Simulating transport through quantum networks in the presence of classical noise using cold atoms

by

Christopher Gill

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Ultracold Atoms Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
The University of Birmingham

Sep 2016

Abstract

In work towards the development and creation of practical quantum devices for use in quantum computing and simulation, the importance of long lived coherence is key to the advantages offered over classical systems. Isolation from environmental sources of decoherence is of integral importance to such platforms. Recent experiments of biological systems seem to indicate the existence of certain structures that utilise the decohering effects of their surrounding environments to enhance performance.

In this thesis, we present a novel experimental set-up used to simulate the effect of decoherence-enhanced systems by mapping the energy level structure of laser cooled Rubidium 87 atoms interacting with electromagnetic fields to a quantum network of connected sites. Classical noise is controllably added to the states to create a tunable system to explore these effects. We present the design and implementation of the system, with results of several technical aspects, along with initial work on the effect of the applied noise on transport through the network. We present a key result that transport through the network is non-trivially enhanced by the presence of an optimum amplitude of noise, pointing to the interplay between the coherent nature of the quantum system and the decoherence provided by the noise.

ACKNOWLEDGEMENTS

I would like to acknowledge the following people:

Person	**Reason**
Dr. Vincent Boyer	For being simply a fantastic supervisor and mentor.
Dr. Plamen Petrov	For his overall support and helping me with the wonderful world of UHV technology.
Andy White	For joining the experiment in my final year and putting up with the stressful pace of learning an experiment built from the ground up with *somewhat* less than perfect documentation.
The Nuclear Physics PhD Students *et al.*	For being reliable lunch companions and great friends.
The Birmingham iSense Team	For sharing a lab with me for several years and allowing me to test the very limits of 'borrowing' equipment.
The Birmingham Cold Atoms Group	For the vibrant group environment and many tech-talks.
Louise	For putting up with occasionally playing second fiddle to my PhD.
My Parents	For the tremendously positive impact they have had on my life.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

OVERVIEW

Typically, when considering desirable characteristics of quantum systems, the presence of interference is key to their operation. The ability to interfere relies on coherence between the different states of the system. For example, the entanglement of qubits in quantum computing [1], the interferometer paths in atom interferometers [2] and Josephson junctions in SQUIDS [3] all operate on the purity of this coherence. In physically implemented systems, the state coherences decay in time via various mechanisms. This normally reduces the efficiency or practical limits of such systems. One of the main mechanisms for decay is the influence of, and interaction with, the surrounding environment of the system. It is believed that these mechanisms are crucial to the transition of quantum mechanics to the classical world that we see on macroscopic scales and at room temperatures [4]. In such quantum systems where coherence is desirable, great care must be taken to isolate them from the thermal environment of the world, through either placing the devices in vacuum systems or in sophisticated cryogenic temperature setups.

Recent observation of quantum coherence in biological structures [5] at physiological temperatures has led to great interest in the surprising possibility that quantum effects might provide an important role in the operation of systems which are coupled to room temperature environments. Theoretical investigations have proposed that the startling efficiency of such systems may actually be a result of the complex interplay between the quantum coherence of the systems and their surrounding thermal environment. Unlike typical quantum systems, their operation is potentially enhanced by the decoherence provided by the environment.

Unfortunately, the nature of these biological systems is such that in-depth testing of these models is not currently practical. Even the simplest of biological examples of such systems require complex methods to detect the coherent properties. Instead, there has been interest in creating highly controllable systems that can mimic the behaviours that are fundamental to these proposed models. The term ‘Simulations’ to many

means the narrow interpretation of using mathematical expressions solved by a classical computer to mimic the behaviour of systems. If we use the broader meaning of any system that mimics the behaviour of another, these controllable systems themselves can be considered simulators [6].

The subject of this thesis is work undertaken to build a simulator capable of investigating the effects of decoherence-enhanced quantum systems. We use dilute clouds of laser-cooled Rubidium atoms subjected to a multitude of externally controlled stimuli to create a system that provides a demonstrably coherent system with the facility to add variable amounts of decoherence to enhance particular dynamical properties.

1.1 Background

This section discusses the background and terminology of these decoherence-enhanced systems.

1.1.1 Biological Motivation

The bulk of the interest in environment-assisted systems stem from biological investigations into photosynthesis and animal magnetic sensing [7]. The more precise motivation for the work in this thesis is from electron beam spectroscopy on Fenna-Matthews-Olson (FMO) complexes which are integral to the process of photosynthesis in green sulfur bacteria [5].

Briefly, photosynthesis is the process used by biological systems to transfer light energy to chemical energy. The absorption of photons is typically done by specific arrays of molecules such as chlorophyll in green plant life. These molecules have large optical absorption cross sections for the process of exciting electrons from the valance band to create an electron-hole pair. Due to Coulombic interaction, the electron is still localised to the hole in a hydrogen-like quasiparticle referred to as an exciton. The lifetime of such excitations is typically on the order of picoseconds [8] which is obviously inadequate for energy storage in physiological systems.

The *reaction centre* is a set of proteins that specialise in utilising the energy stored in these excitons to conduct a chemical process that stores the energy chemically for later use.

In some biological systems, the chlorophylls are closely integrated into the reaction centres but in others, these two elements are separated and there is a third complex that mediates the transfer of excitons from the light-harvesting arrays to the reaction centres. Green Sulphur Bacteria (GSB) are an interesting example of the latter case. GSBs are highly suited to low light environments due to their highly efficient light harvesting chlorosomes. The chlorosomes and reaction centres are composed of different types of bacteriochlorophyll (BChl). The transfer of excitons from the chlorosomes to the reaction centres is mediated by a structure

also composed of BChl molecules called the FMO complex. The FMO complex itself is a trimer with each element identically composed of seven BChl A molecules arranged in a protein scaffolding. An exciton passed from the chlorosome can be localised to one of these molecules but Coulomb coupling between proximal BChl sites provides a path for the exciton to 'move' between the sites. Measurements of the structure of FMO complexes has led to a standard numbering of the different BChl molecules. Excitons from the chlorosomes are generally passed to the FMO at sites 1 and 6 with the reaction centre coupled to site 3. The role of the FMO complex is then to transfer the excitons from their initial sites to the target site before the exciton decays due to its finite lifetime.

The surprising discovery that the FMO complex displays relatively long lived coherent properties at physiological temperatures by Engel *et al* [5] lead to proposals that the surprising efficiency of this exciton transport process is the result of the interplay between the underlying coherence mechanisms with the thermal environment of the complex. Environments of quantum systems are a typical source of decoherence. The effectiveness of such an interplay could have been honed by evolution of the structure of the complex to maximise the transport efficiency and robustness to mutations or damage [9]. These discussions have led to a more general interest in the concept of environment-assisted processes, especially with respect to the transport of an excitation to a target state. The FMO system is typically modelled by treating each of the seven BChl molecules as a discrete site in which a single exciton can be localised. Each site has an associated energy and each pair of sites has a coupling determined by calculations made by Adolphs and Renger [10]. The reaction centre connection can be modelled as an irreversible transition from one of the sites to a 'sink state' in which population accumulates. Discussion is normally limited to a single excitation, i.e. there is never more than one exciton present in the complex. This is justified in the FMO complex because of the relatively small amount of light that it is exposed to in its natural habitats. Whilst the applications of these mechanisms in biological systems provide a fascinating motivation to probe deeper, the biological systems themselves are very difficult to manipulate and test different theories on. From here, we generalise our interest in these mechanisms to a quantum network.

1.1.2 Quantum Networks

Here, we define a quantum network[1] as a set of coupled quantum states with an emphasis on the propagation in time of a prepared initial state, through the network to the other states, specifically to an end 'goal' state. It is natural to visualise these states as spatially separated but this is not necessarily the case. The generality

[1]It should be noted that in the quantum information community, a quantum network is defined as a physical network used to transfer quantum states from one location to the other. We do not use this definition.

of this concept is the key to the work of this thesis. Whilst the motivation for much of the interest in decoherence- or environment-assisted transport is biological in nature, the same principles have to potential to be readily mapped to a wide variety of quantum systems for the purposes of simulation.

Figure 1.1 shows a schematic representation of a quantum network. The topology of a quantum network is specified by the energies of the states and the coupling strengths between them, which govern the ease with which transitions can occur. This is suitably represented for a network with N states by an $N \times N$ matrix, where the diagonal elements, ϵ_i correspond to the energy of the ith state and the off diagonal elements, $V_{i,j}$ correspond to the coupling strengths between the ith and jth states. This matrix conveniently describes the Hamiltonian of the system. The goal state of a network may irreversibly remove population out of the network and to some other system. This is typically modelled using an irreversible decay to a 'sink' state.

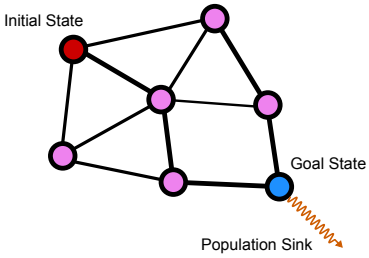

Figure 1.1: An example of a quantum network.

With $N = 2$, the network reduces to a trivial two-level system, exhibiting coherent oscillations between the states in the form of Rabi flopping (see Section 6.4). In the case where the coupling is weak compared to the relative energy differences between the states ($|V_{12}| \ll |E_1 - E_2|$), the system will mostly stay in its initial state in the absence of an environment as the eigenstates of the system are very similar to the uncoupled states. However, the addition of a coupled environment can cause coupling between these system eigenstates, thus allowing significant population to move out of the initial state. This can be thought of as a 'blurring' of the energy levels such that there is more overlap between the sites. The relative simplicity of the $N = 2$ case means it does not contain some of the phenomena that make quantum networks so interesting. Nevertheless, the simplicity provides an excellent test model for coupling to the environment. Many theoretical works looking at advanced environmental interactions beyond the standard approximations use the $N = 2$ case as a testbed for their methods [11, 12].

With $N = 3$, the network starts to exhibit more complex behaviour. The presence of multiple paths between sites means that the relative phase of the couplings becomes important. The different paths to a state coherently interfere with each other, inhibiting or enhancing movement to the state, dependent on the relative phase between the paths. This can obviously cause problems in the inhibiting case if the purpose of the system is to move to some goal state. The introduction of system decoherence by an environment or other source leads to a destruction of this coherent interference between these paths and can lead to an

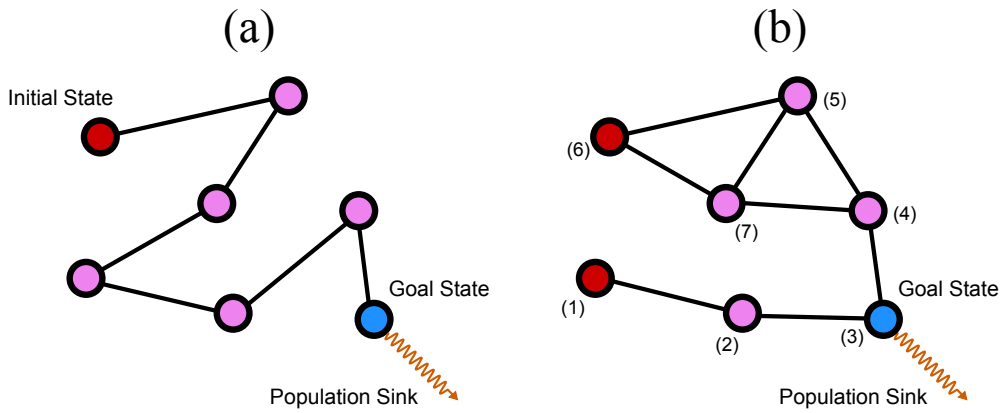

Figure 1.2: (a) An example of a ladder network. (b) A schematic of the FMO complex only showing the strongest couplings.

increase in performance of the system [13].

A particular special case of a quantum network is the ladder system [14] as shown in Figure 1.2a. This is a system with N levels where only adjacent sites are coupled to each other ($V_{i,j} = 0$ if $|i - j| \neq 1$) and the energies are spaced linearly apart. The presence of an environment causes behaviour similar to the $N = 2$ case discussed above where it induces coupling between the eigenstates of the system but does not have the interfering paths behaviour of the general $N = 3$ example discussed above. Almost the antithesis of the ladder network is the fully connected network (FCN), which contains couplings between every site pair. These systems may exhibit strong path interference effects. A special case of the FCN is one where every coupling term is the same as well as all the energies of the sites being the same. Here, interference effects play a hugely fundamental role [14] on the operation of the system and may make such a system an excellent test-bed for investigations into such effects.

Using this network framework, the FMO complex is typically modelled using energy levels and couplings obtained from Adolphs and Renger [10]. Whilst it is technically a FCN, certain strong couplings dominate. Figure 1.2b shows a network schematic of the complex showing only the strongest couplings. The overall structure is of two starting branches coming together near the goal state. The population enters the complex from the end of either of the branches.

The first step of the experimental proposal of this thesis is to simulate these quantum networks in an $N = 3$ ladder system, with further extensions being possible to a fully connected $N = 3$ system and then systems with higher N. However the majority of the work undertaken was to arrive at an $N = 3$ ladder system.

1.1.3 Environmental Coupling vs. Classical Noise

A key discussion to be made with regards to this thesis is the nature of the decoherence that causes the enhancement of particular system functions.

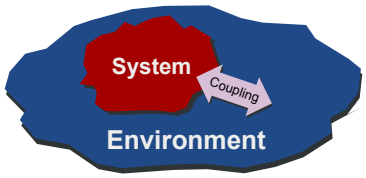

Figure 1.3: A schematic showing a system coupled to an environment. Evolution of the whole system + environment is closed but under certain conditions, the system can be solved alone using open system dynamics.

The most commonly discussed source, especially regarding the biological systems discussed previously, is that of system-environment coupling. The time dynamics of any closed quantum systems are described by the evolution of state vectors in the Schrödinger equation, or when it is required to deal with statistical mixtures, with the evolution of the density matrix and the Liouville-von Neumann equation. Prediction of these time dynamics is 'simply' a case of finding the initial conditions of such a system and the system Hamiltonian to find the unitary evolution of the system. Issues arise in situations where the system of interest is coupled to another, much larger system, often called the environment (see Figure 1.3). Let us start with the Liouville-von Neumann equation ($\hbar = 1$)

$$\frac{d\rho_{se}}{dt} = -i\left[H_{se}, \rho_{se}\right] \tag{1.1}$$

where ρ_{se} is the density matrix of the system plus environment and H_{se} is the Hamiltonian of the system plus environment. In principle, one could solve the unitary dynamics of this large compound system and then trace out the environment's degrees of freedom to just leave the state of the system of interest. However, the solution of the unitary dynamics of large quantum systems is in general very difficult. Instead, the system of interest can be treated as an open quantum system that evolves non-unitarily. In the simplest case, this can be achieved by modifying the Liouville-von Neumann equation via the addition of a Lindblad operator that governs these non-unitary dynamics[1]

$$\frac{d\rho_s}{dt} = -i\left[H_s, \rho_s\right] + \mathcal{L}\rho_s \tag{1.2}$$

where now ρ_s is the density matrix of just the system, H_s is the Hamiltonian of the system in the absence of the environment and $\mathcal{L}\rho_s$ is the Lindblad superoperator acting on ρ_s to model the non-unitary effects of the environment. This can typically only be done under certain assumptions. For particularly large environments with weak coupling to the system, one can take the Born approximation which states that the environment

[1] This is typically done in the interaction picture but the detail is left out for brevity.

stays approximately constant throughout the evolution. For example, a thermal bath of bosons does not behave fundamentally differently with the addition or removal of a microscopic number of bosons from one of its modes. Additional simplification comes from the Markov approximation, which states that the evolution of the system only depends on its current state and not on any information of its state in the past. The practical implication of the Markov approximation is that any coherent relationship between the system and environment must decay on a much faster timescale than the dynamics of the system.

More advanced techniques now attempt to find solutions outside of these limitations. For the FMO complex, it is believed that the surrounding protein scaffolding that forms the immediate environment is not large enough to make the weak coupling approximation and also has strong memory effects that violate the Markov approximation. Work has been done to circumvent the weak coupling for example by using variational polaron transformations [15] and methods developed to deal with non-Markovian environments [16–18], for example by moving the non-Markovian dynamics into an intermediate pseudosystem (for which the dynamics are fully resolved) which is itself coupled as an intermediary between the initial system and a Markovian environment [19] (see Figure 1.4).

An alternative to this environment-system coupling as a source of decoherence is the presence of a stochastic process in the system Hamiltonian. This, in general, makes the Hamiltonian not just time dependent, but also dependent on the sample of the stochastic process. Whilst a stochastic process can be defined or characterised using particular metrics that can be shared between all samples (such as spectral density or correlation time), the resulting state of a system subjected to such a Hamiltonian will necessarily depend on the particular time sample taken by the stochastic process. The evolution of the state along a single trajectory is still unitary, but if we consider an ensemble of state evolutions averaged over many stochastic trajectories, we in general have non-unitary [20] evolution of the ensemble average. This process is represented in Figure 1.5. Such situations arise in quantum systems that are subjected to fields from uncontrolled stochastic stimuli such as in solid-state systems, where we may be measuring the net effect from a collection of particles that experience fields produced from the presence of other types of particles. Whilst in general, each particle in the collection is subjected to the same processes, the exact details of the Hamiltonian experienced by each one will be different. A well known consequence of such dephasing is the relaxation times associated with the free induction signal in Nuclear Magnetic Resonance (NMR) experiments. The magnetic field in the Hamiltonian may be slightly different for each particle in the sample due to local inhomogeneities in the field due to, for example the dipoles of other particles in the sample or other particles in the material. Another example is trapping centres in superconductors interacting with the Fermi sea [21].

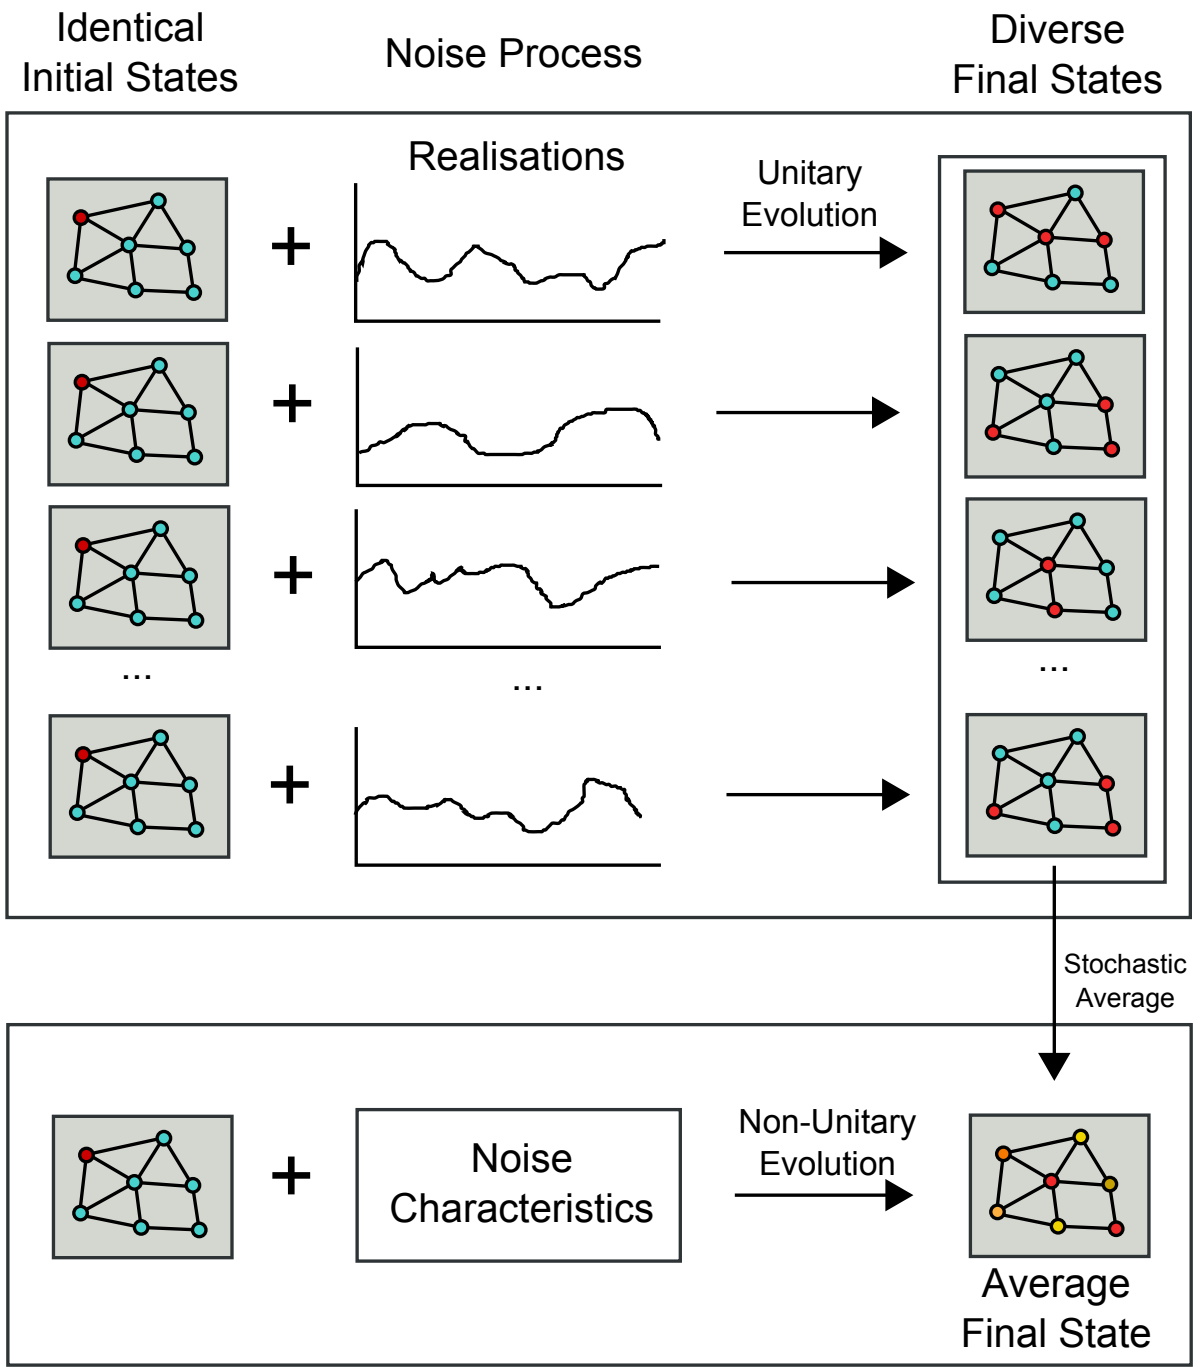

Figure 1.5: A schematic showing the process of stochastic averaging over multiple trajectories of a stochastic process. Whilst each individual trajectory follows the unitary dynamics of a closed quantum system, the average of the final states compared to the initial state is in general a non-unitary evolution. Note: The final states (and indeed, the initial states) are in general superposition states of the basis states. It is hard to get this across diagrammatically.

There are interesting parallels between these two broad approaches. To explore these, let us consider an example of a classical random electro-magnetic field signal coupled to the on-diagonal elements of a qubit. It seems uncontentious to say that this should be the equivalent to the coupling of the system to an 'environment' of a photon bath. To maintain the equivalence, the system may not significantly affect the photon bath as this would be the same as changing the characteristics of the field. This can be satisfied if the photon bath is in a high occupation limit. An interesting characteristic of the photon bath case in open quantum systems is the ability for the system to irreversibly transfer energy into the environment, the mechanism for spontaneous emission of energy states. This cannot happen in the case of a classical field as spontaneous emission is an intrinsically quantum process. The equivalence is maintained if the stimulated absorption and emission processes dominate the spontaneous emission as to make it irrelevant. This is true for extremely large occupation numbers of the bath with a very weak coupling.

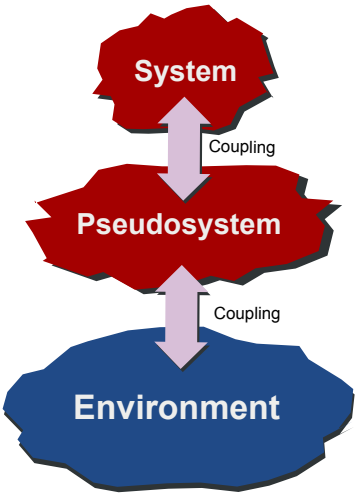

Figure 1.4: A schematic showing an example of a technique for dealing with non-Markovian environments. The non-Markovianity is moved to a new pseudosystem that interacts with the Markovian environment and the system. Both the time dynamics of the original system and new pseudosystem have to be resolved.

These parallels lead to the conclusion that under certain conditions, the dynamics of the stochastic average of particular forms of classical noise are the same as the dynamics of systems coupled to environments in the Born approximation. A specific example of this is shown by Saira *et al.* [20] where the average response of a quantum two-level system to classical Random Telegraph Noise (RTN) is shown to be equivalent to the open quantum system solution using an exponentially decaying environment correlation function. This equivalence allows us to investigate environmental coupling without the experimentally difficult task of coupling a real quantum environment to another controlled system. Whilst this equivalence is confined to these limited cases, there may be possible extensions to non-Markovian or stronger coupling regimes. As each trajectory of the stochastic process can be applied deterministically to the system, the state of the system can be read out progressively and the information used to change the 'future' form of the trajectory. This implicitly introduces dependence of the noise on the system as it evolves in time, which may, when viewed in the context of an open quantum system, imply environmental coupling that does not obey the Born or Markov approximations. A schematic of such a scheme

is shown in Figure 1.6. It should be noted that this scheme bares similarity to quantum control schemes where feedback parameters are changed based on weak non-destructive measurements of the system [22, 23]. Whilst these thoughts potentially provide exciting promise for the direction of this work, the results of the work contained in this thesis does not touch on them further.

1.1.4 Simple Model - Dephasing

A simple model used extensively in the work of Plenio [8, 13, 14] is the use of a dephasing Lindblad operator in the Liouville-von Neumann equation. Accompanying these papers is an intuitive explanation of the effect of the noise, decomposed into two parts. The first is the idea that the application of dephasing is the equivelent to finitely 'blurring' the energy levels, artificially enhancing their 'line-width'. If two coupled levels have a large energy difference, the effect of the coupling will be comparatively smaller than if they had the same energy level. The enhancement of population exchange between levels can therefore be achieved if the levels are blurred to enhance the overlap between them. This is shown in Figure 1.7a. The extent of the widening of the levels is related to the strength of the dephasing, we can intuitively visualise a regime where the mutual overlap between the levels is maximal. This implies a level of dephasing that is optimal for enhancing coupling between the levels. This effect is shown in the Plenio papers cited above. It can be considered that the source of the dephasing provides a mechanism for the system to gain and lose energy to make otherwise non-energy conserving transitions.

The other stated idea is related to potential destructive interference between multiple paths. Consider the situation shown in Figure 1.7b; the fixed phase difference accumulated between the two arms of the Mach-Zehnder interferometer causes destructive interference at one of the outputs and no light is seen. Applying dephasing (shown in Figure 1.7c) to one of these paths reduces the interference effect and allows light to appear at the previously dark port. Here, the dephasing opens transport channels that are otherwise inhibited.

1.1.5 Experimental Simulation

Any proposed system to be used to simulate these decoherence-assisted processes can be considered according to whether they are capable of incorporating many of the different aspects that may be of interest besides their capability to simulate the dynamics of system-environment coupling or the application of a classical stochastic field. Here, we outline and briefly explain a collection of these aspects before discussing proposals that have been suggested with respect to these criteria.

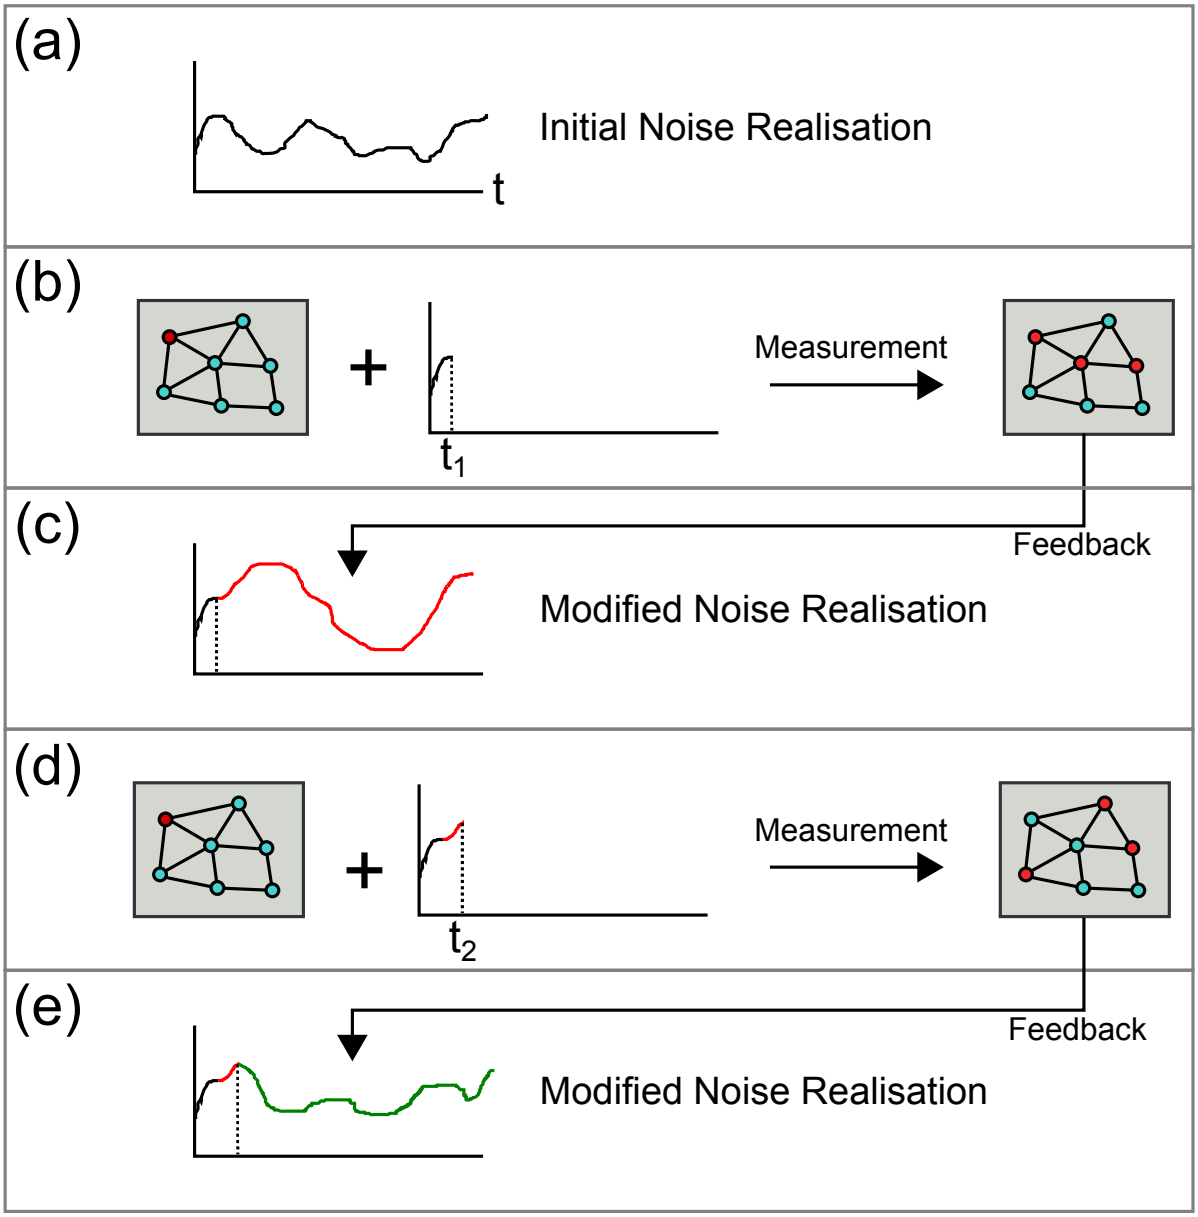

Figure 1.6: Illustration of a potential noise feedback system. (a) A single realisation of a stochastic noise process is applied to the sample. (b) A measurement is made on the sample at time t_1 (c) The noise trajectory is modified as a result of the measurement. (d, e...) This process is iterated at different times leading to evolution that depends not only on the current state of the system but also its past.

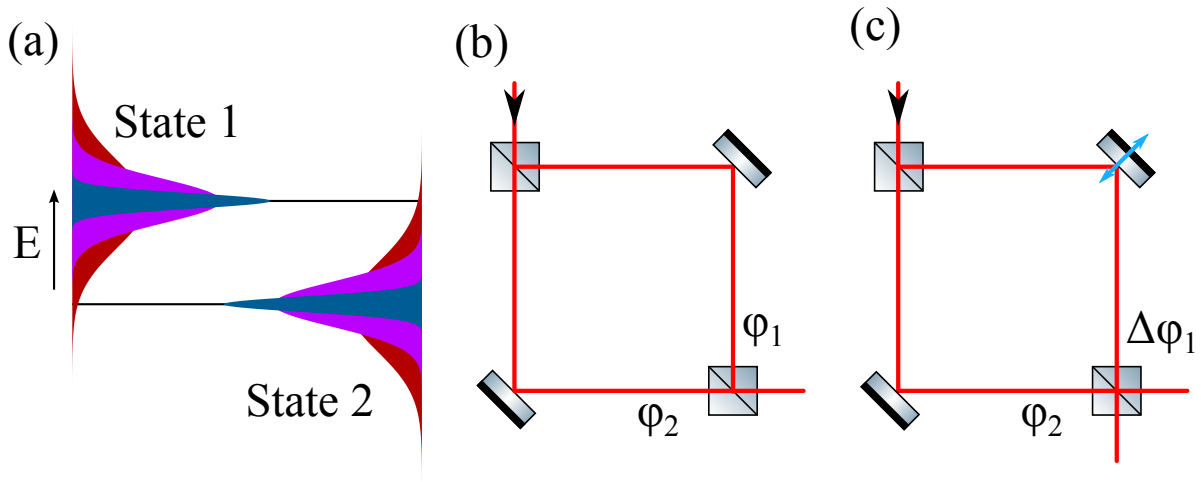

Figure 1.7: Examples of how dephasing can enhance a system. (a) The two states have different energy levels so coupling between them is weak. If the dephasing adds some width to the transitions (shown by purple and red), the overlap between the levels is enhanced. (b) The defined phase difference between the two paths of a Mach-Zehnder interferometer leads to destructive interference of the light at one of the output ports, inhibiting light emitting here. (c) If one of the path length is varied, the destructive interference is reduced and light can appear at both outputs. (a) is heavily inspired by Figure 1a in [8].

The most fundamental requirement is to have a system that is intrinsically quantum in nature. To be able to see the effect that different environments have on a system, it must be capable of providing long lived coherence in the absence of the applied 'testing environment'. Obviously, it is impossible to completely decouple a system from its real environment but it is desirable that any coupling to its 'real' environment is minimal and allows for coherence to exist longer than the models to be tested.

Coupling between states in a proposed system should ideally provide many degrees of freedom to the topology of the network and the strength of the couplings. Specific to the topology is the desire to be able to observe interference between multiple paths to the same site.

Another important property is the method of implementation of the decoherence mechanism intended to enhance a property of the system, and specifically the validity and applicability of such an implementation to the motivating situations. Of specific interest is possibilities to provide couplings that could potentially break the Born and Markov approximations.

An additional desirable property specific to the concept of quantum transport is the ability to mark a state as a target state and for population to irreversibly decay from this site into a 'sink'. The speed and total amount of accumulation in this sink then marks a simple metric for the efficiency of the network. This mimics theoretical models used by, for example Plenio and Huelga [13].

1.2 Simulation Proposals

Whilst there has been much discussion on the theoretical understanding of the mechanisms associated with environment assisted transport, there has been relatively little experimental work done to test or demonstrate these mechanisms in physical systems. Many sources have pointed to the use of cold atomic and ionic systems for their past and current quantum simulation work [24] in the fields of solitons in Bose-Einstein Condensates (BECs) [25] and Anderson localisation in disordered lattices [26–28]. Such systems provide excellent testbeds for such intrinsically quantum mechanical tests as they provide a clean quantum environment with a multitude of tools and methods developed for manipulation.

Previous work regarding BEC's inside of double well potentials [29, 30] could be expanded on to investigate added sources of decoherence [31]. The wells can be created using high intensity, far-detuned optical light beams manipulated by elements such as Spatial Light Modulators (SLMs) or Acousto-Optic Deflectors (AODs) with a large degree of freedom on the potential created. The different potential wells can be mapped under certain approximations to the different sites, with the relative depths giving their energies. Quantum barrier tunnelling provides the coupling between the sites, the strength of which can be manipulated by the height and width of the barrier.

Plenio and Huelga [13] have suggested a 'moving barrier' problem consisting of a three level system with couplings between states $1 \leftrightarrow 2$ and $2 \leftrightarrow 3$ where states 1 and 3 have the same energy but state 2 has a higher energy. They show that adding dephasing noise onto state 2 can greatly increase the transport to the sink from state 3. In the same paper, it is proposed that this system could be implemented using a Raman transition in a cold ion system.

There are also proposals [11] to use cold atom systems coupled to a network of optical cavities. Cavity Quantum Electro-Dynamics (CQED) can be used to describe the system of coupled atoms and cavities and can reach the strong coupling regime. In such systems, the atoms would be considered the system and the cavity the environment. By adding coupling to other cavities, where a photon emitted into the cavity 'environment' could move between cavities and return later, this represents an environment that depends on the past state of the system, testing physics beyond the Markov approximation.

On the more biological side, work is being done to create synthetic complexes that mimic the same behaviour with specific pigmentation applied to allow for efficient measurement of the state of the system [32].

1.2.1 Our Proposal

Our proposal involves a comparatively simple set-up. We model a subset of energy levels in a laser cooled Rubidium 87 atom as the sites of a quantum network. The single valence electron leads to a simple atomic energy structure that is well understood. The energies of, and couplings between, the states can be manipulated using a variety of fields and techniques commonly used in the cold atoms community. This defines the topology of the quantum network. Standard techniques can also be used to prepare the system in a particular state, and read out the populations of the different states after a the system is allowed to evolve. The same techniques used to shape the energy levels of the network sites can be used to create time varying terms in the Hamiltonian, allowing the application of classical noise.

Figure 1.8 illustrates the main concepts behind the system. We select the $F = 1$ ground state of Rubidium, which contains three sub-levels, labelled by $m_F = -1, 0, +1$, each describing a projection of the angular momentum (Figure 1.8a). In the mapping to a quantum network, these atomic states form the sites of the network. The relative energy levels of the states can be changed by the application of a static magnetic field exploiting the Zeeman shift (Figure 1.8b). Couplings between adjacent levels are induced using an oscillating magnetic field (Figure 1.8c). This results in population exchange between the sites based on the frequency of the coupling field compared to the energy level differences. Decoherence is added by varying the static magnetic field in time randomly to 'shake' the energy levels (Figure 1.8d).

The topology of this network state forms a three-site ladder. This can be further complicated by the addition of a two-photon Raman transition (Figure 1.8e), which can couple the two outer states, or more generally the next-nearest neighbour states. An irreversible sink state can also be connected to one of the sites (Figure 1.8f), which will lead to spontaneous decay to a different atomic state. This can be used to test transport properties of the network from an initial state in the presence of the decoherence induced by the shaken energy levels.

Figure 1.8g shows the network topology for this three-level system formed by the $F = 1$ ground states in Rubidium. This could easily be extended to states with larger angular momentum and therefore a larger number of sites. Figure 1.8h for example shows the network topology if these techniques were applied to an $F = 3$ level such as the one found in Rubidium 85.

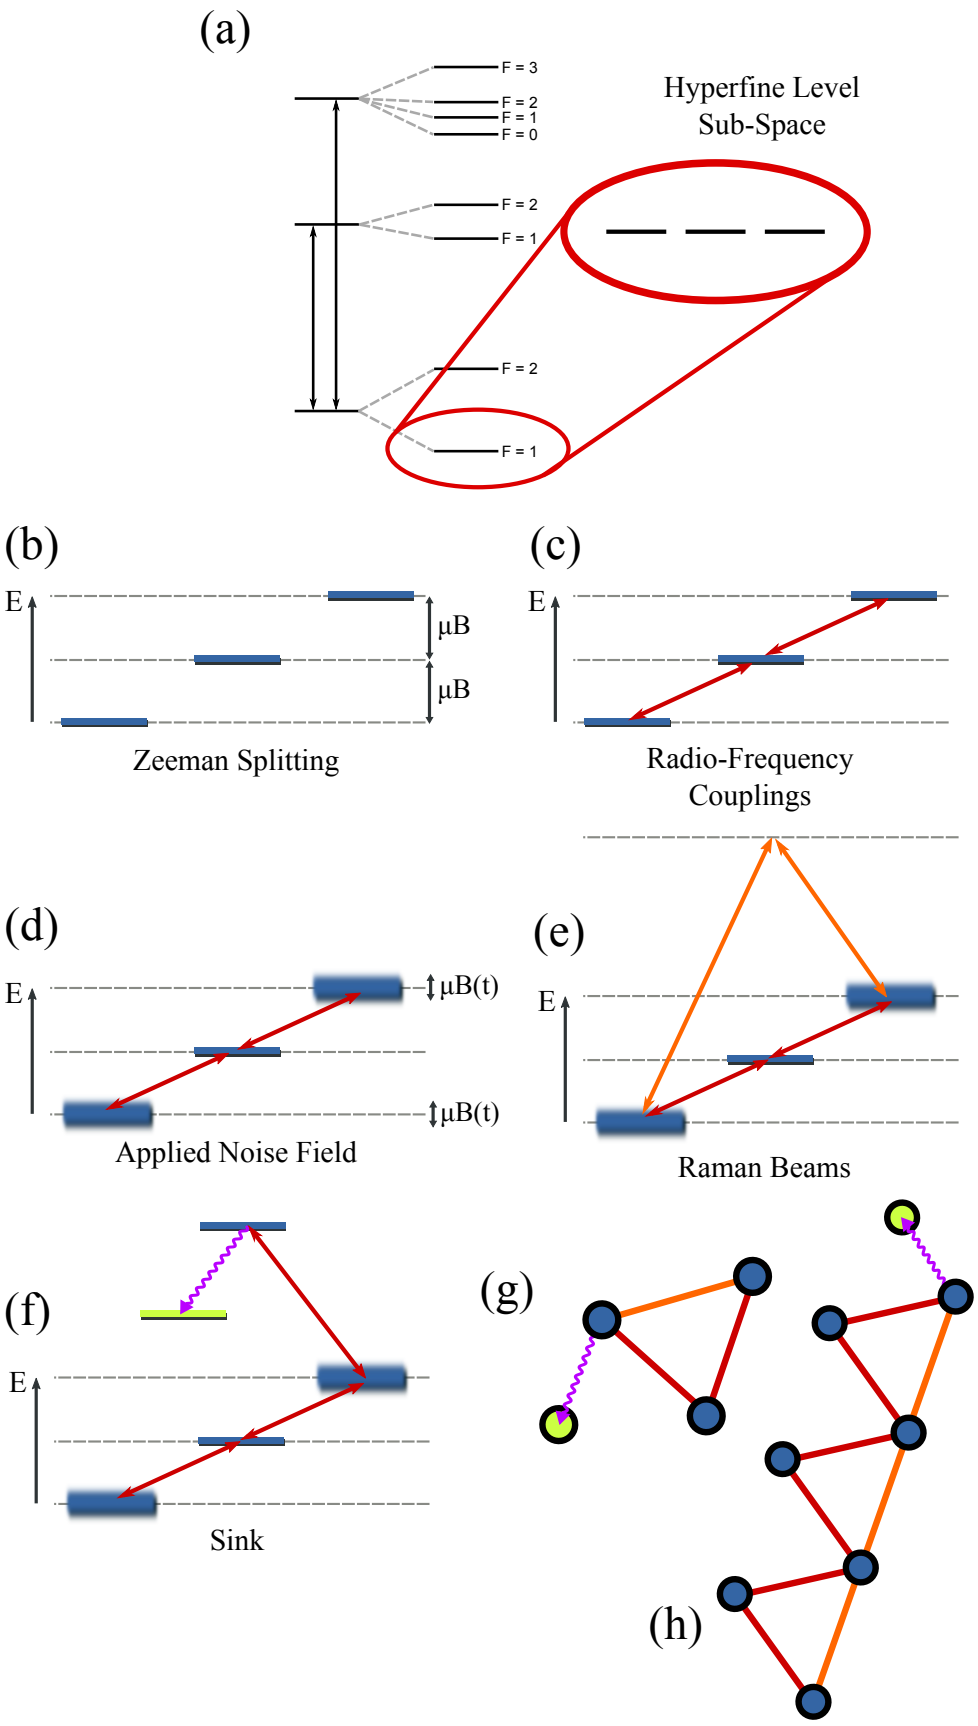

Figure 1.8: Subsequent buildup of the mapping of the energy level states of a Rubidium atom to a quantum network.

A summary of the characteristics of our system are shown below:

- **Quantum Network Sites** - Our system creates a quantum network with each energy level of the atom corresponding to a site on the network. In the presented example, the energy of each level is governed by the linear Zeeman shift.
- **Quantum Network Couplings** - Movement between the states of the network is created by an oscillating magnetic field to produce couplings between adjacent sites. The resulting network topology is a ladder-like system.
- **Coherent Evolution** - The isolated nature, and lack of thermal motion of the cold atoms means the couplings between the states are highly coherent and with little coupling to the surrounding environments.
- **Classical Noise** - The application of decoherence is done via the application of a classical field to shake the energies of the network sites. To get the general behaviour of the system, it is necessary to repeat any experiment with multiple *realisations* of the noise signal generated according to some common behaviour.
- **Non-Trivial Topology** - The topology can be expanded from the simple linear chain or ladder case with two-photon Raman transitions to produce next-nearest neighbour couplings. This allows the investigation of path interference effects between levels.
- **Disordered Structure** - The relative spacing between the site energy levels can be made disordered via application of stronger magnetic fields to access the non-linear response of the Zeeman shift.
- **Population Sink** - A controlled irreversible path out of the network can be implemented via coupling to a level with spontaneous decay.

Added to the required mapping of the atomic states of Rubidium to a quantum network, it is important to consider the practicalities of conducting useful experiments on the network. To be able to see the effect of the network, it is necessary to create apparatus capable of preparing the atoms in an atomic sample into the same initial state before the network couplings are 'switched-on' as well as infer information about the state of the network afterwards. The state preparation is implemented using light to optically pump the atoms in the sample to one of the network states. Once the system has evolved in time, the populations of the network sites are read-out by applying a magnetic gradient to the cloud, spatially separating the atoms in different states for differentiation by an absorption image.

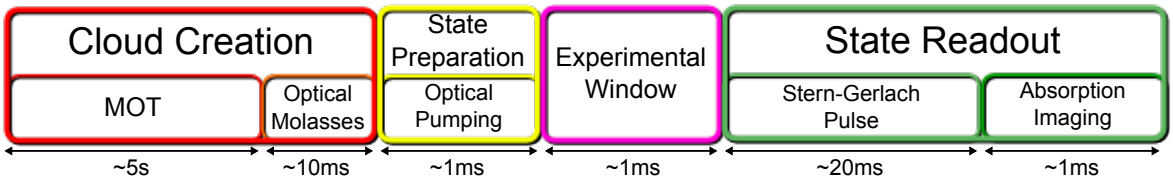

Figure 1.9: The different stages of a single experimental run, consisting of the creation of a cold atom cloud, exposing it to stimuli during the experimental window, followed by destructive imaging of the cloud.

An experimental 'run' consists of the creation, manipulation and subsequent destructive measurement of a single cloud of cold atoms. The components of a run are shown in Figure 1.9.

1.3 Thesis Contents and Structure

The main contents of this thesis are from the work conducted towards building apparatus capable of implementing the proposal outlined above. The vacuum chamber, laser systems and experimental control software were all built by the author during the Ph.D. period of study. As much of this work is not significantly novel, the contents of this thesis contain disproportionately little material compared to the length of time spent on these aspects.

The majority of the work in this thesis was conducted alone under supervision from Dr. Vincent Boyer and Dr. Plamen Petrov. During the final of four years, Mr. Andrew White joined the experiment as a new Ph.D. student.

As well as the main result of the thesis - the observed behaviour of the system enhanced by a classical noise field - we present results from many different aspects of the experiment, many peripheral to the overall experimental aim. The result of this is reasonable amounts of data that have been left with minor discrepancies unsolved. Attempts have been made to provide explanations for such discrepancies but many decisions were made to pursue the overarching goal of the experiment rather than worry about side issues.

1.3.1 Thesis Structure

As much of the work conducted contains many relatively disparate elements, the inclusion of a bulky theory section followed by a results section would be highly cumbersome and experimental chapters would constantly refer to individual sections in the theory chapter. Instead, **Chapter 2** provides an overview of some of the topics that are fundamental to many of the subsequent chapters, particularly the time dynamics of quantum

systems for multiple coupled states and the relevant energy level structure of Rubidium atoms. Many more specific pieces of theory are discussed directly in later chapters when it is relevant to the section being discussed.

After this, one chapter is dedicated broadly to each of these different stages shown in Figure 1.9 as follows:

- **Chapter 3: Cold Atoms Cloud Creation** - The initial task to implement the proposed experiment was the creation of an ultracold atoms set-up to create the laser-cooled clouds of atoms to serve as the system mapped to the quantum network. This chapter includes a brief outline of the general method, technical details of the specific set-up used and some typical results from measurements that characterise the system.

- **Chapter 4: State Readout** - An essential part of the experiment is the ability to observe the quantum states of the atom cloud after it has evolved in the applied stimuli. This is achieved using a magnetic gradient pulse to spatially separate the atoms in different quantum states and measure their relative populations. This chapter includes a detailed overview of the general method and experimental implementation, as well as quantitative results compared with the developed theory.

- **Chapter 5: State Preparation** - To properly infer the effect of the stimuli from the resultant state populations requires knowledge, and preferably repeatable preparation of an initial state. This is achieved using the standard technique of optical pumping. This chapter contains an overview of the scheme used and the use of the state readout apparatus to conduct quantitative analysis of the process.

- **Chapter 6: Nearest Neighbour Coupling** - With the ability to prepare an initial state and readout the populations after exposure to stimuli, the next step is to couple the states together in a chain using magnetic dipole transitions. This chapter contains discussion and results on several coupling schemes.

- **Chapter 7: Hamiltonian Diagonal Time Dependence** - With the states coupled, we introduce the decoherence term using a time varying modulation to the energy level splitting. This chapter includes a discussion and experimental results on the effects of *coherent* modulation of the energy levels as well as numerical simulations and experimental results on modulation resulting from a random process.

- **Chapter 8: Next Steps** - With the majority of the experimental work presented previously, this chapter includes a more in-depth discussion on the next steps that can be taken with the system including the implementation of next-nearest neighbour coupling and the possible implementation of a sink.

- **Chapter 9: Other Work** - At the start of the Ph.D., the apparatus was earmarked for an alternative line of enquiry on strong atom-light coupling in the low photon limit. The main focus of the experiment was switched to the current topic early on, however this initial work led to a publication with collaborators. This chapter includes a *very* brief outline of that research.

- **Chapter 10: Conclusions** - With everything discussed, the work presented is reviewed and concluded.

CHAPTER 2

GENERAL THEORY

In this chapter, we review the mathematics of quantum mechanics in closed systems with finite numbers of states. We discuss the time evolution of such systems, the effects of resonant and non-resonant couplings, the rotating wave approximation and the geometric representation of certain systems on the Bloch sphere. Much of this work can be found in any standard quantum mechanics textbook but Shore's treatment [33, 34] includes more of an emphasis on more complicated systems that the usual two-levels. Following the quantum mechanics review, the atomic physics of Rubidium 87 [35, 36] is discussed in the context relevant to laser cooling and trapping as well as the creation a coupled 3-level closed system for the simulations.

2.1 Quantum Mechanics Review

The general description of a pure state of a quantum system can be expressed as a state vector over a complete basis set. The choice of basis is to a large extent arbitrary but is normally chosen for convenience. Here we discuss generally over a finite basis set labelled with states $|1\rangle, |2\rangle, |3\rangle, ... |N\rangle$. A general state vector $|\psi\rangle$ can be expressed as a superposition over these basis states as

$$|\psi\rangle = c_1 |1\rangle + c_2 |2\rangle + c_3 |3\rangle + ... + c_N |N\rangle \tag{2.1}$$

where the complex coefficients c_i govern the overlap of the overall state with each basis state ($c_i \equiv \langle i|\psi\rangle$). A measurement of state $|\psi\rangle$ in the chosen basis set results in the collapse to one of the basis states $|n\rangle$ with probability $|c_n|^2$. Since probability must sum to unity, a constraint on the state coefficients c_n is given as

$$\sum_n |c_n|^2 = 1. \tag{2.2}$$

The time evolution of such a state is governed by the Schrödinger equation $i\hbar\frac{d}{dt}|\psi\rangle = H|\psi\rangle$. In our representation, this leads to the time dependence of the coefficients as

$$i\hbar \begin{pmatrix} \dot{c}_1 \\ \dot{c}_2 \\ \dot{c}_3 \\ ... \\ \dot{c}_N \end{pmatrix} = H \begin{pmatrix} c_1 \\ c_2 \\ c_3 \\ ... \\ c_N \end{pmatrix} \tag{2.3}$$

where H is the Hamiltonian operator of the system which takes the form of a $N \times N$ matrix. The diagonal elements of the Hamiltonian correspond to the energies of the basis states and the off-diagonal terms correspond to couplings between the different basis states that can result in population exchange between the states.

Solving the dynamics of the system is an exercise in finding solutions to the simultaneous differential equations for the state coefficients c_i.

2.1.1 Mixed States

As well as quantum superpositions between states with which one can produce interference effects, it is also useful to be able to describe a system that is merely in a statistical mixture of different states, i.e. the state is in a distribution of pure states, with each assigned a classical probability. A wavefunction has no method of describing such a *mixed state* of the system and instead a *density matrix* representation is employed to describe statistical mixtures. For a pure state, $|\psi\rangle$, a density matrix ρ is formulated by

$$\rho = |\psi\rangle\langle\psi| \tag{2.4}$$

and a statistical superposition of the pure states $|\psi\rangle$ and $|\phi\rangle$ is given by

$$\rho = P_\psi |\psi\rangle\langle\psi| + P_\phi |\phi\rangle\langle\phi| \tag{2.5}$$

where P_ψ and P_ϕ describe the classical probabilities of being found in the respective pure states. The density matrix takes the form of a $N \times N$ matrix with a trace of 1. The diagonal elements represent the *populations* of each of the basis states and the off-diagonal elements describe the *coherences* between these states. The state expressed by a density operator is a pure state is the trace $\rho^2 = 1$ and a mixed state if the trace of the

density matrix squared is less than 1.

To describe unitary dynamics with the density matrix formalism, we use the Liouville-von Neumann equation

$$i\hbar\frac{d}{dt}\rho = [H, \rho] \tag{2.6}$$

where $[H, \rho] = H\rho - \rho H$. When dealing with non-unitary dynamics, an extra term can be added that governs the non-unitary evolution. The typical case is a Lindblad super-operator that acts on the density matrix

$$i\hbar\frac{d}{dt}\rho = [H, \rho] + \mathcal{L}(\rho). \tag{2.7}$$

A simple application of the Lindblad super-operator is to model *dephasing* by applying an exponential decay to the coherences of the density matrix [13].

We will not go into further detail on the theory of open quantum systems and non-unitary dynamics (for more details see [37]) but we will mention that in general, if we take many pure states that are evolved unitarily under different Hamiltonians, then compare a statistical mixture of the initial states and the final states, the evolution of the statistical mixture will have been non-unitary [20]. This can be shown by considering that the stochastic average as the addition with relevant weighting of density matrices, eg. $\rho = \frac{1}{2}\rho_1 + \frac{1}{2}\rho_2$. If ρ_1 and ρ_2 are the individual results from two different unitary evolutions of the same pure state ρ_0, both ρ_1 and ρ_2 are known to be pure states. Consider the example where ρ_1 has all of the population in state 1 and ρ_2 has all of the population in state 2;

$$\rho_1 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \rho_2 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad \rho = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}. \tag{2.8}$$

It is trivial to see that the trace of ρ^2 is less than 1 and is thus a mixed state.

In this thesis, we will deal with the case where the different Hamiltonians are different realisations of a stochastic process.

2.2 Rotating Wave Approximation

Calculating the dynamics of closed quantum systems is much simpler when the Hamiltonian is explicitly time independent. Sometimes, when faced with time dependent Hamiltonians, a procedure known as the Rotating Wave Approximation (RWA) can be used to remove the time dependence under certain conditions [34].

Consider a system consisting of two states; $|1\rangle$ and $|2\rangle$, separated in energy by ΔE, coupled by a sinusoidal potential of angular frequency ω and strength V. The Hamiltonian of such a system is given by

$$H = \begin{pmatrix} 0 & V\cos(\omega t) \\ V\cos(\omega t) & \Delta E \end{pmatrix}. \tag{2.9}$$

We utilise a change of basis to one where state $|2\rangle$ accumulates phase at the same frequency as the coupling field $|2'\rangle = e^{i\omega t}|2\rangle$. This change of basis can be achieved by multiplication of our state vector by a unitary matrix U to get from the old basis to the new basis and multiplication by its inverse to get from the new basis to the old. In this case, U is given by

$$U = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\omega t} \end{pmatrix} \quad U^{-1} = \begin{pmatrix} 1 & 0 \\ 0 & e^{-i\omega t} \end{pmatrix}. \tag{2.10}$$

Starting from the Schrödinger equation in the old basis

$$i\hbar\frac{d|\psi\rangle}{dt} = H|\psi\rangle \tag{2.11}$$

and then substituting $|\psi\rangle = U^{-1}|\psi'\rangle$ yields,

$$i\hbar\frac{dU^{-1}}{dt}|\psi'\rangle + i\hbar U^{-1}\frac{d|\psi'\rangle}{dt} = HU^{-1}|\psi'\rangle. \tag{2.12}$$

Rearranging forms the Schrödinger equation in the new basis

$$i\hbar\frac{d|\psi'\rangle}{dt} = H'|\psi'\rangle \tag{2.13}$$

where

$$H' = U^{-1}\left(H - i\hbar\frac{dU^{-1}}{dt}\right)U \tag{2.14}$$

which is explicitly

$$H' = \begin{pmatrix} 0 & \frac{V}{2}(1+e^{i2\omega t}) \\ \frac{V}{2}(1+e^{-i2\omega t}) & \hbar\Delta \end{pmatrix} \tag{2.15}$$

where $\Delta = \Delta E/\hbar - \omega$ and is typically referred to as the detuning. Everything that we have done has been exact up to this point. Here, we make the assumption that the dynamics of the system are very slow compared

to the 2ω exponential terms. If this is the case, we can neglect these oscillations. This leaves a new, time independent RWA Hamiltonian

$$H_{RWA} = \begin{pmatrix} 0 & V/2 \\ V/2 & \hbar\Delta \end{pmatrix} \tag{2.16}$$

which is much more straightforward to solve. In this two-level example, the coupling induces Rabi flopping between the states.

2.3 Bloch Sphere Representation

A very commonly used visualisation for two-level systems is the Bloch sphere. Due to a homomorphic transform between the mathematical groups that governs rotations in three dimensions (the SO(3) group) and those which consist of norm-preserving Unitary transforms for 2-level systems (the SU(2) group) [38], it is possible to describe the application of such operators on a state vector in terms of rotations in a 3-dimensional space.

The result of this mapping is the Bloch Sphere, which is a sphere of unit radius. A 2-level pure quantum state is described by a vector pointing from the origin to the surface of the sphere. The vectors pointing to the poles are pure state vectors that only occupy one of the basis states. We will adopt the convention that 'up' represents occupying state $|0\rangle$ and 'down' the occupation of state $|1\rangle$ when discussing 2-level systems. Points along the x-y plane consist of an equal superposition of the 2 basis states, with the angle in the x-y plane determining the relative phase between them. This mapping is shown in Figure 2.1 with a pure state fully defined by the angles ϕ and θ as $|\psi\rangle = \cos(\theta/2)\,|1\rangle + \sin(\theta/2)\exp(-i\phi)\,|2\rangle$. The application of any unitary operator on the 2-state basis can be represented by a rotation about a corresponding axis on the Bloch sphere. Any unitary 2×2 matrices can be generated according to

$$U = \exp\left(i\theta \boldsymbol{n} \cdot \boldsymbol{\sigma}\right) \tag{2.17}$$

where $\boldsymbol{n}$ is a 3D unit vector and $\boldsymbol{\sigma}$ are the three (scaled) Pauli matrices given by

$$\sigma_x = \frac{1}{2}\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \frac{1}{2}\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{2.18}$$

The rotation resulting from the application of the unitary matrix U is of magnitude θ around the axis specified by the unit vector n. The generating term $n \cdot \sigma$ can produce every 2×2 Hermitian matrix with

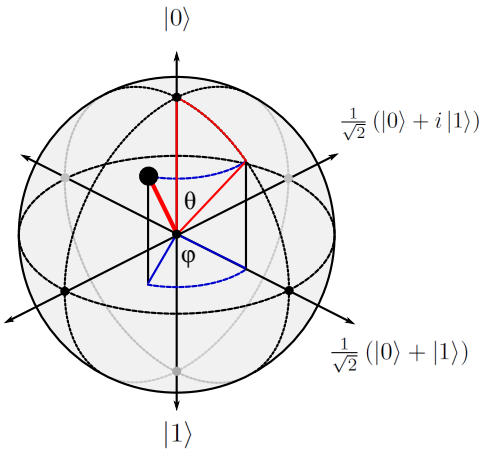

Figure 2.1: Typical Bloch sphere representation. A two-level pure state is fully characterised by a vector from the centre to the surface of the sphere. The two angles θ and ϕ fully characterise the pure state.

determinant -1; multiply the matrix by a constant and we can produce any 2×2 Hermitian matrix. Now consider the general solution of the Schrödinger equation for a time independent Hamiltonian

$$|\psi(t)\rangle = \exp(iHt)\,|\psi(t=0)\rangle = U(t)\,|\psi(t=0)\rangle\,. \tag{2.19}$$

We can see that the application of a time independent Hamiltonian leads to rotation of the state vector in time around the axis given by the Hamiltonian decomposition into the Pauli matrices and with angular frequency given by the determinant of the Hamiltonian. The eigenstates of a Hamiltonian are then the states associated with vectors aligned and anti-aligned to the Hamiltonian rotation axis. Unitary changes of basis are characterised by a rotation of the principal axes used to describe the sphere co-ordinates. In fact, the unitary quantum dynamics of a two-level system are equivalent to the precession of a classical magnetisation vector in the presence of a magnetic field [39].

It should be noted that the Bloch sphere is also useful for the representation of mixed states. Vectors that are contained inside the volume of the Bloch sphere but do not reach the surface are these mixed states. A density matrix formalism can be made with the vector as follows

$$\rho = \frac{1}{2}\left(1 + \boldsymbol{r}\cdot\boldsymbol{\sigma}\right) \tag{2.20}$$

where $\boldsymbol{r}$ is the Bloch vector. The most important relative metric for this thesis is the projection of the Bloch vector in the z-axis which represents the relative populations.

2.3.1 Three-level Bloch Sphere

Whilst Bloch spheres are ubiquitous in the description of two-level systems, they can be partially generalised to systems with more levels. The group of all Hermitian matrices for an N-level system are in general characterised by the group SU(N), which means for $N > 2$, we lose the holomorphic map to SO(3). However, if we consider only the application of Hermitian matrices that are representations of the SU(2) group in the N-dimensional vector space, we preserve this mapping. These matrices are a subgroup of the SU($N > 2$) matrices but manifest themselves quite commonly.

Much of the work in this thesis involves a three-level system composed of angular momentum projection states. For these states, the application of a magnetic field along an axis gives rise to Hamiltonians that preserve this SU(2) nature. This means we can use the intuition, and expect much of the same behaviour as in the two-level case for these systems such as the manifestation of Rabi flopping and Ramsey interferometry. The generators for the three-dimensional representation of SU(2) that we shall use are [38]

$$\sigma_x = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad \sigma_y = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & i & 0 \\ -i & 0 & i \\ 0 & -i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \tag{2.21}$$

This establishes a map from a vector on the Bloch Sphere to a three-level representation as long as the Hamiltonian can be composed of a linear superposition of these three matrices.

In the 2-level case, every possible pure state can be represented by a vector on the same Bloch sphere. This is not true for the 3-level system as we have limited our Hamiltonians to those expressible using SU(2) representations. What this means is that by choosing a particular state to be on a 3-level Bloch sphere surface, the other points on this Bloch sphere are only those accessible from this chosen initial state using the limited Hamiltonian space. Therefore, there are an infinite number of choices of Bloch sphere, simply by choosing a different starting state and none of these different choices of sphere *overlap* in the sense that a state vector cannot belong to two different Bloch sphere representations (assuming we consider rotated Bloch spheres essentially the same).

For the remainder of this thesis, we consider only one of these infinite choices. Using the state labellings $|-1\rangle$, $|0\rangle$ and $|+1\rangle$, choosing the $|-1\rangle$ state vector as the starting vector, we trace out a Bloch sphere that has this $|-1\rangle$ state at the north pole, the $|+1\rangle$ at the south pole and the equator constituting superposition states of the form $\exp(i\phi)\,|-1\rangle + \sqrt{2}\exp(i\phi/2)\,|0\rangle + |+1\rangle$[1] where ϕ is the angle from the x axis in the x-y

[1] Normalisation excluded for clarity

plane. All of these equatorial states correspond to relative populations (1/4, 1/2, 1/4).

Whilst the ability to map our three levels onto a Bloch sphere aides in the intuitive understanding of what can otherwise be a very complicated system, in the end, it alone does not add any more effective complexity to the system. In future work (see Chapter 8), we aim to make Hamiltonians that do not satisfy this SO(3) holomorphism.

2.4 Rubidium 87 Structure

Neutral Rubidium is an atomic element with a single unpaired electron in its outer shell. As such there is little interaction between this outer electron and the other electrons. The inner cloud effectively screens the positive nucleus. The energy level structure is very similar to Hydrogen and the other alkali metals.

2.4.1 Fine and Hyperfine Structure

Within the principal quantum number ($n = 5$) states, the coupling of the electron's orbital angular momentum (denoted $\boldsymbol{L}$, quantum number L) and intrinsic spin angular momentum (denoted $\boldsymbol{S}$, quantum number S) leads to a 'Fine' structure splitting in the energy levels proportional to the term $\boldsymbol{L}.\boldsymbol{S}$. The split levels are labelled further with the quantity $\boldsymbol{J} = \boldsymbol{L} + \boldsymbol{S}$ with quantum number J which represents the total electronic angular momentum. The vector addition of the electronic orbital and spin angular momenta results in the allowed values for J satisfying $|L - S| \leq J \leq |L + S|$. For $L = 0$, $J = 1/2$ is the only allowed value resulting in the $5^2S_{1/2}$ ground state[1]. For $L = 1$, J can equal 1/2 or 3/2. The $J = 1/2$ excited state is denoted as $5^2P_{1/2}$. The transition between this state and the ground state is commonly called the $D1$ line. For $J = 3/2$, the state is denoted $5^2P_{3/2}$ with the transition to the ground state forming the $D2$ line.

Further splitting in the energy levels results from the interaction between the electronic angular momentum and the angular momentum of the nucleus (denoted $\boldsymbol{I}$, quantum number I). Similar to the spin orbit coupling above, the magnitude of this splitting is proportional to the quantity $\boldsymbol{J}.\boldsymbol{I}$. Similar arguments to above lead to total atomic angular momentum $\boldsymbol{F} = \boldsymbol{J} + \boldsymbol{I}$ with quantum number F, where the allowed values are $|J - I| \leq F \leq |J + I|$. In Rubidium 87, $I = 3/2$ and thus the splitting results in a more complicated 'Hyperfine' structure. Rather than exhaustive enumeration of the resulting states, the reader's attention is directed to Figure 2.2.

Discussion so far has revolved around coupling of the magnitude's of the relevant angular momenta. Since the particles in possession of the angular momentum are electrically charged, they form magnetic moments

[1]The notation is typical spectroscopic notation $n^{2s+1}L'_J$ where L' is S, P, D for $L = 0, 1, 2$

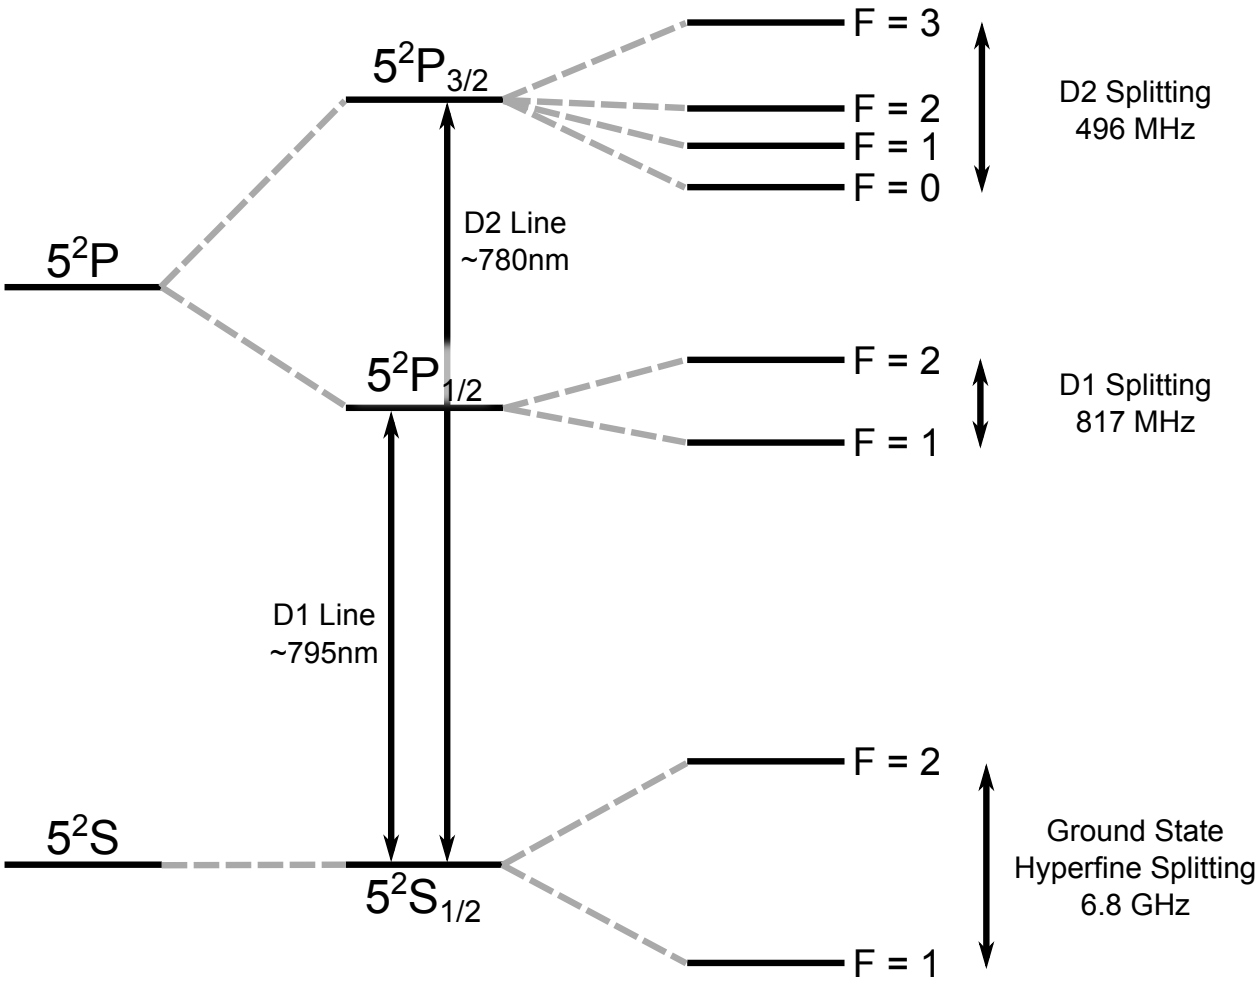

Figure 2.2: Energy level structure of Rubidium 87 relevant to this thesis. The orbital angular momentum of the P orbital couples with the electron spin to produce Fine Structure splitting into the D1 and D2 lines. Interaction between the electronic and nuclear angular momenta produces the Hyperfine Structure splitting. Diagram not to scale.

that interact with external magnetic fields according to their projection along the external field direction. Quantum mechanically, the projection of an angular momentum along a particular axis is quantised. Using standard convention and considering the projection along the z-axis, the projection of angular momentum F_z is given by $m_F\hbar$ where $-F \leq m_F \leq F$ and m_F is integer spaced. This means that each F state is composed of $2F+1$ sub-levels. In the presence of a small magnetic field, the energy degeneracy of these states is lifted according to the linear Zeeman shift equation $\Delta E = \mu_B g_F m_F B_z$ for a magnetic field applied along the z-axis B_z, where the constant μ_B is the Bohr magneton and g_F is the gyromagnetic ratio for the state.

Polarisation Selection Rules

As the angular momentum projection quantum number m_F governs the quantity of angular momentum along the local z-axis, transitions between states with different m_F's must be governed by the conservation of angular momentum. For electric dipole transitions between levels induced by external fields, any change in angular momentum projection must be supplied by the perturbing field. For electromagnetic fields, the angular momentum is governed by the behaviour of the field vector in time, i.e. its polarisation. Field vectors whose magnitude is constant but direction rotates as a function of time are circularly polarised. The direction of rotation determines the handedness of the polarisation (left or right). The photons constituting a circularly polarised field carry $1\hbar$ of angular momentum along the axis of rotation. If this axis of rotation is along the local z-axis of the atom, the absorption of a photon can change the m_F of the state by ± 1 depending on the handedness of the polarisation. Light that, on absorption, induces a transition where $m_F \rightarrow m_F + 1$ is referred to as σ^+ light and $m_F \rightarrow m_F - 1$ as σ^- light. Note that a linearly polarised field (one in which the field direction is constant but whose amplitude oscillations in time) can be considered as a equal superposition of two circularly polarised fields. Linearly polarised fields that oscillate in the $x-y$ plane are capable of exciting both σ^+ and σ^- transitions.

Transitions can also be excited where the m_F value is unchanged. An electromagnetic field that does this must contain some π polarisation. This is where the field has some component in the z-axis. Note that a beam of light in free space propagating along the z-axis cannot contain π polarisation as the propagating field vector must be perpendicular to the propagation axis.

2.4.2 Rubidium 87 3-Level System

To attempt to produce a closed 3-level system for use in quantum simulations from a neutral Rubidium 87 atom, we must consider 3 internal states that can be manipulated and coupled coherently together but

are stationary in time without manipulation. For this, we consider the $5^2S_{1/2}$, $F = 1$ ground state and its associated magnetic projection sublevels. This constitutes three states that we shall henceforth refer to as $|-1\rangle$, $|0\rangle$ and $|+1\rangle$ with the label corresponding to their m_F quantum numbers. We shall assume that with no coupling to other levels, they form a complete basis of three states.

In the absence of external fields, the states are degenerate and have no coupling. In the presence of a small bias magnetic field B_b in the z-axis, the Hamiltonian is modified by the linear Zeeman shift of the energy levels:

$$H = \mu_B g_F \begin{pmatrix} -B_b & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & B_b \end{pmatrix}. \tag{2.22}$$

Application of a magnetic field along an different axis can be easily expressed using a linear combination of the SU(2) generators in the 3D representation (see Section 2.3). i.e. The application of a field $\boldsymbol{B} = B_x\boldsymbol{x} + B_y\boldsymbol{y} + B_z\boldsymbol{z}$ can be expressed with a Hamiltonian $H = \mu_B g_F(B_x\boldsymbol{\sigma_x} + B_y\boldsymbol{\sigma_y} + B_z\boldsymbol{\sigma_z})$.

We can see that in the chosen basis, a static magnetic field in the $x - y$ plane will produce the desired coupling between the states, however it is experimentally very difficult to turn a static field on and off quickly enough to prevent adiabatic following of the spins to the new field. This makes state preparation and readout in the original basis very difficult. Instead, we apply a sinusoidally oscillating magnetic field in the $x - y$ plane. Here, we apply the field $B_c\cos(\omega_{RF}t + \phi)$ in the x-axis. This results in the following Hamiltonian;

$$H_c = \mu_B g_F \begin{pmatrix} -B_b & \frac{B_c}{\sqrt{2}}\cos(\omega_{RF}t + \phi) & 0 \\ \frac{B_c}{\sqrt{2}}\cos(\omega_{RF}t + \phi) & 0 & \frac{B_c}{\sqrt{2}}\cos(\omega_{RF}t + \phi) \\ 0 & \frac{B_c}{\sqrt{2}}\cos(\omega_{RF}t + \phi) & B_b \end{pmatrix}. \tag{2.23}$$

This time dependent Hamiltonian can be simplified under certain assumptions in the Rotating Wave Approximation (RWA) as discussed in section 2.2. We create a new Hamiltonian H' using Equation (2.14) with the judicious choice of U as

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\omega t} & 0 \\ 0 & 0 & e^{i2\omega t} \end{pmatrix}. \tag{2.24}$$

Following the procedure in Section 2.2 leads to the equation for the time evolution of the system in this

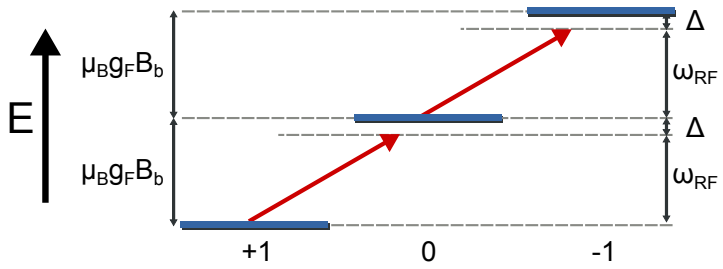

Figure 2.3: Diagram showing the radiofrequency coupling of the magnetic projection sub-levels. The degeneracy of the states is lifted by the bias field B_b and then the states are coupled by an AC magnetic field with angular frequency ω_{RF}.

rotating wave picture

$$i\hbar \begin{pmatrix} \dot{c}_{-1} \\ \dot{c}_0 \\ \dot{c}_1 \end{pmatrix} = \hbar \begin{pmatrix} -\Delta & \frac{\Omega_{RF}}{\sqrt{2}} & 0 \\ \frac{\Omega_{RF}}{\sqrt{2}} & 0 & \frac{\Omega_{RF}}{\sqrt{2}} \\ 0 & \frac{\Omega_{RF}}{\sqrt{2}} & \Delta \end{pmatrix} \begin{pmatrix} c_{-1} \\ c_0 \\ c_1 \end{pmatrix} \tag{2.25}$$

where we have introduced the detuning of the AC field frequency from the energy level splitting; $\Delta = g_F \mu_B B_b - \hbar\omega_{RF}$ and the Rabi frequency $\hbar\Omega_{RF} = g_F \mu_B B_c/2$ which characterises the coupling strength between the levels. The validity of the RWA is maintained only when $2\omega_{RF}$ is significantly larger than the frequencies corresponding to the remaining dynamics of the system (Ω_{RF} and Δ). Figure 2.3 shows the energy level structure of the un-coupled states with the coupling field.

With this in place, it is possible to investigate very familiar coherent schemes such as Rabi oscillations, adiabatic transitions and Ramsey fringes. Investigations into these different areas were conducted and their results are presented in Chapter 6.

The application of time dependence to the diagonal terms to simulate decoherence processes can then be added simply by making the original z-axis bias field time dependent. This has the effect of making Δ time dependent as well. The results of these procedures are presented in Chapter 7.

CHAPTER 3

COLD ATOMS CLOUD CREATION

At the core of the overall experiment is the requirement for the robust and repeatable creation of cold atom clouds. The cooling and trapping of neutral atoms using lasers is a well established field with many excellent sources to draw from [35, 40–42]. Here, we briefly describe the theory of atom cooling and trapping before detailing the experimental apparatus and techniques that were used in the creation of the cold atom clouds in this experiment. We then produce typical characteristic measurements of the clouds including information on their temperature and number. The creation of all of this apparatus was conducted as work for this Ph.D.

Very briefly, the cooling and trapping of neutral atoms makes use of the average momentum transfer from beams of light to an atom in the act of scattering photons. When an atom absorbs a photon, conservation of momentum requires that the atom receives a momentum kick of $\hbar k$ (where k is the wavenumber of the photon) in the same direction as the photon's trajectory. On spontaneous emission of a photon, the atom receives a kick of $\hbar k$ in the opposite direction of the emitted photon's trajectory. As the photon emission direction is random, the average momentum change over many emissions is zero. Successive absorption-emission cycles therefore leads to a net momentum change in the direction of the incoming photons, referred to as a radiation pressure force.

Slowing of an atom can be achieved using a pair of counter-propagating beams with a frequency red-detuned (lower frequency) than the target atomic transition. An atom with zero velocity along the axis of the beams will scatter photons equally from the two beams and the two radiation pressure forces will cancel out having a net zero effect, leading to no net force. However, due to the Doppler effect, an atom moving towards one of the beams will observe the frequency of that beam to be higher, and thus closer to resonance compared to the other beam, which will be observed at a lower frequency. This atom will therefore

preferentially absorb photons from the beam it is travelling towards, leading to a net momentum transfer that opposes the direction of travel. This produces a slowing force that, at sufficiently small velocities, has a magnitude proportional to the atom's velocity. Since the temperature of a gas of atoms is related to their speed ($\frac{1}{2}k_BT = \frac{1}{2}m\langle v\rangle^2$), the reduction in speed of a group of atoms that are then allowed to thermally equilibrate with each other can be thought of as a process that reduces a gas's temperature, or a cooling mechanism.

As cooling is achieved exploiting the Doppler effect to produce a differential scattering probability between two beams of light dependent on the atomic velocity, a similar effect can be exploited for trapping by creating a differential scattering probability dependent on the position of an atom. The linear Zeeman effect is used in conjunction with a magnetic field gradient to shift the energy levels of the atoms dependent on their position in the field. The two beams of light are chosen with circular polarisations such that only select transitions can be excited. These two elements together mean that, when the magnetic field is finite, the different allowed transitions induced by each of the two beams are of different frequencies, creating differential scattering. With the direction of the field inverted, the sign of the differential scattering changes. By setting the polarisations 'correctly' with respect to the field orientation, the result is a force that drives the atoms towards magnetic zero.

The combination of this velocity dependent cooling force and the position dependent trapping force work together to create a Magneto-Optical Trap (MOT) which is the workhorse of neutral atom cooling and trapping. Adding pairs of beams in different directions produces cooling in multiple dimensions. The addition of an anti-Helmholtz coil pair creates a magnetic zero at which the atoms are also trapped by the light pressure force.

In this project, we utilise a very typical scheme, starting with the loading of a 3D MOT from a vapour of Rubidium 87 atoms inside a vacuum chamber. Rubidium is a very commonly used species in cold atom systems due to its uncomplicated cooling scheme and the ready availability of lasers at its D2 transition wavelength (780nm). 3D vapour MOT's are typically used to trap from a few million to billion atoms and cool them to the order of 100 µK.

The remainder of this chapter outlines the experimental set-up used to trap and cool Rubidium 87 clouds followed by a discussion on the methods used to image the clouds and certain techniques used to characterise the process.

3.1 Optical System

Two frequencies of light are typically required in standard laser cooling applications. The first is the 'Cooler', which is red shifted a few linewidths from a cycling transition ($F = 2 \rightarrow F = 3$ in Rubidium 87) and is responsible for the cooling and trapping forces on the atom. Due to the finite linewidth of the excited hyperfine states, the cooling beam has a small probability to excite the atom to the $F = 2$ excited state, from which there is a finite probability to decay into the $F = 1$ ground state which is dark to the cooling beam and therefore the atom would be lost from the captured sample. The second frequency is the 'Repumper', which is set to the $F = 1 \rightarrow F = 2$ transition and provides a path for any atoms suffering this fate to be restored to the $F = 2$ ground state.

In our setup, we used three different lasers to generate all of the beams required to cool, trap and image the atoms. These lasers are named the 'Master', 'Slave' and 'Repumper'. One of the initial design goals of the experiment was to minimise its general footprint. As a consequence, most of the construction of the optical setup is confined to three aluminium breadboards (two of which measure 61 cm $\times$ 30.5 cm and one 45 cm $\times$ 30.5 cm). Each breadboard contains one of the lasers and the relevant optics (The Repumper occupies the smaller board). Light is transferred between the boards using optical fibres which allows the boards to be moved individually without complete misalignment of the entire system. The eventual design iteration was to stack the boards, mounted onto a small rack unit. As the experiment progressed, the required optical system complexity increased and this goal was eventually abandoned for the sake of practicality. A small number of the components are now mounted on the optical table in the near proximity to the boards. All of the optics related to the creation of a MOT are however still mounted on the boards. This makes the entire MOT setup fairly transportable with minimal realignment of the optics required.

In what follows, we describe the optical set-up pertinent to the creation of the cold atom cloud sample used in the rest of the experiment. A schematic diagram showing the general sources of the different beams and relation between the lasers is shown later in Figure 3.3. Chapter 5 discusses a small number of additions used for the sample state preparation. We begin our discussion with the lasers.

Bare laser diodes constituted of a gain material inside of a simple Fabry-Perot cavity are generally unsuitable for laser cooling due to their large linewidths ($\sim$100 MHz) To establish the differential radiation forces for cooling and trapping, the laser linewidths should be smaller than the cooling transition's linewidth, which for the Rubidium D2 line is 6 MHz. We utilise Extended Cavity Diode Lasers (ECDLs) [43, 44], which are bare diode lasers with an optical diffraction grating mounted in the output beam path such that one of the diffraction orders of the grating reflects back into the diode. The effect of this is to produce frequency

selective optical feedback in the diode cavity that serves to narrow the linewidth via the angularly dependent order of the grating and the formation of a free space optical cavity between the laser diode front facet and the grating. The grating is typically mounted with a Piezo-Electric Transducer (PZT) to allow the tuning of the angle of the grating as well as the length of the external cavity. This tuning allows for control of the optical frequency of the feedback and thus the frequency of the diode. We use ECDLs in the Littrow configuration in which the light from the minus first order of the diffraction grating is reflected back into the laser diode to provide optical feedback and the specular reflection is used as the output of the laser. This optical feedback can reduce the linewidth to less than 1 MHz.

Along with the reduced linewidth, reliable and long-term operation of the MOT lasers requires that their frequency does not drift in time. A variety of environmental and electronic fluctuations can cause long and short-term frequency fluctuations. These fluctuations can be combated using a locking system with a frequency dependent signal and feedback loop on the grating position of the ECDL using the PZT. There are a variety of methods used to create the frequency dependent signal which are outlined for each of the lasers below. Electronic feedback is achieved using home built PID lock-boxes.

Master Laser

The Master laser is a home-built ECDL locked to the Rubidium 87 $F = 2 \rightarrow F = 3$ transition using Polarisation Spectroscopy broadly following Pearman et al. [45]. Polarisation spectroscopy relies on circular birefringence induced in a vapour cell by population imbalance created by a circularly polarised pump beam counter-propagating to the probe beam. The resultant rotation of the probe's linear polarisation is detected using a Polarising Beam Splitter (PBS) and difference photodiode detector. The induced circular birefringence has a zero-crossing on the excited hyperfine transitions with the strongest signal given for the $F = 2 \rightarrow F = 3$ transition. Whilst the relatively simplicity of the optical setup is desirable, issues were encountered from residual magnetic fields affecting the signal. The situation was improved by placing the vapour cell inside a mu-metal shield which improved the shape of the signal but experimentation on the MOT later showed a slowly moving offset of the locking point from the $2 \rightarrow 3$ transition of 2 to 5MHz. This offset was calibrated and compensated in the experimental control.

Feedback on the Master laser is enacted using both the grating PZT and on the current of the diode. The ECDL utilises a Thorlabs L785P090 diode which has a natural lasing wavelength of 785 nm which is pulled by the grating to operate at 780 nm. The laser is operated with an output power from the grating of $\sim$30 mW. The light is used for referencing the Repumper and Slave lasers (discussion upcoming) as well as

deriving a variable frequency beam used in the absorption imaging system (see section 3.4).

Repumper Laser

The Repumper laser is a home-built ECDL locked close to the $F = 1 \rightarrow 2$ transition using a simple beat lock with the Master laser that is a slight simplification of the method used by Schünemann et al [46]. In a beat lock, the frequency difference between two lasers is stabilised using the beat signal between the light from the two lasers when overlapped onto a photodiode. Figure 3.1 shows the schematic diagram of this locking scheme.

The two beams are overlapped with the same polarisation onto a biased Hamamatsu G4176-03 photodiode. A bias-tee (Minicircuits ZX86-12G-S+) removes the DC signal from the beat note AC signal, which is amplified using three low power amplifiers (Minicircuits ZX60-8008-E) before being split with a splitter (Minicircuits ZX10-2-722-S+) and then recombined in a mixer (Minicircuits ZX05-73L-S+) to form a radiofrequency interferometer. One of the arms of the interferometer uses a long cable (∼2.5 m) as a delay line. The resultant signal at DC from the mixer is proportional to the cosine of the phase difference of the beat frequency between the arms of the interferometer. The phase delay induced by the fixed length of the

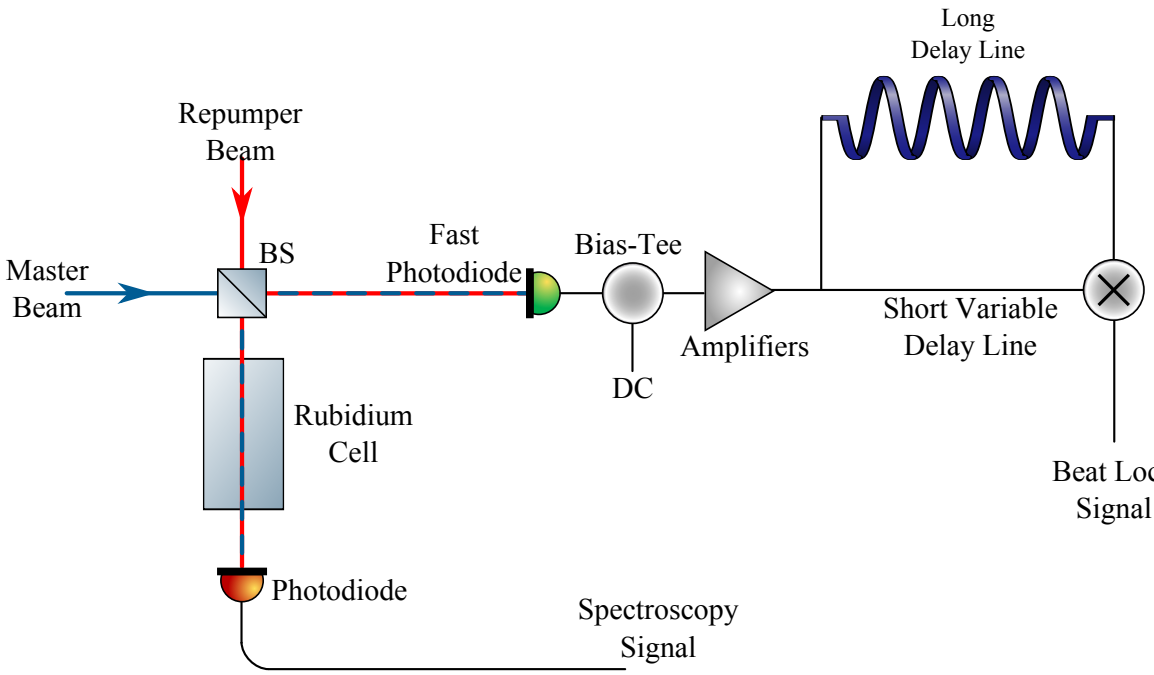

Figure 3.1: Schematic of the Master-Repumper beat lock. This stabilises the difference frequency between the Master and Repumper lasers to the hyperfine ground state splitting of 6.8 GHz.

delay line is proportional to the frequency of the beat note $\phi = \frac{2\pi}{v} f \Delta L$ (with f the beat note frequency and v the speed of the RF in the coaxial cables). This results in a signal that oscillates as the Repumping laser's frequency is swept with the Master locked. With a large number of lock points to chose from, the oscillating lock signal was supplemented with a very simple spectroscopy setup of the overlapped Master and Repumper beams co-propagating through a Rubidium vapour cell and onto a photodiode. Without the Master beam included, sweeping the Repumper simply yields the Doppler broadened $F = 1 \rightarrow F'$ manifold absorption profile, but with the inclusion of the Master beam, the atoms with zero velocity class are pumped by the Master beam into the $F = 1$ ground state, leading to increased absorption for the Repumper beam when it is resonant with this velocity class of atoms. Figure 3.2 shows the signal with and without the Master beam next to the signal from the beat locking apparatus. This gives a simple guideline for chosing which zero point to lock to. The position of the zero crossings of the locking signal can be tuned using a variable delay line in one of the interferometer arms.

Feedback on the repumper is enacted using just the PZT on the grating. The Repumper laser also uses a Thorlabs L785P090 diode and is also operated at 30 mW output power. The light from the repumper not used in the beat lock is used to create an optical pumping beam (see Chapter 5) and for the solitary repumping beam in the MOT. This beam goes through a focus where a physical shutter is situated to allow for beam switching, after which it expands slowly to the chamber where it is approximately 25 mm in diameter.

Slave Laser

As the light from the Master laser is used for its own polarisation spectroscopy lock, the Repumper beam lock and deriving the imaging beam, there is insufficient power to additionally provide the large amount of light required for the cooling beams. To provide this, we use a small amount of power from the master laser to reference a 'Slave' laser. The Slave laser is a bare laser diode injection locked using light shifted up $\sim 2 \times 75$ MHz in frequency from the Master laser with a double pass Acousto-Optic Modulator (AOM). Injection locking is a technique which uses some seed light from a different source sent into the cavity of the laser to be referenced. Like the optical feedback introduced by the grating in an ECDL, this saturates the laser diode gain at the seeded frequency. The diode used for the slave laser is a Roithner LaserTechnik RLT780-150GS with 150 mW nominal power.

The Slave laser's role is to provide the light for the cooling and shift the frequency during sub-Doppler cooling. The elliptical beam emitted from the laser is passed through an anamorphic prism pair to make the beam more circular to improve the AOM and fibre coupling efficiency. It then passes through a Thorlabs

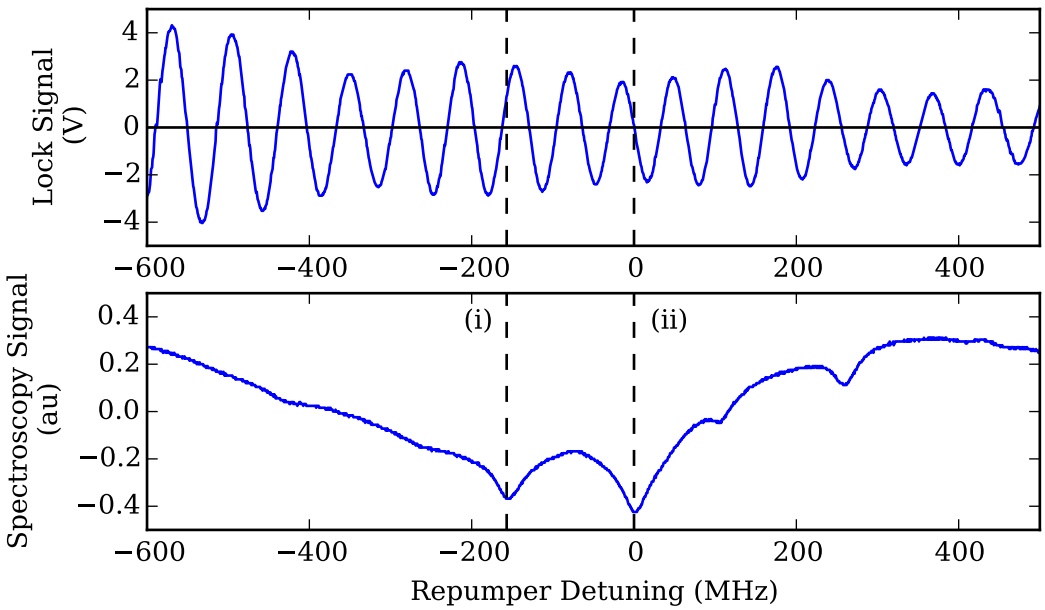

Figure 3.2: Relevant traces of signals used to lock the Repumper, The upper trace shows the output from the beat lock apparatus that is fed into the PID control loop to lock the frequency. The variable delay line must be changed to shift this signal such that a zero crossing occurs at the desired lock point. The lower trace shows the co-propagating spectroscopy setup. As the Master beam is locked to the $F = 2 \rightarrow 3$ transition, atoms with certain velocities are optically pumped to the F = 1 states. When the scanning repumper beam is resonant with a transition of one of these velocity classes, we see enhanced absorption dips. (i) is the line associated with the repumper being resonant with the $F = 1 \rightarrow 1$ transition for the zero velocity class. (ii) is the line associated with the repumper being resonant with the $F = 1 \rightarrow 2$ velocity class. The other enhanced absorption dips are different velocity classes being pumped.

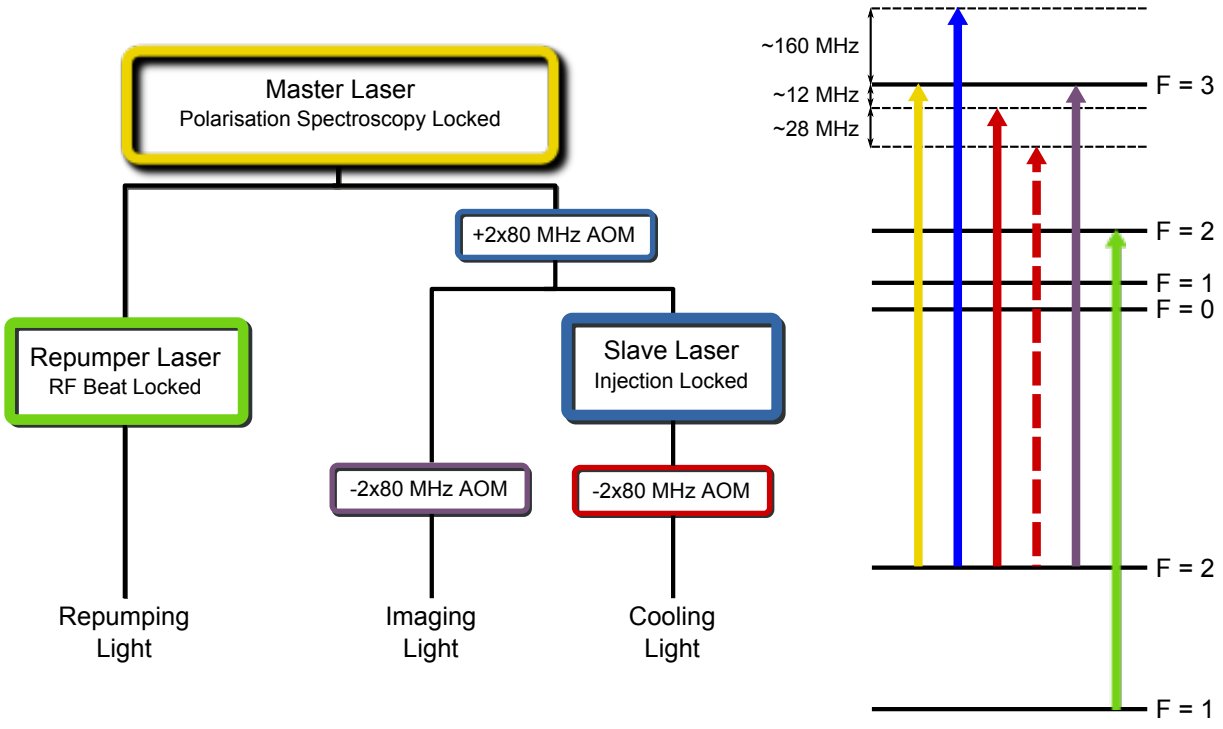

Figure 3.3: Schematics of the lasers and derived beams used in the cooling and trapping of the atoms.

IO-3-780-HP optical isolator, and is frequency down-shifted by $\sim 2 \times 80\,\mathrm{MHz}$ using an AOM and distributed to the MOT beams using optical fibres. The down-shift is actually achieved using two separate AOMs, one feeding two of the MOT beams and the other the remaining four. This allows separate frequency control of these two beams with respect to the others. These two beams were used as the diagonally-upward facing beams for use in an atomic fountain. This was to facilitate experimentation done on a separate branch of research from the main topic of this thesis (for more information of this alternative branch, please see Chapter 9).

At the operating point of the Slave, the optical power after the initial collimation lens is 70 mW. Although significant power is lost ($\approx$30 mW) at this lens due to the high divergence of the beam, the Slave is operated below its maximum power output due to the limitations of the current driver. After the several finite efficiency steps, the total amount of cooling power was about $\sim$18 mW. The injection locking process was monitored using the transmission of a picked off beam through a confocal Fabry-Perot cavity.

3.2 Vacuum System and Magnetic Gradient Coils

The vacuum system is composed of a Science Chamber (Kimball Physics MCF800-SphOct-G2C8) attached using a CF40 to CF16 conic section to a CF40 6-way cross. The cross is also attached to an ion pump (Agilent Technologies Vaclon Plus 20), valve (MDC MAV-150-V) and viewport (MDC CVP-40). The orientation of the viewport is such that it provides line of sight through the conic section to the science chamber. Rubidium is pumped into the chamber by passing current through a SAES Rubidium dispenser electrically connected via an electrical feedthrough (MDC HV1-15C-4-C16). The remaining ports on the science chamber had windows attached (MDC CVP-16 for the CF16 ports, CVP-63 for the CF63 windows).

Cooling beams are orientated through the two large windows of the chamber and through the four small windows 45° from the horizontal. Imaging is done through the horizontal axis through the chamber, conic section and 6-way cross. The small aperture in the conic section provides a limiting pupil for the imaging system which limits the size of the absorption imaging beam restrictively to approximately 1 cm diameter. This has implications for the ease of the state readout discussed in Chapter 4. A schematic diagram of the chamber with the beam orientations is shown in Figure 3.4.

The quadrupole magnetic field for the MOT is created using a pair of anti-Helmholtz coils situated around the large viewports of the chamber. The coils are each 90 turns with a radius of 7.5 cm. They are separated by 7.5 cm leading to a theoretical gradient of 1.7 G/cm A along the central coil axis and 0.85 G/cm A along the radial direction in between the coils. Typical experiments were done at a current of 6 A giving gradients of 10.2 G/cm and 5.1 G/cm accordingly, although this calculation does consider all of the turns to be in the ideal location and neglects the physical extent of the coils so the true gradient is likely lower as many of the turns will be slightly displaced from this position. Switching of the MOT coils is achieved using an Insulated-Gate Bipolar Transitor (IGBT) (International Rectifier IRG4PH50SPbF) driven by an IGBT driver chip (ON Semiconductor MC33153).

Cooling Beams

Of the six cooling beams, two propagate through the large windows in the horizontal plane of the lab, and the remaining four through the smaller windows 45° to the horizontal.

The two horizontal beams are collimated from the fibres using a 10 cm focal length lens to give a beam size approximately 2 cm in diameter. The remaining four beams are collimated using a $f = -25$ mm divergent lens follows by a $f = 40$ mm convergent lens to give a beam size of 15 mm diameter.

An unfortunate issue that was belatedly discovered was the poor anti-reflection coating that was applied

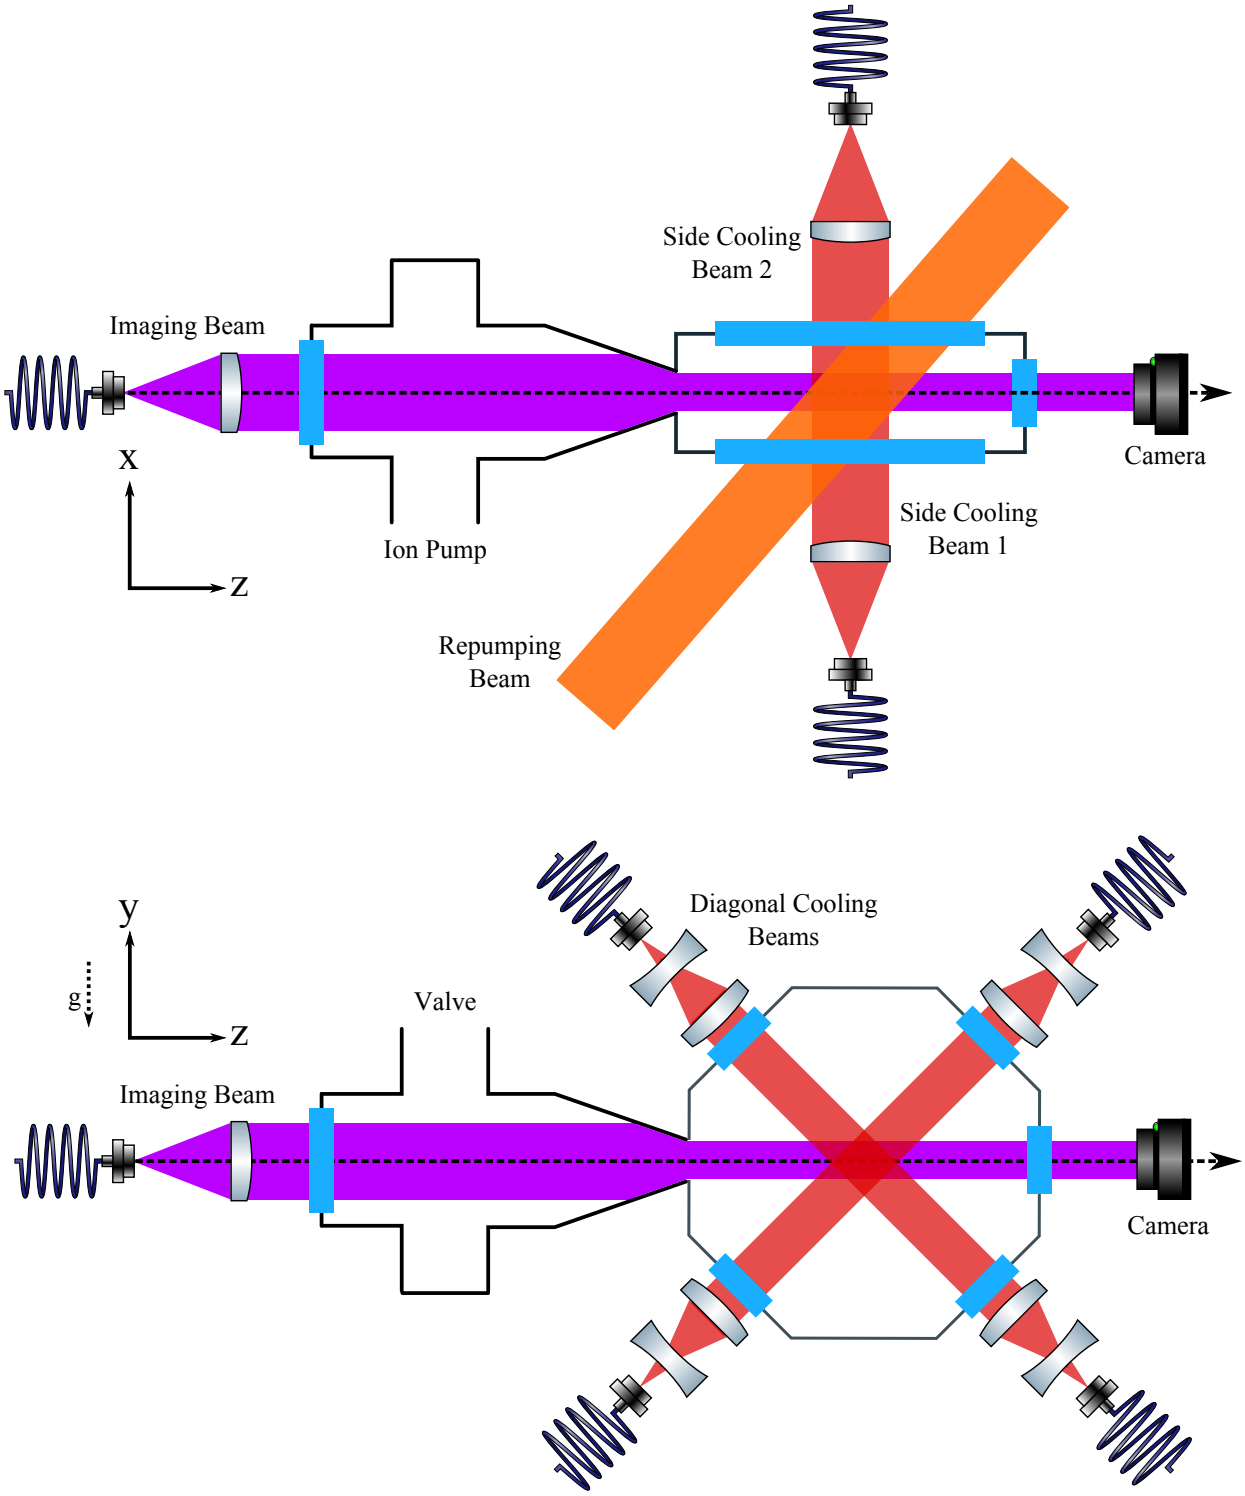

Figure 3.4: Schematic of the vacuum system and the beams used for cooling and trapping the atoms, shown in the z-x plane (top) and the z-y plane (bottom). Not to scale.

to the windows by the supplier. The custom coating was designed to provide very low reflections for both 780 nm and 1550 nm light. This was again to facilitate an alternate experiment (see chapter 9). The coatings were provided with specification testing sheets with agreeable characteristics but it was not discovered until more than three years into the project that each of the small windows had transmissions of approximately 30% and each large window a transmission of 70% at 780 nm. This severely reduces the amount of cooling power sent to the atoms. Before this was discovered, the smaller beam size for the small window beams warranted a lower power per beam than the horizontal beams. A frequent issue was the instability of the sub-Doppler cooling stage over long time frames. With the issue discovered, the beam powers were balanced such that the intensities after the windows were the same. This resulted in much better reliability for the cloud creation.

3.3 Experimental Control

Control of cold atoms experiments typically requires sub millisecond precision in the execution of commands. This is typically achieved using Field Programmable Gate Arrays (FPGAs) which are networks of digital gates that can be programmed to be connected in a particular way.

In our experiment, we utilise a Digilent Atlys board with a Xilinx Spartan-6 FPGA. The board includes a RAM module connected to the built-in Spartan-6 RAM controller which can greatly increase the storage capability of the device. The Atlys also has a built in USB controller chip with provided libraries that facilitate communication of data between the FPGA and a desktop computer.

The FPGA was programmed using Verilog and VHDL Hardware Description Languages (HDL) by the author before starting this Ph.D. work. The result is a device that provides 40 digital channels with 5 ns time resolution (or alternatively 20 digital channels at 2.5 ns time resolution). Data is sent to the FPGA in the form of a list of 'words' that describe the state of the digital channels and an integer number of 5 ns clock cycles, the length of which the state will be held for before moving onto the next word. The built in RAM controller has the facility to store 6 million of these words.

To interface with the experiment in a more intuitive way than manipulating the output states of digital channels directly, a set of Python layers were built which translate intuitive, user-written Python scripts which state the timings of broad *experimental events* such as, turning the MOT on, doing an optical molasses ramp and taking an absorption image. The underlying layers translate these instructions to the required bit pattern for the FPGA outputs as well as interface with other experimental components, the Direct Digital Synthesiser (DDS) that controls the AOM frequencies, the absorption imaging camera and the Arbitrary

Wave Generator (AWG).

3.4 Experimental Procedure

To create a cloud of atoms, the quadrupole coils (henceforth the MOT coils) are turned on and both the Repumper and Cooling light is turned on. The cooling light detuning from the $F = 2 \rightarrow 3$ transition is set using the Master laser AOM to give the maximum cloud atom number (12 MHz red detuned). In this MOT phase, atoms accumulate according to an exponential loading curve. Atoms in the MOT stage are cooled to approximately 100 µK. Further cooling is limited by the random walk of the atoms in momentum space related to to their scattering rate and transition photon momentum. This is called the Doppler Limit [35, 47]. Typical MOT temperatures can be below the Doppler temperature (which for Rubidium is 146 µK) owing to the presence of sub-Doppler cooling mechanisms. Cooling significantly below the Doppler limit requires a more controlled implementation of this same sub-Doppler mechanism.

For significant sub-Doppler cooling, we utilise polarisation gradient cooling. The quadrupole field is turned off before the cooling light detuning is increased to 40 MHz and power decreased, both ramped over 10 ms and then held for 1 ms. Cooling in this larger detuned regime is governed by the superposition of polarisations of the three pairs of counter-propagating beams. For the interfering circular polarisations in our setup, along a single pair of beams, this produces a rotating linear polarisation. Atoms travelling along this path experience a changing quantisation axis that produces a population imbalance from equilibrium. The light scattering that brings it back to equilibrium further cools the atoms [48, 49].

During this detuning ramp and hold, the atoms are cooled to just below 10 µK and ready for experimentation. Efficient polarisation gradient cooling requires very low magnetic fields where the atom cloud is. This is optimised using three large pairs of compensation coils to null the field in each of the three spatial dimensions.

Significant further cooling is prohibited due to the recoil velocity on a single atom due to a single photon. This limit is called the recoil limit [35, 41] and for Rubidium 87 is 362 nK [36]. Cooling below this limit requires evaporation techniques in conservative potentials [50] created using either strong magnetic gradients or optical dipole traps [51] or the use of more complex sub-Doppler laser cooling methods such as Raman cooling [52] or Velocity Selective Coherent Population Trapping (VSCPT) [53]. The temperature of 10 µK is sufficient for our needs so we do not pursue such techniques. This temperature regime is consistent with other published results [40, 54] although direct comparison is difficult as the temperature is inversely related to the atom number which is considerably lower in this experiment.

A typical atom number in this cloud is 10-15 million although this can be increased by pushing more current through the dispenser. This number is more than an order of magnitude lower than possible using vapour cell MOTs [54, 55]. It is believed that the power of the cooling beams is one of the most severe restrictions on increasing the atom number. However an increase in the atom number will likely cause in increase in the cloud temperature, which would also be undesirable.

Absorption Imaging

The atomic clouds are imaged using a very standard absorption imaging technique. A weak beam (henceforth the imaging beam) of light resonant with the $2 \rightarrow 3$ transition is aligned through the atom cloud onto an Imaging Source DMK 21BUC03 camera. The camera focal plane is set to the position of the atom cloud. To take an image, the imaging beam is pulsed during a camera exposure. If there are atoms present, the light beam will be scattered out of the forward direction of the beam by the atoms, leading to a localised decreased intensity in the beam where the atoms are. The absorption of the beam follows the Beer-Lambert law, $I = I_0 \exp(-\sigma N)$ where I is the intensity incident on the camera, I_0 is the intensity with no atoms, σ is the absorption cross section per atom and N is the column density of atoms.

It can be seen from this that the atoms form a 'negative' picture in the beam. This can be turned into a proportional positive picture by taking an image with atoms and another without. Then by taking the logarithm of the pixel-wise division of these two images, we are left with a picture of σN. By dividing by the known cross-section, we are left with an image of the atom number density integrated along the imaging beam direction. Figures 3.5a and 3.5b show typical images obtained both without and with atoms respectively and 3.5c shows the image resulting from the logarithm of the division between them, giving a positive image of the atoms. Practically, slight discrepancies in the imaging beam between the two pictures lead to finite signals where there are no atoms. When quantifying the cloud number, the image is summed along one of the axes and a 1D Gaussian fit is conducted and then analytically integrated over to give a measure for the atom number that does not include these pixel fluctuations where there are no atoms.

3.4.1 Select System Measurements

To demonstrate some of the tools and techniques pertinent to the remainder of the experiment, we include measurements related to the calibration of the Master polarisation lock, the cloud atom number and temperature measurements using a typical Time of Flight (TOF) scheme.

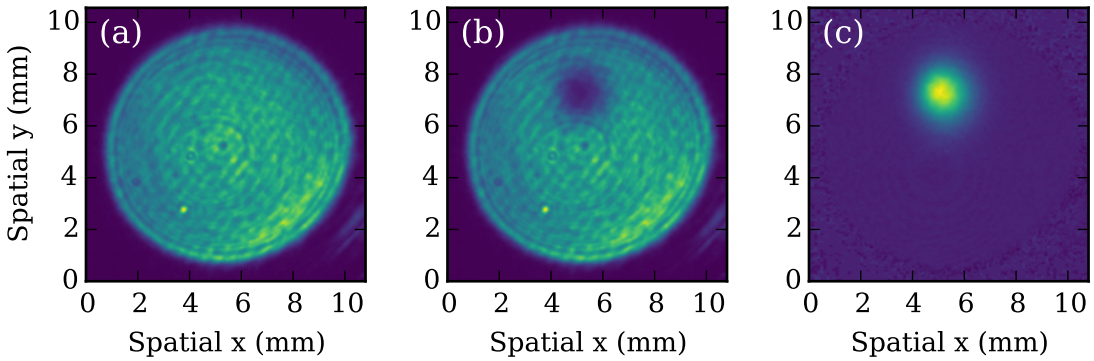

Figure 3.5: Acquisition of an absorption image. (a) The imaging beam passes through the vacuum chamber and onto the camera. The hard edge is from clipping on the conic section of the vacuum chamber. (b) When the beam is pulsed when atoms are present, the atoms scatter light from the imaging beam, leaving a shadow on the captured camera image. (c) By taking the logarithm of the pixel-wise ratio of an image with, and without atoms; an image of the atomic cloud is made.

Master Lock Calibration

As stated in Section 3.1, the polarisation lock is sensitive to external magnetic fields and other environmental fluctuations which, despite efforts of mitigation, cause the lock point to drift slowly in time. The propagation of this frequency drift affects the cooling beams, imaging beam and repumping beams' detunings from their relevant transitions. The small drift of a few megahertz is not an issue for the repumping beam as its large intensity saturates the transition so it is not precisely dependent on detuning. The drift can be compensated for the cooling and imaging beams through the use of an offset programmed into the first master AOM frequency, however this compensation requires a measurement that references the atomic transition. This is done by taking absorption images of the MOT clouds whilst changing the imaging AOM frequency. The resulting signal of the absorption image is σN as stated. The scattering cross-section σ for a very weak light beam has a Lorentzian form given by [36]

$$\sigma = \frac{\sigma_0}{1 + 4\left(\frac{\Delta}{\Gamma}\right)^2} \tag{3.1}$$

where σ_0 is the resonant cross-section, Δ is the detuning of the light from the transition and Γ is the linewidth of the transition. It is clear that $\Delta = 0$ yields the largest signal on the absorption image. By taking images for multiple light frequencies and comparing the signal size, we can calibrate the offset of the polarisation lock from the desired $2 \rightarrow 3$ transition and produce a known offset on the master AOM frequency to compensate the shift. Figure 3.6 shows an example of this process.

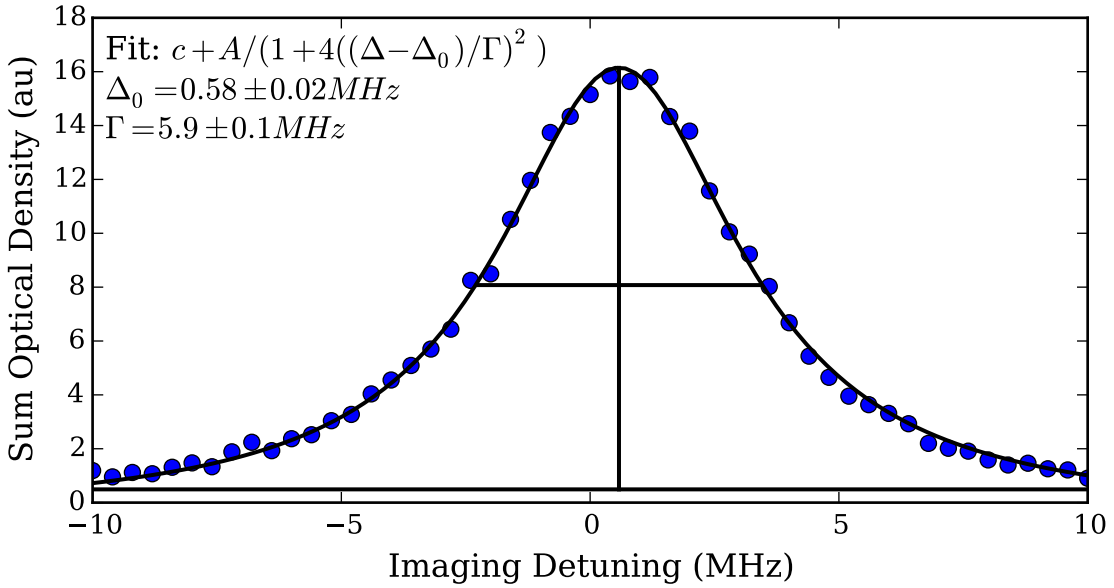

Figure 3.6: Total optical density from many absorption images with different imaging beam detunings. This image was taken with the experimental control not calibrated correctly. This miscalibration is known by the offset of the peak from zero detuning. The fitted Lorentzian linewidth of 5.9 ± 0.1 MHz is consistent with the literature value of 6.065 MHz.

Loading Rate

With the MOT operational, atoms are being continuously trapped from the Rubidium vapour and lost from it via collisions with hot particles. This forms a rate equation for the number of atoms in the cloud which (in the low density limit) has the solution of an exponential loading curve [56]

$$N(t) = N_0 \left(1 - e^{-t/\tau}\right). \tag{3.2}$$

Broadly, the time constant τ is the inverse of the collision rate per trapped atom and the asymptotic atom number N_0 is the product of a loading rate the the time constant. The loading rate is proportional to the density of Rubidium atoms to trap whilst the scattering rate $1/\tau$ has terms proportional to the density of scatterers including un-trapped Rubidium atoms. The largest atom number can then be obtained when the Rubidium density is much higher than the other background constituents of the vacuum pressure. The Rubidium density can be controlled by the current flowing through the Rubidium dispensers. The exact form of the loading rate is complicated but contains dependencies on the cooling beam sizes and powers.

By taking absorption images of the cloud for different MOT loading times, we can trace the exponential

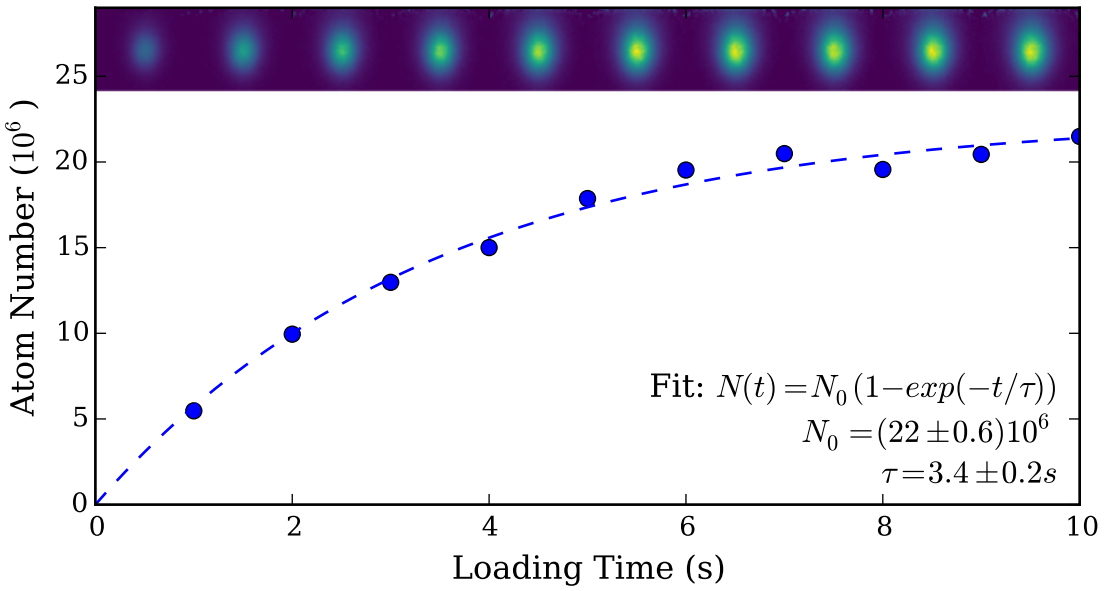

Figure 3.7: Exponential loading curve of atoms in the MOT. Rather than use a photo diode to collect the fluorescence of the MOT, we take a number of discrete absorption images after varying loading times.

loading form above. If a target number is required, the loading time can be set to an appropriate value. Figure 3.7 shows a typical trace for this loading curve.

Temperature Measurements

The temperature requirements for the atoms for the experiment are directly related to the method of temperature measurement. In a cold atom cloud, the temperature is a measure of the velocity spread of atoms in the cloud. When the MOT is turned off, the cloud expands ballistically according to this velocity spread, hotter clouds expand faster. Assuming a Gaussian spatial cloud distribution and a Gaussian 1D velocity distribution in accordance with the 1D Maxwell-Boltzmann distribution, the resulting spatial distribution after a given time t is the convolution of these two distributions (with the velocity multiplied by the time to give a position). The result of this convolution is a Gaussian with a width given by [57]

$$\sigma(t) = \sqrt{\sigma(t=0)^2 + \frac{k_B T}{m} t^2}. \tag{3.3}$$

Using this equation, we can measure the width of the cloud for many different values of t (commonly referred to as the *time of flight*) and appropriately fit this expression to find the temperature. An example of these plotted widths for both the atoms after the MOT and sub-Doppler molasses phase is shown in Figure 3.8.

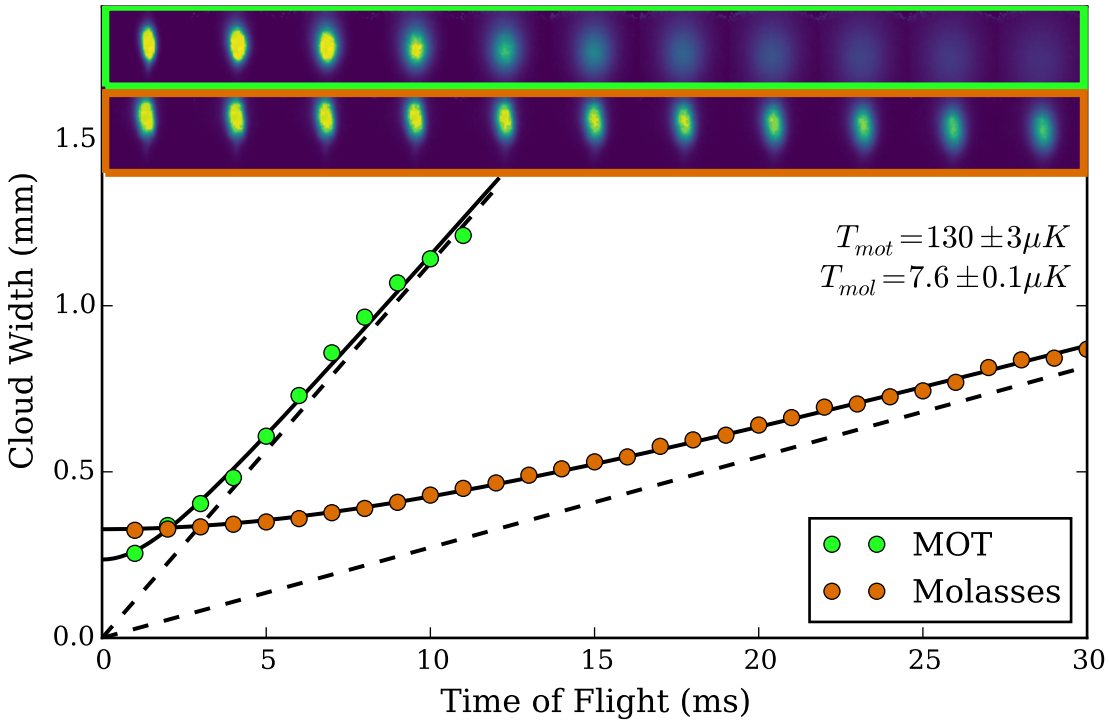

Figure 3.8: Illustration of how the temperature is inferred. The widths of the cloud are measured for multiple times of flight after the cloud is released. Equation (3.3) is then fitted to the data to obtain a temperature. In this graph, we show the widths for a cloud after MOT cooling (green points/top atoms) and for a cloud after the sub-Doppler optical molasses (orange/bottom atoms). The fit to Equation (3.3) is shown (black solid) alongside the asymptotic behaviour in the long time limit (black dashed).

In these cases, the measured temperatures are $130 \pm 3\,\mu\text{K}$ and $7.6 \pm 0.1\,\mu\text{K}$ for the MOT and molasses respectively. It should be noted that the temperature can also be more simply inferred from a single measurement at sufficiently long time where the contribution of the $\sigma(t = 0)$ term in Equation (3.3) is negligible.

3.5 Concluding Remarks

In this chapter, we have given a very brief overview of laser cooling and trapping and discussed the exact implementation for this experiment. We have provided typical characteristics of the cloud and shown the results of typical diagnostic tests conducted frequently throughout the project.

The atom source fulfils its requirements of producing clouds of atoms in a number sufficient to allow for low error in the quantum state readout, and sufficiently cooled to allow for the spatial separation state readout discussed in Chapter 4.

CHAPTER 4

STATE READOUT - STERN-GERLACH SEPARATION

As discussed in Chapter 1, the states of interest in our system are the magnetic projection sub-levels of a particular ground state. After the atoms have experienced the Hamiltonian under test, the populations of these states need to be read out. This is achieved by spatially separating the sub-level populations using an applied inhomogeneous magnetic field over the atom cloud. Due to their differing magnetic dipole moment projections, the different states experience differing accelerations in the presence of the magnetic field gradient as famously conducted using atomic beams by Stern and Gerlach in their seminal experiment investigating the quantisation of angular momentum states [58]. Once the states have been spatially separated, a standard absorption imaging system is used to capture the different clouds in a single shot. Computer fitting then calculates the number of atoms in each cloud.

Spatial separation of m_F states using Stern-Gerlach separation has been achieved before in the cold atoms community but predominantly on Bose-Einstein Condensate (BEC) clouds [59, 60]. The lower kinetic temperature ($\sim$100 nK) of these clouds leads to a much slower expansion than the clouds used in this experiment ($\sim$10 µK) which makes the task of effectively separating the states significantly easier. To apply the technique to atoms of approximately 10 µK was a significant experimental challenge. This chapter will outline the basic theory of the technique used, introduce the physical and technical challenges encountered and then present the results obtained.

4.1 Stern-Gerlach Theory

The interaction of an atom's magnetic moment with a weak external magnetic field causes a linear energy shift of the atomic state as given by the Zeeman splitting formula

$$\Delta E = \mu_B g_F m_F \left|B\right| \tag{4.1}$$

provided the quantisation axis is taken to be in the local magnetic field direction. If the field is inhomogeneous, the spatially dependent energy of the states will lead to a force given by the energy's spatial derivative. If we consider a field that points along the quantisation axis but with a magnitude varying linearly along the x-axis, the force on the atom will be given by

$$F_x = \mu_B g_F m_F \frac{d\left|B\right|}{dx}. \tag{4.2}$$

If this inhomogeneous field is applied for a time t, it will result in a displacement of the atoms that depends on their magnetic moment projection quantum number, m_F

$$\Delta x(t) = \frac{1}{2}\frac{F_x}{m}t^2 = \frac{1}{2m}\mu_B g_F m_F \frac{d\left|B\right|}{dx}t^2. \tag{4.3}$$

This results in a single cloud of atoms splitting into multiple clouds dependent on their state which can be imaged in a single shot. We note that the spatial state separation increases with t^2. As the atoms are not trapped during this separation, the clouds freely expand at a rate related to their temperature. To resolve the spatial splitting of the cloud and thus reliably measure the state populations, the induced separation must outpace the thermal expansion, which as stated in Section 3.4.1, is given as

$$\sigma(t) = \sqrt{\sigma^2(t=0) + \frac{k_B T}{m}t^2}. \tag{4.4}$$

We note that at large t ($t >> \sigma^2(t=0)\cdot\sqrt{m/k_B T}$), this can be approximated to $\sigma(t) = \sqrt{\frac{k_B T}{m}}t$ and thus the width expands linearly with time. Taking the quotient of the cloud separation and cloud widths, we obtain a quantification of the effective contrast; C, between the different clouds

$$C(t) = \frac{\Delta x(t)}{\sigma(t)}. \tag{4.5}$$

Since this ratio is proportional to t at large times, it is clear that applying the pulse for longer times increases the contrast between the clouds.

However, pulsing for long times may not be practical in real systems for several reasons, the principal issue being the reduction in optical density of an expanding cloud. As cold atom imaging systems have a noise floor which places a lower limit on the detectable optical density, this places a limit on the total expansion time. For absorption imaging systems, the noise floor is typically a result of imaging beam fluctuations and camera pixel noise. There may also be spatial constraints of the imaging system originating from a limited field of view of the camera or the limited size of the imaging beam. This places an upper limit on the total separation of the clouds, which in turn, limits the separation time. The 1 cm conic section opening provides the limiting aperture in our apparatus and is a more significant issue than the reduction in optical density.

Surmounting these constraints to achieve sufficient contrast can be facilitated by reducing the kinetic temperature of the cloud to decrease $\sigma(t)$, or increasing $\Delta x(t)$ by increasing the field gradient. Our cloud temperature of $\sim$10 µK was limited by our laser set-up and would require more complex cooling techniques as discussed in Chapter 3 to be reduced significantly. Instead, efforts were focused on increasing the field gradient. The gradient itself however cannot be increased arbitrarily without encountering physical and technical issues. An explanation of these issues follows.

4.2 Creating Magnetic Field Gradients

The simplest method of creating a switchable magnetic field gradient along a particular direction is to place a current carrying coil along the axis of the desired gradient. The axial field gradient from a circular coil with N turns, radius R a distance x along its symmetry axis can be easily derived from the Biot-Savart law as

$$\frac{dB_x}{dx} = \frac{3\mu_0 R^2 I N}{2} \frac{x}{(x^2+R^2)^{5/2}}. \tag{4.6}$$

A large gradient can be achieved using a coil with small radius placed close to the atoms. In cold atom systems however, the proximity to the atoms is generally limited by the size of the vacuum chamber. Atom chips [59] can be used in vacuum chambers to create magnetic field gradients using wire tracks on vacuum compatible substrates however the implementation of such devices requires sufficient expertise as well as significant design efforts for the vacuum system. It is also possible to place more conventional coils in vacuum but these suffer from a lack of convective cooling from the air, so more complicated cooling devices

such as water or oil cooling must be used. Due to the relative complexity of both of these methods, these options were not pursued.

Non-geometric considerations for an increased gradient involve having many turns in the coil (large N) and a high current flowing through (large I).

4.2.1 Finite Coil Time Constant

Starting from no current flowing through a coil and then applying a steady voltage V, the inductance of the coil prevents the current change from occurring instantly. This finite time response can be simply modelled by considering the coil as an RL circuit, where R is the circuit resistance and L is the inductance of the coil which is primarily governed by the geometric size of the coil and the turn number. When the current gate is opened at time $t = 0$, the current as a function of time is given by

$$I = \frac{V}{R}\left(1 - e^{-t/\tau}\right) \tag{4.7}$$

where $\tau = L/R$ is the time constant of the circuit. Unless the time constant is much less than the pulse time, the finite turn-on time degrades the cloud separation as stated in Equation (4.3). Since the separation force $F \propto \frac{dB}{dx} \propto B \propto I$, the force becomes time dependent with the separation given by

$$\Delta x(t) = m_F\ \alpha\ \beta\ \gamma\ \frac{V}{R}\left(\frac{t^2}{2} - \tau t + \tau^2(1 - e^{-\frac{t}{\tau}})\right) \tag{4.8}$$

where $\alpha \equiv \mu_B g_F/m$ is the force exerted per unit of magnetic field gradient produced at the atoms (excluding the state dependence), $\beta \equiv \frac{dB}{dx}/B$ is the field gradient produced per unit of magnetic field at the atoms and $\gamma \equiv B/I$ is the field created per unit current.

The separation given by Equation (4.8) is strictly less than the case with negligible inductance given in Equation (4.3) and this difference increases with larger τ as shown in Figure 4.1. Reducing τ by either increasing the resistance or decreasing the inductance by reducing the turn-number is, however, not without penalty. An increase in the resistance without an increase in applied voltage reduces the asymptotic current which reduces Δx despite the decrease in τ. Reducing the turn number decreases the inductance, which in turn decreases τ. For large turn numbers, the coil inductance scales as $L \propto N^2$. If the resistance of the coil is governed just by the length of wire, then the resistance scales as $R \propto N$. The ratio of these scalings therefore leads to a time-constant that scales linearly with the turn number, $\tau \propto N$. The additional complication is that $\gamma \propto N$ leads to an optimal turn-number for a given pulse time. The calculations are

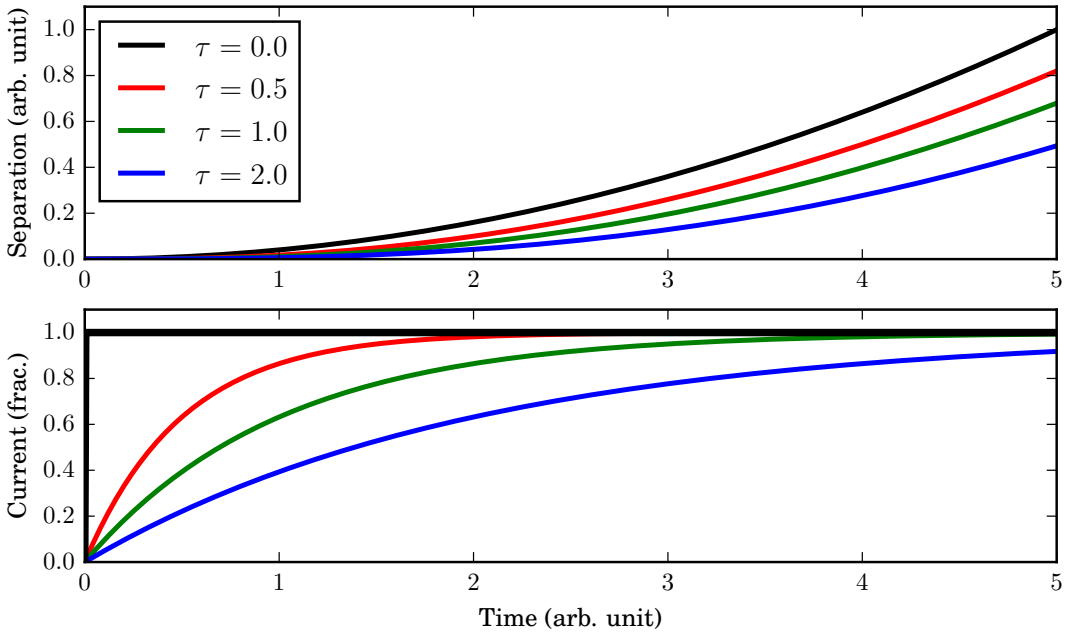

Figure 4.1: Calculation showing the effect of finite time constant on current and resultant separation in a Stern-Gerlach pulse. The separation curves with finite time constants are strictly below the quadratic separation curve for $\tau = 0$

relatively straightforward but are best solved for a specific case numerically.

It is clear from the discussion so far that the design of a optimal Stern-Gerlach coil is an exercise of careful compromise between multiple different parameters. However, there are more effects that must be taken into account to gain a complete understanding.

4.2.2 Coil Switch Off and Eddy Currents

To measure the separated cloud populations, the atoms must be imaged after they have been separated. If there is a residual magnetic field on the atoms whilst they are imaged, the resonant frequency of the imaging cycling transition will be shifted by the Zeeman effect. If this shift is on the order of a linewidth (6 MHz), the optical density of the atoms will be reduced, hampering imaging. Due to reasons outlined below, there may be a residual magnetic field after attempting to stop the current pulse and these can cause issues if the system is not designed to minimise them.

Magnetic Energy Dissipation

When the voltage across the coil is turned off, the inductance of the coil prevents the flowing current from decreasing quickly. Normally, switching is done using a power MOSFET or IGBT semiconductor device. Whilst the switching of these devices stops the current flowing through them, it does little to dissipate the energy stored from the current in the magnetic field of the inductive coil. A varistor or Zener-diode array is typically employed over the IGBT collector-emitter to safely dissipate the energy when the coil is switched off as well as protect the IGBT from the induced *emf* from the coil. A varistor is a bi-directional diode-like device that has close to infinite resistance below a threshold voltage V_{clamp} and a small resistance above. When the switch is opened, the flowing current charges the residual capacitance of the open IGBT until the voltage buildup on the terminals of the IGBT and varistor reaches the varistor conduction threshold.

At this point, the current can then flow through the varistor, dissipating energy at a rate IV_{clamp}. The energy stored in the coil in the form of the magnetic field is given by

$$E = \frac{1}{2}LI^2. \tag{4.9}$$

If we assume that the conducting resistance of the varistor is negligible, a functional form for the current can be derived from the power dissipation equation

$$\frac{dE}{dt} = -V_{clamp}I(t) = \frac{d}{dt}\frac{1}{2}LI^2 = LI(t)\frac{dI}{dt} \tag{4.10}$$

which has the solution for current as

$$I = I_0 - \frac{V_{var}}{L}t. \tag{4.11}$$

With most of the current dissipated, the small amount of energy stored in the IGBT capacitance must be dissipated. At this point, the circuit effectively becomes an RLC circuit starting with a charged capacitor. This results in current oscillations whilst the remaining power dissipates through the resistance of the circuit, which is much slower. To counter this, an RC snubber can be used in parallel with the varistor and IGBT. This constitutes a resistance and capacitance chosen to allow another path for the oscillating current with parameters chosen to provide more effective damping. The amount of energy in the capacitor must normally be significantly large to cause a problem.

To achieve fast switching characteristics, one should chose a varistor or Zener diode with a high conduc-

tion voltage. This allows the bulk of the current to dissipate quickly inside this device. The upper limit on the voltage is the breakdown limit of the IGBT. The clamping device should be chosen to be able to dissipate the expected amount of energy stored in the coil. It is wise to overspecify the device as there is usually characteristic degradation over time. Measuring the coil switch off with the clamp in place can give information on the required characteristics of an RC snubber to create a more critically damped circuit if necessary.

Eddy Currents

Even with ideal current switching behaviour, there can still be a residual magnetic field long after the current in the coil has stopped flowing, due to the presence of conductive materials near the coil. This is because internal eddy current loops will be induced in these materials by the changing magnetic flux. The direction of these induced currents will attempt to stop the magnetic field from changing. Relevant components that cause these problems are copper gaskets used to seal the vacuum chamber components and to a lesser extent, the steel vacuum chamber itself. The eddy current loops will decay according to their geometry and material resistivity. The net effect when the current in the coil has stopped is an exponential decay of some of the field according to this eddy current time constant. This effect can be minimised by using non-conductive components or reducing the mutual coupling of such components to the gradient coil by mismatching their sizes and avoiding orientating them the same way. It should be noted that these eddy currents also affect the field as it is being switched on which can lead to a double exponential growth curve for the field at the atoms.

As discussed, the coil switch off can be optimised but the eddy currents are very difficult to deal with if they are present. For this reason, a time t_{wait} may be introduced between the gradient pulse and the imaging shot to wait for the residual field to decay to a level where the atoms can be imaged. Since the atom cloud is still expanding during this time while the differential acceleration has stopped, this introduces a further worsening of the contrast compared to pulsing for $t_{pulse} + t_{wait}$. The separation in t_{wait} is a linear function of t as it is solely due to the velocity of the clouds at the end of the pulse. As the thermal expansion has the same dependence, the contrast will plateau during t_{wait}.

To incorporate the behaviour of the cloud during t_{wait}, we take Equation (4.8), which is valid until the end of the pulse and simply add an additional term. If we neglect any acceleration due to the decaying eddy currents, this extra term should just be linear in t_{wait} with the velocity accrued from the acceleration in the

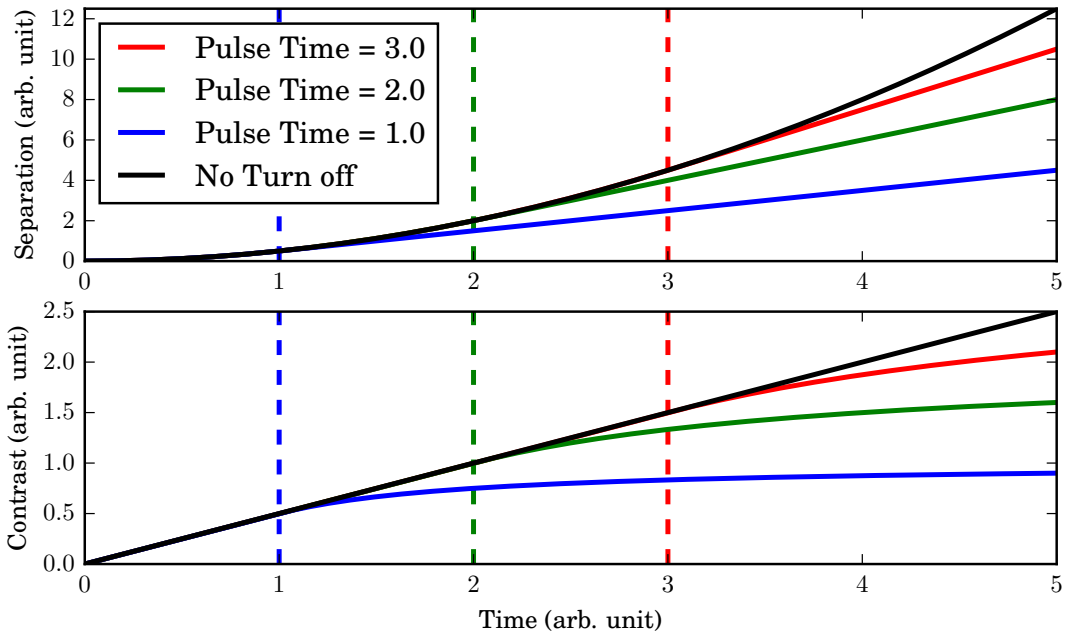

Figure 4.2: Calculation of the evolution of cloud separation and contrast when the pulse is cut short and t_{wait} is introduced. Note that the contrast reaches a plateau after the pulse has switched off. This is due to the t scaling of the separation contribution from Equation (4.12). In this example, $\tau = 0$.

pulse time. This additional term takes the form

$$\Delta x_{off} = v(t_{pulse}) t_{wait} = m_F \ \alpha \ \beta \ \gamma \ \frac{V}{R} \left(t_{pulse} - \tau \left(1 - e^{-t_{pulse}/\tau} \right) \right) t_{wait}. \tag{4.12}$$

A plot showing the form of the total displacement as well as the contrast is shown in Figure 4.2.

4.2.3 Non-Linear Zeeman Shifts

Equation (4.1) shows that the magnetic projection sub-levels experience energy shifts that depend linearly upon the applied magnetic field for small fields. To create large gradients, it may be necessary to create larger fields. For these large fields, the energy shifts become non-linear. This is because in the presence of very strong fields, the hyperfine structure splitting becomes overwhelmed by the Zeeman splitting and the (F, m_F) doublet ceases to be a appropriate description of the atom. For the ground states of Rubidium (and

other alkali metals), the general form of the energy shift is given by the Breit-Rabi formula [36]

$$\Delta E = g_I \mu_B m B \pm \frac{\Delta E_{hfs}}{2}\left(1 + \frac{2mx}{2I+1} + x^2\right)^{1/2} \tag{4.13}$$

where ΔE_{hfs} is the ground state hyperfine splitting, I is the nuclear angular momentum quantum number, g_I is the gyromagnetic ratio for nuclear spin, B is the magnetic field and $m = m_I \pm m_J$ where the sign of the $\pm$ is the same as the formula. The symbol x is a dimensionless parameter for the magnetic field given by

$$x = \frac{(g_J - g_I)\,\mu_B B}{\Delta E_{hfs}} \tag{4.14}$$

where g_J is the gyromagnetic ratio for spin-orbit coupling. A plot of the the Breit-Rabi splitting and energy-field derivative for the Rubidium 87 $F = 1$ ground state is shown in Figure 4.3a.

If we consider the high field region for the $F = 1$ manifold, the field dependence is linear in m_J and thus $\frac{dE}{dB}$ becomes the same for all three m_F states. This means that in this regime, there will be no relative separation between the states. In the intermediate region, there is deviation from the linear regime which reduces the differential force between the states as shown in Figure 4.3b. This starts to reduce the effectiveness of simply increasing the magnetic field created by the separation coil to increase the spatial gradient and thus the force on the atoms. In the results presented in this chapter, we reach a region where these non-linear energy shifts become noticeable.

Rather than incorporate the entire Breit-Rabi formula into our model, we expand it to quadratic order of B. This is generally referred to as the quadratic Zeeman shift. This is done by taking the Taylor expansion of Equation (4.13) around $B = 0$ to second order in B to give an equation of the form $\Delta E = aB + bB^2$. The linear a coefficients are as given in the linear regime and quadratic coefficients given by

$$b = \frac{\Delta E_{hfs}}{2}\left(\frac{1}{2} - \frac{2m^2}{(2I+1)^2}\right)\left(\frac{(g_J - g_I)\,\mu_B}{\Delta E_{hfs}}\right)^2 \tag{4.15}$$

with calculated values for the $F = 1$ and $F = 2$ ground state manifolds given in Table 4.1.

The practical effect of this non-linear shift is that the differential forces between the states become dependent explicitly on the magnetic field as well as the gradient. Since for a given coil geometry, the gradient is proportional to the field itself $\frac{dB}{dx} = \gamma B$, the differential force follows the quadratic form

$$\Delta F_x = (\Delta a + 2\Delta b B)\frac{dB}{dx} = \gamma(\Delta a B + 2\Delta b B^2). \tag{4.16}$$

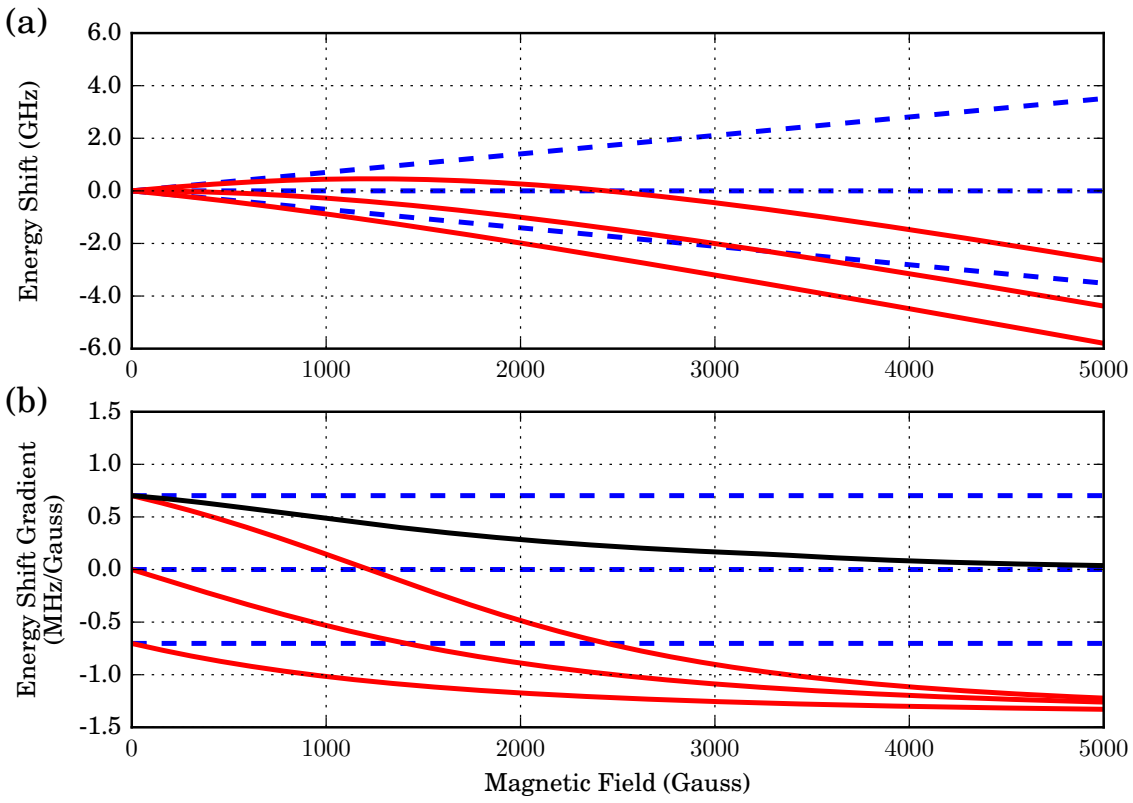

Figure 4.3: Calculation of (a) Briet-Rabi splitting of the energy level and (b) the energy-field derivative as a function of applied magnetic field for the manifold of the Rubidium 87 $F = 1$ ground state. The linear approximation (Blue Dashed) is shown with the full Breit-Rabi model (Red Solid). In addition, the smallest difference between derivatives is shown (Black Solid)

Table 4.1: Derived values for the quadratic Zeeman shift parameters a and b for the Rubidium 87 ground state manifold.

F	**mF**	**mI**	**mJ**	**a (MHz/G)**	**b (kHz/G^2)**
1	-1	-1/2	-1/2	+0.7	-3/8 0.575 = -0.216
	0	+1/2	-1/2	0	-1/2 0.575 = -0.288
	+1	+3/2	-1/2	-0.7	-3/8 0.575 = -0.216
2	-2	-3/2	-1/2	-1.4	0
	-1	-3/2	+1/2	-0.7	+3/8 0.575 = +0.216
	0	-1/2	+1/2	0.0	+1/2 0.575 = +0.288
	+1	+1/2	+1/2	+0.7	+3/8 0.575 = +0.216
	+2	+3/2	+1/2	+1.4	0

For the practical purposes of having sufficient contrast, it is necessary to consider the states with the worst differential force. For the $F = 1$ manifold, this is between the $m_F = -1$ and 0 states and for the $F = 2$ manifold, between the $m_F = 1$ and 2 states. For these states, Δb has the opposite sign to Δa and it is clear that the effect of the quadratic shift is minimised if a particular $\frac{dB}{dx}$ is achieved with a small magnetic field B, i.e. having a large γ.

4.3 Experimental Implementation

To separate and read-out the state populations for a given cloud, there should generally be a bias field applied before the Stern-Gerlach pulse occurs. The states to be read-out will be the m_F levels associated with the quantisation axis defined by the local field. In the apparatus, this bias field is applied along the z-axis (see Figure 3.4) whereas the gradient field is applied along the x-axis. The m_F states adiabatically follow the quantisation axis defined by the local magnetic field providing the angle of the field does not change too quickly [61]. After the gradient pulse, the field from the chamber eddy currents is allowed to settle to zero for a few milliseconds before an absorption image is taken. If an image is to be taken of the $F = 1$ sub-levels, a $F = 1 \rightarrow F' = 1$ beam (see Chapter 5) is turned on during this settle time. This pumps the atoms into the $F = 2$ manifold where normal absorption imaging can occur on the cycling $F = 2 \rightarrow F' = 3$ transition. The atoms will still have the position and velocity corresponding to their state during the gradient pulse.

During the course of the work on this thesis, two different coil configurations were tested. As the MOT coils are ideally placed around the vacuum chamber in close proximity to the atoms, they were used for the Stern-Gerlach pulse. A current switching circuit was made to switch between the current configurations for the MOT and Stern-Gerlach pulse. Initially, the magnetic gradient was created with current flowing through only one of the MOT coils (henceforth configuration A) but as presented later, this led to prohibitive non-linear effects that greatly harmed the contrast. The second configuration (henceforth configuration B) was to pass the current through both of the MOT coils and then a similarly shaped third coil ($r = 5.25\,\text{cm}, N = 35$), situated adjacent to one of the MOT coils. This creates a field profile similar to an anti-Helmholtz coil pair with different turn numbers. This increases the magnitude of the magnetic gradient for a given central magnetic field over configuration A. Two coils with the same turn number would also provide this same effect but the mismatch is necessary to create a non-zero field at the atoms to prevent spin-flips. Figure 4.4 shows the field and gradient profiles for the steady state currents through configuration B.

The Stern-Gerlach current is generated using two 24V lead-acid batteries connected in series for an expected steady state current of 160A in configuration A and 67A in configuration B for respective field

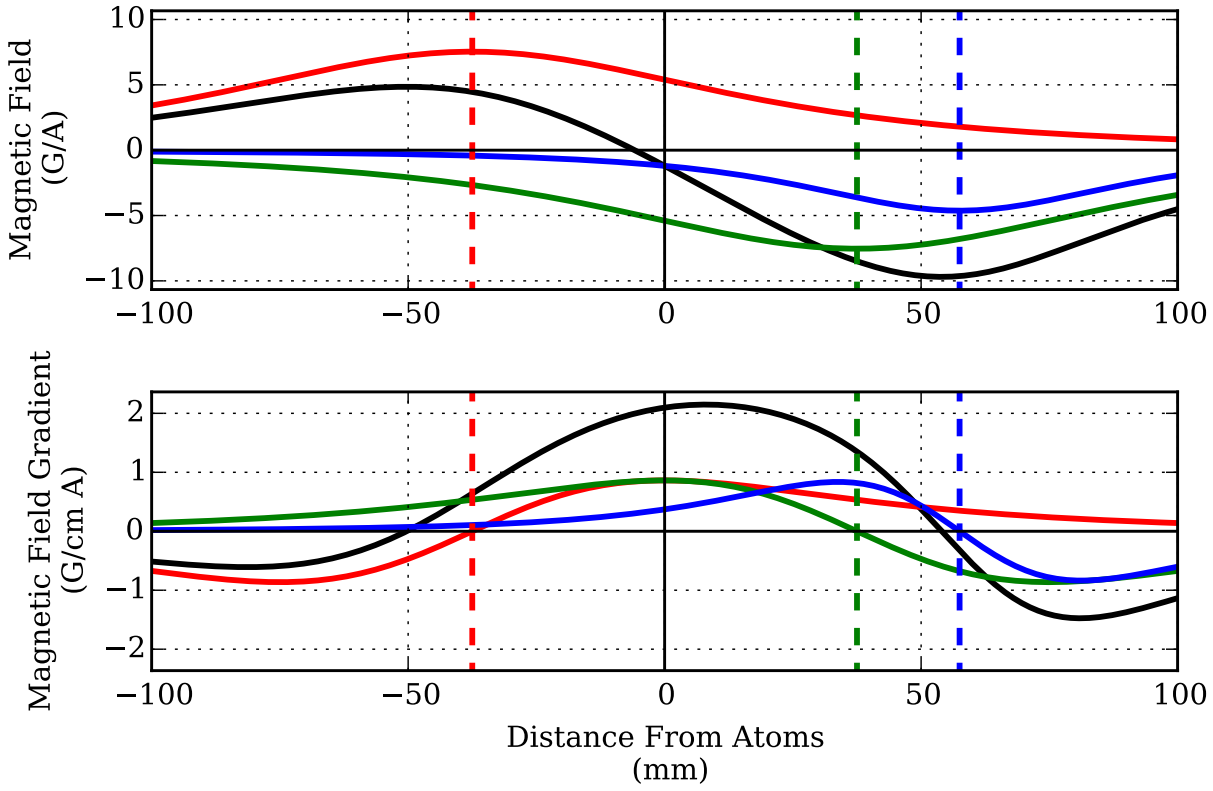

Figure 4.4: Calculations of the field and gradient profiles along the central coil axis of configuration B. The two MOT coils are shown in red and green with the third coil shown in blue. The configuration A profile simply follows the red curve. The total field and the gradient in configuration B are shown in black. The third coil is necessary to create a non-zero field at the location of the atoms.

gradients of 136 G/cm and 140 G/cm. As mentioned, Stern-Gerlach separation in cold atom experiments are typically done on BEC clouds which require significantly lower relative state separations for sufficient contrast. Published examples include DeMarco *et al.* in which a field gradient of 80 G/cm applied for 10 ms is sufficient, Stamper-Kern *et al.* [62] in which a low gradient of a few gauss per centimetre is applied for 40 ms and Leanhardt *et al.* in which they use an atom chip to achieve 180 G/cm for a few milliseconds.

The current through the coils is controlled using a separate IGBT (STMicroelectronics STGE200NB60S) from the MOT gradient switching circuit. A schematic of the coils and their drivers in both configurations is shown in Figure 4.5. In addition, a fuse unit (2x Bourns MF-RHT650-0[1] in parallel) was placed in series, next to the battery to provide a safeguard against a fault condition (eg. the field being left on too long).

Clamping across the IGBT is achieved using 10 Zener diodes (model number 1N3321B) connected in series, placed in parallel across the IGBT collector-emitter terminals. The clamping voltage of this array is approximately 240V.

To capture the dynamics of the current and field during a typical pulse, we use a TTi I-prober-520 magnetic field probe. As the current is too large for the probe's current head to measure directly, a signal proportional to the current was measured by placing the probe head close to the current carrying wire. As the field in the centre of the chamber cannot be measured, we instead measure the field outside of the chamber at a point along the coil axis without the probe overloading. This was done ≈10 cm away from the coils. The total field is a combination of the field produced by the coil currents and any additional eddy currents. At the atoms, the relative strength of the eddy currents will be larger but this method still affords us the ability to extract certain characteristics of them. The scaled results of these traces for a pulse time of 15 ms are shown in Figure 4.6.

The rising time constants are measured roughly by reading when the current/field reaches 63% of its maximum value. These are determined to be 1.0 ms and 1.5 ms for the current and field respectively for configuration A and 1.9 ms in both cases for configuration B. It is consistent with the principles of eddy currents that the field rise is inhibited by their presence. The full details of the field rise are likely to follow a double exponential curve [63]. The contribution of the eddy currents to the field measurements made here is smaller in configuration B as there are conducting coils closer to the probe whereas in configuration A, the single conducting coil is on the opposite side of the chamber. We do not expect that their relative effect at the atoms is smaller.

[1] These fuses are meant to be resettable. They become highly resistive in the presence of large current and then return to their highly conductive state when the current dies. However, the large currents in our system meant that the two fault conditions that were encountered made them exceptionally resistive on a more permanent basis.

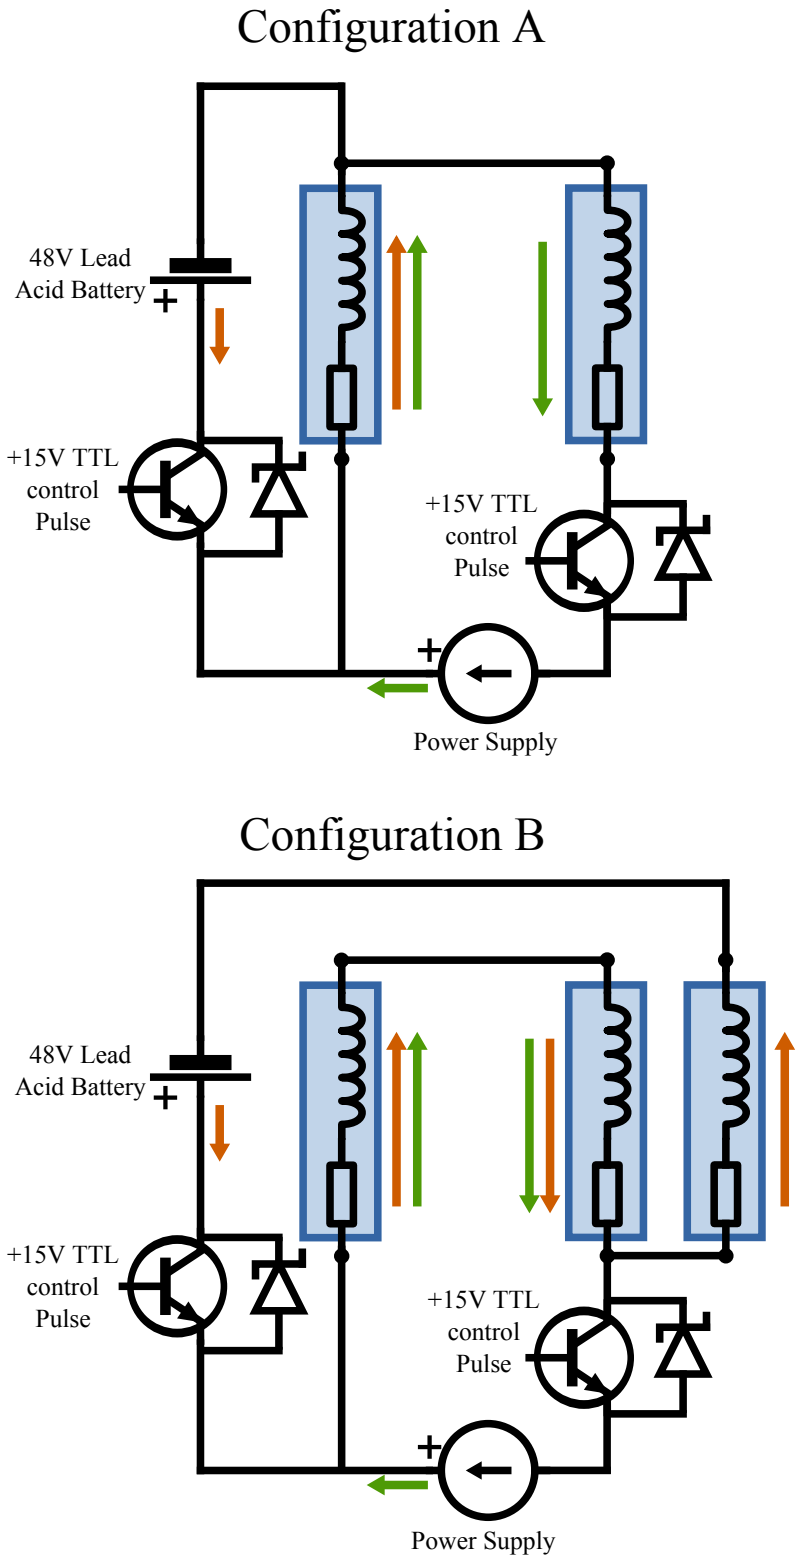

Figure 4.5: Electronic circuitry driving the coils responsible for the MOT quadrupole magnetic gradient and the Stern-Gerlach gradient pulse. Both configurations A and B are shown.

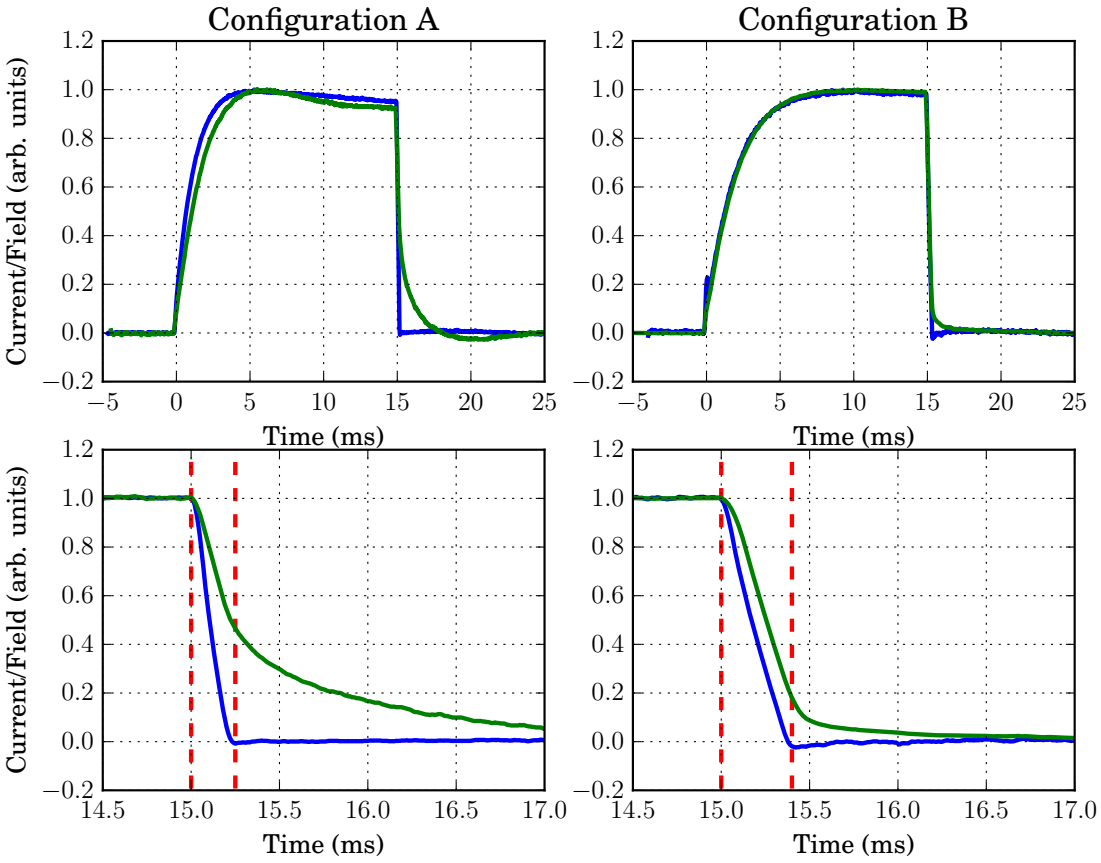

Figure 4.6: Traces of the current and magnetic field dynamics during a 15 ms Stern-Gerlach pulse. The current trace (Blue solid) is shown with the field trace (Green solid) for the entire pulse for configurations A and B in the top-left and top-right respectively. The bottom-left and bottom-right plots show zoomed in dynamics of the switch off for both configurations. The red dashed lines indicate the window determined as the current switching time. All traces are normalised to the maximum value.

The inductance of the coils can be extracted from the current time constants and the measured resistances as 300 µH and 1.36 mH for configurations A and B respectively. The expected substantial increase in inductance in configuration B due to the much larger turn number is tempered due to the mostly anti-Helmholtz configuration of the different coils.

It can also been seen that for configuration A, there is a noticeable decrease in both the field and current during the pulse. This is believed to be caused by the fuses starting to increase their resistance. This could potentially cause issues with inconsistent separations if the fuse is not able to fully reset before the next pulse. This is seen in configuration A but not B because of the larger current.

The field turn-off times can be seen to be longer than the current switching time. The current switching is approximately of linear form as discussed previously. The switching time is measured to be 250 µs and 400 µs for configurations A and B respectively. This is surprising as the power dissipation rate for the two configurations should be the same but we expect that the energy stored in the coils, as calculated from $E = \frac{1}{2}LI^2$ should be 3.8 J for configuration A and the smaller amount 3.0 J for configuration B.

With the current no longer flowing through the coils, the field profile should be reduced to just the eddy current contribution. From this, we can estimate the time constant of the eddy currents as 0.76 ms. This is roughly consistent with the increase in time constant for the turn-on time constant for configuration A.

4.3.1 Image Processing

Once an image of the separated clouds is captured, further image processing steps are required to extract the relative populations of the different states. The main exercise is to fit a Gaussian function to each of the cloud projections. From these fitted Gaussians, the number of atoms can be estimated from the integral over the Gaussian as is very common in cold atom measurements (see Section 3.4).

Fitting multiple, closely spaced superimposed Gaussians provides a large number of degrees of freedom which leads to an error function with many local minima. Basic curve fitting algorithms use gradient descent processes and thus only seek the minimum in the error function that is local to their initial starting point. Since we are interested in the minima associated with our physical model, we must provide good initial guesses for the fitting parameters that match our model.

Whilst the guesses can be entered manually, we use a technique that obtains them using a Blob Detection Algorithm [64]. This is a fairly standard computer vision technique used to detect the location and scale of 'blob-like' features in an image. The process involves taking the convolution of the image with the Laplacian of differently scaled versions of target kernels. In our case, the desired feature is a 2D Gaussian so we use

the Laplacian of Gaussian (LoG) kernel for the convolution. The result of this convolution is a discrete three dimensional scale space function with two spatial dimensions and one scale dimension. When the image contains a blob of a particular size, the result of this convolution operation will produce a local maximum in all three of these dimensions at the position and scale matching the blob. We extract the local maxima inside of this discrete representation to give rough locations and sizes of clouds in the original image. These features are then used as the initial guesses to fit the superposition of Gaussians numbering how many clouds are expected. The scale space limits are chosen to reflect realistic cloud sizes.

Once the initial guesses are obtained, fitting of the Gaussians is done using the very standard Levenberg-Marquardt method. Fitting of the Gaussian clouds could potentially be done with either a 2D Gaussian fit on the 2D atom image or a 1D Gaussian fit on either the sum projection onto the separation axis or a slice through the middle of the clouds in the separation axis. These three techniques were investigated with the projection method seeming to be more robust to irregularities in the cloud shapes. For the remainder of this thesis, unless otherwise stated, all results from such fits are produced using the projection method. The fitted positions and widths of the Gaussians are constrained close to the initial guesses given by the blob detection algorithm and the amplitudes are constrained to be positive.

Note that neither the blob detection nor the fitting routine are flawless and without sufficient contrast between the clouds, both algorithms are prone to failure. The most common fail point occurs when the population in one of the states is significantly depleted. In this case, the cloud associated with this state cannot be detected by the blob detector. To counter this problem, when conducting an experimental run to measure the relative state populations, a calibration stage is conducted where an image is captured with identical timings to the experimental runs but with no state preparation or coupling, such that each of the states is moderately populated. The blob detection and subsequent fit are conducted to provide a reference for the peak positions and widths. As the timings of the experimental runs are identical, the only variables that change between the experimental runs are the state populations which are then fitted to with tight constraints on the peak positions and widths around the reference values.

4.3.2 Separation Quantitative Analysis

With the ability to determine the cloud positions from a shot of given t_{pulse} and t_{wait}, we can test our theoretical model developed in this chapter for the dynamics of the cloud separations by varying these two parameters. In this section we compare the measured peak positions and theoretical expectations of the peak positions from the sum of Equations (4.8) and (4.12).

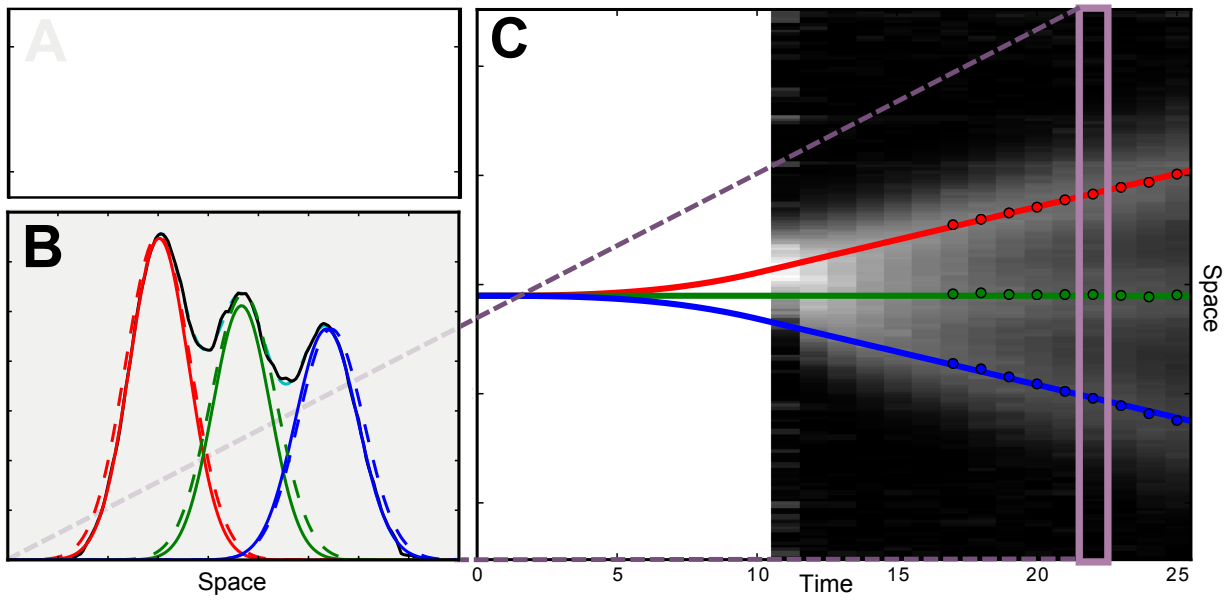

Figure 4.7: Illustration of imaging Stern-Gerlach visualisation. (A) Filtered atom shot for a particular t_{pulse} and t_{wait}. (B) Summed projection of the atom shot onto the separation axis (Black solid) with the decompositions into the separate clouds given by the blob-detection algorithm (Coloured dashed) and the subsequent fitting algorithm (Coloured solid). (C) Build up of different atom shots with fixed t_{pulse} but varying t_{wait}. The projection of a single shot constitutes a single column of data. t_{pulse} is given by the white area where the atoms cannot be imaged. The coloured circles show the fitted centres of the different clouds if successfully found and the coloured lines are an attempt to fit these cloud centres.

As the peak contrast in much of the data is insufficient to reliably locate the peak positions, we will also be using a visualisation method to qualitatively test agreement with the images that were unable to provide quantitative results. This visualisation involves building up a 2D image where in each column of the image data is the projected atom shot data for a single shot. Then movement along the columns involves varying the time. An example of this visualisation is shown in Figure 4.7.

The relevant parameters in Equations (4.8) and (4.12) are α, β, γ, V, R and τ. When comparing the data to the theory, we enter α from literature values [36], β and γ from standard coil calculations, V and R from experimentally measured quantities and allow τ to vary as a fitting parameter to our model as we are unable to measure the time dynamics of the field profile at the position of the atom cloud. As α, β, γ, V and R are all proportional (or inversely proportional for R) to the relative separations, we multiply our separation equation by a new parameter η that is allowed to vary in the fit rather than varying multiple perfectly correlated variables. If our estimated values for all of the parameters are correct, η should evaluate to one. We also allow for an initial spatial offset to the positions of the cloud at time $t = 0$, denoted x_0, as a fitting parameter. This leaves three fitting parameters η, τ and x_0.

Figure 4.8 shows these results for the $F = 1$ states in configuration B for a select number of t_{pulse}'s. The full dataset varies t_{pulse} from 5 ms to 15 ms in increments of 1 ms. For each t_{pulse}, t_{wait} was varied from 1 ms to the point where $t_{pulse} + t_{wait} = 30$ ms in increments of 1 ms. Attempts to fit the model to the different values of t_{wait} for a single t_{pulse} produced wildly disparate values in the fitting parameters η and τ due to the large correlation between them. Only a global fit of the model to all of the data simultaneously reliably converges to more realistic results although the correlation between τ and η remained very high (0.988). The resulting fitting parameters were determined as $\eta = 0.36 \pm 0.01$ and $\tau = 2.6 \pm 0.1$ ms. The large correlation between these states indicates that it is unwise to trust them individually. Fixing the time constant to the value obtained from the current and field measurements yields $\eta = 0.332 \pm 0.001$. This leads to the conclusion that, even with the large correlation between τ and η, the magnitudes of the separations in the model are overestimated by a factor of three.

The model does produce good agreement with the data but to obtain a time constant of the field independent to η from a fit, further investigation would be needed at pulse times more comparable with the time constant. The issue in this regime is having insufficient contrast to extract the cloud positions.

It should be noted that while the batteries providing the current were measured to be close to capacity voltage, they were not tested under load. Subsequent charging of the batteries did seem to yield larger separations (and correspondingly η closer to unity) although results are not presented here. This would explain some fraction of the discrepancy between the fitted η and its expected value of one. Another possible reason for this deficiency is that the magnetic coil modelling has not taken into account the finite size of the coil, such that the turns are spatially distributed. Preliminary calculations imply that the created gradient could be as low as 65% of that expected initially.

Other possible deficiencies of the model include the possibility that the eddy currents produce a field of such magnitude that the field follows a double exponential loading curve rather than the single exponential stated here. Insufficient sensitivity to the single exponential leaves little hope of detecting such effects at present.

Figure 4.9 shows the equivalent results for the $F = 2$ state. For this fit, the parameters were determined to be $\tau = 2.8 \pm 0.3$ ms and $\eta = 0.37 \pm 0.01$ with a correlation of 0.990 meaning that, again, the individual values should not be trusted. We can see clearly the issues regarding the spatial limit of the beam here, as the $m_F = \pm 2$ states fall rapidly off of the beam before a reasonable contrast can be established.

Results for configuration A are shown in Figure 4.10 for the F = 1 state. These results were obtained after one another so any effect of the batteries should be common to both. The limited number of detected

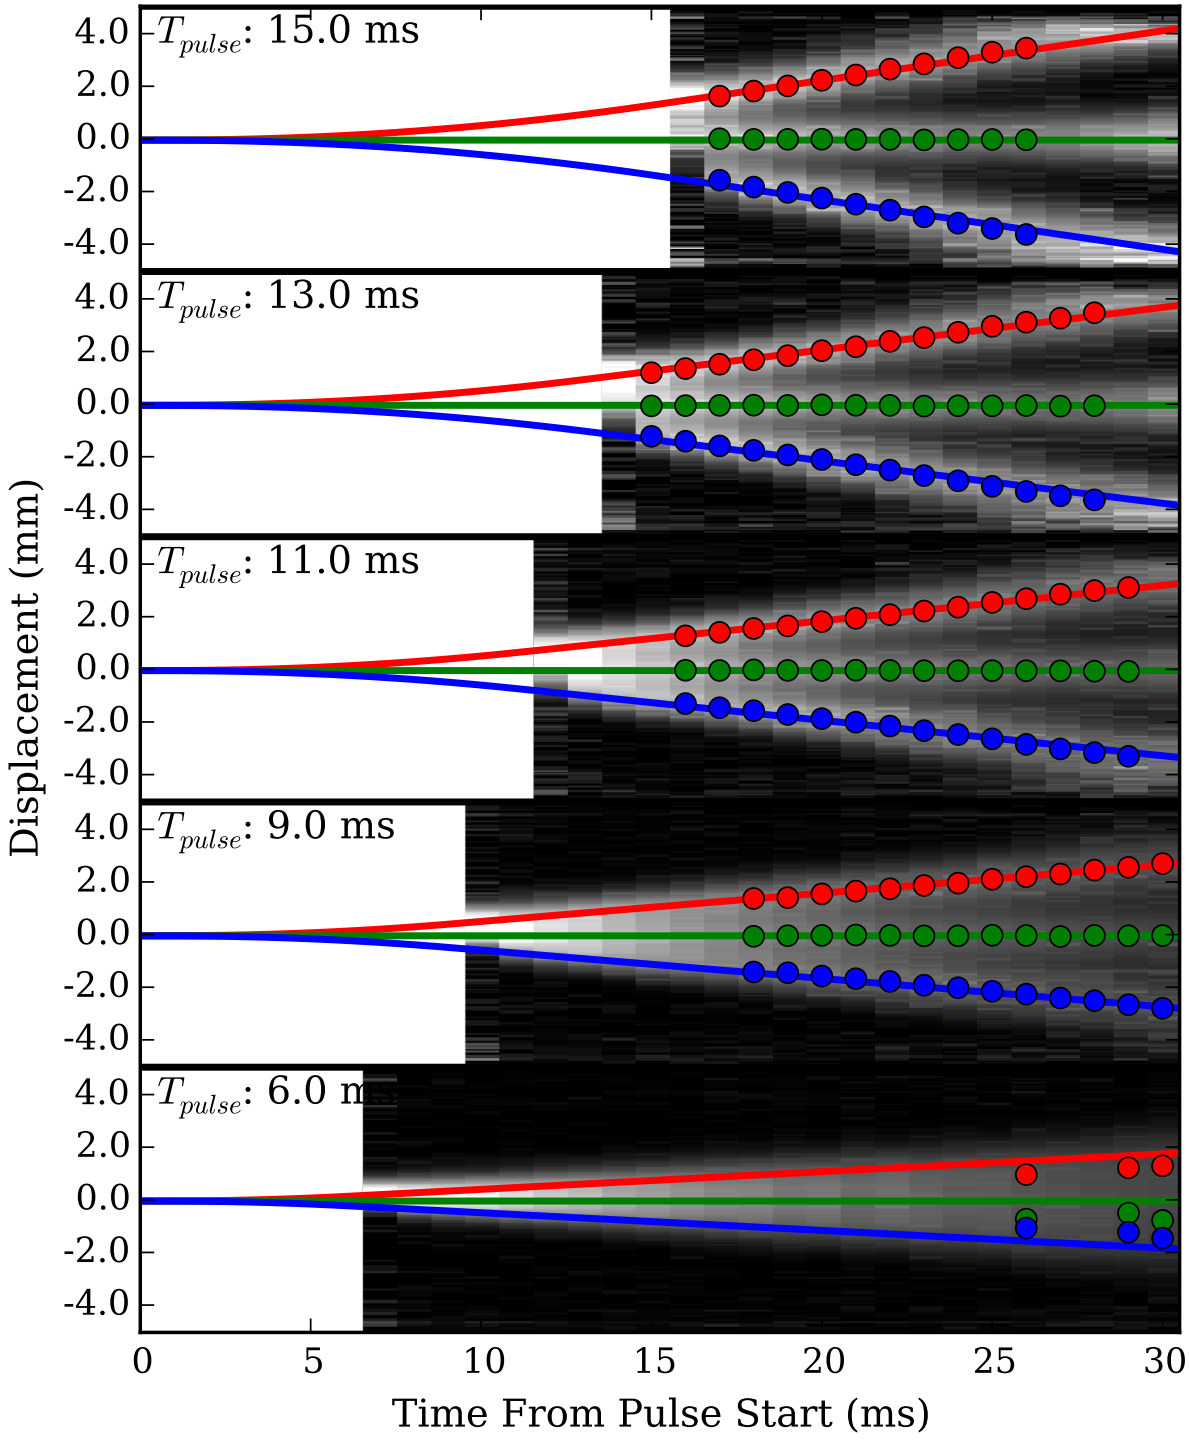

Figure 4.8: Select results from the Stern-Gerlach experiments conducted on the $F = 1$ state in configuration B.

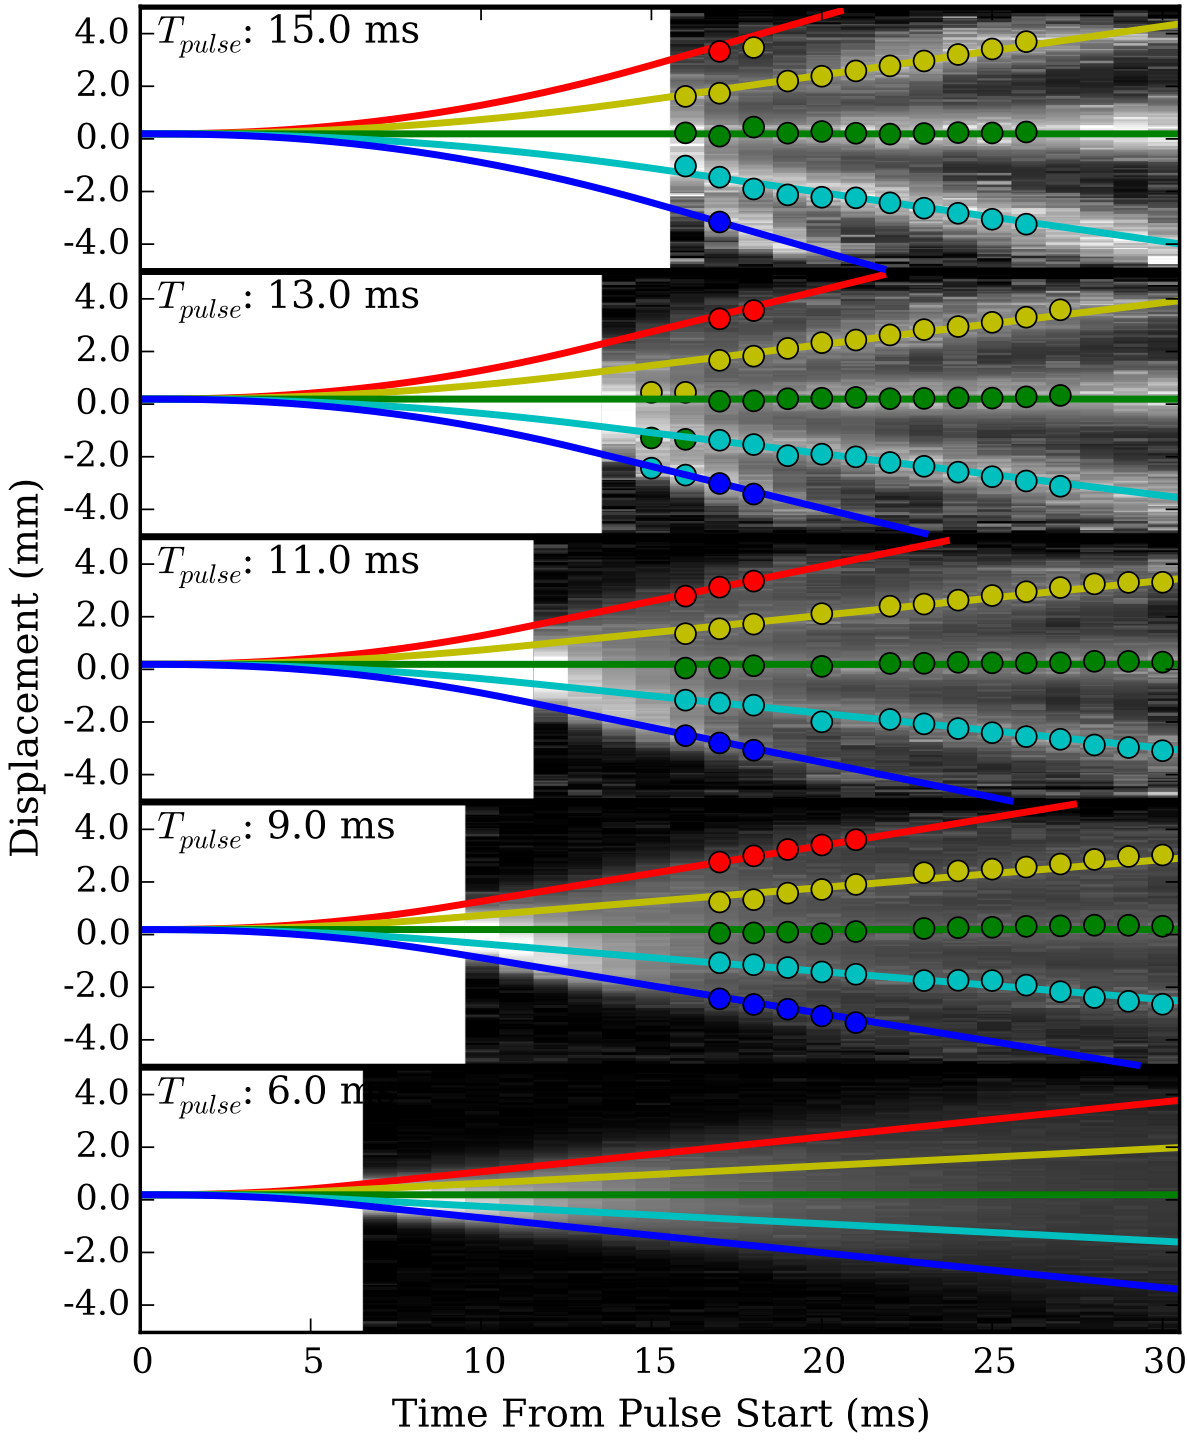

Figure 4.9: Select results from the Stern-Gerlach experiments conducted on the $F = 2$ state in configuration B.

cloud centres means the fit does not converge with both η and τ as free parameters. By fixing τ to 1 ms as obtained earlier, the fitted η is 0.47 times that obtained for configuration B. The expected ratio of these quantities from the calculated field profiles and steady-state currents should be close to one. The most likely explanation seems to be related to the nature of the eddy currents in the chamber.

Quadratic Zeeman Shift Results

Due to the partially drained batteries, the above results for configuration A do not show a significant effect from the quadratic Zeeman shift. Here, we present a limited dataset taken when the batteries were fully charged of the F = 2 states. Figure 4.11 shows the pictures of the atoms and the plotted cloud centres as a function of t_{wait}. Due to insufficient contrast and unequal populations, the fitting routine did not produce satisfactory results so the cloud centres are determined by eye. It is obvious from the images that the cloud spacing is unequal, and this is also seen in the plotted cloud centres. The left-most cloud in the images has a much larger separation from its nearest neighbour than the rest of them. This cloud is therefore easily identifiable as the $m_F = 2$ state, which is unaffected by higher order Zeeman shifts. It can also be seen that the middle state, ($m_F = 0$) has a finite velocity. This is a clear indication that the field used in this pulse produces a significant higher-order Zeeman shift on the clouds. It is clear that as well as hampering the contrast between some of the states as a function of time, the unperturbed trajectory of the $m_F = 2$ state means that the spatial limit of the imaging beam is reached too soon for sufficient contrast for the other states. It is for these reasons that configuration B was adopted.

4.4 Concluding Remarks

In this chapter, we have reviewed the basic theory on the use of a Stern-Gerlach pulse to spatially separate the m_F states of a cold atom cloud. We gave special attention to some of the technical aspects including the ideal coil parameters and the coil switching mechanism. Results were then presented for two configurations, with a single coil and three coil set-up. We conclude that the developed theoretical framework does sufficiently well to qualitatively describe the behaviour of the clouds but more of the model parameters need to be verified to compare the results quantitatively. For the practical purpose relating to the remainder of the experiment, the state separation provides sufficient contrast for accurate readout of the populations in the different states despite the limitations created by the clipped imaging beam.

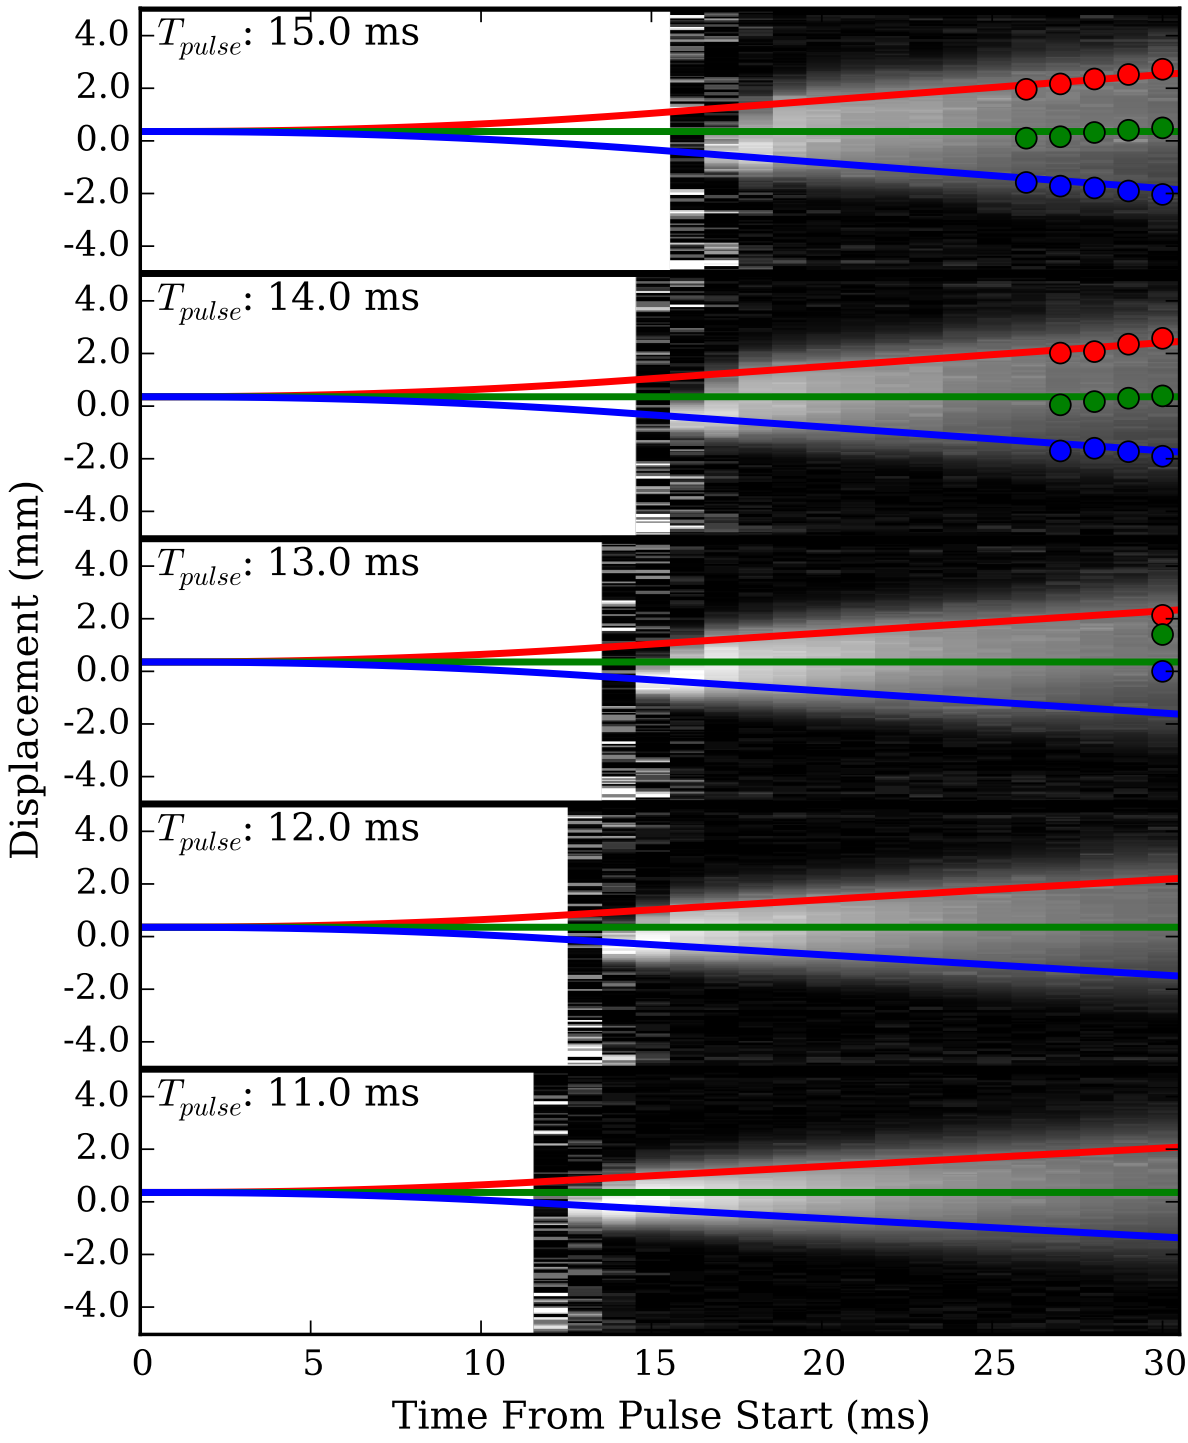

Figure 4.10: Select results from the Stern-Gerlach experiments conducted on the $F = 1$ state in configuration A.

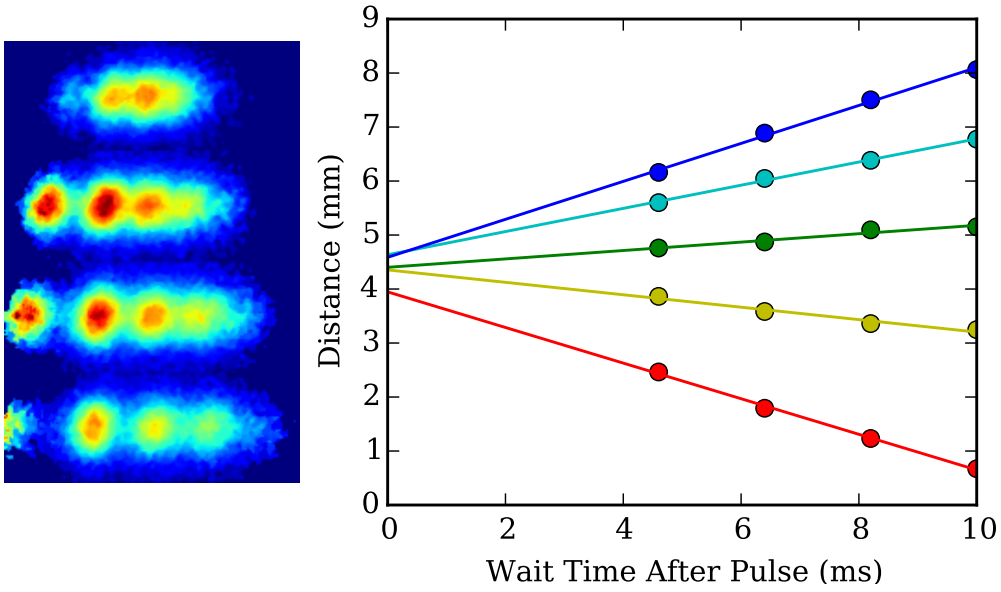

Figure 4.11: Experimental results illustrating the effect of the quadratic Zeeman shift on the state separation. (left) Absorption images of atoms for various t_{wait} after the pulse. (right) Plotted peak centres with a linear fit. The large distance between the left (red) cloud identifies it as the $m_F = 2$ state. Note also, the non-zero gradient for the $m_F = 0$ state (green).

CHAPTER 5

STATE PREPARATION - OPTICAL PUMPING

After the sub-Doppler optical molasses phase, the atoms are roughly evenly distributed between the different m_F states within the $F = 2$ manifold. Since the experimental process outlined in Chapter 1 is to see the resultant transported distribution starting from an initial state, we must be able to prepare the cloud into a single m_F state.

Most specific m_F state preparation schemes in cold atoms systems target either the $|F, m_F\rangle = |F, 0\rangle$ state using π polarised light or $|F, m_F\rangle = |F, \pm F\rangle$ using circularly polarised light. Due to the topology of our network, it is of more interest to look at the movement through multiple sites and as such, we use the latter. We will consider the target state for preparation as $|F, m_F\rangle = |F, +F\rangle$ utilising σ^+ polarisation as the pumping light. However the arguments are general to having a target state of $|F, m_F\rangle = |F, -F\rangle$ utilising σ^- polarisation.

It should be noted that the angular momentum projection states labelled by m_F require a designated quantisation axis, which can be taken arbitrarily. All of the work in this chapter is discussed from the point of view of a quantisation axis along an applied bias field. This makes the m_F states stationary states to the Zeeman Hamiltonian and thus only the light causes exchange between the states. The technical details of the creation of the bias coil are important to the induced radio-frequency (RF) coupling discussed in Chapter 6 and so are left to that chapter. Suffice to say that the field is established during the optical pumping window and has a magnitude of approximately one gauss.

In this chapter, we briefly outline the basic concepts in the preparation of a sample of cold atoms into the $|F, m_F\rangle = |F, F\rangle$ state before looking at our specific schemes used for the $F = 1$ state, which is used in the three level simulation results, and for a basic scheme for $F = 2$. We then present results from these two pumping schemes. The $F = 2$ pumping scheme was not used in the overarching experiment presented in the

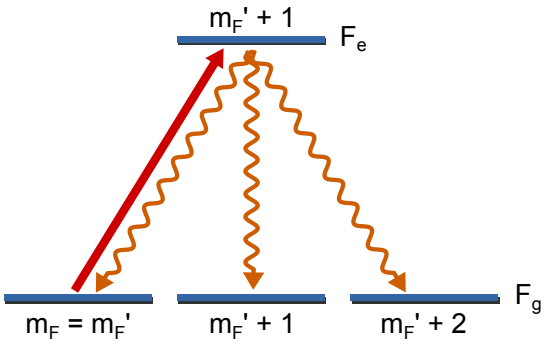

Figure 5.1: A single step of the optical pumping scheme, comprising a directional absorption followed by random emission of a photon.

rest of the thesis but is included for comparison to the $F = 1$ scheme.

5.1 Optical Pumping Overview

The optical pumping scheme requires an excited state manifold with many m_F states and a finite lifetime. A single step involves the absorption of a σ^+ photon and the subsequent spontaneous emission of a randomly polarised photon. Due to the selection rules as a result of conservation of angular momentum, the resultant state from the absorption has an m_F one higher than the starting state. As the upper levels have a finite lifetime, the atom will decay back into one of the lower energy levels, emitting a randomly polarised photon as it does. Again, due to the selection rules, the m_F level of the resulting state will depend on which polarisation was emitted. This process is illustrated in Figure 5.1. The polarisation of the spontaneously emitted photon determines the resultant state of the single step, with σ^-, π and σ^+ polarisations leading to final states $m'_F = m_F + 2, m_F + 1$ and m_F respectively. We can see that subsequent applications of this single absorption-emission process will move the atom to higher m_F states. The relative probabilities of these different polarisation emissions can be determined from considering the relative dipole transition strengths between them. For different state transitions on the same hyperfine transitions, this is given by the particular transitions' Clebsch-Gordan coefficients [36].

An important decision in this pumping scheme is the selection of the excited state manifold. For an electric dipole transition (which is the most commonly used in laser excitation) conservation of angular momentum means that coupling can only be achieved with an excited level with angular momentum $F' = F - 1$, F or $F + 1$. Here, we will briefly discuss these three options in the context of an optical pumping scheme. The schematic representation of each choice is illustrated in Figure 5.2.

The $F' = F - 1$ is an inadequate choice due to the lack of σ^+ absorption path from the $|F, F-1\rangle$ state. This means that this state is 'dark' to the pumping light. This causes atoms to accumulate in this state instead of the desired $|F, F\rangle$ state. The $|F, F\rangle$ state will be populated partially by emission from the $|F', F-1\rangle$ state but not with a large purity.

The $F' = F + 1$ and $F' = F$ states are both sufficient for state preparation, however the latter has the advantage of the target state being a dark state to the pumping light. In the former case, the atoms in the target state will continuously scatter photons on the $|F, F\rangle \rightarrow |F+1, F+1\rangle$ transition (referred to as a cycling transition). This is problematic as each absorption induces a momentum kick to the atom in the direction of travel of the absorbed photon. This leads to a accumulated velocity of the cloud along the axis of the pumping beam. Additionally, each spontaneous emission gains the atom a photon recoil momentum kick in a random direction. This leads to an increase in the velocity spread of the cloud in all three dimensions leading to an increased temperature.

In contrast, for the $F' = F$ case, the target state is a dark state to the light as there is no level to absorb a σ^+ photon to and thus this cycling does not happen. It is clear with the arguments made so far that the $F' = F$ state is the best choice for the σ^+ pumping light.

There is however, an additional consideration. When the atom spontaneously emits from an excited level, there may be a finite probability of emitting to a different set of ground states. In Rubidium for example, the other hyperfine ground state has an energy several GHz away from the intended ground state. If the pumping light is the only light used, any atoms that fall to this state will be lost from the sample and not be further pumped. The possibility of emission to these other levels necessitates extra light sources to return the atoms to the pumping cycle. This is a similar role to the repumper in the MOT cooling mechanism. The selection of the transition and polarisation of this beam is of crucial importance to provide efficient pumping.

5.1.1 Optical Pumping Recovery Beam

To discuss this 'recovery' beam further, we will mainly restrict discussion to our specific situation of Rubidium 87, the target pumping state of $|1, 1\rangle$ and a main pumping beam of σ^+ polarisation on the $F = 1 \rightarrow F' = 1$ transition. For each m_F state in the excited $F' = 1$ manifold, there is a finite probability of emission into the $F = 2$ ground state manifold. We must therefore implement a recovery beam to re-excite the atoms from the $F = 2$ ground state with the aim of repopulating the $F = 1$ ground state.

To achieve this, the selection of excited state for the recovery beam must be electric dipole accessible to both the $F = 2$ ground state (so that it is possible to excite the transition) and the $F = 1$ ground state

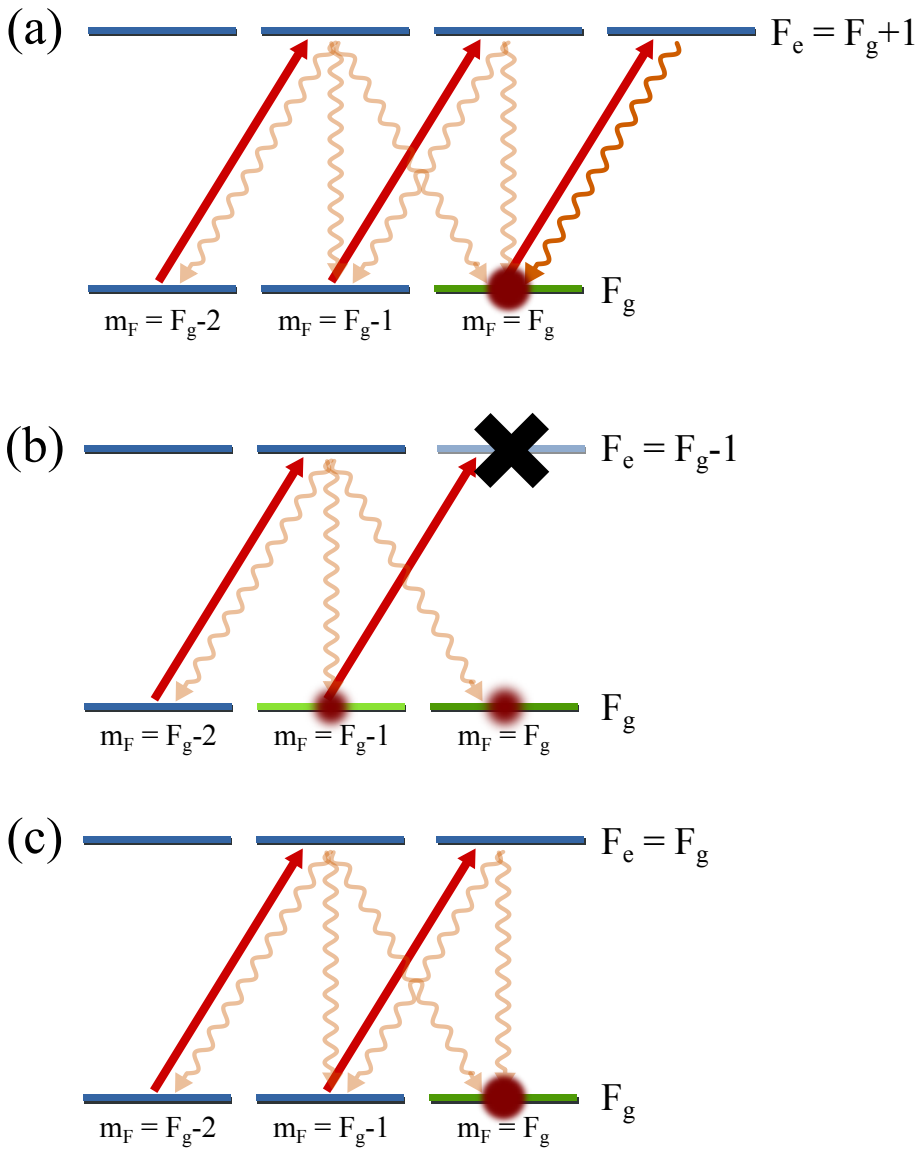

Figure 5.2: Illustration of the different choices for optical pumping levels. (a) The case where $F_e = F_g + 1$. The sample can be pumped but the target state cycles, leading to heating. (b) The case where $F_e = F_g - 1$. The sample will not be successfully pumped into a single state as both the target state $|F_g, F_g\rangle$ and the state $|F_g, F_g - 1\rangle$ are dark to the pumping light. (c) The ideal case where $F_e = F_g$. Here, the atoms are pumped to the dark target state.

(so that the atoms can spontaneously emit into it). This limits the options to $F' = 1$ or 2. An important characteristic to consider is the branching ratio from these upper states. For $F' = 1$, the probability of emission into the $F = 1$ and $F = 2$ ground states is 5/6, 1/6 respectively, compared to 1/2, 1/2 for the $F' = 2$ intermediate state[1]. This larger branching ratio to decay to the $F = 1$ ground state for $F' = 1$ means that, on average, fewer absorption-emission pairs are required to repopulate the $F = 1$ ground state.

With regards to polarisation of the recovery beam, it is tempting to consider choosing the polarisation to help the directionality of the pumping process, however, this is not possible when using an excited manifold of equal F to the recovery level as, using pure σ^+ polarisation, this would lead to a dark state ($|2, 2\rangle$) in the recovery manifold. Equally, if the angular momentum of the excited level is lower than that of the recovery level, then π polarisation is also insufficient as the $|F, \pm F\rangle$ states would be dark to the light due as there are no excited states with equal angular momentum projection. For π polarisation with $F' = F$, the $|F, 0\rangle$ state is dark due to the zero Clebsch-Gordan coefficient as it would require a simultaneous flip of the electronic and nuclear angular momentum projections.

As a result of these restrictions, to avoid having any of the m_F levels be a dark state in the recovery level, the recovery light must contain at least two of σ^+, σ^- and π polarisation[2]. A simple example of how this can be done is by sending a linearly polarised beam at an angle to the quantisation axis between 0 and $\pi/2$. The angle of polarisation should also be set such that the quantisation axis is not within the polarisation plane.

5.2 Rubidium 87 $F = 2$ Optical Pumping

As the atoms are in the $F = 2$ manifold after the optical molasses stage, initial experimental tests of the radio-frequency coupling (see Chapter 6) were conducted within the $F = 2$ manifold before the pumping system for the $F = 1$ state was conceived and constructed. To test the RF system required some rudimentary state preparation which was conducted with this simple but sub-optimal scheme. While no results are presented for the $F = 2$ RF coupling in this thesis, the pumping scheme used is included as a comparison between the cycling and non-cycling schemes discussed above.

The pumping light was provided by the same beam used for the absorption imaging system. This beam was already circularly polarised along the quantisation axis and tuned to the $2 \rightarrow 3$ transition. As mentioned, this choice of excited state levels leads to heating of the cloud during the pumping time as well as 'pushing'

[1] This is found by comparing the sum of the transition strengths for each of the possible decay paths to the different ground states.

[2] Strictly speaking, there will always be a superposition state that is dark to the incident light regardless of the polarisation, however these states are not stationary in the Zeeman Hamiltonian and can therefore be neglected.

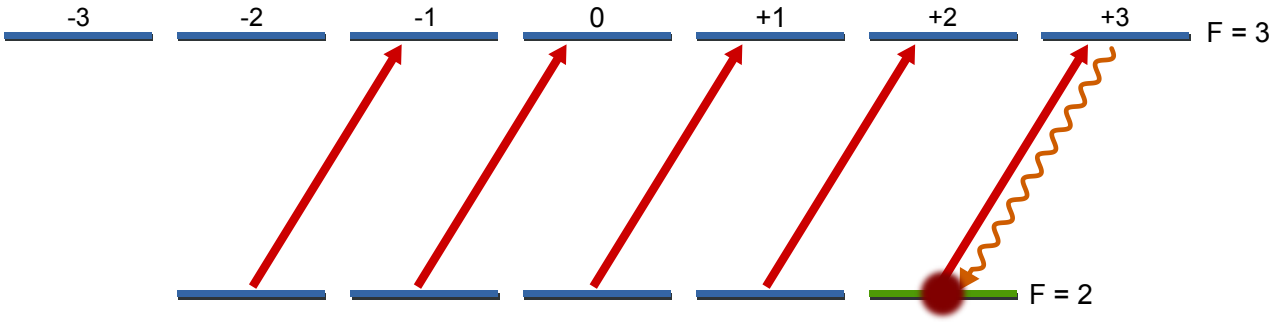

Figure 5.3: The F = 2 pumping scheme. The $|2,2\rangle$ target state is continuously cycled, heating the atoms during the pumping.

the cloud along the axis of the pumping beam. The pushing effect of the cloud is along the imaging axis and is therefore difficult to resolve. We do not see any noticeable effects of it in the imaging.

Whilst the $F = 1$ ground state is not dipole accessible from the excited $F = 3$ state, there is a small chance (one in several thousand) that, due to the linewidth of the transition, the atom will instead be excited to the $F = 2$ excited state, from which it can decay to the $F = 1$ level. Despite the probability being very low, we therefore leave the MOT repumper on as a recovery beam for atoms that end in the $F = 1$ ground state. This is unlikely to be necessary for short pumping times but was done anyway as it had no experimental overhead.

This scheme is very simple as it involves no additional beams than were not already in use for the MOT and imaging set-up. Figure 5.3 illustrates the scheme.

5.2.1 Experimental Results

Qualitative results are presented for the general pumping behaviour as well as more quantitative results on the thermal heating effect of the light on the atoms.

Due to the technical challenges in the Stern-Gerlach approach, accurate measurement of the state populations for the $F = 2$ manifold is difficult as the spatial separation limitation of the small imaging beam quickly becomes an issue with the five states. Coupled with the heating of the cloud via the pumping mechanism, the exact time dynamics of the pumping behaviour cannot be reliably obtained with the current setup. To show qualitatively that the pumping works, Figure 5.4a shows several absorption images for the Stern-Gerlach separated atom clouds with different pumping times ranging from 0 to 1 ms. Figure 5.4b shows the same experimental timings without the Stern-Gerlach pulse for comparison. By measuring the width of the clouds, the temperature can be inferred from the known time of flight from the end of the pumping (16.6ms); we

plot the temperature as a function of pumping time in Figure 5.4c.

The continuous spontaneous emission of photons leads to a spread in velocity space following a random walk. In one-dimension, this spread is proportional to the square root of the number of photons scattered. The number of photons scattered is linearly dependent on the pumping time. This means the spread of the velocity due to the scattered photons is proportional to the square root of the pumping time, $\Delta v \propto \sqrt{t_{pump}}$. As the temperature (inferred from the width of the cloud after expansion) is proportional to the square of the velocity spread, $T \propto \Delta v^2$, we expect the temperature of the cloud to vary linearly with the pumping time. The results presented in Figure 5.4c seem to agree with this. The intensity of the imaging beam incident on the chamber is approximately $23\,\mu W/cm^2$. With an intensity significantly below the saturation intensity of the transition ($1.669\,mW/cm^2$ for the $m_F = 2 \rightarrow m_F = 3$ cycling transition [36]), the scattering rate can be approximated by $(\Gamma/2) \cdot (I/I_{sat})$ which would yield a rate of 262 photons/ms. If we assume the cloud starts at zero temperature, the temperature of the cloud as a function of pumping time follows

$$T = \frac{m v_r^2}{k_B} R t_{pump} \tag{5.1}$$

where R is the scattering rate and v_r is the recoil velocity of a single photon. The expected gradient between the temperature and pumping time is $95\,\mu K/ms$. Whilst this is three times the measured value, it is expected be an overestimate due to a potential overestimation of the intensity of the pumping beam due to absorption from the vacuum chamber window as well as the assumption that all of the atoms are cycling the low saturation intensity $m_F = 2 \rightarrow m_F = 3$ transition at exactly zero detuning. As the atoms start distributed amongst all of the m_F states, some of the initial cycles will not be on the cycling transition and will have a different scattering rate due to the different saturation intensity and different detuning due to the differential Zeeman shift as a result of the bias field. Potential detuning from the cycling transition will also cause a drop in the expected scattering rate.

The general conclusion is that the heating from this rudimentary optical pumping scheme makes accurate quantitative measurements of the state numbers very challenging due to the limited aperture size. If further work were to be carried out on the $F = 2$ levels, a pumping scheme centred around the $F = 2 \rightarrow F' = 2$ transition would be necessary for accurate state readout.

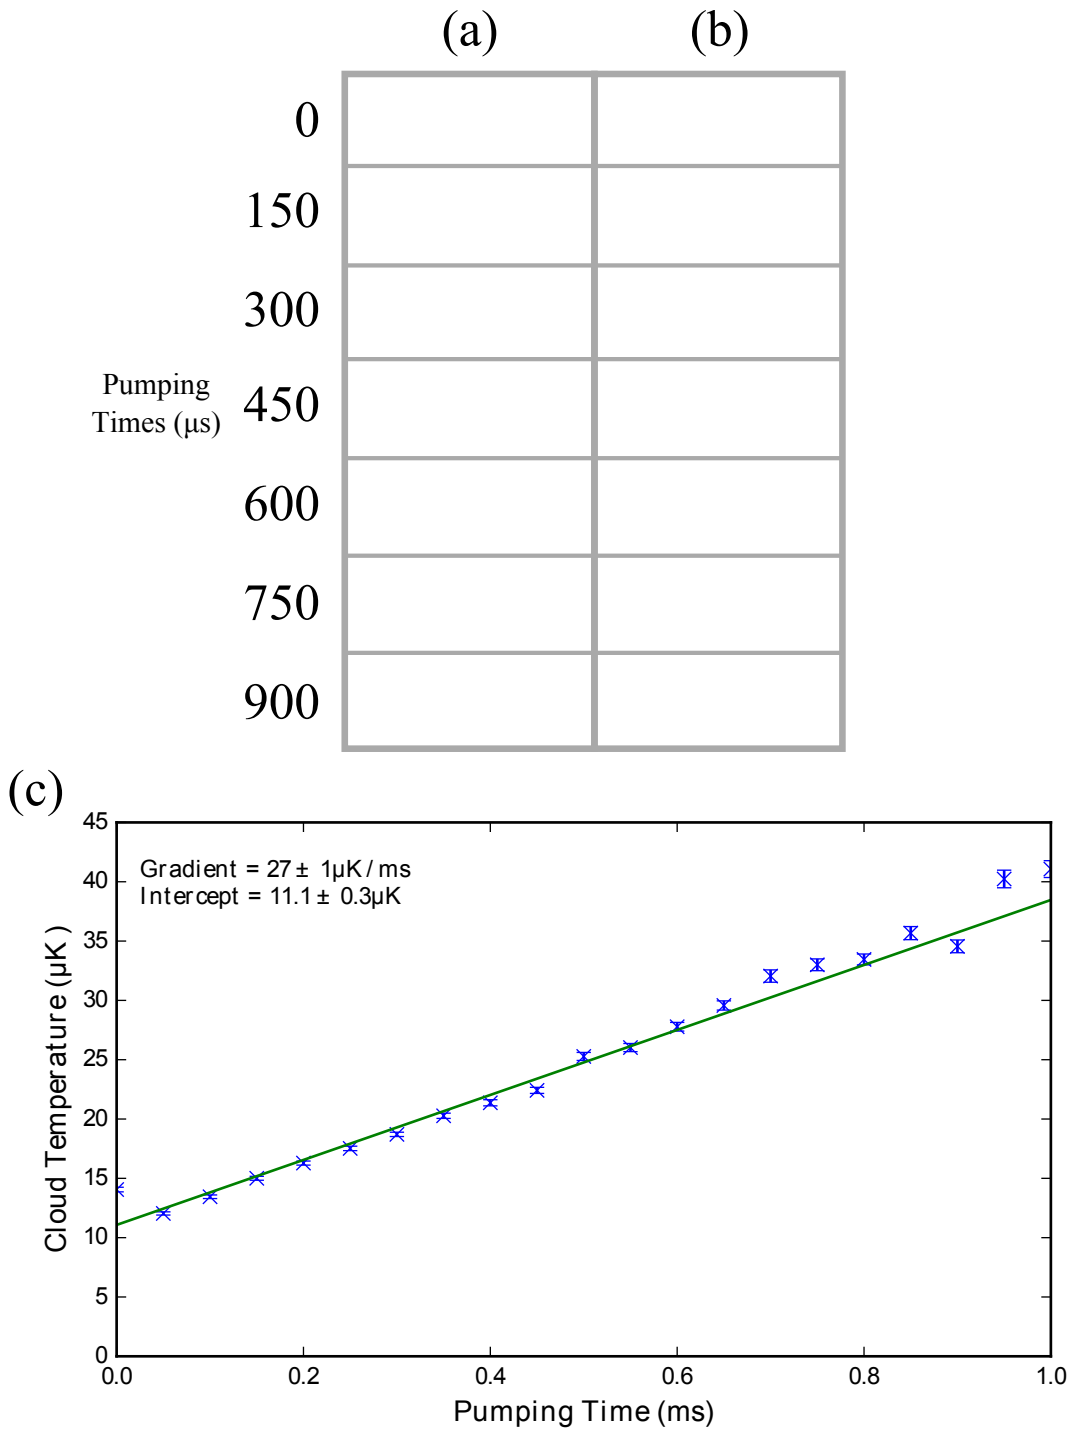

Figure 5.4: Experimental Results of the $F = 2$ optical pumping process. (a) Stern-Gerlach separated images with imperfect contrast showing the general pumping trend. (b) Images of clouds subject to the exact same stimulii and timings as in *(a)* but without the Stern-Gerlach magnetic separation pulse applied. (c) The temperature as a function of the pumping time.

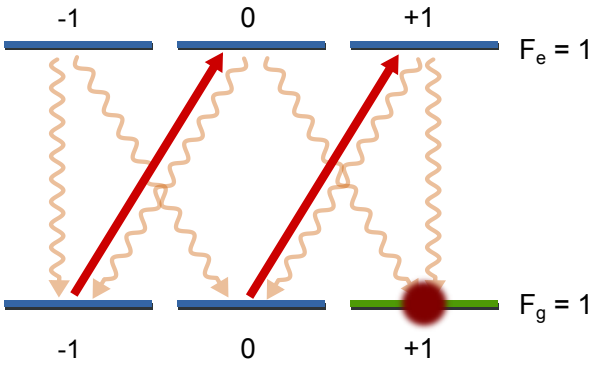

Figure 5.5: Optical pumping scheme for the $F = 1$ states. The target $|1, 1\rangle$ state is dark to the pumping light.

5.3 Rubidium 87 $F = 1$ Optical Pumping

Pumping to the $|1, 1\rangle$ state requires the introduction of two new beams to the setup. The σ^+ pumping beam is set to the $F = 1 \rightarrow F' = 1$ transition and is derived from light shifted using an AOM $-2 \times 78.5\,\text{MHz}$ from the repumping light which is locked to the $F = 1 \rightarrow F' = 2$ transition. The pumping beam is delivered to the apparatus through the same fibre as the imaging beam and expanded to the same size of 2.5 cm diameter and sent along the quantisation axis of the experiment. Like the imaging beam, it is clipped to $\sim$1 cm diameter by the chamber conic section. Circular polarisation is achieved using a Quarter-Wave-Plate (QWP) after the fibre launch.

Recovery and initial transfer of atoms from the $F = 2$ ground state to the $F = 1$ ground state is accomplished using a beam on the $F = 2 \rightarrow F' = 1$ transition. This beam is derived from the master light locked to the $F = 2 \rightarrow F' = 3$ transition, shifted $-2 \times 212\,\text{MHz}$ by an AOM. We refer to this beam as the depumper as it has the opposite function to the MOT repumper. The beam co-propagates with the repumper and is incident on the position of the atom cloud at approximately 45° to the quantisation axis defined by the bias field. This allows the presence of both σ and π polarised light. The pumping scheme is illustrated in Figure 5.5 and the beam derivations are displayed in Figure 5.6. This pumping scheme is used for state preparation for the simulations discussed in subsequent chapters.

Considering just the $F = 1 \rightarrow F' = 1$ transition, the system is quite simple. All of the non-zero Clebsch-Gordan coefficients are $\sqrt{5/24}$ and as mentioned previously, the $|1, 0\rangle \rightarrow |1, 0\rangle$ transition is not electric dipole accessible, meaning that each of the excited $|1, m_F\rangle$ states can decay into two states with equal probability.

We briefly consider the case where there is both σ^+ and σ^- light present. This is possible if we are unable to fully ensure the polarisation of the light propagating along the quantisation axis. We shall also relabel

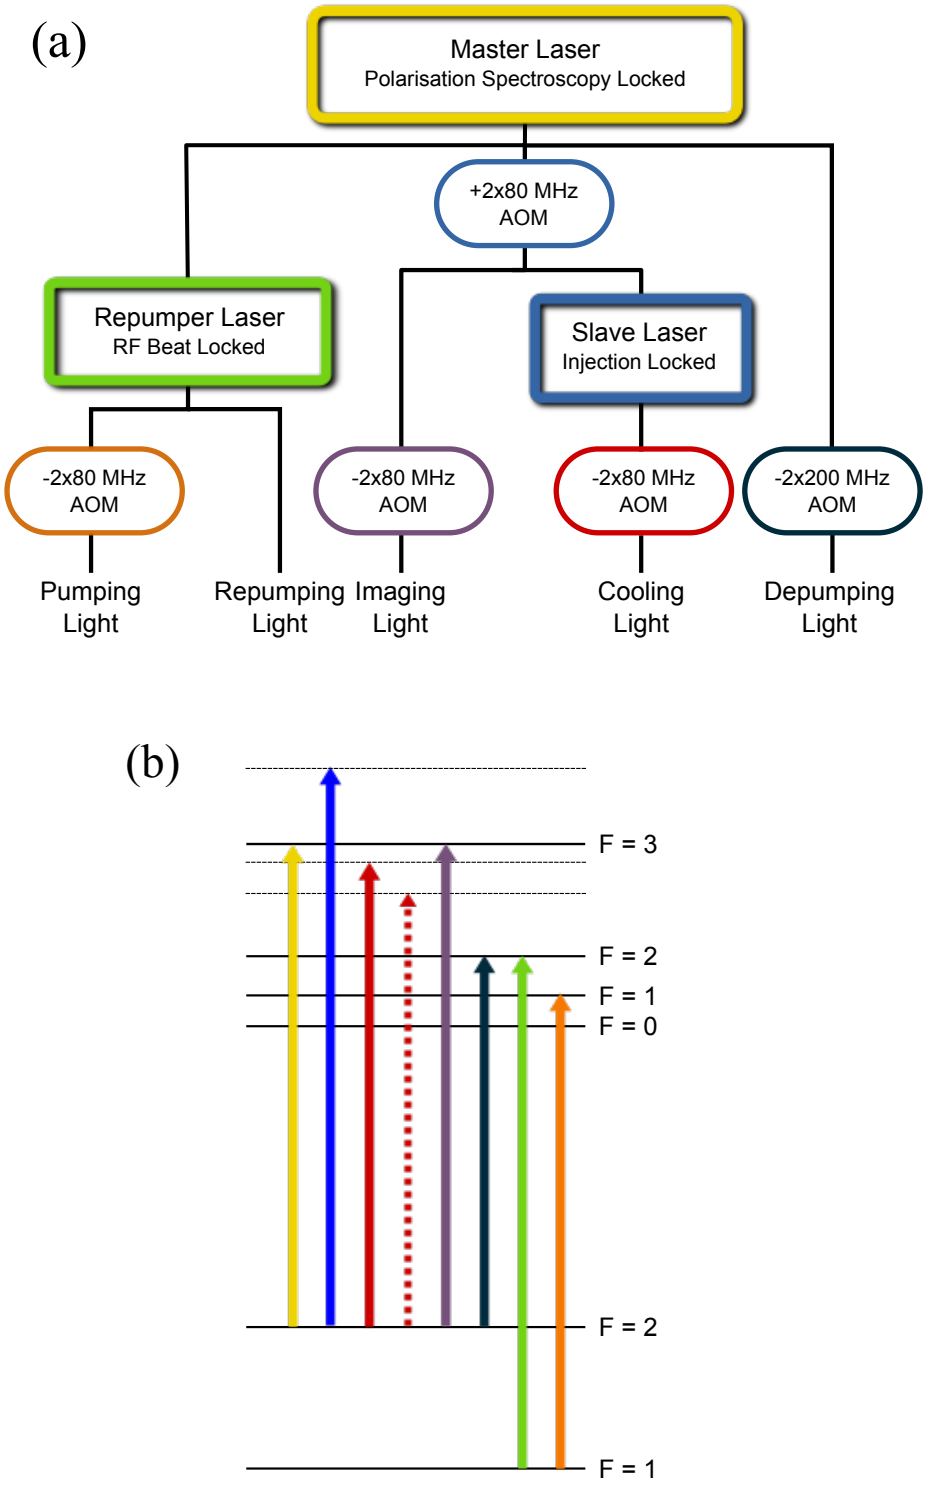

Figure 5.6: Illustration of the different beams derived for the purposes of creating the cold atoms cloud and state preparation. (a) Shows a block diagram of how the different frequency beams are derived with respect to the other beams. (b) Shows the energy level diagram of the Rubidium 87 D2 line with the appropriate beams in place.

the ground state sublevels; $|g, m_F\rangle$ and the excited levels; $|e, m_F\rangle$ to avoid confusion. In the absence of π polarisation, it can be seen qualitatively that the $|g, 0\rangle$ state will be rapidly depleted as the population is pumped to the $|e, \pm 1\rangle$ states and then only decays back to the $|g, 0\rangle$ with 50% probability or falls into the $|g, \pm 1\rangle$ states with 50% probability. Once in the $|g, \pm 1\rangle$ states, there is no path to repopulate the $|g, 0\rangle$ state. With π polarisation present, this will lead to pumping to the $|g, 0\rangle$ state.

The steady state ratio of populations in the $|g, -1\rangle$ and $|g, +1\rangle$ states will depend on the relative strength of the $\sigma^{\pm}$ light.

To quantify the time dynamics of the population, we can use a simplification of the Optical Bloch Equations (OBE) [35] which comes from solving the Liouville-von Neumann equation [37] with a Hamiltonian of optical couplings in the presence of spontaneous emission. As the Rabi frequency of the pumping light is much smaller than the natural linewidth of the excited states ($\sim$100 kHz cf. 6 MHz), there is not a significant build up of population in the excited states. Additionally, this condition means that there is little coherence buildup between the levels. With these simplifications in place, we can reduce the problem to a set of classical rate equations considering just the populations of the ground states where the terms determining the decrease in population are due to coupling to the excited levels and the population increases as decays from these excited levels. If we assume for now that all of the transitions have zero detuning, the rate equation for the $|g, -1\rangle$ level is given by

$$\frac{dN_{-1}}{dt} = -(R_\pi + R_+)N_{-1} + (R_- N_0 + R_\pi N_{-1}) + (R_+ N_{-1} + R_- N_{+1}) \tag{5.2}$$

where we have introduced the quantities $R_{\pm/\pi}$ as the scattering rates of the transitions associated with $\sigma^{\pm}$ and π polarisations and N_{m_F} as the fractional populations of the states. In the low intensity limit, the scattering rate is proportional to the intensity of the light with that polarisation [36] ($R_{\pm/\pi} \propto I_{\pm/\pi}$). The first set of brackets denotes the loss to the excited states, the second set is the contribution emitted from $|e, -1\rangle$ and the third the contribution emitted from $|e, 0\rangle$. Simplifying and using the expression that $N_{-1} + N_0 + N_{+1} = N$, we get the expression

$$\frac{dN_{-1}}{dt} = -(R_\pi + R_+)N_{-1} + R_-(N - N_{-1}). \tag{5.3}$$

Using $R_{\sigma^{\pm}/\pi} = \alpha I_{\sigma^{\pm}/\pi}$, and the sum of the intensities $I = I_{-1} + I_\pi + I_{+1}$, Equation (5.3) has the closed form solution;

$$N_{-1}(t) = a_{-1} + b_{-1} e^{-t/\tau} \tag{5.4}$$

where $a_{-1} = NI_-/I$ gives the steady state solution for the population and $b_{-1} = N_{-1}(t = 0) - a_{-1}$. The

decay time τ is related to the total scattering rate $\tau = 1/R$, where $R = R_{-1} + R_{\pi} + R_{+1}$. Due to the symmetry between the three ground states, the solutions for N_0 and N_{+1} have solutions of the same form with $a_0 = NI_{\pi}/I$ and $a_{+1} = NI_{+}/I$. The fractional purity of the desired prepared state at long times is therefore proportional to the fractional purity of the corresponding polarisation.

5.3.1 Experimental Results

Using the Stern-Gerlach m_F state sensitive imaging for the $F = 1$ state, the efficiency and purity of this scheme can be measured with relative ease. Here, we present experimental results for the $F = 1$ pumping scheme discussed above. We quantitatively show the validity of the derived rate equation as well as draw conclusions for the pumping purity. Note that in practice, these optical pumping scheme setups are very straightforward and are not usually looked into in such detail. We do so here because the Stern-Gerlach state readout allows us to investigate some of the mechanisms behind it.

To characterise the time dynamics of the scheme, the pumping time was varied discretely from 0-300 µs along with three shots at 1 ms to provide a 'long-time' steady state result. The relative populations of the three different m_F states were extracted from the Stern-Gerlach separated clouds. An example of one of these time resolved runs is shown in Figure 5.7. Results were captured for different input polarisations and different powers. The time dynamics of the populations are well fitted by the rate equations determined above.

Quarter Waveplate Angle

In principle, the polarisation of the pumping light is determined by the angle of a QWP to the injected polarisation axis of the optical fibre that delivers the beam to the apparatus. As the beam travels roughly along the experimental quantisation axis, it should consist mostly of σ^+ and σ^- polarisations. The corresponding ratio of σ^+ and σ^- light is determined by the angle of the QWP.

The angle of the waveplate is the first parameter that was changed. The waveplate mount was marked roughly into 5° increments and the state populations as functions of time were recorded for different angles according to these crude markings. These time resolved results were fitted to the rate equation model discussed earlier with the relevant parameters of the total scattering rate R and the fractional intensities I_{+}/I, I_{-}/I and I_{π}/I. The results of this fitting routine as a function of the waveplate angle are shown in Figure 5.8.

It can be seen that the total scattering rate R is broadly constant across the measurements. This is

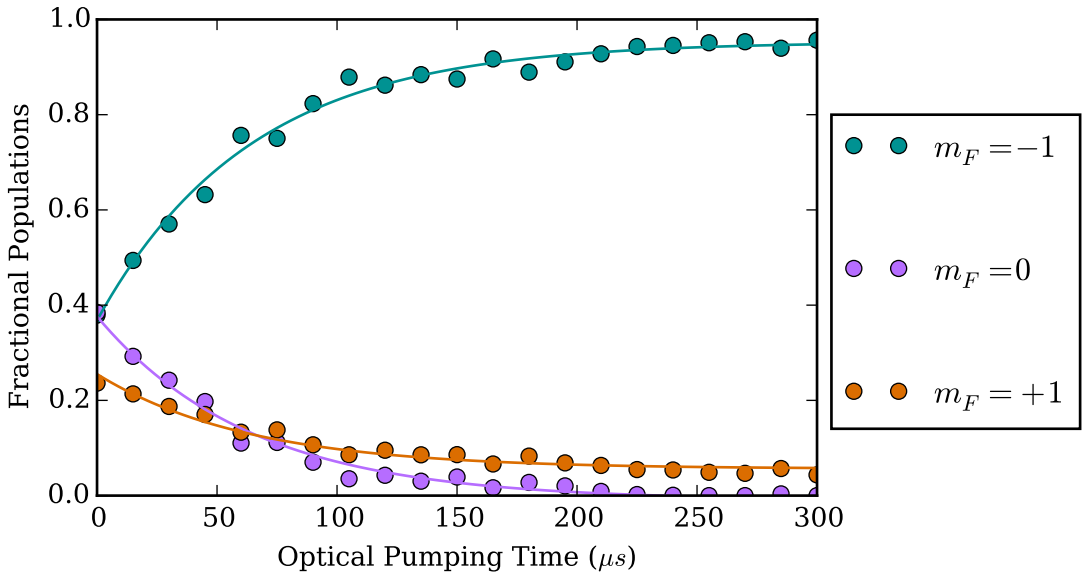

Figure 5.7: Fractional populations of the different states against pumping time with fit of Equation (5.4). Note that the initial state populations are a product of the optical molasses phase and the time range is chosen to illustrate the relevant dynamic behaviour.

expected as the total scattering rate is proportional to the total intensity of light ($R = \alpha I$), which is not varied. The fractional intensities of σ_+ and σ_- light broadly follow the expected behaviour of exchanging between themselves with the QWP angle. However, when the σ^+ light dominates the σ^- light, there also appears to be a small amount of π polarisation which is unexpected.

We attribute this effect to the depumper, which we have neglected in the rate equation formalism above. The depumper has an effectively random polarisation, which can allow decay to the $|g, 0\rangle$ state. We do not have a solid explanation for why this effect seems to be more pronounced when the σ^+ light is strongest. This may be an effect of the depumper polarisation, with an imbalance between the $\sigma^\pm$ polarisations as it is the only element that can break the symmetry. Detailed investigation has not been conducted on this issue.

An additional complication that we have neglected in our treatment is the relative detuning of the transitions. When discussing the relative excitation rates from the ground states, we neglected any detuning of the transitions and just considered the relative powers of the different polarisations. Since the pumping happens in a small bias field ($\sim$1 G), the Zeeman-shift means that each of the transitions has a slightly different frequency. Table 5.1 shows the relative shift of these transition frequencies from the presence of a magnetic field.

With a transition linewidth of 6 MHz, the effect may be present but should be reasonably small. For

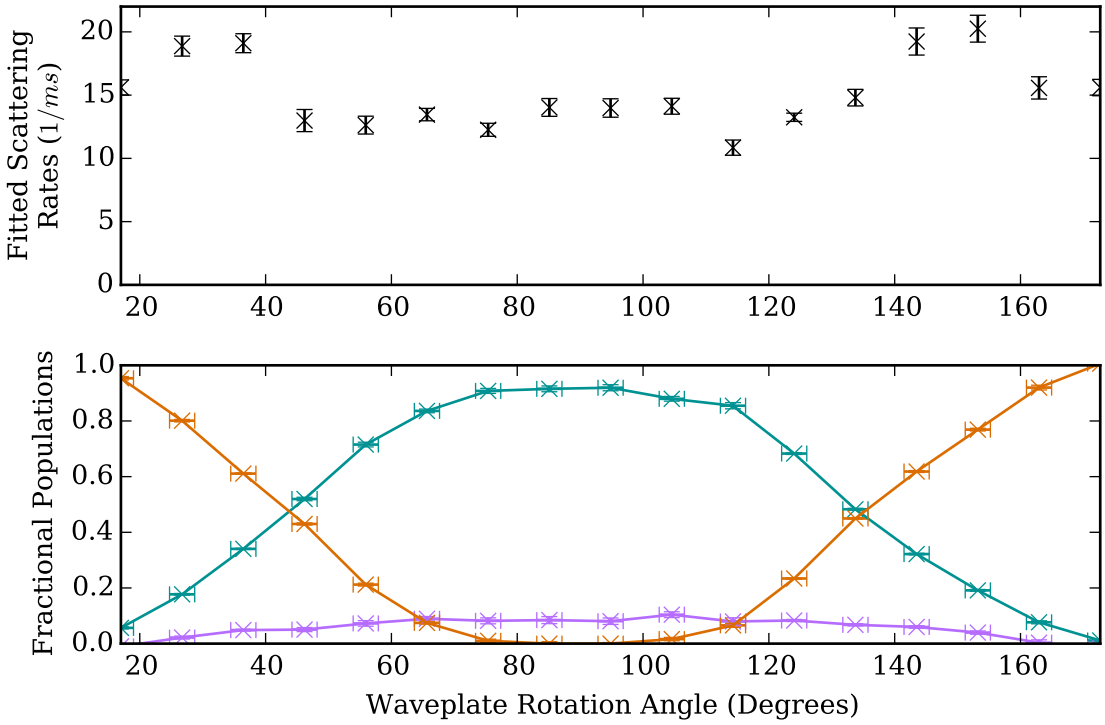

Figure 5.8: Fitted rate equation parameters as a function of waveplate angle. (Upper) The fitted total scattering rate of the pumping light. (Lower) The fitted fractional intensities of σ^- (orange), σ^+ (turquoise) and π (lilac) polarisations.

Ground m_F	**Excited** m_F	**Frequency Shift** (MHz/G)
-1	-1	-1.63
	0	-0.7
0	-1	-0.93
	+1	+0.93
+1	0	+0.7
	+1	+1.63

Table 5.1: $F = 1$ optical pumping scheme transitions frequency shift in the presence of a weak magnetic field.

example, at a field of one gauss, the frequency of the $\sigma^+ : |g, -1\rangle \rightarrow |e, 0\rangle$ transition will be 1.4 MHz detuned from the $\sigma^- |g, +1\rangle \rightarrow |e, 0\rangle$ transition. This will cause a small difference in the scattering rate.

It should be noted that whilst these results show a relatively poor pumping efficiency on the order of 95%, this is unusually poor and mainly an artefact of the time that these results were taken, potentially due to poor fibre polarisation alignment of the pumping beam. During typical operation of the experiment, the state preparation produces purities in excess of 99%. For examples of this, see the results in Chapter 6 and specifically Figure 6.14.

Pumping Beam Intensity

An investigation was also conducted into the effect of varying the intensity of the optical pumping beam. We expect that the relative populations in the steady state should remain unchanged but the scattering rate should increase proportionally to the pumping intensity.

Figure 5.9 shows a selection of the time resolved populations for differing powers whilst Figure 5.10 shows the inverse of the fitted time constants against the intensity of the pumping beam. It should be noted that only the power of the beam is directly measured and the intensity proportionally inferred from the unchanging beam size.

The results broadly follow the expected results with the relationship between $1/\tau$ and P being proportional. When conducting further experiments that require the sample to be pumped, the pumping time is set to be approximately 1 ms to ensure that the pumping is complete.

5.4 Concluding Remarks

In this chapter, we have given an overview of the choices to be made when designing an optical pumping scheme. We have described and presented results from two schemes, one used for the $F = 1$ ground state that is the basis of all later chapters in this thesis, and the other, a basic method for the $F = 2$ ground state, which was used briefly during the building of the apparatus and is presented as a comparison.

We have also presented some quantitative results which exhibit the state readout work presented in the previous chapter. Results were presented for the $F = 1$ scheme in which the atoms were prepared into one of the $|1, \pm 1\rangle$ states with purities in excess of 95%. The purity presented in these results is on the lower end of what is typically possible so we conclude that the state preparation works sufficiently to conduct investigations into the time dynamics of coherent couplings between the states.

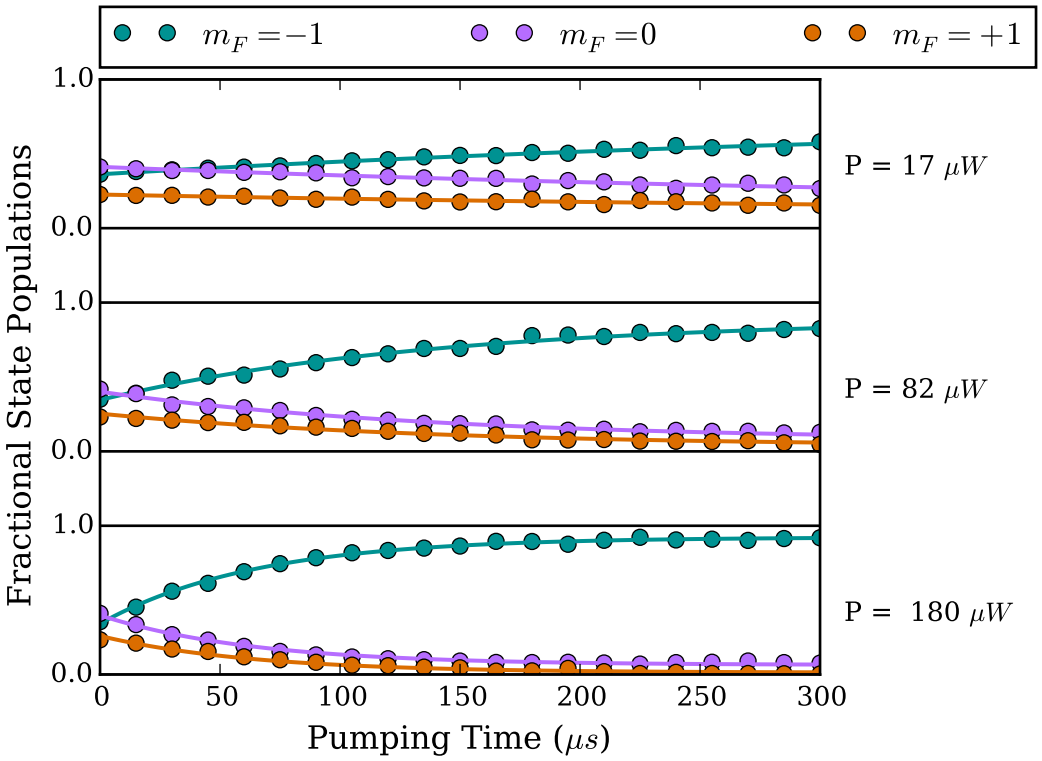

Figure 5.9: Fractional populations as a function of pumping time for varying powers of the pumping beam.

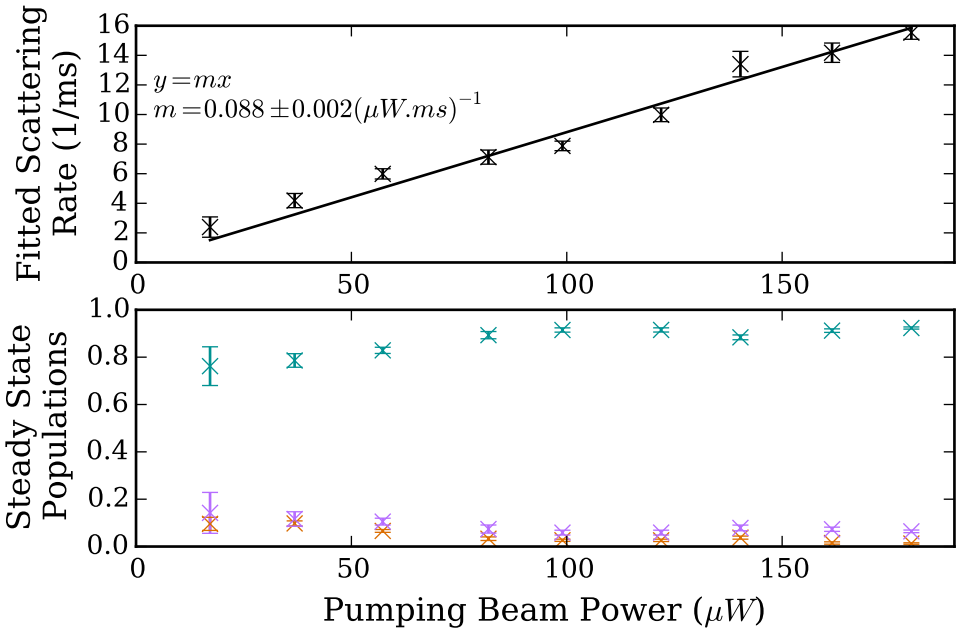

Figure 5.10: Fitted rate equation parameters as a function of pumping beam power. (Upper) The fitted scattering rate with a linear fit. (Lower) The steady state fractional populations, which are roughly independent of the pumping power

CHAPTER 6

NEAREST NEIGHBOUR COUPLING - MAGNETIC DIPOLE TRANSITIONS

With the atoms prepared in the $|F, m_F\rangle = |1, 1\rangle$ state with a bias field on the order of one gauss, an oscillating magnetic field was applied with a frequency close to resonance with the energy level difference between the m_F states. The Hamiltonian for our 3-level system in the rotating frame under the RWA is given by Equation (2.25) but restated here for ease of reference

$$H = \hbar \begin{pmatrix} -\Delta & \frac{\Omega_{RF}}{\sqrt{2}} & 0 \\ \frac{\Omega_{RF}}{\sqrt{2}} & 0 & \frac{\Omega_{RF}}{\sqrt{2}} \\ 0 & \frac{\Omega_{RF}}{\sqrt{2}} & \Delta \end{pmatrix}. \tag{6.1}$$

The time evolution of this Hamiltonian results in the exchange of population between the different magnetic sublevels.

In this chapter, we present the apparatus used for creating both the bias and oscillating magnetic fields that leads to the implementation of this Hamiltonian on the system. We follow this with discussions and experimental results of some fairly typical schemes used in the characterisation of the system and later in the progress of the experiment. The three topics we consider are adiabatic transitions, which allow us to find the resonance with relative ease; Rabi oscillations, where applying the coupling field for finite times leads to coherent population oscillations between the levels; and Ramsey interferometry, which allows for a more precise calibration of Δ compared to the Rabi frequency Ω_{RF}.

Ultimately, the desired quantity for the parameters Δ and Ω_{RF} are derived from the practicalities of the cold atoms setup. The time taken by the Stern-Gerlach magnetic gradient pulse to provide sufficient contrast for state readout puts an effective maximum on the total time of flight of the atoms that can be tolerated.

We look for an experimental time window on the order of 1 ms in which relevant oscillations in the system can occur more than ten times. We therefore aimed for quantities in the tens of kilohertz range. To satisfy the RWA, the coupling frequency must be much larger than the generalised Rabi frequency $\tilde{\Omega} = \sqrt{\Delta^2 + \Omega^2}$. To satisfy this condition with relatively simple apparatus, a coupling field on the order of one megahertz was decided upon. This corresponds to a required bias field on the order of one to two gauss.

6.1 Bias Field Creation

To create the desired Hamiltonian as well as define a sensible quantisation axis for optical pumping, it is necessary to create a bias field. This field has the effect of lifting the degeneracy of the different magnetic projection sub-levels so that they can then be controllably coupled using a radiofrequency field. This field is created with a 'Bias Coil' pair, two Helmholtz orientated square coils placed to create the bias field along the imaging beam axis. For each atom in the cloud to experience the same Hamiltonian, it is necessary to create a homogeneous field. For this reason the coil pair was made to be relatively large, with square side lengths of 32 cm. The separation of the coils is 7 cm[1] and each coil has 25 turns. The coil pair is arranged around the chamber with the atoms at the centre.

The requirements on field stability are determined such that resulting fluctuations of the Hamiltonian do not significantly affect the dynamics of the system. This requires control of the field at the kilohertz level for the energy level spacing which corresponds to the milligauss level or 0.1% of the total field. This required stability must be achieved over the spatial extent of the cloud and temporally over the 1 ms long experimental window.

Due to the restraints on the total time of flight ($\sim$20 ms), coupled with the requirement that there must be no magnetic fields on during the sub-Doppler cooling stage, the bias coil must be capable of establishing its desired steady value very quickly (preferably in a couple of milliseconds) to allow sufficient time for the subsequent experimentation window and then Stern-Gerlach pulse.

Unfortunately, as arose with the Stern-Gerlach coils, the finite time constant of the coil circuitry and eddy currents created by the chamber resist this turn-on process. To decrease the turn-on time of the coils, it is possible to add a resistance to the coil circuit to lower the time constant at the cost of power dissipation (a larger voltage is required to drive the same current), however, this does little to help with the eddy currents. The relative size of the coils means that the generated bias field should be poorly coupled to the chamber,

[1] A more homogeneous field could be produced by a larger separation but this would decrease the magnitude of the field we are able to create.

which houses the eddy current loops. This should reduce the magnitude of the eddy currents, however, given the field stability required, they still play a significant role in degrading the temporal stability.

To facilitate both fast field switch-on and eddy current compensation, we actively drive current through the bias coil with a specified time profile that overshoots the desired steady-state value and then exponentially decays back towards it as illustrated in Figure 6.1. The field at the atoms is the addition of the field created by the bias coil plus the opposing field created by the eddy currents. As the overshot field from the coil is decaying towards equilibrium, the eddy currents decay to zero. By tuning the proportion of the overshoot and the time constant of the decay, we attempt to make these two effects cancel, leaving a field with a short rise time followed by temporal stability. Similar approaches that involve overshooting the intended set-point have been used for eddy current compensation in other cold atom experiments [63], particle accelerators [65] and NMR experiments [66].

The overshoot function is programmed onto a Keysight Arbitrary Wave Generator (AWG) with the following functional form

$$V(t) = V_0(1 + \epsilon\, e^{-t/\tau}). \tag{6.2}$$

This signal is fed to a home-made Voltage Controlled Current Source (VCCS) with a OPA541 power Operational-Amplifier (op-amp) at its core. A functional diagram of the main components of the circuitry are shown in Figure 6.2[1]. The feedback is set to regulate the current flowing through a resistor in the same loop as the bias coil. The large gain-bandwidth product of the op-amp allows the device to force current quickly through the coil by applying larger voltages when a fast current change is required. When connected to the VCCS, the coil has an added 4Ω resistance to reduce its natural time constant of 0.72 ms to 0.32 ms when subjected to a square voltage pulse. However, with a square pulse applied to the input of the VCCS, the rise time of the current in the coils is 0.17 ms.

Since we do not have an easily usable magnetic field probe at the position of the atoms, we use the atoms themselves to optimise the fractional value of the overshoot, ϵ and the time constant τ using schemes that are sensitive to the magnetic field. This was primarily done by monitoring any shift in the resonance of the Rabi oscillations when starting the RF coupling at different times within the experimental window, although in future could be done with the more sensitive Ramsey method. Both of these methods are discussed later.

Using this large coil approach with active eddy current compensation has the penalty of restricting the magnitude of the bias field that can be created. The VCCS is capable of a peak current supply of

[1]The design of this circuit is still in the preliminary stages and there are too many bugs to faithfully call it fully functional including the intermittent presence of high frequency noise on some loads. Because of this, no detailed schematic is included.

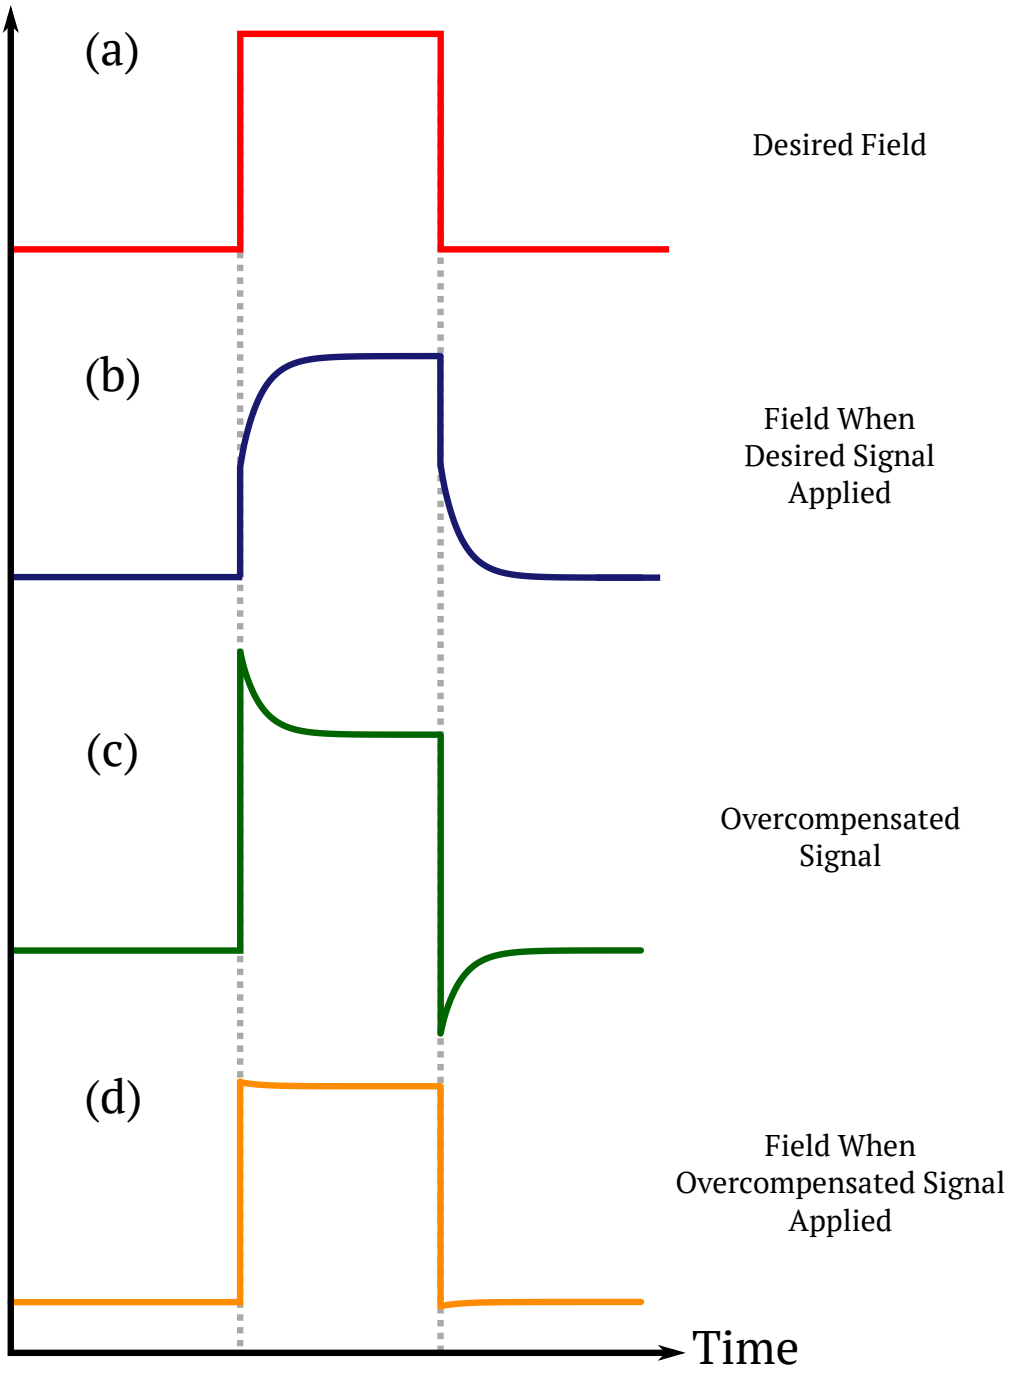

Figure 6.1: Illustration of the exponential overshoot eddy current compensation scheme. In this illustration, the time constant of the coil is ignored and just the effects of the eddy current are shown. The desired field is a square pulse as shown in (a) but if the current in the coil follows this desired profile, the eddy currents lead to an undershoot of the target field, shown in (b). If instead, the current follows an overshoot, followed by an exponential decay as in (c), then the eddy currents are mostly compensated for, leaving a field profile (shown in (d)) much more like the desired shape.

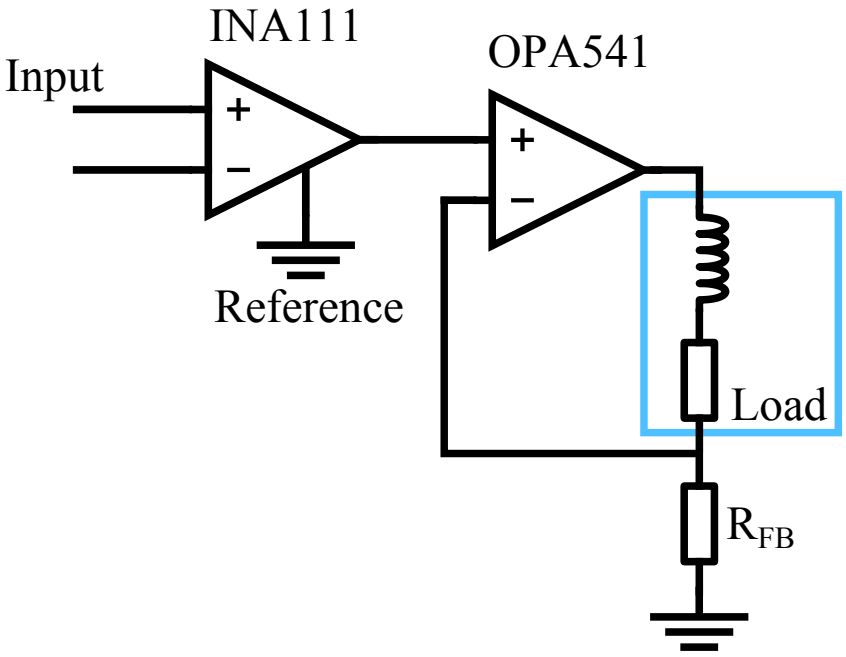

Figure 6.2: Simplified diagram of the circuitry for the VCCS used to drive the bias coil. Intended improvements for future versions include an alternative feedback path to shunt the gain of the OPA541 at high frequencies to reduce high frequency oscillations [67].

approximately 5 A. For the active overshoot, the desired field must be achievable for a fraction of this peak current. As such, with the current hardware, we are limited to the regime of a bias field of a few gauss.

Further improvements could be made to the compensation mechanism by implementing a more complex function on the overshoot, such as a double exponential decay to compensate the initial overshoot.

6.2 Oscillating Magnetic Field Creation

The key properties of the oscillating magnetic field are its frequency, amplitude and homogeneity. The desired frequency of the field is set by the magnitude of the bias field which determines the energy separation of the m_F states. The amplitude determines the strength of the couplings between the m_F levels. Coupling strength is typically parametrised by the Rabi frequency which is measured directly in Section 6.4. As with the bias coil, it is desirable for all of the atoms to experience the same Hamiltonian and this requires a large degree of field homogeneity.

As discussed, to satisfy the RWA, the frequency of the oscillating field must be approximately one megahertz. The desired Ω_{RF} was on the order of tens of kilohertz corresponding to a field amplitude on the order

of tens of milligauss.

The creation of AC magnetic fields is an extensive realm of engineering with many different sources of information spanning many different frequency ranges, geometries and techniques (eg. Nibler [68]). The majority of the literature is directed at the field of antennae. These facilitate the creation of electromagnetic fields that radiate from the source out to infinity. This is not necessary or desirable for our experiment since the desired location of the field is within the chamber. As a result, we need a 'near field coil' as it is commonly referred to in the literature, which is a coil with a negligible radiation component, that can mostly be treated as a DC coil aside from complex impedance issues. When creating such a coil, the field profile may be determined from the Biot-Savart law identically to DC field coil calculations, provided the wavelength of the AC signal is significantly longer than the total length of the coil. Much of the literature on near-field coils concerns their application to Nuclear Magnetic Resonance (NMR) and Magnetic Resonance Imaging (MRI) [69].

For the desired megahertz frequency regime, strong consideration must be taken of the current driving mechanism. The current response of a coil to an voltage oscillating at frequency ω can be roughly determined by treating it as an RL circuit. The response in frequency space is given by

$$|I| = \frac{V/R}{\sqrt{1 + (\omega/\omega_c)^2}} \tag{6.3}$$

where $\omega_c = 1/\tau = R/L$. The field strength and frequency required make driving current through the coil using a simple voltage amplifier a significant challenge due to the current transfer function's low pass form as a large voltage is required to produce significant current. Instead, we add additional components to create resonances in the coil at the desired frequency. This is achieved using two sets of variable capacitors as shown in Figure 6.3.

The two sets of capacitors, added in series and parallel to the coil are referred to as the matching and tuning capacitors respectively. The tuning capacitors are tuned to make the real part of the impedance (the resistance) $50\,\Omega$ at the desired frequency. Once this condition is set, the matching capacitors are tuned to reduce the imaginary part of the impedance (the reactance) to zero at this frequency. These conditions are required for efficient coupling of power from the amplifier (which has a $50\,\Omega$ output impedance) that feeds the coil. Tuning of these parameters was achieved using the Smith plot mode of a HP 3577A network analyser with S-Parameter test set. For our system, three 5.5 - 65 pF variable capacitors operating in parallel were used for the tuning capacitors and a single 5.5 - 65 pF variable capacitor for matching. For the results contained in this chapter and Chapter 7, the resonance was set to 1.13 MHz, which translated to a corresponding magnetic

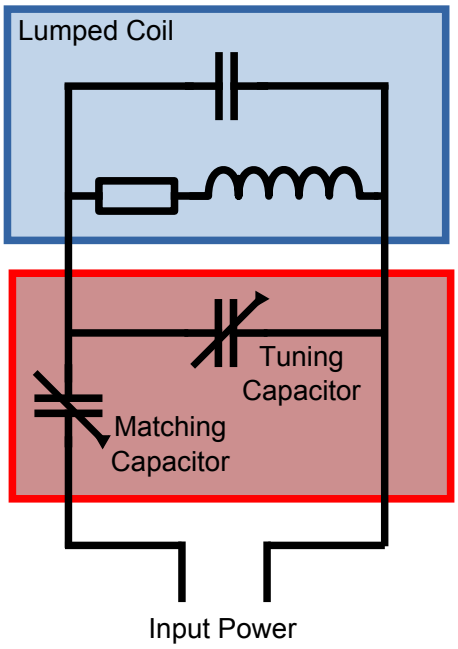

Figure 6.3: Matching network for the radio-frequency coil to make and tune a resonance. The use of a network analyser S-parameter tester allows the tuning capacitor to be used to set the resistance of the system to 50 Ω at the desired resonance frequency before using the tuning capacitor to reduce the reactance to zero.

field of 1.6 G ($\omega = 2\pi$ 0.7 MHz/G).

A further complication of the radio-frequency coil design results from the material constituents of the vacuum chamber. AC fields are prone to being shielded from penetrating the vacuum chamber by (the now famous) eddy currents formed in the chamber material. The steel of the chamber produces a fair amount of shielding but the higher conductance of the copper gaskets produces extremely high shielding. To minimise these effects, the coils were designed to produce the field vertically, perpendicular to the largest gaskets on the chamber that roughly match the shape of the coupling coils but still also perpendicular to the bias field.

The geometry of the coils was chosen to fit around this required orientation; two rectangular coils of dimensions 19 cm × 16 cm with 15 turns orientated in the Helmholtz direction, separated by 7 cm. Theoretical field calculations yield a field amplitude of 1.6 G/A. The relatively small size of the coils is potentially problematic for the homogeneity of the field but with the limited power amplification available, it was necessary to create a coil closer to the atoms.

The coils are driven using a TTi TG5012A signal generator at an output power of 4 dBm. This is amplified using a Minicircuits ZHL 32 A amplifier to give a saturated amplifier power of 29 dBm. This is then passed through a 3 dB attenuator before being connected to the coil. The attenuator protects the amplifier

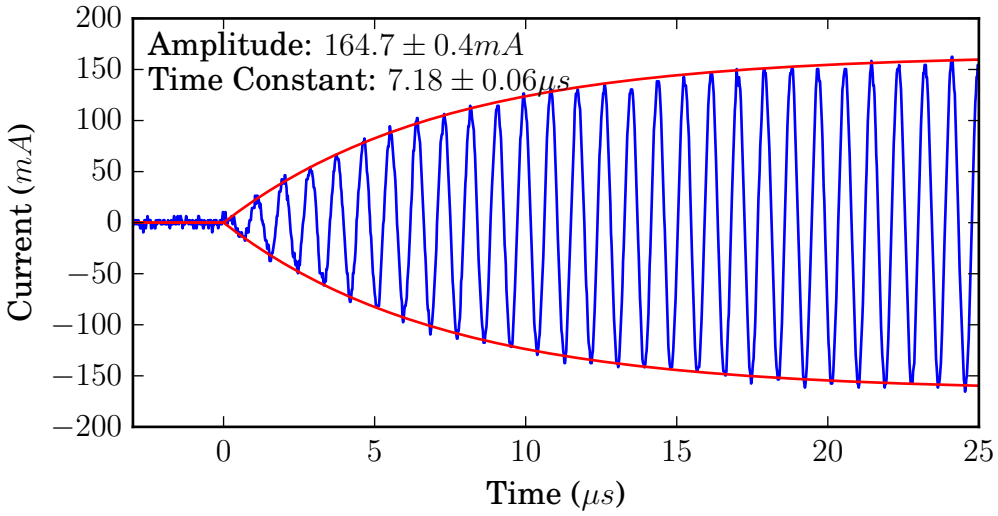

Figure 6.4: Current trace showing the finite switching time of the RF coil.

by attenuating any reflections due to any impedance mismatches with the coil. Since any reflections pass through the attenuator twice, they are suppressed by 6 dB relative to the amplifier output. The final power delivered to the coil is 26 dBm. Delivered to a non-resonant 50 Ω load, this should correspond to a current of 126 mA but this is enhanced by the resonance of the circuit to provide a measured current amplitude of 165 mA. This enhancement is not a significant improvement on the theoretical non-resonant case and is unexpectedly low considering the Q-factor of the resonant circuit. This measured current amplitude would be expected to provide a Rabi frequency of 92 kHz although we do expect the chamber to attenuate the field and reduce this value. The resonant RF coil has the drawback of a finite switching time. When a pulse of RF is applied, the current inside the coil takes several oscillations to reach its peak amplitude. A current trace of the current flowing through the coil is shown in Figure 6.4. The time constant of the oscillation amplitude is evaluated to be 7.18 ± 0.06 µs. As will be seen later, this has distorting effects on the Rabi and Ramsey schemes. This finite switching time also has a ring-down effect when the power is turned off with the same time constant.

It should also be noted that arbitrary increase of the turn number to increase the field resulting from a given current has a drawback in addition to the increased inductance. Adjacent turns in coil windings have a small capacitance between them. The addition of all of these small capacitive elements is often lumped as a parasitic capacitance in parallel to the RL circuit. Large turn numbers lead to a large parasitic capacitance and also the breakdown of this lumped assumption. In this regime, the coil system contains

multiple resonances, even without the matching circuit, of which most cannot be tuned using the added capacitors.

6.3 Finding the Resonance with Adiabatic Transitions

To create resonant coupling between the m_F states, it is essential that the frequency of the coupling field is matched to the energy level separation of the uncoupled states. The absolute frequency of the oscillating magnetic field is well known due to the high accuracy of the frequency generator and can be easily verified using a spectrum analyser or current probe. However, the energy separation of the magnetic states is generated by the bias field inside the chamber, which is harder to characterise accurately. The field is defined by the set-point voltage of the AWG which is then translated into a current by the VCCS, leading to the field governed by the geometry of the coils, which is then attenuated slightly by the chamber eddy currents. As there are many sources of error in each of these conversions, the best calibration of the bias field is achieved by finding the set-point which induces resonant coupling with the known RF frequency. Changing the coupling field detuning requires changing either the bias field or the coupling field frequency. As the coupling field coil is a resonant circuit, changing its frequency also changes the coupling strength which is undesirable. We therefore manipulate the detuning by changing the bias field which is controllable by the experimental control computer through communication with the AWG and determining a voltage set-point for the field. However, without calibration, the mapping of this voltage set-point to the actual field at the atoms is broadly unknown.

In the Rabi and Ramsey schemes, there is very little response when the detuning is large. This makes calibrating the set-point by finding the resonant condition of the atoms using these schemes a significantly laborious process of conducting an experimental run for each individual bias field voltage set-point until there is a response in the state populations. An alternative is to linearly ramp the detuning over a range of voltage set-points with the coupling field on. If the resonant field is within the range of the specified sweep, Landau-Zener theory [70, 71] determines that there will be a change in state population if the system follows an adiabatic transition. This allows us to probe a larger range of the bias field set-points within a single experimental run. The theory behind this mechanism is outlined below.

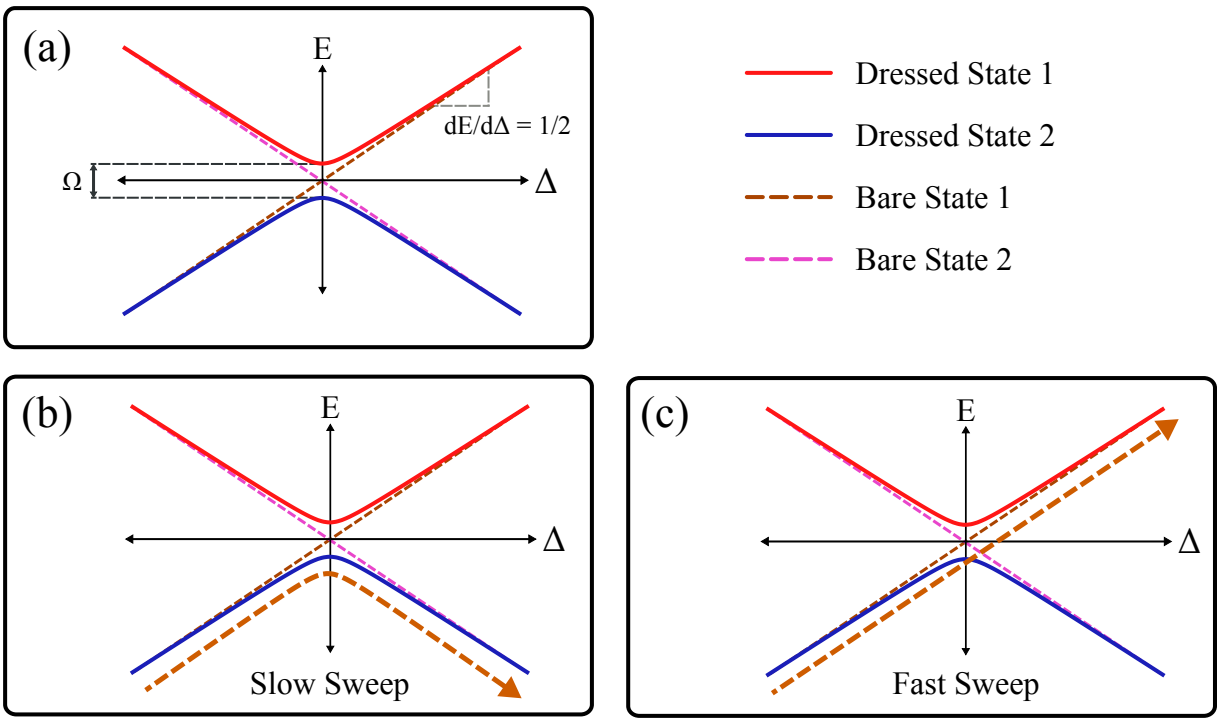

Figure 6.5: Illustration of a two level adiabatic transition scheme. (a) shows the energy levels of the uncoupled states (brown and pink dashed lines) and the eigenstates of the coupling Hamiltonian (red and blue solid lines) as functions of detuning. These states are also known respectively as the *bare* and *dressed* states. The dressed states form an avoided crossing. (b) illustrates the case where the detuning is ramped from $-\infty \rightarrow +\infty$ slowly. The population adiabatically follows the dressed state, illustrated by following the blue line, which represents a change of state in the bare state basis. (c) shows the case where the detuning is swept quickly. In this case, the state population follows the bare state line.

6.3.1 Landau-Zener Theory

The typical picture of an adiabatic transition is a coupled two-level system where, at $t = 0$, the system is in one of the two uncoupled basis states (the uncoupled states are also known as the *bare* states of the system) and the detuning is at $-\infty$. The detuning is linearly swept in time to $+\infty$ at a rate $d\Delta/dt$. The eigenstates of the system (known as the dressed states) as functions of detuning form an avoided crossing as shown in Figure 6.5.

A full treatment of this two-level system is given in [72]. The essence of the result is that when the detuning sweep rate is slow compared to the size of the avoided crossing, the state of the system will adiabatically follow the avoided crossing. Following this dressed state involves interacting with the field and changing state in the non-coupled basis. Conversely, if the sweep rate is fast, then the system will non-adiabatically 'jump' over the avoided crossing and remain in the same uncoupled state.

In the intermediate regime, the final state is a superposition of the two uncoupled states with the probability of measurement in the opposite state given by $P = \exp\left(-\pi\Omega^2/2\frac{d\Delta}{dt}\right)$ where Ω is the two level Rabi frequency and $\frac{d\Delta}{dt}$ is the rate of change of the detuning. This 2-level system is generalised by Carroll and Hioe [73] to the 3-level case. In the 3-level case, there is a probability of transitioning to the $m_F = 0$ state as well as the opposite m_F state. Starting from the $m_F = -1$ state, the expected probability of transitioning to state $m_F = 0$ is given by

$$P_{m_F=0} = 2Q(1-Q) \tag{6.4}$$

and for the fully inverted state $m_F = +1$, is given by

$$P_{m_F=+1} = (1-Q)^2 \tag{6.5}$$

with the probability of remaining in the $m_F = -1$ state given by

$$P_{m_F=-1} = Q^2 \tag{6.6}$$

where, in all three cases, $Q = \exp(-2\pi q)$, where $q = \Omega^2/16\frac{d\Delta}{dt}$. We can see that, as in the two-level case, with a slow ramp $\frac{d\Delta}{dt} \to 0$, the probability of transitioning to the $m_F = +1$ state tends to unity.

6.3.2 Adiabatic Transition Results

Whilst the theoretical treatment describes the detuning ramping from $-\infty$ to $+\infty$, this is obviously impractical in real experiments. As long as the ramp starts and ends significantly far from the resonance (several Rabi frequencies away) and passes through the resonance, the behaviour is the same. The detuning sweep is achieved by programming a linear ramp on the bias field controller between two points for a set period of time. When the resonance point is between the two sweep points, a signal will be manifest as a change in the state populations. The sweep range can then be reduced and moved to narrow the possible position of the resonance. Reducing the sweep amplitude between the same time points should also produce a more complete transfer of states due to the reduction in the detuning sweep rate $\frac{d\Delta}{dt}$. Iterative application of this process leads to a rough position of the resonance and we can then attempt to do Rabi oscillations (see Section 6.4).

Whilst further investigation and analysis are not directly pertinent to the overall experimental goal, investigation of the details of these transitions is possible with our set-up and thus the results of some basic

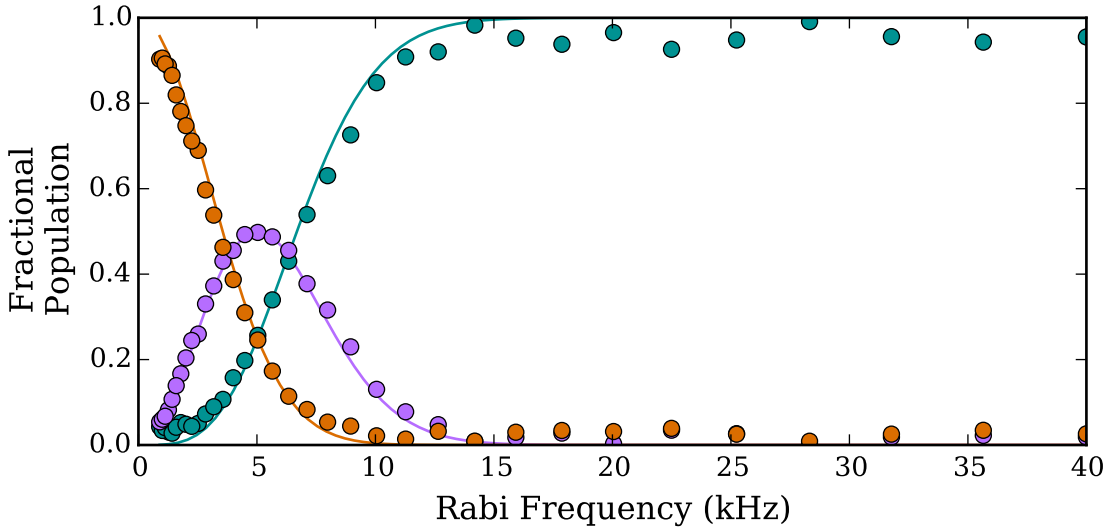

Figure 6.6: Experimental results showing the fractional state populations of the cloud after a detuning sweep for different Rabi frequencies. The points correspond to measured fractional populations of the initial (orange), middle (lilac) and final (turquoise) states. The lines represent a fit to Equations (6.4), (6.5) and (6.6). Note that the Rabi frequencies were calibrated from the RF power using data obtained from the Rabi oscillations discussed in Section 6.4.

investigations follow.

The main parameters of Equations (6.4), (6.5) and (6.6) are the sweep rate $\frac{d\Delta}{dt}$ and the Rabi frequency Ω. The sweep rate is set by the quotient of the field ramp amplitude and the ramp time. The Rabi frequency can be simply varied by changing the RF power output of the signal generator. As the power delivered to the coil is proportional to the square of the coil current, the Rabi frequency is proportional to the square root of the applied power. By varying the RF power for a fixed sweep, we map out the fractional populations of the states after the sweep and fit the functional forms of Equations (6.4), (6.5) and (6.6). An example of these results for a sweep time of 0.5 ms is shown in Figure 6.6. The fitted sweep rate in this case is 359 ± 6 kHz/ms.

The behaviour is as expected. For low Rabi frequencies, the atoms remains in the same uncoupled state and for large Rabi frequencies, the system moves to the opposite uncoupled state.

With the general behaviour observed, we can attempt to vary the sweep rate. This can be achieved by changing the amplitude of the linear ramp or the sweep time. It is to be expected that eddy currents will deform and retard the ramp of the field which may affect the sweep rate in the crucial region of zero detuning.

We repeat measurements of the form shown in Figure 6.6 for different sweep amplitudes centred on the same point. We roughly relate the ramp set points to the field by assuming that the resonant field is in the middle of the sweep. Faster sweeps require a higher Rabi frequency to change states which has the effect of

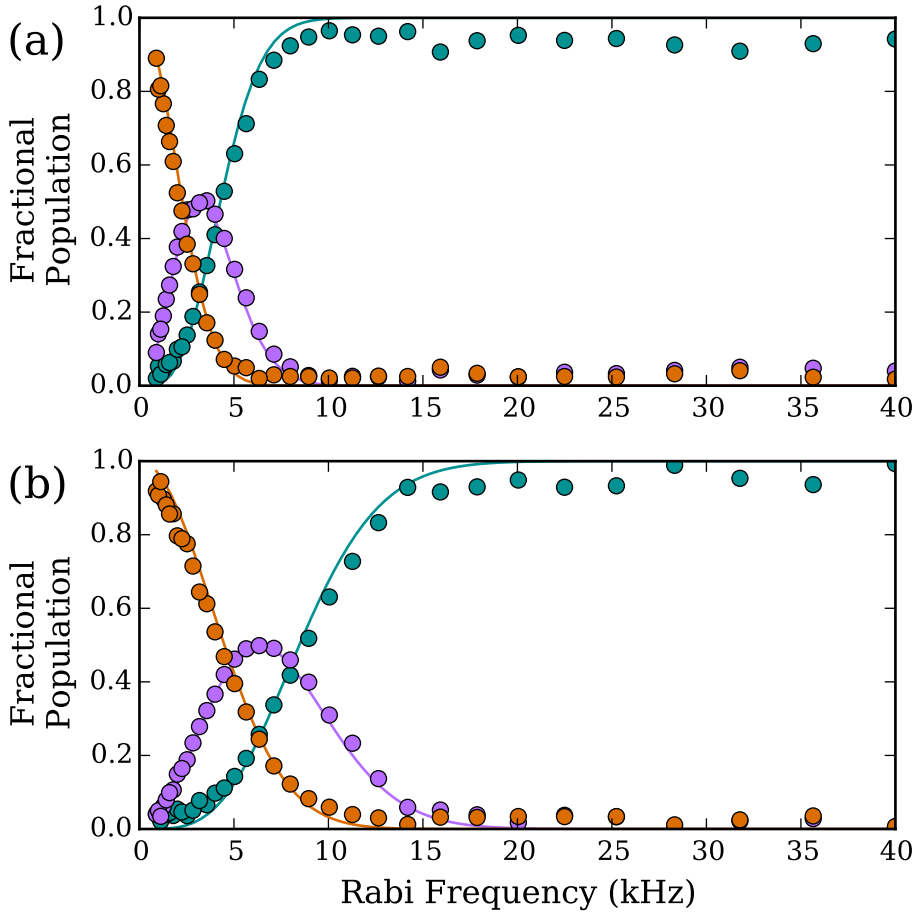

Figure 6.7: Demonstration of the effect of the sweep rate on the adiabatic transitions. A slow ramp rate of 147 kHz/ms shown in (a) allows the transfer of states at lower Rabi frequencies than at the fast ramp rate of 580 kHz/ms shown in (b).

pushing the transition feature in Figure 6.6 towards higher power. Two extreme examples are shown with fitted sweep rates of 147 ± 3 kHz/ms and 580 ± 10 kHz/ms in Figure 6.7.

A comparison of these fitted sweep rates to the expected sweep rates (based on the set-point calibration) is shown in Figure 6.8 with a linear fit. The gradient of 1.11 ± 0.01 indicates a reasonable agreement with our calibrated sweep rate.

With the adiabatic transitions allowing us to easily narrow the bias field set-point for the transition, we look to resolve the state couplings with a constant Hamiltonian using Rabi oscillations.

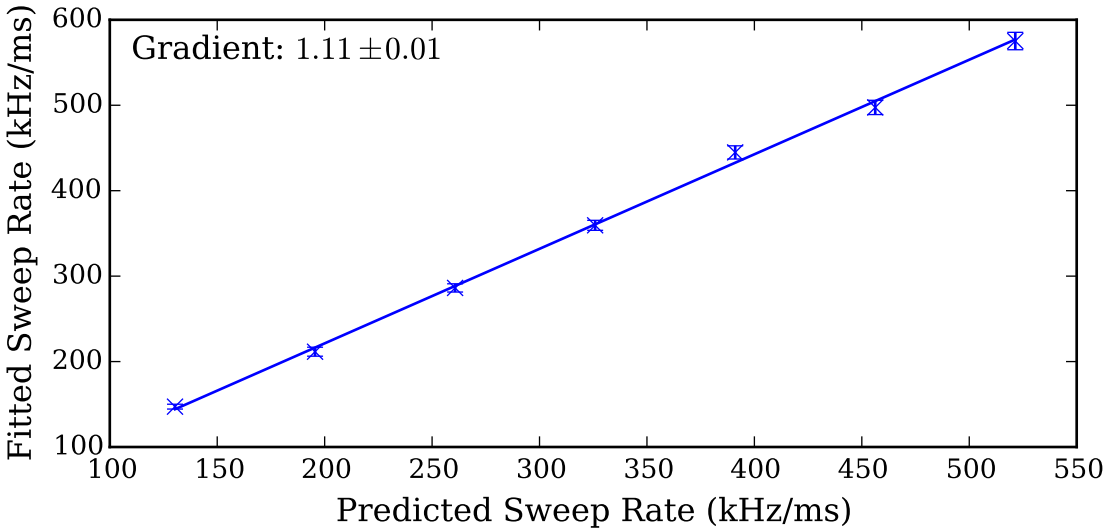

Figure 6.8: Comparison of the fitted sweep rates with those expected from the ramp set-points. The discrepancy is ascribed to eddy currents lowering the gradient of the ramp near to the resonance point.

6.4 Rabi Oscillations

If a 2-level system is prepared in state $|0\rangle$ and a coupling is quickly turned on, the system will oscillate between the two states. The population of state $|1\rangle$ is given by the Rabi formula [35]

$$P_{|1\rangle} = \frac{\Omega^2}{\tilde{\Omega}^2} \sin^2\left(\frac{\tilde{\Omega}t}{2}\right) \tag{6.7}$$

where $\tilde{\Omega} = \sqrt{\Omega^2 + \Delta^2}$ is the *generalised Rabi frequency*, which is the energy level difference between the eigenstates of the coupled Hamiltonian. When the coupling is on-resonance ($\Delta = 0$), Equation (6.7) reduces to $\sin^2\left(\frac{\Omega t}{2}\right)$ and full transfer of the population occurs at the Rabi frequency Ω. Figure 6.9 shows the time dynamics for several values of Δ.

In the Bloch sphere representation, the Hamiltonian can be decomposed into $H = \Delta\sigma_z + \Omega\sigma_x$, which leads to precession of the state around the vector $(\Omega,\ 0,\ \Delta)$. For $\Delta = 0$, the precession axis is along the equator and the state vector passes through the southern-pole of the Bloch sphere, indicating full population transfer. With a finite Δ, the axis of rotation has a component in the polar direction thus the population transfer is reduced as the path of the state vector's precession will no longer take it to the southern pole. Schematics of these cases are shown in Figure 6.10.

The generalisation of these key Rabi oscillation behaviours to three-level systems is, in general, non-trivial.

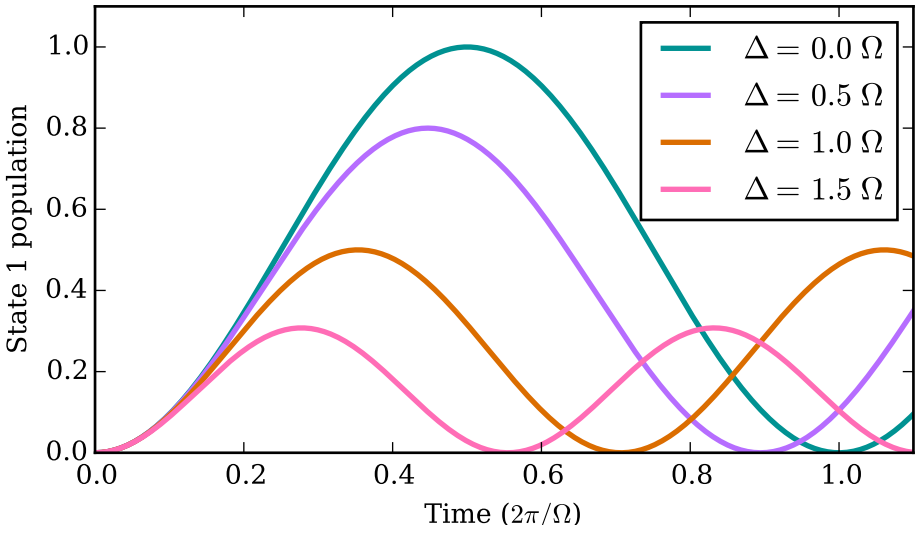

Figure 6.9: Illustration of two-level Rabi formula behaviour. The population of the initially unoccupied state oscillates in time with a frequency given by the generalised Rabi frequency $\tilde{\Omega}$. Increased detuning leads to a quickening of the oscillations and a reduction of the oscillation amplitude.

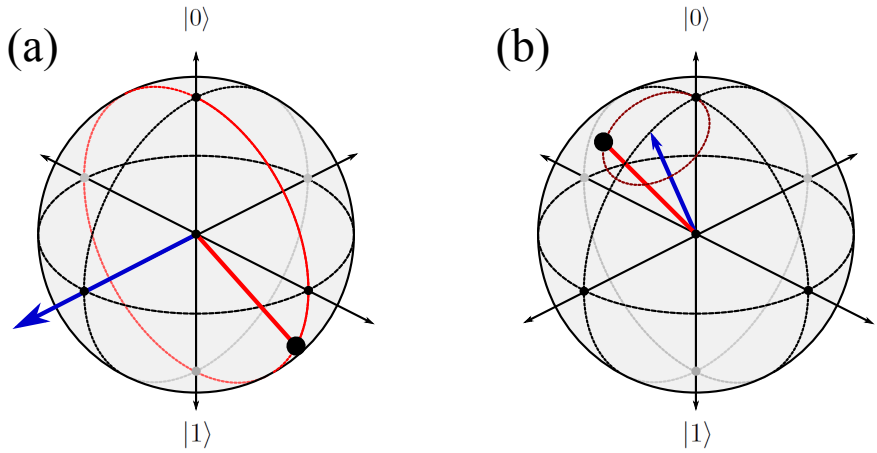

Figure 6.10: Bloch sphere representation of Rabi oscillations. (a) With $\Delta = 0$, the precession axis (blue arrow) is along the equator of the sphere and full population transfer occurs as the precession passes through the southern pole. (b) With finite detuning, the precession axis is in the upper hemisphere and the population transfer is reduced.

However due to the symmetry of our Hamiltonian with regards to the state transitions, the generalisation becomes greatly simplified. Explicitly, the symmetry is the identical couplings and detunings between $m_F = -1 \leftrightarrow m_F = 0$ and $m_F = 0 \leftrightarrow m_F = +1$.

The solution to the application of the coupled Hamiltonian in Equation (6.1) for a time t was obtained analytically through matrix diagonalisation. Starting from the $m_F = -1$ state, the probabilities for each of the three states are given by:

$$P_{m_F=+1} = \frac{\Omega^4}{\tilde{\Omega}^4} \sin^4\left(\frac{\tilde{\Omega}t}{2}\right) \tag{6.8}$$

$$P_{m_F=0} = \frac{1}{2}\left[\frac{4\Delta^2\Omega^2}{\tilde{\Omega}^4} \sin^4\left(\frac{\tilde{\Omega}t}{2}\right) + \frac{\Omega^2}{\tilde{\Omega}^2} \sin^2\left(\tilde{\Omega}t\right)\right] \tag{6.9}$$

$$P_{m_F=-1} = 1 - P_{m_F=0} - P_{m_F=+1} = 1 - \frac{\Omega^2}{\tilde{\Omega}^2}\left(2 - \frac{\Omega^2}{\tilde{\Omega}^2}\right) \sin^4\left(\frac{\tilde{\Omega}t}{2}\right) - \frac{\Omega^2}{2\tilde{\Omega}^2} \sin^2\left(\tilde{\Omega}t\right). \tag{6.10}$$

The behaviour of these equations is shown in Figure 6.11. There are broad similarities to the two-level case. The oscillation frequency increases and amplitude decreases with a larger (magnitude) detuning. The third level adds complications however. The middle state dynamics contain frequencies at the generalised Rabi frequency $\tilde{\Omega}$ and $2\tilde{\Omega}$. Like the two-level case, these dynamics are equivalent to a vector rotating on a Bloch sphere, but with the points of the Bloch sphere corresponding to the state superpositions as discussed in Section 2.3.

The above expressions are valid for 'hard' rectangular coupling pulses, i.e. the Rabi frequency is turned on and off quickly. An example of the comparison of a hard pulse and one with a exponential rise time is shown in Figure 6.12. As discussed, the resonant nature of our RF coil leads to a finite turn-on time of 7 µs. For Rabi frequencies on the order of 10 kHz, this becomes a significant time-scale and the dynamics of the system are not expected to follow the standard formulae above.

To interpret the effect of the finite switching time of the Rabi frequency, consider the Bloch sphere representation; with $\Omega = 0$ and $\Delta \neq 0$, the precession axis lies along the z-axis of the sphere. In a hard pulse, when Ω is turned on to a finite value, the state precession axis moves instantaneously to the vector $(x, y, z) = (\Omega, 0, \Delta)$ and the states precess. With a finite turn-on time, the precession axis rotates slowly towards this vector. If the time associated with the movement of the precession axis is greater than or comparable to the precession period, then the states will spiral around the axis, changing the dynamics.

Analytic forms of the population dynamics resulting from such 'soft' pulses are significantly complicated

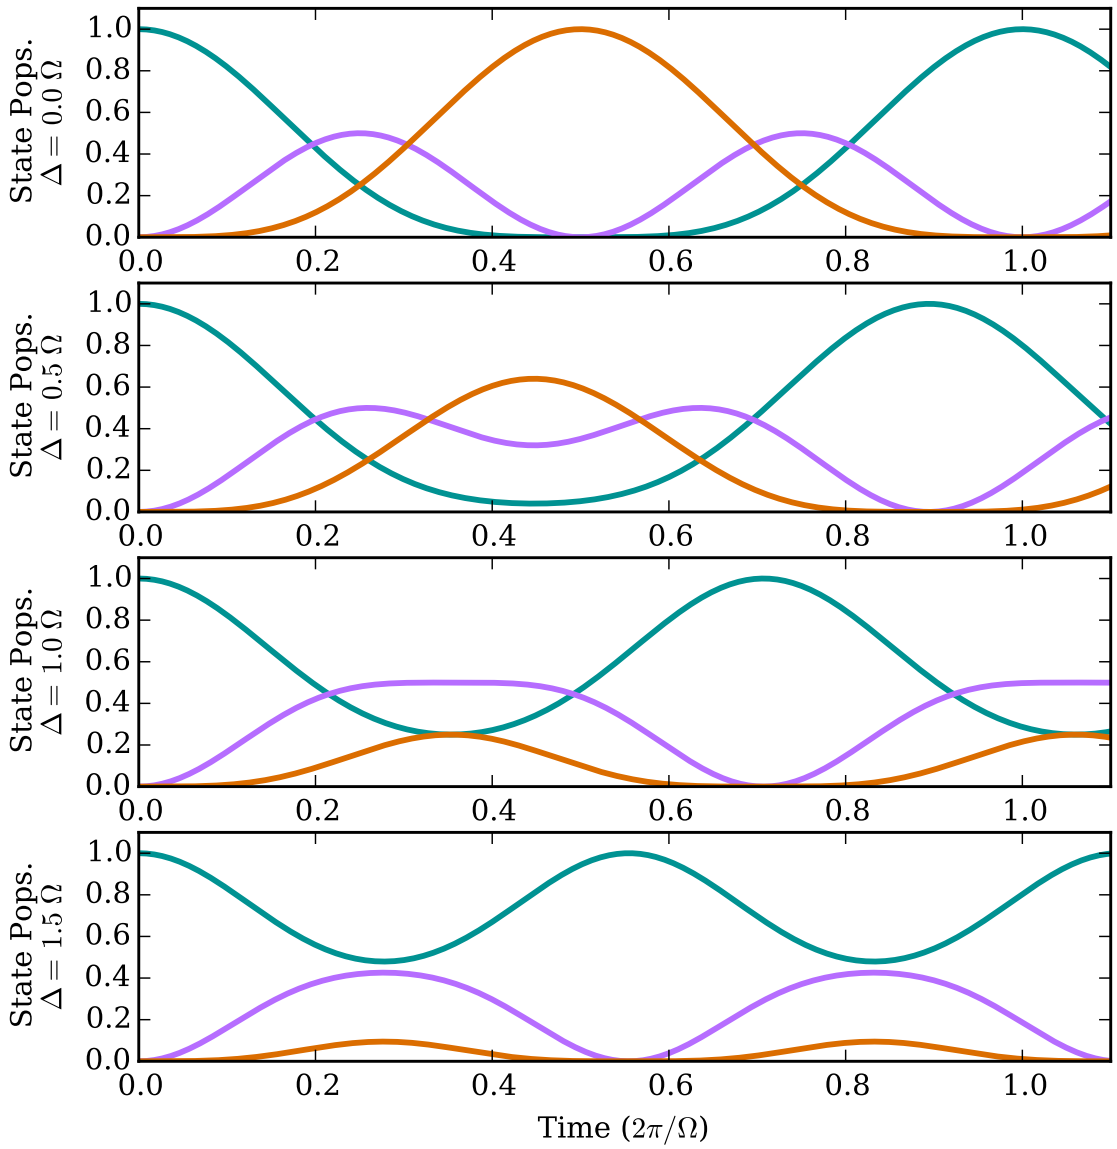

Figure 6.11: Calculation of 3-level Rabi formula behaviour for $\Delta = 0.0\Omega, 0.5\Omega, 1.0\Omega$ and 1.5Ω sequentially from top to bottom. Each plot shots the populations of all three states, $P_{m_F=+1}$ (orange), $P_{m_F=0}$ (lilac) and $P_{m_F=-1}$ (turquoise).

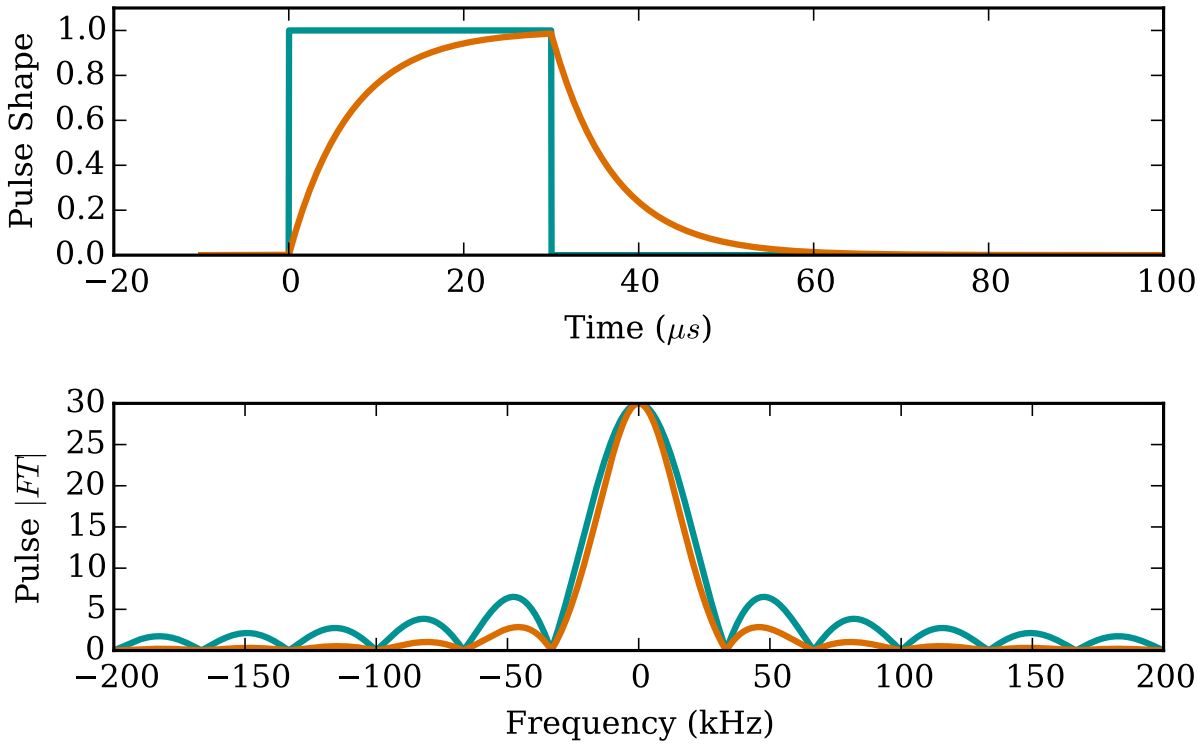

Figure 6.12: Graph showing the frequency components of the pulse shape. The frequency components in the side lobes decay faster with detuning for the pulse with a 7 µs time constant (shown in orange) than for the 'hard' rectangular pulse (shown in turquoise).

and are thus not presented here but results from numerical dynamics of the time dependent Hamiltonian are shown in Figure 6.13 compared to the rectangular pulse. It can be qualitatively seen that the major effect of the soft pulse is a significant reduction in the amplitude of the population oscillations compared to the hard pulse. This effect can be intuitively understood by considering a soft pulse that moves the precession axis adiabatically slowly. In this case, the state vector will stay closely aligned to the precession axis for the duration of the pulse. As the Rabi frequency returns to zero at the end of the pulse, the population will remain unchanged[1].

For the intermediate state, the visualisation becomes less clear but the principle is the same, slow movement of the precession axis leads to a localisation of the state vector to its vicinity. The effect is less pronounced with a small detuning ($\Delta \ll \Omega$) as the angle of the precession axis changes quickly to its resting value (mostly in the x-axis of the Bloch sphere).

An alternative approach to the understanding comes from the solving of the two-level coupled system time dynamics using perturbation theory. This gives the result that, for pulses that only weakly perturb the atomic state, the transition probability is given by the modulus square of the pulse's Fourier transform with $\omega = \Delta$ [74]. The moduli of the Fourier transform of both the hard and soft pulses are shown in Figure 6.12. The magnitude of the Fourier transform at a particular frequency is therefore related to the response of the system with a detuning corresponding to that frequency. This response of the soft pulse is clearly smaller than that of the equivalent hard pulse for all finite detunings. In the NMR spectroscopy community, pulse shaping is used to purposely *reduce* the system response to frequencies other than the resonant one as this aids spectroscopic resolution [75].

It should be noted both from the numerical simulations and the frequency response argument, that at zero detuning, the effect of the pulse shape is reduced. In fact only the time integral of the pulse is relevant [74]. This allows us to measure the resonant Rabi frequency Ω despite the soft pulse shape. With the effects of the finite pulse time on the Rabi oscillation dynamics qualitatively understood, we move to the acquired data.

6.4.1 Rabi Oscillation Results

To make a measurement of the Rabi oscillations, the RF field is pulsed for varying times t during the experimental window and the relative state populations are measured for each. We repeat the process for

[1]This picture can also be used to understand the adiabatic transitions discussed in section 6.3. For a slow sweep, the precession axis moves slowly from one of the poles of the Bloch sphere to the other, with the state vector adiabatically following. The result in this case is the state vector exchanging its polar position.

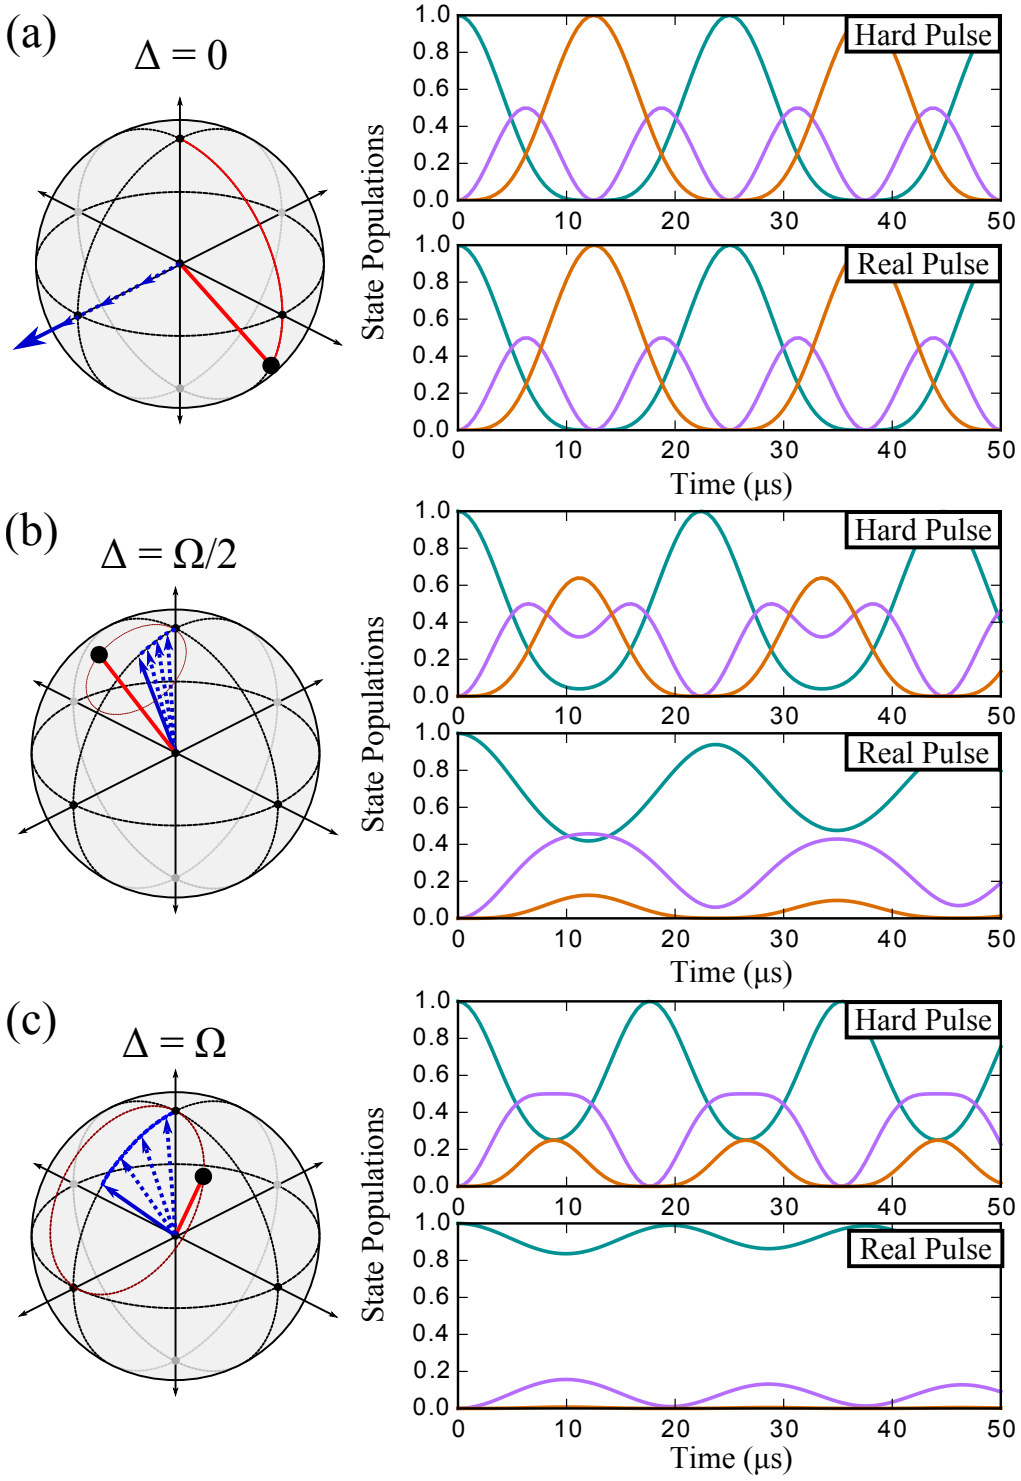

Figure 6.13: Numerical simulations showing the difference in the population time dynamics between a 'hard' square pulse and our exponential increase and decrease for several values of the detuning. It can be seen that the amplitude of the oscillations for the real pulse is reduced by the finite detuning much more than for the hard pulse. The simulations were conducted with $\Omega = 40\,\text{kHz}$ and $\tau = 7\,\mu\text{s}$.

different detunings by varying the bias coil set-point. Typical results of these time resolved Rabi oscillations are shown in Figures 6.14a, b and c for comparatively large, medium and small detunings respectively.

Despite the issues with the finite RF field switching time, the Rabi formulae for hard pulses defined in Equations (6.8), (6.9) and (6.10) provide a good functional fit to the data, albeit with fitting parameters that do not correspond to the physical situation. Results for this fitting are shown in Figure 6.15.

For an applied rectangular pulse, we would expect the fitted Rabi frequency to be constant and the detuning to vary linearly (specifically as $y \propto |x|$ as the Rabi scheme is insensitive to the sign of the detuning) with the changing bias field. From Figure 6.13, we can see that the soft pulse does not significantly affect the frequency of the oscillations compared to the rectangular pulse. In the Rabi formulae, the oscillation frequency is parametrised by $\tilde{\Omega} = \sqrt{\Delta^2 + \Omega^2}$, and the oscillation amplitude by powers of the ratio $\Omega/\tilde{\Omega}$. To fit the reduced oscillation amplitude of the soft pulse whilst maintaining the oscillation frequency, the solution to the fitted equations reduces Ω and increases Δ.

Even accounting for the effects of the soft pulse, the data also has some unexpected discrepancies, the first of which is that the fitted detuning does not reach zero despite the set-point clearly encompassing the resonance and the second is the apparent asymmetry of the response around the resonance.

For the fit to produce a zero detuning measurement requires the amplitude of the oscillations to fully exchange the populations. Any inhomogeneity in the bias or RF field leads to a spread in Δ or Ω across the cloud. As the population read-out is a sum of the population exchanges experienced by each part of the cloud, a complete exchange of populations is impossible with any inhomogeneities. The same effect is also responsible for the eventual time decay of the oscillation amplitudes.

As the Rabi scheme is only sensitive to the magnitude of the detuning, we expect the response of the system to be symmetric around a central point corresponding to $\Delta = 0$. Our reasoning for the apparent asymmetry is the presence of a residual DC magnetic field with a component along the RF field axis. This situation is illustrated in Figure 6.16. If an additional DC field is present, the detuning is defined by the difference between the RF oscillation frequency and the energy level splitting caused by the *net* field. Since the magnitude of such a net field is not linear to the applied bias field, we may expect a non-linear response in the results. In addition to this effect, the Rabi frequency is defined only by the component of the RF field that lies perpendicular to the net DC field. A change in the bias field may cause a small angular shift between the fields, also causing a slight change in the Rabi frequency.

Any residual field in the experiment would have to be small enough to allow the polarisation gradient cooling to operate at optimised efficiency. Measurements using Ramsey interferometry (see section 6.5) lead

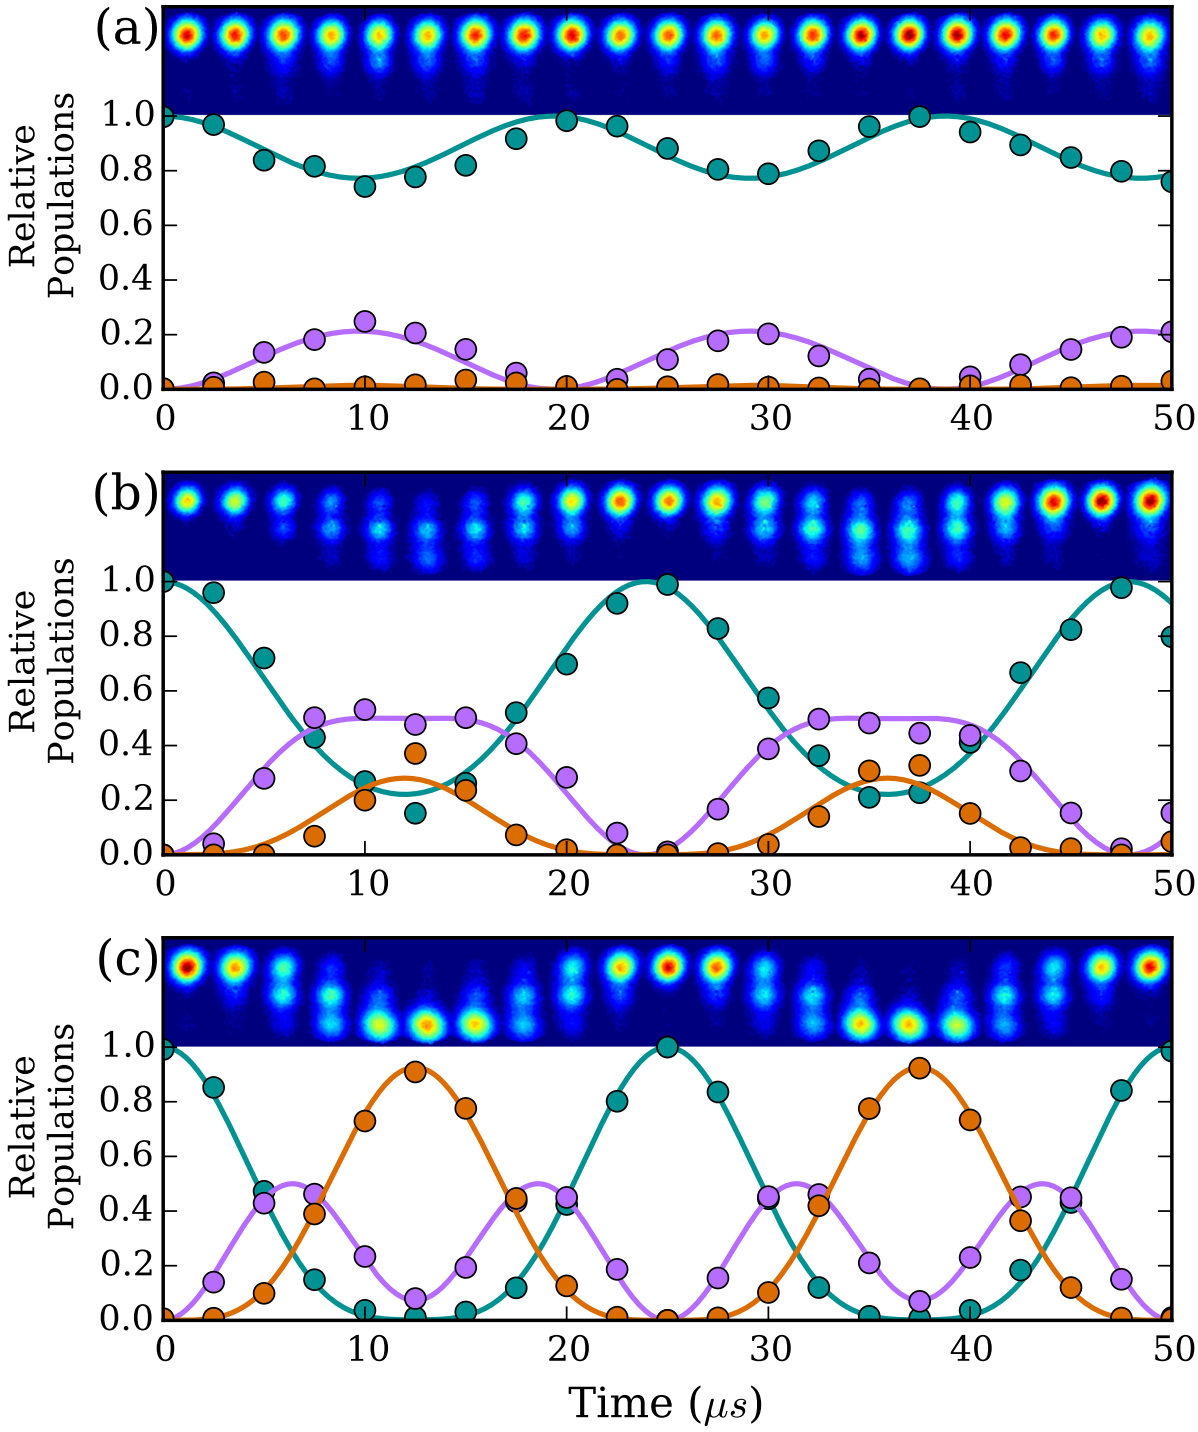

Figure 6.14: Experimental results showing examples of Rabi oscillations with various detunings. The graph shows the relative populations in the initial (turquoise), middle (lilac) and end (orange) states. Each point corresponds to an image on the stitched together array shown above the graph.

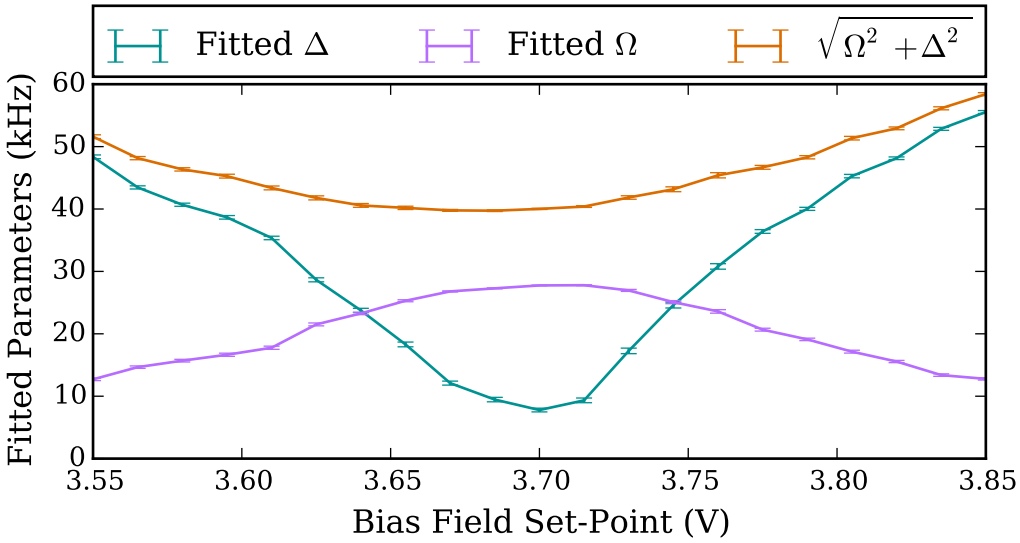

Figure 6.15: Results of fit to Rabi formula for different bias fields. The detuning Δ and Rabi frequency Ω were left to be free parameters. It can be seen that the fitted Ω unexpectedly changes despite being unchanged experimentally. This is a manifestation of the effect of the finite time constant of the RF field switching. Note that the voltage set-point is not directly comparable to the other voltage set-points presented in this chapter as the results were taken at a different stage in the experimental assembly.

to more insight about a potential residual field.

Further quantitative analysis in the time domain was not conducted although it does provide an interesting area that could be looked further into. The concluding remark is that the pulse shape resulting from the resonant RF coil leads to a reduced oscillation amplitude, making the population oscillations difficult to observe for finite detunings. We are roughly content that the system displays Rabi coupling behaviour and the pulse shape issue is unlikely to cause further issue for the noise applied Hamiltonians discussed in Chapter 7.

Resonant Rabi Frequency

As the differences in response between the soft and hard pulses near resonance are small, we can still use the Rabi scheme to characterise the resonant Rabi frequency Ω.

To check the scaling relation of the resonant Rabi frequency against the RF power, the Rabi oscillations were measured close to resonance for varying powers on the RF signal generator. The oscillations are shown in Figure 6.17 using a single figure metric given by $P_{+1} - P_{-1}$ where $P_{\pm 1}$ is the fractional population of the $|\pm 1\rangle$ states.

The resonant Rabi frequency Ω is expected to vary as $\Omega \propto B_{RF} \propto I_{RF} \propto \sqrt{P_{RF}}$. Figure 6.18 shows the

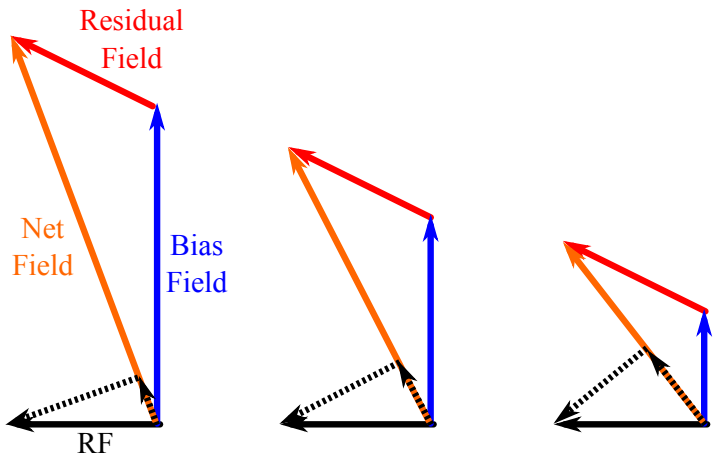

Figure 6.16: Illustration of the effect of a residual field on the Rabi frequency. The presence of a residual field means that varying the magnitude of the bias field changes the field axis which in turn changes the projection of the RF field in the plane perpendicular to the net field, changing the Rabi frequency. The effect here is greatly amplified as we do not expect a residual field so large.

measured population oscillation frequency against the square root of the RF power expressed in milliwatts.

From these results, we can see the expected linear relationship between the resonant Rabi frequency and the square root of the RF power. The small 'turn' at the top of the data is possibly the start of the saturation regime of the RF amplifier. The Rabi frequency of 40 kHz, measured with an RF power of 4 dBm is approximately half of that expected from the calculations conducted in Section 6.2. The discrepancy is attributed primarily to eddy current shielding of the field.

6.5 Ramsey Fringes

For accurate manipulation of the system Hamiltonian, it is important to be able to set, and verify the exact magnetic field at the position of the atom cloud. The control level for this is a voltage set-point on the AWG that ultimately drives the Bias coil. With the Landau-Zener scheme and Rabi oscillations already discussed, we know the approximate set-point required to produce resonant coupling. As the frequency of the coupling field is well known, this allows us to know the bias field at that particular set-point. Making the assumptions that the set-point/bias field transfer function is linear and passes through the origin, this single point is enough for a calibration of the bias field to voltage set-point. The former assumption seems fairly reasonable; the linear electronics and the linear nature of the Biot-Savart law ensure this. The assumption of a zero background field seems reasonable to a degree as the good operation of the polarisation gradient

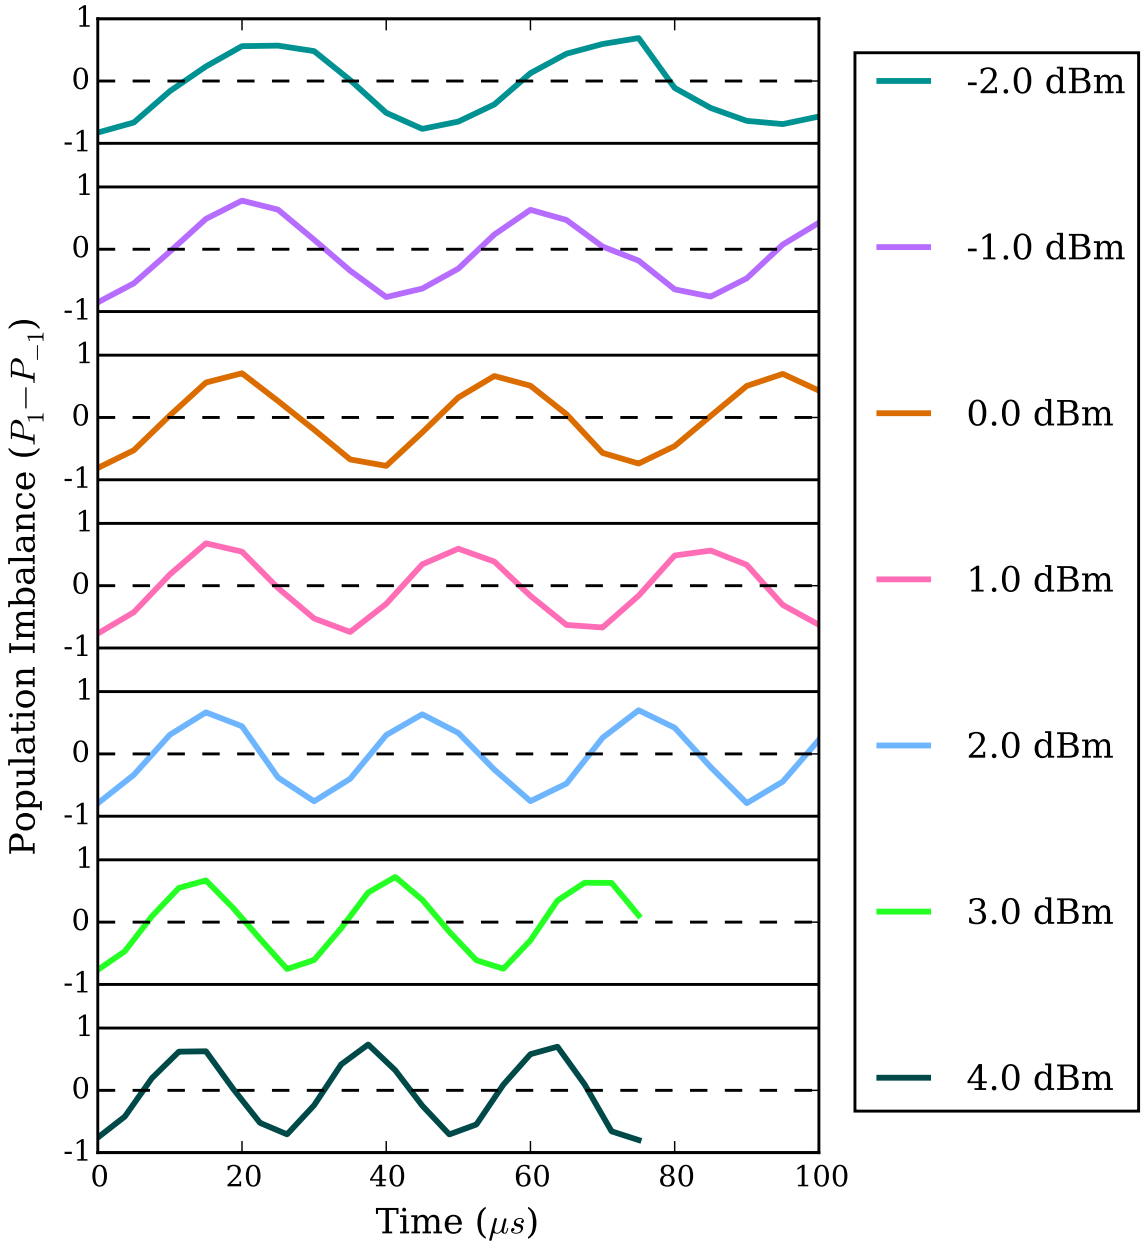

Figure 6.17: Experimental results of the effect of radio-frequency power on the Rabi oscillation frequency. The state populations are reduced to a single figure of merit, $P_{+1} - P_{-1}$ for ease of visualisation.

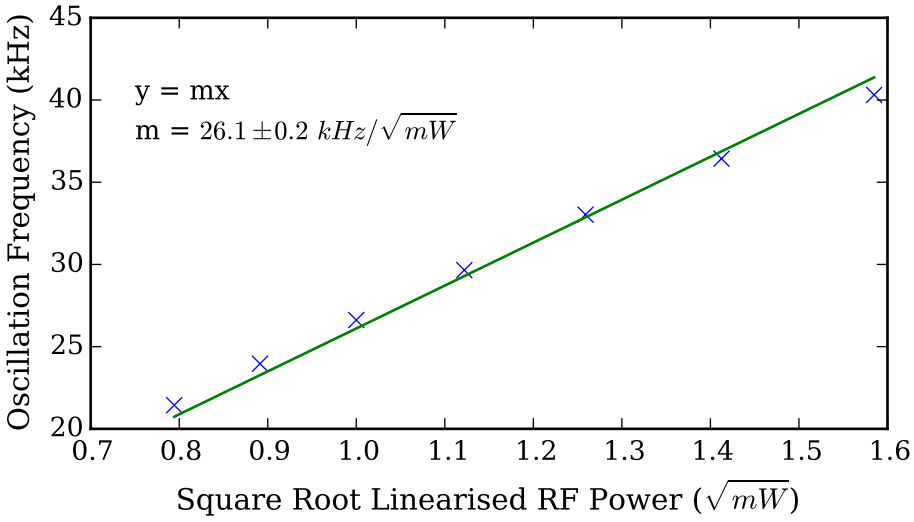

Figure 6.18: Linear fit of the Rabi frequency to the square root of the power supplied by the frequency generator.

cooling, which occurs with the set-point at zero, is reliant on minimal residual magnetic fields. However the extent of this zero field is not easily quantifiable.

In principle, the Rabi scheme can be used to determine the calibration using the population oscillation amplitudes and generalised Rabi frequency as a function of the detuning from the resonance. However, as our pulse shape means the dynamics do not follow the well known formulae, we turn to an alternate method.

Ramsey interferometry [76] is a scheme used as an accurate and precise measure of transition energies by creating a pulse interferometer. It is used in atomic clocks to probe the hyperfine ground splitting of various atomic species. For a review of atomic clocks, please see [77].

Below, we outline the theory behind Ramsey Interferometry, followed by a presentation of results looking to calibrate the bias field.

6.5.1 Ramsey Sequence

We start by considering a two-level coupled system in the RWA with Hamiltonian as in Equation (2.16). The basic Ramsey scheme involves two $\pi/2$ coupling pulses (a pulse that has a length given by a quarter of a resonant Rabi period) with a time interval T between them. The sample is prepared in the initial state of $|0\rangle$. The first $\pi/2$ pulse creates an equal superposition of the $|0\rangle$ and $|1\rangle$ states. During the interval time T, a phase difference accumulates between the eigenstates according to their different energy levels $\phi = 2\pi\ \Delta\ T$

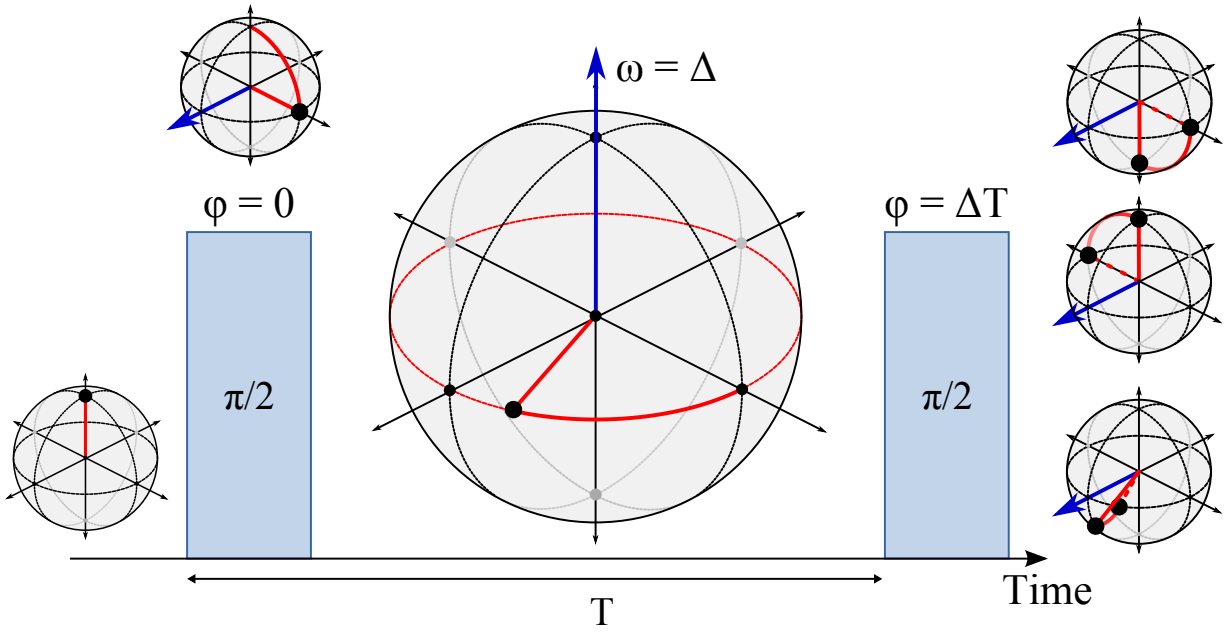

Figure 6.19: Schematic illustration of simple Ramsey scheme. The initial $\pi/2$ pulse moves the Bloch vector to the equator. It then precesses around the z-axis with an angular frequency given by the detuning for the interval time T. The second $\pi/2$ pulse rotates around the same axis with the resultant final position depending on the location of the Bloch vector before the second pulse.

where Δ is the difference in energy (with $\hbar = 1$) between the 2 states. The phase difference between the two levels in the superposition state causes the second $\pi/2$ pulse to either move the sample back to the original state, completely, to the other state or some other superposition. Since the phase accumulation process is periodic with period $1/\Delta$, the resultant populations oscillate with the corresponding frequency; $\propto \sin^2(\Delta\ T)$.

Using the Bloch sphere representation (shown in Figure 6.19), the initial $\pi/2$ pulse rotates the Bloch vector a quarter rotation around an equatorial axis, such that it moves from the north pole to the equator. During the time interval T, the Bloch vector rotates around the vertical axis of the sphere with angular frequency $2\pi\Delta$, the second $\pi/2$ pulse conducts the same rotation about the same axis as the first. The result is either a vector pointing opposite to its original position (if the vector had rotated an integer number of times during the interval), back to its original position (if the vector had rotated an extra half revolution) or somewhere on the surface of the sphere along the plane formed by the vertical and the pulse rotation axis.

Population oscillations can therefore be obtained by varying either Δ or T. Oscillations when varying Δ are referred to as Ramsey Fringes.

When the coupling strength (the Rabi frequency) does not significantly exceed the detuning, the applied pulses cannot be considered true $\pi/2$ pulses as they do not rotate the vector $\pi/2$ radians in a single plane on

the Bloch sphere. Instead, the rotation will be a larger angle about an axis out of the equatorial plane. This means that the second pulse will be unable to provide full population transfer to the other state irrespective of the interval time T, reducing the contrast of the fringes. This modulates the amplitude of the fringes by an envelope function, so they can only be seen strongly in the vicinity of resonance.

It should be noted for completeness that the initial assumption that the pulses are $\pi/2$ pulses does not necessarily have to be true. If we treat them more generally as pulses of length τ, then we obtain further complications to the model. As this Ramsey scheme is not the focus of this work however, we will not elucidate the consequences of these additions and instead refer the reader to work more dedicated to the scheme [78].

As with the Rabi oscillations and adiabatic transitions, our three-level system keeps much of the intuition and principles of the two-level case but with slightly modified mathematics in the populations.

The upshot of this is that the short-oscillations in state population for the Ramsey fringes have a period that is inversely proportional to the interval time. As the interval time is a well controlled experimental parameter, we can use these oscillations to calibrate the bias coil driver to the energy shift of the levels.

Results

Results were collected for variations in both the detuning via the changing of the bias field set-point and the time interval between the pulses. An example of the fractional population exchanges as a function of detuning is shown in Figure 6.20 for the interval times of 16.25 µs and 46.25 µs. The $\pi/2$ pulse time was determined from the resonant Rabi frequency measurements and set to 6.25 ms.

As well as the decreased fringe spacing from the large interval time, it is apparent that there is some reduction in amplitude between the central fringe (which corresponds to a detuning of zero) and the nearest fringes. The envelope function is generally an even power of $\sin(\Delta)/\Delta$ for ideal square pulses but as discussed in Section 6.2, the RF field has a finite switching time constant on this timescale, significantly distorting the pulse. The envelope function would thus be related to the Fourier transform of the pulse envelope as well as due to other uncharacterised sources of dephasing.

To quantify the fringe spacing with this envelope shaping, we fit a phenomenological function to the population depletion of the initial state $(1 - P_i)$ with P_i the fractional population in the initial state. The equation is of the form

$$1 - P_i = A\cos^2\left(\pi\frac{(B - B_0)}{B_{env}}\right)\cos^2\left(\pi\frac{(B - B_0)}{B_{fr}}\right) + c \tag{6.11}$$

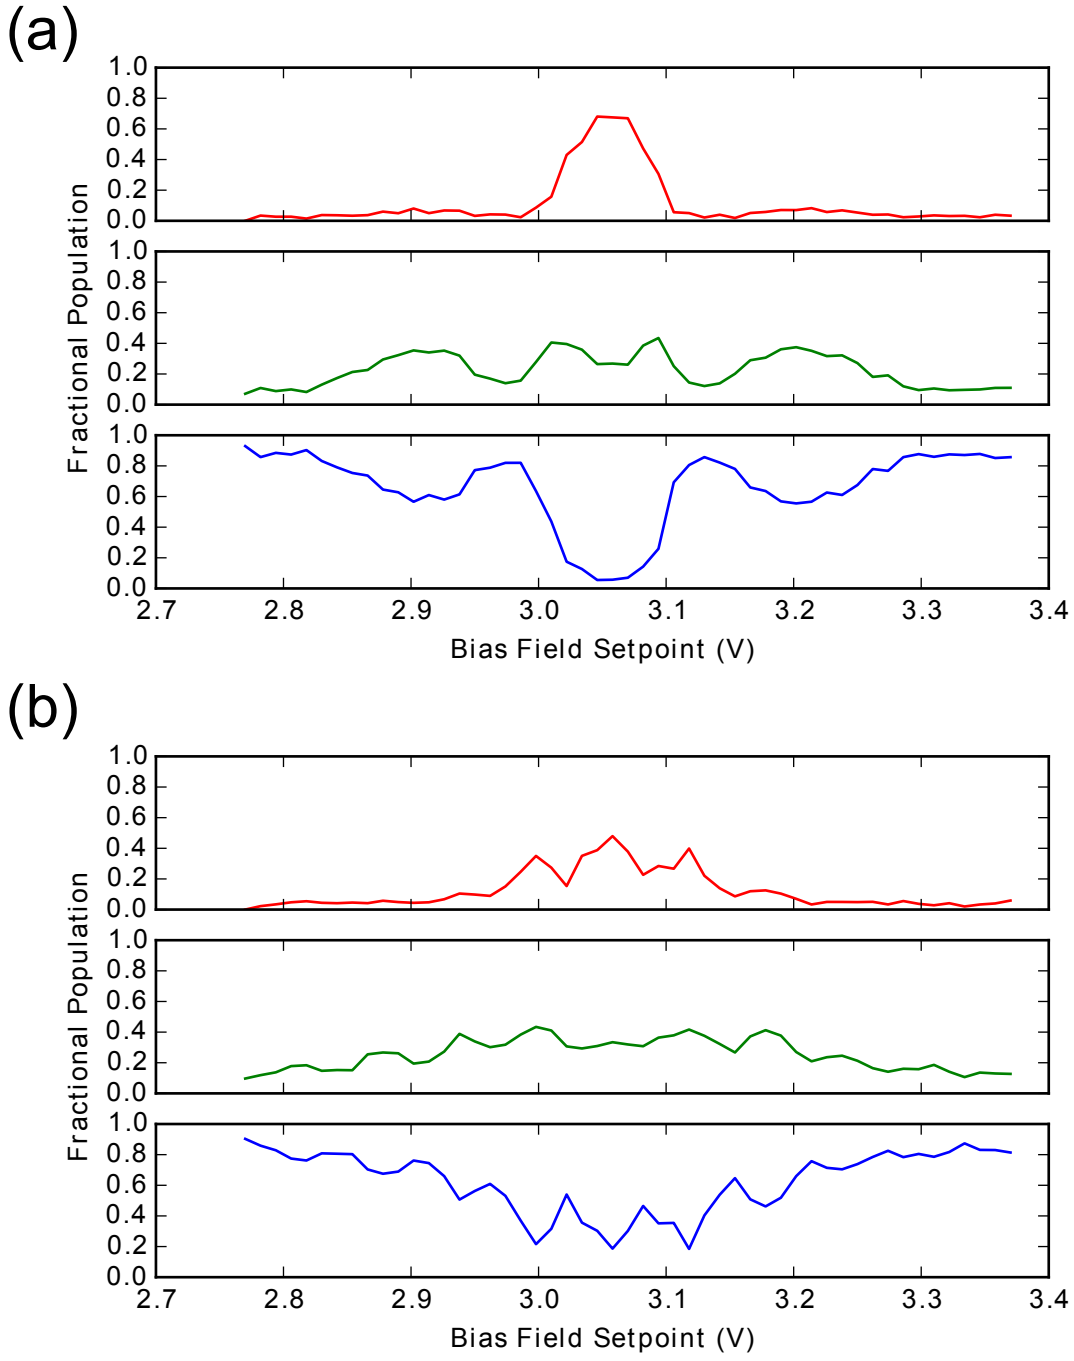

Figure 6.20: Experimental results showing the fractional populations after a Ramsey sequence with interval times of (a) 16.25 µs and (b) 46.25 µs. The population traces for the the far (top), middle (middle) and initial (bottom) states. Note the reduction in fringe spacing for the longer interval time.

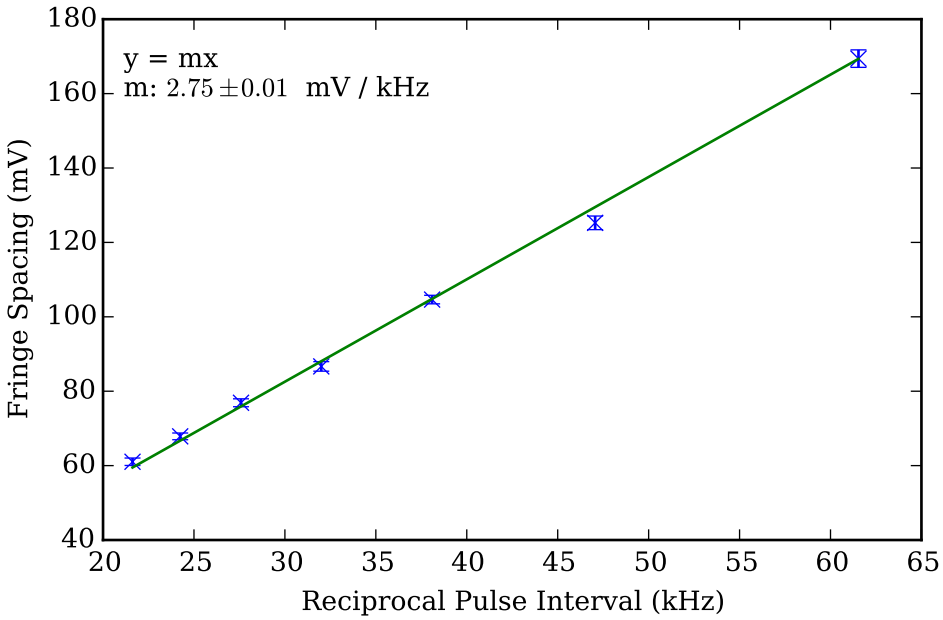

Figure 6.21: Linear fit of the Ramsey fringe spacing against the reciprocal of the interval time. If the fringe spacing is measured as a function of detuning, the gradient is expected to be one.

where we have introduced an amplitude (A), resonance offset (B_0), envelope width (B_{env}), fringe spacing (B_{fr}) and constant (c).

The parameter of most interest is the fringe spacing. As the fringe spacing is measured in the arbitrary units of a voltage set-point, we can compare this measured value to the expected value expressed as a function of detuning, which is $1/T$. Since T is well determined, we can attempt to extract a calibration of the voltage set-point in terms of detuning. As the fringe spacing is a difference between voltage set-points, any offset is cancelled and thus should be directly linearly comparable to $1/T$ with no offset. The result of this linear fit is shown in Figure 6.21.

The gradient of $2.75 \pm 0.01\,\mathrm{mV/kHz}$ gives the field as a function of detuning, but the inverse $0.364 \pm 0.001\,\mathrm{kHz/mV}$ allows us to calculate the change in detuning, and thus bias field per unit set-point. The central fringe in the Ramsey fringes corresponds to zero detuning or the resonant condition and should be the same value of the set-point independent of the pulse spacing. This is measured to be $3056.9 \pm 0.4\,\mathrm{mV}$, which, using the gradient calibration, corresponds to a Zeeman splitting of $1111 \pm 3\,\mathrm{kHz}$. Comparing this value to the RF frequency applied (1130 kHz), implies that a small background field is present that provides the 19 kHz difference. This would correspond to a background field of 27 mG if aligned along the bias field.

6.6 Concluding Remarks

In this chapter, we have discussed the method of implementation and the measurement of coupling to the adjacent m_F states of our system. We conclude that our system broadly follows the outlined theoretical frameworks for adiabatic transitions, Rabi oscillations and Ramsey fringes with the exception of the complications caused by the finite switching time of the RF field. The Rabi oscillations have a resonant flopping frequency in the tens of kilohertz range, which was the desired target. The system is then ideally placed for the application of a time dependent term to the diagonal elements of the Hamiltonian.

CHAPTER 7

HAMILTONIAN DIAGONAL TIME DEPENDENCE

With the system exhibiting expected coherent coupling behaviour with the RWA Hamiltonian, adding a time dependence to the diagonal elements of the Hamiltonian presents a very open field of investigation as this dependence can have many degrees of freedom. As discussed in Section 1.2.1, we attempt to interpret the time dependence of the on-diagonal elements of the Hamiltonian to noise or environmental coupling. The main focus of this chapter is on the Hamiltonian presented as:

$$H = \hbar \begin{pmatrix} -\Delta - \delta(t) & \frac{\Omega_{RF}}{\sqrt{2}} & 0 \\ \frac{\Omega_{RF}}{\sqrt{2}} & 0 & \frac{\Omega_{RF}}{\sqrt{2}} \\ 0 & \frac{\Omega_{RF}}{\sqrt{2}} & \Delta + \delta(t) \end{pmatrix} \tag{7.1}$$

where we are interested in the responses to different forms of the function $\delta(t)$. Initially, investigations were carried out using single frequency (pure tone) oscillations. This was expected to produce multiple frequency dependent responses in the system in a form analogous to parametric resonances [79–81]. This effect was observed and the parameters compared to theoretical models.

Since the system responds to multiple pure tone stimuli, it is expected that the application of a continuous frequency spectrum spanning these multiple excitation frequencies should also yield a response with potential for complex interplay between the different resonances. Initial investigations into this large parameter space are presented and compared to numerical simulations.

7.1 Experimental Implementation

The application of time dependence to the diagonal elements of the Hamiltonian is explicitly achieved using amplitude modulation of the magnetic bias field that lies in the z-axis. Since, in the RWA RF coupled Hamiltonian (see Equation (6.1)), the diagonal elements are $-\Delta, 0$ and Δ, where $\Delta = \mu B - \omega_{RF}$, introducing a time dependence to the Bias field $B \to B(t)$ leads to a time dependence of Δ. We treat this time dependence separately to the DC detuning as $\Delta \to \Delta + \delta(t)$. We refer to this $\delta(t)$ term as the 'noise field' even when the functional form is not considered a noisy waveform.

Adding time dependence to the bias field is implemented using a coil pair, dubbed the 'noise coils', driven, in series with a 4Ω resistor, by a VCCS of the same design as for the bias field as outlined in Section 6.1. The current driver is fed by the second channel of the same arbitrary wave generator used for the bias coil. The noise coil pair is formed of two rectangular coils in Helmholtz orientation situated close to the chamber with dimensions 18.5 cm × 13 cm, separated by 7 cm orientated to create the field along the same direction as the bias field (lab frame z-axis). Each of the coils has only 2 turns. A new coil was implemented rather than add the time dependence to the bias coil due to the different requirements of the fields. The bias field needs to quickly establish a static field on the order of one gauss whereas the noise field should be an AC field capable of a relatively large frequency response (∼100 kHz) at a magnitude on the order of the Rabi frequency (40 kHz ≈ 60 mG).

From the computer control, the specified noise field is defined at 1000 sample points within the 1 ms experimental window. This provides a sampling rate of 1 MHz, which at the noise coil frequency cut-off of 100 kHz provides 10 points per oscillation, with linear interpolation implemented by the AWG between these points. The number of points can be increased at the expense of more protocol overhead.

Calibration of the field produced by the noise coil is challenging for the same reasons discussed with respect to the RF and bias coils. Calibration at DC using the atoms is possible using the Rabi scheme by moving the bias coil voltage set-point away from resonance and finding a corresponding set-point for the noise coil that compensates this shift. This lead to a very rough calibration of ≈80 kHz/V(114 mG/V). Direct field measurement using a field probe next to the chamber window yielded ≈87.5 kHz/V(125 mG/V). Calibration at AC is more problematic as the eddy current magnitude may depend on the frequency of the oscillation due to the time constant of an eddy current loop given by its resistance and inductance. Using the probe at the chamber window to measure the sum of the noise-field and the eddy currents for various noise field frequencies, there was a noticeable drop-off at approximately 60 kHz, reducing the amplitude of the oscillations to roughly half at 100 kHz. It is difficult to know the exact effect on the atoms as the effect

of the eddy currents is likely weaker at the centre of the chamber than at the edge. We must therefore keep in mind that there is a weak low pass filter on the noise field that is felt by the atoms. This could be overcome with a frequency dependent calibration by applying a corresponding high pass filter to the applied noise signal such that the product of this filter and the low pass filter formed by the eddy currents is a flat spectral response, however this was not done for the work presented in this thesis.

More rigorous calibration could be conducted using a varying AC field during a Ramsey sequence (see Section 6.5) as the phase accumulated during the interval time is related to the integral of the detuning. For the remainder of this chapter, even though it is not accurate for all frequencies, we shall use the DC Rabi calibration of 80 kHz/V to provide a more physical discussion of the relevant quantities of the field rather than arbitrary voltage set-points.

The limit on the field amplitude is ultimately set by the current limit of the driving op-amp; 5 A, but we limit this further to 3.5 A (which corresponds to a voltage set-point of 5 V) to reduce the impact of a fault condition. This allows us to provide fields of ∼0.57 G which is approximately one third of the bias field, allowing us to reach a region of very strong modulation.

7.2 Pure Tone Perturbation

The application of a single frequency amplitude modulation to the bias field leads to a Hamiltonian in the radiofrequency RWA of

$$H = h \begin{pmatrix} -\Delta - \delta_0 \cos(\omega_N t) & \Omega/\sqrt{2} & 0 \\ \Omega/\sqrt{2} & 0 & \Omega/\sqrt{2} \\ 0 & \Omega/\sqrt{2} & \Delta + \delta_0 \cos(\omega_N t) \end{pmatrix} \tag{7.2}$$

where the modulation frequency is given by ω_N and the modulation amplitude by δ_0. The time periodicity of the Hamiltonian means the general approach to solving the dynamics is the use of Floquet states [82, 83]. These are similar in nature to Bloch states which manifest as solutions to space-periodic Hamiltonians and are ubiquitous in solid-state systems. Rather than detailing Floquet theory, we can initially gain a more intuitive picture of the behaviour that follows using a subsequent application of the rotating wave approximation.

Let us first reconsider the case of no modulation on the detuning ($\delta_0 = 0$). The time dynamics of this system are given by the Rabi formulae discussed in Section 6.4. The eigenstates of the system can be explicitly determined by finding the eigenvectors of the Hamiltonian matrix, although the Bloch sphere allows a more

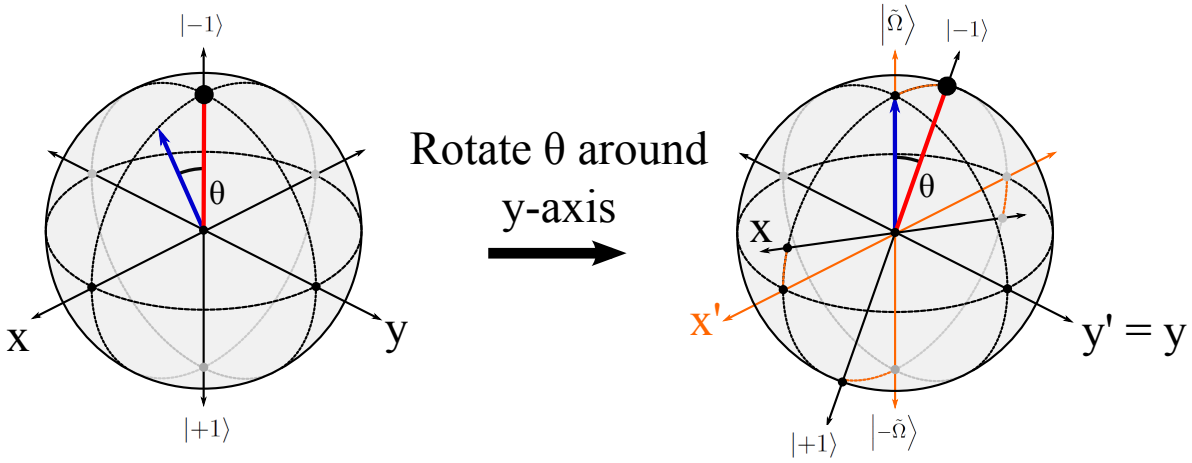

Figure 7.1: Illustration of the change to the eigenbasis of the RF coupled Hamiltonian by rotation of the Bloch sphere. The sphere is rotated such that the precession axis is orientated towards a pole.

intuitive picture; as an eigenstate is (aside from a global phase) unchanged by time evolution, the Bloch vectors parallel and anti-parallel to the precession axis set by the Hamiltonian will satisfy this condition and can be deduced simply to be eigenstates. As the matrix is 3×3, there is a third eigenstate but this does not lie within the sub-group of states described by the Bloch sphere representation. This is also the case for the undressed states $|m_F = -1\rangle, |0\rangle$ and $|+1\rangle$. The $|\pm 1\rangle$ states are the polar states on the Bloch sphere but the state associated with sole occupation of the $|0\rangle$ state does not have a corresponding position on the sphere.

However derived, the eigenbasis states $\left|+\tilde{\Omega}\right\rangle, \left|-\tilde{\Omega}\right\rangle$ and $\left|\tilde{0}\right\rangle$ (where the label is given by their corresponding energy) of the coupled basis is related to the uncoupled basis by

$$\begin{pmatrix} \left|-\tilde{\Omega}\right\rangle \\ \left|\tilde{0}\right\rangle \\ \left|+\tilde{\Omega}\right\rangle \end{pmatrix} = \begin{pmatrix} \frac{\cos(\theta)+1}{2} & \frac{\sin(\theta)}{\sqrt{2}} & \frac{\cos(\theta)-1}{2} \\ -\frac{\sin(\theta)}{\sqrt{2}} & \cos(\theta) & +\frac{\sin(\theta)}{\sqrt{2}} \\ \frac{\cos(\theta)-1}{2} & \frac{\sin(\theta)}{\sqrt{2}} & \frac{\cos(\theta)+1}{2} \end{pmatrix} \begin{pmatrix} |-1\rangle \\ |0\rangle \\ |+1\rangle \end{pmatrix} \tag{7.3}$$

where we have used the convenient angle θ which is defined as $\arctan(\Omega/\Delta)$. On the Bloch sphere, this change of basis is a rotation of the axes that define the polar states on the sphere. For a Hamiltonian of the form $H = \Delta\sigma_z + \Omega\sigma_x$, the basis change is a rotation around the y-axis. This is shown in Figure 7.1. Now the polar states are the $\left|-\tilde{\Omega}\right\rangle$ and $\left|+\tilde{\Omega}\right\rangle$ states and the $\left|\tilde{0}\right\rangle$ state does not have a corresponding position on the Bloch sphere representation.

The Hamiltonian in this dressed eigenbasis is given by the diagonal matrix

$$\tilde{H} = \hbar \begin{pmatrix} -\tilde{\Omega} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & +\tilde{\Omega} \end{pmatrix} \tag{7.4}$$

with the state ordering as in Equation (7.3). This is generally referred to as the *dressed basis*. In the limit of $\Delta \gg \Omega$, the eigenstates of the coupled Hamiltonian are very similar to the uncoupled states of the system that we readout (i.e. a small rotation on the Bloch sphere). This is why, in the large detuning regime, the oscillation of populations is small as the precession axis is very close to the state vector. If we reintroduce the detuning modulation $\delta_0 \cos(\omega_N t)$ in this basis, we introduce time-dependent terms along both the diagonal and on the off-diagonal elements:

$$\tilde{H} = \hbar \begin{pmatrix} -\tilde{\Omega} - \cos(\theta)\delta_0 \cos(\omega_N t) & \frac{\sin(\theta)}{\sqrt{2}}\delta_0 \cos(\omega_N t) & 0 \\ \frac{\sin(\theta)}{\sqrt{2}}\delta_0 \cos(\omega_N t) & 0 & \frac{\sin(\theta)}{\sqrt{2}}\delta_0 \cos(\omega_N t) \\ 0 & \frac{\sin(\theta)}{\sqrt{2}}\delta_0 \cos(\omega_N t) & +\tilde{\Omega} + \cos(\theta)\delta_0 \cos(\omega_N t) \end{pmatrix}. \tag{7.5}$$

The off-diagonal elements make it clear that this is no longer the eigenbasis of the system. This Hamiltonian looks similar to the non-rotating frame Hamiltonian with the RF coupling in Equation (2.23) with the addition of the on-diagonal time dependence. We can remove the time dependence of the off-diagonal elements of this Hamiltonian by a subsequent application of a rotating wave approximation as before, this time rotating with the frequencies $-\omega_N$, 0, and ω_N for the states respectively. Neglecting the on-diagonal time dependence for now, this results in a new, time-independent Hamiltonian

$$H = \hbar \begin{pmatrix} -\Delta' & \Omega' & 0 \\ \Omega' & 0 & \Omega' \\ 0 & \Omega' & \Delta' \end{pmatrix} \tag{7.6}$$

where $\Delta' = \tilde{\Omega} - \omega_N$ is the difference between the pure tone 'noise' frequency and the generalised Rabi frequency, and $\Omega' = \sin(\theta)\delta_0/2$ is a new coupling term between the dressed states. This Hamiltonian obviously mimics the Rabi Hamiltonian of Equation (6.1) but coupling the dressed basis states instead. A Bloch sphere visualisation of this process is shown in Figure 7.2. We justify the neglect of the on-diagonal time dependence terms from Equation (7.5) by stating that for valid RWA, the fast oscillations with frequency ω_N average out

quickly to zero over the dynamic timescales of Δ' and Ω'.

The direct analogy to the RF coupled RWA Hamiltonian means that if the frequency of the pure tone noise that is applied to the undressed energy levels is close to the generalised Rabi frequency ($\omega_N \approx \tilde{\Omega}$), the dressed basis states will start to significantly exchange population. Again considering the case where $\Delta \gg \Omega$, the dressed states are very similar to the bare states which means that we should see significant population transfer in the measurable bare states as well.

This can be seen as an example of the situation discussed in Section 1.1.4 where the noise provides or absorbs the energy to transfer to states that they would otherwise be unable to move to.

An alternative viewpoint to consider these oscillations from is to consider the behaviour from the perspective of frequency modulation. Since, in the undressed basis, the unperturbed detuning Δ determines the rate at which a relative phase is accumulated between the undressed states, it occurs in exponential form $\exp(i\Delta t)$. By incorporating the time varying portion $\delta(t)$, we have terms varying as

$$\exp\left(i\left(\Delta + \delta_0 \cos\left(\omega_N t\right)\right) t\right) \tag{7.7}$$

which is an expression of frequency modulation. Basic frequency modulation theory [84] tell us that at weak modulation ($\delta_0 \ll \Delta$) the frequency spectrum of such a function has a large amount of power at the *carrier* frequency Δ, with some power in sidebands a frequency ω_N either side of the carrier. Using this picture, one can see that the time varying detuning directly allows for manipulation of the relative phase differences between the states, which instantaneously determines the direction of population transfer. The modulation of this phase difference allows beat patterns to form, creating both high and low frequency oscillations.

Experimentally, we show this effect by setting the detuning such that $\Delta \approx 2-3\Omega$ and applying a pure tone oscillation to the noise field with frequency equal to the observed weak flopping frequency at this detuning in the absence of a noise field. In the population oscillations, we expect to see a small amplitude, high frequency oscillation due to the difference between the measured bare state basis and the coupled dressed state basis, superimposed on a larger but slower oscillation caused by the coupling of the dressed basis eigenstates by the noise field. The addition of the small oscillations coupled with noticeable inhomogeneities in the cloud makes detailed fitting of an analytic form challenging so we resort to more qualitative measures. By taking the quantity $r_z = (-1 \cdot P_{m_F=+1}) + (0 \cdot P_{m_F=0}) + (1 \cdot P_{m_F=+1}) = P_{m_F=+1} - P_{m_F=-1}$ where the P's represent the fractional populations of the states, we provide a robust number that corresponds to projection of the Bloch vector onto the z-axis of the Bloch sphere. Results of the time resolved r_z for the pure tone noise field are shown Figure 7.3.

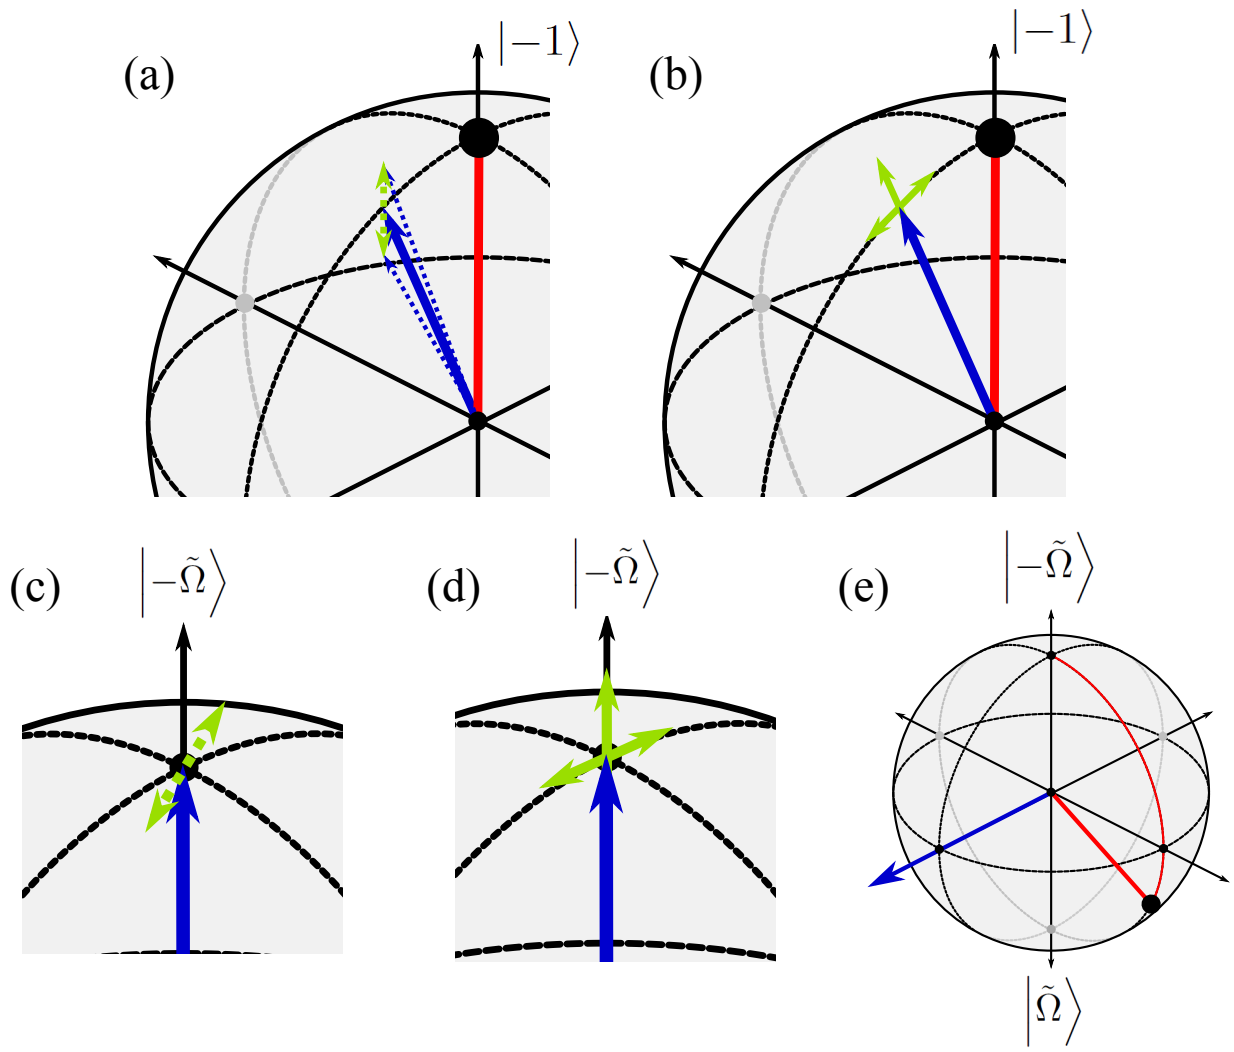

Figure 7.2: Illustrations of the effect of the noise when moving to the dressed basis states. (a) The noise field (green) adds time dependence to the precession axis (blue) in the vertical direction. (b) The noise field can be decomposed into terms parallel and perpendicular to the precession axis. (c) After rotating to the dressed basis, the noise field is similarly rotated to act along both the z-axis and the x-y plane. (d) Decomposition of the noise field into these fluctuations in the z-axis (on-diagonal terms) and the x-y plane (off-diagonal terms). (e) The off-diagonal terms allow a further rotating wave approximation to be made which can induce precession in the dressed basis.

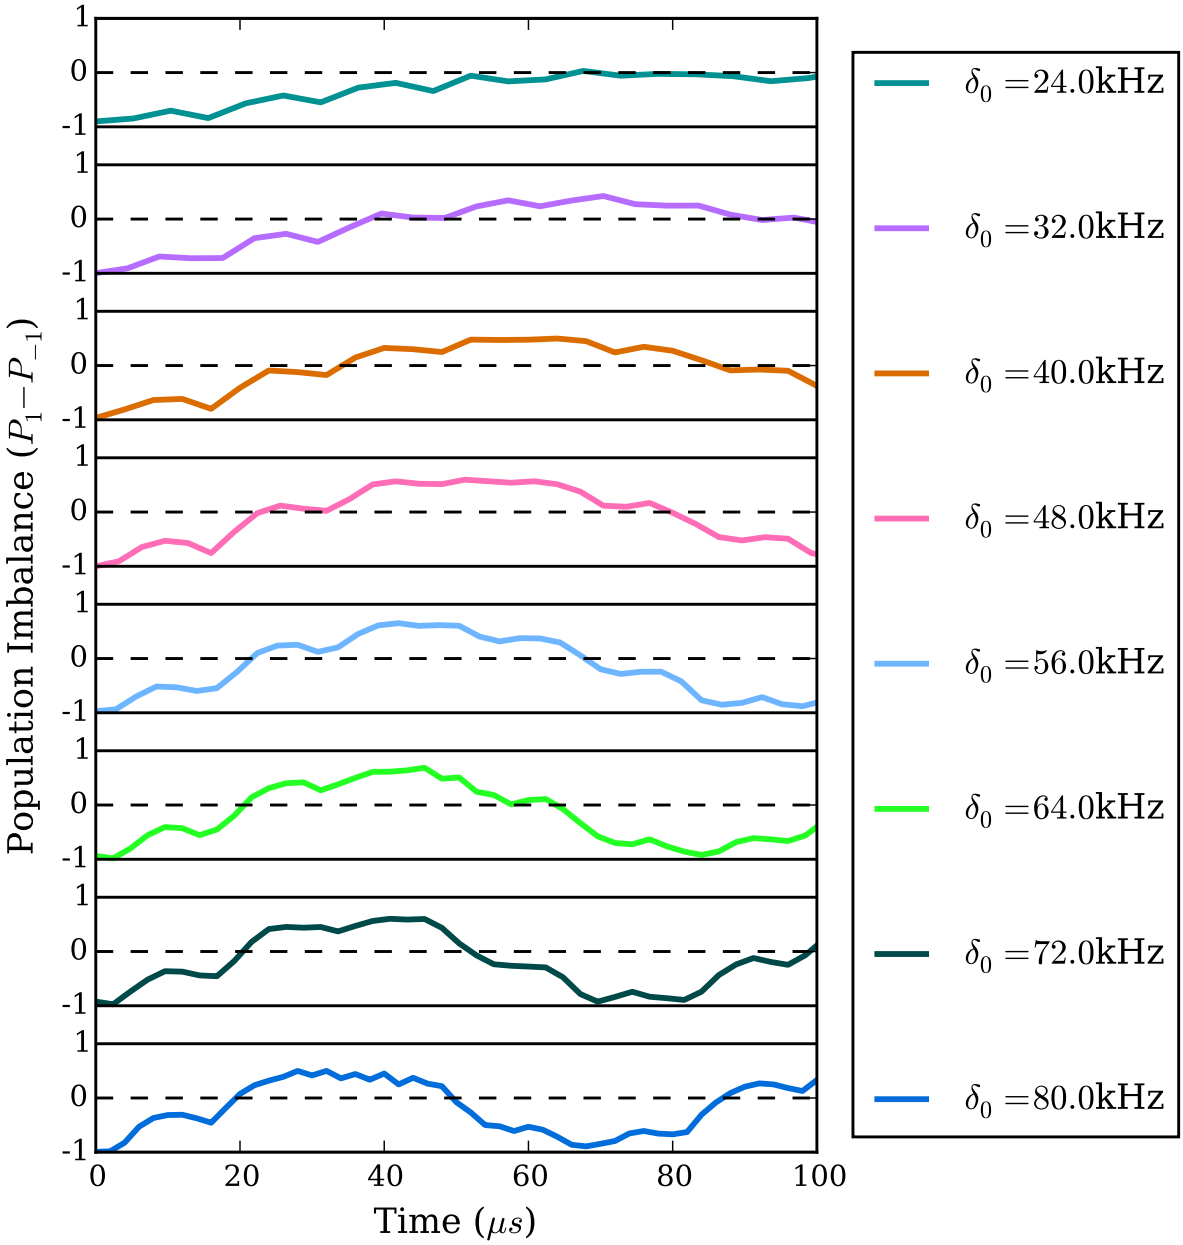

Figure 7.3: Experimental results showing the oscillations in the population imbalance (r_z) for a variety of noise-field amplitudes. In general, it can be seen that there is a slower oscillation from the noise field superimposed upon the smaller and faster oscillations from the RF coupling. The frequency of the oscillations due to the noise-field increases with the noise-field strength δ_0. If $\Delta' = 0$, we expect a full amplitude oscillation regardless of the field strength. The lessening of the oscillation amplitude for smaller δ_0 implies that for this dataset, $\Delta' \neq 0$.

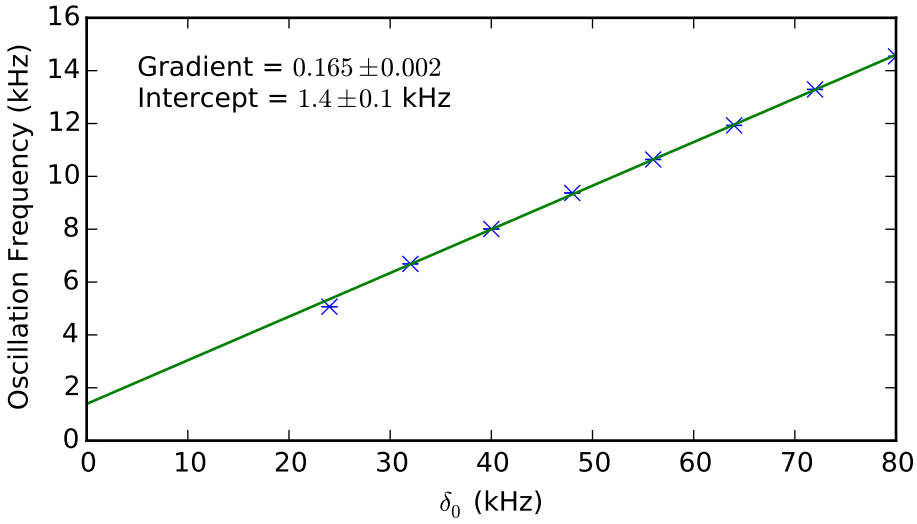

Figure 7.4: Relationship between parametric oscillation frequency and the amplitude of the noise field. The behaviour can be seen to be roughly linear as predicted.

Due to the similarity of the Hamiltonian in Equation (7.6) to the RF coupled Hamiltonian, we anticipate that the behaviour of the states should follow particular properties of Rabi flopping including the linear dependence of the resonant oscillation frequency on the amplitude of the applied coupling field which in this case is proportional to the noise field amplitude δ_0. Figure 7.4 shows the results of this test and it can be seen to roughly agree with the linear relationship. The frequency was estimated from fitting r_z to a phenomenological decaying sine oscillation.

Until now, little comment has been made on the validity of the second rotating wave approximation made to arrive at these parametric oscillations. Validity of the approximation requires that the discarded frequency ω_N is much larger than the inverse of the time dynamics of the remaining system, which is governed by the coupling strength Ω'. From the latter traces of Figure 7.3, we can see that this is not the case as the flopping frequency reaches parity with the faster oscillations of $\tilde{\Omega}$. The breakdown of the RWA for classical coupling fields can be treated with Floquet formalism such as the treatment by Shirley [85].

To describe the consequences of the RWA breakdown, let us consider a generic AC coupled two-level system with energy splitting $\hbar\omega_0$, associated Rabi frequency Ω and coupling field frequency ω_c. The first effect of the RWA breakdown that is usually discussed is a small shift in the resonance frequency required for the Rabi flopping. This is referred to as the Bloch-Siegert shift [86]. This shift is the consequence of an additional dressing of the states by the counter-rotating field which is the term neglected in the RWA. Beyond this, more exotic behaviour is manifest by the presence of population oscillation when the coupling frequency

is an integer fraction, or sub-harmonic, of the energy level splitting; $\omega_c = \omega_0/n$ with n as an integer (for simplicity, we neglect the difference resulting from the Bloch-Siegert shift). This corresponds to the co- and counter-rotating terms interacting to form a 'multi-photon resonance'[1]. We refer to these as sub-resonances with the most simple example given for $n = 3$.

For the normal 'single photon' resonance, the system moves from the ground state to the excited state, $|g\rangle \rightarrow |e\rangle$ whilst absorbing a co-rotating photon (the exponential term that moves with the rotating frame) and the corresponding reverse process. However, when taking the counter-rotating term into account (the exponential term that rotates counter to the rotating frame), consideration must be given to the transition path of the atom de-exciting, $|e\rangle \rightarrow |g\rangle$ whilst *absorbing* a counter-rotating photon from the field. Obviously, the massive difference in energy ($\Delta E = \omega_c - \omega_0 \approx -\omega_c$) between the start and end of the transition makes staying in this state untenable, however it can serve as a valid intermediate state. If we consider our two-level model system as a spin 1/2 system interacting with radiofrequency photons, the co-rotating and counter rotating terms correspond to the polarisation of the photon, which governs the restrictions between changing from state and the absorption of a photon. For example, if $|g\rangle \equiv |-1/2\rangle$ and $|e\rangle \equiv |+1/2\rangle$, then the co-rotating term corresponds to σ^+ polarised photons and the counter-rotating term corresponds to σ^- polarised photons.

This situation is illustrated in Figure 7.5a. Starting from the ground state,

- $|g\rangle \rightarrow |e\rangle$ — Absorb co-rotating
- $|e\rangle \rightarrow |g\rangle$ — Absorb counter-rotating
- $|g\rangle \rightarrow |e\rangle$ — Absorb co-rotating

is a path that consists of the absorption of three photons and the movement of the state from ground to excited. The total energy change for this process is $\Delta E = \omega_0 - 3\omega_c \approx 0$ so the sum of these processes is energy conserving. As this chain of state changes relies on non energy-conserving intermediate states, the total transition probability is very low. The result is, however, the same as the single photon transition case, population flopping between the undressed states. However the coupling strength varies cubically [85] with the amplitude of the coupling field rather than linearly in the case of normal one-photon absorption.

To excite a sub-harmonic transition that excites the atom using the absorption of an even number of photons, we must be able to include a photon whose absorption does not change the atomic state. Whilst the intermediate states do not have to satisfy energy conservation, they must comply with the selection rules

[1] Although the description comes from a general solution to Floquet theory with no mention of quantised fields, the multi-photon description is an intuitive picture. More on the equivalence in [82]

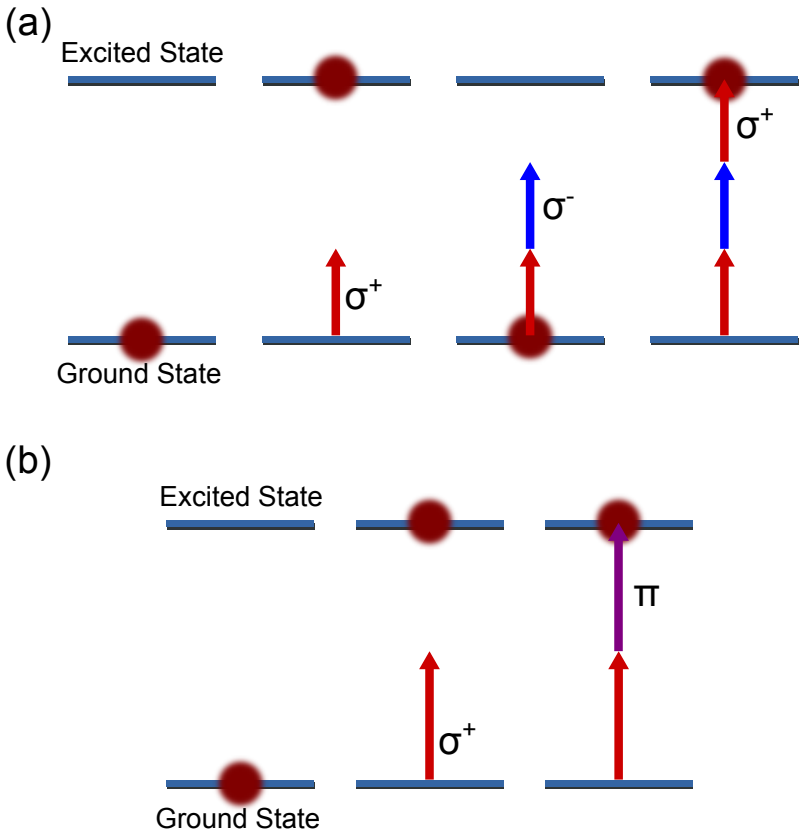

Figure 7.5: Schematic example of the 'intermediate' transitions for a (a) 3-photon and (b) 2-photon transition. The atom must go from ground to excited state on absorption of a σ^+ photon, from excited to ground for a σ^- photon and remain unchanged for a π photon. Note that this figure does not show all possible paths, for the 2-photon case, the order can be reversed and with the inclusion of π photons, the 3-photon case, absorption of one σ^+ and two π photons in any order is also possible.

of the transition. For the three-photon transition above, the co-rotating term allowed the atomic state to go from the excited state to the ground state whilst absorbing a photon only because of the opposite polarisation of the counter-rotating field.

To do a transition with an even number of photons, with a change in the atomic state, one of the intermediate transitions must be the absorption of a photon with no change in the atomic state. In the spin 1/2 example discussed above, this would be achievable with a π polarised photon. This requires some of the coupling field to be along the axis of the energy level splitting. This occurs in our case from the previously disregarded on-diagonal time dependence of Equation (7.5). This on-diagonal term can also be seen in Figures 7.2b and 7.2d as the component of the noise field vector parallel to the precession axis. As we do have these on-diagonal terms in our system, we should expect to see both even and odd sub-resonances. The case for a two-photon transition is shown in Figure 7.5b. The introduction of this term also modifies the calculated flopping frequencies for the odd photon number transitions as well, as there are more paths available (eg. absorption of two π photons and one σ^+ photon). The exact forms are given in Shirley's thesis [82] but the main result is that for an n photon transition, the coupling strength scales as $\delta_0^n/\omega_N^{n-1}n!$. This scaling is the reason such transitions are not seen in the valid rotating wave approximation where the ratio of the Rabi frequency to the oscillation frequency is very small.

This behaviour can also be explained using the frequency modulation argument discussed previously. When the modulation index (in this case δ_0/Δ) is increased [84], more sidebands start appearing in the frequency spectrum $2\omega_N$ away from the carrier, and then $3\omega_N$. The strength of the sidebands in this regime is governed by the relevant Bessel functions of the first kind $J_n(\beta)$ where n is the order of the Bessel function and β is the modulation index. By setting ω_N to half the required frequency, the 2nd order sideband will have the correct frequency required to induce the parametric oscillations. At small β, these approximate to linear, quadratic, cubic etc. as functions of β which is consistent with the treatments outlined above.

An interesting property of frequency modulation is that at high modulation indices, the Bessel functions start to oscillate, this means that as the modulation index is increased, the frequency scaling $\propto \delta_0^n$ will start to breakdown and at some point start to decrease. Preliminary numerical simulations of a two-level system have been conducted that agree with this prediction and is a topic for further study.

Experimental observations of the sub-resonances can be seen by setting the noise field oscillation frequency to a corresponding sub-harmonic frequency. Because of the relative narrowness and scaling of the sub-resonances, they are much more sensitive to the detuning field and thus homogeneities in the bias field significantly reduce the contrast of the oscillations.

We show examples for the time dynamics of the two photon transitions in Figure 7.6. The scaling relations for this flopping frequency against the coupling field amplitude were tested and the results shown in Figure 7.7. It is consistent with a quadratic form.

Aside from these different forms of the resonant coupling frequency, the off-resonant case follows the same behaviour as the fundamental case, i.e. the amplitude of the population oscillations decreases with increased detuning and the frequency of the oscillations increases. The relative width of the response of each transition is governed by the resonant flopping frequency.

7.3 Random Broadband Spectrum

With the understanding provided by the previous section, we can conceive a situation where, if $\delta(t)$ is a signal containing many different frequencies, we may expect a response if the frequency content spans at least one of these parametric resonances. It can easily be conceived that there is potentially complicated interplay between different frequencies within the width of a single resonance as well as the interplay between different excited resonances. Also considering the previously discussed frequency modulation arguments, there is potential for any response to not just increase as a function of the strength of the noise field, but to also decrease.

The tremendous degree of freedom for functional forms for $\delta(t)$ is a dizzying prospect. Here, we attempt to identify particular regimes that may be of interest for the form of $\delta(t)$ using the original motivation developed in Chapter 1 of transport in quantum networks in the presence of decoherence or environmental coupling.

While it is possible to use $\delta(t)$ for coherent manipulation, we are primarily interested in the effects of non-deterministic $\delta(t)$'s. We consider $\delta(t)$ to be the result of a random, or stochastic process. There is extensive literature of stochastic processes (see Hoel [87], Taylor [88] or Grimmett [89] for introductions). We will consider a random process simply as some physical mechanism that creates a stochastic signal in time that affects our system Hamiltonian through the time dependent detuning $\delta(t)$. The nature of the stochasticity is that there are many different realisations of the process (also known as samples or trajectories). If we take many realisations of this stochastic process, we get a set of functions in time that are linked by some underlying characteristics of the process. The possible complexity of general stochastic processes puts a very weak limit on the properties of this relationship.

We narrow this generality by considering the physical situation of a system embedded in some environment that is unaffected by the system[1]. The environment creates a random field that couples to our system via the

[1] Although please refer to section 1.1.3 for a possible way to consider systems that can affect the stochastic process

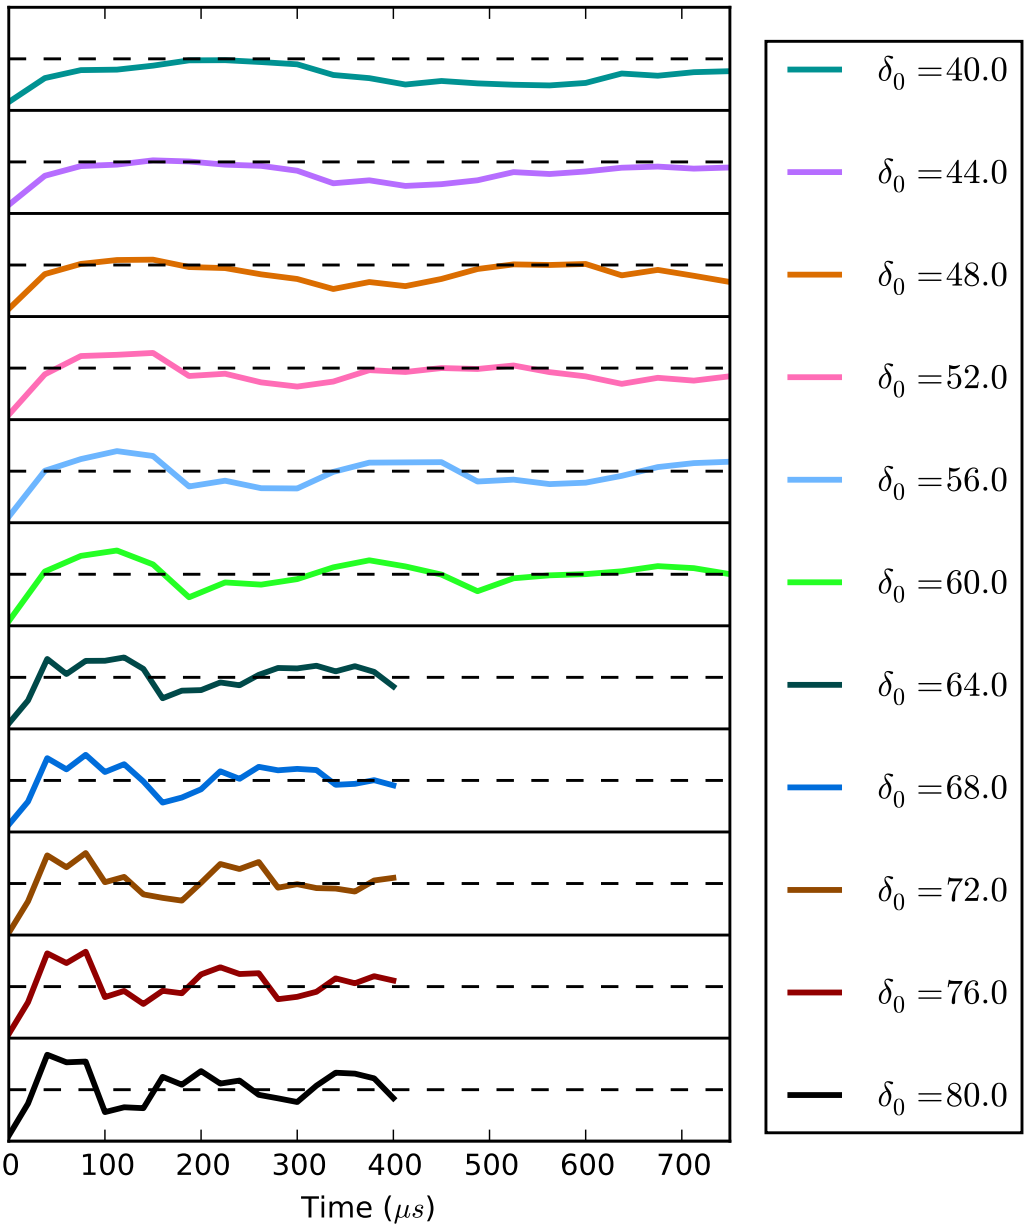

Figure 7.6: Experimental results showing oscillations in the population imbalance (r_z) for a variety of noise-field amplitudes for the $n = 2$ sub-resonance. The lessening of the oscillation amplitude for smaller δ_0 implies that $\Delta' \neq 0$.

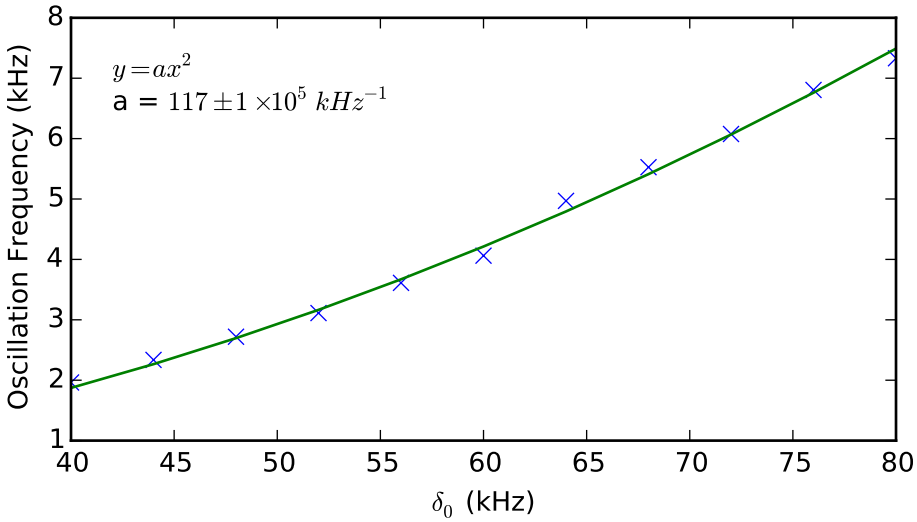

Figure 7.7: Relationship between the $n = 2$ parametric oscillation frequency and the amplitude of the noise field. The behaviour can be seen to be roughly quadratic as predicted.

$\delta(t)$ terms. We consider that the system has some 'task' or purpose that it is to fulfil in the presence of this fluctuating field. The task is of the nature that it is expected to repeat it many times. For the FMO complex example discussed in Chapter 1, this is the propagation of an exciton to a target site within a particular time frame. This is expected to occur many times in a single complex, each time with a reset starting state (the exciton arrival). We consider the analogue to this for our system to be the change from an initial state in which one of the end m_F states is wholly occupied (which we will call the *Initial State*), to the opposite m_F state (which we will call the *Target State*) within a particular time frame. It should be capable of doing this task well for many realisations of the noise process.

We restrict the fluctuating field by the following conditions:

- **Expectation of Zero-Mean** - Any mean term for the fluctuating field $\delta(t)$ can be taken into account with the stationary detuning Δ.
- **No definite start point** - The characteristics of the noise process should not be relevant to time.

An important characteristic pertinent to this discussion is the comparison between the *average response* of the system to a stochastic process to the *individual response* of the system to each realisation of the process. The average response is some ensemble average of the responses to different realisations of a given 'type' of noise. Considering the situation above of the fractional population of a target state at a given time, this could simply be the mean of the state population over the different realisations. We can conceive a system

that might conduct its task very well for a limited number of realisations and poorly for the remainder of the realisations compared to a system that performs moderately well for all realisations. This comparison in some way reflects the 'mean' and 'variance' of the system's response. A small variance of the response implies, in some sense, a robustness of the system to the statistical fluctuations associated with the stochastic noise. It is worth reiterating however that the objective is not to consider systems that are resilient to the noise field but which instead utilise the noise field to enhance their performance of the stated task.

In the remainder of this section, we outline the computational methods used in the simulation of a stochastic process, numerical simulations of the dynamics of the system described above in the presence of band-limited white noise, the corresponding experimental implementation of such a process and the comparison between these experimental results and the numerical simulations. The work done in this section is of a preliminary nature, so as of writing there has been little extension of the small parameter space discussed. We follow the results with a brief discussion of the possible parameter spaces of interest. Further extensions of the work are discussed more extensively in Chapter 8.

7.3.1 Spectral Shaping

One of the primary areas of interest in environmental coupling and decoherence processes is the spectral shape of the noise that it generates. There are many examples in different systems of the differing types of noise spectrums including 1/f [90, 91], Ohmic [21, 92] (as well as sub- and super-ohmic) and Lorentzian [93].

To generate the points of the random signals with a defined power spectrum, we use a very elementary approach of taking N samples of zero-mean Gaussian-distributed random values to create a sample of white noise (i.e. noise with a flat power spectrum), we then take the Fast Fourier Transform (FFT) of this time series. The frequency spectrum of these uncorrelated values forms a flat frequency response when averaged over many realisations. Spectral shaping is achieved by multiplying this FFT result by a spectral transfer function and then taking the inverse Fourier transform of the shaped spectrum. This process leads to a set of points in the time domain that correspond to a particular realisation of a random process with the specified spectral shape. This process of spectral shaping is shown in Figure 7.8.

It should be noted that noise generation is typically not done using this method and that more complex but accurate methods are used [94–96]. This method was used as it was simple and intuitive to the concept of applying multiple single frequency oscillations of the kind discussed in Section 7.2

It is also worth noting that this method of noise generation implies the application of a non-causal filter to the initial white noise. This is because the multiplication of the white noise spectrum and a window

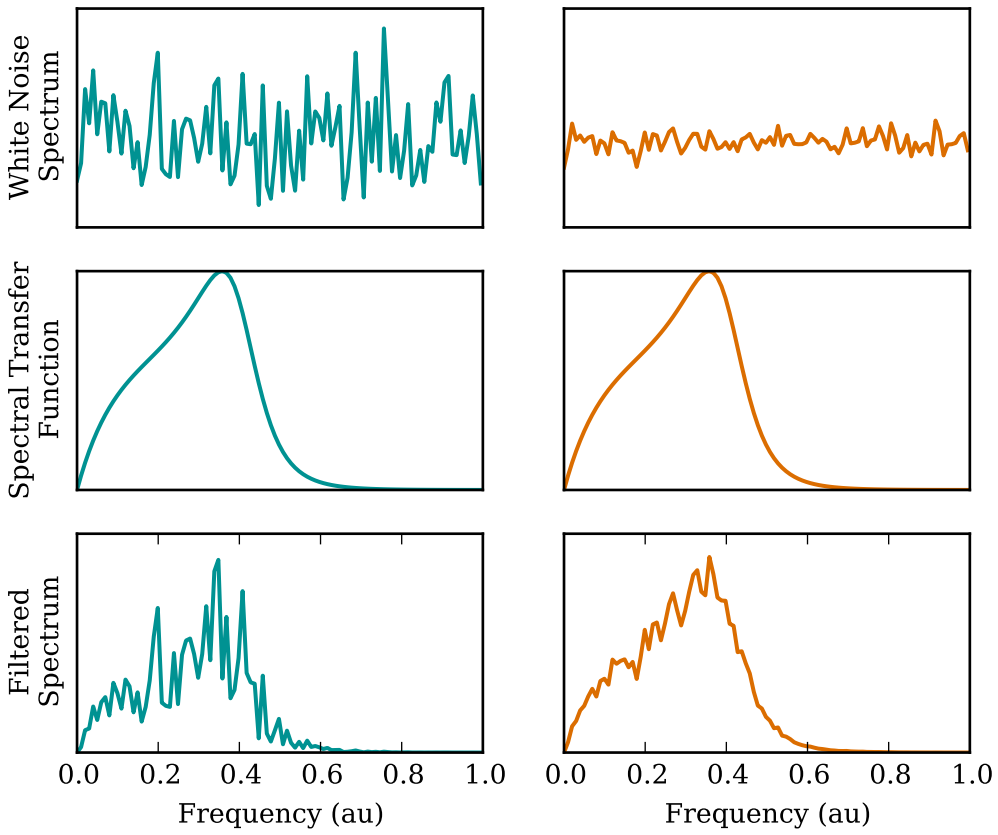

Figure 7.8: Example of the spectral shaping, White noise (spectrum shown in upper) is multiplied by a frequency space transfer function (middle) to leave the filtered spectrum (lower). The results for a single realisation of white noise are shown on the left and the mean over 25 realisations on the right.

function in frequency space is equivalent to a convolution in the time domain of the white noise and the Fourier transform of the filter. If this convolution filter has elements at positive time, then the filtered noise signal has been affected by the white noise at future time. We do not anticipate this to be an issue or have repercussions regarding Markovian approximations etc. as the system is never exposed to the original white noise.

We now consider the simple case of Band Limited White Noise (BLWN). The contribution from all of the frequencies up to a specified cutoff frequency is finite and equal and the contribution from frequencies above this cutoff is zero. A good parameterisation space for such a process is the cutoff frequency and the spectral density of the signal at frequencies below the cutoff frequency or alternatively, the cutoff frequency and the RMS value of the noise signal.

7.3.2 Computer Simulations

Computer simulations were conducted for the Hamiltonian shown in Equation (7.1) with δ being the form of band-limited white noise to identify regimes of interest. The numerical simulations solve the time dependent Schrödinger equation using standard ordinary differential equation numerical integration techniques. The result is the time dynamics of the state populations. We are particularly interested in non-resonant coupling ($\Delta > \Omega_{RF}$) as, in the absence of noise, the transport performance to the target state is low. We attempt to see an increase of the transfer with the presence of the noise. Different values of Δ/Ω were trialled which had similar behaviour. We will first focus on the work done for $\Delta = 5\sqrt{2}\ \Omega$ and then broadly discuss comparisons to other detunings. The white noise cut-off frequency was set to 12 times $\tilde{\Omega}$ for this case. The noise amplitude is set by the amplitude of the spectral transfer function which we translate to an expected RMS value, δ_{RMS}. This was varied from $\delta_{\mathrm{RMS}} = 0$ to $\delta_{\mathrm{RMS}} = 100\sqrt{2}\ \Omega = 40\Delta$. This end point is a regime where the noise totally dominates the on-diagonal term and the energy level splittings intrinsic to the system, Δ should be irrelevant. The time dynamics were solved to $t = 40/\tilde{\Omega}$.

For these conditions, 50 different realisations of the noise process were simulated and the populations at different time steps recorded for each. Examples of the population dynamics over a few different realisations, along with their collective average and variation are shown in Figures 7.10, 7.11 and 7.12 for the noise RMS values of 0.4Δ, 3.6Δ and 19.6Δ respectively. The behaviour of the state populations in the absence of noise is shown in Figure 7.9 for comparison.

We can see that any coherent oscillations are washed out by the noise in the individual realisations and also by the averaging over the realisations. For the average response, what remains is a general convergence

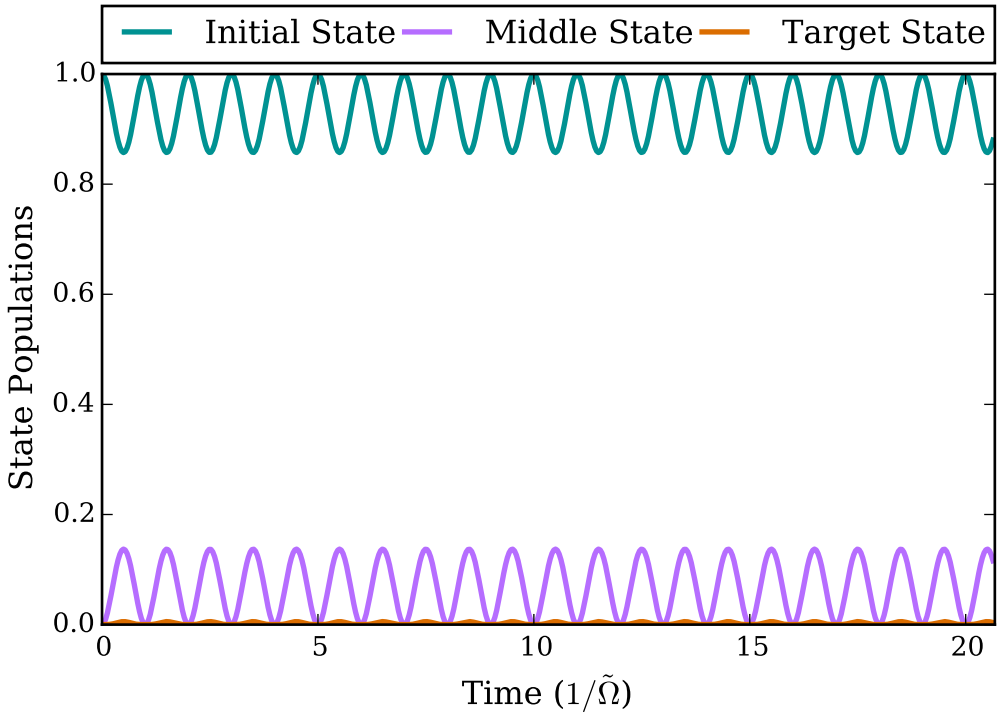

Figure 7.9: Numerical simulations of the time dynamics of the states in the absence of noise for the case of $\Delta = 5\sqrt{2}\ \Omega$.

of the ensemble-averaged state populations towards an equal distribution, where the probability of being in each of the three states is 1/3. We see that for the large noise amplitudes (those with a δ_{RMS} value several times larger than the DC splitting Δ), this trend is slow compared to the case of moderate noise RMS. This is an interesting phenomenon that demonstrates an example where the transport efficiency to the target state is enhanced by the interplay between the quantum coupling and classical noise, that there is a 'sweet spot' in the amplitude of the noise to enhance the task performed by the system. This phenomena of a 'good' amount of noise is also found in dephasing simulations by Plenio [13]. The concept of having a non-trivial optimum noise value is key to the idea of the increased efficiency being an interplay between the quantum coupling and the decoherence of the environment. If the efficiency of the system only increases as the level of decoherence does, then the importance of the quantum nature is greatly diminished and equally if the best transport occurs in the absence of noise, then the system is better off isolated.

These results are consistent with investigations into the power spectrum of frequency modulation with band-limited white noise; the power within a specified frequency band close to the carrier increases and then decreases as a function of the magnitude of the noise.

The nature of the convergence of the population averaging to 1/3 each is consistent with the idea of

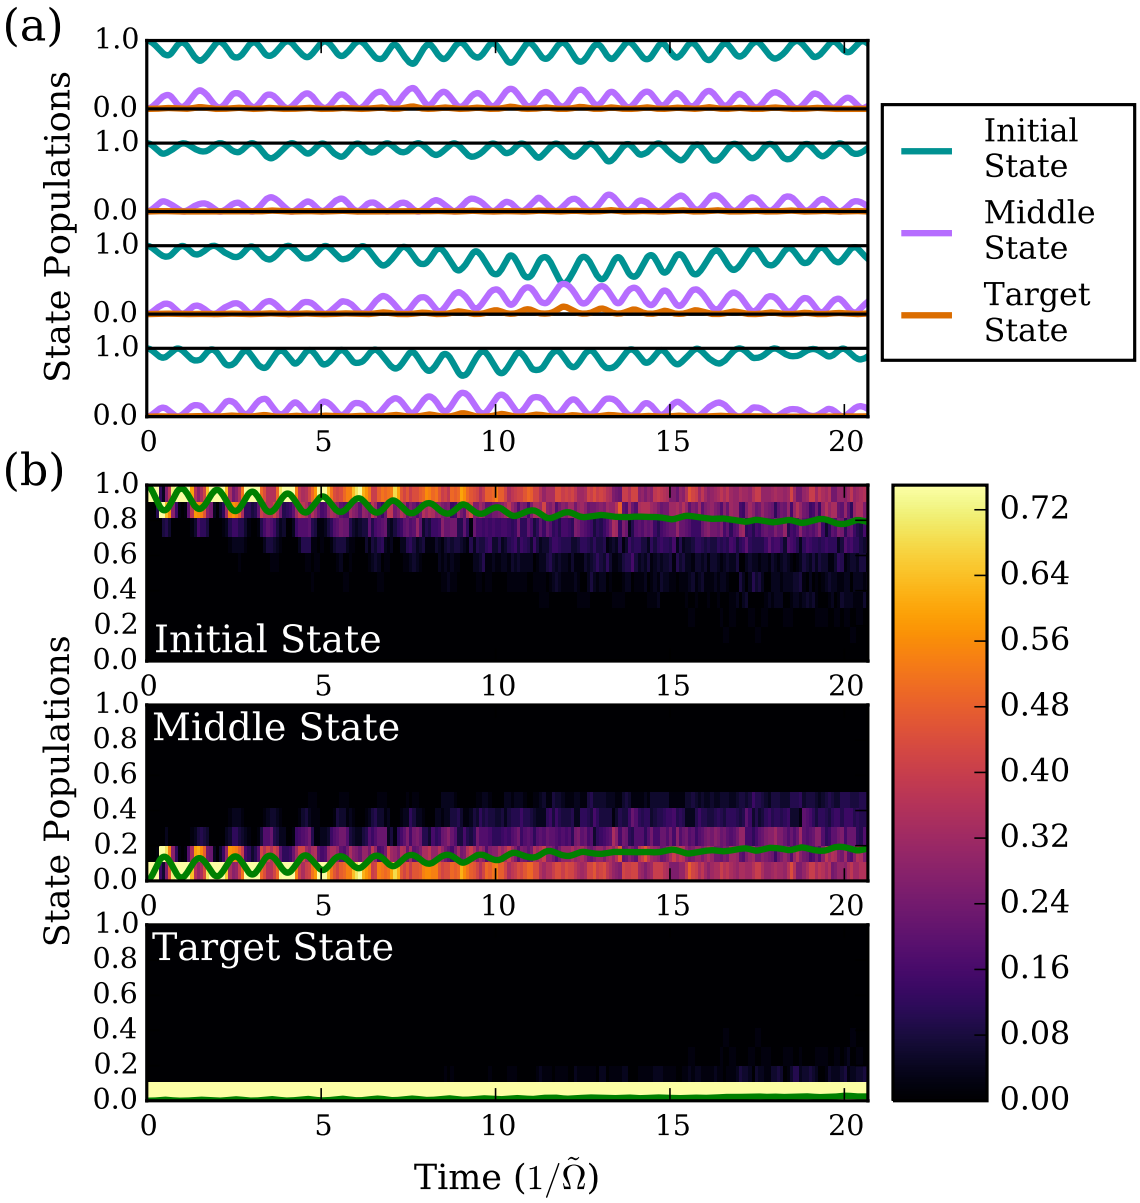

Figure 7.10: Numerical simulations showing the state time dynamics with a noise RMS of 0.4Δ. (a) shows the time dynamics of the state populations for several realisations of the noise. (b) shows the average of the state populations over 50 realisations (green line) on top of a binned distribution of the realisations. i.e. Each vertical stripe is a histogram of the relative state populations at a single time-point. Note the colour scale has been clipped at 0.75 for greater contrast in the less occupied bins. A bin value of 1.00 means that every realisation is in that histogram bin.

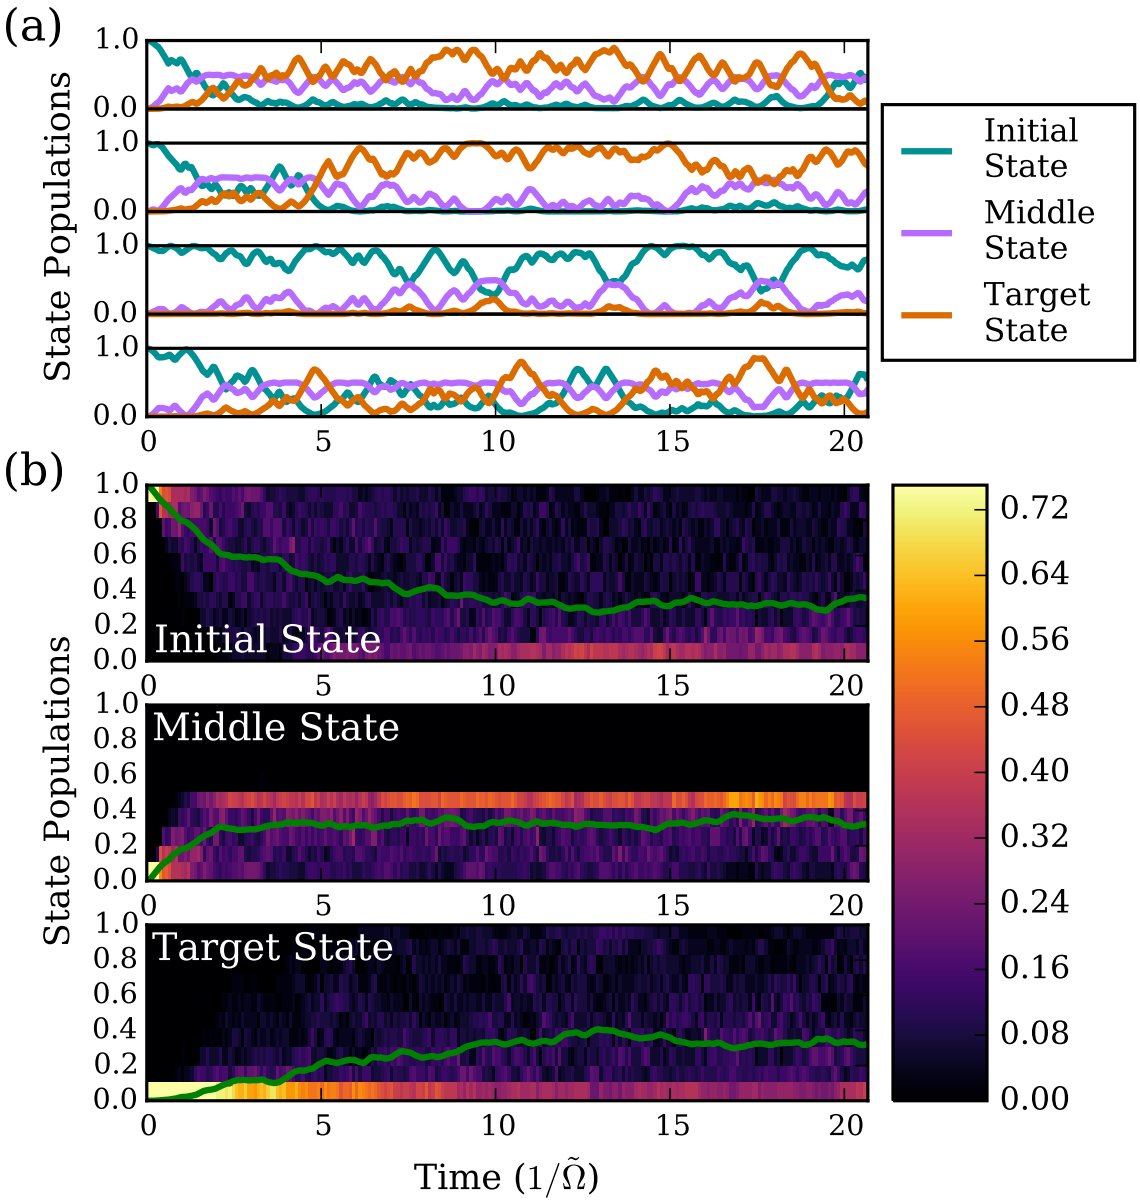

Figure 7.11: Numerical simulations showing the state time dynamics with a noise RMS of 3.6Δ. (a) shows the time dynamics of the state populations for several realisations of the noise. (b) shows the average of the state populations over 50 realisations (green line) on top of a binned distribution of the realisations. i.e. Each vertical stripe is a histogram of the relative state populations at a single time-point. Note the colour scale has been clipped at 0.75 for greater contrast in the less occupied bins. A bin value of 1.00 means that every realisation is in that histogram bin.

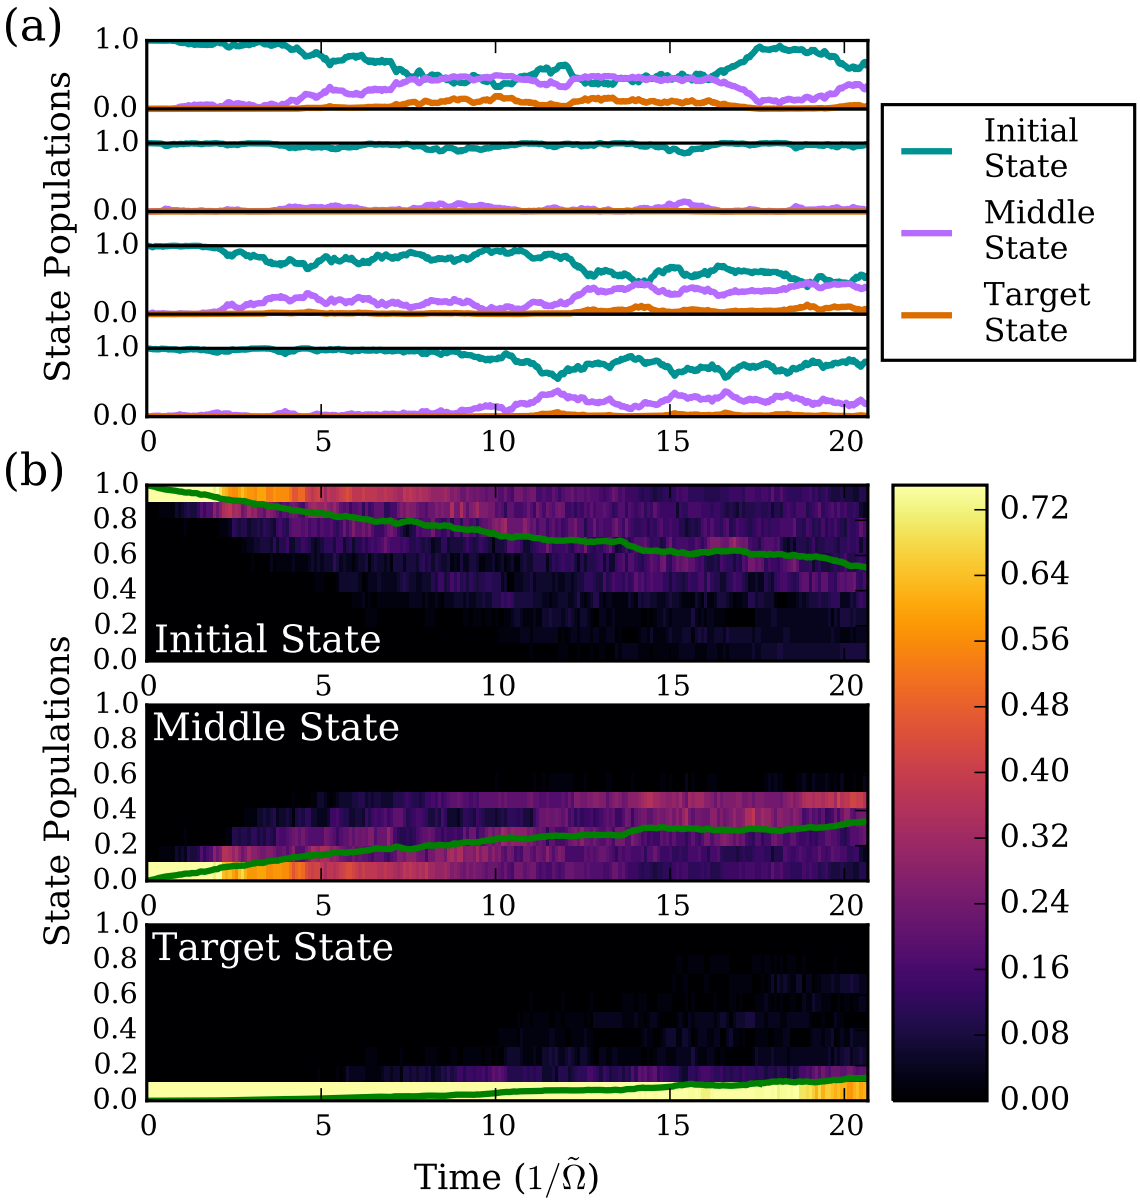

Figure 7.12: Numerical simulations showing the state time dynamics with a noise RMS of 19.6Δ. (a) shows the time dynamics of the state populations for several realisations of the noise. (b) shows the average of the state populations over 50 realisations (green line) on top of a binned distribution of the realisations. i.e. Each vertical stripe is a histogram of the relative state populations at a single time-point. Note the colour scale has been clipped at 0.75 for greater contrast in the less occupied bins. A bin value of 1.00 means that every realisation is in that histogram bin.

the classical noise field behaving as a 'thermalising bath' at infinite temperature. The distribution of the realisations has much in common with what would be expected distributing the states evenly over the surface of the Bloch sphere. Averaging the state populations over the entire surface of the Bloch sphere also yields the same population averages and population distributions. This implies that the noise process acts as a classically diffusive process.

We can visualise the noise enhancement feature more clearly by considering a fixed time t_0 and showing the average state populations for differing noise amplitudes. The result for a single time $t_0 = 20/\tilde{\Omega}$ is shown in Figure 7.13 with the population realisation distributions and with multiple times for comparison in Figure 7.14. Common to all is that increasing the noise from zero leads to a relatively rapid increase towards the 1/3 equilibrium value before either a turning point or plateau followed by a comparatively slow decrease back towards remaining in the initial state. The population distributions seem to agree with the corresponding distributions in the time series results, i.e. the results of the different realisations are not particularly localised. The turning point at which the noise RMS could be considered 'optimal' is approximately $2-3\Delta$ which implies the intrinsic energy level structure of the system still plays a role in this regime.

The equivalent graphs of the population distributions for the different ratios of $\Delta = 10\sqrt{2}\ \Omega$ and $2.5\sqrt{2}\ \Omega$ are shown in Figures 7.15 and 7.16 respectively. Broadly, the behaviour is very similar. We can see that the tendency to the equilibrium populations is significantly faster with a smaller Δ which is not surprising as there is already significant coupling between the bare states in the absence of noise. The optimal noise point also seems approximately the same relative to Δ, which is consistent with the 'blurring' of the energy levels argument made in Section 1.1.4.

7.3.3 Experimental Results

With all previous points in place, we implemented the band-limited white noise discussed above on the experiment. Due to the 100 kHz limit of the noise-field driver, the resonant Rabi frequency Ω was reduced to 10 kHz by decreasing the supplied RF power. This was to allow testing of the regime where the noise amplitude is several times larger than the energy level splitting of the bare states to have a large noise cut-off frequency compared to the Rabi frequency. The detuning was not accurately calibrated for the data taken but is believed to be roughly 30 kHz based on previous resonance measurements, which therefore leads to a generalised Rabi frequency also of approximately 30 kHz. The current limit of the noise-field current driver determines the upper limit on the amplitude of noise that can be applied. As the signal is the form of white noise, this limit is reached intermittently for RMS values that approach the limit but qualitatively, this does

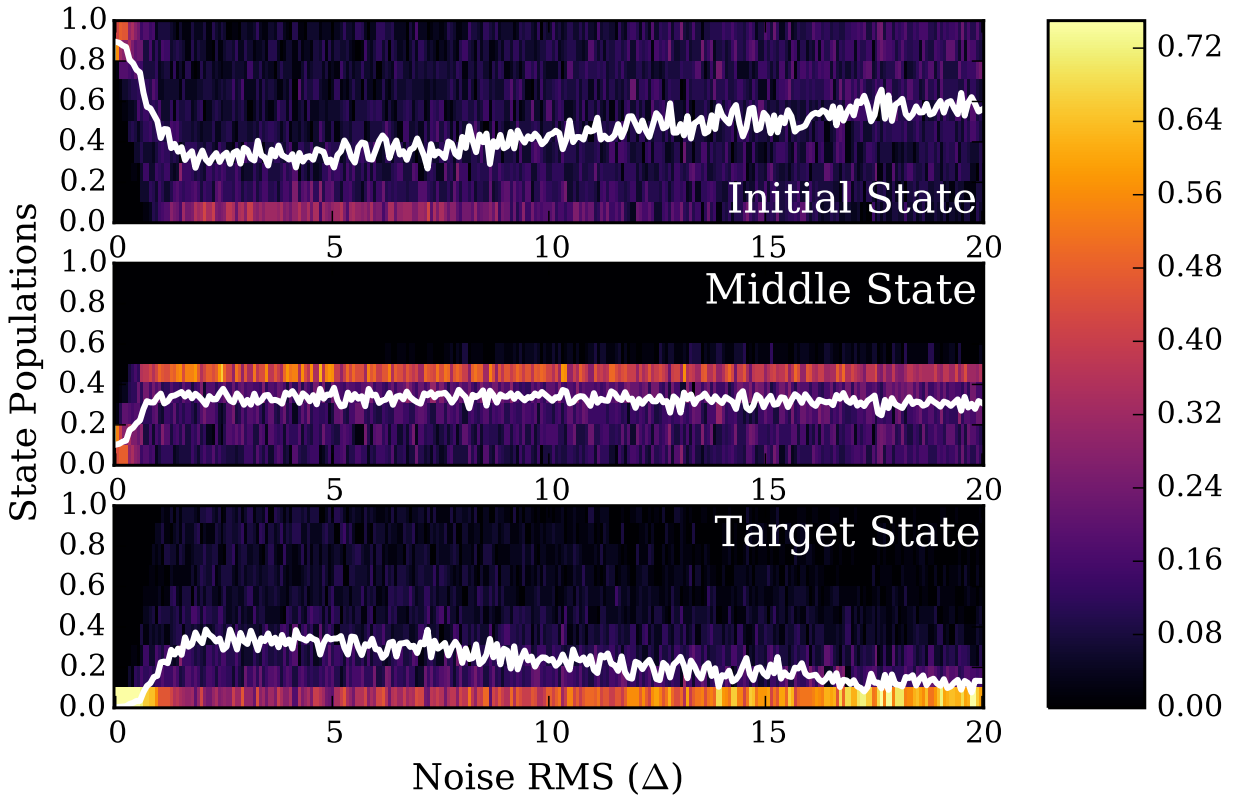

Figure 7.13: Numerical simulation results directly showing the effect of the noise amplitude for a fixed evolution time of $20/\tilde{\Omega}$. The population transfer process has a clear optimum value for the noise RMS.

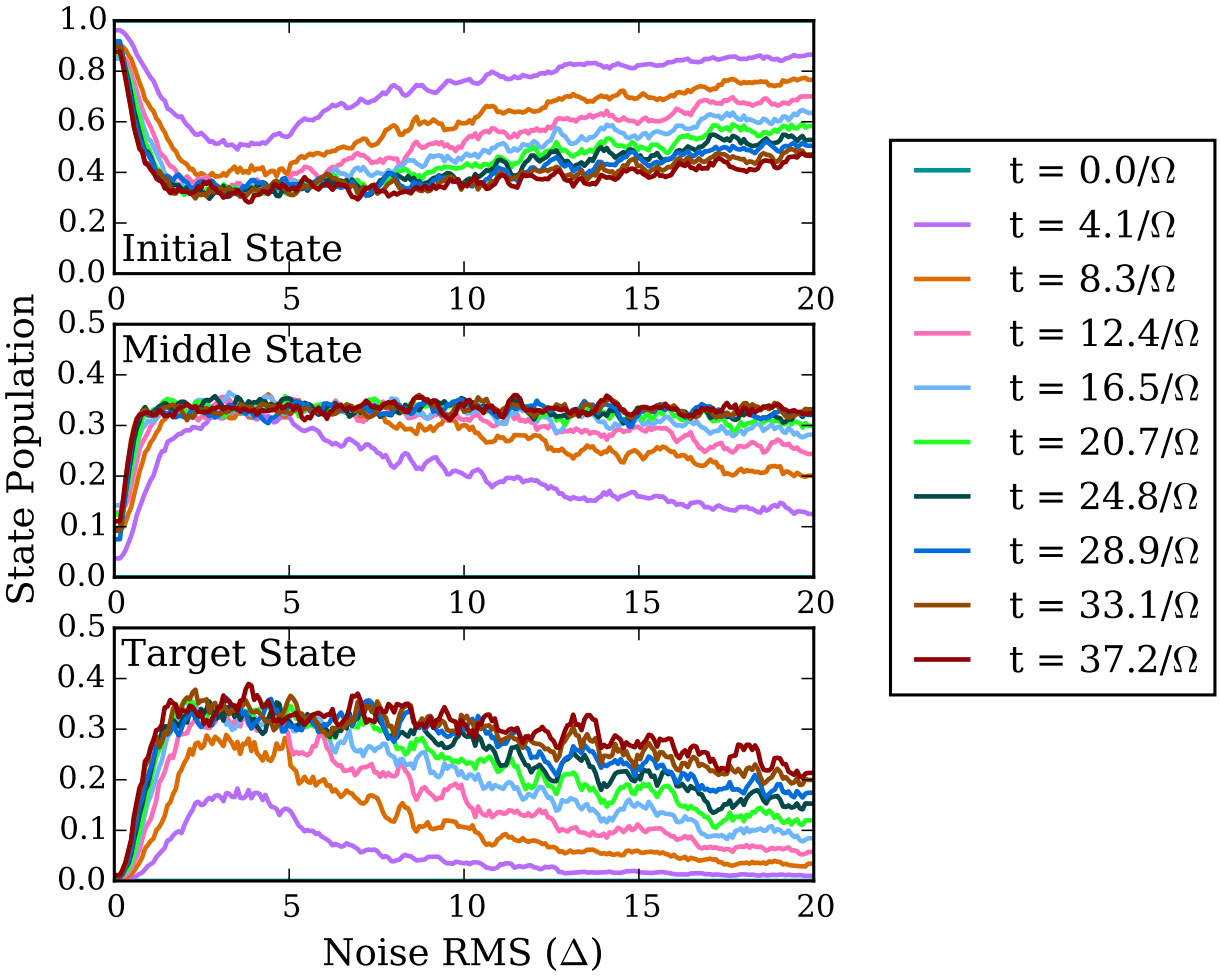

Figure 7.14: Numerical simulation results of the average populations for different evolution times and noise field amplitudes for the different states. Increasing time leads to a movement towards the equilibrium population distribution of 1/3 in each state. Note that the datasets have been filtered by a 5-point moving average for clarity of comparison.

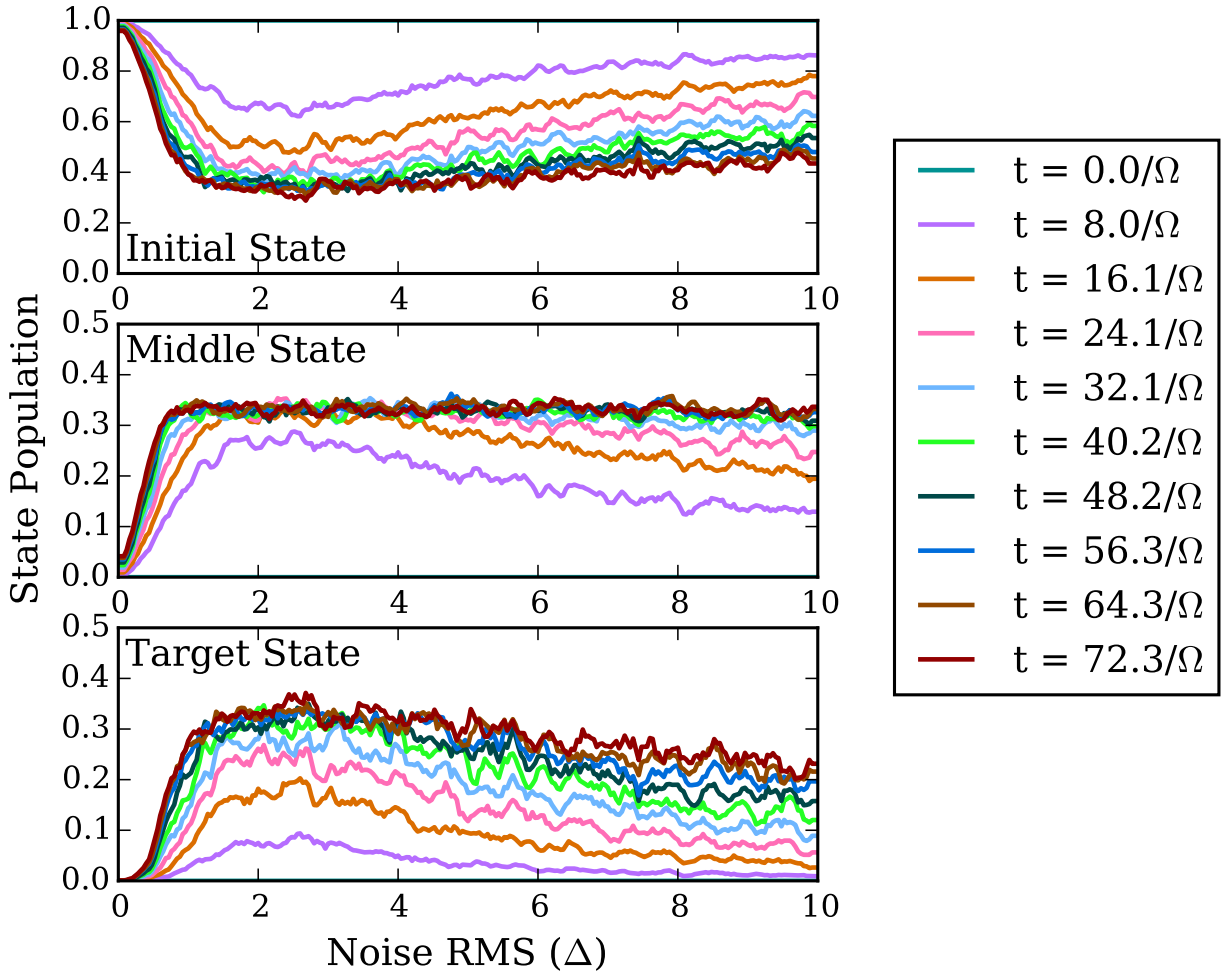

Figure 7.15: Numerical simulation results showing the same metric as in Figure 7.14 with $\Delta = 10\sqrt{2}\Omega$. It can be seen that the population tends more slowly towards the equilibrium position.

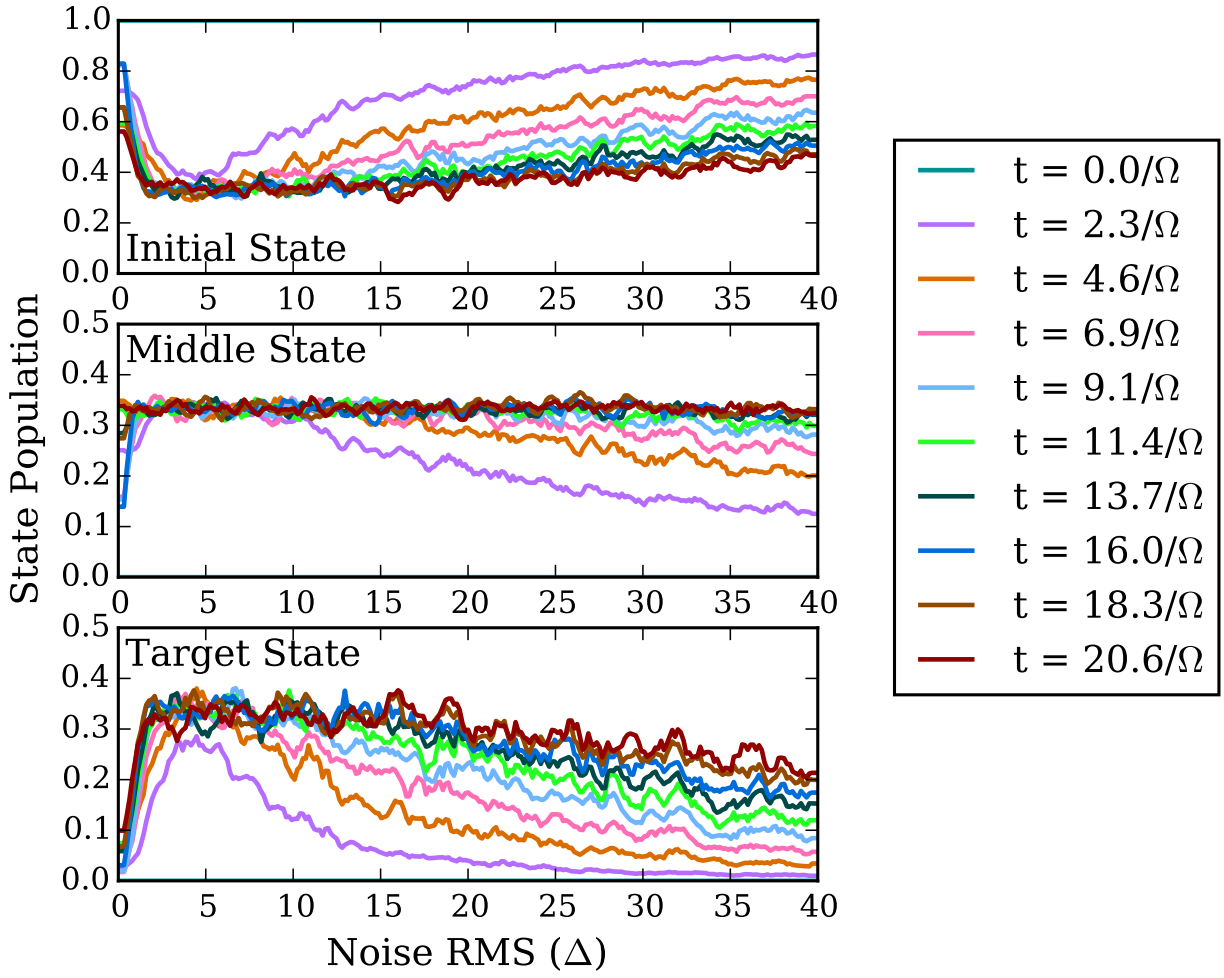

Figure 7.16: Numerical simulation results showing the same metric as in Figure 7.14 with $\Delta = 2.5\sqrt{2}\Omega$. It can be seen that the population tends more quickly towards the equilibrium position.

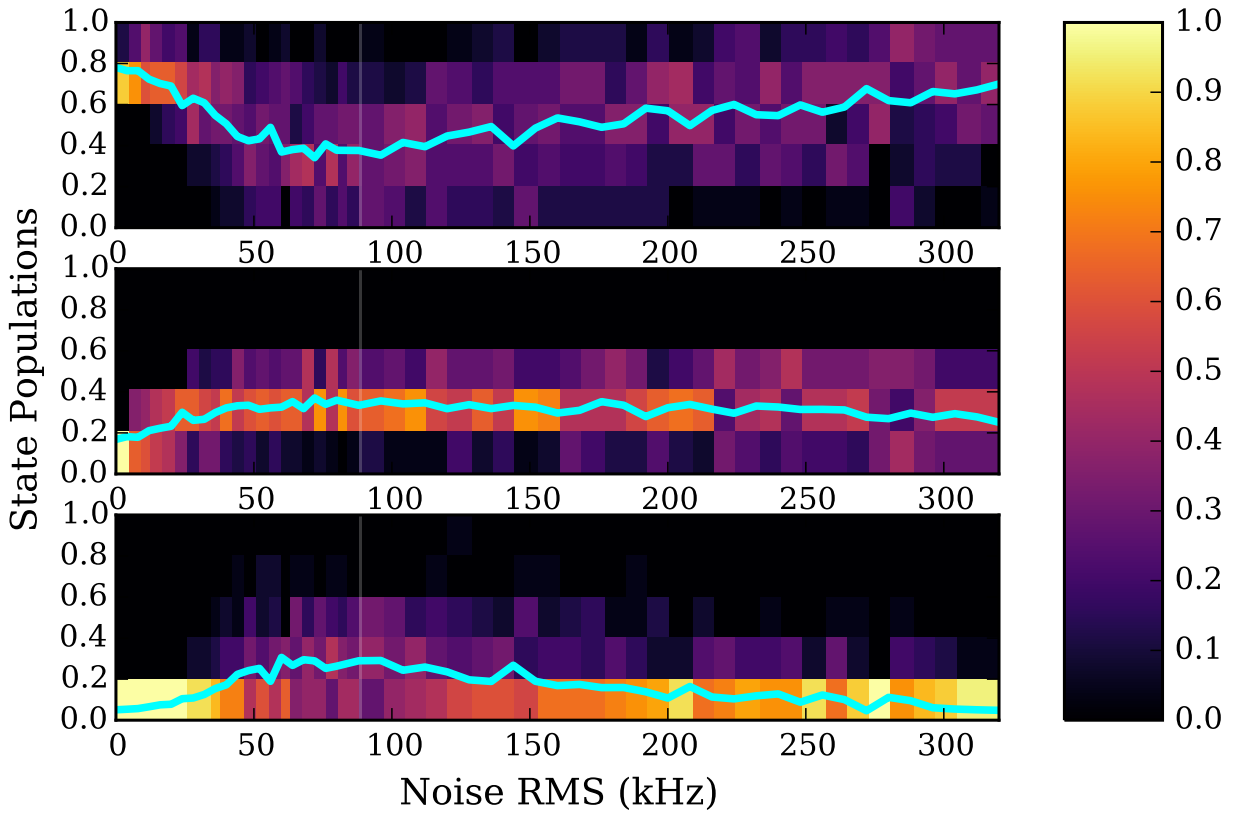

Figure 7.17: Experimental results for the system under the application of band-limited white noise fields of different strengths for $t_0 = 100\,\mu s$. The averaged initial (upper), middle (middle) and target (lower) states' populations broadly follow those predicted by numerical simulation but the distribution of the realisations (shown by the histograms) is much more localised that those of the numerical simulations.

not appear to significantly affect the waveform. The noise RMS was ranged from zero to 350 kHz, each time measuring the state populations after 100 µs of evolution. For each of these parameters, 25 different realisations of the noise were measured. The results are shown in Figure 7.17.

It seems apparent at first inspection that the experimental data matches the simulation in a few qualitative ways, including the presence of a 'sweet spot' for the noise RMS that produces the greatest response in the system. The general shape of the population averages rise close to an equilibrium position and then decays slowly at higher amplitudes. In this case, the turning point noise RMS is approximately 70 kHz, which is approximately twice as large as the estimated detuning. This is consistent with the presented numerical simulations and again implies that the intrinsic energy level structure of the system still plays a role.

An interesting difference that can be seen between the experimental results and the classical simulations is in the distribution of the populations of the different noise realisations. It appears that in the experimen-

tal implementation, the variance of the populations between the different realisations is much smaller and localised to the mean in contrast to the relatively scattered results from the numerical simulations. More explicitly, in the experimental results, the resulting populations seem to be affected less by the differences between each noise realisation. Leaving aside possible reasons for the difference between the simulated and experimental results for now, the observation of collective behaviour of the different realisations is very interesting. The system's sensitivity to the noise leads to changes in the occupied states whilst being seemingly resilient to the stochastic nature of the individual realisations of the noise process. This is exactly the behaviour that would be considered desirable in a 'practical' decoherence-assisted quantum system that had a dedicated task to conduct.

It should be re-stressed that these results are preliminary and that little work has been undertaken to extensively vary parameters and test to see if this effect is caused by some experimental imperfection not present in the computer simulation. The following aspects have been considered but not tested:

- **Hamiltonian Inhomogeneities** - As discussed previously, the small physical size of the noise coil along with the alignment issues of the bias coil have caused some inhomogeneity in the Hamiltonian across the cloud for the detuning and noise field. The principal consequence that could be related to this measurement is that it will generally lead to clouds less likely to have a high concentration of atoms in a single state at a point in time as the evolution of all the different Hamiltonians experienced would have to lead to the same end state from the initial prepared state, which is unlikely. Ultimately, our cloud population measurements cannot distinguish between the quantum projection of a pure state onto the basis vectors and the presence of many different pure states in the sample.

- **RF Finite Switching Time** - As discussed in Chapter 6, the finite switching time of the radio-frequency coils has a complicated effect on the propagation of the state vector during the turn-on and turn-off times. This effect would be further complicated by the presence of the noise field during the switching process. In future experiments, the coupling field could be turned on and off before and after the noise window respectively to minimise any additional complications.

- **Noise Spectrum** - As discussed in section 7.1, the eddy currents in the chamber mean that the exact applied spectrum to the atoms is unknown. Rather than a flat frequency spectrum, the noise likely has an element of low pass filtering. However it seems unlikely that this would so drastically change the behaviour of the response to the noise.

- **Noise Waveform Clipping** - As discussed earlier, the current limit of the noise coil driver causes

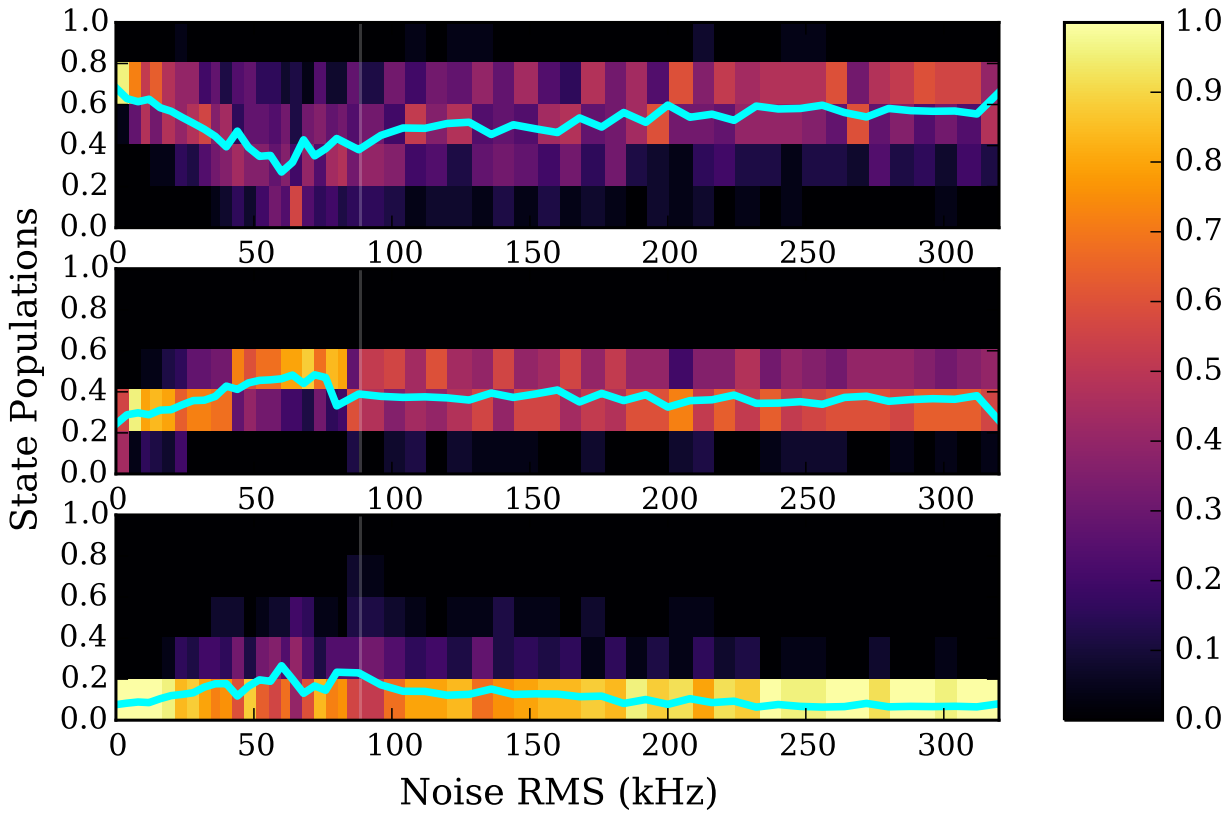

Figure 7.18: Experimental results for the same scheme as in Figure 7.17 but with a evolution time of 75 µs.

clipping of the noise waveform for moderate to large RMS values. This should only affect the data where the noise amplitude is large and is thus unlikely to cause the difference in behaviour at all noise amplitudes.

With the understanding of these potential flaws with the current experimental apparatus, we consider more data taken at different evolution times of the system of 75 µs and 150 µs. These are shown in Figures 7.18 and 7.19 respectively. Note that all three of the datasets presented were taken on different days and whilst we believe that there should be no significant variations, caution should be considered when comparing them directly.

Their behaviour broadly follows the $t_0 = 100\,\mu s$ case except for the expected behaviour of the three states being closer to equilibrium for the $t_0 = 150\,\mu s$ case and closer to the initial state for the $t_0 = 75\,\mu s$ case.

Whilst formal data is not presented for time resolved dynamics, initial investigations confirmed that the system was deterministic in the sense that the response to a particular realisation of the noise-field did replicate roughly the same distribution of states. It is therefore possible to build up time resolved dynamics

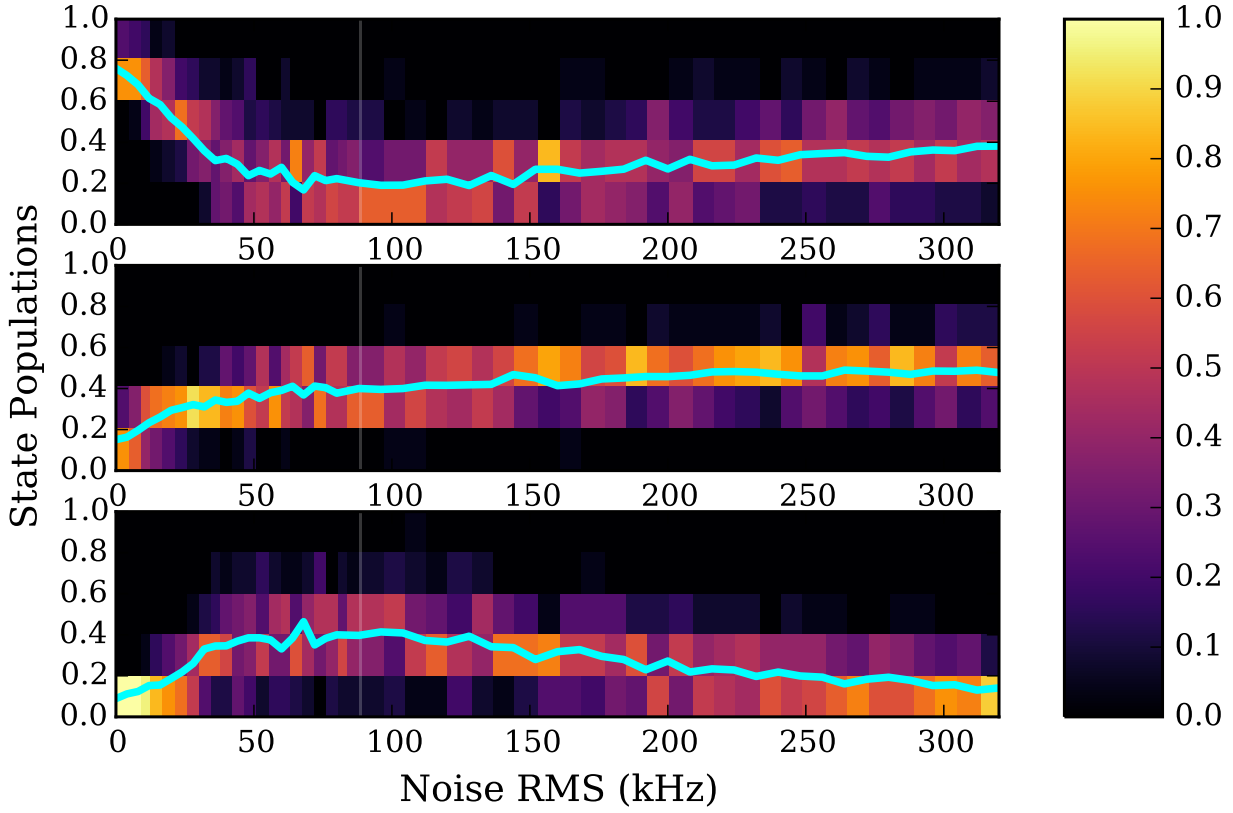

Figure 7.19: Experimental results for the same scheme as in Figure 7.17 but with a evolution time of 150 µs. Note that population settling to a ratio above (1/3, 1/3, 1/3) is likely a cloud fitting issue as the Stern-Gerlach batteries were partially drained and not providing sufficient contrast for reliable fitting.

to see the tendency to equilibrium position and then remain there but sufficient data was not taken for these results.

7.4 Concluding Remarks

In this chapter, we have discussed and shown results of coherent manipulation of the populations using parametric oscillations via modulation of the detuning at the generalised Rabi frequency and its sub-harmonics. We have used this to elicit a further understanding behind the expectation for population dynamics on application of a broad-band noise-field to the detuning.

Results of numerical simulations were presented and discussed for the case of band-limited white noise and preliminary experimental data was presented containing both similarities and differences to the computed results and the possible causes briefly discussed. Nevertheless, it appears that at least qualitatively, the system does have a response to the applied noise-field which depends non-trivially on its characteristics.

While the experimental work conducted in this area is still in its infancy, it raises intriguing questions on the nature of our system's ability to respond to stochastic fields in a way that can be framed in the discussion of decoherence-assisted quantum systems.

CHAPTER 8

NEXT STEPS

With the work done in designing, building and characterising the experimental apparatus along with the preliminary results obtained, the experiment is at an exciting stage where interesting studies can be undertaken. Along with the additional tasks that have been outlined in the previous chapters, we have a host of possibilities for future improvements to make to the apparatus and investigations to be done. In this chapter, we will outline some of these ideas that have been explored but not yet experimentally implemented.

8.1 Breaking the Symmetry

As has been stated throughout the thesis, the use of angular momentum sub-levels within the same hyperfine manifold with a weak field means that there is a mapping of the system dynamics onto the Bloch sphere which itself has equivalent classical dynamics of rotations around particular axes. Whilst we believe this equivalence does not diminish the results achieved, the fact that our system Hamiltonian is constrained to this mapping space does limit some of the more interesting network topologies that do not have a mapping to such an easily solvable classical system. This is essential to the premise of a useful quantum simulator. The mapping is only valid if the Hamiltonian remains a representation of the SU(2) group. This could be broken by introducing couplings to non-adjacent states (see Section 8.2) or by removing the symmetry of the energy level differences between the bare states. In this section, we discuss options for this experimental possibility using higher order magnetic shifts or light shifts.

8.1.1 Higher Order Magnetic Shift

The identical energy level differences between the states are a consequence of the linear Zeeman regime. As discussed in Chapter 4 with respect to the state readout, when the magnitude of the energy level splitting

increases, higher order effects come into play. If we consider the quadratic order for the $F = 1$ ground state, the splitting is given using $\Delta E = aB + bB^2$ with values for a and b taken from Table 4.1. This leads to energy shifts of

$$E_{|+1\rangle} = 0.7\,\mathrm{MHz/G} - 0.216\,\mathrm{MHz/G^2} \tag{8.1}$$

$$E_{|0\rangle} = -0.288\,\mathrm{MHz/G^2} \tag{8.2}$$

$$E_{|-1\rangle} = -0.7\,\mathrm{MHz/G} - 0.216\,\mathrm{MHz/G^2}. \tag{8.3}$$

The quadratic shift causes the energy of the $|0\rangle$ state to move relative to the half way point between the $|+1\rangle$ and $|-1\rangle$ energies. This is shown in Figure 8.1. For the experiments conducted so far at a field of 1.6 G, this difference is given by 188 Hz, which is very small compared to the coupling strength of 40 kHz.

For this quadratic shift to noticeably affect the dynamics of the system, it should be close to the order of the coupling term. To be at the level of 40 kHz would require a field of 25 gauss. This is experimentally achievable with modest changes but a lot of experimental effort would be required to match the performance of the fast rise time of the bias field associated with the current set-up as well as a significant change in the RF coupling frequency as the energy level splittings would now be on the order of 8 MHz.

A major issue with increasing the magnitude of the bias field is the amplification of inhomogeneities in the detuning. As we would want to approximately preserve the detuning, which is defined by the difference between the field and the RF frequency, the field profile would have to be much more fractionally homogeneous across the area of the atoms to maintain the same absolute homogeneity in the detuning.

Whilst using this quadratic shift would allow for greater freedom of the Hamiltonian, including the ability to create unevenly spaced levels and even Vee systems (see Figure 8.1), the resonant Rabi frequencies of the $|-1\rangle \leftrightarrow |0\rangle$ and $|0\rangle \leftrightarrow |+1\rangle$ remain the same. However, this does not affect the ability to remove the mapping to the Bloch sphere, making the system much more interesting to investigate.

8.1.2 Decreasing Noise Correlations with Light Shift

As well as the symmetry of the energy level splittings provided by the linear Zeeman Hamiltonian, the response to the noise field also contains the same symmetry. Oscillations of the energy level of the $|-1\rangle$ state correlate exactly with the oscillations experienced by the $|+1\rangle$ levels. We have briefly considered the use of a laser beam

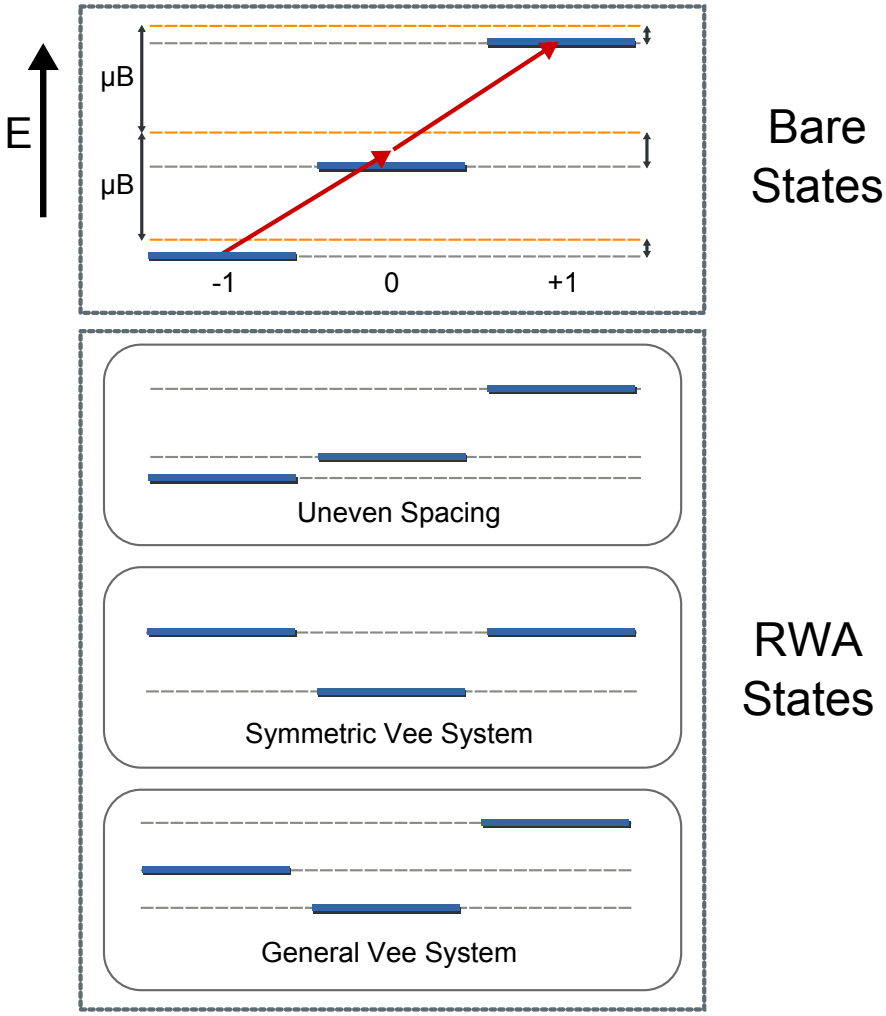

Figure 8.1: Possible energy level spacings in the presence of the quadratic Zeeman effect. The dashed orange lines show the linear Zeeman shifts are reduced in energy by the quadratic shift toward the dashed grey lines. As the $|0\rangle$ state is more strongly affected by this quadratic shift, the spacing between the $|-1\rangle$ and $|0\rangle$ states becomes smaller than the spacing between the $|0\rangle$ and $|+1\rangle$ states, breaking the symmetry. Moving to the rotating wave approximation with a single coupling field, it is possible to create more complex energy schemes.

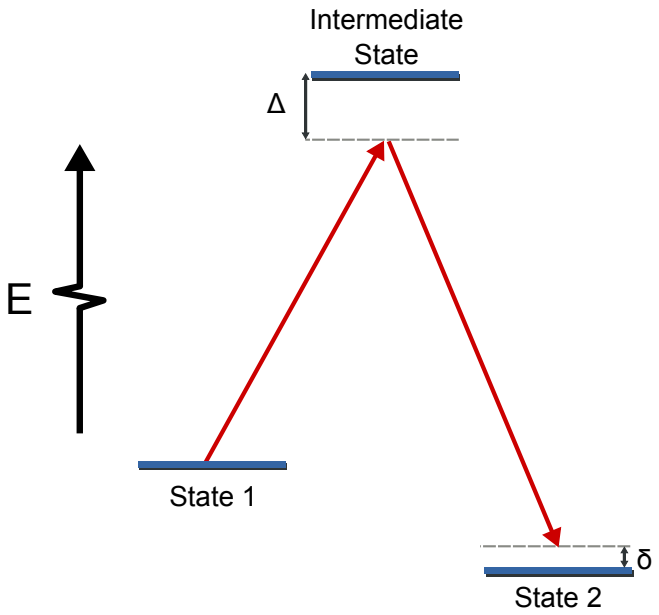

Figure 8.2: General schematic of a two-photon Raman transition. The two beams couple State 1 and State 2 via the Intermediate State. Transfer between the states involves the simultaneous absorption of a photon from the first beam and stimulated emission of a photon into the second beam. The *one photon detuning* Δ is normally set to be many linewidths to avoid significantly populating the intermediate state. The *two-photon detuning* δ acts much like the detuning in Rabi oscillations.

induced AC Stark shift to provide a time dependent energy shift on the atom by modulating the intensity of the coupling beam. Detailed calculations are not presented here but the large possible modulation bandwidth of the laser amplitude has a significant advantage over the technical difficulties inherent in producing a large bandwidth oscillating magnetic field.

8.2 Bridging the Gap using Raman Beams

Our currently implemented system Hamiltonian leads to the most simple non-trivial topology of a three site network; a linear chain or ladder. As discussed in Chapter 1, one of the interesting properties of transport assisted by decoherence is the role of interfering paths, which are not present in the current system.

We have considered and begun implementation of a system that allows the $|-1\rangle$ state to couple directly to the $|+1\rangle$ using Raman transitions. Briefly, Raman transitions use two-photons to couple two states via an intermediary state. A schematic example is shown in Figure 8.2.

Raman transitions are used extensively in the cold atoms community, especially for their uses in atom interferometers [97] for coherently coupling the two ground states via an excited state. In atom interferom-

eters, the coupling via the Raman scheme is set-up to impart the photon recoil velocities onto the atoms to create their interferometer sequence but we do not need, or use, this effect. It is typically undesirable to populate the intermediate state and as such, the one photon detuning Δ_{Raman} (completely unrelated to the Δ in our system Hamiltonian) is made to be large compared to the linewidth of the excited level. With this achieved, the excited state can be *adiabatically eliminated* [98] from the relevant population dynamics and the remainder is a two-level system with coupling strength given by

$$\Omega_{Raman} = \frac{\Omega_1 \Omega_2}{2\Delta_{Raman}} \tag{8.4}$$

where Ω_1 and Ω_2 are the Rabi frequencies associated with each of the two coupling beams which are each proportional to the square root of their corresponding beam intensities. The spontaneous scattering rate from the initial state to the intermediate state in the large-detuning limit is given by [36]

$$R_{sc} = \frac{\Gamma \Omega_1}{4\Delta_{Raman}^2} \tag{8.5}$$

where Γ is the linewidth of the transition. Scattering from the excited state is generally undesirable for most applications as it causes significant decoherence of the atom sample. For our system, this is also true as we are only interested in the Raman transitions as a coherent coupling method.

The Raman coupling is usually applied for a finite amount of time to complete some operation with a relevant timescale of $1/\Omega_{Raman}$. It can be seen that the ratio between the Raman Rabi frequency and the scattering rate is proportional to Δ_{Raman}. This means that a larger detuning minimises the number of photons spontaneously scattered (and therefore minimises the induced decoherence) during a set number of Raman Rabi oscillations. A corresponding increase in the beam intensities can recover the same coupling strength at a larger detuning.

This simple three-level model of the Raman transition is complicated by the real energy level structure of atoms. Consideration must be given to all of the possible intermediate states that coherently interfere to form a net effective coupling for the Raman transition. This is calculated via a simple sum of the Raman Rabi frequencies from each of the individual states, taking into account the relative signs:

$$\Omega_{Raman} = \sum_{F'} \frac{\Omega_{1,F'} \Omega_{2,F'}}{2\Delta_{F'}} \tag{8.6}$$

where we have added the subscript F' to label the contributions from each intermediate state. For coupling

states on different hyperfine levels with a large detuning, this sum of the coupling strengths ($\sum_{F'} \Omega_{1,F'}\Omega_{2,F'}$) is a finite constant and the situation reduces to a single effective intermediate state like the simple case discussed above. However, for our desired situation of coupling two levels with $\Delta m_F = 2$ in the same hyperfine level, the paths through the different intermediate levels destructively interfere, leaving the sum of the coupling strengths as zero. This means that for a detuning much larger than the difference in energies between the intermediate states, the Ω_{Raman} scales as $1/\Delta^2$, as all of the Δ_{Fi}'s are approximately the same. An example of a Raman scheme with multiple intermediate levels is shown for the Rubidium 87 D1 line in Figure 8.3.

The consequence of this for our desired coupling states is that the detuning cannot be made arbitrarily large to reduce the scattering effect. The one-photon detuning must be set so that the contribution for each of the excited states is weighted differently. This can be achieved with the beam close to the hyperfine structure or indeed in between the levels, where the relative signs of the detunings will be different. These transitions have been used on the D1 line before in Rubidium 87 BECs [99], although little discussion is made of the effect of spontaneous emission. We primarily consider the D1 line due to its large hyperfine splitting of 817 MHz, allowing us to sit between the levels to avoid the destructive path interference whilst maintaining a moderate detuning to avoid spontaneous scattering.

To calculate the effects of spontaneous emission, we find the Raman Rabi frequency and the spontaneous scattering rates from the ground states via all possible excited levels. The full energy level structure of the Rubidium 87 D1 line is shown in Figure 8.4. Particularly problematic is spontaneous scattering from the $|1, \pm 1\rangle \rightarrow |2, \pm 2\rangle$ states as they have a coupling strength six times larger than the $|1, \pm 1\rangle \rightarrow |2, 0\rangle$ levels used for the Raman coupling.

We calculated the Raman coupling strength and scattering rates of this process as functions of the one-photon detuning for the excited states of the D1 line. This allows us to be detuned significantly compared to the natural linewidth of the transitions while satisfying the condition that the different $\Delta_{F'}$'s are not the same. The primary results of these calculations are shown in Figure 8.5.

Note that the above formulations are only valid when the excited state can be adiabatically eliminated and as such, the results close to the levels (indicated by the dashed lines) are not physically valid. For the best compromise between a large Raman coupling frequency and low spontaneous scattering rate, the detuning should be set roughly halfway between the two levels. At the exact half-way point, the contribution from each F' level to the Raman Rabi frequency is equal in strength and in sign due to the opposite detunings cancelling the opposing Clebsch-Gordan coefficient products. The main source of incoherent scattering is coupling to the

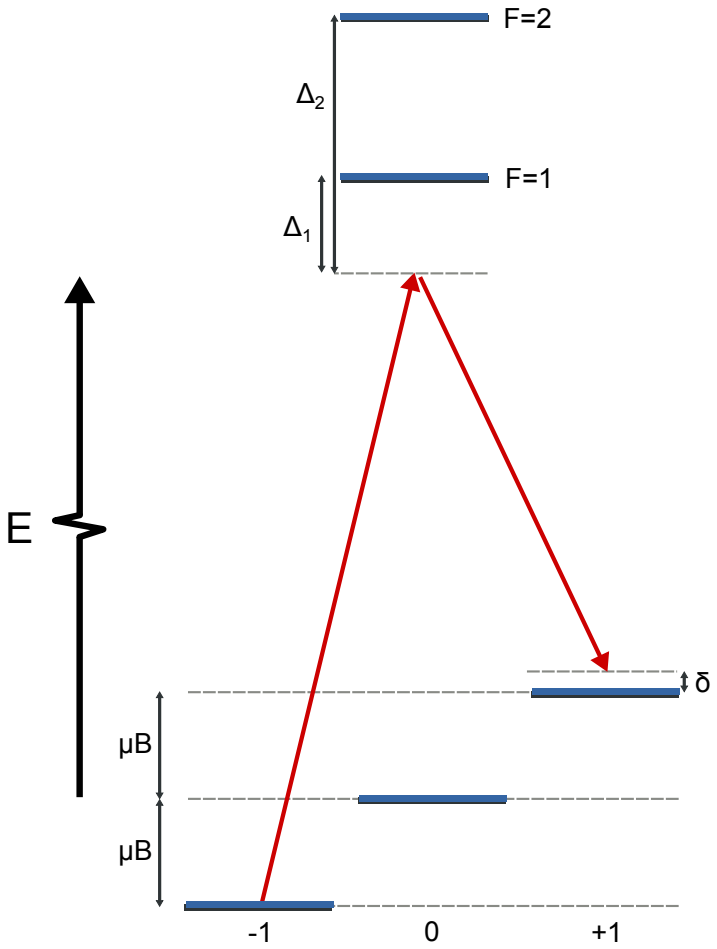

Figure 8.3: Effect of the Raman scheme for our ground state levels. The product of the Rabi frequencies with the F=1 intermediate state is equal and negative to those using the F=2 intermediate state. If the detuning from these states is large $\Delta_1 \approx \Delta_2$, the sum presented in Equation (8.6) is zero. To get a reasonable Raman coupling strength, the detuning must be small such that the detunings from the individual intermediate states are sufficiently different.

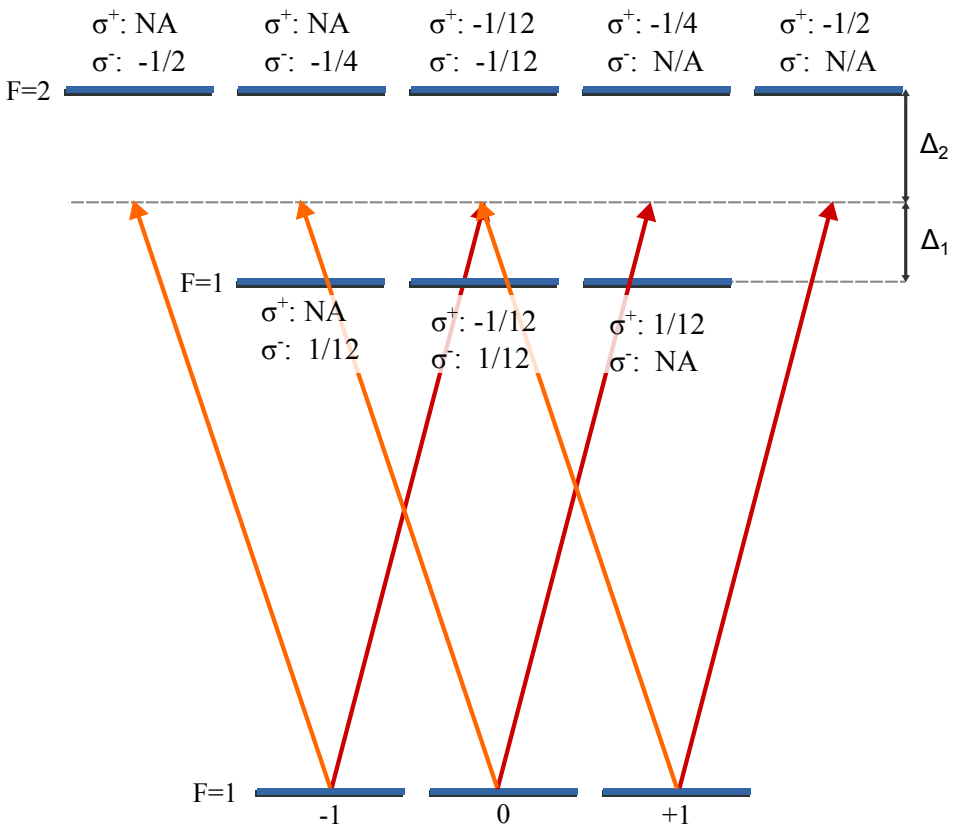

Figure 8.4: Diagram showing all of the states that are coupled by the presence of the σ^+ (red) and σ^- (orange) beams. Attention must be given to the resulting light shifts and scattering rates from all of these levels. The relative coupling strengths (the square of the Clebsch-Gordan coefficients) are shown for each beam and each excited state. The Zeeman effect and the difference in the beams' frequencies has been omitted for clarity. Note that the sign shown is actually the sign of the Clebsch-Gordan coefficients not their square. Representing them like this allows us to see the interfering paths without a multitude of square roots everywhere.

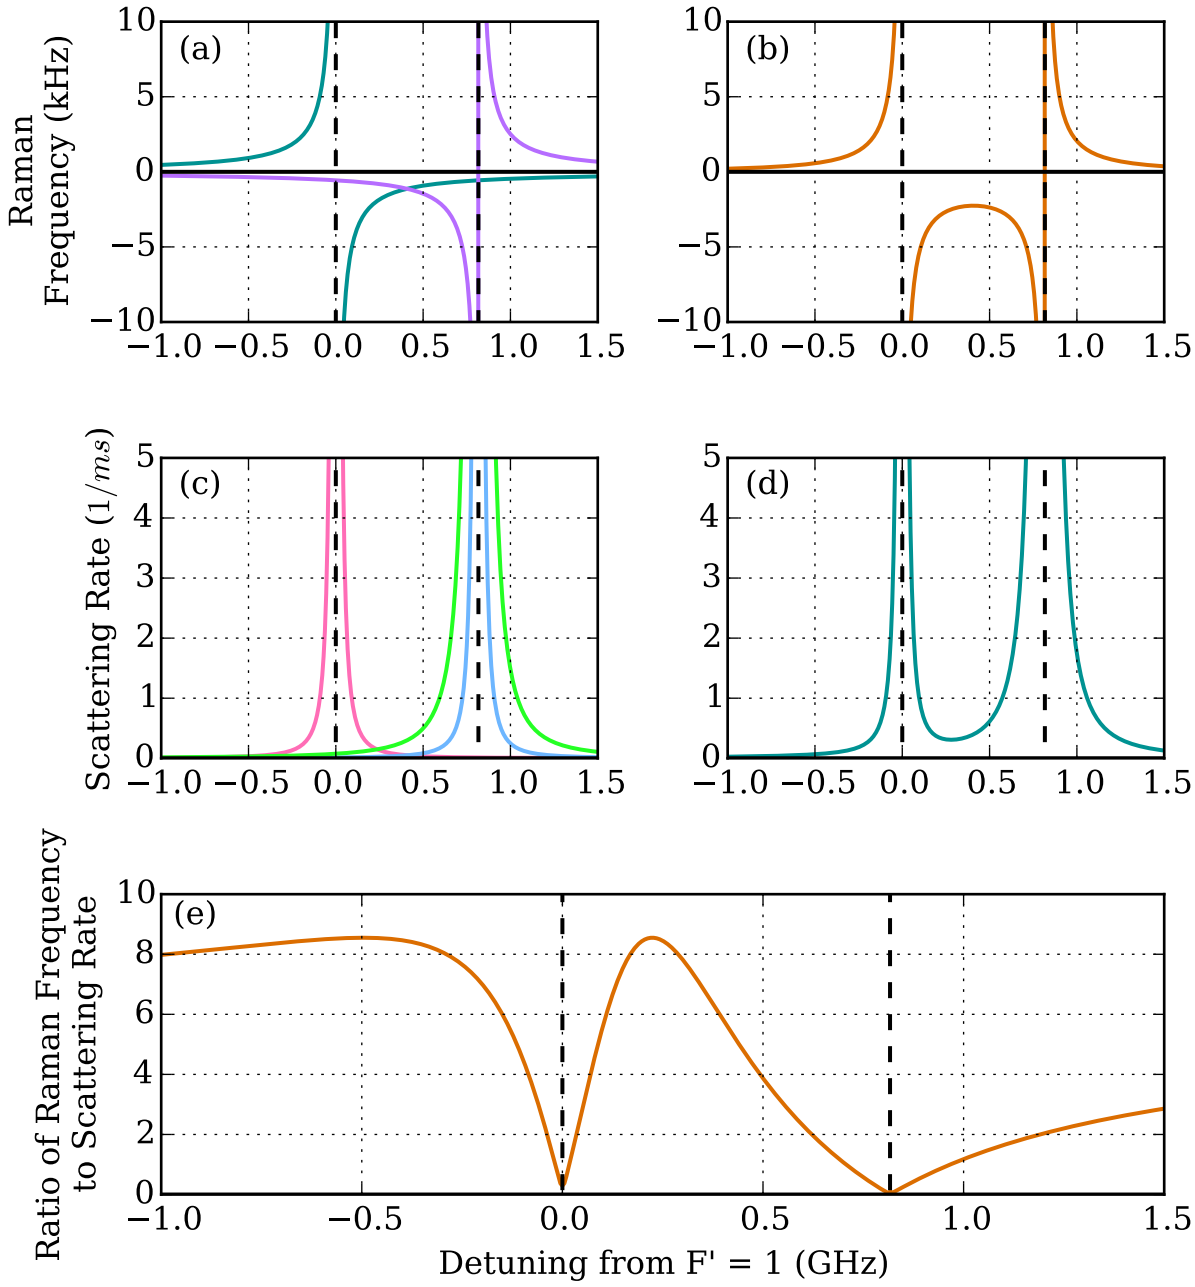

Figure 8.5: Calculated results for the expected Raman Rabi frequency and scattering rates as functions of detuning. These were calculated for equal beam intensities of $1\,\mathrm{mW/cm^2}$. (a) The contributions to the Raman transition from the F=1 excited state (turquoise) and the F=2 excited state (purple): note the sign means that they interfere destructively when the laser is not in-between them. (b) The sum of the two contributions to give the Raman Rabi frequency. (c) The scattering rates from the initial ground state $|1,-1\rangle$ to the $|1,0\rangle$ (pink), $|2,0\rangle$ (light blue) and $|2,-2\rangle$ (green). (d) The sum of all of the scattering rates gives a general rate for the decoherence caused by the scattering. (e) The ratio of the Raman Rabi frequency to the scattering rate. The presence of the strongly scattering $|2,-2\rangle$ state pushes the optimum position towards the F = 1 excited state.

stretched $|F' = 2, m'_F = \pm 2\rangle$ states from which the $|F = 1, m_F = \pm 1\rangle$ states strongly scatter as described by the larger Clebsch-Gordan coefficients. For this reason, the best ratio of Raman Rabi frequency to incoherent scattering rate is found with the one photon detuning set 224 MHz blue detuned from the $F' = 1$ state.

For the full connection of the network, it is desirable that the Raman coupling strength be comparable to the Rabi frequency of the RF couplings. For this to be achieved with a one photon detuning set to the optimum value, we require an intensity of approximately $14\,\mathrm{mW/cm^2}$ in each beam (as $\Omega_{1/2} \propto \sqrt{I_{1/2}}$, if the beams are of equal intensity, the Raman Rabi frequency scales linearly with intensity the beams).

The presence of the Raman beams also causes a change in the energy levels due to the AC Stark effect [51] (a.k.a. light shift). With equal power in the two Raman beams, the light shift on the $|1, \pm 1\rangle$ states is equal but a differential shift arises between these states and the $|1, 0\rangle$ state. The light shifts for the three states are shown in Figure 8.6 for both Raman beams with a beam intensity of $1\,\mathrm{mW/cm^2}$. We can see that the differential shift is on the order of 10 kHz for this small power. This relatively large shift for such a weak coupling will undoubtedly lead to complications for incorporating the RF field with the Raman beams. This actually has the effect of breaking the symmetry of the Hamiltonian similarly to the use of the quadratic Zeeman shift discussed in Section 8.1.1.

A potential issue with all of this discussion of the Raman scheme is the homogeneity of the beams. As the coupling strength is proportional to each beam intensity, the requirement for homogeneity necessitates a very flat beam intensity profile over the cloud. As beam profiles are normally Gaussian, flatness can be achieved by making the beam larger, however this requires a larger amount of total power to conserve the required intensity for the coupling strength. It is unclear for now whether the creation of a sufficiently homogeneous beam is possible without extreme experimental effort. Additionally, fringing effects from multiple sources cause periodic structures in the beam profiles (see for example the fringing of the imaging beam in Figure 3.5). To reduce fringing issues, the use of high quality anti-reflection coated optics may be used. Currently, the window used for the imaging beam and Raman beams is not AR coated, leading to a large amount of fringing.

8.2.1 Experimental Implementation

We have built a setup to investigate the use of Raman beams to provide a coupling strength comparable to the RF Rabi frequency of 40 kHz. The Raman beams are derived from a Toptica DL Pro ECDL laser operating at 795 nm. The beam is split into two beams and each passed through a separate double pass AOM setup, with nominal frequency shift of $+2 \times 80\,\mathrm{MHz}$. These two beams produce the orthogonal σ^+ and

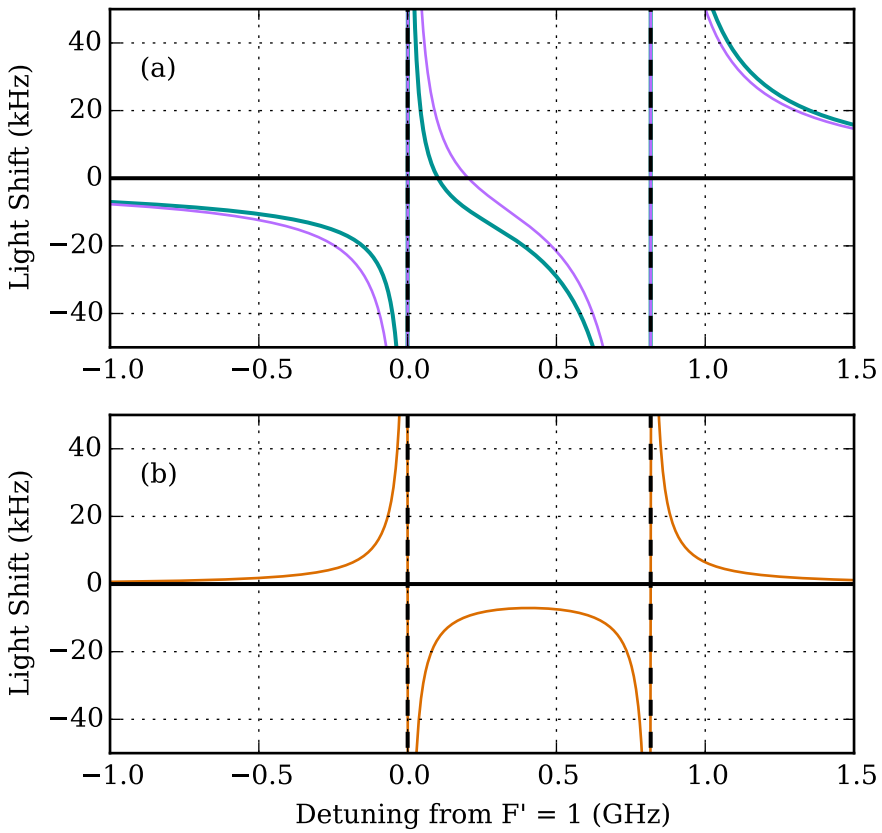

Figure 8.6: Calculated light shifts caused by the Raman beams, each with a $1\,\mathrm{mW/cm^2}$ intensity. (a) The light shifts on the $|1,\pm1\rangle$ states (turquoise) and the $|1,0\rangle$ state (lilac). (b) The differential light shift between these levels.

σ^- light by recombining them on a Polarising Beam Splitter (PBS) with orthogonal polarisations. They are then sent through an optical fibre to the vacuum chamber apparatus. A $\lambda/4$ waveplate circularly polarises them at the fibre launch. They are collimated to a beam diameter of approximately 2.5 cm with a single lens and then launched through the chamber, closely aligned to the imaging and pumping beams along the bias field direction. With respect to the quantisation axis of the experiment, the two beams have σ^+ and σ^- polarisations and their relative frequency can be shifted by the different AOM paths. The power in each polarisation is approximately 5 mW. As the intensity varies over the beam due to its Gaussian nature, we can only say the intensity is on the order of 1 mW/cm^2.

The AOM's serve to switch the beams on and off as well as allow computer control of their relative frequencies and intensities. The frequency of the laser is determined using saturated absorption spectroscopy of the D1 line in a Rubidium vapour cell but for this preliminary investigation, the laser is left free-running and is not locked.

To observe a signal, the main experimental procedure was followed but the experimental window did not involve any RF coupling. Instead, the Raman beams were pulsed for varying periods to see the Raman transition in a Rabi scheme. The bias field was applied at the same level that created resonant Rabi flopping for the RF field at 1.13 MHz, and the two-photon detuning δ was set to zero by setting the difference of the AOM frequency shifts equal to twice the resonant RF frequency. In these initial investigations, no consideration was taken for any differential light shift of the two levels as a result of any inequalities in the powers in each of the Raman beams.

Figure 8.7 shows the results from these initial investigations of Raman transitions in the Rabi scheme. There is clearly a movement to the opposing state and then the start of the return however this is evidently short lived compared to the RF Rabi oscillations. A decaying cosine was fitted which yielded an oscillation frequency of 4.7 ± 1.0 kHz and a 1/e decay time of 68 ± 1 µs. A key factor in the fast decay is likely the inhomogeneous nature of the beam as a result of fringing and uneven profile. Previous experiments done using this transition [99] did so on a BEC expanded to the size of 70 µm compared to our cloud with a size of ≈1 mm. To overcome these homogeneity issues, more experimental effort is required to create a flat beam front free from fringing. To ensure that the signal was not a result of optical pumping, the relative frequencies of the σ^+ and σ^- beams were switched and no population exchange was observed. If the population exchange were the result of a Raman transition, this exchange would make the two-photon detuning δ equal to 2.26 MHz, which is much larger than the coupling strength evidenced from the results. The 5 kHz oscillation frequency is roughly consistent with the developed theory. For a 1 mW intensity with

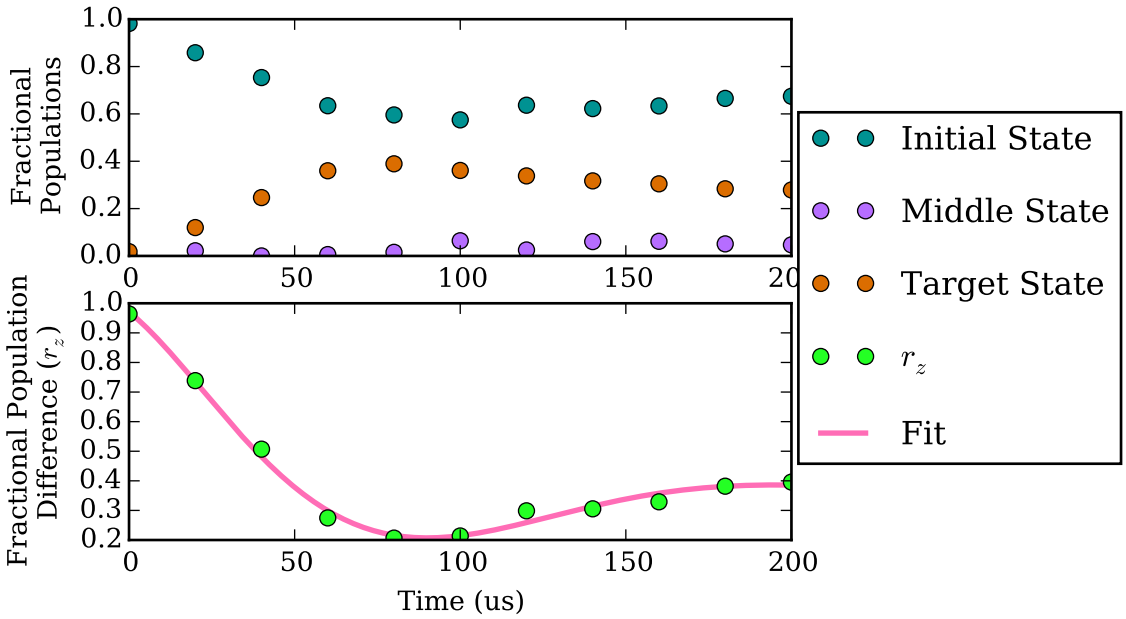

Figure 8.7: Preliminary experimental results from Raman beam coupling. It can be seen that population is exchanged directly from the initial state to the target state, bypassing the middle state (upper). The quantity r_z is plotted (lower) with a fit of a decaying cosine oscillation.

the detuning positioned in the middle of the F = 1 and F = 2 hyperfine levels, the expected Raman Rabi frequency is approximately 2.5 kHz. The intensity at the atoms may be larger than 1 mW as the beam is aligned with its centre (which has a higher intensity) through the atomic cloud and it is possible that the Raman laser had drifted in frequency closer to one of the hyperfine excited states, leading to an increased Raman Rabi frequency.

The measured Raman Rabi frequency of 5 kHz is a factor of 8 lower than the peak RF Rabi frequency[1]. To increase the Raman coupling strength, the intensity of the Raman beams can be increased. More precise optical alignment of the Raman setup could yield a factor of two more power but significantly more than that is unlikely with the current set-up. The preliminary nature of these results means that there was no attempt to characterise the light shift of the levels to determine their influence on RF Rabi oscillations.

In conclusion, these preliminary results are promising but much more work is required to be able to incorporate the Raman transitions into the general simulations.

[1] Although it is only a factor of two lower than that used during the noise measurements in section 7.3

8.3 Implementation of a Sink

A key component in many of the theoretical treatments of the quantum transport problem is a mechanism for the population of a state to be irreversibly moved ('sunk') to another site otherwise detached from the network (the 'sink') [13]. In what has been considered so far in our system, we look for a trend of moving towards the 'thermal' (with $T \to \infty$) distribution and quantify the transport as the speed at which this tendency to equilibration occurs. This is in comparison to sink based models that look at the population accumulated in the sink. For a fully connected network, it has been suggested that these sink based models lead to the surprising result that as $t \to \infty$, only a finite proportion of the system population ends up in the sink in the absence of decoherence [14].

To implement irreversible movement to another site in our model, we must consider a spontaneous decay mechanism. We have considered the use of a weak circularly polarised laser beam to excite atoms from the (for example) $|1,+1\rangle$ state to the $|2,0\rangle$ state, from where the atom will potentially decay to a state in the $F = 2$ ground state manifold. The population of the states in this manifold can then be measured independently of the remaining population in the $F = 1$ ground state. Figure 8.8 illustrates this process. The advantage of using σ light on a $|F = Fg, m_F = Fg\rangle \to |Fe, Fe\rangle$ transition is that decay from the excited state cannot reach the other m_F ground states used in the simulation without violating conservation of angular momentum selection rules.

An issue with such a scheme is the coupling of this 'Sink Beam' to other states, for example coupling the $|1,0\rangle \to |2,+1\rangle$ and $|1,-1\rangle \to |2,+2\rangle$ transition. This will lead to population sinking from these states as well, which is generally not present in the theoretical models and serves to introduce uncontrolled decoherence all over the network. These effects are mitigated 'for free' by the reduced intrinsic coupling strengths for these transitions compared to the stretched state transition of $|1,+1\rangle \to |2,+2\rangle$, but this can be further mitigated by judicious selection of the detuning of the sink beam. As the g factors (which determine the change in energy of the magnetic projection states in the linear Zeeman regime), are of different sign for the ground $F = 1$ manifold ($g_{Fg=1} = -0.7MHz/G$) and the excited $F = 2$ manifold ($g_{Fe=2} = +0.93MHz/G$), the applied magnetic field leads to shifted resonance transition frequencies for the σ^+ transitions discussed above. The formulas are given as

$$\frac{d\delta(|1,+1\rangle \to |2,+2\rangle)}{dB} = 2 \times g_{Fe=2} - 1 \times g_{Fg=1} = 2.56\,\mathrm{MHz/G} \tag{8.7}$$

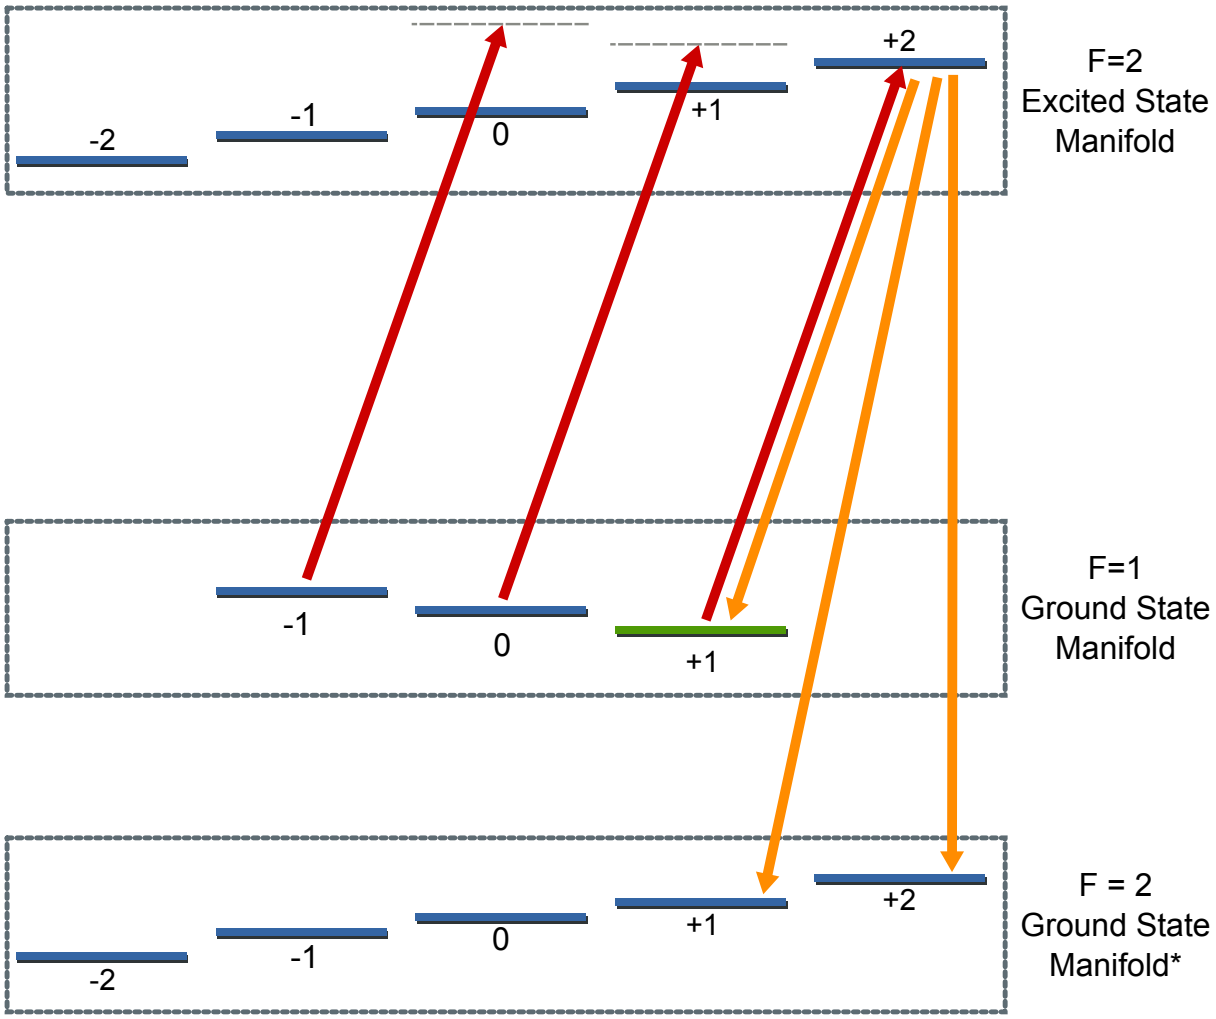

Figure 8.8: Possible implementation of a sink state in our system. A weak σ^+ beam couples the $|1,+1\rangle$ to the $|2,+2\rangle$ from which it can decay into the F = 2 ground states. The σ^+ light will also couple other transitions however the coupling strengths can be minimised with judicious choice of the relative detuning of the beams. $*$ Note that the F = 1 and F = 2 ground states are inverted with respect to energy in this diagram for clarity.

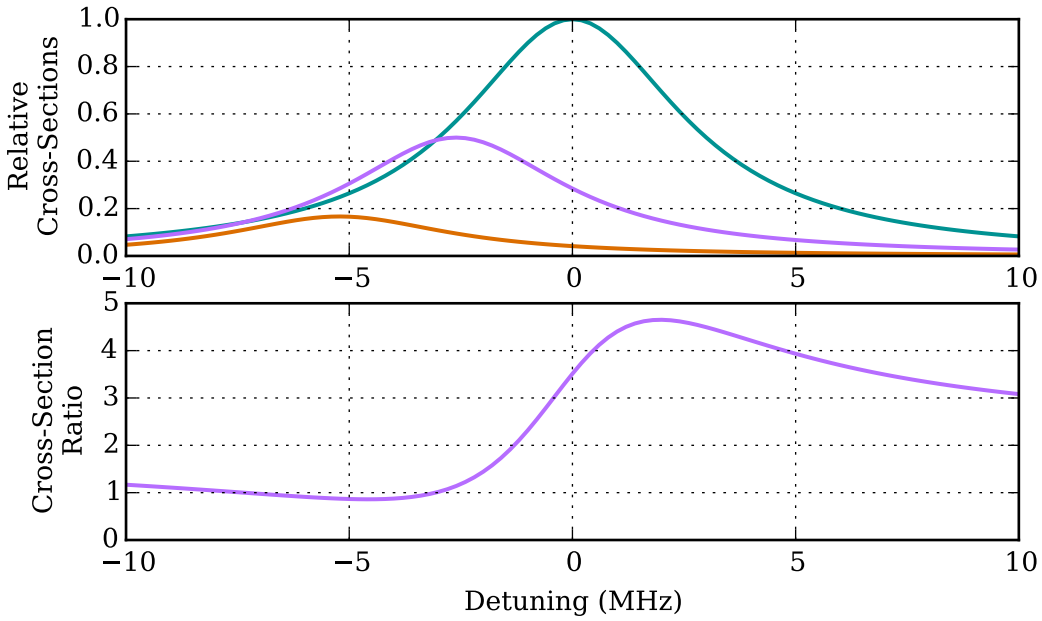

Figure 8.9: (Upper) Dependence of transition cross-sections on sink beam detuning for a fixed magnetic field of 1.6 G for the $|1,+1\rangle \rightarrow |2,+2\rangle$ transition (turquoise), $|1,0\rangle \rightarrow |2,+1\rangle$ transitions (lilac) and $|1,-1\rangle \rightarrow |2,0\rangle$ transition (orange). (Lower) The ratio of the cross sections for the $|1,+1\rangle \rightarrow |2,+2\rangle$ state over the $|1,0\rangle \rightarrow |2,+1\rangle$ state. For an effective sink that does not disrupt the rest of the system, this ratio should be large.

$$\frac{d\delta(|1,0\rangle \rightarrow |2,+1\rangle)}{dB} = 1 \times g_{Fe=2} - 0 \times g_{Fg=1} = 0.93\,\mathrm{MHz/G} \tag{8.8}$$

$$\frac{d\delta(|1,-1\rangle \rightarrow |2,0\rangle)}{dB} = 0 \times g_{Fe=2} + 1 \times g_{Fg=1} = -0.70\,\mathrm{MHz/G}. \tag{8.9}$$

This means the differences in the resonant frequencies of these transitions can be made on the order of the natural linewidth $\Gamma = 6\,\mathrm{MHz}$ with a field of a few gauss. Combined with the different resonant coupling strengths, we can potentially reach a regime where the transitions for the desired sink transition $|1,+1\rangle \rightarrow |2,+2\rangle$ is coupled much more strongly than the other transitions.

For example, if we use the bias field value of 1.6 G used throughout much of the work, the transitions will be separated as shown in Figure 8.9. The maximum ratio of ≈ 4.5 can be bettered with a higher field. Figure 8.10 shows this ratio plot for a range of fields that are potentially experimentally available with modest equipment improvements.

We can see that it is therefore possible to implement a mechanism that irreversibly sinks population from the 'target state' whilst leaving the rest of the system mostly unperturbed. Implementation would require the

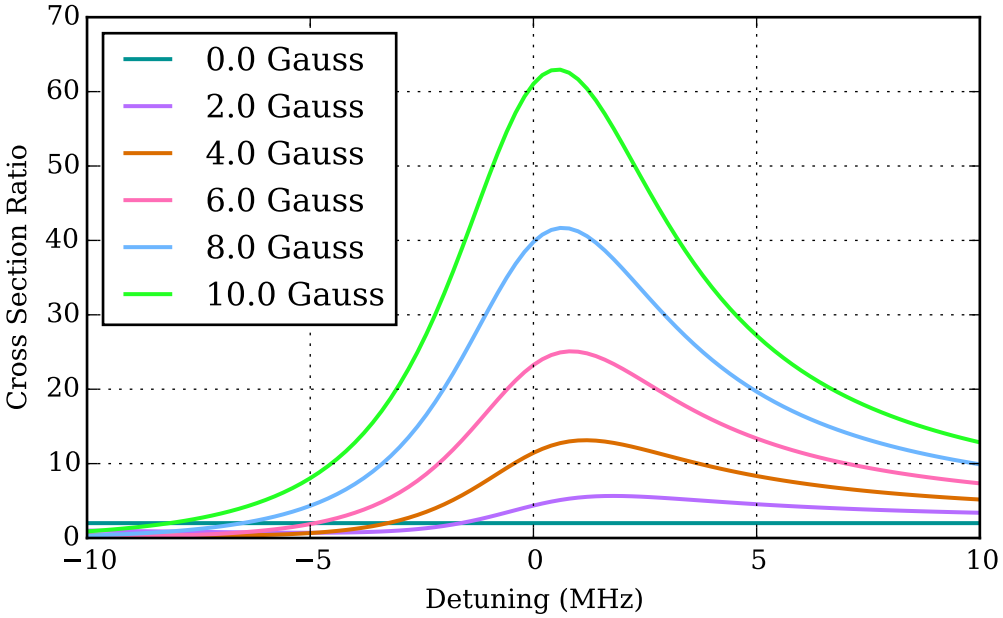

Figure 8.10: Calculated coupling strength ratios between the desired $|1,+1\rangle \rightarrow |2,+2\rangle$ state and the undesired $|1,0\rangle \rightarrow |2,+1\rangle$ state as a function of detuning for a range of magnetic fields.

ability to create larger bias fields on the order of five gauss, which would also require a higher frequency RF coupling field, however the presence of a sink allows us to compare our results with some of the approaches taken by published theoretical material.

8.4 More Noise Spectra

In Section 7.3, we looked at the effects of the application of noise and discussed the way in which noise waveforms can be made to adhere to some given arbitrary frequency spectrum. However, we have thus far only considered the relatively simple case of band-limited white noise. Other spectral shapes of noise could be investigated such as power law scalings, where the power spectral density scales as $J(f) \propto f^n$. As discussed in section 7.3, special cases of this are $n = -1$, which has the general name $1/f$ noise, $n = 1$, so called Ohmic noise, $n > 1$, super-Ohmic noise and $0 < n < 1$, sub-Ohmic noise. This could be of interest to probe whether the different frequency scales of the noise field affect different aspects of the system's dynamics.

These effects could be investigated without any additional hardware added to the experiment, and is the likely immediate source of further investigation.

CHAPTER 9

OTHER WORK - COUPLING ATOMS TO PHOTONIC CRYSTAL WAVEGUIDES

The initial intent of the cold atoms apparatus discussed in this thesis was for an experiment into coupling cold atoms to the evanescent fields of Photonic Crystal Waveguide (PCW) structures. However due to other commitments from collaborating parties, the PCW fabrication process was delayed indefinitely and the project focus switched. A new collaboration since that point resulted in a publication [100]. This chapter is therefore included as a stand-alone chapter very briefly discussing this work.

PCWs are refractive media that have been structured by the introduction of spatially periodic modulation of the refractive index. The well known solutions to the Schrödinger equation for a periodic potential are the Bloch states, which are well known to have energy gaps where no propagating states exist. As the propagation modes of light through a medium are also solutions to a wave equation, an analogue exists when the refractive index of a medium contains a periodic structure; the result is a medium in which light with a frequency within a frequency band gap cannot propagate in the medium. This is referred to as a photonic crystal. For an excellent overview of photonic crystals, see Molding the Flow of Light [101].

Typical geometries for planar structures involve a 2D array of 'holes' etched into the refractive medium. Of great interest is the ability to locally remove a line of these holes to create a path for light in the bandgap to propagate. This is referred to as a photonic crystal waveguide. Due to the strong dispersive effect of the photonic crystal that bounds the light in this spatial gap, the light mode can be engineered to have a very low group velocity. In addition, the light carries a significant portion of its energy outside of the waveguide medium as an evanescent field.

If an atom is brought close to the surface of such a waveguide, such that it is in the evanescent field, the atom can interact with the guided light. Due to the spatial compression of the light and its low group

velocity due to the guiding effect of the photonic crystal, the coupling cross section can be made very large for a small amount of light. This is of great interest for using atoms to mediate interactions between small numbers of photons [102].

Ongoing experiments of Lukin [103] and Kimble [104], have successfully coupled atoms to the guided light in waveguide structures that have relatively weak periodicity. A collaboration with the Light in Complex Nanostructures group at the LP2N-IOGS laboratory in Bordeaux yielded a publication [100] on a novel geometry that allows light to be guided between a photonic crystal and the edge of a waveguide, dubbed the W1 cut. This geometry allows for a slow light mode guided between the photonic crystal array and the edge of the waveguide with a large spatial extension of the field into the space outside of the waveguide (several hundred nanometres). This large extension into free space makes coupling to the atoms much easier as they can remain far from the waveguide surface, where van-der-Waals interactions and the Casimir-Polder force are significant. An additional advantage of this new geometry is robustness of the group velocity of the slow-light mode against inevitable minor manufacturing defects of the crystal structure.

While the numerical calculations and modelling were conducted by the group in Bordeaux, the initial geometric concept and much of the atomic motivation was provided by Birmingham.

The initial planning stages (laser system and vacuum chamber design) of the experiment discussed in this thesis were based around the ability to couple cold atoms to the evanescent field of a waveguide. The waveguide was to be mounted inside the chamber with coupling access through the large side windows. The atoms were to be transported from the MOT to the PCW via a moving molasses atomic fountain. This basically involves the normal optical molasses polarisation gradient cooling mechanism but implemented with uneven frequencies in opposing beams in a single direction. This means the atoms are cooled toward zero velocity in this moving frame, leading to a launch of the atoms. By tuning the launch velocity such that the turning point of the atom cloud trajectory is in the vicinity of the waveguide, some atoms will spend sufficient time in the evanescent field to interact.

CHAPTER 10

CONCLUSIONS

In this thesis, we have introduced the topic of quantum processes that are assisted by some form of decoherence. We have specifically looked at the case of transport through quantum networks, with initial interest in the field coming from the recent discoveries of quantum coherence present in biological structures at physiological temperatures.

We have presented a mapping of a single excitation quantum network onto the energy level structure of an atom. Practical applications of this mapping allow for the simulation and testing of decoherence-assisted processes on samples of laser cooled atoms capable of reaching the tangible quantum regime of energy level couplings. Decoherence was then added via the application of a time dependent energy term to the Hamiltonian governed by a stochastic process. The total response of the system is measured as the average over multiple realisations of the stochastic process. This form of decoherence is limited to investigations of classical noise, which is short of the expected quantum noise that is thought to govern many of the biological processes.

Much of the thesis has detailed the technical work conducted in the construction of the apparatus with a mixture of straightforward and more novel techniques. We have presented results of parametric oscillations of state populations as a result of coherently modulating the energy level structure of the Rubidium atoms including sub-resonance behaviour. The results of simple numerical simulations were compared against preliminary experimental data, testing the modulation of the energy level structure by a broadband field resulting from a stochastic process and in both cases a 'noise assisted' feature was identified and discussed.

The apparatus now has the ability to more extensively investigate the regime of band-limited white noise but also has many additional regimes that can be probed without the addition of more hardware in the form of different spectral densities of the noise field. There are also many additional extensions of the model that

can be made to break the symmetry of the applied system Hamiltonians which have been discussed.

LIST OF REFERENCES

[1] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien. Quantum computers. *Nature*, 464(7285):45–53, 2010.

[2] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard. Optics and interferometry with atoms and molecules. *Reviews of Modern Physics*, 81(3):1051–1129, 2009.

[3] R. L. Fagaly. Superconducting quantum interference device instruments and applications. *Review of Scientific Instruments*, 77(10):101101, 2006.

[4] W. H. Zurek. Decoherence and the transition from quantum to classical - revisisted. *arXiv preprint quant-ph/0306072*, 2003.

[5] G. S. Engel, T. R. Calhoun, E. L. Read, T.K. Ahn, T. Mančal, Y.C. Cheng, R. E. Blankenship, and G. R. Fleming. Evidence for wavelike energy transfer through quantum coherence in photosynthetic systems. *Nature*, 446(7137):782–786, 2007.

[6] R. P. Feynman. Simulating physics with computers. *International journal of theoretical physics*, 21(6):467–488, 1982.

[7] N. Lambert, Y.N. Chen, Y.C. Cheng, C.M. Li, G.Y. Chen, and F. Nori. Quantum biology. *Nature Physics*, 9(1):10–18, 2013.

[8] A. W. Chin, S. F. Huelga, and M. B. Plenio. Coherence and decoherence in biological systems: principles of noise-assisted transport and the origin of long-lived coherences. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1972):3638–3657, 2012.

[9] G.Y. Chen, N. Lambert, C.M. Li, Y.N. Chen, and F. Nori. Rerouting excitation transfers in the Fenna-Matthews-Olson complex. *Physical Review E*, 88(3), 2013.

[10] J. Adolphs and T. Renger. How proteins trigger excitation energy transfer in the FMO complex of green sulfur bacteria. *Biophysical Journal*, 91(8):2778–2797, 2006.

[11] Z.X. Man, N.B. An, and Y.J. Xia. Non-Markovian dynamics of a two-level system in the presence of hierarchical environments. *Optics Express*, 23(5):5763, 2015.

[12] H.S. Zeng, N. Tang, Y.P. Zheng, and T.T. Xu. Non-Markovian dynamics for an open two-level system without rotating wave approximation: indivisibility versus backflow of information. *The European Physical Journal D*, 66(10), 2012.

[13] M. B. Plenio and S. F. Huelga. Dephasing-assisted transport: quantum networks and biomolecules. *New Journal of Physics*, 10(11):113019, 2008.

[14] F. Caruso, A. W. Chin, A. Datta, S. F. Huelga, and M. B. Plenio. Highly efficient energy excitation transfer in light-harvesting complexes: The fundamental role of noise-assisted transport. *The Journal of Chemical Physics*, 131(10):105106, 2009.

[15] D. P. S. McCutcheon and A. Nazir. Consistent treatment of coherent and incoherent energy transfer dynamics using a variational master equation. *The Journal of Chemical Physics*, 135(11):114501, 2011.

[16] H.P. Breuer, E.M. Laine, J. Piilo, and B. Vacchini. Non-Markovian dynamics in open quantum systems. *Reviews of Modern Physics*, 88(2), 2016.

[17] C. Karlewski and M. Marthaler. Time-local master equation connecting the Born and Markov approximations. *Physical Review B*, 90(10), 2014.

[18] W.M. Zhang, P.Y. Lo, H.N. Xiong, M. W.Y. Tu, and F. Nori. General non-Markovian dynamics of open quantum systems. *Physical Review Letters*, 109(17), 2012.

[19] J. Iles-Smith, A. G. Dijkstra, N. Lambert, and A. Nazir. Energy transfer in structured and unstructured environments: Master equations beyond the Born-markov approximations. *The Journal of Chemical Physics*, 144(4):044110, 2016.

[20] O.-P. Saira, V. Bergholm, T. Ojanen, and M. Möttönen. Equivalent qubit dynamics under classical and quantum noise. *Physical Review A*, 75(1), 2007.

[21] R. de Sousa, K. B. Whaley, F. K. Wilhelm, and J. von Delft. Ohmic and step noise from a single trapping center hybridized with a fermi sea. *Physical Review Letters*, 95(24), 2005.

[22] S. Lloyd and J.J. E. Slotine. Quantum feedback with weak measurements. *Physical Review A*, 62(1):012307, 2000.

[23] G. G. Gillett, R. B. Dalton, B. P. Lanyon, M. P. Almeida, M. Barbieri, G. J. Pryde, J. L. OBrien, K. J. Resch, S. D. Bartlett, and A. G. White. Experimental feedback control of quantum systems using weak measurements. *Physical Review Letters*, 104(8), 2010.

[24] I. Bloch, J. Dalibard, and S. Nascimbne. Quantum simulations with ultracold quantum gases. *Nature Physics*, 8(4):267–276, 2012.

[25] S. Burger, K. Bongs, S. Dettmer, W. Ertmer, K. Sengstock, A. Sanpera, G. V. Shlyapnikov, and M. Lewenstein. Dark solitons in Bose-Einstein condensates. *Physical Review Letters*, 83(25):5198, 1999.

[26] L. Sanchez-Palencia, D. Clément, P. Lugan, P. Bouyer, G. V. Shlyapnikov, and A. Aspect. Anderson localization of expanding Bose-Einstein condensates in random potentials. *Physical Review Letters*, 98(21), 2007.

[27] G. Roati, C. D'Errico, L. Fallani, M. Fattori, C. Fort, M. Zaccanti, G. Modugno, M. Modugno, and M. Inguscio. Anderson localization of a non-interacting Bose-Einstein condensate. *Nature*, 453(7197):895–898, 2008.

[28] J. Billy, V. Josse, Z. Zuo, A. Bernard, B. Hambrecht, P. Lugan, D. Clment, L. Sanchez-Palencia, P. Bouyer, and A. Aspect. Direct observation of Anderson localization of matter waves in a controlled disorder. *Nature*, 453(7197):891–894, 2008.

[29] Y.I. Shin. *Experiments with Bose-Einstein condensates in a double-well potential*. PhD thesis, Department of Physics - Massachusetts Institute of Technology, 2005. Available from `http://www.physics.utoronto.ca/~jht/pub/theses/Shin2006.pdf`.

[30] M. Albiez. *Observation of nonlinear tunneling of a Bose-Einstein condensate in a single Josephson junction*. PhD thesis, Combined Faculties for the Natural Sciences and for Mathematics of the Ruperto-Carola University of Heidelberg, 1993. Available from `http://archiv.ub.uni-heidelberg.de/volltextserver/5869/`.

[31] Y. Japha, S. Zhou, M. Keil, R. Folman, C. Henkel, and A. Vardi. Suppression and enhancement of decoherence in an atomic Josephson junction. *New Journal of Physics*, 18(5):055008, 2016.

[32] R. A. Miller, A. D. Presley, and M. B. Francis. Self-assembling light-harvesting systems from synthetically modified tobacco mosaic virus coat proteins. *Journal of the American Chemical Society*, 129(11):3104–3109, 2007.

[33] B. W Shore. *The Theory of Coherent Atomic Excitation*, volume 1. Wiley Interscience, 1st edition, 1990.

[34] B. W. Shore. *The Theory of Coherent Atomic Excitation*, volume 2. Wiley Interscience, 1st edition, 1990.

[35] C. Foot. *Atomic Physics*. Oxford University Press, 1st edition, 2005.

[36] D. A. Steck. Rubidium 87 D Line Data, 2001. Available from `http://steck.us/alkalidata/rubidium87numbers.1.6.pdf`.

[37] H.P. Breuer and F. Petruccione. *The Theory of Open Quantum Systems*. Oxford University Press, 1st edition, 2002.

[38] S. Haywood. *Symmetries and Conservation Laws in Particle Physics*. Imperial College Press, 1st edition, 2011.

[39] E. Paladino, Y.M. Galperin, G. Falci, and B.L. Altshuler. 1 / f noise: Implications for solid-state quantum information. *Reviews of Modern Physics*, 86(2):361–418, 2014.

[40] W. D. Phillips. Nobel lecture: Laser cooling and trapping of neutral atoms. *Reviews of Modern Physics*, 70(3):721, 1998.

[41] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer-Verlag, 1st edition, 1999.

[42] C. E. Wieman, D. E. Pritchard, and D. J. Wineland. Atom cooling, trapping, and quantum manipulation. *Reviews of Modern Physics*, 71(2):S253, 1999.

[43] A. S. Arnold, J. S. Wilson, and M. G. Boshier. A simple extended-cavity diode laser. *Review of Scientific Instruments*, 69(3):1236, 1998.

[44] 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(12):4477, 2001.

[45] 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(24):5141, 2002.

[46] U. Schünemann, H. Engler, R. Grimm, M. Weidemüller, and M. Zielonkowski. Simple scheme for tunable frequency offset locking of two lasers. *Review of Scientific Instruments*, 70(1):242, 1999.

[47] Y. Castin, H. Wallis, and J. Dalibard. Limit of doppler cooling. *Journal of the Optical Society of America B*, 6(11):2046, 1989.

[48] P. J. Ungar, D. S. Weiss, E. Riis, and S. Chu. Optical molasses and multilevel atoms: theory. *Journal of the Optical Society of America B*, 6(11):2058–2071, 1989.

[49] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the doppler limit by polarization gradients: simple theoretical model. *Journal of the Optical Society of America B*, 6(11):2058, 1989.

[50] W. Ketterle and N. J. Van Druten. Evaporative cooling of trapped atoms. *Advances In Atomic, Molecular, and Optical Physics*, 37:181 – 236, 1996.

[51] R. Grimm, M. Weidermller, and Y. B. Ovchinnikov. Optical dipole traps for neutral atoms. *Advances in Atomic, Molecular and Optical Physics*, 42:95–170, 2000.

[52] M. Kasevich and S. Chu. Laser cooling below a photon recoil with three-level atoms. *Physical Review Letters*, 69(12):1741, 1992.

[53] A. Aspect, E. Arimondo, R. Kaiser, N. Vansteenkiste, and C. Cohen-Tannoudji. Laser cooling below the one-photon recoil energy by velocity-selective coherent population trapping. *Physical Review Letters*, 61(7):826, 1988.

[54] H. J. Davies. *An optically guided atomic fountain*. PhD thesis, Durham University, 1999. Available at Durham E-Theses Online: `http://etheses.dur.ac.uk/4580/`.

[55] A. G. Sinclair, E. Riis, and M. J. Snadden. Improved trapping in a vapor-cell magneto-optical trap with multiple laser frequencies. *Journal of the Optical Society of America B*, 11(12):2333–2339, 1994.

[56] J.P. Yuan, Z.H. Ji, Y.T. Zhao, X.F. Chang, L.T. Xiao, and S.T. Jia. Simple, reliable, and nondestructive method for the measurement of vacuum pressure without specialized equipment. *Applied Optics*, 52(25):6195, 2013.

[57] T. M. Brzozowski, M. Maczynska, M. Zawada, J. Zachorowski, and W. Gawlik. Time-of-flight measurement of the temperature of cold atoms for short trap-probe beam distances. *Journal of Optics B: Quantum and Semiclassical Optics*, 4(1):62, 2002.

[58] W. Gerlach and O. Stern. über die richtungsquantelung im magnetfeld. *Annalen der Physik*, 379(16):673–699, 1924.

[59] S. Machluf, Y. Japha, and R. Folman. Coherent SternGerlach momentum splitting on an atom chip. *Nature Communications*, 4, 2013.

[60] C.A. Käfer. *Stern-Gerlach experiments with Bose-Einstein condensates and the introduction of a new thermometry method in an optical dipole trap.* PhD thesis, Albert-Ludwigs Universitat Freiburg, 2010. Available from `https://www.freidok.uni-freiburg.de/fedora/objects/freidok:7695/datastreams/FILE1/content`.

[61] V. G. Minogin, J. A. Richmond, and G. I. Opat. Time-orbiting-potential quadrupole magnetic trap for cold atoms. *Physical Review A*, 58(4):3138, 1998.

[62] D. M. Stamper-Kurn, M. R. Andrews, A. P. Chikkatur, Stenger Inouye, H.-J. Miesner, J. Stenger, and W. Ketterle. Optical confinement of a Bose-Einstein condensate. *Physical Review Letters*, 80(10):2027, 1998.

[63] C. L. G. Alzar, P. G. Petrov, D. Oblak, J. H. Mueller, and E. S. Polzik. Compensation of eddy-current-induced magnetic field transients in a MOT. *arXiv preprint physics/0701251*, 2007.

[64] A. Kaspers. Blob detection. Master's thesis, UMC Utrecht, 2011. Available from: `http://dspace.library.uu.nl/handle/1874/204781`.

[65] J. T. L. Thong and F. Li. Eddy current compensation for magnetic electron lenses. *Measurement Science and Technology*, 7(11):1583, 1996.

[66] V. J. Schmithorst and B. J. Dardzinski. Automatic gradient preemphasis adjustment: A 15-minute journey to improved diffusion-weighted echo-planar imaging. *Magnetic Resonance in Medicine*, 47(1):208–212, 2002.

[67] Apex Microtechnologies. Voltage to Current Conversion. Technical report, 2013. Available from: `www.apexanalog.com/resources/appnotes/an13u.pdf`.

[68] F. Nibler. *High Frequency Circuit Engineering*. The Institution of Electrical Engineers, 1st edition, 1996.

[69] D. F. Doty, G. Entzminger, J. Kulkarni, K. Pamarthy, and J. P. Staab. Radio frequency coil technology for small-animal MRI. *NMR in Biomedicine*, 20(3):304–325, 2007.

[70] L. D Landau. Zur theorie der energieubertragung. II. *Physics of the Soviet Union*, 2(2):46–51, 1932.

[71] C. Zener. Non-adiabatic crossing of energy levels. *Proceedings of the Royal Society of London. Series A*, 137(833):696, 1932.

[72] J. R. Rubbmark, M. M. Kash, M. G. Littman, and D. Kleppner. Dynamical effects at avoided level crossings: A study of the Landau-Zener effect using rydberg atoms. *Physical Review A*, 23(6):3107, 1981.

[73] C. E. Carroll and F. T. Hioe. Transition probabilities for the three-level Landau-Zener model. *Journal of Physics A: Mathematical and General*, 19(11):2061, 1986.

[74] C. W. S. Conover. Effects of pulse shape on strongly driven two-level systems. *Physical Review A*, 84(6), 2011.

[75] J. Keeler. *Understanding NMR Spectroscopy*. John Wiley & Sons, 2nd edition, 2011.

[76] N. F. Ramsey. A molecular beam resonance method with separated oscillating fields. *Physical Review*, 78(6):695, 1950.

[77] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P.O. Schmidt. Optical atomic clocks. *Reviews of Modern Physics*, 87(2):637–701, 2015.

[78] D. Seidel and J. G. Muga. Ramsey interferometry with guided ultracold atoms. *The European Physical Journal D*, 41(1):71–75, 2007.

[79] L. D. Landau. *Mechanics*. Oxford: Pergamon, 2nd edition, 1969.

[80] William B. Case. The pumping of a swing from the standing position. *American Journal of Physics*, 64(3):215–220, 1996.

[81] Alejandro Jenkins. Self-oscillation. *Physics Reports*, 525(2):167–222, 2013.

[82] J. H. Shirley. *Interaction of a quantum system with a strong oscillating field.* PhD thesis, California Institute of Technology, 1963. Available from `http://thesis.library.caltech.edu/1805/1/Shirley_jh_1963.pdf`.

[83] V. Novičenko, E. Anisimovas, and G. Juzeliūnas. Floquet analysis of a quantum system with modulated periodic driving. *Physical Review A*, 95(2), 2017.

[84] J. B. Hagen. *Radio-Frequency Electronics.* Cambridge University Press, 1st edition, 2009.

[85] J. H. Shirley. Solution of the Schrödinger equation with a hamiltonian periodic in time. *Physical Review*, 138(4):B979, 1965.

[86] F. Bloch and A. Siegert. Magnetic resonance for nonrotating fields. *Physical Review*, 57(6):522, 1940.

[87] P. G. Hoel, S. C. Port, and C. J. Stone. *Introduction to Stochastic Processes.* Houghton Mifflin Company, 1st edition, 1972.

[88] H. M. Taylor and S. Karlin. *An Introduction to Stochastic Modeling.* Academic Press, Inc., 1st edition, 1984.

[89] G. Grimmett and D. Stirzaker. *Probability and Random Processes.* Oxford University Press, 3rd edition, 2001.

[90] Y. M. Galperin, B. L. Altshuler, J. Bergli, D. Shantsev, and V. Vinokur. Non-Gaussian dephasing in flux qubits due to 1 / f noise. *Physical Review B*, 76(6), 2007.

[91] J. Eroms, L. C. van Schaarenburg, E. F. C. Driessen, J. H. Plantenberg, C. M. Huizinga, R. N. Schouten, A. H. Verbruggen, C. J. P. M. Harmans, and J. E. Mooij. Low-frequency noise in Josephson junctions for superconducting qubits. *Applied Physics Letters*, 89(12):122516, 2006.

[92] J.H. An and W.M. Zhang. Non-Markovian entanglement dynamics of noisy continuous-variable quantum channels. *Physical Review A*, 76(4), 2007.

[93] W. Guo, B. Cretu, J.-M. Routoure, R. Carin, E. Simoen, and C. Claeys. Temperature impact on the Lorentzian noise induced by electron valence-band tunneling in partially depleted SOI p-MOSFETs. *Solid-State Electronics*, 51(9):1180–1184, 2007.

[94] N. J. Kasdin. Discrete simulation of colored noise and stochastic processes and 1/f power law noise generation. *Proceedings of the IEEE*, 83(5):802–827, 1995.

[95] R. F. Fox, I. R. Gatland, R. Roy, and G. Vemuri. Fast, accurate algorithm for numerical simulation of exponentially correlated colored noise. *Physical Review A*, 38(11):5938, 1988.

[96] N. J. Kasdin and T. Walter. Discrete simulation of power law noise [for oscillator stability evaluation]. In *Frequency Control Symposium, 1992. 46th., Proceedings of the 1992 IEEE*, pages 274–283. IEEE, 1992.

[97] M. Kasevich and S. Chu. Measurement of the gravitational acceleration of an atom with a light-pulse atom interferometer. *Applied Physics B*, 54(5):321–332, 1992.

[98] M.P. Fewell. Adiabatic elimination, the rotating-wave approximation and two-photon transitions. *Optics Communications*, 253(1-3):125–137, 2005.

[99] K. C. Wright, L. S. Leslie, and N. P. Bigelow. Raman coupling of Zeeman sublevels in an alkali-metal Bose-Einstein condensate. *Physical Review A*, 78(5), 2008.

[100] X. Zang, J. Yang, R. Faggiani, C. Gill, P. G. Petrov, J.P. Hugonin, K. Vynck, S. Bernon, P. Bouyer, V. Boyer, and P. Lalanne. Interaction between atoms and slow light: A study in waveguide design. *Physical Review Applied*, 5(2), 2016.

[101] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade. *Photonic Crystals. Molding the Flow of Light*. Princeton University Press, 2nd edition, 2008.

[102] D. E. Chang, V. Vuletić, and M. D. Lukin. Quantum nonlinear optics photon by photon. *Nature Photonics*, 8(9):685–694, 2014.

[103] J. D. Thompson, T. G. Tiecke, N. P. de Leon, J. Feist, A. V. Akimov, M. Gullans, A. S. Zibrov, V. Vuletić, and M. D. Lukin. Coupling a single trapped atom to a nanoscale optical cavity. *Science*, 340(6137):1202–1205, 2013.

[104] A. Goban, C.-L. Hung, J. D. Hood, S.-P. Yu, J. A. Muniz, O. Painter, and H. J. Kimble. Superradiance for atoms trapped along a photonic crystal waveguide. *Physical Review Letters*, 115:063601, 2015.