$ is multi-valued as the system becomes bistable.

there are always multiple atoms contributing to the signal. This indicates another point of interest; the cooperativity of the atom-cavity system can be increased by increasing the number of interacting atoms. This fact will be exploited in experiments presented in Chapter 5.

2.3.7 High intensity limit and optical bistability

Up to now, the limit of weak driving has been assumed, enabling the derivation of an analytical expression for intra-cavity field. Outside of this limit, Eq.(2.29) can be rearranged to find an expression for α in terms of $|\alpha|^2$:

$$\alpha = \left(\frac{\eta}{\tilde{\kappa}}\right) \frac{1 + 2g^2|\alpha|^2/|\tilde{\gamma}|^2}{1 + 2\tilde{C}_N + 2g^2|\alpha|^2/|\tilde{\gamma}|^2} \quad (2.36)$$

In general, this equation must be evaluated numerically. One exception occurs for the all-resonant case, when $\Delta_A = \Delta_C = 0$. For this situation, α is real and Eq.(2.36) can be rearranged to find the cubic equation:

$$\alpha_{res}^3 - \alpha_{res}^2 \left(\frac{\eta}{\kappa}\right) + \alpha_{res}(2C_N + 1)\frac{\gamma^2}{2g^2} - \frac{\eta}{\kappa} \frac{\gamma^2}{2g^2} = 0 \quad (2.37)$$

Figure 2.6 shows $|\alpha_{res}|^2$ as a function of pumping rate, η , plotted for $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz. The lighter green, solid line is calculated for a single atom, corresponding to $C_N = 0.3$. This function is single-valued for all η , with $|\alpha_{res}|^2$ increasing monotonically with increasing η . The darker green, dashed line is calculated for 13 maximally coupled atoms,

corresponding to $C_N = 4$. In contrast, at $\eta \approx 1000$ photons/sec, $|\alpha_{res}|^2$ becomes multi-valued, with multiple stable solutions, indicating bistability. It can be shown that the onset of optical bistability occurs when $C \geq 4$ [97] and the effect has been seen experimentally [28, 84, 98].

As is will be shown in Chapter 4, atomic saturation has an adverse effect on the atom signal reflected from the cavity — for very high pumping rates, it ultimately leads to the complete extinction of the signal. The theoretical form for the reflection signal is presented in the following section.

2.4 Cavity signals

In our experiment there are two types of signals which are caused by the presence of atoms inside the cavity; referred to throughout as the ‘reflection’ signal and the ‘fluorescence’ signal. Here, the analytical form for each signal is presented.

2.4.1 Cavity reflection signals

A cavity reflection signal is produced when the cavity is pumped at a rate η and the signal we measure is the light reflected back. As shown in Eq.(2.3), the reflected field, E_R has two contributions; one from the light reflected from the first cavity mirror and the other from the light transmitted through the first mirror from inside the cavity:

$$E_R = -r_1 E_{in} + E_{circ} r_2 t_1 \zeta \quad (2.38)$$

The first term on the RHS describes the light reflected back from the first cavity mirror. This is just the product of the incoming field, E_{in} and the reflection coefficient of mirror 1. Here, E_{in} is the time-averaged field of the incoming beam, where $E_{in} = \sqrt{J_{in}}$ and J_{in} is the incoming rate of photons impinging on cavity mirror 1. For these purposes, the units of E_{in} are thus $\sqrt{\text{photons/sec}}$.

The second term describes the light transmitted through the first mirror from inside the cavity. This is a product of the time-averaged field circulating inside the cavity, E_{circ} and terms denoting the efficiency with which light is coupled from the cavity: the reflection coefficient of mirror 2 describing its reflection at the second mirror inside the cavity, r_2 ; the transmission coefficient describing the transmission out of the cavity through mirror 1, t_1 ; and the factor ζ which quantifies the mode overlap between the mirror and the optical fibre.

The time-averaged cavity field is given in terms of the field amplitude, α , as $E_{circ} = \alpha/\sqrt{\tau_{rt}}$. From §3.4 we have an expression for α , given by Eq.(2.36). For simplicity, the reduced

expression in the limit of weak pumping given by Eq.(2.32) will be used here, although in principle the full expression could be used to describe the field in all regimes. Making these substitutions and assuming the case of high reflectivity of mirror 1, so that $r_1 \approx 1$, yields:

$$E_R = -\sqrt{J_{in}} + \frac{r_2 t_1 \zeta}{\sqrt{\tau_{rt}}} \times \left(\frac{\eta}{\tilde{\kappa}} \right) \frac{1}{1 + 2\tilde{C}_N} \quad (2.39)$$

This can be written in a more meaningful way by finding an expression for η in terms of J_{in} . We know that for an empty cavity, the steady state photon number, $n = \eta^2/\kappa^2$. The rate at which photons enter the cavity is given by the incoming rate, J_{in} , multiplied by the efficiency of photon transfer into the cavity from outside, given by $(t_1 \zeta)^2$. In the steady state, this equals the rate at which photons are lost from the cavity, given by $n\varrho^2/\tau_{rt}^2$, leading to:

$$n = \frac{\eta^2}{\kappa^2} = \frac{J_{in} t_1^2 \zeta^2 \tau_{rt}}{\varrho}$$

Rearranging and remembering from §1.1 that $\kappa = \varrho/\tau_{rt}$ gives an expression for η :

$$\eta = \sqrt{\frac{J_{in} t_1^2 \zeta^2}{\tau_{rt}}} \quad (2.40)$$

Including this in Eq.(2.39) leads to an expression for the outcoming reflected photon rate, J_R in terms of J_{in} :

$$J_R \propto |E_R|^2 = J_{in} \left| -1 + \frac{r_2 t_1^2 \zeta^2}{\varrho} \frac{1}{1 + 2\tilde{C}_N} \right|^2 \quad \text{or} \quad \frac{J_R}{J_{in}} = \left| -1 + \frac{r_2 t_1^2 \zeta^2}{\varrho} \frac{1}{1 + 2\tilde{C}_N} \right|^2 \quad (2.41)$$

On resonance, Eq.(2.4) from earlier in this chapter gives an expression for the ratio of the incoming and reflected intensities in terms of the fringe contrast, θ :

$$\theta = 1 - \frac{I_r^{res}}{I_{in}} = 1 - \frac{J_r^{res}}{J_{in}}$$

We can therefore simplify Eq.(2.41) by defining the fringe visibility, v , where $\theta = 1 - v^2$, so that:

$$-1 + \frac{r_2 t_1^2 \zeta^2}{\varrho} = -v \quad \text{giving} \quad 1 - v = \frac{r_2 t_1^2 \zeta^2}{\varrho} \quad (2.42)$$

Generally v could be positive or negative, depending on whether the cavity is over- or under-coupled, respectively. For our cavity, v is non-zero and is undercoupled so that we use $-v$. This choice means that the reflected signal increases with increasing cooperativity. This allows us to reach a final, simplified form for the reflected photon rate:

$$J_R = J_{in} \left(-1 + \frac{1 - v}{2\tilde{C}_N + 1} \right)^2 = J_{in} \left(\frac{2\tilde{C}_N + v}{2\tilde{C}_N + 1} \right)^2 \quad (2.43)$$




Figure 2.7: The impact of different fringe visibilities, v , on the reflection signal. LHS: The normalised reflection signal, J_R/J_{in} as a function of cooperativity, plotted for three different values of v . From darker to lighter green, $v^2 = 0.01, 0.25, 0.5$. The reflection signal saturates with increasing C_N , saturating more quickly for smaller values of v . The grey dots indicate how an increase in cooperativity has an equivalent effect to a reduction in fringe contrast. RHS: The differential of J_R/J_{in} as a function of atom number, N for three different cavity contrasts. From darker to lighter green correspond to $v = 0.1, 0.2$ and 0.3 . As v increases (and thus the fringe contrast decreases), the change in signal caused by a single atom decreases, making it harder to discriminate between atom numbers.

When no atom is present, $C_N = 0$, so that $J_R/J_{in} = v^2$. An atom inside the cavity will lead to a positive value for C_N , producing an increase in the signal which ‘looks’ like an increase in v . The presence of an atom in the cavity is therefore manifest as a *reduction in cavity contrast*. Physical, the reasoning behind this is that an atom in the cavity scatters light which would otherwise destructively interfere with the field reflected from the cavity mirror.¹¹

The LHS of Figure 2.7 shows J_R/J_{in} as a function of C_N (treating the resonant case), plotted for three different values of v ; from darker to lighter green, $v^2 = 0.01, 0.25, 0.5$. This shows that increasing C_N results in change to the signal which is equivalent to an increase v , as indicated schematically by the example shown in grey. Also evident is the saturation of the signal with higher C_N , tending towards 1 as $C_N \rightarrow \infty$. How quickly this saturation occurs is determined by the size of v ; for higher fringe contrasts, leading to smaller values of v , the saturation of the reflection signal occurs more slowly.

The value of the fringe contrast and saturation of the reflection signal also impacts our ability to discriminate between different atom numbers. This can be seen in the RHS of Figure 2.7, which shows $d(J_R/J_{in})/dN$ as a function of atom number, N , for three different values of v ; from lighter to darker green, $v = 0.1, 0.2$ and 0.3 . It is clear that the differential increases with decreasing v , so that it is easier to discriminate between different atom numbers as v decreases. With all values of v , however, saturation is evident. In certain circumstances, the fluorescence signal will also saturate with increasing cooperativity, as discussed in the following section.

2.4.2 Cavity fluorescence signals

The cavity fluorescence signal is the result of atoms emitting photons into the cavity mode, enhanced by means of the Purcell effect. Their emission is triggered by a pump beam of Rabi frequency Ω , which drives the atom-cavity system from the side, transverse to the longitudinal cavity axis. The cavity mode is therefore not being pumped directly, so $\eta = 0$. This situation is shown schematically in Figure 2.8.

The addition of the transverse pump beam can be included in the Hamiltonian of Eq.(2.25) by making the replacement $g\hat{a} \rightarrow (g\hat{a} + \Omega/2)$. For a single atom this yields:

$$H = -\Delta_C \hat{a}^\dagger \hat{a} - \Delta_A \sigma^+ \sigma^- - g \left(\hat{a} \sigma^+ + \hat{a}^\dagger \sigma^- \right) - \frac{1}{2} (\Omega \sigma^+ + \Omega^* \sigma^-)$$

¹¹Were the light pumping the cavity far-detuned, the scattering rate of the atom would be substantially reduced and the dominant effect of the atom is to shift the resonance frequency of the cavity [65], inducing a phase shift in the transmitted and reflected light.




Figure 2.8: Schematic of fluorescence experiment. An atom inside the cavity is pumped into an excited state by transverse pump beams, passing vertically through the cavity from the top and/or the bottom. The atom's decay rate has two components: γ into free space and γ_{cav} into the cavity mode. Here, the cavity is formed between an optical fibre and a silicon wafer, as will be described in Chapter 3.

Following the method of §3.3 and assuming low excitation, the intra-cavity field in the case of transverse driving, α_F , is given by:

$$\alpha_F = -\frac{2\tilde{C}}{1 + 2\tilde{C}} \frac{\Omega}{2g} \quad (2.44)$$

The resulting fluorescence signal, J_{fl} , is then:

$$J_{fl} = 2|\alpha_F|^2 \tilde{\kappa} \quad \text{giving} \quad J_{fl} = \frac{\tilde{\kappa}\chi}{2} \left(\frac{\Omega}{g}\right)^2 \left|\frac{2\tilde{C}}{2\tilde{C} + 1}\right|^2 \quad (2.45)$$

The effect of imperfect signal collection between the cavity mode and the detector, the case for any real experiment, is taken into account by including the factor χ , which gives the total collection efficiency.

We can use the optical Bloch Equations, see for example [99], to derive this expression in another way. The steady excited state population, ρ_{ee} , of an atom driven by a laser field with coupling, Ω , detuned by Δ_a is given by:

$$\rho_{ee} = \frac{\Omega^2}{4\Delta_a^2 + 2\gamma'^2 + 2\Omega^2}$$

Here, $2\gamma'$ is the excited state spontaneous decay rate. We know that this is modified by the presence of the cavity, becoming:

$$\gamma' = \gamma \left(1 + \frac{2C}{(1 + \Delta_c^2)}\right)$$

Of this total scattering rate, the partial rate $4\gamma C/(1 + \Delta_c^2)$ goes into the cavity. The rate of photons scattered by the atom into the cavity is therefore:

$$J_{fl} = \frac{4\gamma C}{(1 + \Delta_c^2)} \rho_{ee} = \frac{4\gamma C}{(1 + \Delta_c^2)} \frac{\Omega^2}{4\Delta_a^2 + \left[\gamma \left(1 + \frac{2C}{(1 + \Delta_c^2)} \right) \right]^2 + 2\Omega^2}$$

This can be rewritten in terms of the saturation parameter, $s = \Omega^2/2\gamma^2$ and, for the all-resonant case when $\Delta_c = \Delta_a = 0$, J_{fl} becomes:

$$R_{scatt} = 2C\gamma \frac{s}{(1 + 2C)^2 + s}$$

In the limit of small s this equation is equal to Eq.(2.45).

As with the reflection signal, this treatment can be extended to include multiple atoms. Following the procedure of §3.6. gives:

$$\alpha_F = -\frac{2C}{1 + 2N_{\text{eff}}C} \sum_j^N \frac{g_j}{g_0} \frac{\Omega_j}{2g_0} \quad \text{so that} \quad J_{fl} = 2\kappa\chi \left| \frac{2C}{1 + 2N_{\text{eff}}C} \sum_j^N \frac{g_j}{g_0} \frac{\Omega_j}{2g_0} \right|^2 \quad (2.46)$$

Unlike the equivalent case in the reflection, the total contribution from multiple atoms cannot simply be absorbed into an effective atom number as individual atoms are not necessarily pumped in phase. Instead, we can separate the summation in the expression of J_{fl} into two parts:

$$J_{fl} = 2\kappa\chi \left| \frac{2C}{1 + 2N_{\text{eff}}C} \right|^2 \left(\sum_{i \neq j} \frac{g_i g_j^*}{2g_0} \frac{\Omega_i \Omega_j^*}{2g_0} + \sum_j \frac{|g_j|^2}{2g_0} \frac{|\Omega_j|^2}{2g_0} \right)$$

We can identify the first summed term as being the result of individual, uncorrelated scattering by different atoms inside the cavity, whilst the second term represents ‘ordered’ scattering. This ‘ordered limit’ occurs if we imagined N atoms inside the cavity mode, spatially ordered so that they scattered in phase, or, equivalently, with a uniform pump beam, Ω . In this case Eq.(2.46) becomes:

$$J_{fl} = \frac{\tilde{\kappa}\chi\Omega^2}{2g_0^2} \left| \frac{2\tilde{C}N_{\text{eff}}}{1 + 2\tilde{C}N_{\text{eff}}} \right|^2$$

In this special case atoms are scattering cooperatively and this effect leads to Dicke superradiance [100]. However, we have no reason to believe that in our experiment there should be any spatial ordering, as atoms fall through the cavity taking random positions. In this case, we’d expect the uncorrelated scattering to be the dominant effect and the fluorescence described by the ‘disordered limit’, when the atoms scatter with random phases relative to one another. In this limit, the summation over coupling rates has an average value of $\sqrt{N_{\text{eff}}}$, so that Eq.(2.46) becomes:

$$J_{fl} = \frac{\tilde{\kappa}\chi\Omega^2}{2g_0^2} \left| \frac{2\tilde{C}\sqrt{N_{\text{eff}}}}{1 + 2\tilde{C}N_{\text{eff}}} \right|^2 \quad (2.47)$$




Figure 2.9: The fluorescence signal as a function of N_{eff} , in the ‘ordered’ and ‘disordered’ limits. This is plotted for the system parameters $(g, \kappa, \gamma, \Omega) = 2\pi \times (100, 5200, 3, 10)$ MHz. In the ‘ordered limit’, shown in dark red, J_{fl} initially increases quadratically with atom number, but for higher atom numbers the signal saturates. In contrast, the ‘disordered limit’, shown in light red, initially increases linearly with atom number and then rolls over, peaking when $N_{\text{eff}} = 1/2C$. As $N_{\text{eff}} \rightarrow \infty$, $J_{fl} \rightarrow 0$.

Figure 2.9 shows the qualitative differences between the fluorescence signals in the two scattering limits as they scale with N_{eff} . In the ‘ordered limit’, shown in dark red, J_{fl} initially increases quadratically with atom number and then saturates. In contrast, the ‘disordered limit’, shown in light red, initially increases linearly with atom number and then rolls over, peaking when $N_{\text{eff}} = 1/2C$. In this case, as $N_{\text{eff}} \rightarrow \infty$, $J_{fl} \rightarrow 0$, so that eventually the cavity goes dark.

Whilst these limits provide helpful insight into the behaviour of the fluorescence signal, in general the equation must be solved numerically. Appendix 1 presents some Monte Carlo simulations of the fluorescence signal, although for most situations, Eq.(2.47) is sufficient to understand our experiment.

Having given a theoretical treatment of our experiment, we now turn to practical issues. The following chapter describes our experimental apparatus.

Chapter 3

Experiment set up

The optical microcavities used throughout this thesis had already been used in a successful ‘proof of principle’ experiment [43]. This setup did, however, involve sharing an optical table and laser system with another experiment and it was decided that the only viable means to continue was to become independent. Thus, for the purposes of the work undertaken during this thesis, the cavities were installed on a new optical table and the laser system and computer control system completely re-designed. This chapter describes the new experimental apparatus.

3.1 High level description of experiment procedure and layout

The cavity experiment and supporting laser system is located on a 3×1.2 m floating optical table. The cavity is held inside a vacuum chamber (Kimble spherical square), at a pressure of $< 10^{-9}$ Torr. This pressure is continuously maintained with an ion pump (Varian Vacion 20 Diode).

Inside this vacuum chamber is the optical cavity and a mirror and two magnetic field coils used to make a magneto-optical trap (MOT). These are all rigidly held using two metal posts attached to a vacuum flange. This flange also contains feed-throughs for the MOT coils and the optical fibres which form our cavities. Two rubidium dispensers are fixed on a separate vacuum flange, to allow ease of access and reduce thermal contact with the cavities.

Using this system, we load ^{87}Rb into a MOT and then allow them to expand in a optical molasses, cooling them to $< 40 \mu\text{K}$. Hereafter, the atoms are either pushed using a resonant beam or fall freely under gravity, passing through the cavity positioned 5.1 mm below. Using this procedure we can demonstrate the single atom sensitivity of our cavity system, using either reflection or fluorescence measurements. Figure 3.1 shows a photograph of the experiment




Figure 3.1: The optical microcavity experiment, showing all the major components. The typical position of the MOT is indicated by a red dot.

inside the vacuum chamber, with all the key components labelled.

3.2 The cavities

Central to our experiment is the optical microcavity, which is unique in its design. The cavity is formed between a dielectric mirror glued onto the front facet of a single-mode optical fibre and a micromirror etched into a silicon wafer. This arrangement allows direct optical access to the cavity mode via the optical fibre, which also automatically mode-matches the incoming beam. In addition, the etched micromirror is one of an array microfabricated onto the silicon surface, naturally lending itself to scalability. Figure 3.2 shows a schematic of the cavity.

The cavity design and fabrication, characterisation and testing for the purposes of atom detection was the subject of [97] and is also presented in [43, 86]. Whilst the essentials of the design and fabrication procedure remained the same, the cavity was re-built during the course of this thesis. This section provides a brief overview of the design and characterisation of the cavity. It also includes the scheme used for measuring cavity signals and frequency-locking the cavity.

3.2.1 Design and fabrication

Our initial plan when building the experiment was to use the optical microcavities built as part of [97]. Unfortunately, however, these cavities were mis-aligned whilst being installed in the new setup. For this reason, the cavities were re-built following the procedure outlined in detail

in [97], although the supporting structure which holds the cavity inside the vacuum chamber was altered to reduce thermal contact between the cavities and the MOT coils, for reasons which will be discussed later.

As described, one end of the cavity is a plane, dielectric mirror fixed with index-matching glue onto the front facet of a single-mode, optical fibre. This dielectric coating is custom-made by OIB and was specified to have a reflectivity of 98.5%. The coating received had a reflectivity of $98 \pm 0.5\%$, although subsequent measurements indicate that it is closer to 98%. This was an intended departure from the earlier design, which used a coating with 99.2% reflectivity and was done to improve the collection efficiency between the cavity mode and the fibre.

The optical fibre sits inside a groove which has been machined into a piece of Macor. The fibre is manually aligned to a micromirror on the silicon wafer and then permanently fixed in position with Epoxy (Bylapox, Byla 8725¹). The Macor piece is glued onto a glass block which is held by in aluminium stand, as shown schematically in Figure.3.2.

The other end of the cavity is formed by a micromirror which has been etched onto a silicon wafer. This micromirror is wet-etched using a mixture of HF, HN3 and acetic acid and is then coated with a dielectric, found to have a reflectivity of 99.9%, contingent on finding a smooth micromirror and using a small beam spot size so that only the bottom of the mirror surface is sampled. The radius of curvature of the micromirrors, R , was measured in [97] and found to be $185 \pm 3 \mu\text{m}$.

The silicon wafer is glued onto a piece of Macor and held on a piezo stack, allowing the length of the cavity to be adjusted. The piezo is held in place by two pieces of Macor screwed together.

The use of the silicon wafer in the design means all measurements from this cavity are done in reflection, measuring light reflected back from the cavity through the fibre. This is distinct from most other cQED experiments and has implications for the layout of the detection optics, which are discussed in more detail later in this section.

An important factor for the experiment is the contrast of the cavity fringes, θ . The new cavity, once built and installed inside the vacuum chamber, had a measured fringe contrast of $100 + 0 - 1\%$, representing a significant improvement on the previous cavity, for which $\theta \sim 50\%$. However, the fibre and silicon wafer are extremely sensitive to misalignment, particularly as a result of heating. Inside the vacuum chamber there are two main sources of heat; the MOT coils and the dispenser used as rubidium source.

¹This Epoxy is no longer in production




Figure 3.2: Schematic of the optical microcavity, showing the key features. The optical fibre is glued onto a piece of Macor, which itself is glued onto a glass block held underneath the gold mirror used to form the MOT, as discussed later in this section. A dielectric mirror is glued onto the front facet of the fibre to form one end of the cavity. The other end of the cavity is formed by one of an array of micromirrors etched onto the surface of a Silicon wafer. The wafer is glued onto a piece of Macor held on top of a piezo stack, which allows the length of the cavity to be adjusted. The white arrow indicates the direction travel of the piezo. The piezo stack is held in place using Macor pieces screwed together.

As previously mentioned, the structure supporting the cavities was altered to minimise the effects from heating from the MOT coils. Previously, both the cavities and the coils had been held on the same supporting metal post, so that they were ‘good’ thermal contact. This post, ~ 10 cm long and ~ 5 cm thick, was screwed onto a vacuum flange and pointed diagonally so that the cavities and coils were centred in the chamber. In an attempt to reduce heating via conduction through the post, the design used in this experiment has the cavities and the coils held on separate posts, but arranged to fit together identically as before. Both posts are screwed into the same vacuum flange, allowing the chamber to act as a heat sink. This flange also contains feed-throughs to supply current to the MOT coils and optical feed-throughs for the cavity fibres. Despite this alteration, running current through the MOT coils still has an appreciable effect on the cavities, which is manifest as a slow frequency drift of the fringes. This appears to be caused by radiative heating from the coils and, for this reason, the currents run through them during the experiment are kept quite low, at < 1.5 A, leading to a maximum axial magnetic field gradient at the MOT location of 1.6 G/cm.

It also appears that using the dispensers has a similar thermal effect. To minimise this effect, a copper shield was added to the dispenser flange which blocks the line of sight between the dispenser and cavities whilst still allowing Rb to enter the chamber from the sides.




Figure 3.3: The reflected cavity fringes. LHS: The reflected cavity fringes for the 804 nm light used for locking and trapping (shown in light green) and the 780 nm light used for atom detection (shown in dark green), taken whilst the cavity length is being scanned. The relative amplitudes should not to be directly compared; the 780 nm is measured using a low pass filter to convert APD counts into an analogue voltage for the purposes of locking whilst the 804 nm is measured using a photodiode. RHS: A ‘zoomed’ scan of just one cavity fringe. The dashed line indicates an inverted Lorentzian lineshape, which is a good approximation for the cavity fringes.

With these alterations made the loss of contrast has appeared to stabilise at $75 \pm 1\%$. It should also be noted that the misalignment caused by heating has also caused a reduction by a factor of ~ 2 to the cavity finesse.

3.2.2 Characterisation of properties

For us to fully understand the signals from our cavity several properties must be measured.

The length, d , of the cavity was measured following the method outlined in [97]. The cavity was pumped with light of two colours; 780 nm and ~ 804 nm. The wavelength of the 780 nm was kept constant, locked on resonance with the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition of the D_2 manifold of ^{87}Rb , whilst the wavelength of the other colour was changed. Pumping the cavity in this way produces two sets of reflection fringes, as shown in Figure 3.3. In this scan the lighter green corresponds to the ~ 804 nm light and the darker green to the 780 nm light. The wavelength of the ~ 804 nm was varied by changing the current to and temperature of the laser diode producing it. This changes the position of its fringes relative to those from the 780 nm light. The separation between the lowest order fringes of two colours, λ_0 and λ_i , is given (in

terms of cavity length) as²:

$$\Delta d = \frac{m_i \lambda_i - m_0 \lambda_0}{2}$$

Here, m_0 and m_i signify the longitudinal mode indices. Throughout the measurement, $d \sim \frac{m_0 \lambda_0}{2}$, so that:

$$\Delta d = \frac{m_i \lambda_i}{2} - d \quad (3.1)$$

The changes in λ_i for the ~ 804 nm light were kept small to ensure that the same fringe moved between the m_0 and $(m_0 + 1)$ fringe of the 780 nm light, allowing d and m_i to be found from a least squares fit to Eq.(3.1). This method produced an estimate of $d = 139 \pm 1 \mu\text{m}$, which is in line with the stability condition for a plano-concave cavity.

The cavity mode waist is determined by d and the radius of curvature, R , of the silicon micromirror, using the expression from [97]:

$$w_0 = \sqrt{\frac{\lambda}{\pi}} \sqrt[4]{Rd \left(1 - \frac{d}{R}\right)}$$

From the same text, R was found to be $185 \pm 3 \mu\text{m}$, making $w_0 = 4.5 \pm 0.1 \mu\text{m}$. Combined, these values give a cavity mode volume, $V = \pi w_0^2 d / 4 = 2.17 \times 10^{-15} \text{ m}^3$, which in turn allows us to estimate the maximum atom-cavity coupling rate using Eq.(2.20). This gives $g = 2\pi \times 97.5 \pm 2.3 \text{ MHz}$. Comparing this estimate of g to the atomic transition lifetime, $\gamma = 2\pi \times 3 \text{ MHz}$, indicates that we fulfil one criterion of cQED, namely that $n_{\text{sat}} \ll 1$.

Another factor determining the nature of the interactions inside the cavity is the cavity linewidth, κ . The best way to measure this is to reference a cavity fringe to a known frequency separation, in this case the 6.8 GHz frequency difference between the ‘trapping’ and ‘repumping’ light used to produce a MOT (discussed later). To ensure accuracy, the measurement is done at a slow piezo scan rate, well under the bandwidth of the peizo. The scans are shown on the LHS of Figure.3.4; one with only trapping light circulating inside the cavity (the lower right scan); one with only repumping light circulating inside the cavity (the lower left scan); and one where both trapping and repumping light are circulating inside the cavity (the upper scan). To make an estimate of κ , two types of Lorentzian fit are made to these scans. The first type involves individually fitting the separate scans to single Lorentzians, with κ as a free parameter. This method produces an estimate of $\kappa = 5.27 \pm 0.05 \text{ GHz}$. The second type involves making a double Lorentzian fit to the upper scan, but constraining the amplitudes, fringe visibilities and lineshape centres to be those found from the individual fits, so that κ for each lineshape

²This assumes that $m_0, m_i \gg 1$ and that $\lambda_0 \sim \lambda_i$.




Figure 3.4: Measuring cavity parameters. LHS: Three scans taken to measure the cavity linewidth, κ . The bottom scans show cavity fringes for only trapping light (RHS) and only repumping light (LHS) circulating inside the cavity. The top scan is for both trapping and repumping light, where the slight notch in the bottom of the fringe and increased width indicates that this fringe is the sum of two fringes of different frequency. RHS: The fringe contrast measured as a function of cavity length before gluing.

is the only a free parameter. This produces an estimate of 5.12 ± 0.08 GHz. These fits are shown alongside the scans in the darker green. Taking the mean of these two estimates gives us a final estimate of $\kappa = 5.20 \pm 0.1$ GHz.

The final cavity parameter to characterise is the fringe contrast. This was initially measured as a function of cavity length whilst the cavity was being made — data from this measurement are shown on the RHS of Figure.3.4. This data was used to choose the position at which the fibre was glued, with the aim of maximising the contrast. This is important, as from the figure one can see variations in the contrast by $\sim 15\%$ between $40 - 180 \mu\text{m}$. From this, the cavity was glued with a length of $130 \mu\text{m}$, although the recent measurement discussed earlier indicates that this has since changed. As discussed, at this length we were able to achieve $\approx 100\%$ contrast once installed inside the chamber, although this has since fallen as a result of heating.

3.2.3 Cavity pumping and stabilisation

Our cavity design is such that coupling light into the cavity mode requires only that light is coupled into the single mode optical fibre. During reflection measurements, the cavity mode is pumped by light which is resonant or near-resonant with the $|F = 2\rangle \rightarrow |F' = 3\rangle$. Before being coupled into the fibre, the light used for this purpose is double-passed through an acousto-optic modulator (AOM), allowing the frequency and intensity of the light to be controlled. There is

also a shutter placed in the beam path, allowing the light to be blocked when necessary (for example, when collecting atomic fluorescence emitted into the cavity mode).

The cavity signals, which consist either of the pumping light reflected back from the cavity or fluorescence emitted into the cavity mode, exit the fibre and are measured at an avalanche photodiode, or APD (Perkin Elmer SPCM AQR15). This detector counts single photons with a quantum efficiency, $\eta_{QE} = 60\%$, producing a 5 V, TTL signal for each photon counted. The APD is specified to have a dark count rate, $N_{dark} = 200\text{s}^{-1}$ and a dead time, $\tau_d = 46.5\text{ ns}$, although this latter value was found empirically to be 44 ns [97]. The effect of efficiency, dark counts and saturation can be taken into account by correcting the measured counts per integration time, τ_{int} , according to³:

$$N_{corrected} = \frac{1}{\eta_{QE}} \left(N_{meas} \times \frac{1}{1 - \frac{\tau_d N_{meas}}{\tau_{int}}} - N_{dark} \right) \quad (3.2)$$

In addition to near-resonant light at 780 nm , far-detuned light at 804 nm is also coupled into the mode for the dual use of frequency-locking and atom trapping. The light reflected back from the cavity at this wavelength is detected at a separate photodiode (Femto LCA-S-400K) with transimpedance gain of $1 \times 10^7\text{ V/A}$ and bandwidth DC-400 kHz.

The cavity is actively stabilised during the experiment by frequency-locking to the side of one of the fringes produced by the 804 nm light and providing feedback to the piezo actuator controlling the cavity length. When locking, the location of the 804 nm fringes relative to the 780 nm fringes can be altered by making fine adjustments to the current and temperature of the 804 nm diode laser. To lock the cavity on resonance with the 780 nm light, fine adjustments are made until the bottom of a 780 nm fringe is coincident with the HWHM point of a 804 nm fringe.

The 804 nm diode laser is not actively stabilised, which means that the cavity length is actually being locked to an error signal whose frequency could drift. However, the required performance of the cavity lock is determined by the cavity linewidth; i.e. a narrower linewidth cavity requires a tighter lock. In our case, the cavity linewidth is 5.2 GHz and small changes in the frequency output of the laser have a negligible effect on the cavity. Generally, once settled and without large disturbances, the cavity will remain on resonance for several hours.

In the scheme described thus far, both 804 nm and 780 nm light are circulating inside the cavity and both exit the fibre in the same mode. However, for the purposes of measurement, these two wavelengths need to be separated as stray 804 nm light measured at the APD increases the noise floor of any signals and makes it difficult to lock the cavity at the correct frequency.

³Specified by the manufacturer when $N_{meas}\tau_d/\tau_{int} \ll 1$




Figure 3.5: A schematic showing the arrangement of optics prior to the cavity fibre. This setup is designed to filter 804 nm light from the 780 nm which is detected at the APD. Four dichroic filters are used; three in reflection, labelled R_1 , R_2 and R_3 and one in transmission, T_1 .

This need is compounded by the fact that the 804 nm light will also be used to produce an intra-cavity dipole trap, discussed in Chapter 6, so that there is substantially more of it than 780 nm light (with typically $\sim 500 \mu\text{W}$ compared to $\sim 500 \text{ fW}$). Given this, the layout of the optics prior to the fibre is designed to allow large amounts of 804 nm light into the cavity, whilst filtering the APD from its backward reflection.

A schematic of the optical layout is shown in Figure 3.5. All the optics, the cavity fibre input, Femto and APD are housed inside the black cardboard box to minimise stray light reaching the APD. The 804 nm and 780 nm light are initially combined outside the box at a polarising beam splitter (PBS) and then coupled into the same polarisation-maintaining fibre. The output of this fibre is inside the box. The two colours are immediately separated by polarisation at a PBS 1.

The 780 nm light follows path 1 in Figure 3.5. It is reflected from a 90:10 beam splitter, followed by three dichroic filters, labelled R_1 , R_2 and R_3 (Semrock Razor Edge LP02-785 RU-25), which are designed to reflect 780 nm light and transmit 804 nm light. The 780 nm light then passes through a $\lambda/4$ plate which allows small adjustments to the incoming light's polarisation, and is then coupled into the cavity fibre using two mirrors. Ideally, this polarisation optic would be placed earlier in the path, so that only the 780 nm light passed through it (as it is,




Figure 3.6: The transmission (solid line) and reflection (dashed line) through the dichroic as a function of angle. The same units are used for transmission and reflection amplitudes, so that the scales are one-to-one. This indicates that a loss in transmission was accompanied by an equivalent gain in reflection.

804 nm light passes through it as well). It was placed here because of space limitations.

The 780 nm light reflected back from the cavity follows the reverse of path 1. At the 90:10 beam splitter, 90% of the light is transmitted and then focussed onto the APD with a lens. Between the beam splitter and the APD there is also a transmission notch filter, labelled as T_1 in Figure.3.5 (Semrock MaxLine LL01-780-12.5) designed to pass only 780 nm. Each optical element, particularly the dichroics, result in some losses — although this must be tolerated as the result of separating the wavelengths. The final transmission from the cavity fibre to the APD for 780 nm was measured to be 69%.

Consider now the path of the 804 nm light, after the first PBS this light follows path 2. It is reflected at PBS 2 and then passes through a $\lambda/4$ plate. It is then transmitted through the dichroic filter closest to the cavity fibre (which is designed to transmit 804 nm light) and is coupled into the cavity with the same two mirrors as the 780 nm light.

The 804 nm reflected back from the cavity reverses this path, but having double-passed the $\lambda/4$ it gets transmitted at PBS 2 and is detected at the Femto detector. The signal measured here is used for locking the cavity. Inevitably, the reversal of this path isn't perfect and some 804 nm light leaks into path 1. This justifies the use of three consecutive reflection dichroic filters, followed by a transmission dichroic filter.

The reflection dichroics are extremely sensitive to angle and polarisation. The sensitivity to angle was measured and is shown in Figure 3.6. The 'Dichroic angle' shown on the horizontal axis is just that measured on the rotation stage and has an arbitrary relation to the actual

angle of the dichroic to the incoming beam — but it is shown here to give a sense of the size of the features. In addition, the same units are used for the transmission and reflection on the y axis, so that the changes are one-to-one. It is clear that a loss in transmission translates to an equivalent gain in reflection. The reflection efficiency shows a clear edge at 162° , where the filter begins to reflect 780 nm light. However, beyond this edge there is still substantial structure scanning a range of 61 – 97%. Once this angle was optimised for each filter, however, the combined affect is to attenuate the locking/trapping light by a factor of 10^{11} .

3.3 Atom preparation

This section describes the main features of our atom of choice, ^{87}Rb and our methods for cooling and moving these atoms into the cavity.

3.3.1 Rubidium

The atom used for the experiments discussed here is rubidium. This element has only one stable isotope, ^{85}Rb , although the half-life of a second isotope, ^{87}Rb is sufficiently long for it to be treated as stable for our purposes [101]. The atomic structure of rubidium makes it a common choice of element for atomic physics experiments, as the presence of a single, free, outer electron allows it to be approximated to a ‘hydrogen-like’ atom, simplifying the level scheme. In addition, the relevant atomic transitions can be addressed by relatively cheap and readily available commercial laser diodes.

The relative natural abundances of the ^{85}Rb and ^{87}Rb isotopes are 72.17% and 27.17% respectively, favouring the use of ^{85}Rb in experiments, as was the case in the previous incarnation of the experiment. In this experiment, however, ^{87}Rb is used. The experiment is, however, in the progress of being adapted to allow it to switch between isotopes.

The level scheme for the D_2 manifold (representing the $5^2\text{S}_{1/2} \rightarrow 5^2\text{P}_{3/2}$ transitions) for ^{87}Rb is shown in Figure 3.7. Each hyperfine, F , level is shown with the relative frequency separation between them, although the relative separations are not drawn to scale. On the right of the schematic are two arrows indicating the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cycling transition and the so-called repumping transition using $|F = 1\rangle \rightarrow |F' = 2\rangle$. The cycling transition is used to trap, cool and probe the atoms. The name derives from the weak coupling the atom has to other states once in the $|F = 2\rangle$ state; when scattering a photon the probability of exciting the atom to the $|F' = 3\rangle$ state is 70%. Despite this, 25% of the time the atom will de-excite to the $|F = 1\rangle$ state and the repumping light is used to ensure atoms in this level aren’t lost from the




Figure 3.7: A schematic of the D_2 manifold of ^{87}Rb . The upper and lower level hyperfine states are shown with the splitting between consecutive levels. The relative splitting is not to scale. To the right are arrows indicating the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cycling transition and the $|F = 1\rangle \rightarrow |F' = 2\rangle$ repumping transition.

cooling process.

As mentioned, the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cycling transition is used to probe atoms inside the cavity. As an additional complication, the rate at which atoms couple to the cavity is proportional to the dipole matrix element, μ_{ge} which describes the transitions for different magnetic hyperfine, m_F , states. This value varies across the m_F states, given by the Clebsch-Gordan coefficients [101]. Assuming an average distribution of atoms across the 5 states (i.e. $m_F = -2, -1, 0, 1, 2$), the mean value of $g^2 \sim 0.45g_0^2$, although in principle this could be increased by optically pumping atoms into a maximally coupled state.

3.3.2 The MOT

Rubidium atoms are initially trapped and cooled using a MOT. MOTs exploit a spatially-dependent scattering force to confine atoms in a small region in space. The force is the result of atoms scattering photons from laser beams and its spatial dependence comes from the inclusion of a magnetic quadrupole field, which causes a Zeeman-shift to the relevant atomic transition and makes the scattering rate a function of position. The standard configuration uses three orthogonal pairs of counter-propagating beams, confining atoms in the region where they intersect. In this experiment, however, the construction of the cavity restricts our optical access and so a mirror MOT is used instead [102]. This uses a chip surface as a mirror to reflect two mutually orthogonal, co-planar beams, incoming at 45° to the chip surface. These




Figure 3.8: A schematic for the arrangement for the mirror MOT. A pair of beams diagonal to the mirror surface are reflected to form counter-propagating pairs. Two further beams directed in/out of the page, indicated by the dashed circle, are used to form a MOT. The MOT coils are aligned diagonally either side of the mirror.

two reflected beams have their polarisations flipped from $\sigma^{\pm}$ to $\sigma^{\mp}$ on reflection at the mirror surface, producing two pairs of oppositely-directed and orthogonally-polarised beams. Two beams orthogonal to the reflected beam plane and parallel to the mirror surface provide a third dimension of optical confinement. The final arrangement is shown in Figure 3.8. This arrangement is in common use on atom chips as it allows a MOT to form close to the surface of a chip and, for practical purposes, reduces the number of beams needed from the usual six to four.

To form a mirror MOT requires four correctly polarised beams, frequency-locked to be resonant or near-resonant with the cycling transition, the so-called ‘trapping light’. As previously noted, repumping light is also used to avoid loss of atoms via de-excitation to other states.

In our system, the trapping is frequency-locked ~ 260 MHz red-detuned from the cycling transition and is brought into/close to resonance using a double-passed AOM. The trapping light and re-pumping light are then combined and coupled into a tapered amplifier, where they are amplified to a combined power of ~ 250 mW. The two-colour output beam is then coupled into a Schafer Kirchhoff Fibre-Port. This commercially-available part has four fibre output ports, each of which have output couplers producing collimated 2 cm diameter beams with the appropriate polarisations for a mirror MOT. This is an extremely quick, convenient and space-saving method for producing the MOT beams, but does have the detrimental result

of a substantial power loss, as the Schiefter Kirchoff system is relatively lossy. Approximately $\sim 50\%$ of the light is lost coupling into the Fibre-Port and a further $\sim 50\%$ in the Fibre-Port part itself.

To form an atom trap, light must be used in combination with a magnetic quadrupole field. In this case it is produced by a pair of near-identical coils arranged in an anti-Helmholtz configuration. These MOT coils are held on a rigid, custom-made mount inside the vacuum chamber, aligned along the 45° angle of the laser beams, as was shown in Figure 3.8. The coils are wound from kepton-coated copper wire to ensure vacuum compatibility and each consist of approximately 60 turns. To make a MOT, 1.4 A is run in series through the coils, generating a magnetic field gradient of approximately 1.5 G/cm. This is a relatively small field gradient and does restrict the degree of confinement available but, as was discussed earlier, a restriction to low currents is made to minimise the known heating effect and outgassing in the chamber.

There are also three bias field coils of varying diameters and number of turns held outside the chamber. These coils are aligned with their centres along three orthogonal directions, allowing some fine tuning of the MOT position. The MOT position is critical to us, as once released, atoms fall through a 1 mm diameter hole in the mirror to get to the cavity. The bias coils are also used to optimise the expansion of the cloud once the MOT coils are switched off (i.e. during the ‘molasses stage’, discussed below). The currents used in these coils have changed during the course of the experiment, but are generally $< 1\text{A}$, corresponding to $< 0.1\text{G/cm}$.

We can make an estimate of the number of atoms trapped in the MOT using fluorescence imaging. The camera used for this purpose (Allied Vision Technologies Marlin) is outside the chamber, 20 cm away from the MOT. The collection optics in front of the camera use a 50 mm diameter lens with a focal length of 60 mm, so that the image spans a fraction of solid angle, $d\Omega = \pi 25\text{mm}^2 / 4\pi 200\text{mm}^2$. By measuring the intensity, I , and detuning, Δ , of the trapping beams, the atomic scattering rate can be estimated. This can be used to convert between measured camera counts, N_{counts} and the number of atoms, N_{atoms} , if the image exposure time, t_{exp} and camera efficiency, η_{count} are known, using [101]:

$$N_{\text{atoms}} = \frac{8\pi \left[1 + 4 \left(\frac{\Delta}{\Gamma} \right)^2 + 4 \left(\frac{I_0}{I_{\text{sat}}} \right) \right]}{4 \left(\frac{I_0}{I_{\text{sat}}} \right) \Gamma t_{\text{exp}} \eta_{\text{count}} d\Omega} N_{\text{counts}} \quad (3.3)$$

In the above, I_{sat} is the isotropic saturation intensity for ^{87}Rb [101]. This gives an estimate of the atom number to be 19.4 ± 2.0 million, the large error being dominated by the uncertainty in the distance from the MOT to the image plane. The LHS of Figure 3.9 shows an image of




Figure 3.9: CCD images of the MOT (LHS) after 3 s of loading and the atom cloud after 30 ms of molasses (RHS). The bright disk at the bottom of each image is scattered light from the 1 mm diameter hole in the mirror through which the atoms fall. The LHS and RHS pictures have spatial dimensions 6.5×6.2 mm and 5.5×5.2 mm, respectively, found by calibrating the distance per pixel to the known size of the hole.

the MOT after 3 s of loading.

3.3.3 The molasses

After trapping the atoms in a MOT, we use a further stage of cooling via optical molasses. During this stage, the magnetic quadrupole field is switched off and the frequency of the trapping beams further red-detuned by ~ 50 MHz using the AOM. This results in additional ‘polarisation gradient cooling’, caused by converting the kinetic energy of the atoms into angular momentum. When the atom scatters a photon, transitions between magnetic hyperfine states occur with differing probabilities. Because the atom’s polarisation vector has a spatial dependence (due to the spatial variation of the beams) this can result in a transitions which favour a net loss of kinetic energy, effectively cooling the cloud, see for example [2].

The optimal duration, beam detuning and beam power for the molasses stage were found empirically. To select the duration for the molasses stage, the ratio of the cloud width after 5 ms of free expansion to the cloud width after varying durations of molasses plus 5 ms of free expansion was measured using fluorescence imaging. After the molasses stage, all the trapping light was switched off to allow the cloud to freely expand and, after 5 ms it was turned back on, whilst simultaneously a fluorescence image was taken of the cloud. The respective widths were found by fitting Gaussian profiles to a pixel slice from each image. Ratios < 1 indicate net cooling of the cloud from the original MOT temperature. The results are shown on the LHS of Figure 3.10, where cooling is evident after only 3 ms of molasses, when the first image was taken. The ratio appears to stabilise quickly, with no apparent benefit gained by using durations > 15 ms. Alongside the data is an exponential fit which is constrained to be one for zero duration of molasses stage. The fit indicates a $1/e$ lifetime of 1.7 ms — although this dependence would only be confirmed by further data at sub-3 ms molasses durations. From this, the molasses duration was always kept between 15 – 30 ms. A CCD image of the cloud after 30 ms of molasses is shown on the RHS of Figure 3.9.

The detuning of the trapping beams used for optical molasses was optimised by maximising the amplitude of the normalised reflection signal producing when dropping atoms through the cavity. This changed during the course of the experiment, but was always between 40 – 65 MHz. For the majority of experiments discussed here, the molasses beams were detuned by 62 MHz.

Another factor impacting the effectiveness of the molasses stage was the power in the molasses beams. Again, the optimal power was chosen by maximising the amplitude of the normalised reflection signal, discussed in more detail in the following chapter, the results of




Figure 3.10: Optimising the molasses stage. LHS: The ratio of cloud width with and without a molasses stage, measured as a function of molasses stage duration. The ratio appears to stabilise ~ 0.88 , with no apparent benefit from long durations. An exponential fit is shown alongside the data, with a $1/e$ lifetime of 1.7 ms. RHS: The normalised reflection signal amplitude, measured as a function of molasses beam power for two different frequencies; 47 MHz red-detuned (dark green); and 62 MHz red-detuned (light green). The maximum achievable power, set by the efficiency of the AOM at that frequency, is indicated by the dashed lines. The power is measured using a photodiode located inside the fibre port, which measures the pick-off from a 99:1 beam splitter, the output of which is measured with a voltmeter. In the case of the smaller detuning a peak of ~ 1.85 in the drop signal is visible for powers corresponding to ~ 10 mV. Higher drop signals can be achieved with the larger detuning, although no maximum was measured due to the inability to test higher powers.

which are shown on the RHS of Figure 3.10. The darker green shows the results for molasses beam detuning of 47 MHz, whilst the lighter green shows the results for a detuning of 62 MHz. In both cases, the dashed lines indicate the maximum achievable powers, set by the efficiency of the AOM used to frequency shift the beam. The dark green data shows a peak in the drop signal of ≈ 1.85 for a power (as measured from a photodiode) corresponding to ≈ 10 mV. The light green data indicates that detuning further to the red improves the effectiveness of the molasses stage, with a highest signal of ≈ 2.1 . Showing a maximum at this frequency appears to require more power than was available, the result of the AOM being less efficient at the outer edges of its operating range.

Fluorescence imaging can also be used to measure the temperature of the atoms, again by monitoring the free expansion of the cloud. The MOT is loaded for 3 s, followed by 30 ms of molasses and then the atoms are released and allowed to freely expand. At a time, t , later the trapping beams are switched back on and, simultaneously, an image of the cloud is taken with the camera. By taking images at different times, the expansion of the cloud can be monitored by fitting a pixel slice from the fluorescence image to a Gaussian profile and assuming that the width, σ of the cloud expands according to:

$$\sigma = \sqrt{\sigma_0^2 + \frac{k_B T}{m} t^2} \quad (3.4)$$

Here, m , is the mass of an ^{87}Rb atom, T is the cloud temperature, σ_0 the initial width of the cloud and k_B the Boltzmann constant. The fitted widths are shown in Figure 3.11. The maximum duration of the expansion which can be monitored is limited by the time it takes for atoms to fall onto the gold mirror, which is approximately 20 ms. Fitting to Eq.(3.4) gives a temperature of the cloud of $T = 38.7 \pm 2.4 \mu\text{K}$. A separate measurement of the cloud temperature was made after optimising the beam power during molasses. This was done by fitting a Gaussian function to the ‘atom drop signal’, a method which will be outlined in the following chapter. The method estimated $T = 16 \pm 0.5 \mu\text{K}$ — although here the fit uncertainty doesn’t provide an accurate measure of the true error limits of this measurement, as it relies on the assumption that the cloud is truly Gaussian. That this measurement is over a factor of two lower can be attributed to a further optimised molasses, but is also caused by the nature of the measurement. By using the cavity signal, it is only looking at the portion of atoms which arrive at the cavity — these atoms are likely to be colder as they are not part of the cloud which expands enough, when released from the molasses, that it does not fit through the hole on the mirror which separates the MOT from the cavities.




Figure 3.11: The width of the atom cloud with time. Each point is found by fitting a Gaussian profile to the fluorescence image taken by the camera. Shown alongside the widths is a fit to Eq.(3.4), which estimates $T = 38.7 \pm 2.4 \mu\text{K}$.

3.3.4 Transport and driving beams

Later in this thesis experiments are presented which involve driving the atoms to emit fluorescence using beams transverse to the cavity (in Chapters 4 and 5) and pushing atoms from the MOT to the cavity using a resonant beam (Chapter 5). These processes use beams located vertically above and below the MOT, shown schematically in Figure 3.12.

These beams are delivered to the chamber using polarisation-maintaining, single-mode optical fibres. The output beams are focussed into the chamber using a lens of focal length 150 mm and designed to achieve a $60 \mu\text{m}$ waist at the cavities. The beams are aligned to each other, although this process is awkward due to various obstacles inside the chamber limiting optical access. There are also $\lambda/2$ plates immediately after the fibre outputs to allow control over the polarisation.

Repumping light and tunable, near-resonant light are both coupled into the fibres. The near-resonant light for each fibre is double-passed through an AOM, allowing control of their frequency and amplitude. The AOMs are both driven using RF between 80 – 150 MHz, at powers of up to 1 W. One RF source is a 4-channel digital signal generator (Novatech Instruments 409A) which is capable of providing 4 independent phase synchronous RF outputs at frequencies of up to 150 MHz with powers of up to 4 dBm. A second RF source is from a Versatile Function Generator (Toptica VFG-150). These sources are synchronised to each other using an external quartz oscillator from RS. To drive the AOMs, both RF signals are amplified and pass through RF switches (Mini Circuits ZAS WA-2-50DR+), the latter part allowing them to be switched on and off using TTLs signals.




Figure 3.12: Schematic of the setup for the beams used for pushing atoms and driving fluorescence. The beams are located vertically above and below the MOT and cavities and are delivered to the chamber using single mode, polarisation maintaining fibres.

3.4 Computer control

Control of the experiment and data collection is automated using a data acquisition card (National Instruments PCI 6221) in tandem with a computer control program written in the *LabView* package from National Instruments. The data acquisition card has 10 digital input/output lines and 2 correlated counters with a specified maximum clock rate of 1 MHz.

The program is structured to contain ‘good’ and ‘bad’ timing sections. During the ‘good’ timing section, one of the on-card counters is used as a clock to synchronise the output of eight digital waveforms which control the experiment with up to 500 kHz precision. For a typical experiment, the good timing section will include the final 10 ms of MOT loading, ~ 30 ms of optical molasses and then 100 ms for the atoms to fall through the cavity. Throughout the good timing section, the second counter is used to count TTLs produced by the APD, which are counted in time bins with a duration set by the integration time, τ_{int} . The minimum τ_{int} available is $2 \mu\text{s}$, set by the maximum clock rate used to synchronise the digital waveform output. The length of cabling between the APD and the computer was minimised, as long distances (> 2 m) appeared to result in signal reflections which cause the counter to duplicate the number of counts.

The ‘bad’ timing sections occur at the beginning and end of the experiment sequence,

during which the program runs on computer hardware timing. This section is used for all non-time-sensitive actions, such as initialising equipment, downloading the digital waveforms onto the data acquisition card, loading the MOT and collecting images from the camera and data from the APD.

As mentioned, the active part of the experiment was controlled by digital waveforms outputted by the data acquisition card. Eight parallel waveforms describe high/low TTL sequences for different parts of the experiment. One such waveform controlled the MOT B-field, which was switched on and off using a TTL to a solid-state relay switch on the power supply driving the coils. TTLs also trigger the camera, the shutter to block and unblock the light driving the cavity mode and the AOM controlling the MOT trapping light. For the latter, the TTL was fed into a multiplexer, triggering pre-set analogue voltage levels to control the AOM's frequency and intensity (i.e. allowing the frequency to be red-detuned for the molasses stage). In the case of the top and bottom beams just discussed, RF switches were triggered by TTLs to switch these beams on and off.

3.5 Laser sources

The experiment needs light of controllable frequency and amplitude for a multitude of different functions already described; trapping and cooling atoms in the MOT and optical molasses; pumping the cavity mode and driving fluorescence; and for frequency-locking the cavity and dipole trapping. During the course of this thesis, the laser system providing this light was been completely rebuilt to provide more optical power, improved frequency stability and greater functionality (e.g. more flexibility to detune and switch beams on/off). This section provides an overview of the system and discusses each laser and its respective locking system.

3.5.1 Overview of laser sources

Of the four lasers currently in regular use on our experiment only one, referred to here as the 'Repump' laser, remains from the earlier incarnation of the experiment. One of the new lasers, the Toptica DL100, an extended-cavity diode laser (ECDL), was bought 'off-the-shelf' and is used alongside a new tapered amplifier, the Toptica BoosTA, with the aim of ensuring sufficient optical power. Of the remaining two new lasers, one, the 'Cavity' laser, is a free-running laser made from laser module parts bought from Thorlabs. The other, referred to here as the 'Reference' laser, is an ECDL and was built in-house.

To improve the frequency stability of the experiment, the DL100 is frequency-locked to the

Purpose:	Transition:	Detuning:	Functionality:	Source:
Reference	$F = 2 \rightarrow F' = 3$	$\Delta = 0$	None	Reference laser
MOT, molasses	$F = 2 \rightarrow F' = 3$	$\Delta \leq \pm 100 \text{ MHz}$	Switchable	DL100, double-passed AOM
Repump	$F = 1 \rightarrow F' = 2$	$\Delta = 0$	CW, phase locked to DL100	Repump laser
Top/bottom beams	$F = 2 \rightarrow F' = 3$	$\Delta \leq \pm 30 \text{ MHz}$	Switchable, phase coherent	DL100, double-passed AOM
Cavity pump	$F = 2 \rightarrow F' = 3$	$\Delta \leq \pm 30 \text{ MHz}$	Switchable	DL100, double-passed AOM
Cavity lock	$\lambda = 804 \text{ nm}$	Tunable	Switchable	Cavity laser

Table 3.1: A summary of the light needs on the experiment.

purpose-built Reference laser, whilst the Repump laser is phase-locked to the DL100. All three lasers are ECDLs in the Littrow configuration [103], where optical feedback is provided to a free-running laser diode by reflecting back the first order reflection of the diode from a diffraction grating. This provides an external cavity in which the diode can lase and results in a reduced linewidth by virtue of the increased length of the external cavity. Overall, the linewidth can be reduced from a free-running value of several tens of MHz to $< 5 \text{ MHz}$. Our lasers include the addition of an extra plane mirror [104] which fixes the direction of the output-coupled beam, making it independent of the grating angle. Further linewidth reduction is achieved by actively stabilising the laser frequency, either by controlling the diode current or the length of the external cavity using a control voltage to a piezo actuator.

These lasers are used alongside several double-passed AOMs, which have been introduced in key positions in the experiment to allow individual light paths to be frequency and amplitude controlled. Table 3.1 provides an overview of the different light needs and how they are met in our system.

3.5.2 Reference laser

As its name suggests, the Reference laser is used to provide an accurate frequency reference for locking the DL100. The laser was built in-house and substantial design effort was made to ensure its overall frequency and mechanical stability.

The design of Reference laser uses a mixture of commercially available optics and mounts,




Figure 3.13: TOP: A photograph of the Reference laser in situ. BOTTOM: A photograph of the same laser, but with the aluminium box open to show the internal parts.

all housed inside a machined aluminium box. The laser uses a solid-state laser diode (Sharp GH0781JA2C) with a free-running wavelength of $\sim 784\text{ nm}$ at room temperature. This diode is rigidly held inside collimating lens tube (Thorlabs LT110P-B), which in turn is held inside a fine adjustment mirror mount (Newport Ultima). To form the extended cavity for the diode, a diffraction grating (Newport 53-*-33OH) and prism mirror are arranged in the Littrow configuration in front of the laser diode by gluing them to an aluminium block screwed onto the front of the mirror mount.

The mirror mount holding the laser diode has two fine-adjustment tuning screws for adjusting the horizontal and vertical position of the diode relative to the grating. Adjusting the horizontal position changes the wavelength supported by the cavity. A piezo actuator (Thorlabs AE0203D04F) is also placed inside one of the horizontal screw holes of the mirror mount, allowing the cavity length to be adjusted by applying a voltage across it. Adjusting the vertical position of the diode tunes of the feedback from the first grating order to the diode.

The mount sits on a flat peltier cooler (Thorlabs TEC3-6) which is connected to a temper-

ature controller (Wavelength Electronics MPT2500), stabilising the temperature of the diode. The peltier is glued onto an aluminium block inside the laser box and measurements of the temperatures are fed back to the controller from a thermistor located in a small hole inside it. All of the homebuilt lasers are kept at below room temperature to lower their free-running wavelengths and reduce their tendency to run multi-mode. In the case of the Reference laser, the temperature is controlled at $13.5 \pm 0.1^\circ\text{C}$.

Current for the diode, the temperature control and voltage for the piezo control are fed into the aluminium box housing the mirror mount. This box stands on a thick aluminium plate with four solid steel legs. To reduce the impact of vibrations, the steel legs and aluminium plate are separated by Sorbothane washers around the screws. Figure 3.13 shows photographs of the reference laser in-situ, with the key parts indicated.

The electronics for the piezo and current drivers were built in-house. The design for the current driver is based on [105], modified so that is powered from a single supply. To reduce circuit noise, the circuit is powered by a car battery which is continuously charged from an external power supply. With this system, a stable power output of $\sim 20\text{ mW}$ at currents of $\sim 80\text{ mA}$ is achieved.

To provide a stable and precise frequency reference we use servo-electronics on the current control to the laser diode and the voltage control to the piezo actuator controlling the laser cavity length. The laser is locked on the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cycling transition. The controlling error signal is produced by polarization spectroscopy [106, 107], which results in a sub-Doppler, high-gradient dispersive feature with a zero-crossing near resonance with this transition.

Sub-Doppler polarisation spectroscopy uses two counter-propagating beams, a ‘pump’ and a ‘probe’, passing through a rubidium cell, using the optical setup shown in Figure 3.14. The ‘pump’ beam is circularly polarised and, as such, when passing through the rubidium vapour cell it drives σ^+ transitions between m_F states. This produces an anisotropy in the distribution of atomic population across m_F states, making the vapour birefringent. The probe beam, which is linearly polarised, samples this birefringence, which has the effect of rotating its polarisation. The extent of the rotation is analysed at a polarising beam splitter (PBS) located after the cell, which separates the beam into two polarisation components. Their relative intensities are measured at two subtracting photodiodes (PD 1 and PD 2). The plane of polarisation of the probe beam is initially aligned by a $\lambda/2$ plate before the cell with no pump beam present, so that the net signal from the photodiodes is zero.

The degree of rotation of the polarisation is frequency-dependent and, close to resonance,




Figure 3.14: Schematic of the optical layout for polarisation spectroscopy, with all the relevant parts labelled. $\frac{\lambda}{2}$ and $\frac{\lambda}{4}$ refer to a half-wave and quarter-wave plate respectively. ‘PBS’ and ‘NPBS’ refer to polarising beam splitters and non-polarising beam splitters. ‘PD 1’ and ‘PD 2’ are the two photodiodes. The difference signal from these is used as an error signal.

relatively large rotations in the polarisation are caused as the counter-propagating pump and probe beams are interacting with the same atoms. As the frequency of the beams is scanned across the atomic transition, the net signal from the subtracting photodiodes maps out a dispersive feature, with a zero-crossing on resonance. A typical polarisation spectroscopy feature is shown on the LHS of Figure 3.15.

The zero-crossing on resonance is caused by equal and opposite signals at each photodiode and, as such, is particularly robust to intensity fluctuations in the probing light. In contrast, however, the error signal is extremely sensitive to external magnetic fields, as this can change the transitions driven by light of a given polarization. To reduce the impact of any magnetic field fluctuations, the rubidium cell is held inside two nested shields; a hollow μ -metal tube inside a thick steel box with small slits for the beams to pass through. Following the work of [108], the cell is also placed inside a solenoid, which provides an extremely small ‘aligning’ magnetic field of < 0.1 mT, acting to cancel the Earth’s magnetic field. The RHS of Figure 3.15 shows the impact that small changes to this field can have on the error signal. Here, the signal for three values of the current running through the solenoid is shown, from top to bottom: $I = 0$ mA, $I = 10$ mA and $I = 20$ mA. The middle graph, corresponding to $I = 10$ mA, produces the flattest spectroscopy away from the atomic features and is the current used in the experiment.

As previously discussed, this spectroscopy results in a large-gradient dispersive feature on resonance with the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition. To achieve a fast-response, ‘tight’ frequency lock requires high gain on the original error signal. In any electronic feedback loop, however, an optimum must be reached between the gain and the bandwidth, with an increase in electronic




Figure 3.15: Polarisation spectroscopy, showing the dispersive features across the $|F = 2\rangle \rightarrow |F' = 1, 2, 3\rangle$ of ^{87}Rb and $|F = 3\rangle \rightarrow |F' = 2, 3, 4\rangle$ of ^{85}Rb . LHS: A typical polarisation spectrum, with the lock point is indicated. RHS: Polarisation spectroscopy for three different values of the current through the solenoid surrounding the cell; $I = 0$ mA (TOP); $I = 10$ mA (MIDDLE); and $I = 20$ mA (BOTTOM). The top and bottom graphs have been shifted vertically on the graph for clarity.

gain necessarily resulting in a loss of lock bandwidth. This explains the benefit of a large-gradient error signal; an increase in gradient acts like an increase in gain without compromising on lock bandwidth. It has been previously demonstrated [109] that the amplitude of the feature is dependent on the angle of overlap between the pump and probe beams, with smaller angles resulting in larger signals. As shown in Figure 3.14 this setup uses a non-polarising beam splitter to retro-reflect the pump beam through the cell, allowing the beam overlap to be fine-tuned until the overlap angle is effectively zero. In addition, the ratio of the pump-probe beam powers was varied to optimise the gradient. For our locking system this lead to the using a total power of 2 mW, split equally between pump and probe.

This error signal is used as an input for the locking electronics which servo the laser current and cavity piezo. These electronics were designed and built in-house. The error signal is fed into an error signal circuit which provides some buffering and allows an external offset voltage to be added. The output of this circuit is fed separately to a current servo circuit and a piezo servo circuit, both containing feedback loops with proportional and integral gain. The servo outputs are fed to the respective current and piezo drivers which are then fed into the aluminium box housing the diode and piezo. This system achieves a servo bandwidth of approximately 1 MHz.

The Reference laser has successfully provided a stable frequency reference on our experi-




Figure 3.16: LHS: Saturated absorption spectroscopy for the $|F = 2\rangle \rightarrow |F' = 2, 3\rangle$ transitions using the DL100. RHS: The error signal used to frequency lock the DL100 to the Reference laser. The lock point is indicated and corresponds to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition, 266 MHz red-detuned from the Reference laser.

ment. One point of interest which was noticed towards the end of the project (and with the installation of the phase lock on the repumping laser, which will be discussed later in this section) was a systematic offset of 3 MHz between the position of the polarisation spectroscopy zero-crossing and transition resonance. Unfortunately we cannot comment on the long term stability of this offset and, at this point, haven't identified whether it is an intrinsic shift or an artefact of our apparatus.

3.5.3 DL100

Laser light for trapping, cavity pumping and fluorescence is derived from the Toptica DL100. This laser is accompanied by a Toptica electronics rack, which includes a temperature controller module, a current controller module, a peizo controller module and peizo servo loop module. Additional control of the current is achieved via the modulation input on the laser head, where a FET shunts current away from the diode. This laser is typically run at a current of 190 mA, producing ~ 135 mW.

The DL100 is frequency locked to the Reference laser, offset by ~ 260 MHz, putting it close to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition. The error signal used for locking the laser is shown on the RHS of Figure 3.16, alongside the saturated spectroscopy for the ^{87}Rb $|F = 2\rangle \rightarrow |F' = 2, 3\rangle$ on the LHS. Choosing to lock on this transition, ~ 260 MHz red-detuned from the cycling transition allows us to use double-passed AOMs to independently shift the frequency of certain beams, enabling the necessary flexibility for different functions in the experiment.

The error signal shown above is produced using a frequency offset lock based on [110]. The principle for operating the lock is to stabilise an input signal to the 3 dB point of a steep high pass filter. The lock uses a locking circuit whose first action is to split the input signal, of frequency ω_{in} , equally to travel on two separate branches. The first branch contains one 1 dB attenuator and one 3 dB attenuator in series (MiniCircuits VAT-1+ and VAT-3+ respectively), attenuating the signal by a total of 4 dB. This is followed by a rectifying diode, which produces a positive DC signal.

The second branch contains a 1 dB attenuator and the high pass filter (MiniCircuits SHP-100+) in series. The high pass filter has 3 dB frequency point at ω_{3db} . This is followed by an oppositely-poled rectifying diode, producing a negative DC signal. The two branches are combined and this summed output is used as the error signal. To understand the nature of this signal, it is instructive to consider the output in different frequency limits. When $\omega_{in} < \omega_{3db}$, the signal will be most heavily attenuated on the second branch by the high pass filter, such that when the branches are recombined the positive DC signal from the first branch will be larger than the negative DC signal on the second branch and the net result will be a positive signal. When $\omega_{in} > \omega_{3db}$ will be most heavily attenuated on the first branch, as the high pass filter on the second branch will be passing these frequencies. Thus, the combined signal will be negative. When $\omega_{in} = \omega_{3db}$ the two branches will exactly balance and produce a zero-crossing.

A schematic of the locking circuit and other components is shown in Figure 3.17. As the DL100 is being frequency locked to the Reference laser, the input signal is the mixed-down optical beatnote between the two. To create a beatnote, beams from both lasers are coupled into a single mode optical fibre. The lower frequency component (the difference between the Reference and DL100 laser frequencies) is detected on an AC-coupled fast photodiode (Newport 818-BB-21A). Because the frequency at which the 3 dB point of the filter occurs is fixed by the choice of commercial filter used (in this case at 100 MHz) the input signal is mixed down to ~ 80 MHz by mixing with a tunable voltage controlled oscillator (VCO) (MiniCircuits ZOS-300) at a mixer (MiniCircuits ZX05-1LHW-S). This mixed-down signal is then divided into two branches at a splitter (MiniCircuits ZFSC-2-1-S+) and enters the locking circuit.

As is shown in Figure 3.17, additional components are required for useful operation of the lock. Immediately after the photodiode detecting the beatnote is a power limiter (MiniCircuits VLM-52-S+). This component has a maximum constant output power of 10 dBm when saturated. The optical power in the beatnote and additional amplification before this part is sufficient to ensure it is saturated, to minimise the variations in the power of the signal




Figure 3.17: A schematic of the locking system used to frequency lock the DL100 to the Reference laser. The Reference and DL100 beatnote at $\sim 260\text{MHz}$ is detected at an amplified fast photodiode. This signal has an amplitude sufficient to saturate the limiter, which then provides a constant beatnote signal at 10dBm . The signal gets split, with one branch used as a monitor. The other branch is amplified and then mixed down to $\sim 80\text{MHz}$ at a mixer by mixing with a $\sim 180\text{MHz}$ signal from the VCO. The mixer output is amplified and then split into two branches, producing an error signal which is sent to the piezo driver for the DL100.

entering the locking system. This is important, as any changes in the input signal power are immediately translated to the frequency position of the zero-crossing. This is also shown in Figure 3.18. The LHS shows the full error signal for 8 different input powers; the darkest green corresponding to -9dBm and the lightest green corresponding to -30dBm , in steps of 3dBm . The shape and the frequency of the zero crossing vary with power, with a larger gradient at the zero crossing for larger input powers. The RHS shows the frequency position of the zero crossing, taken from the traces on the LHS, with power. As the power increases the position appears to stabilise at $\sim 81\text{MHz}$, but there is still a variation of $\sim 1.5\text{MHz}$ around this level, which would add directly to the laser linewidth.

To minimise the impact of the power variations, a further experiment was conducted to find the input frequency at which the zero crossing was least sensitive to changes in signal power. The amplitude of the input signal was modulated and the amplitude of the resulting modulations on the error signal were measured as a function of input signal frequency. The results are shown in Figure 3.19; where the darker green data correspond to the measured peak-to-peak amplitude of the error signal modulations and the lighter green data correspond to the average error signal offset. As the input signal frequency is swept, the offset goes through




Figure 3.18: The DL100 error signal as a function of input signal power. LHS: Oscilloscope traces of the error signal shown for 8 different input powers. From dark to light green corresponds to $-9 \rightarrow -30$ dBm in 3 dBm steps. As the power increases the shape of the error signal and position of the zero crossing change, with the gradient around the zero crossing increasing with increasing power. RHS: The frequency position of the error signal as a function of power, for the data shown on the LHS.

zero at 80.6 MHz, as expected. The initial assumption might be to lock at this zero-crossing frequency. However, a minimum in amplitude variations occurs at a slightly higher frequency of 81.2 MHz and locking here minimises the impact of power variations. The offset can be counteracted by adding a DC voltage to the error signal to shift the zero-crossing.

Returning to the lock schematic, after the power limiter the beatnote signal is split; one part is used as a monitor, the second is mixed down by mixing with the signal from a VCO. A low-pass filter (MiniCircuits VLFX-300) is installed immediately after the VCO to reduce the amplitude of higher harmonics from it. The mixed-down signal is amplified (MiniCircuits ZHL-3A) and then split into the two branches for the locking circuit. The output of the locking circuit is the final error signal used as an input to the piezo driver.

This locking system uses high frequencies and can be sensitive to signal reflections and resonances along the length of the cables between components. To reduce stray cabling, where possible direct connections between components are used and, where this wasn't feasible, there are tailor-made SMA cables for the length needed. The entire system is located inside one shielded box.

This system has proved highly robust to disturbances and drift. Standard operating performance means the laser remaining locked and on target throughout the course of a day of experiments. When locked, the beatnote between the Reference laser and DL100 has a FWHM of 1.5 ± 0.1 MHz.




Figure 3.19: The modulation amplitude (dark green) and offset (light green) of the error signal as a function of input frequency. Whilst the zero crossing (i.e. position when the error signal equals zero) occurs at 80.6 MHz, the minimum in the modulation amplitude occurs at the higher frequency of 81.2 MHz, implying that this is where the error signal is least sensitive to power variations.




Figure 3.20: Schematic of the optical layout for the DAVLL, with all the relevant parts labelled. $\frac{\lambda}{2}$ and $\frac{\lambda}{4}$ refer to a half-wave and quarter-wave plate respectively. ‘PD 1’ and ‘PD 2’ are the two photodiodes. The difference signal from these is used as an error signal.

3.5.4 Repump laser

The Repump laser was also built in-house and was inherited from the previous incarnation of the experiment. The mechanical and electronic design are similar to the Reference laser, with the latter being an ‘improved’ version. Its temperature stabilisation system is the same and the laser also uses the same brand of laser diode. The standard operating power achieved with this system is ~ 20 mW at a current of ~ 70 mA.

The Repump is frequency-locked to the $|F = 1\rangle \rightarrow |F' = 2\rangle$ repumping transition. For much of the early work undertaken on this experiment, this was achieved using a dichroic atomic vapour laser lock (DAVLL) [111]. An illustration of the optical set-up for this is shown in Figure 3.20. The principle behind the DAVLL is to magnetically induced dichroism in




Figure 3.21: LHS: Saturated absorption spectroscopy for the $|F = 1\rangle \rightarrow |F' = 0, 1, 2\rangle$ transitions. The dashed line indicates the repumping transition on $|F = 1\rangle \rightarrow |F' = 2\rangle$. RHS: The DAVLL error signal, with the lock point indicated.

an atomic vapour. In this experiment, a Rb cell is subject to a magnetic field produced by permanent magnets held directly around the cell. This B-field lifts the degeneracy of the magnetic sublevels of the hyperfine levels, shifting each symmetrically about the centre frequency. Linearly polarized light, which can be decomposed into σ^+ and σ^- polarizations, is passed through the vapour. If this light is on resonance, each circular polarization will be equally off resonant with the magnetic sublevel it excites and the net absorption of the two combined polarizations is zero. If the light is off resonance, the relative absorption of each polarization will be different. The result is a dispersion curve with a zero crossing at the centre frequency where the laser is locked, as shown on the RHS of Figure 3.21, alongside the saturated absorption on the LHS for the $|F = 1\rangle \rightarrow F = |0, 1, 2\rangle$ transitions. The dashed lines indicate the lock point.

To increase the gradient of the dispersion signal's slope, the vapour cell is heated with a foil heater (MINCO), which is controlled by a bench top temperature controller (Thorlabs TC200). Increasing the temperature of the vapour cell increases the amount of absorption, which in turn increases the size of the DAVLL feature. This can be seen in Figure 3.22, where the gradient of the DAVLL dispersion feature at the zero-crossing is shown for different temperatures, ranging from 20 °C to 60 °C.

One disadvantage of the DAVLL is the need for large magnetic fields; chosen so that the splitting of the magnetic sublevels is of the order of their Doppler widths, thus requiring relatively large magnetic field ($\sim 150\text{G}$). Because other parts of the experiment are extremely sensitive to magnetic fields (such as the polarization spectroscopy), the rubidium cell and




Figure 3.22: The gradient of the DAVLL error signal about the zero-crossing for different vapour cell temperatures.

permanent magnets are housed in a nested shield, identical to that used for polarization spectroscopy. In addition, the position of the zero crossing is sensitive to the input polarisation and, as such, we use a high quality polariser (Thorlabs GTH 10N).

In contrast to polarization spectroscopy, the DAVLL spectroscopy produces a Doppler-broadening feature which has a lower gradient (and hence a lower optical gain) around the zero crossing where the laser is locked. The result is a broader frequency profile, although with the compensating feature of increasing the lock's capture range, improving its general robustness to perturbations. This feature and its relative simplicity (it does not, for example, require any modulation) were the reasons for choosing this locking scheme.

Despite these benefits, we found the frequency drift in the repump frequency had a detrimental impact on the experiment, particularly during long data runs. For this reason the lock was replaced by an optical phase locked-loop (PLL) which phase locks the Repump laser to the DL100. This change not only substantially improved the frequency stability of the repumping light, but it also enables us to coherently address the lambda system formed by the $|F = 1\rangle$, $|F' = 2\rangle$, $|F = 2\rangle$ states. Potential applications which exploit this technology are discussed further in Chapter 6.

A schematic of the phase lock is shown in Figure 3.23. At its heart is an evaluation board which contains a PLL chip (Analog Devices EVAL-ADF4007EB1). This chip includes a phase-frequency detector (PFD) and charge pump which, combined, produce current pulses proportional to the phase difference between two input signals. Also located on the board is a configurable loop filter, which integrates these current pulses. The two signal inputs on the evaluation board come from the 'VCO' and the 'Reference clock'. In this case, the VCO




Figure 3.23: Schematic of the components used in the phase lock. The top left hand side shows the set up for a phase-locked voltage controlled oscillator, whose signal, at 6.035 GHz, is used to mix down the optical beatnote signal to ~ 800 MHz. The top right hand side shows the set up for detecting the optical beatnote, at 6.835 GHz. The mixed down signal is the ‘VCO’ input to the Evaluation Board, where it is compared to a ‘Reference clock’ signal, which is provided by the Meguro Signal Generator. The Evaluation Board contains a phase-frequency detector and charge pump, which outputs current pulses at a frequency proportional to the phase difference between the two input signals. The loop filter integrates these currents pulses, providing a voltage signal used for the peizo and current servo loops.

corresponds to the optical beatnote signal between the Repump and DL100.

The beatnote, which has a frequency of ~ 6.8 GHz, is produced by coupling a beam from each laser into an optical fibre and measuring the mixed output at a high bandwidth photodiode (Thorlabs PDA8GS). This signal passes through a Bias-T (ZX85-12G-S+), which allows the DC power level to be monitored. The AC component is amplified and enters a mixer (ZMX-8GH), where it is mixed down to ~ 800 MHz.

The beatnote is mixed down by mixing with the signal from a phase and frequency locked voltage controlled oscillator (VolCon). The VolCon (Hittite HMC-C028) is locked to a frequency of 6.035 GHz (800 MHz less that the frequency of the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition) using a source-locking microwave counter (EIP 575), which has course and fine modulation outputs. These outputs are summed and are inputted to the voltage control of the VolCon. The output of the VolCon passes through a Directional Coupler (TGD-A1203-20), which taps

off a 20 dB attenuated portion of the signal and feeds back to the EIP 575, thus closing the feedback loop and locking the VolCon.

The mixed down beatnote is one input to the Evaluation Board (AD4007). The second, a Reference clock input, is produced by a Signal Generator (Meguro MSG-2580) frequency generator. This frequency generator produces a signal at ~ 100 MHz which is scaled up by $\times 8$ by a prescaler on the Evaluation Board, producing a 800 MHz signal to compare with the VCO input.

The PFD and charge pump use a comparator to convert the VCO and Reference clock signals into digital signals. The edges of the two signals are compared to measure the relative phase between the signals — the output of this measurement is a series of current pulses which are produced at a frequency proportional to the phase difference between the two signals. The on-board loop filter is used to integrate these signals, producing error signals which are fed back to the current and piezo drivers. The loop filter is formed by an op-amp in the standard inverting transimpedance configuration, with $R \times C_1$ and C_2 in parallel in the feedback path where $C_1 > C_2$. The R_1 branch forms the integrator, and C_2 smoothes ripple. This loop filter has $R = 580 \Omega$, $C_1 = 200$ pF and $C_2 = 20$ pF.

This system provides extremely high gain and requires attenuation of the output error signal by 20 dB to avoid oscillations on the piezo branch. This is evidenced by the error signal, shown in Figure 3.24 alongside the corresponding spectroscopy, which is essentially a step function with an extremely high gradient at the lock point. The lack of structure aside from the lock point is the result of this being a digital lock. The corollary of this is a tight lock, with the beatnote frequency of $799,999,025 \pm 10$ Hz over a one second integration time, measured with an independent microwave counter. It is not known whether this ± 10 Hz limit is set by the stability of the beatnote or the measuring counter but, from this, when locked the beatnote is more stable than the AOM drivers used in our experiment. In addition, the measurement indicates a 75 Hz discrepancy between intended ‘set’ frequency of the mixed-down beatnote and the measured value, although a potential explanation could be that the internal clocks of the source and measuring counter are unsynchronised.

3.5.5 The Cavity laser

The Cavity laser was also a new addition to the experiment and is distinct from the other lasers in not having an automated servo loop to stabilise its frequency. The purpose of this laser is to provide light far off-resonance from the ^{87}Rb $D_{1,2}$ transitions, to be used for locking




Figure 3.24: LHS: Saturated absorption spectroscopy for the $|F = 1\rangle \rightarrow |F' = 0, 1, 2\rangle$ transitions. The dashed line indicates the repumping transition on $|F = 1\rangle \rightarrow |F' = 2\rangle$. RHS: The phase lock error signal.

the cavity and, at some later point, for trapping atoms inside the cavity. The wavelength chosen was $\lambda = 804\text{ nm}$, as a balance between being far enough away in wavelength to be effectively filtered from 780 nm light, whilst still being close enough to produce a relatively deep trap. Calculations following [25] indicate that a far off-resonance trap inside the cavity at this wavelength using $400\text{ }\mu\text{W}$ of power (and taking into account the optical build-up inside the cavity) would produce a trap with a depth of $\sim 1\text{ mK}$.

The Cavity locking laser consists of a solid-state laser diode mounted inside a laser cooled diode mount (Thorlabs LDM21). Temperature and current control are provided by commercially available bench top controllers (Thorlabs LDC205C and Thorlabs TED200C, respectively). The outgoing beam is coupled through a Faraday isolator to protect the diode from back reflections, followed by a polariser. With this system, measured after the polariser, the power achieved is $\sim 25\text{ mW}$ at a current of $\sim 120\text{ mA}$.

After the polariser, the beam passes through an AOM which is used to actively stabilise the intensity of the beam. This is done using only a single pass of the AOM, to reduce any power loss. The frequency shift of several tens of Megahertz caused by the AOM has, for these purposes, a negligible effect and any ‘noticeable’ frequency adjustments are made, as previously mentioned, by controlling the laser diode’s current and temperature. The first shifted beam order exiting the AOM is split at a PBS, where one output is measured by a photodiode and the resulting signal used as an error signal for a locking circuit. The locking signal is applied to the RF amplitude driving the AOM. This system has a bandwidth of $\sim 1\text{ MHz}$ and it appears to be limited by the speed of sound through the crystal inside the AOM. For this reason, the




Figure 3.25: The performance of the BoosTA. LHS: Measured output power versus current, for a constant injection power of 2 mW. At this injection power a maximum output power of 210 mW can be produced. RHS: Measured output power versus injection power, for a constant current of 1500 mA. This indicates that the saturated output power for this current is 375 mW.

beam is focussed and located as closely as possible to the AOM's RF electrodes.

3.5.6 Tapered amplifier system (Toptica BoosTA)

Our final laser is the BoosTA, from Toptica, a commercially available tapered amplifier. It is not, by itself, a standalone laser as it requires injection — however, it is an integral part of the laser system, justifying its inclusion in this section. The BoosTA is used to ensure adequate power for the MOT and is injected with by a Master Oscillator beam, providing an amplified output of identical frequency. The 6 dB gain profile for the BoosTA is ± 10 nm, allowing it to be used to amplify both the trapping and repumping beams used in the MOT.

The BoosTA is specified to operate at a current of up to 2000 mA, with a maximum injection power of 40 mW. With these parameters, a maximum specified output power of 1 W can be achieved. To preserve the lifetime of the diode for this experiment, however, the BoosTA is run well below this range; at a current of 1500 mA and injection powers of < 25 mW. Measurements of the BoosTA's response to current and injection power are shown in Figure 3.25.

The BoosTA used in this experiment is fitted with the Toptica FibreDock component which couples the output beam into a single mode fibre, with an efficiency of $\sim 70\%$. Trap and repump light are coupled into the BoosTA and the output coupled into a Shafter-Kirchoff fibre port which is used for produce a MOT.

This completes the discussion of our laser system and other experimental apparatus. Having discussed this, we now progress to the signals that we measure and their potential for atom

detection.

Chapter 4

Atom detection

There are two types of signals that we can measure from our cavity: ‘reflection’ or ‘fluorescence’ signals. Reflection signals are produced by pumping the cavity mode with resonant or near-resonant light and monitoring the light reflected back. In contrast, fluorescence signals are produced without the cavity being directly pumped and, instead, the coupled atom-cavity system interacts with a pump beam transverse to the cavity axis, triggering the emission of photons. The first part of this chapter describes how these signals are produced and characterises them with respect to pump beam intensity, detuning and polarisation. The second part of this chapter focuses on using these signals for atom detection. To begin with, the ability of our system to measure $N_{\text{eff}} = 1$ is tested; calculating the detection confidence, fidelity and dynamic range of our signals. This is followed by a discussion of the prospects for detecting individual atoms falling through the cavity mode and stationary atoms trapped inside the cavity mode. Initially, however, we return to understanding these signals.

4.1 Producing cavity signals

We begin by presenting the experimental procedure. The initial stages of our method for producing the two types of signals are identical. ^{87}Rb atoms are trapped in a MOT, located approximately 5 mm vertically above the cavity. The MOT loads for up to 4 seconds, trapping approximately 10^7 atoms. Once loaded, the magnetic fields confining the atoms are switched off and, simultaneously, the cooling beams are further red-detuned, allowing the atoms to expand in an optical molasses for sub-Doppler cooling to $< 40\ \mu\text{K}$. After 30 ms the cooling beams are switched off and the atoms fall freely under gravity towards the cavity mode, arriving there approximately 30 ms later. At this point, the method for producing the reflection and




Figure 4.1: The recorded counts from the APD for a single shot of a reflection experiment, taken with $\tau_{int} = 50 \mu\text{s}$. At time $t = 0$ on the graph the MOT has already loaded for 3 seconds and the dashed lines indicate other key points in the experiment. At $t = 10 \text{ ms}$, the MOT field is switched off and the cooling light is red-detuned for the optical molasses. At $t = 40 \text{ ms}$ the trapping light is switched off, allowing the atoms to fall. The signal produced by the atom cloud as it falls through the cavity mode is centred at $t \sim 75 \text{ ms}$, reaching a peak count level of 25 counts in $50 \mu\text{s}$ (or 540 kCounts/sec) compared to a background level of 12 counts in $50 \mu\text{s}$ (or 240 kCounts/sec).

fluorescence signals diverge and the final results are quite different. The reflection signal is considered first.

4.1.1 Cavity reflection signals

Reflection signals are produced by pumping the cavity, typically at a rate of $\approx 10^7$ photons/sec. The light reflected back is monitored at an APD. As discussed in Chapter 2, when the cavity and the pumping light are both on resonance with the free-space atomic transition, atoms inside the cavity absorb light which, for our cavity mirror parameters, leads to an increase in the reflected light level.

We refer to one cycle of loading the MOT, optical molasses and dropping the atoms through the cavity as a single ‘shot’ of the experiment. Figure 4.1 shows an APD trace for a single shot, taken at with an APD integration time, $\tau_{int} = 50 \mu\text{s}$. The ‘atom drop signal’ is visible as an increase in APD counts and lasts $\approx 10 \text{ ms}$. This signal is approximately Gaussian, following the density profile of the dropped atom cloud. Typically, the results of between 10 and 100 shots are averaged to improve the signal-to-noise ratio.

From Eq.(2.43) in Chapter 2 we know that the reflection signal depends on cooperativity,

C_N , which in turn is proportional to the effective number of atoms in the cavity, N_{eff} , with $C_N = N_{\text{eff}}C$. Estimating N_{eff} as a function of time is thus a starting point for understanding these signals. In this experiment, a cloud of atoms of varying density, $D(t)$, falls through a cavity whose field extends to infinity in the radial direction and, under these conditions, neither the number of atoms in the cavity or the volume over which they interact is well defined. Instead, N_{eff} is approximately the product of $D(t)$ and the integrated volume of the cavity mode:

$$N_{\text{eff}}(t) = D(t) \times \frac{d}{4}\pi w_0^2 \quad \text{so that} \quad C_N = C \times D(t) \times \frac{d}{4}\pi w_0^2$$

$D(t)$ is calculated by assuming a Gaussian cloud of fixed atom number, N_A and temperature, T , with a width, $\sigma(T, t)$ which expands with time¹. At the point of its release, referred to now as $t = 0$, the cloud is a distance y_0 above the cavity. The cloud then accelerates towards the cavity at g_{grav} . Given this, N_{eff} becomes:

$$N_{\text{eff}} = \frac{d}{4}\pi w_0^2 \frac{N_A}{[\pi\sigma(T, t)]^{\frac{3}{2}}} e^{-\left(\frac{y_0 - \frac{1}{2}g_{\text{grav}}t^2}{\sigma(T, t)}\right)^2} \quad (4.1)$$

The LHS of Figure 4.2 shows a plot of N_{eff} inside the cavity mode as a function of time, calculated using Eq.(4.1). In this plot, $N_A = 19.7 \times 10^6$ and $T = 38.7 \mu\text{K}$ — values measured previously in Chapter 3 — and $y_0 = 5.1 \text{ mm}$, using the kinematic argument that $y_0 = \frac{1}{2}g_{\text{grav}}t^2$ and knowing t from the arrival time of the peak signal from the APD trace. As expected, N_{eff} has an approximately Gaussian form, with a slight asymmetry (looking closely, the RHS of the Gaussian has a slightly slower slope) caused by the clouds' continued expansion as it traverses the mode. The peak of N_{eff} in the cavity is ~ 1 and represents an upper limit using these initial conditions, as it assumes no atom loss between the MOT and the cavity. This is unlikely to be the case, as the structure supporting the cavity is likely to block the direct path to the cavity for some atom trajectories.

The reflection signal shown in Figure 4.1 thus follows the Gaussian form set by N_{eff} . From Eq.(2.43) in Chapter 2 we know that, in the limit of weak pumping and ignoring detection losses and detection efficiency, the reflection signal, J_{sig} is given by:

$$J_{\text{sig}} = J_{\text{input}} \left(\frac{2C_N + v}{2C_N + 1} \right)^2 \quad (4.2)$$

When the pumping light is resonant with the cavity and no atoms are present there is a minimum in the reflected light, the level of which is determined by the fringe visibility;

¹This follows the same reasoning as that used to estimate the temperature of the cloud in Chapter 3.




Figure 4.2: Understanding the reflection signal. LHS: N_{eff} as a function of time, calculated using Eq.(4.1) and assuming $N_A = 19.7 \times 10^6$, $T = 38.7 \mu\text{K}$ and $y_0 = 5.1 \text{ mm}$. This implies that at the peak in atom density, we would expect there to be ~ 1 atom averaged across the cavity mode. RHS: The APD counts averaged over ten shots of the experiment. The Gaussian-like signal caused by the falling atoms is visible, centred at $\sim 30 \text{ ms}$. Shown alongside the data is a fit, described in the main text, which produces estimates of $N_A = 1.89 \pm 0.08 \times 10^7$, $y_0 = 5.12 \pm 0.02 \text{ mm}$ and $T = 29.7 \pm 0.8 \mu\text{K}$.

$J_{\text{res}} = J_{\text{input}} v^2$. Unless stated, these experiments use resonant light² and, given this, it can be more convenient to use a normalised version of Eq.(4.2), where J_{sig} is expressed in terms of units of the reflected cavity signal on resonance. This is referred to throughout as the ‘normalised reflection signal’, given in ‘lab units’ as:

$$\frac{J_{\text{sig}}}{J_{\text{res}}} = \frac{1}{v^2} \left(\frac{2C_N + v}{2C_N + 1} \right)^2 \quad (4.3)$$

This expression can be used alongside the function for N_{eff} given by Eq.(4.1) to fit the reflection drop signal, with y_0 , T and C_{nn} as free parameters, where $C_{nn} = N_A/C$. One such fit is shown on the RHS of Figure 4.2, alongside the average over ten APD traces of atom drop experiments. The fit gives a peak cooperativity of $C_N = 0.37 \pm 0.02$. Using the measured value of $C = 0.30 \pm 0.01$ from Chapter 3, this implies that $N_{\text{eff}} = 1.2$ at the peak of the drop signal. This is slightly higher than expected, although it can perhaps be explained by the lower temperature estimate of the cloud from the fit. A lower cloud temperature would reduce the rate of expansion, thereby increasing the atom density.

²Actually, throughout these experiments the reference used to lock the cavity was the cavity pump beam. Thus, when the pump beam was detuned from the atomic resonance, the cavity was detuned by the same amount. However, for our system $\kappa \gg \gamma$ and the effect of detuning the cavity by $2 - 3\gamma$ is negligible. In Chapter 5 the full impact of detuning the cavity is considered.

We can use the value for C_{nn} found from the fit to estimate the initial number of atoms in the cloud, finding $N_A = 1.89 \pm 0.08 \times 10^7$. This value falls within the range of the MOT number estimates made in Chapter 3 using fluorescence imaging. The fit produces an estimate for the cloud temperature of $29.7 \pm 0.8 \mu\text{K}$, which lies within the temperature range set by two separate measurements of the cloud temperature discussed in Chapter 3. The distance between the MOT and the cavity mode centre is found to be $y_0 = 5.12 \pm 0.02 \text{ mm}$.

From this, given that 1.2 atoms, when averaged across the cavity mode, lead to $\sim 100\%$ difference between the signal floor and signal peak, this indicates the experiment is sensitive to single atoms. The general agreement between the fit, the free parameters and the various independent measurements indicates that we understand the form of the reflection signal. We now turn to the fluorescence signal.

4.1.2 Cavity fluorescence signals

Unlike reflection, fluorescence signals are measured without pumping the cavity. Instead, top and bottom drive beams (the source of which are located vertically above and below the cavity respectively) pump the atoms from the side, transiently exciting them into the $|F' = 3\rangle$ excited state. From here, atoms decay to the $|F = 2\rangle$ ground state, emitting atomic fluorescence into the cavity mode at a maximum rate of $2C_N\gamma$.

The top and bottom pump beams can be used together or separately and the frequency and intensity of each are independently controlled. The fluorescence into the cavity is collected by the fibre and counted at the APD, with a total collection efficiency χ . The cavity pump beam used in reflection measurements is used here before and after the atom transit to ensure that the cavity is on resonance, but is switched off for the duration of the atoms' flight using a shutter or AOM, which is applied 10 ms after the atoms are released from the MOT.

The results of a single shot of a fluorescence experiment using only one beam are shown in the LHS of Figure 4.3. Here, the cavity is initially pumped at a mean rate (measured at the APD) of 28 counts per $50 \mu\text{s}$. When $t = 20 \text{ ms}$, a shutter blocks the pump beam, so that the count rate falls to ~ 1 count per $50 \mu\text{s}$ within 2 ms. This background level is higher than the APD dark count of $200 \text{ counts s}^{-1}$ and has multiple sources, such as scatter from the room lights and equipment screens. The dominant source however, appears to be from the MOT beams as whilst the cooling beams are switched off to release the atoms, the repumping beams are on continuously. In addition, the background scatter increases in the vicinity of the fluorescence signal, this extra contribution arising from the pump beam itself. Careful alignment of the MOT

beams and pump beams and masking the lower portions of the MOT beam lenses enabled us to reduce the background scatter substantially, to a level of ~ 0.004 counts per $2\mu\text{s}$, or, for comparison, 0.1 counts per $50\mu\text{s}$, as is the case in the RHS figure to be discussed in more detail later.

Returning to the LHS figure, after the shutter is blocked the counts inside the cavity remain at the background level until $\sim 73\text{ ms}$, when a narrow spike in the signal is visible — this is fluorescence. Its amplitude is sufficiently small that the cavity pump beam used in reflection experiments must be switched off in order to see it. The count rate quickly falls back to the background level, however, and at $t = 125\text{ ms}$ the probe shutter is opened, with the light level returning to that at the beginning of the experiment, indicating that the cavity remained on resonance throughout.

The RHS of Figure 4.3 shows the same experiment, but taken at a higher resolution of $\tau_{\text{int}} = 2\mu\text{s}$ and averaged across 250 shots. The figure is ‘zoomed’ in, showing only the fluorescence signal itself. The shape of the signal is well approximated by an exponential decay, a fit to which is shown alongside the data and which gives a $1/e$ lifetime for the signal of $114.0 \pm 0.4\mu\text{s}$. This figure highlights the key benefit of the fluorescence signal, as it produces a signal with a high signal-to-noise ratio when $N_{\text{eff}} \approx 1$. For this plot, the ratio of the peak signal amplitude to the background noise is 135:1.

The fluorescence signal is clearly quite different than the reflection signal. From §4.2 of Chapter 2, we know that the expected fluorescence signal rate measured at the APD, in the limit of weak excitation, is given by:

$$J_{fl} = 2\kappa\chi|\alpha_F|^2 = 2\kappa\chi \left| \frac{2C}{1 + 2C_N} \sum_j \frac{g_j}{g_0} \frac{\Omega_j}{2g_0} \right|^2 \quad (4.4)$$

This indicates that we should expect the signal to last as long as atoms are present in the cavity. Looking again at the RHS of Figure 4.3, we see that when the pump beam is switched on, at $t = 73\text{ ms}$, this quickly translates into a peak in the photon counts at the APD. Within $500\mu\text{s}$, however, the peak has decayed to the mean background count level. The timescale for the duration of this signal is much shorter than the free transit of the cloud. There are two potential causes; the push felt by the cloud when scattering photons from the pump beam, accelerating them out of the cavity mode; or heating.

An estimate of the timescale on which the effect of pushing the cloud will be manifest can be made by estimating how long it takes an atom on the edge of the cloud when the beam is turned on to be pushed through the cavity mode. When centred on the cavity, the Gaussian width




Figure 4.3: Understanding the fluorescence signal. LHS: The counts recorded by the APD for a single shot of the experiment. At $t = 0$ the MOT has been loading for 3 seconds. Dashed lines indicate key points in the experimental; at $t = 10$ ms the MOT field is turned off and the atoms are released into the optical molasses; at $t = 20$ ms the shutter for the cavity pump light is closed and, following a ~ 2 ms delay, the cavity count rate falls to a low background level of ~ 1 count per $50 \mu\text{s}$; at $t = 40$ ms the cooling light is switched off and the atoms fall towards the cavity. At $t \sim 73$ ms a narrow peak is visible, corresponding to the collected fluorescence produced by the atoms as they fall through the cavity. At $t = 125$ ms the pump light is turned on to check that the cavity remained on resonance. RHS: A section of a recorded trace averaged across 250 fluorescence experiments. This data was taken using $\tau_{\text{int}} = 2 \mu\text{s}$ (as opposed to $\tau_{\text{int}} = 50 \mu\text{s}$ shown on the LHS), allowing the structure of the fluorescence signal to be visible. Shown alongside the data is an exponential fit, which gives an estimate of the $1/e$ lifetime of $114.0 \pm 0.4 \mu\text{s}$. The background counts have also been reduced to 0.004 counts in $2 \mu\text{s}$, the result of improved beam alignment.

of the cloud is approximately 5mm. Thus, a late fluorescence signal might be expected from an atom ~ 2.5 mm away from the cavity mode being pushed back in. With the fluorescence beam on resonance, atoms feel a force from the beam equal to $\hbar k R_{scatt}$, where R_{scatt} is the atomic scattering rate. For a beam power of $200 \mu\text{W}$ focussed to a waist of $60 \mu\text{m}$ at the cavities, this produces an instantaneous acceleration of over 100 km/s^{-2} . Assuming constant acceleration at this rate³, it takes an atom approximately $210 \mu\text{s}$ to travel across the cavity mode. Despite these simplifications, this order of magnitude estimate corresponds well to the actual duration of the signal, which has a FWHM of $170 \mu\text{s}$ and has a fitted half-life of $114 \mu\text{s}$. In this model, however, one would expect the density profile of the cloud to be mapped out in the fluorescence signal, albeit in a shorter time frame. The effect of heating may explain why this is not observed.

One can estimate the timescale on which heating will be an important effect by calculating how long it takes for the density of the portion of the cloud addressed by the pump beam to fall substantially, say, to one half of its original value. For an atom performing a random walk caused by repeated photon emissions in random directions, the expectation value of its position obeys:

$$\frac{d^2}{dt^2} \langle x^2 \rangle = \frac{2}{3} v_r^2 R_{scatt} t$$

In the above, v_r is the atomic recoil velocity, in this case equal to 5.88 mm/sec . If one assumes no correlation between the position, x and the velocity, dx/dt , of the atom, one finds that $\langle x^2 \rangle \propto \frac{1}{3} v_r^2 R_{scatt} t^3$. The time taken for the density to fall by one half will be the time taken for the area to double. If we only treat the atoms inside the area addressed by the pump beam, equal to $\pi \times (60 \mu\text{m})^2$, the time taken for this area to double is approximately $540 \mu\text{s}$. This is a similar order of magnitude as the $1/e$ duration of the signal, indicating that heating of the cloud is another important factor.

Thus far, only the case for one beam has been considered. In the case for two resonant, power-matched and counter-propagating beams, one would expect there to be no net pushing force on the atom and that, as a consequence, the signals would last longer. One would also expect heating to be the limiting factor, so that the signal would be expected to decay in a few hundred microseconds. In addition, because the atoms are not accelerated as quickly, they may be expected to stay on resonance with the fluorescence beams for longer, leading to an overall increase in the amplitude of the signal. These qualitative differences are both evident

³This assumption substantially simplifies the calculation. In reality the atoms' accelerations are a function of their detunings, which, due to the Doppler shift, will change with their velocities.




Figure 4.4: Fluorescence signal produced using one beam (shown in the darker red) and two counter-propagating beams (shown in the lighter red). The y-axis corresponds to the number of counts at the APD per $100\ \mu\text{s}$ time bin, summed for five experimental shots. The two-beam signal has an amplitude of nearly three times that of the one-beam case. In addition, the two-beam signal decays more slowly, having a FWHM of $500\ \mu\text{s}$ compared to $300\ \mu\text{s}$ in the one-beam case.

in Figure 4.4, where fluorescence signals for one beam (darker red) and two beams (lighter red) are shown. The two-beam signal has a peak amplitude nearly three times larger than that of the one-beam signal. For these shots, both beams have a power of $315\ \mu\text{W}$ which, from data discussed in the following section, implies that the atoms should already be saturated by one beam. This implies that the amplitude increase is caused by atoms spending longer in the cavity mode, rather than having a higher scattering rate. This is corroborated by an increase in FWHM of the signal, going from $300\ \mu\text{s}$ for the one-beam case to $500\ \mu\text{s}$ for the two-beam case.

The discussion so far indicates that the basic features of the fluorescence signal can be understood in terms of the dynamic response of the atoms to the pump beam. The next section goes a step further, characterising the fluorescence and reflection signal's response to changing experimental conditions.

4.2 Characterising the cavity signals

This section characterises the reflection and fluorescence signal under a variety of experimental conditions. The parameters tested here are the intensity, detuning and polarisation of the pumping beams and the size of N_{eff} .

4.2.1 Intensity

The first measurement treated is how the reflection and fluorescence signals vary with pump beam intensity. One might expect differences to be caused by the onset of atomic saturation, which occurs as the pump beam Rabi frequency, Ω , becomes comparable to the spontaneous decay rate. In the case of the reflection signal, when the cavity is directly pumped, $\Omega = \sqrt{2}g\alpha$, so that the number of photons required to saturate the atom is given by:

$$n_{sat} = |\alpha_{sat}|^2 = \frac{\gamma^2}{2g^2}$$

Thus far in reflection we have always been in the limit of weak pumping, so that the intra-cavity photon number, $n \ll \gamma^2/2g^2$. This allows us to assume weak atomic excitation and to use the simplified expression for α given by Eq.(2.32). When this isn't the case, α is given by Eq.(2.36), which must be evaluated numerically. The normalised reflection signal becomes:

$$\frac{J_{sig}}{J_{res}} = \frac{1}{v^2} \left([1 - v] \frac{1 + 2ng^2/\gamma^2}{(2C + 1) + 2ng^2/\gamma^2} - 1 \right)^2 \quad \text{with } n = \left[\left(\frac{\eta}{\kappa} \right) \frac{1 + 2g^2n/\gamma^2}{1 + 2C_N + 2g^2n/\gamma^2} \right]^2 \quad (4.5)$$

This solution leads to the reflection signal decreasing as the atom saturates. Data demonstrating this effect can be seen in the LHS of Figure 4.5, which show the normalised reflection signal measured as a function of input pump count rate. For this experiment, a neutral density filter with a measured power attenuation of 113:1 was used to shield the APD from the high count rates. Ten shots were taken at each input count rate and the final normalised signal found by averaging. The data in light blue show the case with no repumping light circulating inside the cavity, whilst the data in dark blue show the case with repumping light. Intuitively, one might expect that the onset of saturation would be slower with repump present, although this effect isn't obvious from the data and the fits confirm this. The confidence intervals for the free parameters overlap, indicating that the presence of repumping light does not appear to affect the onset of saturation.

The fits shown are to Eq.(4.5), with the free parameters being g and a parameter quantifying the pumping efficiency, converting input pump rate to cavity pump rate from Eq.(2.40). The fit estimates $g = 2\pi \times (82.4 \pm 1.6)$ MHz. Assuming $(\gamma, \kappa) = 2\pi \times (3, 5200)$ MHz, this makes $C_N = 0.22 \pm 0.01$. This is smaller than the value found earlier, but can be accounted for by two possible factors; a change in fringe contrast, θ , or a change in the initial MOT number. The former definitely has some impact here, as the slow thermal drift of the cavity alignment did cause θ to change over time — for these experiments $\theta = 0.3$, compared to later measurements

of $\theta = 0.24$ shown earlier. Additionally, changes in the overall MOT number would affect this by changing N_{eff} — this is also possible, as the MOT shape and size varied substantially over the course of the experiments and a change of $N_{\text{eff}} = 1.2 \rightarrow 0.8$ would account for this. This is corroborated by unsaturated drop signals taken on the same day which have a smaller overall peak level, reduced by $\sim 25\%$ to ~ 1.6 lab units. Assuming $\theta = 0.3$, this implies a cooperativity of 0.24, reasonably consistent with that found from the fit.

The RHS of Figure 4.5 shows a plot of the normalised reflection signal in lab units plotted for cavity input count rates spanning 0.1 – 500 MCounts/sec, using the fit parameters found above. The light blue data from the LHS are shown alongside and map out the regime where atomic saturation effects begin to take place. It shows that when pumping at rates, $J_{\text{in}} \sim 1$ MCounts/sec, we are safely out of the saturation zone. This is important for detection, as the largest signal amplitude will be achieved when the atom isn't saturated. This will also lead to a necessary balance in the appropriate input powers to use, as in the absence of effects such as saturation, signal-to-noise generally increases with increasing signal power. This issue is discussed in a later section.

Turning now to fluorescence signals, Eq.(4.4) assumes that the atoms are weakly pumped. As the intensity of the pumping beam increases however, this is no longer true and the expression for the fluorescence signal (on resonance) becomes:

$$\alpha_F = -2C \sum_j \left(\left| \frac{g_j}{g_0} \right|^2 \alpha_F + \frac{\Omega_j}{2g_0} \right) \frac{1}{1 + |2g_j\alpha_F + \Omega_j|^2/\gamma^2} \quad (4.6)$$

Again, this must be solved numerically, with the additional complication that each atom ‘feels’ a different g_j and a different pump beam Rabi frequency, Ω_j , so that their contributions will not necessarily add in phase. Numerical simulations of this behaviour are presented in Appendix A, the results of which indicate that we should expect the signal to increase and saturate at intensities $I/I_{\text{sat}} \sim 100$.

To measure the effect of saturation, the number of fluorescence counts, N_{fl} , per integration bin at the APD were measured as a function of pump beam power, P . The beam was kept on resonance and was switched on when the atom density in the cavity peaked, corresponding to $N_{\text{eff}} \approx 1$. The pump beam power was varied by changing the amplitude of the RF power to the AOM controlling the beam. Ten shots were taken for each power and the peak counts calculated by averaging the counts in the first time bin after the beam was switched on for each shot. The counts were corrected for APD saturation and efficiency using Eq.(3.2).

Figure 4.6 shows the corrected fluorescence counts, N_{fl} as a function of pump beam power,




Figure 4.5: The effect of atomic saturation on the reflection signal. LHS: The normalised reflection signal measured for increasing cavity pump rates, taken with (in dark blue) and without (in light blue) repumping light circulating inside the cavity. There appears to be no significant difference between the two and fits to Eq.(4.5) confirm this. Shown alongside the data are fits to Eq.(4.5), where the atom-cavity coupling rate, g and a factor parameterising the pumping efficiency are free parameters. The fit to the case without repump estimates $g = 2\pi \times (82.4 \pm 1.6)$ MHz. RHS: The normalised reflection signal calculated for pump rates between 0.1–500 MCounts/sec using the fit parameters found from the LHS data. The effect of saturation is clear, with the reflection signal reducing substantially for pump rates > 10 MCounts/sec. The vertical, dashed line indicates the typical upper pumping rate limit used for reflection experiments.




Figure 4.6: Peak fluorescence counts, N_{fl} , measured as a function of pump beam power, P . The measured data is shown alongside a fit to Eq.(4.7) which gives $N_{sat} = 34.8 \pm 1.9$ and $P_{sat} = 50.4 \pm 10.6 \mu\text{W}$.

P . From this, there appears to be two regimes of operation; for low pump beam powers ($P \leq 50 \mu\text{W}$) N_{fl} appears to linearly increase with increasing P , whilst at higher powers ($P \geq 200 \mu\text{W}$) N_{fl} appears to saturate. This behaviour can be described by the empirical saturation formula:

$$N_{fl} = N_{sat} \frac{P/P_{sat}}{1 + P/P_{sat}} \quad (4.7)$$

The equation above characterises the dependence of N_{fl} with P in terms of two parameters; the saturated count level, N_{sat} , and the power, P_{sat} , at which $N_{fl} = N_{sat}/2$. Shown alongside the data in Figure 4.6 is a fit to Eq.(4.7), from which we find that $N_{sat} = 34.8 \pm 1.9$ and $P_{sat} = 50.4 \pm 10.6 \mu\text{W}$.

The value of P_{sat} corresponds to a peak intensity of 890 mWcm^{-2} , or $270 I_{sat}$ and is over two orders of magnitude larger than that expected for an atom in free space. This is, however, of a similar order of magnitude to the results from Appendix 1. The difference of a factor of ~ 2.5 between our measured saturation intensity and the value found from simulations could be accounted for by systematic errors in the experiment. One significant effect is the error in the focussing of the pumping beam with respect to the cavity. The lens used to focus the pumping beam into the cavity is positioned to produce a waist of $60 \mu\text{m}$ at the cavity and all the intensity calculations have assumed this value. The exact position of the lens with respect to the cavity is, however, unknown and could be incorrect by as much as 1 cm , as the position of the lens against the vacuum chamber is awkward and difficult to access for accurate measurement. A 1 cm error in the distance between lens and cavity would increase the pumping beam waist to $72 \mu\text{m}$, reducing the estimated saturation intensity from $250 I_{sat}$ to $170 I_{sat}$. The mis-alignment

of the beam would also lead to an error in the estimated saturation intensity due to power loss. The optical path through to the cavities is restrictive and misalignment would lead to power loss from beam clipping and scattering from other surfaces by as much as 50%. Combined, these errors could easily account for the difference.

This saturation experiment can also be used to characterise our collection efficiency, χ . Assuming that when saturated the atom emits fluorescence into the cavity mode at a full Purcell-enhanced rate of $2C_N\gamma$, the saturated count rate, N_{sat} , measured at the APD should be $N_{sat} = 2C\gamma\chi$. Taking $C = 0.3$, the total scattering rate into the cavity mode is 11.4×10^6 photons/sec. This can be compared to the estimated count rate collected at the APD when saturated of N_{sat}/τ_{int} , equal to 696×10^3 photons/sec, implying a total collection efficiency, $\chi_{tot} = 6\%$.

At first sight this appears quite low. However, by comparing the measured count rate to $2C\gamma$, we have implicitly assumed that the pump beam is optimally conditioned to achieve this full scattering rate. In reality, the effect of pumping with light of the wrong polarisation can have a negative effect on the scattering rate; firstly by not driving the atoms with the optimal polarisation; but with a secondary effect that the atoms are driven into states where they don't optimally couple to the cavity. This effect is discussed in more detail in a later section, but from simulations it is expected that the combined effect of these two factors⁴ is to reduce the expected scattering to $\sim 80\%$ of its peak value, with this latter factor being dependent on the value of I/I_{sat} :

$$\text{Reduction factor} = 0.78 \times \frac{1}{2} \frac{0.55I/I_{sat}}{1 + 0.55I/I_{sat}}$$

Assuming this worst case, comparing the measured and expected scattering rates leads to an estimate of χ of 8%. Conversely, were the atoms optimally pumped, we might expect our measured count rate to increase to 870×10^3 photons/sec.

This experiment indicates that, for powers readily available to us, increasing the pumping rate will lead to an increase in fluorescence signal. This is in contrast to the reflection signal, where saturation results in a monotonically decreasing signal, as measured in lab units. These differences will have to be considered when exploring the optimum conditions for detection. Next, the impact of detuning the pumping beams on the reflection and fluorescence signals will be discussed.

⁴The first factor of 0.78 describes the reduction from the optimal signal ($= 1$) as a result of polarisation effects. The second factor describes how much harder it is to saturate as a result of polarisation effects.

4.2.2 Detuning

In this section we return to the limit of weak pumping and test the effect of detuning the pump beams. Here, the normalised beam detuning, Δ_a , is given in units of γ as:

$$\Delta_a = \frac{1}{\gamma} (\omega_L - \omega_{atom})$$

In the above, ω_L is the angular frequency of the pumping beam and ω_{atom} is the angular frequency of the $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition. Scanning Δ_a effectively maps out the lineshape of the atom whilst in the presence of the cavity. For an atom in free space (and neglecting any broadening mechanisms other than natural broadening) this lineshape is a Lorentzian with a FWHM of $\Gamma = 2\gamma = 2\pi \times 6.06$ MHz. From Chapter 2, however, we expect the presence of the cavity to increase the decay rate of the atom by a factor $(2C_N + 1)$. This Purcell enhancement should be manifest in a Purcell-broadened atomic lineshape.

In reflection, the effect of detuning can be taken into account through a modified version of Eq.(4.3):

$$\frac{J_{sig}}{J_{res}} = \frac{1}{v^2} \left(v^2 + \frac{1}{1 + \left(\frac{\Delta_a}{(2C_N + 1)} \right)^2} \left[\left(\frac{2C_N + v}{2C_N + 1} \right)^2 - v^2 \right] \right) \quad (4.8)$$

This function maps out a Lorentzian of width $(2C_N + 1)$ and peak height given by Eq.(4.3) when $\Delta_a = 0$. The detuning of the cavity is $\Delta_c/\kappa \ll 1$, so that we can treat the cavity as being on resonance. The Purcell broadening of the lineshape is visible in the plots on the LHS of Figure 4.7, which show the normalised reflection signal in lab units as a function of Δ_a , for three different values of κ and calculated using Eq.(4.8). (g, γ) are kept constant at $2\pi \times (98, 3)$ MHz respectively. The darkest blue shows the case of $\kappa = 2\pi \times 5200$ MHz or $\kappa \gg g$, which is the regime we work in; the medium-toned blue shows $\kappa = 2\pi \times 141$ MHz, or $\kappa = \sqrt{2g^2 - \gamma^2}$ and is the threshold for the strong coupling regime; the light blue shows $\kappa = 2\pi \times 50$ MHz or $\kappa < g$, well into the strong coupling regime. It is clear that the lineshape gets taller and wider as κ decreases and thus C_N increases. This effect is corroborated by the data shown on the RHS of Figure 4.7, measured by detuning the pump beam by various amounts between $-6\gamma \leq \Delta_a \leq +7\gamma$ using an AOM. At each detuning, ten drop signals were taken and the average signal height calculated using Eq.(4.3). The cavity contrast was measured to be 76%. The data are plotted alongside a fit to Eq.(4.8), with free parameters being C_N and the frequency position of $\Delta_a = 0$. This estimates $C_N = 0.39 \pm 0.01$, in agreement with our measurement at the beginning of the chapter. The plot also contains two sets of dashed




Figure 4.7: The reflection signal as a function of pump beam detuning. LHS: The reflection signal, calculated using Eq.(4.8), shown for three different coupling regimes. (g, γ) are kept constant at $2\pi \times (98, 3)$ MHz respectively, whilst the value of κ is changed. The dark blue shows $\kappa = 2\pi \times 5200$ MHz or $\kappa \gg g$, which is the regime we work in; the medium-toned blue shows $\kappa = 2\pi \times 141$ MHz, or $\kappa = \sqrt{2g^2 - \gamma^2}$ and is the threshold for the strong coupling regime; the light blue shows $\kappa = 2\pi \times 50$ MHz or $\kappa < g$, well into the strong coupling regime. The increased height and width of the lineshape as κ reduces (and thus C_N increases) is visible. RHS: The normalised reflection signal measured as a function of pump beam detuning (shown in units of γ). A fit to Eq.(4.8) is shown alongside the data, which estimates $C_N = 0.39 \pm 0.01$. The inner pair of dashed lines indicate the bare-atom linewidth of γ . The outer pair of dashed lines indicate the Purcell-enhanced linewidth of $2C_N\gamma$.

lines to emphasise the broadening; the inner pair indicate the bare-atom FWHM, located at $\pm\gamma$ and the outer pair indicate the cavity-enhanced FWHM, located at $\pm 1.78\gamma$.

Before firmly attributing this broadening to Purcell enhancement, however, other mechanisms should be eliminated. Collisional broadening can quickly be dismissed as the atom cloud is dilute, with a collision rate < 1 mHz⁵. From the previous section we know that power broadening can also be eliminated, as these experiments were taken in the weak pumping regime, with an intra-cavity photon number typically two orders of magnitude below n_{sat} . Transit time broadening, caused by the finite interaction time between the atom and cavity, would have a broadening effect of ≈ 35 kHz — assuming the cavity mode waist of $\pm 4.5 \mu\text{m}$ to be the interaction region and an atom velocity of 32 cm s^{-1} . This is clearly too small to account for

⁵The collision rate equals $1/\tau_{coll}$, where τ_{coll} is the mean time between collisions. $1/\tau_{coll} = n\sigma\bar{v}$, where n is the atom number density, σ is the collision cross section and $\bar{v}$ is the mean relative velocity. Taking $n = 1.3 \times 10^{13} \text{ atoms/m}^3$, measured using fluorescence imaging; $\sigma = 8\pi a^2$, where $a = 100a_0$ and a_0 is the Bohr radius; and $\bar{v} = 4\sqrt{k_B T/\pi m}$; a cloud with temperature $30 \mu\text{K}$ will have a collision rate of 0.27 mHz .

the broadening seen in Figure 4.7. Another potential source might be magnetic field variations, which would cause differential Zeeman shifts of the atoms' magnetic hyperfine levels. Typically Zeeman shift energies are $\sim \mu_B B$, where B is the magnetic field and μ_B the Bohr magneton. This amount of broadening would require a B-field of ~ 8 Gauss, which is too large to be caused by a stray field or our MOT coils, particularly as our MOT field gradients are small, at 1.5 G/cm. The broadening is also larger than the laser linewidths used in the experiment, which typically have linewidths < 1 MHz.

We therefore interpret the broadening evidenced in Figure 4.7 as the result of a modified decay rate of the atom, caused by the presence of the cavity. It is demonstrated in the following section that the reflection signal linewidth is also proportional to N_{eff} , confirming this conclusion.

We can test for the equivalent effect in fluorescence by conducting a similar experiment and where now Δ_a refers to the detuning of the transverse pump beam. A modified version of Eq.(4.4) takes into account the effect of detuning:

$$J_{fl} = \frac{\Omega^2}{4g^2} \left(\frac{2C_N}{2C_N + 1} \right)^2 \frac{1}{1 + \left(\frac{\Delta_a}{(2C_N + 1)} \right)^2} 2\kappa\chi \quad (4.9)$$

As with the reflection signal, the expected fluorescence signal is a Lorentzian with width $(2C_N + 1)\gamma$. This is plotted on the LHS of Figure 4.8, for the same three regimes shown earlier. Light to dark red corresponds to the transition from strong coupling to the case when $\kappa \gg g$ and $\Omega = 2\pi \times 10$ MHz for each. From the plot it is clear that the fluorescence follows the same pattern as in reflection, showing a taller and wider Lorentzian as C_N increases.

Measurements of the fluorescence signal as a function of pump beam detuning are shown on the RHS of Figure 4.8. The frequency of the pump beam was changed by varying the frequency of the RF power to the AOM controlling the beam. The beam power was $10 \mu\text{W}$. Ten shots were taken for each detuning and N_{fl} is the average of the first ten bins after the beam is switched on, averaged across the ten shots, after having been corrected for APD dead time and quantum efficiency using Eq.(3.2). From the figure it is clear that the data are consistent with a Lorentzian lineshape.

Shown alongside the data is a fit to Eq.(4.9), where the lineshape width and a factor quantifying the pump rate and total collection efficiency are free parameters. Assuming $(\kappa, \gamma) = 2\pi \times (5200, 3)$ MHz, this estimates $C_N = 1.01 \pm 0.16$, equivalent to a lineshape broadening of $2\pi \times 3.02$ MHz. This is substantially larger than values of C_N found through other methods, suggesting that there are other broadening mechanisms contributing to the spectral width.




Figure 4.8: The fluorescence signal as a function of pump beam detuning. LHS: The calculated fluorescence signal plotted for three coupling regimes, using Eq.(4.9). (g, γ, Ω) are kept constant at $2\pi \times (98, 3, 10)$ MHz respectively, whilst the value of κ is changed. The dark red shows $\kappa = 5200$ MHz or $\kappa \gg g$, which is the regime we work in; the medium-toned red shows $\kappa = 141$ MHz, or $\kappa = \sqrt{2g^2 - \gamma^2}$ and is the threshold for the strong coupling regime; the light red shows $\kappa = 50$ MHz or $\kappa < g$, well into the strong coupling regime. The increased height and width of the lineshape as κ reduces (and thus C_N increases) is visible. RHS: The average peak fluorescence counts, N_{fl} , collected at the APD in a $50 \mu s$ time bin versus detuning, Δ_a of the pumping beam. The fit shown is to Eq.(4.9), which estimates $C_N = 1.01 \pm 0.16$.

Broadening mechanisms such as laser linewidth and fluctuating B-fields have already been discussed and ruled out as affecting the reflection scan and there is no reason to suspect them to have a greater effect here. In this case, the transit time broadening is expected to be larger, as the atoms are being pushed out of the cavity mode by the pump beam. Despite this, their expected transit time is $\sim 10 \mu\text{s}$, leading to a potential broadening of $\sim 100 \text{ kHz}$, which isn't large enough to account for this discrepancy. Another potential broadening mechanism is power broadening. For this data, the pump beam power was $10 \mu\text{W}$, which should be compared to saturation intensity empirically found for fluorescence in the earlier section of $50.4 \pm 10.6 \mu\text{W}$. Given that the lineshape broadening effect goes as $\Gamma \sqrt{(1 + I/I_{\text{sat}})}$, this would have a broadening effect of 10 % to the bare atom lineshape, which again appears to be insufficient in this case.

We might also expect Doppler broadening caused by the acceleration of the atoms by the pump beam. The magnitude of this effect can be estimated by estimating the maximum expected Doppler shift. The first ten bins after the beam is averaged to find the average counts, so that the data shows evidence of the atoms being subject to the beam for up to $500 \mu\text{s}$. From the previous section, for the beam power used here of $10 \mu\text{W}$, we'd expect a total scattering rate of $2\pi \times 800 \text{ kHz}$, so that we'd see the average Doppler shift from 400 scatters. If each scatter leads to a shift of $2\pi \times 7.54 \text{ kHz}$ [101], we'd expect the largest Doppler shift to be $2\pi \times 3 \text{ MHz}$, which is sufficient to explain our data.

These experiments have confirmed that, at least in reflection, we see Purcell enhancement of the atomic lineshape caused by the presence of the cavity. As an aside, the reflection and fluorescence spectra appear qualitatively similar — the following Chapter will explore conditions when this isn't the case.

Considering how this will impact detection, in both cases detuning the atom leads to a monotonically decreasing signal, which would necessarily result in a decrease in signal to noise. From this, assuming that we're detecting atoms by only monitoring changes in signal amplitude (as opposed to a measuring a dispersion), detuning clearly does nothing to improve our sensitivity.

Thus far, however, we've assumed that C_N increases linearly with N_{eff} — this is tested in the following section.

4.2.3 Effective atom number

We have assumed throughout that C_N is linearly proportional to N_{eff} . This can be tested in reflection by measuring the spectral width, $\propto (2C_N + 1)$, for different values of N_{eff} . This




Figure 4.9: The reflection and fluorescence signals as a function of N_{eff} . LHS: The reflection spectral width as a function of atom number density at the cavity. There appears to be a linear dependence and a linear fit shown alongside the data gives a y -intercept of 1.09 ± 0.04 and a gradient of $(0.06 \pm 0.01) \times 10^8 \text{ cm}^3$. RHS: The fluorescence counts measured in $5 \mu\text{s}$ as a function of N_{eff} , calibrated using reflection signals. A fit to the function $y = ax^b$, shown alongside the data, gives an estimate of $b = 1.05 \pm 0.06$, confirming the linear dependence of fluorescence signal with N_{eff} in the regime of low atom number.

should result in a linear dependence with a y -intercept of one.

For this experiment, N_{eff} was varied by changing the number of atoms loaded into the MOT using different MOT fill times. This changes the value of N_A in Eq.(4.1), thereby changing the atom density at the cavity mode and consequently N_{eff} . The fill times ranged between 100 ms and 6 s and, for each one, a pump beam lineshape was taken to calculate the spectral width. The atom density is estimated in two steps: firstly, fluorescence imaging of the MOT is used to ascertain N_A and, secondly, drop signal traces taken on resonance are used to estimate the size of the cloud at the cavity. The calculated atom densities are quite rough, due to large systematic uncertainties in the fluorescence imaging, the assumption that there is no atom loss during the molasses stage or when falling and the assumption of a simple Gaussian density profile. We suspect, however, that these effects depend only weakly on the MOT number and should not affect any functional dependence on the lineshape width.

The data are shown in the LHS of Figure 4.9, alongside a linear fit which gives a y -intercept of 1.09 ± 0.04 and a gradient of $(0.06 \pm 0.01) \times 10^8 \text{ cm}^3$. We can relate the gradient directly to the cooperativity using the cavity mode volume, V :

$$\text{Gradient} = \frac{\delta(1 + 2C_N)}{\delta(N_{\text{atoms}})} = 2C \frac{\delta(N_{\text{eff}})}{\delta(N_{\text{atoms}})} = 2CV \frac{\delta(N_{\text{atoms}})}{\delta(N_{\text{atoms}})} = 2CV$$

Given our measured cavity volume from Chapter 3, this gives $C = 0.14 \pm 0.03$, which is in

reasonable agreement with other previous measurements.

As discussed, in the absence of other affects we would expect a vertical offset on the y -intercept of one. The discrepancy in this case is equivalent to lineshape broadening of 270 ± 60 kHz and, from earlier discussions, potential sources of broadening of this scale include transit time broadening, fluctuations in the magnetic field and laser linewidth. In addition, this calculation has implicitly assumed that all the MOTs were the same temperature. This is unlikely to be the case as smaller number clouds are colder and larger MOTs are known to have areas of leakage into diffuse outer shells. This could lead to a systematically overestimating the atom density for larger MOTs, which would lead to a larger gradient and smaller y -intercept.

Knowing that we have a linear dependence, we can now use the reflection signal to investigate the relation between the fluorescence signal and N_{eff} . Because N_{eff} is a function of time, switching on the pump beam at different times allows different values of N_{eff} to be sampled. The pump beam switch-on times are varied between 71.5 ms and 80 ms, corresponding to a range of $1 \geq N_{\text{eff}} \geq 0$. At each switch-on time, ten shots were taken with $\tau_{\text{int}} = 5 \mu\text{s}$ and the peak counts found by averaging the first bin after the beam was switched on for each shot. To minimise the impact of atoms being pushed by the beam back into the cavity, the tail end of the profile is sampled. The results are shown in the RHS of Figure 4.9, alongside a fit to the function $y = ax^b$. This gives an estimate of $b = 1.05 \pm 0.06$, indicating that in the regime of low atom number there is a linear relationship between N_{eff} and the fluorescence signal. From Chapter 2, this confirms our hypothesis that we are in the limit of ‘disordered’ scattering — the equivalent dependence in the ‘ordered’ limit would be quadratic. In addition, the fit indicates that we would expect to measure approximately 2.6 fluorescence counts per atom.

4.2.4 Polarisation

We can also characterise the effect of changing the polarisation of the pump beams. As discussed earlier, this is expected to have a substantial impact on the signals. For these experiments, the polarisation of the pump light is linear, but the angle of orientation with respect to the cavity is varied; rotating between being linearly polarised along the cavity axis, (corresponding to π -polarised light) and linearly polarised transverse to it (corresponding to superpositions of $\sigma^{\pm}$ -polarised light).

The experiment was only conducted for the fluorescence signal, as our setup makes the equivalent reflection experiment awkward. In reflection, the plane of polarisation of the light pumping the cavity mode is uncertain due to the presence of the optical fibre forming part




Figure 4.10: The fluorescence counts as a function of pump beam polarisation. LHS: The peak fluorescence counts recorded at the APD in $50 \mu\text{s}$, measured as a function of waveplate angle. The data show a sinusoidal variation. Fitting to a sinusoidal function of fixed period 180° estimates a peak/trough contrast of 1.81 ± 0.14 , which compares well to master equation simulations discussed in the main text. It should be noted that the zero in degrees is arbitrary and just corresponds to $2 \times$ the waveplate angle. RHS: Fluorescence lineshapes taken at polarisation angles corresponding to the maximum (290°) and minimum (200°) respectively). Shown alongside the data are the respective fits to Eq.(4.9). For the ‘max’ data, this fit estimates $C_N = 1.01 \pm 0.16$. For the ‘min’ data the fit estimates $C_N = 0.77 \pm 0.24$. Dashed lines indicate the FWHM for each dataset; the outer pair show the FWHM for the ‘max’ data, the middle pair the FWHM for the ‘min’ data and the inner pair for the atom in free space.

of the cavity; this fibre is not polarisation-maintaining and causes an unknown and variable amount of polarisation rotation. This problem does not occur in fluorescence, as the waveplate used to adjust the orientation of the plane of polarisation sat after the beam’s output from a polarisation-maintaining fibre and all other elements in the optical path between the waveplate and the cavity (such as the vacuum chamber window) would produce a fixed amount of rotation.

One complicating factor in performing this experiment was that rotating this waveplate leads to small changes in the beam direction which, as a result of the relatively long optical path of ~ 50 cm between the waveplate and the cavity, causes mis-alignment. To mitigate this effect, the beam was manually re-aligned between each run, although re-producing the exact alignment between waveplate angles was impossible. This is perhaps one explanation for the high degree of fluctuations evident in the data presented here.

As mentioned, the polarisation of the pump beam is varied by rotating a half-wave plate

placed before the beam enters the vacuum chamber. At each waveplate angle ten shots were taken and the average peak count used to measure the fluorescence counts, N_{fl} . This is shown on the LHS of Figure 4.10. As before, these counts are corrected for APD efficiency and saturation using Eq.(3.2). The power in the pump beam was kept constant throughout. It should be noted that the zero in degrees is arbitrary and just corresponds to $2\times$ the waveplate angle, although afterwards we checked the angles with respect to the cavity axis to confirm that the angles at which the signal was minimised corresponded to π -polarised light.

The data shown on the LHS of Figure 4.10 corresponds to one full rotation of the waveplate, equivalent to a $2 \times 360^\circ$ rotation in the polarisation angle of the beam. Evident in the data is a sinusoidal variation, with a period of 180° ; every 90° the count level moves from a peak to a trough, as the light rotates from having its plane of polarisation transverse to the cavity axis to having it parallel to the cavity axis. Fitting this data to a sinusoid with a period of 180° estimates a peak/trough contrast of 1.81 ± 0.14 .

It is interesting to consider the physical explanation for this. For an atom with no internal degrees of freedom, conservation of momentum requires that the atom re-emit light with the same polarisation as it absorbed. The cavity does not support a π -polarised mode and so a spinless atom cannot scatter light into the cavity from a π -polarised pump beam. In this case, as the plane of polarisation of the pumping light is rotated, the fluorescence collected by the cavity will vary from a minimum of zero to some maximum value, with a period of 180° .

A rubidium atom has more freedom owing to its internal structure; quantising its orbital, spin and nuclear angular momenta result in discrete magnetic hyperfine levels, $|m_F|$ states. Transitions between these levels can involve a change in angular momentum, Δ , of $\Delta = 0, \pm 1$, with probabilities given by Clebsch-Gordan coefficients for each transition. Pumping with π -polarised light favours transitions involving no net change in angular momentum, resulting in the atomic population congregating in the central, low $|m_F|$ states. Pumping with $\sigma^\pm$ -polarised light favours transitions with a net angular momentum change of ± 1 , which optically pumps the atom into the high $|m_F|$ states.

The emission of fluorescence into the cavity by the atom is still restricted to being $\sigma^\pm$, but the presence of internal structure means that there is no longer complete extinction of the fluorescence when the atom is pumped with π -polarised light. Some variation in the collected fluorescence, again with the period of 180° , would still be expected.

To further corroborate the physical explanation for this variation, a master equation simu-




Figure 4.11: The excited state populations as a fraction of total population for the magnetic hyperfine levels $m_F = -3, -2, -1, 0, 1, 2, 3$, calculated from the master equation simulation discussed in the main text. LHS: Fractional populations in the case of pumping by $\sigma^\pm$ -polarised light. The atomic population builds up in the higher magnitude m_F levels as transitions with a net angular momentum change of ± 1 are favoured. RHS: Fractional populations in the case of pumping by π -polarised light. Population builds up in the lower magnitude m_F states as transitions with no net angular momentum change are favoured.

lation was performed⁶ which treated atoms in free space and follows these steps:

1. The excited state fractions, ρ_{e,m'_F,e,m'_F} , for each of the $F' = 3$ hyperfine levels are calculated for an atom in free space,
2. The scattering rate produced from radiation of a given polarisation of light, $-q$, is calculated, using:

$$R_{scatt,-q} = \sum_{m_F m'_F} |\langle F' m'_F | F 1 m_F q \rangle|^2 \rho_{e,m'_F,e,m'_F} \quad (4.10)$$

where $\langle F' m'_F | F 1 m_F q \rangle$ gives the Clebsch-Gordan coefficient for the transition $|F, m_F\rangle \rightarrow |F', m'_F\rangle$,

3. The relative scattering rate into the cavity is calculated by assuming the cavity will only ‘accept’ $\sigma^\pm$ polarisations.

This simple model of the experiment predicts a contrast between minimum and maximum scattering rates of 1.75, which corresponds well with the value found experimentally. The excited state populations for the magnetic hyperfine levels calculated in Step 1 are shown in Figure 4.11. Both the LHS and RHS show the excited state populations as a fraction of total

⁶This simulation was performed by J. Goldwin, then a post-doctoral researcher at CCM.

population across the magnetic hyperfine levels $m_F = -3, -2, -1, 0, 1, 2, 3$. On the LHS, the populations are calculated for the case of $\sigma^\pm$ -polarised light and it is clear that the population tends towards the higher magnitude m_F levels, as transitions with a net angular momentum change of ± 1 are favoured. The RHS shows the populations for the case of π -polarised light, where the population congregates at lower magnitude values of m_F , as transitions with no net angular momentum change are favoured.

These polarisation effects should also be manifest in fluorescence lineshapes taken with pumping light of different polarisations, as each m_F state will couple to the cavity at a different rate. Two fluorescence lineshapes are shown in the RHS of Figure 4.10; the data in darker red taken using a polarisation corresponding to a minimum in the fluorescence counts (200°), the ‘min data’; and the data in lighter red taken at a polarisation corresponding to a maximum in the fluorescence counts (290°), the ‘max’ data. Alongside both datasets are their respective fits to Eq.(4.9), used to compare their relative widths.

The larger amplitude evident in the max data appears to confirm that there is a larger cooperativity in this case. This does not, however, extend to comparisons of the lineshape widths, which both cover a large range consistent with being equal. The max data, presented earlier in §2.2., was found from the fit to have $C_{N,max} = 1.01 \pm 0.16$. A fit to the min data gives an estimate for $C_{N,min} = 0.77 \pm 0.24$. These give atom-cavity coupling rates of $g_{max} = 2\pi \times (178 \pm 28)$ MHz and $g_{min} = 2\pi \times (155 \pm 48)$ MHz. Both are substantially larger than other measurements of C and g , indicating, as discussed earlier, that an additional broadening mechanism is at work. From that discussion, Doppler broadening was found to be the most likely candidate and it is possible that this is masking any differences in cooperativity between the polarisations.

What is demonstrated by these experiments is the importance of polarisation to ensure a maximum signal. Throughout the discussions of part two of this chapter, it will be assumed that the optimal polarisation (in this case that which corresponds to the ‘max’ data) is being used.

4.3 Single atom detection

From the data and discussion presented thus far, it is clear that this experiment is *sensitive* to single atoms; when $N_{\text{eff}} = 1$ there is a substantial impact on the reflection signal and a clearly detectable fluorescence signal. One application for this single-atom sensitivity is as a chip-based atom detector. The next stage is to attempt to quantify the single atom detection

confidence, in varying detection conditions.

4.3.1 Signal to noise ratio and Poissonian statistics

A standard method for quantifying the confidence with which a signal can be detected is via the signal-to-noise-ratio (SNR), which is a measure of the size of the signal compared to fluctuations in the signal. Both the signal and noise are assumed to have Poissonian fluctuations about some mean level and, to take this into account, a ‘pooled’ version of the SNR is used here, defined as:

$$\text{SNR} = \frac{N_{sig} - N_{bg}}{\sqrt{N_{sig} + N_{bg}}} \quad (4.11)$$

Here, N_{sig} refers to the number of signal counts; in our case this is the number of counts measured at the APD as a result of atoms being in the cavity. N_{bg} is the number of noise counts measured at the APD and is a sum of all sources of background noise. For our experiment, this latter quantity will include scattered photons from the light pumping the atoms, from repumping light used for trapping and from other miscellaneous sources such as equipment screens and room lights. There will also be a contribution from the off-resonant light circulating inside the cavity, used for producing an error signal for locking the cavity.

As mentioned, this expression for SNR is only valid when the signal and noise follow Poissonian statistics, one characteristic of which is that the signal variance is equal to the signal mean, or the equivalent for the background noise. We can test how reasonable this assumption is by comparing the fluctuations of reflection and fluorescence signals across a large number of experiments. The LHS of Figure 4.12 shows the variance/mean as a function of time, calculated per integration bin for 300 reflection measurements, with $\tau_{int} = 2 \mu\text{s}$. These 300 consecutive shots were chosen from a selection of 500 taken on the same day as having the least change in cavity pumping rate across them. Generally, this data shows fluctuations about the line of variance/mean equal to one, although at $t \sim 70 \text{ ms}$, which corresponds to atoms being in the cavity, there does appear to be a slight dip. However, the white line shows the variance/mean calculated for groups of summed counts, 200 integration time-bins long (i.e. of $400 \mu\text{s}$ duration) and shows small fluctuations about the line of variance/mean equal to one, but no dip. In addition, a linear fit comparing the variance and mean gives a gradient of 1.002 ± 0.004 , indicating that our reflection signal is consistent with a Poissonian and shot-noise-limited signal.

The equivalent dataset for the fluorescence signal is shown on the RHS of Figure 4.12. Here, the variance versus the mean for each $2 \mu\text{s}$ time bin is calculated across 250 shots. A linear




Figure 4.12: Comparison of the signal variance and mean. LHS: The variance/mean as a function of time for the reflection signal, calculated for 300 shots. The white line shows the variance/mean calculated for groups of integration time-bins, 200 long. RHS: The variance versus the mean for the fluorescence signal, calculated for 250 shots. The linear fit with zero intercept shown alongside the data produces a gradient of 1.04 ± 0.01 .

fit to this data, with zero intercept, gives a gradient of 1.04 ± 0.01 , again indicating that the fluorescence signal is nearly photon shot-noise-limited.

One factor determining the size of N_{sig} is τ_{int} , which defines how long we measure counts at the APD before binning the signal. The size of τ_{int} relative to the duration of the signal, τ_{sig} , is also relevant; for a signal of finite duration, increasing τ_{int} will naturally lead to a SNR optimum as N_{bg} , which is generally independent of the signal, will continue to increase whilst N_{sig} will not. This also implies that the optimum SNR will occur when $\tau_{int} \sim \tau_{sig}$.

The following sections test the functionality of our system for measuring $N_{eff} = 1$. This discussion is then extended to consider the prospects for detecting individual atoms falling through the cavity and individual atoms trapped inside the cavity. The limit for confident detection of an atom will be taken as when the $SNR = 1$.

4.3.2 Detecting $N_{eff} = 1$

We already know that we are sensitive to single atoms. The LHS of Figure 4.13 shows the mean signal from the 300 drops discussed earlier, alongside a Gaussian fit to the signal indicated by the white line. The peak in the signal corresponds to a normalised reflection signal of 2.00 ± 0.01 lab units, which assuming $C = 0.3$, gives $N_{eff} = 1.1$. This signal thus corresponds, approximately, to the signal one would obtain from a single atom optimally coupled to the cavity. We now use this data to calculate the SNR for different values of τ_{int} .

The RHS of Figure 4.13 shows the SNR as a function of τ_{int} , where in each case the period of integration begins at the centre of the drop signal and expands symmetrically with increasing time. The blue data points are discrete calculations of the SNR made from the averaged data by summing over increasing numbers of time bins. The solid line shows a theoretical calculation of the expected behaviour of SNR with τ_{int} , assuming the fit parameters found for the average signal shown on the LHS of the same figure and Poissonian photon statistics. It matches the data points well; both show a clear rise in SNR with τ_{int} , peaking at SNR = 10 when $\tau_{int} = 6.65$ ms. The Gaussian fit from the LHS estimates a FWHM of 8.6 ms. This means that the optimum value of τ_{int} is approximately $1.5\times$ the FWHM of the signal, broadly in line with our earlier expectation.

Much faster detection can be achieved, however, when setting a minimum threshold of SNR = 1. From the figure, this occurs at $\tau_{int} = 37 \mu\text{s}$, meaning that we can confidently detect $N_{\text{eff}} = 1$ in $37 \mu\text{s}$, or, equivalently, at a rate of 27 kHz.

The dashed line in the RHS Figure 4.13 is also a theoretical calculation of the expected SNR, but now calculated for the peak signal, assuming no Gaussian time dependence (i.e. for a stationary atom). This line follows the data and solid line for small integration times, but once $t \sim 4$ ms, they diverge as a larger portion of the Gaussian profile is sampled.

We can make an alternative measure of the detection confidence by using a threshold count level, u , to distinguish between signal counts, N_{sig} , and background counts, N_{bg} and then comparing the relative probabilities for detecting both at once. The probability of detecting both $N_{sig} \geq u$ AND $N_{bg} < u$ can be calculated analytically by knowing that our background and signal are well approximated by Poissonian distributions.

The mean signal and background count levels are calculated by fitting a Poisson distribution to histograms of the relevant sections of data. For the signal counts, the appropriate section is the time segment of 8.6 ms corresponding to the FWHM region of the averaged Gaussian signal. The number of counts per $2 \mu\text{s}$ time bin are counted in this time segment for all 300 drops, giving a distribution of signal counts, shown in the lighter blue on the LHS of Figure 4.14. The data are plotted with statistical error bars, assuming a Poisson distribution of counts — although on the scale of the graph these aren't possible to resolve.⁷ Fitting a Poisson distribution to this dataset gives the signal count mean, $\bar{x}_{sig}^{fit} = 0.3008 \pm 0.0002$, which completely describes the distribution. This compares well to the actual mean, $\bar{x}_{sig}^{act} = 0.3011$ — providing further

⁷The random errors are presented in this plot because we are trying to characterise the distribution of detected counts. The primary contribution to any systematic error would be the APD dead time, which is essentially a problem of how accurately we can count photons leaving the cavity.




Figure 4.13: The average drop signal in reflection and calculated SNR. LHS: The average of 300 reflection experiments. The white line shows a Gaussian fit to the average, from which the peak normalised reflection signal is 2.00 ± 0.01 lab units, equivalent to an $N_{\text{eff}} = 1.1$, and the FWHM is 8.6 ms. RHS: The SNR as a function of integration time, calculated for the average drop signal shown on the LHS. The blue data points show calculations of the SNR from the data and using Eq.(4.11). The solid and dashed lines are theoretical calculations of the SNR as a function of τ_{int} . The solid line takes into account the Gaussian time dependence of the signal, whilst the dashed line does not. From this, the SNR rises quickly, equalling 1 after $37 \mu\text{s}$ and peaking at 10 after 6.65 ms. This latter roll-over occurs when τ_{int} equals $1.5 \times$ the FWHM of the Gaussian signal.

evidence that our signal has Poissonian fluctuations.

The background count mean, $\bar{x}_{bg}^{fit} = 0.1621 \pm 0.0003$, is calculated in a similar fashion. The number of counts per integration time bin are counted for the final 30 ms segment of each of the 300 drops (when no atoms are present), resulting in the count distribution shown in darker blue on the LHS of Figure 4.14. As with the signal counts, statistical error bars are included. The actual mean is $\bar{x}_{bg}^{act} = 0.1625$, again showing good agreement.

Both sets of means (signal and background) are calculated for one time bin, corresponding to $2\mu\text{s}$, although the equivalent for larger integration times can be made by multiplying by a factor $\tau_{int}/2$ as this doesn't alter the statistics. Using these means, the joint probability of counting $N_{sig} \geq u$ and $N_{bg} < u$ is given by:

$$P_{\geq u, sig} \times P_{< u, bg}$$

where

$$P_{\geq u, sig} = 1 - \sum_{j=0}^{j=u-1} e^{-\bar{x}_{sig}} \frac{\bar{x}_{sig}^j}{j!} \quad \text{and} \quad P_{< u, bg} = \sum_{j=0}^{j=u-1} e^{-\bar{x}_{bg}} \frac{\bar{x}_{bg}^j}{j!} \quad (4.12)$$

This joint probability is shown as a function of τ_{int} on the RHS of Figure 4.14. Three values of the threshold count level, u are shown; $u = 3$ is shown in the darkest blue; $u = 2$ is shown in the medium-toned blue; and $u = 1$ in the lightest blue. From this, as u increases, the peak detection fidelity decreases and occurs at longer values of τ_{int} . As will be seen later in this section, this is contrast to the fluorescence signal. The best achievable fidelity is 37 %, which occurs for $\tau_{int} = 7\mu\text{s}$, using $u = 1$. Practically, this implies that one can be most confident that an atom was detected when the background signal count is < 1 and the signal count level is ≥ 1 . That the fidelity is so low is, again, not surprising; the count distributions for the signal and the background are quite similar, which makes it difficult to discriminate between them.

Another detection parameter of interest is the dynamic range of our signals; i.e. the range of N_{eff} over which we can confidently distinguish between different atom numbers. From Chapter 2, we know that the reflection signal saturates with higher N_{eff} . One would therefore expect the detector to fail to discriminate between different values of N_{eff} when the change in signal caused by an increase/decrease of one equals the fluctuations in the signal. This is shown in Figure 4.15, where the solid lines show the differentiated signal (i.e. the differential of Eq.(4.2) with respect to N_{eff}) as a function of N_{eff} , for two different pumping rates; the lighter blue corresponds to a pumping rate of 810 kCounts/sec and the darker blue to 81 kCounts/sec. The Poissonian fluctuations about the mean are indicated by the dashed lines, as a function of N_{eff} . The value of N_{eff} at which the solid and dashed lines of each colour meet (as shown in the




Figure 4.14: Using Poisson statistics to calculate the detection fidelity. LHS: The distribution of counts per $2\mu\text{s}$ integration time bin for the signal counts (lighter blue) and the background counts (darker blue). Fitting these distributions to a Poisson distribution gives estimates of the respective means as $\bar{x}_{sig}^{fit} = 0.3008 \pm 0.0002$ and $\bar{x}_{bg}^{fit} = 0.1621 \pm 0.0003$. RHS: The joint probability for detecting a signal count level $\geq u$ and a background count level $< u$ as a function of integration time. Three values of u are shown; $u = 3$ is shown in the darkest blue; $u = 2$ is shown in the medium-toned blue; and $u = 1$ in the lightest blue. As u increases the peak fidelity decreases and occurs for larger τ_{int} .

figure with solid points in each colour) indicate the upper limit to the detection range. It is clear that using higher pumping rates enables one to discriminate between higher values of N_{eff} : the upper limit for the higher pumping rate occurs at $N_{\text{eff}} = 26.2$, whilst for the lower pumping rate it occurs at $N_{\text{eff}} = 14.1$. This makes sense, as whilst the atoms aren't saturated, we would expect the SNR to increase with increasing pumping rate.

Thus far we characterised the quality of our system using the reflection signal as an atom detector by calculating the achievable SNR, detection speed, detection fidelity and dynamic range. These findings are summarised in Table 4.1. We can use a similar rationale to calculate the equivalent parameters for the fluorescence signal, allowing us to directly compare the two methods. For all the following calculations, we use the set of 250 shot-noise-limited fluorescence drops discussed earlier.

As before, the first parameter to be considered is the SNR as a function of τ_{int} . This is shown in Figure 4.16. The dark red data points are discrete calculations of the SNR made from the averaged fluorescence signal, calculated as a function of τ_{int} by summing over increasing numbers of time bins. In each case, τ_{int} begins at the first time bin after the fluorescence pump beam is switched on.




Figure 4.15: The range of N_{eff} which can be discriminated by our atom detector. The solid lines show the differentiated signal with respect to N_{eff} as a function of N_{eff} , for two different pumping rates: 810 kCounts/sec (in lighter blue) and 81 kCounts/sec (in darker blue). The dashed lines show the signal fluctuations as a function of N_{eff} . Where the solid and dashed line of each colour meet, shown by the solid points, indicates the upper limit of the detection range. For the higher pumping rate this occurs at $N_{\text{eff}} = 26.2$ and for the lower pumping rate this occurs at $N_{\text{eff}} = 14.1$.




Figure 4.16: The SNR as a function of integration time, calculated for a fluorescence signal averaged over 250 drops. The period of τ_{int} always begins for the time bin after the pump beam is switched on. The dark red data points are discrete calculations of the SNR from the averaged data, calculated by summing over increasing numbers of time bins. The solid line is the SNR, calculated assuming an exponential signal, with parameters found by fitting to the average signal. It shows good agreement to the data points. The dashed line is the SNR calculated assuming a constant, peak signal — this diverges from the solid line and data points once $\tau_{\text{int}} \geq 20 \mu\text{s}$. SNRs > 1 are easily achievable, with the SNR equalling 1 after only $4 \mu\text{s}$.

The solid and dashed lines shown in Figure 4.16 are calculations of the expected SNR. The solid line shows the case of a signal which exponentially decays with increasing time — this was shown earlier to be a good approximation of the actual fluorescence signal and fits the data points well. The dashed line shows the case for a stationary N_{eff} ; i.e. the SNR which would be achieved were the mean peak fluorescence signal to be constant. This matches both the solid line and the data points well whilst $\tau_{\text{int}} < 20 \mu\text{s}$, but they diverge as more of the exponentially decreasing signal is sampled. This does however indicate that for the first few tens of microseconds, the signal can be well approximated as a constant signal — a point which will be used later when calculating the fidelity of atom detection using this method.

From this we know that, as with the reflection signal, high SNRs are easily achievable. The SNR peaks at 5.4 when $\tau_{\text{int}} = 560 \mu\text{s}$ and one can achieve a $\text{SNR} = 1$ after only $3.9 \mu\text{s}$. This is an order of magnitude faster than the equivalent case in reflection, enabling faster detection rates of up to 250 kHz. From the figure it is also clear that integrating for significantly longer times does not substantially reduce the SNR, despite the short duration of the signal. This is because the background noise is more than two orders of magnitude smaller than the peak amplitude of the signal. In fact, using $\tau_{\text{int}} = 5 \text{ ms}$ only reduces the SNR to 5.0.

We can also calculate the detection fidelity when using the fluorescence signal, following the same method of calculating the joint probability for detecting $N_{\text{sig}} \geq u$ AND $N_{\text{bg}} < u$. As before, the distribution of the signal counts and background counts are used to define Poisson distributions with means, $\bar{x}_{\text{sig}}^{\text{fit}}$ and $\bar{x}_{\text{bg}}^{\text{fit}}$. The signal counts used here are from the first $2 \mu\text{s}$ time bin across the 250 shots. The background counts are from the final 15 ms of each of the 250 drops. The LHS of Figure 4.17 shows the signal (lighter red) and background (darker red) count distributions. Unlike the case in reflection, these distributions are quite different, particularly for count levels ≥ 1 , suggesting that it will be possible to discriminate between them with a high degree of accuracy, which in turn should lead to a high detection fidelity. As with the reflection signal and background count distributions analysed previously, the statistical error bars are shown here. The much larger uncertainty in the signal counts arises from the substantially smaller sample tested (i.e. only 250 samples compared to nearly 2 million background samples). The calculated means of these two count distributions are $\bar{x}_{\text{sig}}^{\text{fit}} = 0.4410 \pm 0.0867$ and $\bar{x}_{\text{bg}}^{\text{fit}} = 0.0039 \pm 0.0001$, which can be compared to the actual means of $\bar{x}_{\text{sig}}^{\text{act}} = 0.5$ and $\bar{x}_{\text{bg}}^{\text{act}} = 0.0040$. Again, that these two estimates of the means are well matched is further confirmation that it is reasonable to use Poisson distributions to approximate our signal and background.




Figure 4.17: Calculating the detection fidelity using the fluorescence signal. LHS: The distribution of signal counts, shown in the lighter red, and background counts, shown in the darker red. Unlike reflection, the count distributions are quite different. RHS: The joint probability to detect $N_{sig} \geq u$ AND $N_{bg} < u$ as a function of τ_{int} . Three threshold values are shown; $u = 1$ is shown in the darkest red; $u = 2$ is shown in the medium-toned red; and $u = 3$ in the lightest red. As u increases so too does the fidelity of the measurement, although a corollary of this is having to wait longer. A fidelity of 95 % can be achieved after $10.3 \mu s$ if we use $u = 1$ counts as the counting threshold. This can be increased to 99 % after $15.6 \mu s$ using $u = 2$ and even further increased to 99.7 % after $22.6 \mu s$ using $u = 3$.

As before, we can extend to larger values of τ_{int} by multiplying by $\tau_{int}/2$. In the case of N_{bg} , this approximation is valid for τ_{int} of several milliseconds, as the background counts fluctuate about this constant level over this time period. For N_{sig} , this approximation is only valid for $\tau_{int} \sim 20 \mu s$, where the signal is approximately equal to its peak level.

The joint probabilities, found using Eq.(4.12), are plotted as a function of τ_{int} on the RHS of Figure 4.17. Three values of u are shown; $u = 1$ is shown in the darkest red; $u = 2$ is shown in the medium-toned red; and $u = 3$ in the lightest red. From this we see that choosing a larger discriminator value for u means having to wait longer to maximise the detection fidelity, but the resulting fidelity is improved. If we use the lowest discriminator of $u = 1$, we can detect an atom with 95 % fidelity after $10.3 \mu s$. The fidelity can be increased by waiting longer and using $u = 2$ count as a discriminator, whereby a 99 % fidelity is achieved after $15.6 \mu s$. Evenmore, setting a discriminator of $u = 3$ counts enables a fidelity of 99.7 % to be achieved after $22.6 \mu s$, although this is at the upper limit of our constant signal level approximation. This demonstrates that we can achieve detection fidelities $> 99\%$ whilst also reaching detection rates of up to 50 kHz.

Whilst the fluorescence signal compares favourably with the reflection signal with regards to the detection speed and fidelity, the dynamic range available is smaller. From Chapter 2 we expect the average fluorescence signal from atoms scattering in the disordered limit to be $\propto C\sqrt{N_{eff}}/(1 + 2CN_{eff})$, whilst the atoms aren't saturated, although for high saturation, the signal $\propto 2CN\gamma$. This function peaks when $N_{eff} = 1/2C$, rolling over and tending to zero with increasing N_{eff} . This feature implies a fundamental difficulty for atom counting, as one could never be absolutely certain which side of the maximum one is on. Despite this, were we to assume low atom numbers (i.e. to the left of the signal maximum), the peak value of N_{eff} which could be measured, ignoring any uncertainty caused by fluctuations, would be $1/2C$. In our system this equals 1.6.

The following chapter presents simulations exploring the effect of detuning the fluorescence pump beam for higher N_{eff} . These indicate that using a detuned pump beam can shift the roll-over to larger values of N_{eff} , increasing the available dynamic range. This will, however, have a concomitant effect of reducing the atomic scattering rate and resulting signal.

Table 4.1. presents the key results from this section, comparing the key detection parameters for the reflection and fluorescence signal. These findings imply that, whilst the reflection signal has a larger dynamic range and can ultimately achieve higher SNRs, the fluorescence signal is more suited to fast detection and has a higher fidelity. Whilst this appears to make

Detection parameter:	Reflection:	Fluorescence:
Peak SNR	10	5.4
Required τ_{int}	6.65 ms	400 μ s
τ_{int} required for SNR= 1	37 μ s	3.9 μ s
Peak fidelity	37 %	99 %
Required τ_{int}	7 μ s	15.6 μ s
Dynamic range	$0 \leq N_{\text{eff}} \leq 15 - 25$	$0 \leq N_{\text{eff}} \leq 1.6$

Table 4.1: Summary of the key detection parameters for when $N_{\text{eff}} = 1$ for the reflection and fluorescence signals.

the fluorescence signal more suitable for single atom detection, another point to consider is how destructive each measurement is. This is the main negative aspect of the fluorescence signal, which can be quite destructive; in the process of measuring the atoms can be pushed out the cavity, whilst in reflection the only limiting factor is the atoms' transit-time.

From this, however, we know that we can confidently detect $N_{\text{eff}} = 1$. The following section considers the case of detecting individual atoms, either as they fall through the mode or when they are trapped.

4.3.3 Detecting individual atoms falling through the cavity mode

It is interesting to consider how efficient our system is at detecting an individual atom as it falls through the cavity mode. There are two factors to consider here; how large a signal do we expect from one atom and how long should this signal be integrated for to maximise the signal to noise. The question of optimal τ_{int} is addressed first.

We expect the optimal τ_{int} to equal the signal duration, although in these circumstances this isn't well-defined. After falling from the MOT, the atom will be travelling at ~ 32 cm/sec, taking approximately 28 μ s to cross from $-w_0 \rightarrow w_0$ of the cavity mode. If we assume that the atom follows a straight line trajectory, this will equal the HWHM of the atom's signal and, from previous analysis, we might expect the best choice of τ_{int} to be similar.

A more rigorous method for determining the best τ_{int} to use can be made by simulating the reflection counts measured at the APD as an atom falls through the cavity and re-binning the data with different τ_{int} to find the maximum SNR. Initially, this is done for an atom which passes through the centre of the cavity mode, so that it has a peak coupling rate of $g = 98$ MHz. First, a 'no-atom' trace is simulated by producing a fake APD trace which has a Poissonian

distribution of counts spanning $250 \mu\text{s}$ in 5 ns time bins, with a mean of $0.0025 \text{ counts/bin}$, equivalent to 500 kCounts/sec . This is the input count trace. The effect of the atom is taken into account for each time bin using Eq.(4.3), with $C = 0.3$ and an atom number, $N_A(t)$ which is a Gaussian centred at $t = 95 \mu\text{s}$, with $\text{HWHM} = 28 \mu\text{s}$ and a peak = 1, mimicking the atom falling through the mode. This produces an output count trace. For an individual time bin, the input counts, N_{in} are related to the output counts, N_{out} by:

$$N_{out} = \frac{N_{input}}{v^2} \left[\frac{2CN_A(t) + v}{2CN_A(t) + 1} \right]^2 \quad \text{where} \quad N_A(t) = e^{-\frac{\text{velocity}^2 t^2}{w_0^2}} \quad (4.13)$$

An example simulated trace, with and without an atom and before re-binning, is shown on the LHS of Figure 4.18. The ‘with atom’ case, shown on top, can now be used to determine the optimum τ_{int} by performing box-car sums across the trace for different box car lengths between $0.5 - 50 \mu\text{s}$. Examples of the traces which result from this are shown on the RHS of Figure 4.18, for $\tau_{int} = 500 \text{ ns}$ (on the top) and $\tau_{int} = 10 \mu\text{s}$ (on the bottom).

The SNR is found using Eq.(4.11), where N_{sig} is the mean peak count level found between $t = 75 \rightarrow 125 \mu\text{s}$ for each box-car trace and N_{bg} is the mean count level, taken across the whole trace. This procedure is repeated 100 times with fresh input traces and the final SNRs at each τ_{int} are found by averaging. The results are shown in Figure 4.19, where SNR is shown as a function of τ_{int} . An optimum SNR of ≈ 2.1 occurs when $\tau_{int} = 20 \mu\text{s}$. This value of τ_{int} will be used henceforth.

It should be noted that by virtue of only using the mean pumping rate across the ‘without atoms’ trace to calculate N_{bg} , any other contributions to the background noise, such as the far-detuned light used for locking the cavity, have been ignored. This is reasonable, as typically these contribute $< 10 \text{ kCounts/sec}$, compared to the pumping rate of 500 kCounts/sec .

This analysis already indicates that we’d expect to detect a maximally coupled atom with an $\text{SNR} \geq 1$. In reality, however, most atoms will not fall through the centre of the mode and will follow a random trajectory. One can estimate the distribution of coupling rates by randomly assigning an atom coordinates inside the cavity multiple times and looking at the average effect. To do this, an atom is randomly positioned on a plane centred in the cavity one million times. For each ‘throw’ of an atom, the peak coupling, which occurs when the atom is vertically centred on the mode, is calculated, using the atom’s x and z coordinates. The longitudinal coordinate, z , can take any value between $0 \leq z \leq d$, where d is the cavity length, whilst the horizontal radial coordinate, x , is taken to have a range $-2w_0 \leq x \leq 2w_0$, where w_0 is the cavity waist. The distribution of coupling rates is shown on the LHS of Figure




Figure 4.18: Box-car analysis of simulations of the signal from a single atom passing through the cavity. LHS: An example input (bottom) and output (top) trace used to determine the optimal τ_{int} . The input trace is $250 \mu\text{s}$ long, in $50,000 \times 5 \text{ ns}$ time bins. The counts across the trace have a Poissonian distribution, with a mean of 0.0025 counts/bin. The output trace is calculated, per bin, using Eq.(4.13). RHS: The same output traces, but now re-binned with $\tau_{int} = 500 \text{ ns}$ (top) and $\tau_{int} = 10 \mu\text{s}$ (bottom).




Figure 4.19: The SNR of the reflection signal as a function of τ_{int} , produced using simulations of atom traces as detailed in the main text. A SNR optimum of ≈ 2.1 occurs at $\tau_{int} = 20 \mu\text{s}$, indicating that this is the best value of τ_{int} to use.




Figure 4.20: The probability of an atom having a given peak coupling rate, calculated by randomly assigning coordinates inside the cavity mode to one atom, one million times. The bounds for each coordinate are set by the cavity length and cavity waist. LHS: The probability of occurrence of a given peak coupling rate, g . The coupling rates are grouped into 1 MHz wide bins spanning a range $-98 \rightarrow 98$ MHz. RHS: The probability of occurrence of a given peak C . The peak cooperativities are grouped into 0.005 wide bins spanning a range $0 \rightarrow 0.31$.

4.20, where the one million coupling rates have been grouped into bins of width 1 MHz, from $-98 \rightarrow 98$ MHz. The y -axis shows the probability of finding an atom of a given peak coupling rate. Integrating the distribution indicates that 28 % of atoms will have a peak coupling rate, $|g| < 10$ MHz. The mean peak coupling rate is $\bar{g} = 26.3$ MHz and the median is 20.9 MHz.

The RHS of Figure 4.20 shows the same set of atoms, but instead the probability of an atom having a peak C of a given value is shown, now grouped into bins of width 0.005, spanning a range from $0 \rightarrow 0.31$. These values are calculated assuming $(\kappa, \gamma) = 2\pi \times (5200, 3)$ MHz. From this one finds that the 89.1 % of atoms have a cooperativity, $C < 0.1$, whilst the average cooperativity is $\bar{C} = 0.037$. Only 0.03 % of atoms have $C = 0.3$, as assumed in the earlier analysis.

Whilst this calculation of the cooperativity distribution can be helpful, it is relatively arbitrary as the cavity mode volume is not well defined. Thus, the final distribution ultimately depends on how large a volume one ‘throws’ atoms into. Instead, one can calculate the number of single atoms that would be detected by treating an atom density, multiplied by the volume of space in which the cooperativity is large enough to produce a detectable signal.

We can use the earlier analysis to calculate the threshold cooperativity required to achieve an $\text{SNR} \geq 1$. Now, τ_{int} is fixed at $\tau_{int} = 20 \mu\text{s}$ and the value of C used in Eq.(4.13) is varied. The results of this are shown in Figure 4.21. This calculation indicates that a threshold peak




Figure 4.21: The SNR of the reflection signal calculated as a function of peak atom cooperativity, assuming $\tau_{int} = 20 \mu\text{s}$. This indicates that a threshold of $C \geq 0.13$ produces an SNR ≥ 1 .

value of $C \geq 0.13$ results in a SNR ≥ 1 . From our distribution, 6.8% of atoms, or equivalently, 6.8% of the cavity mode area fulfil this (we refer to the ‘mode area’, which is from the plane of $y = 0$, as the above distribution calculates the *peak* C which assumes $y = 0$). The calculations of SNR above assume an atom travelling across the mode from $-w_0 \rightarrow w_0$, so the relevant area becomes a volume by multiplying by $2w_0$. Thus, the actual volume with sufficient C for an atom to be detected is $1.53 \times 10^{-15} \text{m}^3$.

This volume can be multiplied by a atom density to calculate the number of single atom events we would detect. From Eq.(4.1) we have an expression for the density at the cavity as a function of time. At the peak in the atom drop signal, we have a density of $5.83 \times 10^{14} \text{atoms/m}^3$. This implies that at the peak in the drop signal, we’d expect to be able to detect 0.89 atoms, i.e. less than one atom. This is corroborated by analysing experimental traces of atom drops.

An initial response to this might be to increase the atom density, although this isn’t practical for single atom detection. At this density, the number of atoms in the cavity mode volume (i.e. $\pi w_0^2 d/4$) is 1.3, so increasing the density further would mean that there was more than one atom in the cavity at once.

We can try to improve the situation by increasing the input pumping rate, although atomic saturation means that this cannot be done indefinitely. How much the input pumping rate can be increased before saturation starts to reduce the signal is shown in Figure 4.22 for three different cooperativities; $C = 0.1$ in dark blue, $C = 0.2$ in the medium-toned blue and $C = 0.3$ in light blue. The onset of saturation occurs $\sim 50 \text{MCounts/sec}$ for all three cases, although




Figure 4.22: The normalised reflection signal calculated as a function of input counting rate, for three values of the cooperativity; $C = 0.1$ in dark blue, $C = 0.2$ in the medium-toned blue and $C = 0.3$ in light blue. For all three cases it appears as though saturation begins to take place at ~ 50 MCoun/sec, but the larger the cooperativity the larger the rate of decrease in the signal.

as C increases it is clear that the rate at which it impacts the signal increases. From this, one might assume we could increase our rate to, say 20 MCoun/sec and increase our SNR. In reality, however, we must also consider the health of our detector. Our APD has an absolute damage threshold of 10 MCoun/sec, although we would want to keep below a count rate of 5 MCoun/sec for safety⁸.

In this situation, using an input counting rate of 2.6 MCoun/sec (chosen so that an atom coupled with the maximum possible cooperativity of $C = 0.3$ produces a signal count rate of 5 MCoun/sec), does not reduce the threshold C required to detect an atom with $\text{SNR} \geq 1$.

We can make similar estimates for the fluorescence signal and begin by defining the contributions to Eq.(4.11). N_{sig} will be the number of fluorescence photons emitted by the atom into the cavity mode and collected by the APD. This is the product of the atom's mean duration in the cavity, its scattering rate and the total collection efficiency, χ . In §3.1. χ was found to be 0.06 and this is used here.

As soon as the pump beam is switched on, the atom will feel a force pushing it out of the cavity. Our chances of detecting the atom are better if it spends longer in the cavity, so the beam should come on when the atom is one mode-waist, w_0 , past the cavity centre.

⁸In principle, we could replace our APD with a detector with a much higher damage threshold. Whilst this might help the cause of single atom detection using the reflection signal, it isn't a practical option for us as we would lose the functionality of the APD which we need for other experiments; namely, the ability to count single photons.

The relevant time is then the time taken for the atom to cross $2w_0$. In this time, the mean cooperativity is the average over the mode profile, from $-w_0 \rightarrow w_0$, equalling $0.6 \times$ the peak cooperativity. This mean value will be used to calculate the scattering rate.

The total atomic scattering rate, when saturated, is $(2C + 1)\gamma(P)$. $\gamma(P)$ is a function of power and is adjusted using Eq.(4.7), assuming that the atom is half-saturated for a beam power of $20 \mu\text{W}$. This value is lower than that found in §2.1., but the beam was subsequently re-aligned to optimise the fluorescence signal. Each scatter will lead to a momentum kick, $\hbar k$ and, assuming the atom scatters at a constant rate, this leads to an acceleration, $\hbar k/m$, so that we can calculate how long the atom is in the cavity. With a pump beam power of $100 \mu\text{W}$ and assuming peak cooperativity of 0.3, it takes the atom $6.3 \mu\text{s}$ to travel out of the cavity. In this case, $N_{sig} = 2C\gamma(P) \times 6.3 \mu\text{s} \times 0.06$.⁹

The contributions to N_{bg} should also be defined. These will have several sources; far-detuned light used for frequency locking the cavity; stray light from the repumping beams, equipment screens and laboratory room lights; and scattered light from the pump beam itself. We can estimate the size of this background scatter using the average fluorescence trace from §3.1. Excluding scattering from the pump beam, the total background count rate is 920 counts/sec. With the pump beam switched on, the ratio of peak signal to background signal was 135:1, at a power of $20 \mu\text{W}$. Assuming the pump beam contribution to N_{bg} scales linearly with power, this ratio can be used to find the expected background noise for a given beam power.

Given these assumptions, the SNR is shown as a function of τ_{int} on the LHS of Figure 4.23, calculated for a beam power of $170 \mu\text{W}$. Results for two cooperativities are shown; $C = 0.3$ in light red and $C = 0.1$ in dark red. As expected, for both, the best SNRs occur when τ_{int} equals the atom transit time, indicated by the dashed lines in each colour. The higher cooperativity case yields a peak $\text{SNR} = 1$, indicating that only the most largely coupled atoms would be detected. The sharp point on this curve is a manifestation of defining an exact boundary at which the signal ends (i.e. within w_0) — in reality this point would be ‘rounded off’ because the point at which the atom leaves the cavity isn’t well defined.

We would expect that calculating SNR as a function of pump beam power will show a maximum, caused when the added background noise from increasing the beam power negates any increase in signal. This is evident on the RHS of Figure 4.23, which shows the peak SNR (i.e. assuming τ_{int} equals the transit time), calculated as a function of pump beam power. Again, the lighter (darker) red shows an atom with a peak cooperativity of 0.3 (0.1). The SNR

⁹This assumes the same polarisation conditions as found in §2.1.




Figure 4.23: Detecting a single falling atom with the fluorescence signal (1). LHS: The signal to noise ratio as a function of integration time, calculated for a beam power of $170 \mu\text{W}$. Two peak cooperativities are shown; $C = 0.3$ in light red and $C = 0.1$ in dark red. The dashed line indicates the average time the atom spends in the cavity before being pushed out by the fluorescence beam. This is where the SNR peaks, at $\text{SNR} = 1.01$ and $\text{SNR} = 0.57$, respectively. RHS: The peak signal to noise ratio as a function beam power. The same two cases as for the LHS are shown here. For $C = 0.3$ shown in light red, a peak of $\text{SNR} = 1.01$ occurs for a beam power of $180 \mu\text{W}$ - although this value varies by only 1% between $80 - 350 \mu\text{W}$. For $C = 0.1$ shown in dark red, the peak is $\text{SNR} = 0.57$ and has a similarly large range over which it barely changes.

peaks at a higher power of $180 \mu\text{W}$ (compared to $100 \mu\text{W}$) for the higher cooperativity case, although in both cases, there is large range of powers over which the SNR is within 1% of its the peak value.

Given the small variation of the SNR for pump beam powers within $\sim 20\%$ of the optimum power, we can directly compare the SNR achieved at different cooperativities. This is shown in Figure 4.24, where the peak SNR using a pump beam power of $200 \mu\text{W}$ is calculated as a function of cooperativity. From this, the threshold cooperativity required to have a $\text{SNR} \geq 1$ is $C \geq 0.29$, indicating that, as with the case in reflection, at our peak density we would not expect to detect a single atom.

To summarise, it is unlikely that either the reflection or fluorescence signal would be able to detect a single atom falling through the cavity. The main limiting factor is that the atom isn't localised in the cavity mode, so that its coupling is generally quite weak. This indicates that perhaps our best option is to trap the atom. This is a future aim for the experiment and, for this reason, it is interesting to extend the earlier analysis to treat this case.




Figure 4.24: Peak fluorescence detection SNR as a function of cooperativity, calculated assuming a pump beam power of $200 \mu\text{W}$. This indicates that atoms with $C \geq 0.29$ will produce a signal with $\text{SNR} \geq 1$.

4.3.4 Detecting individual atoms trapped inside the cavity

A future aim of the experiment is to trap atoms inside the cavity with an optical dipole trap. The findings for $N_{\text{eff}} = 1$ when assuming a stationary atom already indicate that we'd confidently detect a trapped atom. Here, the SNR should be improved as the atom spends longer in the cavity and it will be localised at antinodes in the potential. In practice, $\sim 50\%$ of the trapping light antinodes will overlap with the cavity mode antinodes, but a controllable loading mechanism could be optimised to ensure atoms are loaded into positions which maximise their coupling to the mode. We'll assume that this is the case for this calculation.

As before, we begin by defining contributions to Eq.(4.11). Treating first the case of atoms being detected in reflection, $N_{\text{sig}} = J_{\text{sig}}\tau_{\text{int}}$, where J_{sig} is given by Eq.(4.2) and is proportional to C . Whilst the trap substantially improves the atom's coupling, the atom will not be static. Typically, it will undergo oscillations across the cavity mode at a frequency $\omega_{\text{trap}} = k(\frac{2U}{m})^{\frac{1}{2}}$ where U is the depth of the trapping potential. This will result in a time variation of C , although the oscillations will only sample the lower portion of the cavity antinode and so an average coupling of $0.9g_{\text{max}}$ is assumed, or $0.8C_{\text{max}}$, where $C_{\text{max}} = 0.3$.

A negative aspect of dipole trapping is that it involves far off-resonant light circulating inside the cavity. Despite filtering by polarisation and dichroic filters, some proportion of this light will leak through to the APD, contributing to the background noise. Given this, N_{bg} will have two main contributions:

$$N_{\text{bg}} = \frac{J_{\text{in}}\tau_{\text{int}}}{v^2} + N_{\text{dipole}}$$

The first term is from the resonant light pumping the cavity. The second term represents




Figure 4.25: Detecting a single trapped atom in reflection (1). LHS: The mean lifetime, τ_{mean} for an atom inside a dipole trap in the cavity, calculated as a function of J_{in} . The primary cause of atom loss out of the dipole trap is assumed to be from scattering cavity pumping light. It is evident that τ_{mean} is inversely proportional to J_{in} ; for $J_{in} = 1$ MCOUNTS/sec, $\tau_{mean} = 6.7$ ms, whereas for $J_{in} = 0.5$ MCOUNTS/sec, $\tau_{mean} = 13.3$ ms. RHS: The reflection measurement SNR as a function of τ_{int} for $J_{in} = 0.25$ MCOUNTS/sec, so that $\tau_{mean} = 26.6$ ms. As expected, the peak of SNR = 64 occurs when $\tau_{int} = \tau_{mean}$.

the scattered light from the dipole trap, taken to be 0.2 MCOUNTS/sec when using an input trapping power to the cavity of $500 \mu\text{W}$, which is typical for our experiment. At this input power and including the effect of cavity build up, light of wavelength 804nm will produce a trap depth of $920k_B \mu\text{K}$.

The timescale for detecting a trapped atom will be set by its mean lifetime inside the trap, τ_{mean} . We can estimate this by assuming that the primary reason for atom loss is due to heating as a result of scattering light from the cavity pump beam. We assume that an atom leaves the trap having scattering sufficient photons that its additional energy is equal to half the trap depth, with each scatter transferring 362 nK of energy. The mean lifetime is shown as a function of J_{in} on the LHS of Figure 4.25. It is inversely proportional to J_{in} , so that a typical pumping rate of $J_{in} = 1$ MCOUNTS/sec results in $\tau_{mean} = 6.7$ ms, whilst using $J_{in} = 0.5$ MCOUNTS/sec gives $\tau_{mean} = 13.3$ ms.

Given these assumptions one can calculate the SNR for reflection detection of a single atom trapped inside the cavity. The RHS of Figure 4.25 shows the SNR as a function τ_{int} , calculated for $J_{in} = 250$ kCOUNTS/sec. High SNRs are achievable, with a peak value of 64 occurring when τ_{int} equals the mean trapping lifetime, in this case 26.6 ms. This is a step change improvement on the case for detecting falling atoms in reflection. It should be noted that here, the sharp




Figure 4.26: Detecting a single trapped atom in reflection (2). LHS: The highest achievable SNR as a function of J_{in} . τ_{mean} is calculated for each value of J_{in} and it is assumed that the highest achievable SNR occurs when $\tau_{int} = \tau_{mean}$. Higher J_{in} results in higher values of the SNR, until $J_{in} \sim 1$ MCOUNTS/sec, when the SNR begins to saturate at ~ 75 . This is because the constant dipole trap light contribution to N_{bg} becomes negligible in comparison to that from the cavity pumping light and, once in this regime, any increase in N_{sig} is counteracted by a corresponding increase in N_{bg} . RHS: The minimum value of τ_{int} required to achieve an SNR ≥ 1 , calculated as a function of J_{in} . As J_{in} increases, this minimum time reduces, enabling faster detection. A corollary of this is that the atom spends less time in the trap.

point in the curve is the result of assuming that the atom has a fixed position (and thus a fixed cooperativity) for the duration of its time in the trap. In reality, as the atom becomes increasingly energetic it will sample a larger portion of the cavity mode and this would ‘round off’ the sharp tip of the curve.

The LHS of Figure 4.26 shows the highest achievable SNR for different values of J_{in} , assuming in each case that the peak occurs when $\tau_{int} = \tau_{mean}$, as confirmed by the previous figure. The SNR rises quickly with J_{in} , but begins to saturate when $J_{in} \sim 1$ MCounts/sec. This saturation occurs as, increasingly, the contribution to the noise from the dipole trap becomes less important, so that increasing J_{in} increases the signal, but it also equivalently increases the background noise through the J_{in}/v^2 term. Here, the maximum value of $SNR \sim 75$, with this value being entirely determined by the cooperativity.

In addition to the effect of saturation, we are again limited by the damage threshold of our APD. The upper limit of safe operation corresponds to $N_{sig} + N_{bg} = 5$ MCounts/sec, which occurs when $J_{in} = 1.7$ MCounts/sec, making our best possible $SNR = 73.5$. At these count rates, one would also expect to see the affects of the finite dead time of the detector.

From this we know we can expect to achieve high SNRs, but it is also interesting to know the speed with which we could detect a single trapped atom. As well as increasing the SNR, increasing J_{in} reduces the τ_{int} required to confidently detect an atom with an $SNR \geq 1$. This is evident on the RHS of Figure 4.26, which shows the minimum value of τ_{int} needed to achieve an $SNR \geq 1$ as a function of J_{in} . From this, detection rates > 1 MHz can be achieved using a $J_{in} = 1.3$ MCounts/sec, whilst still having a trapping lifetime of several milliseconds. Alternatively, one might aim to trap an atom for ~ 30 ms, setting an upper limit of $J_{in} \sim 0.2$ MCounts/sec. In this case, one could still achieve a high SNR (> 60) at a rate of ~ 250 kHz.

Turning now to the case of fluorescence, N_{sig} is calculated following the same method as for a falling atom, although now the increased localisation means that the atom is assumed to have a cooperativity of $0.8C_{max}$, for the duration of its time in the trap. N_{bg} is also similar, but with the additional contribution from the trapping light.

The timescale for detection is again limited by the time it takes an atom to be heated out of the trap. This is determined by the scattering rate and, given the higher pump rates used in fluorescence, the result is substantially shortened trap lifetimes compared to those in reflection, as shown on the LHS of Figure 4.27. A pump beam power of $100 \mu\text{W}$ will lead to a mean trap lifetime of $56 \mu\text{s}$, with lower power beams leading to longer lifetimes; a $50 \mu\text{W}$ pump




Figure 4.27: Detecting a single trapped atom using the fluorescence signal (1). LHS: The mean lifetime for an atom inside the dipole trap when being detected using the fluorescence signal. The atom's scattering rate increases for higher pump beam powers, so that the time taken to scatter sufficient photons to be lost from the trap reduces. For a beam power of $100 \mu\text{W}$, the mean lifetime is $56 \mu\text{s}$. The curvature is caused by saturation of the atomic scattering rate. RHS: The SNR as a function of τ_{int} , calculated for a beam power of $100 \mu\text{W}$. The dashed line indicates the average lifetime for an atom in the trap of $56 \mu\text{s}$. This is the point at which the SNR peaks, at a value of 2.2.




Figure 4.28: Detecting a trapped atom using the fluorescence signal (2). LHS: The SNR as a function of beam power, calculated assuming $\tau_{int} = \tau_{mean}$. As for the case in reflection, the SNR increases with increasing pump beam power, but when the power $\geq 100 \mu\text{W}$ the SNR saturates. The saturated SNR is ~ 2.3 . RHS: The minimum τ_{int} required to achieve an SNR ≥ 1 , calculated as a function of pump beam power. The minimum required τ_{int} decreases with increasing beam power. However, the rate at which it reduces saturates quickly; only $7 \mu\text{W}$ of power are needed for a minimum τ_{int} of $10 \mu\text{s}$, but reducing this further to $1 \mu\text{s}$ requires beam powers of several watts.

Detection parameter:	Reflection:	Fluorescence:
Peak possible SNR	73.5	2.3
Required τ_{int}	$3.9 \mu s$	$48 \mu s$

Table 4.2: Summary of the results from analysing the prospect for detecting a single, trapped atom in our cavity using the reflection and fluorescence signals.

beam increases the mean trap lifetime to $66 \mu s$. In the absence of saturation, the lifetime will fall inversely with increasing pump powers, but the effect of saturation is to produce the bend in the curve visible in the figure for powers $> 20 \mu W$.

The RHS of Figure 4.27 shows the SNR as a function of τ_{int} , calculated for a beam power of $100 \mu W$. The SNR peaks at a value of 2.2 when $\tau_{int} = \tau_{mean}$, as expected. This is substantially smaller than the SNRs achieved using reflection detection, the primary cause being the short amount of time that the atom spends inside the cavity, compared to mean durations on the order of 10 ms in the case of reflection detection.

Increasing the beam power only produces small improvements in the SNR. The peak SNR is shown as a function of beam power on the LHS of Figure 4.28. Again, the SNR saturates with increasing beam powers, with no substantial improvement achieved beyond a beam power of $\sim 100 \mu W$. One point to note, however, is that despite the atom being pumped at a high rate (when compared to reflection), the total signal + background count rate at the APD is still much smaller, as the cavity mode is not being directly pumped. Even with a beam power of 1 mW, the total rate is < 1 MCounts/sec, so we would be in no danger of damaging the APD.

We can also calculate the minimum τ_{int} required to achieve a $SNR \geq 1$. This is shown on the RHS of Figure 4.28 as a function of pump beam power, where, like reflection, higher pump beam powers enable atom detection at faster rates. $\tau_{int} = 10 \mu s$ is needed when using a beam power of $7 \mu W$, implying a possible detection rate of 100 kHz. These possible rates saturate quickly, however, so that increasing by an order of magnitude requires a beam power of several watts, which isn't experimentally practical. A summary of the findings of this section are given in Table 4.2.

This discussion has highlighted some of the differences between the fluorescence and reflection signal for use in atom detection. It appears that whilst in our current mode of operation, fluorescence signals offer the best route to detecting a single atom. Were an atom to be trapped in the cavity however, reflection detection would be preferable as providing a potentially non-destructive detection method. These calculations indicate that it would be possible to detect

an atom with an $\text{SNR} = 1$ after $250\,\mu\text{s}$, minimising the heating of the atom and the potential for loss out of the trap. Other qualitative differences between reflection and fluorescence signals are explored further in the following chapter, where a regime of higher N_{eff} is accessed.

Chapter 5

Increasing the atom number

For the experiments discussed up to now it has generally been the case that $N_{\text{eff}} \sim 1$ and the cavity is resonant with the atoms. This chapter begins by presenting a method for increasing N_{eff} , thereby enabling us to examine the impact of higher collective cooperativities on the reflection and fluorescence signals. This is followed by experiments which explore the dispersive nature of the atom-cavity interaction by measuring the excitation spectrum of the system in reflection and fluorescence. In doing so, we demonstrate that, for weak coupling, these two signals are manifestly different.

5.1 Motivation for increasing N_{eff}

The atom-cavity cooperativity is a direct measure of how strongly the system is coupled and, ideally, we'd like to increase the cooperativity as much as possible. Whilst remaining with ^{87}Rb , there are two options for doing this; increasing g or decreasing κ . This latter option isn't practical without a structural change to the cavity, as one would need to reduce the internal mirror roughnesses, thereby increasing the cavity finesse. There are also benefits to working with a large κ ; one can efficiently couple photons in and out of the system and the rate at which this can be done sets an upper speed limit to the operation of any cavity QED-based quantum logic gate.

Instead one can try to increase g . Reducing the mode volume would achieve this, but again this would require a structural change. Instead, large increases can be gained by increasing N_{eff} . From Chapter 2 we know that for N atoms coupled identically to the cavity at a rate g , the atom-cavity coupling rate becomes $\sqrt{N}g$ [40, 96]. This effect has been exploited experimentally with cold atoms [39, 38] and ions [112] and leads naturally to the concept of a *collective* atom-

cavity cooperativity, C_N , given by:

$$C_N = \frac{Ng^2}{2\kappa\gamma} = NC \quad \text{or} \quad C_N = N_{\text{eff}}C = \frac{\sum_{i=1}^N g_i^2}{2\kappa\gamma} \quad (5.1)$$

This latter expression takes into account the variation in g 's across the atoms using the spatially-dependent local coupling strength, g_i . We can thus expect to increase C_N by increasing N_{eff} . Our method for increasing N_{eff} is discussed in the following section and with it we achieve an increase of up to $N_{\text{eff}} \approx 11$ atoms, corresponding to increasing C_N by ~ 3.5 . This enables us to see substantially enhanced Purcell broadening of the atomic lineshape and dispersive effects.

5.2 Increasing N_{eff}

Our method for increasing N_{eff} uses a resonant beam to push atoms out of the MOT. This reduces the atom cloud's travel-time between the MOT and the cavity, thereby reducing the absolute expansion of the cloud and increasing the atom number density.

During this PhD project several alternatives were tried, all of which aimed either to reduce the clouds' expansion or increase the clouds' initial density. To achieve the latter effect we tried increasing the magnetic field gradient (thereby increasing the atom confinement) and switching off the repumping light (to encourage atoms to stop radiating and therefore reduce outward radiation pressure). Neither of these methods showed any notable increase in C_N , indicating that we were limited by the clouds' expansion. Attempts to increase the magnetic field were particularly thwarted by the need to ensure the coils didn't outgas into the vacuum chamber, so that any increase had to be transitory.

Alternatively, we tried including a far-detuned beam along the vertical axis to provide a confining dipole potential to reduce expansion of the atom cloud as it fell. Again, there appeared to be no increase in density, which was attributed to insufficient beam power. Finally, however, the push beam was settled on. This section describes the experimental procedure for this method and examines the resulting signals.

5.2.1 Experimental method

The dominant factor limiting the atom number density at the cavity is the ballistic expansion of the cloud in the $\sim 30\text{ms}$ it takes for them to travel between the MOT and the cavity. During this time, the radius of the cloud nearly doubles, so that the cloud density reduces by a factor

of 8. Compared to the push beam case, where typical travel times are ≤ 0.5 ms, the cloud radius changes by $< 1\%$, resulting in a negligible reduction in density. This implies that we should be able to increase N_{eff} by at least a factor of 8, with an equivalent increase of $\sqrt{8}$ in the coupling.

The ‘push’ beam reduces the cloud expansion by accelerating the atoms towards the cavity. With the beam switched on, atoms feel a net force caused by scattering push-beam photons, pushing them towards the cavity. The frequency and intensity of the push beam are controlled by an AOM, whose drive signal can be switched on and off using an RF switch. The optimal timing for the push beam was found empirically to be during the optical molasses stage, although there isn’t a strong dependency on exactly when during this phase it occurred, as the cloud number and size vary relatively slowly. For the majority of experiments discussed here the beam was switched on 20 ms into the molasses, for a duration of a couple of hundred microseconds.

A typical reflection signal which results from pushing can be seen in Figure 5.1, which shows the mean counts recorded at the APD over 10 shots of the experiment. Clearly visible is a tall, narrow signal at $t = 35$ ms, caused by the pushed atoms passing through. This beam is narrowly focused, forming a waist of $\sim 60 \mu\text{m}^1$ at the MOT. When on continuously it destroys the MOT, but its transient effect is to punch a column of atoms out of the cloud, leaving the rest intact. This explains the presence of the residual atom drop signal also visible in Figure 5.1, centred at $t = 75$ ms. Using Eq.(4.2) to compare the signal amplitudes we find that the pushed signal is nearly 40% higher, at a normalised value of 2.70 ± 0.09 (in lab units), compared to 1.75 ± 0.02 achieved for dropped atoms. Using the signal height to estimate a value for the C_N (and assuming $\theta = 75\%$) gives 1.75 ± 0.39 for the pushed-atom signal and 0.31 ± 0.01 for the dropped-atom signal. Given that in the dropped atom case, $N_{\text{eff}} = 1$, this implies that N_{eff} increases to 5.5 atoms.

We would expect atoms pushed out of the MOT to be travelling much faster than those falling freely under gravity and their respective arrival times at the cavity can be used to estimate their speeds. In Figure 5.1, the push beam is switched on for $300 \mu\text{s}$ from 33.0 ms. The peak of the resulting reflection signal occurs at 33.4 ms, so the atoms take up to $400 \mu\text{s}$ to travel the 5.1 mm between the MOT and the cavity. This implies a minimum speed of 13.5 m s^{-1} . In comparison, the falling atom signal is centred at 74 ms. Assuming the atoms were

¹This is estimated assuming a Gaussian beam propagating from the output collimator of an optical fibre, through a lens and through the chamber towards the MOT.




Figure 5.1: The counts measured at the APD averaged over ten shots of a push beam experiment. The reflection signal resulting from pushed atoms is clearly visible as the tall, narrow feature, centred at approximately 35 ms. The residual reflection signal from falling atoms is centred at approximately 75 ms. The relative signal heights are 2.70 ± 0.09 for the pushed atoms and 1.75 ± 0.02 for the dropped atoms (both in lab units), corresponding to collective cooperativities of 1.75 ± 0.39 and 0.31 ± 0.01 respectively. This implies that the pushed signal corresponds to $N_{\text{eff}} = 5.5$ atoms.

dropped at $t = 40$ ms (i.e. at the end of the molasses stage), this implies a speed of 0.15 m s^{-1} , slower by two orders of magnitude.

Pushing atoms into the cavity also has an equivalently strong effect on the fluorescence signal. Typically, $N_{\text{eff}} = 1$ (from dropped atoms) would produce a signal of ~ 10 photons/ $50 \mu\text{s}$, compared to ~ 100 photons/ $50 \mu\text{s}$ when the atoms are pushed.²

In addition to individual pushes, an alternative timing sequence can be used to produce multiple pulses of atoms in the cavity in a single experimental run. This ‘multi-push’ method involves pushing directly out of the MOT with the push beam and allowing the MOT to re-load in between beam pulses. This method can increase the rate at which data is gathered, although the quality of pushed signals are dependent on the push beam power and repetition rate, both of which set the extent to which the MOT can fully re-load. Generally, lower powers and repetition rates produce a series of signals which decay slowly and/or reach a stable level, making them suitable for gathering multiple shots for averaging. Conversely, using higher powers and repetition rates quickly depletes the MOT, so that the signals decay as a function of push number, allowing a range of N_{eff} to be tested at once.

The plots shown in Figure 5.2 show example APD traces using this method. On the top

²These count rates are those measured at the APD.




Figure 5.2: Example experiment traces using the multi-push method. TOP LHS: The reflection signal, showing 30 pushes out of the MOT at a repetition rate of 2 kHz. Each pulse contains two peaks; one is scattered light from the push beam itself and the other is the reflection signal caused by atoms. This latter signal decays exponentially with time, as indicated by the decreasing, dashed line. The horizontal, dashed line indicates the constant level of the scattered light pulses. TOP RIGHT: Zoomed-in version of the trace on the left, from 32 ms to 54 ms. At this scale it is possible to clearly distinguish between the scattered light pulses and the reflection signal pulses. BOTTOM: The normalised reflection signal in lab units for the data shown on the LHS as a function of pulse number. This demonstrates the utility of this method, which allows us to test amplitudes ranging from $\sim 2.5 - 1$ in one experimental run.




Figure 5.3: Example trace of a multi-push experiment with a sequence alternating between reflection and fluorescence measurements. The trace shows 40 pushes; 20 fluorescence pulses interspersed with 20 reflection pulses. To collect fluorescence the cavity drive beam is shuttered off from the cavity, explaining the alternating height of the background.

LHS is a multi-push reflection signal trace; in this case there are 30 pushes of the beam taken at a repetition rate of 2 kHz. Each pulse shown in the data consists of two peaks. The first peak is caused by scattered light from the push beam and stays constant throughout the trace, as indicated by the level, dashed line. The second peak is the reflection signal from pushed atoms. It begins substantially higher than the scattered light signal, but decays with time, broadly following an exponential decay envelope as indicated by the decreasing, dashed line in the figure. This is the result of the decay of the number of atoms being re-loaded into the MOT. The individual peaks are evident in the zoomed-in version of the trace on the top RHS. Alongside these graphs, shown on the bottom of Figure 5.2 is the normalised reflection signal (or lab unit) as a function of pulse number, which is calculated for the previous data. This shows the amplitude range of the reflection signals, between $\approx 2.5 - 1$.

Figure 5.3 shows a multi-push experiment, alternating between fluorescence and reflection, with 20 pulses each. The fluorescence signals stand out with very high count rates, at ≈ 150 counts collected in $50 \mu\text{s}$ at the beginning of the trace. These fluorescence shots are taken with the cavity pumping light switched off; as evidenced by the drop in background level around the peaks. Either side of the fluorescence peaks are reflection signals, which are only clearly visible at the beginning due to the large scale. As will be discussed later, the reflection signals either side of the fluorescence signal can be used to characterise N_{eff} in the cavity at that point in time. Both signals are subject to the same decaying envelope, again caused by the depletion of the MOT.

5.2.2 Characterisation

One of the first parameters to optimise was the duration of the push beam. Durations ranging between $50 - 300 \mu\text{s}$ were tested, the longest time being set by the arrival of atoms in the cavity. For longer times, this occurred between $300 \mu\text{s}$ and $350 \mu\text{s}$. The peak height at each duration is the average of ten experiment runs and the results are shown on the LHS of Figure 5.4. It appears that the reflection signal amplitude (shown here in lab units) saturates with increasing push beam duration. Alongside the data is a saturation fit, the background level of which is fixed at the reflection signal amplitude achieved with no push (i.e. for dropped atoms). This fit estimates the fully saturated reflection signal amplitude to be 3.04 ± 0.33 .

Physically, the effect of saturation appears sensible as the signal height should be limited by at least two factors. Firstly, the reflection signal is expected to saturate with increasing cooperativity (and thus increasing N_{eff}), given by:

$$\begin{aligned} \text{Normalised reflection signal} &= \frac{1}{v^2} \left(\frac{2C_N + v}{2C_N + 1} \right)^2 \\ \lim_{C_N \rightarrow \infty} (\text{Normalised reflection signal}) &= \frac{1}{v^2} \end{aligned} \quad (5.2)$$

As N_{eff} increases the reflection signal is expected to tend towards a value set by the cavity fringe visibility, v . Here, $v^2 = 0.25$, implying that the signal's saturated value is four lab units. This is too high to be responsible for the saturation seemingly evident in the figure.

We found empirically that the push beam was sensitive to alignment and initial MOT number. Sensitivity to the former seems sensible, as if the beam isn't properly directed to hit the MOT and pass through the cavity mode then the atoms won't a) be pushed out of the MOT by the beam and b) won't pass through the cavities. Sensitivity to MOT number also makes sense, as the push beam can only push out as many atoms as are in the column of the MOT through which it passes. Later, after changing Rb dispensers and optimising the MOT beam alignment and power — resulting in an increased MOT number, the push beam was found to reach higher atom reflection signals of ≈ 3.5 . Whilst this only corresponds to an increase in lab units of $< 20\%$, saturation means that this is actually a large change in N_{eff} . Infact, we'd need to reach an $N_{\text{eff}} > 650$ to be within a quarter of a percent of the limiting value of four, suggesting that it is MOT number which sets the limit on the highest reflection signal we can achieve.

The push beam duration chosen for later experiments was $200 \mu\text{s}$, as here the reflection signal is already close to its saturated value, but it allows a $100 \mu\text{s}$ window between the push beam being switched off and the atoms arriving in the cavity, reducing the possibility of stray




Figure 5.4: Characterising the effect of the push beam. LHS: The normalised reflection signal for different durations of the push beam. Shown alongside the data is a saturation fit, which estimates the fully saturated reflection signal to be 3.04 ± 0.33 . RHS: Fluorescence lineshapes taken for atoms pushed for $300 \mu\text{s}$ (in lighter red) and $100 \mu\text{s}$ (in darker red), alongside Lorentzian fits to the datasets. Both lineshapes are red-shifted with respect to the fluorescence beam (which is located below the MOT and cavity), with their lineshapes shifted by several MHz. As might be expected, the effect is more pronounced for longer durations. The fits produce central frequencies which are red-shifted from resonance by $-9.82 \pm 0.86 \text{ MHz}$ and $-30.79 \pm 1.06 \text{ MHz}$ for $100 \mu\text{s}$ and $300 \mu\text{s}$, respectively. These correspond to velocities of $7.67 \pm 0.67 \text{ m s}^{-1}$ and $24.0 \pm 0.83 \text{ m s}^{-1}$, respectively.

light reaching the APD.

We already have evidence suggesting that the pushed atoms are travelling significantly faster than their dropped counterparts, but we can make a separate estimate of their velocities by measuring their fluorescence lineshapes. Assuming negligible heating, the lineshape would still be expected to be Lorentzian, but now Doppler-shifted from resonance by a frequency, Δf , related to their velocity, v , by:

$$\Delta f = -\frac{v}{c} f_0 \quad (5.3)$$

In the above, c is the speed of light, f_0 the frequency of the transition and the minus sign indicates that the atoms are moving towards the beam driving fluorescence. The RHS of Figure 5.4 shows two fluorescence lineshapes taken for atoms subject to the push beam for two different durations; $300 \mu\text{s}$ (in lighter red) and $100 \mu\text{s}$ (in darker red). Lorentzian fits to these datasets estimate the two velocities to be $7.67 \pm 0.67 \text{ m s}^{-1}$ and $24.0 \pm 0.83 \text{ m s}^{-1}$, respectively. The earlier estimate of the speed based on atom arrival times sits within this range and was calculated for a beam duration of $200 \mu\text{s}$.

A corollary of this is that the interaction time between the atom and cavity is reduced. An attempt to reduce the atom's velocity was made by reducing the push beam power. The data shown above was taken with push beam powers of $> 200 \mu\text{W}$ and it was found that using less power did have some effect. The best result doubled the FWHM of the signal, but at the same time the signal amplitude was reduced to the level of a regular drop, negating any positive effect.

Returning to the lineshapes on the RHS of Figure 5.4, we can also use their width and height to estimate the difference in N_{eff} achieved using different push beam durations. We'd expect a taller and wider lineshape for larger N_{eff} , going as:

$$\text{Height} \propto \left| \frac{2C}{1 + 2CN_{\text{eff}}} \sum_j \frac{g_j}{g_0} \frac{\Omega_j}{2g_0} \right|^2 \quad \text{and} \quad \text{Width} \propto (2CN_{\text{eff}} + 1)\gamma$$

The Lorentzian fit gives a width of $(4.37 \pm 0.59)\gamma$ and $(3.39 \pm 0.38)\gamma$ for the lighter and darker red lineshapes, respectively. This corresponds to $C_N = 1.67 \pm 0.30$ and $C_N = 1.20 \pm 0.18$ respectively and $N_{\text{eff}} = 5.67 \pm 1.02$ and $N_{\text{eff}} = 4.00 \pm 0.61$, so that the extra $200 \mu\text{s}$ of push increases the N_{eff} by 40%. The fluorescence lineshapes found when $N_{\text{eff}} \sim 1$ in the previous chapter had a narrower width, indicating that we can attribute this increase to an increase in N_{eff} , as opposed to other broadening mechanisms.

Alternatively, we can corroborate the expected increase in C_N by comparing the peak amplitudes of the lineshapes. The fits estimate the lighter red peak to be 65.6 ± 4.3 and the darker red to be 47.5 ± 2.3 , equivalent to scattering rates of $949 \pm 45 \text{ kCounts/sec}$ and $1311 \pm 87 \text{ kCounts/sec}$. The pump beam power used in each case was the same (at $37 \mu\text{W}$), so the only differences in scattering rate would arise from different N_{eff} and Doppler shifts, δ . The atomic scattering rate is given by:

$$R_{\text{scatt}} = 2CN_{\text{eff}}\gamma \frac{I/I_{\text{sat}}}{(1 + 2C)^2 + I/I_{\text{sat}} + (\delta/\gamma)^2}$$

As it happens, R_{scatt} is not particularly sensitive to δ and we can make the approximation that $R_{\text{scatt}} = 2CN_{\text{eff}}\gamma$. Comparing the ratios of the scattering rates for each case gives $N_{\text{eff}}^{\text{light}}/N_{\text{eff}}^{\text{dark}} = 1.4$. Again, this implies that the longer push corresponds to an increase in N_{eff} of 40%, consistent with that found when comparing their widths.

We can perform the equivalent reflection lineshape to analyse the increase in N_{eff} and measure the enhanced Purcell broadening. The Lorentzian lineshape produced by scanning the detuning of the cavity pump beam is described by Eq.(4.8), with a width and height given by:




Figure 5.5: The normalised reflection signal with cavity pump beam detuning. The lighter blue dataset corresponds to a detuning scan across the first pulse in a multi-push experiment and the darker blue dataset corresponds to the same scan for the seventeenth pulse. Lorentzian fits to each dataset are shown alongside, producing estimates of $C_N = 3.55 \pm 0.11$ and $C_N = 0.40 \pm 0.04$, respectively.

$$\text{Height} \propto \left(\frac{2C_N + v}{2C_N + 1} \right)^2 \quad \text{and} \quad \text{Width} \propto (2C_N + 1)$$

Reflection lineshape data are shown in Figure 5.5. This scan was taken using the multi-push method, with 100 pulses scanning a range of N_{eff} so that different pulse numbers are associated with different values of N_{eff} . Twenty repetitions of each multi-push experiment were made, so that each final pulse is an average of twenty pulses taken at the same time. This was repeated for various pump beam detunings. The lighter blue data set corresponds to a detuning scan across the first pulse, whilst the darker blue data set corresponds to the seventeenth pulse. Both datasets broadly follow a Lorentzian lineshape and the earlier pulse, which should correspond to a higher N_{eff} , produces one that is both broader and taller. Shown alongside each dataset are Lorentzian fits, which estimate $C_N = 0.40 \pm 0.04$ for the light blue dataset (pulse 17) and $C_N = 3.55 \pm 0.11$ (pulse 1) for the dark blue dataset. These give $N_{\text{eff}} = 11.10 \pm 0.35$ atoms for pulse 1 and $N_{\text{eff}} = 1.24 \pm 0.14$ atoms for pulse 17, indicating that the collective cooperativity of the system has increased by a factor of 10^3 . We can use this large increase to explore how the reflection and fluorescence signals vary as a function of N_{eff} .

³These data were taken after replacing Rb dispensers and re-aligning the push beam, explaining the larger increase in N_{eff} achieved.

5.3 Reflection and fluorescence signals at higher N_{eff}

The presence of multiple atoms in the cavity effects the reflection and fluorescence signals differently. In the case of reflection (and when on resonance) one expects to see the signal saturate with higher N_{eff} . In fluorescence, the situation is complicated by the random phase associated with the light scattered by different atoms, which initially results in an increase in the signal, but ultimately leads to the cavity going dark as N_{eff} increases further. This section tests these effects.

5.3.1 Reflection signals and N_{eff}

Given that we can increase N_{eff} by an order of magnitude, we can also expect to see the effect of the reflection signal saturating. Unfortunately, our only reliable way to calculate N_{eff} is from the reflection signal itself — so when plotting the reflection signal versus N_{eff} it is important that the latter is calculated correctly. To this end, we’ve tested two methods of calculating it:

1. ‘ C_N from height’ This method uses the reflection signal amplitude and the measured cavity fringe contrast, $\theta = 1 - v^2$, to calculate C_N , using Eq.(5.2).
2. ‘ C_N from width’ This method calculates C_N from the spectral width of a reflection signal lineshape, taken using different detunings of the cavity pump beam.

The LHS of Figure 5.6 shows the results of calculating C_N via the two methods. Ideally, one would expect an exact one-to-one mapping between the them, as indicated by the dashed line of unity gradient. The dark blue data show calculations made for dropped atoms and the light blue data show the equivalent for pushed atoms. There is broadly a linear relationship between the two, although with some clear discrepancies. Firstly, the vertical stratum of dropped and pushed atom data are caused by the ‘ C_N from width’ method failing at low values of C_N , even in some cases producing negative values. This is caused by the lower signal to noise ratios at low C_N , which makes it difficult to correctly identify a Lorentzian lineshape to fit to. Conversely, at higher C_N the ‘ C_N from height’ method begins to fail, causing substantial deviations from the unity gradient line. This appears to be due to saturation, which makes it difficult to discriminate between signal heights. This is further corroborated by looking only at the dropped atom data, which samples a much lower C_N regime and shows a much smaller deviation about the line.

These characteristics are, to some extent, reproduced in the RHS of Figure 5.6, which shows the equivalent procedure, but now calculated using ‘synthetic’ data. This is produced




Figure 5.6: Comparison of the different methods for estimating C_N . LHS: A graph of ‘ C_N from width’ versus ‘ C_N from height’, where the dropped atom data is shown in dark blue and the pushed atom data is shown in light blue. Ideally, these two methods should produce identical results and all the data points would lie on a straight line of unity gradient, as indicated by the dashed line. The data do broadly follow this pattern, with some clear exceptions caused by failings in the fitting procedure, discussed in the main text. RHS: A graph of ‘ C_N from width’ versus ‘ C_N from height’, this time calculated using ‘synthetic’ data produced by randomly distributing atoms inside the cavity mode. Some of the features of the real data are replicated here, such as the scatter about the unity gradient line for higher C_N and, to a lesser extent, the stratum of data at very low C_N .




Figure 5.7: A comparison of the ‘known C_N ’ with the ‘measured C_N ’ for the synthetic data, calculated using the single-parameter fitting method. The stratum at low C_N has been eliminated and the scatter about the unity-gradient line is substantially reduced.

by randomly positioning atoms throughout the cavity and making draws of atoms, aiming for a particular value of C_N , ‘known C_N ’. This process is repeated to account for many ensembles of atoms arriving and leaving in a single bin time that we take data and each draw of atoms is used to simulate a reflection signal, from which the C_N is ‘measured’ using the height and the width. Comparing the two datasets, the real data show much worse stratum and appear systematically above the unity-gradient line until $C_N > 4$. Photon shot noise is already included in the synthetic data, but this could be the effect of technical noise which would reduce the signal to noise (i.e. effecting the Lorentzian fits at low C_N) and wouldn’t be reproduced in the synthetic data. Potential sources could be oscillations in the cavity length or variations in the cavity pump beam power.

Alternatively, a single-parameter fit can be used to reduce the impact of these failings in calculating C_N at either end of the cooperativity scale. This optimises the estimate of C_N to fit the height and width simultaneously. The results of using this method are shown in Figure 5.7, which plots the ‘known C_N ’ of the synthetic data against the C_N found using the single-parameter fit. The data now follow the unity-gradient line well, with substantially reduced scatter and no stratum at low C_N , indicating that this method more accurately determines C_N .

We can now employ this method when analysing the reflection signal at high C_N . Data to this effect are shown Figure 5.8, where the normalised reflection signal is plotted as a function of C_N , with data shown for dropped atoms (the dark blue circles) and pushed atoms (the light blue squares). Alongside this are fits to Eq.(5.2). The dashed line in the figure indicates $C = 0.3$, the cooperativity regime accessible when the atoms are dropped and $N_{\text{eff}} = 1$. The push beam makes it possible to achieve higher cooperativities and the saturation of the reflection signal is clearly visible.

There is a slight discrepancy between both the fits and the datasets. This can be accounted by differences in the fringe visibility, v . The data sets are calculated using the measured values, which were $v = 0.27$ for the pushed atoms and $v = 0.26$ for the dropped atoms. The fits also produce slightly different estimates of $v = 0.275$ for the pushed atoms and $v = 0.270$ for the dropped atoms.

We can now use the reflection signal as a way of calibrating the fluorescence signal as a function of C_N .




Figure 5.8: The normalised reflection signal as a function of collective cooperativity, C_N , calculated using the single parameter fitting method. The dropped atom data are shown in dark blue circles, whilst the pushed atom data are shown in light blue squares. Fits to Eq.(5.2) are shown alongside. The discrepancies between the data are caused by differences in fringe visibility.

5.3.2 Fluorescence signals and N_{eff}

At a given point in time, the N_{eff} inside our cavity consists of atoms randomly distributed across the mode. This had no tangible effect on the reflection signal, as the cavity mode was being directly pumped, forming a standing wave inside the cavity which ensured that the atoms scattered in phase and that their effects summed ‘coherently’. This is not the case in fluorescence, as the coupled system is driven from the side. From Chapter 2, this results in the fluorescence signal being given by:

$$J_{flr} = 2\kappa\chi \left| \frac{2\tilde{C}}{1 + 2\tilde{C}N_{\text{eff}}} \sum_j^N \frac{g_j}{g_0} \frac{\Omega_j}{2g_0} \right|^2 \quad (5.4)$$

This equation doesn’t have an analytical solution, although in Chapter 2 the average effect was calculated, with all the atoms being identically pumped at a rate Ω . This showed that J_{flr} initially increased with increasing N_{eff} , but reached a maximum when $N_{\text{eff}} = 1/2C$, whereupon it decreased, tending towards zero as $N_{\text{eff}} \rightarrow \infty$. Effectively, with enough atoms randomly scattering into the cavity, the cavity goes dark, an effect which has been observed experimentally in [113].

To gain some insight into the expected behaviour in our system, we can use a Monte Carlo simulation to randomly position atoms inside the cavity and treat them as scattering independently, using their positions to determine their coupling rate, $g(x, y, z)$ and their pumping rate, $\Omega(x, y, z)$. Their total scattering rate can then be summed and the fluorescence signal




Figure 5.9: Simulations of the behaviour of the fluorescence signal with increasing N_{eff} . LHS: Comparing the effects of having a uniform Ω for all atoms (in light red) to the real case, where Ω_j varies with the position of the j th atom (in dark red). The non-uniform case is simply reduced by $\sim 1/2$, equivalent to the ratio of mean-to-peak intensity, with the rollover position unchanged. RHS: The effect of detuning the driving beam, with a position-dependent Ω_j . Four detunings are shown, with darker to lighter red corresponding to $\Delta_a = 2\pi \times (0, 3, 6, 9)$ MHz, respectively. As the detuning increases, the rollover appears to shift to higher values of N_{eff} .

calculated using Eq.(5.4). In the simulation, 100 throws of atoms into the cavity are made for each value of N_{eff} . Each simulation produces the total fluorescence count rate at the APD and assumes a 10% collection efficiency and drive beam Rabi frequency of $2\pi \times 10$ MHz. The results are shown in Figure 5.9, where two different effects are tested.

The LHS shows a comparison of the fluorescence signals for the case of a uniform beam (in light red) with a position-dependent beam (in dark red), so that the atoms are either driven at an identical rate Ω or at a position-dependent rate Ω_j . The effect appears to shift down the amplitude of the curve, but the general shape and, more specifically, the position of the rollover of the signal, remains unchanged. As a point of interest, a similar curve to the position-dependent case can be produced using a uniform beam with amplitude of $\Omega/2$, equivalent to the spatial average over the cavity mode.

The RHS of Figure 5.9 shows the results of simulations where a position-dependent drive beam is detuned by an amount, Δ_a . Four values of Δ_a are shown; from $\Delta_a = 0$ to $\Delta_a = 2\pi \times 9$ MHz in steps of $2\pi \times 3$ MHz. Increasingly lighter-red data points correspond to larger values of Δ_a . Detuning appears to reduce the amplitude of the curve, but also broadens and shifts the rollover to higher values of N_{eff} . This effect is of particular interest to us because in our experiment the atoms are pushed by the pump beam as they fluoresce, so that we might




Figure 5.10: The fluorescence signal with increasing N_{eff} . LHS: Experimental data showing the fluorescence signal measured for different N_{eff} (bottom horizontal axis) and C_N (top horizontal axis). The dashed line at $N_{\text{eff}} = 1/2C$ indicates where the rollover in fluorescence signal amplitude would be expected to occur. The data appear to increase linearly with N_{eff} at first, consistent with the ‘disordered scattering’ limit - but no rollover is apparent. RHS: Comparison of simulated data (the dark red line) and experimental data from the LHS (the light red points). The simulated data assumes a drive beam Rabi frequency of $2\pi \times 10$ MHz and a detuning of $2\pi \times 5.4$ MHz. There is a good match between them, indicating that detuning may well be the reason for the lack of rollover in the signal.

expect them to be detuned due to the Doppler effect.

These simulations can be compared to the LHS of Figure 5.10, which shows experimental data of the fluorescence signal as a function of N_{eff} (bottom axis) and C_N (top axis), assuming $C = 0.3$ to convert between them. These data are taken using the multi-push method, with 100 pulses alternating between reflection and fluorescence signals. In this way, the ‘equivalent’ reflection signal can be calculated for a given fluorescence signal by taking the mean of the reflection signals either side of it. Multi-push shots were then taken at nine different pump beam detunings, so that each fluorescence signal is associated with a reflection signal height and lineshape, allowing C_N to be calculated using the single-parameter fit. The fluorescence signals are the summed counts for ten shots taken with $\tau_{\text{int}} = 50 \mu\text{s}$.

The data shown here at first appear to increase linearly with N_{eff} , which is consistent with the expected behaviour in the ‘disordered scattering’ limit. In the absence of other effects, we’d expect this signal to peak at $N_{\text{eff}} = 1/2C = 1.6$, but no rollover is evident and the signal still appears to be increasing at $N_{\text{eff}} \sim 3$. This cannot be firmly tested unless higher values of N_{eff} are achieved, but we can gain more insight using our simulations. The RHS of the same figure

compares these data (the light red circles) with a simulation (the dark red line), which includes a detuned, position-dependent beam; $\Omega = 2\pi \times 10$ MHz and $\Delta_a = 2\pi \times 5.4$ MHz. Different combinations of Ω and Δ_a were tested, but these values appear to show the best agreement. This suggests a potential reason for our apparent lack of rollover, which might be caused by a detuning between the pump light and the atoms⁴.

One potential source of detuning is Stark-shifting of the atoms caused by the presence of locking light circulating inside the cavity. This topic is discussed in more detail in the following chapter, but the degree of shifting seen here would require a minimum input power of $100 \mu\text{W}$, which is much higher than that used in this experiment. More likely is that this is caused by the Doppler shift as the atoms are being pushed and we know from earlier characterisation of the push beam that Doppler shifts of several tens of MegaHertz are possible. To confirm this hypothesis the experiment should be repeated, but with the resonance frequency of the atoms found for each value of N_{eff} beforehand.

These experiments have given us an insight into the behaviour of the reflection and fluorescence signals at higher N_{eff} . In the following section higher values of N_{eff} are used to enhance the cavity interaction, measured using two-dimensional scans of reflection and fluorescence signals.

5.4 The dispersive atom-cavity interaction: 2D scans of reflection and fluorescence signals

We can measure the excitation spectrum of the atom-cavity system by taking two-dimensional scans of the reflection and fluorescence signals. Such scans involve detuning both the pumping beam driving the atoms by an amount Δ_a and the cavity by an amount Δ_c , where these detunings are defined as:

$$\Delta_a = \frac{1}{\gamma} (\omega_L - \omega_A) \quad \text{and} \quad \Delta_c = \frac{1}{\kappa} (\omega_L - \omega_C)$$

Here ω_L , ω_A and ω_C correspond to the angular frequencies of the drive beam, free-space atom and cavity respectively. The signals are now given as a function of these detunings as:

⁴This detuning does not effect the estimate of N_{eff} , which is made using the reflection signal and the 'single-parameter' fitting method. Because the reflection signal involves only weakly pumping the atoms, the light-force effects on the atoms as they pass through the cavity is negligible. In contrast, the beam which produces fluorescence signals drives the atoms relatively strongly and accelerates the atoms.

$$\text{Reflection signal} = \left| 1 + \frac{1 - \sqrt{v}}{(1 + i\Delta_c)(1 + \frac{2CN_{\text{eff}}}{(1+i\Delta_c)(1+i\Delta_a)})} \right|^2 \quad (5.5)$$

$$\text{Fluorescence signal} = 2\kappa\chi \left[\frac{\Omega\sqrt{N_{\text{eff}}}}{2g} \right]^2 \left| \frac{\frac{-2C}{(1+i\Delta_c)(1+i\Delta_a)}}{1 + \frac{-2CN_{\text{eff}}}{(1+i\Delta_c)(1+i\Delta_a)}} \right|^2 \quad (5.6)$$

From this, we know that a finite detuning between the atom and the cavity field will result in both the fields corresponding to the signals having an imaginary part. This implies that the interaction is dispersive, producing a phase-shift in the internal cavity field which makes the frequency of the resonance change. When we take two-dimensional scans of the reflection signal we are mapping out the excitation spectrum of the coupled atom-cavity system [68], so that we can observe this dispersion and measure its change with increasing collective cooperativity. Because measuring in reflection one automatically drives the coupled atom-cavity system (because the cavity mode is being pumped), one expects to see evidence of the coupling in the excitation spectrum for any finite cooperativity.

Interestingly, the same isn't always true for the fluorescence signal. This can lead to a qualitative difference between the signals which, up to now, hasn't been evident. When the coupling is sufficiently weak, the driven atom emits a photon which is almost immediately 'lost' by the cavity and measured at the APD. To this end, the atom and cavity are still behaving as separate entities. We shouldn't expect to see strong evidence of the dispersive interaction in the two-dimensional scan of the excitation spectrum because essentially, we are just measuring the excitation spectrum of the atom. As the interaction strength increases, however, one is increasingly driving the coupled system, so that the spectrum should change.

This qualitative difference is shown in Figure 5.11, which shows both types of signal as a function of detuning, Δ , where now $\Delta = \Delta_a = \Delta_c$, so that the pumping beam and the cavity are detuned by the same amount (relative to their linewidths). The LHS (RHS) shows the reflection (fluorescence) signal as a function of Δ . All three sets of graphs are calculated assuming $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz, with N_{eff} changing for each set.

The top graphs show the case of $N_{\text{eff}} = 1$. On the LHS, the reflection signal shows the Lorentzian profile of the cavity. Close to $\Delta = 0$, however, there is a small perturbation from the lineshape, which is caused by the presence of a single atom. Thus, even one atom has a noticeable effect on the cavity lineshape, which shows the emergence a small splitting. The RHS shows the fluorescence signal. The presence of the cavity is, in this case, manifest in the lineshape broadening, caused by the Purcell effect, but, in contrast to the reflection signal, the




Figure 5.11: The reflection (LHS) and fluorescence (RHS) signals, calculated as functions of detuning, Δ , where $\Delta = \Delta_c = \Delta_a$. These are shown for three different values of N_{eff} , with the cavity parameters kept constant at $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz respectively. TOP: $N_{\text{eff}} = 1$. For the reflection signal, the Lorentzian cavity profile is visible, with a small perturbing feature caused by the presence of an atom, centred at $\Delta = 0$. In contrast, the fluorescence signal is a Purcell-broadened Lorentzian, but shows no other features caused by being coupled to the cavity. MIDDLE: $N_{\text{eff}} = 1.6$, which corresponds to $N_{\text{eff}}g^2/\kappa\gamma = 1$. The perturbation at the centre of the reflection signal has increased, but now the fluorescence signal is qualitatively different, with the lineshape flattened in the region of $\Delta = 0$ and the second derivative of the signal equalling zero when $\Delta = 0$. BOTTOM: $N_{\text{eff}} = 80$. The coupling is now substantially larger, although it is not yet into the strong coupling regime, where $g\sqrt{N_{\text{eff}}} \gg (\kappa, \gamma)$. The reflection and fluorescence signals now both show a clearly resolved splitting.

shape of the signal is identical to a free-space atom being driven by a beam of varying detuning. The lineshape is a Lorentzian, centred at $\Delta = 0$.

The fluorescence signal shape noticeably changes, however, for the slightly larger value of $N_{\text{eff}} = 1.6$, shown in the middle graphs. This corresponds to $N_{\text{eff}}g^2/\kappa\gamma = 1$, so that the coherent interaction rate is equal to the loss rates. Now, in the region of $\Delta \sim 0$, the Lorentzian profile of the fluorescence signal has been flattened at the peak, with the second derivative going to zero when $\Delta = 0$. This coupling is now sufficient to change the shape of the fluorescence signal. In reflection, the perturbation close to $\Delta = 0$ has increased, increasing the splitting of the cavity lineshape.

Once more strongly coupled, both the reflection and fluorescence signals show a resolved splitting, with each peak corresponding to the energy of an eigenvalue of one of the system's dressed states. This is clear from the bottom graphs, which are calculated for the case when $N_{\text{eff}} = 80$. This splitting is still well-resolved, despite not being in the strong coupling regime, when $g\sqrt{N_{\text{eff}}} \gg (\kappa, \gamma)$.

As discussed in Chapter 2, the splitting evident in Figure 5.11 arises from the mixing of the atom and cavity states into ‘dressed’ states. The splitting can only be resolved when the atom and cavity are strongly coupled and is manifest as an ‘avoided crossing’. We cannot expect to resolve this for our system parameters, but we can hope to observe the underlying interaction through a splitting and shifting of the cavity resonance when probed in reflection. In addition, we can take advantage of our relatively high fluorescence collection efficiency to show the qualitative differences between the reflection and fluorescence signals.

5.4.1 2D scan in reflection

As mentioned, conducting a two-dimensional scan of the reflection signal should demonstrate the underlying interaction between the atom and the cavity field. The structure we’d expect to observe is shown in Figure 5.12. For simplicity, these plots are calculations of the transmission spectrum with detuning, although the same features will be observed in reflection. The two smaller plots on the LHS show the spectra for an empty cavity (top) and a single, free-space atom (bottom). Since here they are treated as separate entities, they only show a frequency dependence with Δ_a (the horizontal axis) in the case of the atom and Δ_c (the vertical axis) in the case of the cavity. For the atom the spectrum has a band of frequencies where the transmission is minimised, as a result of absorption. The cavity has a band of frequencies which fulfil the resonance condition of the cavity, maximising its transmission.




Figure 5.12: The two-dimensional transmission spectrum of our cavity, with and without an atom. For each plot, Δ_a is given by the horizontal axis and Δ_c is given by the vertical axis. LHS: The separate transmission spectra for the cavity (top) and the atom (bottom). Each show a narrow band of frequencies when light is either transmitted (by the cavity) or absorbed (by the atom). MIDDLE: The transmission spectrum of light passing consecutively through an atom and a cavity. The cavity resonance is split where the atom absorbs strongly, but there is no dispersive effect as the atom and cavity do not interact. RHS: The transmission spectrum from a coupled atom-cavity system. The cavity resonance is still split, but now the dispersive effect of their interaction causes the resonance to shift in frequency. The white, dashed line indicates the value of Δ_c corresponding to the maximum in transmission as one sweeps Δ_a .

The middle plot shows the transmission spectrum which would result from the light passing consecutively through the cavity and the atom, so that there is no interaction between them. This is just the product of the separate atom and cavity spectra shown on the LHS. The cavity transmission band is now split for the band of frequencies where the atom strongly absorbs.

The plot on the RHS shows the transmission spectrum which results from an atom and cavity which are coupled. As in the uncoupled case, the cavity resonance is still split, but now the dispersive nature of the interaction results in the cavity resonance either side of the splitting being frequency-shifted. The white, dashed line shown alongside the spectrum plots the value of the cavity detuning corresponding to the transmission maximum, as a function of the atomic detuning and it maps out a dispersion as one sweeps through the atomic resonance. This dispersion should increase with increasing coupling.

From a practical perspective, a two-dimensional scan requires the detuning of the cavity and the pump beam driving the atoms to be varied independently. For these experiments

the pump beam detuning was varied as before, using an AOM. To independently detune the cavity, the wavelength of the cavity laser (used to produce an error signal for the cavity lock) was varied by changing the current to the laser diode. This changes the lock point relative to the pump beam resonance, detuning the cavity from the pump beam by a controllable amount. For each pump beam detuning, the cavity itself was scanned through 21 detunings, from one side of the cavity resonance to the other.

One such scan of cavity detunings for one frequency of the pump beam takes approximately 20 minutes to complete and, during this time, the cavity can drift. In addition, the change in the current to the laser diode does not produce a linear change in the cavity lock point. Both of these effects mean that the cavity detunings for each data point had to be empirically determined afterwards. To this end, a record was made of the pump beam level inside the ‘empty’ cavity each time the laser current was changed. Having scanned the cavity across 21 detunings, this produced a set of 21 different ‘empty’ cavity pump beam count levels, effectively mapping out a cavity lineshape. Fitting a Lorentzian of fixed FWHM, 2κ to this lineshape allows the cavity detuning for each point to be determined.

One such set of cavity detunings were produced for each pump beam detuning, resulting in 14 cavity detuning scans. These were then ‘lined up’ by lining up the centres of each Lorentzian. In addition, the multi-pushing technique was used for each run, so that each data point is an average over 30 reflection pulses. The final datasets consisted of 294 three-dimensional data points: (pump beam detuning, cavity detuning, normalised reflection signal averaged over 30 pushes). These datasets can then be presented in contour plots by extrapolating the signal level between each discrete point. Contour plots for different values of N_{eff} are shown on the LHS of Figure 5.13. For each, the value of N_{eff} is found by fitting each dataset to Eq.(5.5). The top plot is for $N_{\text{eff}} = 0$; the middle plot is for $N_{\text{eff}} = 1.1$; and the bottom plot is $N_{\text{eff}} = 7.1$.

For comparison, the RHS of Figure 5.13 show Eq.(5.5) plotted for the same cavity contrast of $\theta = 76\%$, fixed cavity parameters $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz, using the values of N_{eff} found from fitting to the LHS data.

The first set of contour plots are displayed on the top of Figure 5.13. They correspond to an empty cavity, with $N_{\text{eff}} = 0$. Both the theory and the data plots show a dark, horizontal stripe along the line of $\Delta_c = 0$, showing the empty cavity profile.

The middle plots shown in Figure 5.13 show a later time when there are atoms inside the cavity. Fitting to Eq.(5.5) estimates $C_N = 0.34 \pm 0.04$, which assuming a single atom cooperativity of 0.3 means that $N_{\text{eff}} = 1.1 \pm 0.1$. The equivalent theory plot of the RHS shows




Figure 5.13: Two dimensional scans of the cavity reflection signal with respect to the cavity and pump beam detuning. The LHS of each plot shows real data, the RHS shows theory plots to Eq.(5.5). Three different values of N_{eff} are shown. TOP: $N_{\text{eff}} = 0$. MIDDLE: $N_{\text{eff}} = 1.1$. BOTTOM: $N_{\text{eff}} = 7.1$.




Figure 5.14: The value of Δ_c at which the reflection signal is minimised, plotted as a function of Δ_a . Each plot corresponds to a contour plot shown above; darker to lighter blue gives of $N_{\text{eff}} = 0, 1$ and 7 . The dispersion clearly increases with larger coupling.

the case of $N_{\text{eff}} = 1.1$. In both plots the cavity resonance is still evident as a dark horizontal stripe along the line of $\Delta_c = 0$, but the cavity resonance is clearly split close to $\Delta_a = 0$. Looking closely, there is already evidence of the resonance being shifted, even for $N_{\text{eff}} \sim 1$.

The bottom plots shown in Figure 5.13 show an even later time in the experiment. Fitting to Eq.(5.5) estimates $C_N = 2.1 \pm 0.1$, corresponding to a $N_{\text{eff}} = 7.1 \pm 0.5$. The theory plot for this value of N_{eff} is shown on the RHS. The splitting and shifting of the cavity resonance now dominates both plots, with the unperturbed resonance only being recovered for detunings $> 10\gamma$. The resonance is now obviously shifted, with strong evidence of the dispersive effect of the atom-cavity interaction.

One feature which is common to all the data plots is that the cavity fringe appears narrower than their theory counterparts. This is a feature of the data presentation — the contour plots extrapolate between points and, because of the limited resolution of the data scans, only a few points sample the deep sections of the cavity fringes, making them appear narrower.

As discussed, for a coupled atom-cavity system, the minimum in the reflection signal maps out a dispersion as one sweeps through the atomic resonance. This is evident in Figure 5.14, which shows the value of Δ_c at which the reflection signal is minimised as a function of Δ_a . Darker to lighter blue correspond to increased values of N_{eff} ; $0, 1$ and 7 , as calculated from the contour plots already shown. The dispersion increases as the ratio of coupling to losses increases, so that in the limit of zero losses there is a discontinuity at $\Delta_a = 0$.

From this, we know that even with one atom in the cavity there is a clear dispersive effect which is readily apparent in the reflection spectrum. The equivalent experiment in fluorescence

is the topic of the following section.

5.4.2 2D scan in fluorescence

We expect that measuring the excitation spectrum of the atom-cavity system with the fluorescence signal will result in spectra qualitatively different from the reflection signal. We use a similar method for taking the two-dimensional fluorescence scan as was used before in the equivalent experiment in reflection, but now the cavity probe is turned off and the atoms are pumped from the side. The cavity lineshape is scanned in the same way, with measurements taken at 20 different cavity detunings for 11 detunings of the pump beam. At each cavity and pump beam detuning point, 10 shots of the experiment were taken and averaged to give a final signal. In the experiment, the pump beam came vertically, from above the cavity, with a beam power of $40\ \mu\text{W}$, chosen to strike a balance between producing a relatively large fluorescence signal whilst not saturating the atoms.

For the beginning of the experiment and whilst the atoms were falling through the mode, the cavity pump beam was blocked by a shutter so as to not saturate the APD and make the fluorescence emitted by the photons impossible to detect. It was then unblocked for the final 50 ms of each experiment run, to allow us to back-calculate the cavity detuning, following the method described earlier.

The counts recorded at the APD are shown in the two contour plots at the top of Figure 5.15, taken for dropped atoms with $N_{\text{eff}} \approx 1$. Each plot shows the counts averaged over different time durations, all starting at the time when the pump beam is switched on. The top left plot shows the counts averaged over $10\ \mu\text{s}$, corresponding to one integration bin. The top right plot shows the counts averaged over $100\ \mu\text{s}$, or ten integration time bins. Comparing the two plots, whilst the main features of the data remain the same, the second is smoother as more time bins are averaged. Much of the original ‘spikyness’ of the data comes from the limited resolution of the scans; were another such scan to be taken the step-size between detunings should be reduced. To ensure that summing over multiple integration bins maintained the integrity of the data, however, the atomic and cavity linewidths at $\Delta_a = 0$ and $\Delta_c = 0$ were compared for summed and non-summed data. It was found that, within the bounds of statistical uncertainty, the results were equivalent.

In addition, the process of centering each drive beam detuning scan means that the range of Δ_c tested are different. This makes the data appear patchy towards the edges of the scan. To present a fuller picture for ready comparison, each corner has been set to have zero counts.




Figure 5.15: Two-dimensional scans of the fluorescence signal. TOP LEFT: The counts recorded at the APD as a function of pump beam detuning, Δ_a , and cavity detuning, Δ_c , integrated over $10\mu\text{s}$, or one time bin. TOP RIGHT: The same as on the LHS, but now integrated over $100\mu\text{s}$, or ten integration time bins. BOTTOM: A theory plot of the same two-dimensional scan, calculated for our cavity parameters and $N_{\text{eff}} = 1.1$ and using Eq.(5.5).

We don't expect to be able to see any splitting in the fluorescence signal at this value of N_{eff} . This is confirmed by the data plots, which show only a single peak. Fitting the $100\ \mu\text{s}$ binned data to Eq.(5.5) estimates $N_{\text{eff}} = 1.09$ and $\Omega = 2\pi \times 46\ \text{MHz}$, assuming a total collection efficiency of 10 %. The value of Ω is particularly sensitive to the value of the total collection efficiency used in the fit, however the estimate for N_{eff} is not.

These plots can be compared to the bottom plot in the same figure, which shows a theory plot of the expected fluorescence signal as a function of detuning, calculated using Eq.(5.5), assuming cavity parameters $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)\ \text{MHz}$ and $N_{\text{eff}} = 1.1$. The main feature of the real data are replicated here; with a single peak close to $\Delta_a = 0$, in contrast to the equivalent two-dimensional reflection scan. The theory plot in two-dimensions does show evidence of the dispersive interaction which couldn't be seen in the plots of Figure 5.11; the Lorentzian is clearly distorted along the line of $\Delta_a = \Delta_c$. Unfortunately, this is not clearly evident in the real data, although this could be caused by noise. Despite this, however, these data do demonstrate a clear distinction between reflection and fluorescence signals which can only be viewed for weak coupling.

We do expect, however, that this will change as N_{eff} increases. The contour plots shown in Figure 5.16 show the expected fluorescence signal as a function of Δ_a and Δ_c for the cases of $N_{\text{eff}} = 1.6$ (LHS) and $N_{\text{eff}} = 7$ (RHS). For $N_{\text{eff}} = 1.6$, the distortion of the central peak is enough to produce the flat-topped function, shown in Figure 5.11. By, $N_{\text{eff}} = 7$ however, two distinct peaks are quite clear and the signal becomes qualitatively similar to the reflection signal. Unfortunately, technical limitations at the time meant that we were unable to reproduce this result experimentally.

One question arising from this experiment is the role of dynamics. We know from earlier measurements that the fluorescence beam can cause the atoms' resonance frequency to shift. Looking at the RHS plot of Figure 5.15, a slight shift of the location of zero detuning on the Δ_a axis is already discernable. This effect becomes increasingly apparent as more integration bins are averaged over. Figure 5.17 shows the same 2D scan, but now after averaging over $500\ \mu\text{s}$. It is now readily apparent that the atomic resonance frequency has been shifted by $\sim 2.9\gamma$.

This effect can be quantified by analysing the frequency shift as a function of integration time. To do this we sum the data for varying integration times and then fit a Lorentzian lineshape to the data corresponding to $\Delta_c = 0$. The results are shown in Figure 5.18. For averaging times $< 100\ \mu\text{s}$ the shift is relatively constant at $\gamma/2$. That this does not appear to vary significantly with averaging time indicates that it is not a result of dynamics, but perhaps




Figure 5.16: Theory plots showing two-dimensional scans of fluorescence signal, calculated for our cavity parameters and using Eq.(5.5). LHS: For $N_{\text{eff}} = 1.6$. The atom-cavity interaction has sufficiently distorted the Lorentzian peak along the line of $\Delta_a = \Delta_c$ to produce the flat-topped function shown earlier. RHS: For $N_{\text{eff}} = 7$. The atom-cavity interaction is now sufficient to split the resonance, becoming qualitatively similar to the reflection signal.

an offset caused by a mis-estimate of the zero in atomic detuning, or the result of some time-independent frequency shift, such as a Stark shift from the far-detuned light inside the cavity used for locking. For longer timescales, however, the shift increases quickly, tending towards a limiting shift of $\sim 2.9\gamma$ on timescales of the order of $400\ \mu\text{s}$. Physically, this appears to indicate that dynamics only play a role for longer timescales and further corroborates the decision to sum over integration bins for times of $< 100\ \mu\text{s}$.

The discussions of this chapter have highlighted at least two areas for further experiments; in the near future we should work towards experimentally re-producing the roll-over in the fluorescence signal with higher N_{eff} and measuring two-dimensional scans at higher N_{eff} so as to witness the resolved splitting of the resonance peak. The following chapter discusses other immediate goals for our experiment.




Figure 5.17: The recorded APD counts measured as a function of detuning. The counts are averaged over 50 integration bins, equal to $500 \mu\text{s}$. The shift of the atomic resonance with respect to the pump beam frequency is equal to $\sim 2.5\gamma$.




Figure 5.18: The shift of the atomic resonance from $\Delta_a = 0$ as a function of averaging time. The value for each shift was produced by a making a Lorentzian fit along the line of $\Delta_c = 0$ for each averaging time. On timescales of $100 \mu\text{s}$ or less, the shift remains constant at $\gamma/2$, likely due to a mis-estimate of where zero detuning actually was. For timescales longer than $100 \mu\text{s}$, the shift increases and quickly tends towards a limiting value of 2.9γ .

Chapter 6

Future directions

There are currently two areas of immediate focus for our experiment. The first area of work involves the installation of an intra-cavity optical dipole trap and the provision of controllable transport to load atoms from the MOT into said trap. The beginning of this Chapter briefly reviews optical dipole traps, discusses the progress made thus far on the installation of such a trap in our system and presents evidence of the trap inside the cavity in the form of a Stark shift on passing atoms. This is followed by a discussion of the optical ‘conveyor’ planned and being built as the transport mechanism. The second area of work is to coherently control an atom whilst in the cavity — preparing the way to use our system as a means to store and transfer quantum information, or as a deterministic source of single photons. One such coherent effect resulting from the interaction between a ‘three level atom’ and two coherent light fields is electromagnetically induced transparency (EIT). The second part of this Chapter will provide a brief review of EIT and will discuss the progress thus far towards viewing the effect.

6.1 Conveying and trapping atoms inside the cavity

In this section the progress towards controllably loading and trapping an atom inside the cavity is presented. The use of an optical dipole trap is discussed first.

6.1.1 An intra-cavity dipole trap

In the near future goal we want to trap and store an atom inside the cavity. A standard way of doing this uses an optical dipole trap, see for example, [25]. The force used when optical trapping can be either attractive or repulsive, depending on the relative detuning between the

dipole trapping light and the relevant atomic transition. One can either implement an intra-cavity dipole trap [26, 27], where the dipole trap has a standing wave along the longitudinal axis of the cavity, or a trap aligned transverse to the longitudinal axis, allowing the trap to also be used for transport in and out of the cavity [28, 31].

Optical dipole traps exploit the interaction between the induced atomic dipole moment caused by the presence of a light field and the intensity gradient of that light field. The trapping light will have a spatial intensity distribution $I(\mathbf{r})$ and a detuning, Δ_{trap} , from the atomic resonance at ω_A such that $\Delta_{trap} = \omega - \omega_A$. In the limit of large detuning¹, the interaction results in the atom feeling a pseudo-conservative potential, U_{dip} and scattering photons from the field at a rate, Γ_{scatt} given by:

$$U_{dip}(\mathbf{r}) = \frac{3\pi c^2}{2\omega_A^3} \frac{\Gamma}{\Delta_{trap}} I(\mathbf{r}) \quad \text{and} \quad \Gamma_{scatt}(\mathbf{r}) = \frac{3\pi c^2}{2\hbar\omega_A^3} \left(\frac{\Gamma}{\Delta_{trap}} \right)^2 I(\mathbf{r}) \quad (6.1)$$

U_{dip} becomes a truly conservative potential in the limit when $\Gamma_{scatt} \rightarrow 0$. These equations demonstrate some generic features about dipole traps. For red-detuned light (i.e. $\Delta_{trap} < 0$), atoms feel an attractive potential, leading them to regions of high intensity. In this case, the potential minima occur at the intensity maxima of the trapping light. The reverse is true for blue-detuned light (i.e. $\Delta_{trap} > 0$), when atoms are attracted to regions of low intensity. From Eq.(6.1) we know that $\Gamma_{scatt} \propto I(\mathbf{r})/\Delta_{trap}^2$, whereas $U_{dip} \propto I(\mathbf{r})/\Delta_{trap}$, leading to a natural balance, when using red-detuned light, between achieving high trap depths whilst also having low scattering rates, which result in atom loss through heating. Given this, it is generally better to operate far-detuned with high intensities.

This balance can be avoided by using a blue-detuned trap, as atoms are trapped at regions of low intensity and thus naturally have a lower scattering rate. This has already successfully implemented in a cavity QED experiment [114]. This cannot be used in our system, however, as it requires a trapping beam configuration which provides full three dimensional confinement. In our experiment, the output mode of the single mode fibre forming one end of our cavity limits the spatial distribution of light inside the cavity, so that we only excite the TEM₀₀ mode.

Eqs.(6.1) are valid for a two level atom with no internal structure. This isn't the case for a ⁸⁷Rb, which has fine and hyperfine structure which should be taken account using state-dependent ground-state potentials [115]. This approach can be simplified for the case of large trap detunings, when the hyperfine splitting, Δ_{HFS} , is not resolved. This is justified for the

¹By 'large' we mean large enough that the atom can be treated as having no internal structure, i.e. $\Delta_{trap} \gg$ the fine and hyperfine structure splitting.




Figure 6.1: The dipole trap depth and atomic scattering rate as a function of trap wavelength, calculated using Eq.(6.2) and assuming an input power of $500 \mu\text{W}$. The trap depth, shown in the lighter green, is given by U_{dip}/k_B and for red-detuned light gives a negative potential, attracting the atoms to regions of high intensity. The scattering rate, Γ_{scatt} , shown in the darker green shows two poles in Γ_{scatt} , corresponding to the D_2 and D_1 transitions.

case of a dipole trap with $\lambda_{trap} = 804 \text{ nm}$, which is detuned by $\Delta_{trap} \sim 4$, and 11 THz from the D_2 and D_1 lines, respectively. This should be compared to a largest hyperfine splitting of $\Delta_{\text{HFS}} \sim 7 \text{ GHz}$. In contrast, Δ_{trap} is of a similar size to the fine structure splitting, $\Delta_{\text{FS}} \sim 7 \text{ THz}$, so that it is appropriate for this level of internal structure to be taken into account, via a modified form of Eq.(6.1):

$$U_{dip}(\mathbf{r}) = \frac{\pi c^2 \Gamma}{2\omega_A^3} \left(\frac{2}{\Delta_{2,F}} + \frac{1}{\Delta_{1,F}} \right) I(\mathbf{r}) \quad \text{and} \quad \Gamma_{scatt}(\mathbf{r}) = \frac{\pi c^2 \Gamma^2}{2\hbar \omega_A^3} \left(\frac{2}{\Delta_{2,F}^2} + \frac{1}{\Delta_{1,F}^2} \right) I(\mathbf{r}) \quad (6.2)$$

These equations are valid for linearly polarised light.

As discussed, the choice of λ_{trap} requires balancing the desire for a deep trap with that for a low scattering rate. Figure 6.1 shows the dipole trap depth, U_{dip}/k_B in milliKelvin (lighter green) and Γ_{scatt} in kHz (darker green), calculated as a function of λ_{trap} using Eq.(6.2). Clearly visible are the two poles of Γ_{scatt} at $\lambda = 780 \text{ nm}$ and $\lambda = 795 \text{ nm}$, corresponding to the D_2 and D_1 transitions respectively. For a red-detuned dipole trap we need a negative value of U_{dip} and from the figure it appears that the best balance occurs for $800 \leq \lambda_{trap} \leq 820 \text{ nm}$. As discussed in Chapter 3, the cavity laser, which is also used for dipole trapping, has $\lambda = 804 \text{ nm}$, chosen to sit within this region.

The cavity laser providing the dipole trapping light can supply up to $500 \mu\text{W}$ of input power to the cavity. Including the effects of optical build-up, this leads to trap depths of

up to $920\text{ }\mu\text{K}$. At this power, the scattering rate is 130 photons/sec. Assuming each atomic scattering event leads to an energy increase of 362 nK [101] and that an atom is lost when the total additional energy from scattering equals $U_{dip}/2$, one would expect trap lifetimes of $\sim 10\text{ s}$. This is optimistic however, as it ignores other heating effects, such as stochastic heating caused by intensity fluctuations in the trapping light.

A large scattering rate is not the only reason, however, to compromise on trap depth. One technical limitation is that our signals are small amounts of resonant or near-resonant light leaving the cavity and measured at an APD. Large amounts of trapping light, which co-circulates inside the cavity, can obscure these signals. This is mitigated to some extent by using several dichroic reflection and transmission filters in the optical path between the cavity and the APD, as discussed in Chapter 3. The combined attenuation factor is $\sim 10^{11}$, which means high powers still result in a visible signal at the APD. For example, $500\text{ }\mu\text{W}$ of trapping light entering the filtering system would result in 0.1 counts per $5\text{ }\mu\text{s}$ at the APD. Given that our peak fluorescence signal is ~ 1 counts per $5\text{ }\mu\text{s}$, we would not want to use an input power substantially larger than this.

Another negative effect resulting from high trapping powers inside the cavity is optomechanical. Input powers of only several hundred microwatts can result in a sizeable radiation force on the cavity mirrors, given by:

$$F_{rad} = \frac{IA}{c} \quad \text{where} \quad A = \pi w_0^2$$

For our cavity parameters and with a circulating power of 15 mW (assuming an input power of $250\text{ }\mu\text{W}$ and a cavity buildup factor of $\sim \mathcal{F}/\pi$), this results in a force of $\sim 210\text{ pN}$. Whilst on sight this appears small, only a small change in cavity length is needed to have an effect — in our case, the effect of input powers at this level are clearly visible in scans of the cavity fringes, which become sharply asymmetric. Considering the cavity design, it seems most likely that the fibre-tip mirror is moving, as this is glued onto a fibre which overhangs the supporting macor structure. In contrast, the other end of the cavity is glued to the peizo and held by a large macor block. In addition, from observing the fringes it seems unlikely that this is a thermal effect, as they do not slowly drift in frequency. This does indicate a potential area for improvement in the cavity design for future evolutions of the experiment; ideally the cavity stiffness should be increases to reduce any length changes to sub-picometre scales.

The presence of the trapping light also has an affect on the ground and excited atomic states, which are shifted by an equal and opposite amount $\Delta E = U_{dip}$ as a result of their dipole interaction with the light. This is known as the AC Stark shift. This effect results in




Figure 6.2: The Stark shift caused by increased trapping light circulating inside the cavity. LHS: Atomic lineshapes measured for different intra-cavity trap powers. The detuning of the pump beam is varied with respect to the cavity resonance. Lighter to darker blue correspond to lower to higher trap powers. The shift of the resonance is visible, but is accompanied by a reduction in the amplitude and an increase of the width as a result of different atoms being Stark-shifted by different amounts. RHS: A comparison of the increase in measured Stark shift as a function of the increase in power, taken from the lineshapes centres on the LHS. The increase in power in the dark blue dataset is estimated by measuring the reflected trap power and using this to estimate a cavity pumping rate. The increase in power in the light blue dataset is estimated from the counts at the APD when the cavity is only pumped with trap light. For both, the Stark shift increases with increasing power, as would be expected.

a combined frequency shift of $2\Delta E$, detuning the atom from the pump beam and the cavity. Using Eq.(6.2), a circulating power of 15 mW is expected to produce a maximum light shift of 19.2 MHz for our atom-cavity parameters, assuming that the atom is static and positioned at an antinode of the trapping light. This is inconvenient, but not impossible to overcome as the cavity and pumping light can be detuned relative to the atom to accommodate this change².

Other experiments have noted the affect of the light shift. One complicating factor is the dependence of the shift on the atom's magnetic hyperfine internal state. This could be overcome by using a dipole trap at a 'magic wavelength' [116], whereby the trapping light is chosen to have a wavelength at which the excited and ground state Stark shifts are equal and are relatively constant across m_F states.

The affect of the light shift is already visible on atoms falling through the cavity. The LHS of Figure 6.2 shows reflection lineshapes taken for dropped atoms, at four different levels of

²As a point of interest, it should be noted that other cQED experiments use the Stark shift as a way of controllably detuning the atom from the cavity/pumping light.

intra-cavity power, where low to high powers go from dark to light blue. As the power increases, the lineshapes are shifted further blue (evident by comparing the centres of each lineshape) and the overall amplitude of the lineshape is reduced. This latter effect wouldn't be anticipated for a single atom trapped in a dipole trap, however here there are several atoms contributing to the signal, each Stark-shifted by different amounts. This effect reduces N_{eff} , as at any given frequency some proportion of the contributing atoms are Stark-shifted. This effect also leads to 'smearing' of the lineshape, as each atom's contribution is centred at a different frequency. The net result is a reduction in amplitude, but accompanied by an increase in width. This effect complicates the data analysis as it becomes difficult to use the measured shift from resonance to back-calculate the total power in the dipole trap.

The RHS of Figure 6.2 shows the relative increase in measured by a Stark shift for the data shown on the LHS of the same figure, as a function of the relative increase in power. Two datasets are shown. The first dataset, in the dark blue, shows the measured increase in Stark shift as a function of the increase in power. The increase in power is estimated by measuring the trapping power reflected back from the cavity (remembering to record the power at the lock point rather than the fringe minimum or maximum) and using this to estimate the cavity pumping rate. As expected, the shift increases linearly with increasing power.

The second dataset, in the light blue, shows the measured increase in Stark shift as a function of the increase in power, but now the increase in power is measured directly from the APD traces of the atom drops (when no atoms are present) — this is just the trap beam count rate measured at the APD during each drop experiment. As with the other dataset, the Stark shift does increase with increasing power, but there isn't a clear linear dependence. This could be due to fluctuating background counts from another source included in the trapping light count level, or that the filtering optics have a nonlinear attenuation with power.

The data shown in Figure 6.2 does indicate that the Dipole trap is present and having an effect. Up to now, however, it has not been possible to load the trap directly with atoms whilst they fall past. This is because, on the timescale of an atom falling through the cavity, the dipole force is conservative, so that there is no net energy loss for atoms crossing the trap. An atom travelling at 30 cm s^{-1} has a mean temperature of $\sim 1 \text{ mK}$, which is above the maximum trap depth we can achieve. A solution to this is to ensure that atoms reach the cavity with negligible velocities. To do this, we plan to introduce a soft-loading mechanism, which will controllably transport atoms between the MOT and the cavity.

6.1.2 Conveying atoms to the trap

We are currently adapting the experiment setup to introduce a second optical dipole trap which will act as an atom conveyor, allowing us to controllably transfer atoms between the MOT and the cavity. Such conveyors have already been successfully implemented in cavity QED experiments [28, 31].

The conveyor will be formed by two counter-propagating and overlapping beams, located vertically above and below the cavity (in the positions currently taken by the ‘top’ and ‘bottom’ beams used for pushing and transverse pumping). Two counter-propagating, mode-matched beams derived from the same source will form a standing wave as long as their path difference is less than their coherence length. Given this, a laser with linewidth, $\Delta\nu = 1$ GHz will have a coherence length of 300 mm, so that we could reasonably produce a standing wave over a range of several millimetres with free-running cavity-extended diode lasers. In addition, their maxima will constructively interfere such that the peak intensity will be $4\times$ the peak from one beam, having obvious benefits for optical trapping. If one beam is detuned by an amount δ_ν , or angular frequency $\delta = 2\pi\delta_\nu$, then the lattice will move at a velocity:

$$v = \frac{\delta}{2k} = \frac{1}{2}\delta_\nu\lambda$$

From this, two beams with $\lambda = 800$ nm and $\delta_\nu = 10$ MHz will result in a lattice velocity of 4 mm/ms. Assuming a MOT-cavity distance of 5.1 mm, it would take 1.3 ms for an atom in the lattice to move between them, assuming a constant transport velocity. A more accurate scenario, however, would not involve transporting at a constant velocity but, instead accelerating the atoms from $v \sim 0$ at the MOT and then decelerating them so that they returned to $v \sim 0$ for loading into the cavities.

During their transport, a major source of atom loss from the trap will be from scattering the trapping light. For the two 800 nm beams with a Rayleigh range of 10 mm and power of 250 mW, the achievable trap depth at the MOT is ~ 540 μ K, assuming the beam waist is located at the cavities. This is an adequate trapping potential, being $> 10\times$ the measured temperature of the MOT. In this trap, atoms will scatter photons at a rate of ≈ 140 photons/sec and, using the same assumptions as earlier, one can expect trapping lifetimes ~ 10 s. This indicates that travel times of a few milliseconds would not result in significant losses.

Atoms will, however, be lost from the trap if they are accelerated too quickly. We can quantify this affect by treating an ideal atom, with zero thermal energy located at the maximum potential in the trap. The atom’s potential, $V(y)$ is then given by:

$$V(y) = may - A_0 \sin(2ky)$$

Here, a is the atom's acceleration, m is the atom's mass, y is the distance and A_0 is set to be the half the maximum achievable trap depth. Taking $A_0 = 270 \mu\text{K}$ and conveyor beams of 800 nm, solving for $dV(x)/dx = 0$ gives $a_{\max} = 405 \text{ mm/ms}^2$. This is the maximum possible acceleration allowed, but assuming an acceleration and deceleration of $a = a_{\max}/10$, this gives us a total MOT-cavity journey time of 0.7 ms, well within the bounds set to limit scattering losses and, unsurprisingly, quicker than that calculated using a constant velocity.

The conveyor beams will be derived from a single 800 nm laser and their relative detunings and intensities will be controlled by passing both beams through AOMs. The outputs from the two AOMs will pass through two tapered amplifiers (Toptica TA800s) to increase the beam power to 500 mW. Both beams exiting the TAs will be separately coupled into fibres to achieve identical mode matching and deliver them to the chamber. The result, assuming 50% fibre coupling efficiency, will be 250 mW of power per beam, producing a maximum trap depth of $540 \mu\text{K}$. From considerations such as those above, the expected transfer time will be ~ 1 ms. The AOM driver sources will be those currently in use for the top and bottom beams. As discussed in Chapter 3, these are digitally synthesized and synchronised to one another, so there should be no drift in their relative frequency.

This set up is in the process of being built, but once completed should allow controllable loading of atoms into the intra-cavity dipole trap, allowing us to trap single atoms, or small clouds of atoms, inside the cavity mode. Trapping clouds of N atoms will allow us to further increase our cooperativity and there are two obvious bounds of interest; for $N > 13$, we would move into a regime where we might expect effects such as optical bistability; and for $N > 2800$, we would move into a regime of strong coupling. Alternatively we could explore the opportunity for coherent control.

6.2 Coherent control of atoms inside the cavity

One application of cavity QED is to use the atom-cavity interaction to coherently control the internal state of an atom. In this way, the cavity provides an interface which can be used to map information onto the atom and then store and retrieve it via emission of a photon into the cavity mode. Several experiments have demonstrated this effectively [76, 78]. A generic atomic system used in such processes is a three-level 'lambda' system, as shown in Figure 6.3.




Figure 6.3: The three-level ‘lambda’ system, consisting of levels $|1\rangle$, $|2\rangle$ and $|3\rangle$. The transitions are addressed by two coherent laser fields; Ω_c and Ω_p , with respective detunings $\Delta_p = \omega_{31} - \omega_p$ and $\Delta_c = \omega_{32} - \omega_c$.

The lambda system contains three atomic states; $|1\rangle$, $|2\rangle$ and $|3\rangle$, which can be addressed by two fields; Ω_p connecting the $|1\rangle \rightarrow |3\rangle$ states and Ω_c connecting the $|2\rangle \rightarrow |3\rangle$ states. In ^{87}Rb , the $|F=1\rangle$, $|F=2\rangle$ and $|F'=2\rangle$ hyperfine states form such a system.

Typically, when using such an atom in conjunction with a cavity, the atom-cavity interaction, with $\Omega = 2g\sqrt{n}$, can take the place of one of the addressing fields. Crucial then, to coherently controlling the atom is the ability to address only one of the transitions. We can already demonstrate this, via the state selectivity of our detection process.

6.2.1 Incoherent control: state selectivity of atom detection in the cavity

Thus far, we have demonstrated ‘incoherent’ control of the atomic state; atoms can be triggered into emitting a photon into the cavity mode by transiently exciting them into the $|F'=3\rangle$ state, producing a fluorescence signal. This emission, however, is probabilistic. In addition, using the repumping light we can optically pump atoms between the $|F=1\rangle$ and $|F'=2\rangle$ states and, by probing on a particular transition, demonstrate the state selectivity of our atom detection. Figure 6.4 shows the results of such an experiment. The cavity and pumping light are resonant with the $|F=2\rangle \rightarrow |F'=3\rangle$ transition. Atoms in the $|F=2\rangle$ state can therefore be detected through their resonant interaction with the pump light inside the cavity. Conversely, atoms in the $|F=1\rangle$ state will not interact with the pump and will therefore not be detected.

Atoms can be collected in the $|F=1\rangle$ state by switching off the repumping light during the optical molasses stage of the experiment. Entering the cavity in this state, they won’t interact with the pumping light and there will be no change in the light level. This can be seen from Figure 6.4, where the profile of the atom cloud would normally be evident from




Figure 6.4: The normalised reflection signal in lab units, averaged for ten shots. Initially, atoms entering the cavity are in the $|F = 1\rangle$ state, caused by switching off the repumping light during the optical molasses phase of the experiment. In this state they will not interact with the light pumping the cavity. At $t = 69\text{ms}$, indicated by the vertical dashed line, the repumping light is switched back on. This pumps the atoms into the $|F = 2\rangle$ state, where they can be detected through their interaction with the pumping light. This is evidenced by the sudden retrieval of the drop signal.

$t \simeq 60\text{ms}$ onwards. At $t = 69\text{ms}$, indicated by the dashed line, the repumping light is switched back on. This pumps atoms into the $|F' = 2\rangle$ state, from which they can decay to $|F = 2\rangle$ state and interact with the light pumping the cavity. This process is evidenced by the sudden reappearance of the drop signal for $t > 69\text{ms}$. From this, it is clear that the cavity detection is state selective. This is, however, an incoherent process, relying on the probabilistic nature of the interaction between atoms and light. The next stage for the experiment would be to coherently control the atom, an example of which is electromagnetically induced transparency (EIT), see for example [117].

6.2.2 Coherent control: EIT

To describe EIT we can consider the dressed states of the lambda system already shown in Figure 6.3. Typically, one refers to the Ω_p field as the ‘probe’ laser and the Ω_c field as the ‘coupling’ laser, where each laser has a detuning, $\Delta_p = \omega_{31} - \omega_p$ and $\Delta_c = \omega_{32} - \omega_c$. Here, the transition $|1\rangle \rightarrow |2\rangle$ is forbidden by dipole selection rules. Diagonalising the interaction Hamiltonian describing the dipole interaction between the fields and atomic states gives three eigenstates [117] for the case of a ‘two-photon resonance’, when $\Delta_p = \Delta_c$:

$$\begin{aligned}
|a^+\rangle &= \sin \theta \sin \phi |1\rangle + \cos \phi |3\rangle + \cos \theta \sin \phi |2\rangle \\
|a^-\rangle &= \sin \theta \cos \phi |1\rangle - \sin \phi |3\rangle + \cos \theta \cos \phi |2\rangle \\
|a^0\rangle &= \cos \theta |1\rangle - \sin \theta |2\rangle
\end{aligned}$$

These eigenstates are given in terms of the single photon detuning, $\Delta = \Delta_p$ and the mixing angles θ and ϕ given by:

$$\tan \theta = \frac{\Omega_p}{\Omega_c} \quad \text{and} \quad \tan 2\phi = \frac{\sqrt{\Omega_p^2 + \Omega_c^2}}{\Delta}$$

It is clear that one eigenstate, $|a^0\rangle$, has no contribution from the bare atomic state, $|3\rangle$. Atoms in this state are effectively decoupled from the light fields, as there is no possibility for excitation to $|3\rangle$ and subsequent spontaneous emission. For this reason, $|a^0\rangle$ is known as a ‘dark state’.

Treating the case of a weak probe field, $\Omega_p \ll \Omega_c$, in this limit, $\sin \theta \rightarrow 0$ and $\cos \theta \rightarrow 1$, so that the application of these fields transforms an atom in the ground state $|1\rangle \rightarrow |a^0\rangle$. Once in the dark state, the atom can be rendered transparent to the probe field, as no absorption can take place. This leads to a ‘transparency window’ in the absorption profile, which is visible whilst the atomic dephasing is small.³ An additional technical limitation is set by the necessity to minimise fluctuations in amplitude and phase of the probe and coupling fields, which can reduce the visibility of EIT by causing additional decoherence.

We have observed EIT in a room temperature vapour cell of Rubidium, using phase-locked repumping and trapping light, as discussed in Chapter 3. The trapping light acts as the coupling field, tuned to the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition, whilst the repumping light is the probe field, tuned to the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. The experiment uses a modified saturated absorption spectroscopy set up, with the beams co-propagating through the vapour cell and the transmission of the probe field measured at a photodiode. The transmission profile of the probe beam as Δ_p is tuned across the ^{87}Rb $|F = 0\rangle \rightarrow |F = 0, 1, 2\rangle$ manifold is shown in the LHS of Figure 6.5. The darkest green shows the Doppler-broadened absorption profile. The medium-toned green shows the saturated absorption spectrum of the same transitions, where now the probe beam is retro-reflected through the cell. The lightest green shows the same saturated absorption spectrum, but now with the coupling laser present. The narrow EIT feature is clearly visible at zero detuning. A ‘zoomed in’ version of this feature is shown

³This is satisfied when the lifetime of $|2\rangle$ is substantially longer than $|3\rangle$, so that $\gamma_{21} \ll \gamma_{31}$ from $|2\rangle$.




Figure 6.5: Electromagnetically induced transparency in an Rb vapour cell. LHS: The transmission profile of the probe beam is shown as Δ_p is tuned across the ^{87}Rb $|F=0\rangle \rightarrow |F=0, 1, 2\rangle$ manifold. The darkest green shows the Doppler-broadened absorption profile. The medium-toned green shows the saturated absorption spectrum of the same transitions, where now the probe beam is retro-reflected through the cell. The lightest green shows the same saturated absorption spectrum, but now with the co-propagating coupling laser present. The narrow EIT feature is clearly visible at zero detuning. RHS: A ‘zoomed in’ version of the EIT feature. Shown alongside the scan is the fitted Lorentzian lineshape, which estimates the HWHM to be 357 ± 4 kHz.

on the RHS of Figure 6.5, where it is fitted to a Lorentzian lineshape. This gives a sub-natural HWHM of 357 ± 4 kHz. This linewidth is likely to be limited by the transit time for atoms travelling across the beams; the expected broadening for a room temperature vapour of rubidium atoms interacting with a 1 mm wide beam is 300 kHz.

This demonstrates the ability to coherently control a room temperature sample of Rb atoms. The next stage will be to demonstrate equivalent coherence for atoms inside the cavity mode. The cavity mode will be driven by the probe laser, with the coupling laser driving the atoms transversely from the side. The apparatus has been set up and used with atoms falling through the cavity mode, but there were no conclusive signs of EIT, attributed to the short transit time of atoms through the cavity mode and low atom signal level at the time. In the future, this can be attempted with a single trapped atom in the cavity or being conveyed.

Successfully demonstrating EIT inside our cavity will put us a step closer to using our cavity for quantum information processing (QIP) applications. Firstly, EIT and the closely related technique of stimulated Raman scattering by adiabatic passage (STIRAP) [118] can be used to produce single photons, as discussed in detail in [91]. In this application, the atom-cavity interaction takes the place of the coupling field, so that $\Omega_c = 2\sqrt{n}g$. With STIRAP, the procedure involves adiabatically transferring the atomic population from $|1\rangle \rightarrow |2\rangle$ by varying the ratio of $|\Omega_c|^2/|\Omega_P|^2$ with time. At the end of one STIRAP cycle the cavity photon number will have increased by one. One cycle of EIT, however, involves moving from $|1\rangle \rightarrow |2\rangle \rightarrow |1\rangle$, so that the cavity photon number is only transiently increased before the atom re-absorbs the photon towards the end of the cycle. This transient increase can, however, correspond to a photon being emitted from the cavity if κ is faster than the timescale of the EIT pulse.

A second application of EIT is to use it to slow down or stop a photon inside an atomic vapour [119], which could subsequently be retrieved — offering the possibility of a storage mechanism. All of these experiments are just one area for future exploration. The following chapter concludes this thesis and presents some further ideas for the future of our experiment.

Chapter 7

Conclusions and outlook

7.1 Conclusions

This thesis has presented experiments conducted using an optical microcavity designed and built within the Centre for Cold Matter. The cavity's design, being formed between a single mode fibre and a micromirror etched from a silicon wafer, is such that there is direct optical access to the internal mode and aspects of the device are scalable and could be integrated within other atom chip experiments. These factors make it well suited to applications within quantum information processing and atom detection.

In Chapter 2, the theory underlying the optical cavity and the interaction between atoms and the cavity field was discussed. From this, a theoretical description of the reflection and fluorescence signals arising from our cavity was presented. This could be used to understand later experiments.

Chapter 3 contained a detailed description of the apparatus used to perform our cavity experiments. Nearly every aspect of the experiment was rebuilt during the course of this thesis, the aim being to increase the overall stability and increase the range of experiments which could be performed. Having re-built the experiment's laser system, an optical dipole trap is now ready to trap atoms inside the cavity and the technology is available for coherent control. In addition, the optical system used to collect atom signals has been improved to reduce scatter and thus increase the signal-to-noise.

The first part of Chapter 4 contained a comprehensive characterisation of the reflection and fluorescence signals with respect to pump beam detuning, intensity and polarisation. Assuming changes in the reflection signal are used to detect atoms, these experiments indicated that the largest signals would be achieved using a resonant beam with an input pumping rate of up

to 5 MCounts/sec. In fluorescence it was found that the largest signals involved a resonant, linearly polarised beam, with the polarisation axis transverse to the cavity axis and pump beam powers of several hundred microwatts.

The second part of Chapter 4 tested the functionality of our system as a chip-based atom detector. It was shown that the fluorescence signal could be used to detect an effective single atom in our cavity (i.e. $N_{\text{eff}} = 1$) with a signal-to-noise ratio of 1 in $4\mu\text{s}$ and achieving a 99% detection fidelity in $15.6\mu\text{s}$. Using the reflection signal, signal-to-noise ratios of 1 can be achieved in $37\mu\text{s}$. We then explored the prospect for detecting individual atoms falling through the cavity and trapped inside it. It appeared unlikely that we would be able to detect individual atoms as they fell through the cavity, primarily due to their lack of localisation in the mode. In contrast, it was found that a single trapped atom could be confidently detected using both the reflection and fluorescence methods. The reflection signal would likely be the method of choice, as the atom can be detected whilst also retained in the trap. In contrast, fluorescence detection would lead to heating out of the trap in $\sim 50\mu\text{s}$.

Chapter 5 began by presenting a procedure for increasing the number of atoms inside the cavity. This method uses a resonant push beam to push atoms from the MOT to the cavity, reducing the atom travel time and thereby increasing the atom density at the cavity. Using this method we are able to increase N_{eff} from ~ 1 to ~ 7 , with a largest recorded $N_{\text{eff}} = 11$. This increases the collective cooperativity by an order of magnitude. Using the increase in N_{eff} , enhanced Purcell broadening is demonstrated, alongside saturation of the reflection and fluorescence signal. In the latter case, the fluorescence signal was expected to reduce beyond $N_{\text{eff}} \sim 2$, but later simulations indicated that the lack of reduction might be caused by the pump beam being detuned.

The latter part of Chapter 5 explored the dispersive interaction between the atom and the cavity. The reflection spectrum of the cavity was measured for different values of N_{eff} , with clear evidence of dispersion apparent when $N_{\text{eff}} \sim 1$. We were able to exploit the high collection efficiency to make a direct measurement of the equivalent spectrum in fluorescence, previously only inferred in the strong coupling regime. In contrast with the strong coupling case, no obvious splitting is apparent in the signal when $N_{\text{eff}} = 1$, demonstrating a qualitative difference between the types of signals.

Finally, Chapter 6 focussed on the immediate future for our experiment; optical dipole trapping an atom inside the cavity and coherent control of an atom using electromagnetically induced transparency. Progress towards the former goal was reported, including evidence of the

Stark shift caused by the presence of the optical dipole trap inside the cavity. An experiment showing electromagnetically induced transparency (EIT) of rubidium inside a vapour cell was also presented, demonstrating that we have the technology to reproduce the experiment inside the cavity. Such a technique could also be used to deterministically produce single photons.

7.2 Outlook

There is an exciting future for our optical microcavity experiment. In achieving the goals outlined in Chapter 6, a single atom will have been successfully inside the cavity and EIT/STIRAP used to produce single photons, all with a system outside of the strong coupling regime. As a next step, it will be important to demonstrate the statistical characteristics of our single-photon source; showing photon antibunching and sub-Poissonian statistics. A second APD has already been installed, thus enabling Hanbury-Brown-Twiss coincidence-counting experiments, see for example [120]. As an intermediate step, it would be interesting to perform such an experiment on the existing photon streams produced by falling atoms.

There are also several areas of future development. Firstly, a new cavity with increased finesse could be built. Methods for reducing the mirror roughness via new etching methods have already been discussed in [97]. In this way, we can hope to trap an atom and combine with cavity cooling techniques to increase storage times for atoms inside the cavity.

Secondly, the scalability of our device could be exploited to build a multi-cavity experiment. At present there are two cavities currently installed in our chamber, although only one can be used as both cavities are tuned using the same piezo stack so that they cannot simultaneously be resonant. Modifying the design of the cavity would overcome this limitation. With two cavities, we could explore atom-light, atom-atom or light-light entanglement by, for example, using the interaction in both cavities and mixing their outputs on a beam splitter.

Thirdly and finally, we could look to further integrate our cavity within atom chip experiments already in use in CCM. Recent successes within the group include atom detection with optical waveguides integrated on an atom chip and small atom number etched pyramid MOTs. One could combine a waveguide and silicon mirror to form an all-chip optical cavity, or use a pyramid MOT as a source of cold atoms, both moving a step closer towards the concept of a ‘lab on a chip’.

Bibliography

- [1] E. T. Jaynes and F. W. Cummings. Comparison of Quantum and Semiclassical Radiation Theories with Application to the Beam Maser. In *Proc. IEEE*, pages 51–89, 1963.
- [2] Christopher J. Foot. *Atomic Physics*. Oxford University Press, 2005.
- [3] P. Meystre and M. Sargent III. *Elements of Quantum Optics*. Springer, third edition edition, 1999.
- [4] E. M. Purcell. Spontaneous Emission Probabilities at Radio Frequencies. *Physical Review*, 69:681, 1946.
- [5] P. R. Berman, editor. *Cavity Quantum Electrodynamics*. Academic Press, 1993.
- [6] S. Kuhr, S. Gleyzes, C. Guerlin, J. Bernu, U. B. Hoff, S. Deléglise, S. Osnaghi, M. Brune, J.M. Raimond, S. Haroche, E. Jacques, P. Bosland, and B. Visentin. Ultrahigh Finesse Fabry–Perot Superconducting Resonator. *Applied Physics Letters*, 90:164101, 2007.
- [7] M. A. Armen. *Bifurcations in Single Atom Cavity QED*. PhD thesis, California Institute of Technology, 2009.
- [8] J. M. Raimond, M. Brune, and S. Haroche. Manipulating Quantum Entanglement with Atoms and Photons in a Cavity. *Review of Modern Physics*, 73:565582, 2001.
- [9] A. B. Mundt, A. Kreuter, C. Becher, D. Leibfried, J. Eschner, F. Schmidt Kaler, and R. Blatt. Coupling a Single Atomic Quantum Bit to a High Finesse Optical Cavity. *Physical Review Letters*, 89:103001, 2002.
- [10] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. Stable Long–Term Coupling of a Single Ion to a Cavity Mode. *Journal of Modern Optics*, 54:1607–1617, 2007.
- [11] H. Mabuchi and A. C. Doherty. Cavity Quantum Electrodynamics: Coherence in Context. *Science*, 298:1372–1377, 2002.

- [12] J. P. Reithmaier, G. Seogonk, A. Löffler, C. Hofmann, S. Kuhn, S. Reitzenstein, L. V. Keldysh, V. D. Kulakovskii, T. L. Reinecke, and A. Forchel. Strong Coupling in a Single Quantum Dot Semiconductor Microcavity System. *Nature*, 432:197–200, 2004.
- [13] T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H. M. Gibbs, G. Rupper, C. Ell, O. B. Shchekin, and D. G. Deppe. Vacuum Rabi Splitting with a Single Quantum Dot in a Photonic Crystal Nanocavity. *Nature*, 432:200–203, 2004.
- [14] M. Brune, J. M. Raimond, P. Goy, L. Davidovich, and S. Haroche. Realization of a Two-Photon Maser Oscillator. *Physical Review Letters*, 59:1899 – 1903, 1987.
- [15] G. Rempe, H. Walther, and N. Klein. Observation of Quantum Collapse and Revival in a One Atom Maser. *Physical Review Letters*, 58:353–356, 1987.
- [16] F. Bernardot, M. Brune, J. M. Raimond, and S. Haroche. Vacuum Rabi Splitting Observed on a Microscopic Atomic Sample in a Microwave Cavity. *Europhysics Letters*, 17:33 – 38, 1992.
- [17] M. Brune, F. Schmidt-Kaler, A. Maali, J. Dreyer, E. Hagley, J. M. Raimond, and S. Haroche. Quantum Rabi Oscillation: a Direct Test of Field Quantization in a Cavity. *Physical Review Letters*, 76:1800 – 1803, 1996.
- [18] R. J. Thompson, G. Rempe, and H. J. Kimble. Observation of Normal Mode Splitting for an Atom in an Optical Cavity. *Physical Review Letters*, 68:1132–1135, 1992.
- [19] J. J. Childs, K. An, M. S. Otteson, R. R. Dasari, and M. S. Feld. Normal Mode Lineshapes for Atoms in Standing-Wave Optical Resonators. *Physical Review Letters*, 77:2901–2904, 1996.
- [20] William D. Phillips. Nobel Lecture: Laser Cooling and Trapping of Neutral Atoms. *Review of Modern Physics*, 70:721–741, 1998.
- [21] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59:2631–2634, 1987.
- [22] H. Mabuchi, Q. A. Turchette, M. S. Chapman, and H. J. Kimble. Real-Time Detection of Individual Atoms Falling through a High-Finesse Optical Cavity. *Optics Letters*, 21:1393–1395, 1996.

- [23] P. Münstermann, T. Fischer, P. W. H. Pinkse, and G. Rempe. Single Slow Atoms from an Atomic Fountain Observed in a High Finesse Optical Cavity. *Optics Communications*, 159:63–67, 1999.
- [24] P. Münstermann, T. Fischer, P. Maunz, P. W. H. Pinkse, and G. Rempe. Dynamics of Single Atom Motion Observed in a High Finesse Cavity. *Physical Review Letters*, 82:3791–3794, 1999.
- [25] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. Optical Dipole Traps for Neutral Atoms. *arXiv:physics/9902072*, 1999.
- [26] J. Ye, D. W. Vernooy, and H. J. Kimble. Trapping of Single Atoms in Cavity QED. *Physical Review Letters*, 83:4987–4990, 1999.
- [27] P. W. H. Pinkse, T. Fischer, P. Maunz, T. Puppe, and G. Rempe. How to Catch an Atom with a Single Photon. *Journal of Modern Optics*, 47:2769–2787, 2000.
- [28] J. A. Sauer, K. M. Fortier, M. S. Chang, C. D. Hamley, and M. S. Chapman. Cavity QED with Optically Transported Atoms. *Physical Review A*, 69:051804, 2004.
- [29] K. M. Fortier, S. Y. Kim, M. J. Gibbons, P. Ahmadi, and M. S. Chapman. Deterministic Loading of Individual Atoms to a High Finesse Optical Cavity. *Physical Review Letters*, 98:233601, 2007.
- [30] S. Nuß man, K. Murr, M. Hijlkema, B. Weber, A. Kuhn, and G. Rempe. Vacuum–Stimulated Cooling of a Single Atom in Three Dimensions. *Nature Physics*, 1:122–126, 2005.
- [31] M. Khudaverdyan, W. Alt, T. Kampschulte, S. Reick, A. Thobe, A. Widera, and D. Meschede. Quantum Jumps and Spin Dynamics of Interacting Atoms in a Strongly–Coupled Atom–Cavity System. *Physical Review Letters*, 103:123006, 2009.
- [32] M. Keller, B. Lange, K. Hayasaka, W. Lange, and H. Walther. Continuous Generation of Single Photons with Controlled Waveform in an Ion–Trap Cavity System. *Nature*, 431:1075–1078, 2004.
- [33] P. Maunz, T. Puppe, I. Schuster, N. Syassen, P. W. H. Pinkse, and G. Rempe. Cavity cooling of a single atom. *Nature*, 428:50 – 52, 2004.

- [34] P. Horak, G. Hechenblaikner, K. M. Gheri, H. Stecher, and H. Ritsch. Cavity Induced Atom Cooling in the Strong Coupling Regime. *Physical Review Letters*, 79:4974–4977, 1997.
- [35] A. D. Boozer, A. Boca, R. Miller, T. E. Northup, and H. J. Kimble. Cooling to the Ground State of Axial Motion for One Atom Strongly Coupled to an Optical Cavity. *Physical Review Letters*, 97:083602, 2006.
- [36] S. Gupta, K. L. Moore, K. W. Murch, and D. M. Stamper-Kurn. Cavity Nonlinear Optics at Low Photon Numbers from Collective Atomic Motion. *Physical Review Letters*, 99:213601, 2007.
- [37] J. Simon, H. Tanji, J. K. Thompson, and V. Vuletić. Interfacing Collective Atomic Excitations and Single Photons. *Physical Review Letters*, 98:183601, 2007.
- [38] Y. Colombe, T. Steinmetz, G. Dubois, F. Linke, D. Hunger, and J. Reichel. Strong Atom–Field Coupling for Bose Einstein Condensates in an Optical Cavity on a Chip. *Nature*, 450:272–276, 2007.
- [39] F. Brennecke, T. Donner, S. Ritter, T. Bourdel, M. Köhl, and T. Esslinger. Cavity QED with a Bose Einstein Condensate. *Nature*, 450:268 – 271, 2007.
- [40] F. W. Cummings and M. Tavis. Exact Solution for an N Molecule Radiation Field Hamiltonian. *Physical Review*, 170(2), 1968.
- [41] H. W. Chan, A. T. Black, and V. Vuletić. Observation of Collective Emission–Induced Cooling of Atoms in an Optical Cavity. *Physical Review Letters*, 90:063003, 2002.
- [42] A. T. Black, H. W. Chan, and V. Vuletić. Observation of Collective Friction Forces due to Spatial Self–Organization of Atoms: from Rayleigh to Bragg Scattering. *Physical Review Letters*, 91:203001, 2003.
- [43] M. Trupke, J. Goldwin, B. Darquié, G. Dutier, S. Eriksson, J. Ashmore, and E. A. Hinds. Atom Detection and Photon Production in a Scalable, Open, Optical Microcavity. *Physical Review Letters*, 99:063601, 2007.
- [44] A. Haase, B. Hessmo, and J. Schmiedmayer. Detecting Magnetically Guided Atoms with an Optical Cavity. *Optics Letters*, 37:268–270, 2005.

- [45] I. Teper, Y-J Lin, and V. Vuletić. Resonator–Aided Single Atom Detection on a Micro-fabricated Chip. *Physical Review Letters*, 97:023002, 2006.
- [46] E. A. Hinds and I. G. Hughes. Magnetic Atom–Optics: Mirrors, Guides, Traps, and Chips for Atoms. *Journal of Physics D: Applied Physics*, 32:119–146, 1999.
- [47] R. Folman, P. Krüger, D. Cassettari, B. Hessmo, T. Maier, and J. Schmiedmayer. Controlling Cold Atoms using Nanofabricated Surfaces: Atom Chips. *Physical Review Letters*, 84:4749–4752, 2000.
- [48] T. P. Purdy and D. M. Stamper Kurn. Integrating Cavity Quantum Electrodynamics and Ultracold Atom Chips with On–Chip Dielectric Mirrors and Temperature Stabilization. *Applied Physics B*, 90:401405, 2008.
- [49] T. Aoki, B. Dayan, E. Wilcut, W. P. Bowen, A. S. Parkins, T. J. Kippenberg, K. J. Vahala, and H. J. Kimble. Observation of Strong Coupling between One Atom and a Monolithic Microresonator. *Nature*, 443:671 – 674, 2006.
- [50] S. Gleyzes, A. El Amili, R. A. Cornelussen, P. Lalanne, C. I. Westbrook, A. Aspect, J. Estève, G. Moreau, A. Martinez, X. Lafosse, L. Ferlazzo, J. C. Harmand, D. Mailly, and A. Ramdane. Towards a Monolithic Optical Cavity for Atom Detection and Manipulation. *The European Physical Journal D*, 53:107111, 2009.
- [51] S. van Enk and H. J. Kimble. Strongly Focused Light Beams Interacting with Single Atoms in Free Space. *Physics Review A*, 63:023809, 2001.
- [52] D. J. Wineland, W. M. Itano, and J. C. Bergquist. Absorption Spectroscopy at the Limit: Detection of a Single Atom. *Optics Letters*, 12:389–391, 1987.
- [53] P. A. Quinto-Su, M. Tschernneck, M. Holmes, and N. P. Bigelow. On–Chip Optical Detection of Laser Cooled Atoms. *Optics Express*, 12:5098–5013, 2004.
- [54] M. Kohnen, M. Succo, P. G. Petrov, R. A. Nyman, M. Trupke, and E. A. Hinds. An Integrated Atom–Photon Junction. *arXiv:0912.4460v1*, 2009.
- [55] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum Dynamics of Single Trapped Ions. *Review of Modern Physics*, 75:281–324, 2003.
- [56] Z. Hu and H. J. Kimble. Observation of a Single Atom in a Magneto–Optical Trap. *Optics Letters*, 19:1888–1890, 1994.

- [57] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier. Sub-Poissonian Loading of Single Atoms in a Microscopic Dipole Trap. *Nature*, 411:1024–1027, 2001.
- [58] T. Bondo, M. Hennrich, T. Legero, G. Rempe, and A. Kuhn. Time-Resolved and State-Selective Detection of Single Freely Falling Atoms. *Optics Communications*, 264:271 – 277, 2006.
- [59] A. Takamizawa, T. Steinmetz, R. Delhuille, T. W. Hänsch, and J. Reichel. Miniature Fluorescence Detector for Single Atom Observation on a Microchip. *Optics Express*, 14:10976–10983, 2006.
- [60] A. Ottl, S. Ritter, M. Köhl, and T. Esslinger. Correlations and Counting Statistics of an Atom Laser. *Physical Review Letters*, 95:090404, 2005.
- [61] C. J. Hood, T. W. Lynn, A. C. Doherty, A. S. Parkins, and H. J. Kimble. The Atom-Cavity Microscope: Single Atoms Bound in Orbit by Single Photons. *Science*, 287:1447–1453, 2000.
- [62] M. Hijlkema, B. Weber, H. P. Specht, S. C. Webster, A. Kuhn, and G. Rempe. A Single Photon Server with just One Atom. *Nature Physics*, 3:253–255, 2007.
- [63] J. McKeever, A. Boca, A. D. Boozer, R. Miller, J. R. Buck, A. Kuzmich, and H. J. Kimble. Deterministic Generation of Single Photons from One Atom Trapped in a Cavity. *Science*, 303:1992–1994, 2004.
- [64] M. L. Terraciano, R. Olson Knell, D. G. Norris, J. Jing, A. Fernández, and L. A. Orozco. Photon Burst Detection of Single Atoms in an Optical Cavity. *Nature Physics*, 5:480–484, 2009.
- [65] P. Horak, B. G. Klappauf, A. Haase, R. Folman, J. Schmiedmayer, P. Domokos, and E. A. Hinds. Possibility of Single Atom Detection on a Chip. *Physical Review A*, 67:043806, 2003.
- [66] T. Steinmetz, A. Balocchi, Y. Colombe, D. Hunger, T. W. Hnsch, R. J. Warburton, and J. Reichel. A Stable Fibre-based Fabry-Perot Cavity. *Applied Physics Letters*, 89:111110, 2006.
- [67] et. al. J. Goldwin. To be published. 2010.

- [68] T. Puppe, I. Schuster, P. Maunz, K. Murr, P. W. H. Pinkse, and G. Rempe. Light Force Fluctuations in a Strongly Coupled Atom–Cavity System. *Journal of Modern Optics*, 54:1927–1937, 2007.
- [69] P. Goy, M. Gross J. M. Raimond, and S. Haroche. Observation of Cavity–Enhanced Single–Atom Spontaneous Emission. *Physical Review Letters*, 50:1903, 1983.
- [70] D. Kleppner. Inhibited Spontaneous Emission. *Physical Review Letters*, 47:233–236, 1981.
- [71] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum State Transfer and Entanglement Distribution among Distant Nodes in a Quantum Network. *Physical Review Letters*, 78:3221–3224, 1997.
- [72] A. Kuhn, M. Hennrich, and G. Rempe. A Deterministic Single Photon Source for Distributed Quantum Networking. *Physical Review Letters*, 89:067901, 2002.
- [73] T. Legero, T. Wilk, M. Hennrich, G. Rempe, and Axel Kuhn. Quantum Beat of Two Single Photons. *Physical Review Letters*, 93:070503, 2004.
- [74] T. Wilk, S. C. Webster, H. P. Specht, G. Rempe, and A. Kuhn. Polarization–Controlled Single Photons. *Physical Review Letters*, 98:063601, 2007.
- [75] H. P. Specht, J. Bochmann, M. Mücke, B. Weber, E. Figueroa, D. L. Moehring, and G. Rempe. Phase Shaping of Single–Photon wavepackets. *Nature Photonics*, 3:469–472, 2009.
- [76] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe. Single Atom–Single Photon Quantum Interface. *Science*, 317:488 –490, 2007.
- [77] B. Weber, H. P. Specht, T. Müller, J. Bochmann, M. Mücke, D. L. Moehring, and G. Rempe. Photon–Photon Entanglement with a Single Trapped Atom. *Physical Review Letters*, 102:030501, 2009.
- [78] A. D. Boozer, A. Boca, R. Miller, T. E. Northup, and H. J. Kimble. Reversible State Transfer between Light and a Single Trapped Atom. *Physical Review Letters*, 98:193601, 2007.
- [79] E. Hagley, X. Maître, G. Nogues, C. Wunderlich, M. Brune, J. M. Raimond, and S. Haroche. Generation of Einstein–Podolsky–Rosen Pairs of Atoms. *Physical Review Letters*, 79:1–5, 1997.

- [80] B.T.H. Varcoe, S.Brattke, M.Weidinger, and H.Walther. Preparing Pure Photon Number States of the Radiation Field. *Nature*, 403:743–746, 2000.
- [81] J. Simon, H. Tanji, S. Ghosh, and V. Vuletić. Single-Photon Bus Connecting Spin-Wave Quantum Memories. *Nature Physics*, 3:765–769, 2007.
- [82] J. K. Thompson, J. Simon, H. Loh, and V. Vuletić. A High Brightness Source of Narrowband, Identical Photon Pairs. *Science*, 303:74–77, 2006.
- [83] H. Tanji, J. Simon, S. Ghosh, B. Bloom, and V. Vuletić. Heralded Atomic Ensemble Quantum Memory for Photon Polarization States. *Physica Scripta*, 135:014010, 2009.
- [84] G. Rempe, R. J. Thompson, R. J. Brecha, W. D. Lee, and H. J. Kimble. Optical Bistability and Photon Statistics in Cavity Quantum Electrodynamics. *Physical Review Letters*, 67:1727–1730, 1991.
- [85] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble. Measurements of Conditional Phase Shifts for Quantum Logic. *Physical Review Letters*, 75:4710–4713, 1995.
- [86] M. Trupke, E. A. Hinds, S. Eriksson, E. A. Curtis, Z. Muktadir, E. Kukharenska, and M. Kraft. Microfabricated High Finesse Optical Cavity with Open Access and Small Volume. *Applied Physics Letters*, 87:211106, 2005.
- [87] S. Pollock, J. P. Cotter, A. Laliotis, and E. A. Hinds. Integrated Magneto-Optical Traps on a Chip using Silicon Pyramid Structures. *Optics Express*, 17:14109–14114, 2009.
- [88] Eugene Hecht. *Optics*. Addison Wesley, 2002.
- [89] A. E. Siegman. *Lasers*. University Science Books, 1986.
- [90] Howard Carmichael. *An Open Systems Approach to Quantum Optics*. Springer-Verlag, 1993.
- [91] A. Kuhn and G. Rempe. *Optical Cavity QED: Fundamentals and Application as a Single Photon Light Source*. Societa Italiana di Fisica.
- [92] D. J. Heinzen, J. J. Childs, J. E. Thomas, and M. S. Feld. Enhanced and Inhibited Visible Spontaneous Emission by Atoms in a Confocal Resonator. *Physical Review Letters*, 58:1320–1323, 1987.

- [93] D. Kleppner. Inhibited Spontaneous Emission. *Physical Review Letters*, 47:233, 1981.
- [94] G. Gabrielse and H. Dehmelt. Observation of Inhibited Spontaneous Emission. *Physical Review Letters*, 55:67, 1985.
- [95] R. Hulet, E. Hilfer, and D. Kleppner. Inhibited Spontaneous Emission by a Rydberg Atom. *Physical Review Letters*, 55:2137, 1985.
- [96] K. R. Brown, D. M. Stamper Kurn, K. B. Whaley, S. Leslie, and N. Shenvi. Transmission Spectrum of an Optical Cavity containing N Atoms. *Physical Review A*, 69:034805, 2004.
- [97] M. Trupke. *Microcavities for Atom Chips*. PhD thesis, Imperial College London, 2007.
- [98] J. Gripp, S. L. Mielke, L. A. Orozco, and H. J. Carmichael. Anharmonicity of the Vacuum Rabi Peaks in a Many-Atom System. *Physical Review A*, 54:R3746–R3749, 1996.
- [99] H. J. Metcalf and P. van der Straten. *Laser cooling and trapping*. Springer, 1999.
- [100] R. H. Dicke. Coherence in Spontaneous Emission Processes. *Physical Review*, 93:99, 1954.
- [101] D. A. Steck. Rubidium 87 D Line Data. <http://george.ph.utexas.edu/dsteck/alkalidata/>, 2008.
- [102] J. Reichel, W. Hänsel, and T. W. Hänsch. Atomic Micromanipulation with Magnetic Surface Traps. *Physical Review Letters*, 83:33983401, 1999.
- [103] A. S. Arnold, J. S. Wilson, and M. G. Boshier. A Simple Extended-Cavity Diode Laser. *Review of Scientific Instruments*, 69:1236, 1998.
- [104] C. J. Hawthorn, K. P. Weber, and R. E. Scholten. Littrow Configuration Tunable External Cavity Diode Laser with Fixed Direction Output Beam. *Review of Scientific Instruments*, 72:4477, 2001.
- [105] K. G. Libbrecht and J. L. Hall. A Low-Noise High-Speed Diode Laser Current Driver. *Review of Scientific Instruments*, 64:2133, 1993.
- [106] C. Wieman and T. W. Hänsch. Doppler-Free Laser Polarization Spectroscopy. *Physical Review Letters*, 36(20):1170–1173, March 1976.

- [107] R. E. Teets, F. V. Kowalski, W. T. Hill, N. Carlson, and T. W. Hänsch. Advances in Laser Spectroscopy I. *Proceedings of the Society of Photo-Optical Instrumental Engineering*, 113:80, 1977.
- [108] C. P. Pearman, C. S. Adams, S. G. Cox, P. F. Griffin, D. A. Smith, and I. G. Hughes. Polarization Spectroscopy of a Closed Atomic Transition: Applications to Laser Frequency Locking. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 35:5141–5151, 2002.
- [109] Paul F. Griffin. *Laser Cooling and Loading of Rb in a Large Period, Quasi-Electrostatic, Optical Lattice*. PhD thesis, University of Durham, 2005.
- [110] G. Ritt, G. Cennini, C. Geckeler, and M. Weitz. Laser Frequency Offset Locking using a Side of Filter Technique. *Applied Physics B: Lasers and Optics*, 79:363–365, 2004.
- [111] K. L. Corwin, Z. Lu, C. F. Hand, R. J. Epstein, and Carl E. Wieman. Frequency-Stabilized Diode Laser with the Zeeman Shift in an Atomic Vapour. *Applied Optics*, 37(15):3295–3298, 1998.
- [112] P. F. Herskind, A. Dantan, J. P. Marler, M. Albert, and M. Drewsen. Realization of Collective Strong Coupling with Ion Coulomb Crystals in an Optical Cavity. *Nature Physics*, 5:494–498, 2009.
- [113] M. L. Terraciano, R. Olson Knell, D. L. Freimund, L. A. Orozco, J. P. Clemens, and P. R. Rice. Enhanced Spontaneous Emission into the Mode of a Cavity QED System. *Optics Letters*, 32:982–984, 2007.
- [114] T. Puppe, I. Schuster, A. Grothe, A. Kubanek, K. Murr, P.W.H. Pinkse, and G. Rempe. Trapping and Observing Single Atoms in a Blue-Detuned Intracavity Dipole Trap. *Physical Review Letters*, 99:013002, 2007.
- [115] J. Dalibard and C. Cohen-Tannoudji. Dressed-Atom Approach to Atomic Motion in Laser Light: the Dipole Force Revisited. *Journal of the Optical Society of America B*, 2:1707–1720, 1985.
- [116] J. McKeever, J. R. Buck, A. D. Boozer, A. Kuzmich, H. C. Nagerl, D. M. Stamper-Kurn, and H. J. Kimble. State Insensitive Cooling and Trapping of Single Atoms in an Optical Cavity. *Physical Review Letters*, 90:133602, 2003.

- [117] M. Fleischhauer, A. Imamoglu, and J. P. Marangos. Electromagnetically Induced Transparency: Optics in Coherent Media. *Reviews of Modern Physics*, 77, 2005.
- [118] K. Bergmann, H. Theuer, and B. W. Shore. Coherent Population Transfer among Quantum States of Atoms and Molecules. *Review of Modern Physics*, 70:1003 – 1026, 1998.
- [119] M. D. Eisaman, A. André, F. Massou, M. Fleischhauer, A. S. Zibrov, and M. D. Lukin. Electromagnetically Induced Transparency with Tunable Single Photon Pulses. *Nature*, 438:837–841, 2005.
- [120] C. C. Gerry and P. L. Knight. *Introductory Quantum Optics*. Cambridge University Press, 2005.

Appendix A

The fluorescence signal and increasing drive beam

Fluorescence experiments, where the atom-cavity system is driven by a transverse pump beam from the side, are generally more complicated to describe theoretically than their reflection counterparts. This is because we need to take into account both the spatial variation of their coupling to the cavity, g_j and the spatial variation of the Rabi frequency from the pump beam, Ω_j . Including both these factors leads to the general expression for the intra-cavity field, α_F , in the limit of weak pumping:

$$\alpha_F = -\frac{2C}{1 + 2N_{\text{eff}}C} \sum_j^N \frac{g_j}{g_0} \frac{\Omega_j}{2g_0} \quad (\text{A.1})$$

In this Appendix we consider the behaviour of α_F as Ω_j increases. For this situation, the weak pumping limit is, obviously, not a reasonable approximation and we cannot use Eq.(A.1). Instead, finding the equation of motion for $\langle \hat{a} \rangle$ from the fluorescence Hamiltonian given in Chapter 2 results in the implicit expression:

$$\alpha_F = -2C \sum_j \left(\left| \frac{g_j}{g_0} \right|^2 \alpha_F + \frac{\Omega_j}{2g_0} \right) \frac{1}{1 + |2g_j\alpha_F + \Omega_j|^2/\gamma^2} \quad (\text{A.2})$$

This must be solved numerically, with the additional complication that each atom ‘feels’ a different g_j and a different pump beam Rabi frequency, Ω_j , so that their contributions will not necessarily add in phase.

Our first approximation is to treat the resonant case (i.e. $\Delta_a = \Delta_c = 0$). For this case, were the placement of the atoms in the cavity such that they ‘felt’ identical coupling factors and pumping rates, we could take $\Omega_j = \Omega$ and $g_j/g_0 = 1$ and rearrange Eq.(A.2) to find:

$$\frac{4g_0^2}{\gamma^2}\alpha_F^3 + \frac{4g_0\Omega}{\gamma^2}\alpha_F^2 + \left(1 + 2CN_{\text{eff}} + \frac{\Omega^2}{\gamma^2}\right)\alpha_F + 2CN_{\text{eff}}\frac{\Omega}{2g_0} = 0 \quad (\text{A.3})$$

This equation can be solved numerically and two solutions are shown in Figure.A.1. Here, the fluorescence signal rate $J_{fl} = 2\kappa|\alpha_F|^2$ is plotted as a function of the transverse pumping rate, Ω for two values of N_{eff} . The LHS shows the case of $N_{\text{eff}} = 1$, which would be the case in our cavity when the atoms are dropped. Here, J_{fl} increases linearly and then smoothly rolls over when $\Omega = \sqrt{2}\gamma$, which corresponds to the case of $I = I_{\text{sat}}$ for an atom in free space. J_{fl} then monotonically decreases, with a fully saturated atom causing the cavity to go dark.

The RHS of the same figure shows the case for $N_{\text{eff}} = 13$. This value corresponds to $C_N = N_{\text{eff}}C = 4$, assuming our cavity parameters of $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz. When pumping the cavity ‘directly’, this value of C_N described the onset of optical bistability. The figure indicates that this is replicated in fluorescence, with J_{fl} becoming multivalued close to $\Omega = \sqrt{2N_{\text{eff}}}\gamma$. This seems reasonable, as the features of directly pumping the cavity are present here; the atoms are being driven in phase with one another, the only difference is the source of the pumping light.

There is no reason to believe, however, that this accurately represents our experiment; atoms falling through the cavity will take random positions with respect to the pumping beam and the cavity. A more reasonable starting assumption, therefore, is to assume the atoms have different coupling and pumping rates and scatter randomly out of phase. In this case we return to Eq.(A.2), which must now be evaluated numerically. Given this, it is interesting to ask whether this leads to different behaviour.

The results of Master Equation simulations of this regime are shown in Figure.A.2. Ten simulations are presented, for each of which the intra-cavity photon number, $n = |\alpha_F|^2$, is calculated as a function of the saturation parameter, I/I_{sat} . In each simulation, three atoms are assigned random positions throughout a cylinder of radius $2w_0$ and length d , centred on the cavity mode, such that $N_{\text{eff}} = 0.2$ ¹. Eq.(4.6) is then solved, given each atom’s coupling rate, g_j and beam intensity, I_j .

Whilst each ‘throw’ of atoms is different, corresponding to a different configuration of atoms with respect to the pump beam and cavity mode, there do appear to be some generic features. In every case, $|\alpha_F|^2$ increases with increasing I/I_{sat} , tending towards a limiting value

¹A relatively small value for N_{eff} was chosen for these simulations to reduce the dimensions of the Hilbert space used. 3 atoms are used to produce a $N_{\text{eff}} = 0.2$, corresponding to 16 basis states. To calculate for $N_{\text{eff}} = 1$ requires 16 atoms, leading to over 130,000 basis states




Figure A.1: The fluorescence signal, J_{fl} as a function of transverse pumping rate Ω , where $J_{fl} = 2\kappa|\alpha_F|^2$ and α_F calculated using a numerical solution to Eq.(A.3). The cavity parameters used were $(g, \gamma, \kappa) = 2\pi \times (98, 3, 5200)$ MHz and it is assumed that for all atoms, $g_j = g$ and $\Omega_j = \Omega$. LHS: $N_{\text{eff}} = 1$. At this value, J_{fl} increases steadily and then smoothly rolls over when $\Omega = \sqrt{2}\gamma$, from whence it monotonically decreases - leading to a dark cavity when the atoms are fully saturated. RHS: The case of $N_{\text{eff}} = 13$, which corresponds to $C_N = 4$. This shows similar behaviour, with a smooth roll over in J_{fl} , then decreasing to a dark cavity for fully saturated atoms. However, when $\Omega = \sqrt{2N_{\text{eff}}}\gamma$, J_{fl} is multi-valued, showing optical bistability in the cavity.




Figure A.2: The results of 10 simulations calculating the intra-cavity photon number as a function of the saturation parameter, I/I_{sat} . For each simulation, three atoms are assigned random positions throughout the cavity volume, such that $N_{\text{eff}} = 0.2$. Eq.(A.2) is then solved, given each atom's coupling rate, g_j and beam intensity, I_j . Whilst each 'throw' of atoms is different, corresponding to a different configuration of atoms, the general features for each simulation are the same; $|\alpha_F|^2$ initially increases with increasing I/I_{sat} ; then reaches a maximum, which covers several orders of magnitude in I/I_{sat} ; and then rolls over so that at as $I/I_{sat} \rightarrow \infty$, $|\alpha_F|^2 \rightarrow 0$. This is qualitatively similar to the case of purely coherent scattering. The dashed line represents the upper limit of $|\alpha_F|^2$, which occurs at $|\alpha_F|^2 = N_{\text{eff}}C\gamma/\kappa$.

of $N_{\text{eff}}C\gamma/\kappa$. The initial gradient of this slope appears to be accurately described by the intuitive expression:

$$|\alpha_F|^2 = 2C_N \frac{s'}{1 + s'} \quad \text{where} \quad s' = \frac{s}{(1 + 2C_N)^2}$$

The spread in the slopes of increasing $|\alpha_F|^2$ is caused by the beam profile of the pump beam, which had no effect in the earlier calculation. Also unlike the earlier case, this limit extends over several orders of magnitude and it isn't until $I/I_{sat} \sim 10^6$ that $|\alpha_F|^2$ rolls over begins to decrease. From this point increasing I/I_{sat} reduces $|\alpha_F|^2$, so that $|\alpha_F|^2 \rightarrow 0$ as $I/I_{sat} \rightarrow \infty$ and the cavity goes dark. This behaviour appears to be qualitatively similar to the case of purely coherent scattering.

For reasons already outlined, it is difficult to increase the numbers of atoms contributing to the signal in this simulation.