Ultracold atoms in dressed potentials

Tiffany Laura Harte
Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity Term 2017

Acknowledgements

I would like to begin by thanking my supervisor, Chris Foot, for giving me the opportunity to work in the field of cold atoms, and for guiding the project while giving me the freedom and independence to develop my own research interests and take the experiment in an exciting new direction. I would also like to thank the original team: Dimitri Trypogeorgos, who was a fantastic source of ideas and inspiration, Ed Owen, who introduced me to all the quirks of the experiment, and Graham Quelch, whose technical support and sense of humour were invaluable to the smooth running of the basement.

Special thanks go to Elliot Bentine, who joined the experiment in my second year and was the best lab partner I could have wished for. His determination, creativity and ingenuity turned every challenge we faced into an opportunity for growth and improvement, and this attitude helped us develop an ever stronger apparatus after weathering each experimental storm. Our friendship is one of the most valuable results of my DPhil, and I hope this is just the first of many projects that we will work on together.

It has also been a pleasure to work with Kathrin Luksch and Adam Barker, who have provided a fresh outlook and enthusiasm for the project. It’s nice to know that the experiment is being left in your very capable hands! Thanks also to Sean Ravenhall, Rowan Moore and Leo Xu for making the basement such a positive and cheerful place to work. My original inspiration to enter the field of cold atoms came from undergraduate projects at the University of St Andrews: my sincere thanks go to Donatella Cassettari and Graham Bruce for their encouragement and advice, and for giving me the opportunity to continue working on a series of fun and interesting projects with them. I would also like to thank Dr Louise Durning and Carmella Elan-Gaston at Lincoln College, for their support and advice, and for making college such a welcoming and friendly environment.

Thanks to all the amazing dancers in the Oxford University Scottish Dance Society for making me feel so welcome when I was new to Oxford, for having a sensible attitude towards cake, and for providing countless evenings of fun and distraction when I needed it most. Above all, dancing has given me invaluable friendships, and I look forward to dancing with you all soon. Special thanks for ongoing friendship and support, even when I didn’t reply to messages for weeks (or months), to Laura Taylor, David Richards and Varun Ramraj. I would also like to thank Sam Denny for his kindness, patience, support and understanding over the last four years, and particularly for his forbearance over these last few months of writing. My greatest thanks are reserved for my mother, Linda Harte, without whose unwavering love and support none of this would have been possible.

List of publications

The following publications have arisen from work discussed in this thesis:

- D. Trypogeorgos, T. Harte, A. Bonnin, C. Foot, *Precise shaping of laser light by an acousto-optic deflector*, Optics Express **21**, 24837-24846 (2013).
- T. Harte, E. Bentine, E. Owen, D. Xu, B. Yuen, C. Foot, *An introduction to dressed-atom adiabatic potentials for ultracold atoms*, Instrumentation **2**, 76-84 (2015).
- E. Bentine, T.L. Harte, K. Luksch, A. Barker, B. Yuen, J. Mur-Petit, C.J. Foot, *Species-selective confinement of rf-dressed atoms*, J. Phys. B **50**, 094002 (2017).

The following publications are in preparation:

- T.L. Harte, E. Bentine, K. Luksch, D. Trypogeorgos, B. Yuen, C.J. Foot, *Multiple-radiofrequency dressed adiabatic potentials*, arXiv:1706.01491 [cond-mat.quant-gas].
- R. Moore, B. Soueres, D. Trypogeorgos, C.J. Foot, T.L. Harte, *Modulation transfer spectroscopy on the potassium ground state crossover*.

Publications not discussed in this thesis:

- T. Harte, G.D. Bruce, J. Keeling, D. Cassettari, *A conjugate gradient minimisation approach to generating holographic traps for ultracold atoms*, Optics Express **22**, 26548-26558 (2014).
- S.J. Thomson, L.S. Walker, T.L. Harte, G.D. Bruce, *Measuring the Edwards-Anderson order parameter of the Bose glass: a quantum gas microscope approach*, Phys. Rev. A **94**, 051601 (R) (2016).
- D. Bowman, T.L. Harte, V. Chardonnet, C. de Groot, S.J. Denny, G. Le Goc, M. Anderson, P. Ireland, D. Cassettari, G.D. Bruce, *High-fidelity phase and amplitude control of phase-only computer generated holograms using conjugate gradient minimisation*, Optics Express **25** 11692-11700 (2017).

Ultracold atoms in dressed potentials

Tiffany Laura Harte

Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity Term 2017

Abstract

Time-varying fields are widely used to extend the accessible range of trapping potentials for ultracold atoms. This work explores two very different examples of such fields, in the radiofrequency and optical regimes, whose interactions with trapped atoms can both be described in terms of the dressed atom picture.

Forming the basis of this work are radiofrequency dressed adiabatic potentials based on macroscopic trapping coils. Atoms are confined at the south pole of the resultant oblate spheroidal trapping surfaces. This work describes the extension of these potentials by two different methods: the application of multiple dressing radiofrequencies, and addition of a rapidly-scanned optical dipole trap.

This is the first experimental demonstration of a multiple-radiofrequency dressed adiabatic potential, explored using ultracold ^{87}Rb atoms confined in a highly configurable double well. Due to the independent generation of each constituent dressing frequency, the depth of each trapping well and the height of the barrier are easily manipulated, enabling precise and reliable transfer of atoms between the available trapping geometries. Experimental work includes an exploration of the potential-shaping capabilities of the three-radiofrequency system, and characterisation of the potential landscape using radiofrequency spectroscopy with good agreement to the eigenvalues numerically calculated using Floquet theory. This initial exploration of multiple-radiofrequency techniques lays the groundwork for applications in studying double well physics in a two-dimensional system, and independent state or species selective manipulation of trapped atoms. The potential shaping capabilities of this method can also be extended by applying additional trapping frequencies.

In a supplementary line of experimental work, an optical dipole trapping system has been constructed, and the trapping beam aligned to the lower surface of the radiofrequency dressed trapping shell in order to sculpt the radial confinement. Beam shaping is achieved using an acousto-optic deflector, which can be used to produce either a composite array of static deflected beams, a rapidly-scanned painted potential, or some combination of the two approaches.

The development and extension of the experimental apparatus required to implement these enhanced dressed state potentials is explored, and the challenges of their experimental implementation considered.

Contents

List of Figures

Chapter 1

Introduction

The initial experimental demostration of a Bose-Einstein condensate (BEC) in 1995 [1] was facilitated by the evolution of techniques in both laser cooling [2] and radiofrequency (RF) evaporation [3], that allowed the confinement and cooling of a dilute atomic vapour to quantum degeneracy. In the two decades that have followed this early work, the development of increasingly sophisticated methods by which to confine and manipulate ultracold atoms has allowed research into cold atom systems to expand across a diverse range of experimental directions.

A system cooled to quantum degeneracy provides an ideal testing ground for quantum simulation [4], in which a cold atom analogue of a complicated condensed matter system can be tested at various levels of complexity, and the system observables probed with single-site precision [5]. The configurability of cold atom systems, coupled with the absence of thermal effects, have also made BECs ideal candidates for interferometry and precision measurements [6], while the emerging field of atomtronics [7–9], in which atomic analogues of electronic systems are engineered, requires precision waveguiding of a trapped condensate and exploits features such as its long-range phase coherence. One common theme of each of these emerging technologies is their critical reliance on our ability to precisely engineer the potential landscape experienced by the trapped atom cloud.

One method of sculpting a confining potential for ultracold atoms is by the application of a radiofrequency (RF) field to a static magnetic trap. This dramatically changes the character of the magnetic potential, introducing features on small length scales and providing additional parameters by which to tune the potential and manipulate the confined atoms. This work concerns the augmentation of such an RF dressed potential using two different methods: applying multiple RF fields to create concentric trapping spheroids, and developing a rapidly-scanned optical dipole trap to sculpt the radial confinement on these spheroidal shell surfaces. The experimental methods developed and characterised in this thesis provide a strong foundation for the development of highly configurable potentials with the benefits of macroscopic RF dressing techniques, namely an intrinsic state- and species-selectivity, robust construction, and the generation of smooth potentials with low heating rates.

This chapter begins with a discussion of the motivation for developing these experimental techniques and the history of the experiment apparatus in Sec. 1.1, and is followed by an overview of the thesis material in Sec. 1.2.

1.1 Background and motivation

Static magnetic fields are used to create trapping potentials for ultracold atoms and are crucial for the creation of Bose-Einstein condensates. However, these traps are constrained by Maxwell's equations: the number of field minima cannot exceed the number of field-generating wires used [10], and no field maxima exist in a region with no source. Creating a trapping field with the fine features required for a number of applications including beamsplitting for precision measurements or atomtronics thus requires multipole terms of the hexapole order or higher, with an associated reduction in confinement strength associated with losses and non-adiabaticity [11].

The application of an oscillating field to couple the internal states of a trapped atom augments this static trap, introducing small-scale features to the potential landscape. These fields can either be off-resonance, achieving dipole force trapping using either optical [12] or microwave fields [13], or resonant, typically RF, radiation that couples internal states to create a dressed-state adiabatic potential (AP) [14]. Both methods can be intuitively described in terms of the dressed-atom picture of the interaction between atoms and radiation [15].

This experiment was developed with two distinct applications in mind: working towards simulation of strongly correlated electrons, and developing methods to non-destructively probe quantum systems. Although these applications diverge significantly in the details of their implementation, they share a number of common requirements and both draw strength from the various aspects of the experimental apparatus developed during this work. The foundation for both experimental goals is the RF-dressed quadrupole, which forms an oblate spheroidal trapping surface. For all purposes considered in this work, the atoms are allowed to congregate on the lower surface of the spheroid under the influence of gravity. Under standard experimental conditions, this trap has a radial trapping frequency (along the direction tangential to the trapping surface) of a few Hz, with an axial trapping frequency (perpendicular to the trapping surface) of tens to hundreds of Hz.

1.1.1 Simulating strongly correlated electrons

The ability to use the RF-dressed AP surface to create a 2D gas motivates its use in observing strongly correlated states analogous to the fractional quantum Hall effect (FQHE) that arises in 2D electron systems with discrete energy levels. Examples of these systems include Landau-quantised states such as those formed in graphene [16, 17], at the interface between a semiconductor and insulator under an applied magnetic field, or upon rapid rotation of an atomic BEC [18]. In the FQHE, the linear relationship between Hall resistance and magnetic field is disrupted by quantisation of the Hall resistance, producing a stepwise increase with increasing magnetic field, and corresponding magnetoresistance minima. One model for this effect is a uniform sea of charge into which the magnetic field is incorporated by introducing vortices of displaced charge, each carrying one flux quantum. A uniform vortex distribution corresponds to composite particles consisting of electrons associated with flux quanta. A single flux quantum per electron results in the integer quantum Hall effect whereby quantised Hall resistance values $R_H = \frac{h}{ne^2}$ arise at magnetic fields associated with filling of an integer number of Landau levels n. In the fractional case, multiple flux quanta are associated with each electron and a rational fraction of energy levels, characterised by filling factor ν, are filled. This implies particles with fractional charge and quantum numbers, supported by a topological order [19, 20].

The idea of using atoms to simulate strongly correlated electron systems exploits the mathematical similarity between the Lorentz force on charged particles $F_{\text{Lorentz}} = q\vec{v} \times B$, and the Coriolis force on neutral rotating particles $F_{\text{Coriolis}} = 2m\vec{v} \times \vec{\Omega}$. In both equations, $\vec{v}$ is the particle velocity; q and $\vec{B}$ are the particle charge and applied magnetic field, while m and $\vec{\Omega}$ are the particle mass and angular frequency of rotation. However, a critical difference lies in the centrifugal force arising during rotation, without analogue in the effect of magnetic field experienced by charged particles. This can destabilise the rotating system: the radial trapping frequency defines the centrifugal deconfinement limit for the rotation rate, beyond which atoms are ejected. However, irrespective of the value $\omega_\perp$ itself, the angular rotation frequency necessary for the analogy to hold is on the order of $0.98\omega_\perp$. This condition allows the centripetal and radial trapping terms in the Hamiltonian to cancel, leaving only the artificial magnetic field [18, 21].

Angular momentum can be imparted to a BEC by rotation of optical lattices [22] or stirring with laser beams [23], or directly imprinting phase shifts [24] and exploiting properties of systems such as optical flux lattices [25] in which transitions are resonantly driven between atomic states. This is a challenging experimental problem to tackle, but one that is of great interest due to the insights it could hold into topological states [26, 27].

Our approach to this problem is to supplement the shell trap with an optical dipole trapping beam focussed on the lower surface to increase the radial trapping frequency to a few kHz. This dipole trap will increase the radial trapping frequency, correspondingly increasing the deconfinement limit and rotation rate necessary to access the strongly correlated regime. The conditions for entering the FQH regime are [21]:

$$\frac{\omega_\perp - \Omega}{\omega_\perp} < \frac{8a}{Nl_z} \tag{1.1}$$

$$\mu \approx \hbar\omega_\perp \frac{a}{l_z}\nu. \tag{1.2}$$

Ω and $\omega_\perp$ indicate the rotation and radial trapping angular frequencies, with a the scattering length, ν the Landau level filling factor, $l_z = \sqrt{\hbar/m\omega_z}$ the axial harmonic oscillator length, and μ the chemical potential. An increase in Ω facilitated by an increased $\omega_\perp$ allows access to FQH states with a higher atom number, chemical potential and temperature. Relaxing the temperature criterion makes experimental implementation easier, while an increased atom number makes correlation detection feasible[1].

To make the trap smooth enough for rotation, the dipole trap is scanned in a circular orbit using an acousto-optic deflector, to create a time-averaged potential. The design and implementation of this hybrid dipole trap/RF-dressed potential is discussed in Chapter 8. This system can be used more generally to sculpt the radial potential on the canvas provided by the AP shell trap. This constitutes a first step towards the infrastructure required for simulating strongly correlated states, but must be supplemented with methods to control, rotate and image small numbers of atoms. In preparation for imaging, a lock-in imaging lattice has been partially constructed [29];

---

[1] A preprint submitted to ArXiv in 2010 describes the observation of strongly correlated atoms in the FQH regime, though with a small number of atoms ($N < 10$) [28].

based on the principle that correlated states are simply rotated and rescaled under time of flight [30], this lattice can be turned on after some time of flight to preserve the atom distribution for fluorescence imaging.

1.1.2 Manipulation of atoms in multiple potential wells

The second experimental goal for which this apparatus is well suited, and the one on which we ultimately concentrated experimental effort, is that of developing independently-manipulable species-selective wells by the application of an RF dressing field with multiple frequency components. This multiple-RF (MRF) approach has a variety of applications including developing methods for fully quantum probes that would allow observation and probing methods for cold atom systems to keep pace with the developments in quantum simulation experiments.

As discussed in Chapter 2, RF-dressed potentials hold a number of advantages over conventional static macroscopic magnetic or optical traps: with low heating rates [31] and greater resilience to field fluctuations [32], they introduce spatial features on length scales far smaller than the trapping wire configuration, while maintaining the benefits of magnetic trapping including smoothness, robustness, and state- or species-selectivity depending on the value of the Landé g_F factor. This latter feature arises due to the dependence of the AP characteristics on the Zeeman splitting of the trapped atomic species, and means that the same dressing RF traps atoms with different g_F values at different locations in space. A high level of dynamic control over these potentials is also possible, with a large number of experimental handles including the RF amplitude and polarisation, and the properties of the underlying static field. By applying a carefully chosen selection of multiple independently-generated RF fields, it would therefore be possible to achieve independent control and overlap of two atomic states or species.

One example of a useful scheme exploiting the species-selective properties of the

MRF AP is to create a single well for a bulk condensate of one species spatially overlapped with a double well for a second species [33], which contains only a small number of atoms to act as an impurity used to probe the bulk system properties. By observing the dynamics of an impurity acting as a quantum probe to infer the characteristics of its environment, as previously demonstrated using a single atom [34], a more detailed level of knowledge of both local and global system properties could be obtained than with conventional classical probes. These methods could also be adapted to perform nondestructive measurements of the bulk quantum system. In the experiment towards which this work progresses, details of which can be found in [33, 35], the tunnelling dynamics of an impurity trapped in a double-well immersed in a bulk condensate could, for example, be used to act as a thermometer of the surrounding quantum system or to investigate non-Markovianity and dissipation of quantum information.

To this end, a multiple-frequency approach to the RF dressing method was developed, in the first experimental demonstration of this technique. This was used to create a double well trapping potential for a single species cloud of ^{87}Rb atoms. This method can be extended to investigate 2D tunnelling dynamics and interferometry and applied to the impurity probe scheme as outlined above. Additional frequencies can also be added to create more dressing wells, and the independent tuneability of constituent RF amplitudes and polarisations fully explored to create a range of trapping geometries.

1.1.3 The combined experimental vision

This work therefore began with two experimental visions: quantum simulation of correlated electron states, and developing methods to fully probe quantum simulation experiments. The former focusses on augmenting a single-RF adiabatic potential with a rapidly-scanned optical dipole trap, and the latter on extending the method of RF dressing by applying multiple dressing frequencies to create multiple concentric trapping surfaces: in this case focussing on demonstrating and characterising a double well potential created using three independent trapping frequencies.

However, these goals share a number of common elements. These include developing techniques to manipulate and image small numbers of atoms, and in work tangential to this thesis, a lock-in imaging lattice has been constructed. A complementary high resolution imaging setup is discussed in Sec. 4.5.4.2. They can also be used in conjunction to augment the properties of each other: the second RF shell would be a useful phase reference or could be used to assist the rotation of small atom numbers; the addition of the dipole trap with an arbitrary intensity distribution to two closely spaced MRF wells could shape the radial confinement or realise, for example, a Josephson junction array.

Most importantly, at the heart of both routes of experimental work is the RF-dressed quadrupole. Throughout this work, techniques have been developed to better understand and work with this potential and its hardware. The rapidly-scanned optical dipole trap is focussed onto the lower surface of this shell trap to sculpt the radial potential, whereas the multiple-RF method creates a series of concentric shell traps that can be used to shape the potential along the vertical axis. These two approaches can also be combined to achieve potential shaping in both directions; their combination, together with other available techniques including time-averaging, thus grants a detailed level of independent control over all aspects of the potential with scope for expansion to additional trapping geometries. General applications of this hybrid trapping potential include potential shaping for atomtronics, rotation and 2D interferometry, in addition to the schemes discussed above. While the two experiments diverge significantly in their later implementation, a robust, reliable and flexible apparatus has been constructed that is capable of accommodating not only these investigations but also retains the flexibility to be turned to a variety of alternative experiments.

1.2 Thesis overview

This thesis concerns the development of an apparatus and the associated experimental techniques to create a multiple-RF dressed double well trapping potential, and the design and construction of a rapidly-scanned optical dipole trap. I begin by describing the theory behind magnetic trapping and RF-dressed adiabatic potentials in Chapter 2, introducing the dressed atom picture that intuitively describes both trapping methods experimentally explored in this work. An introduction to optical dipole trapping in the context of laser beam shaping methods is presented in Chapter 3. With this experimental motivation in mind, I will then describe the experimental apparatus used and developed to create a BEC of ^{87}Rb atoms in a TOP trap in Chapter 4. The main adaptation necessary to implement multiple-RF dressing, namely a wideband impedance matching network, is discussed in detail in Chapter 5, where an introduction to RF impedance matching methods is also presented as this useful and straightforward technique is relatively little known in the atomic physics community. This sets the stage for the introduction of experimental RF-dressing techniques in Chapter 6, where I discuss the experimental methods by which cold atoms are cooled in a time-averaged adiabatic potential and controllably loaded into the MRF double shell configuration. In Chapter 7 I present the results of experimental investigations into the potential shaping capabilities of the three-RF potential, and show an experimental verification of the calculated eigenenergies in the multiple-RF system. I then return to the optical dipole trap, describing its design and preliminary results on beam shaping with the AOD, before discussing the experimental implementation of the rapidly-deflected beam and its alignment to a trapped atom cloud. To conclude, I assess the progress so far, imminent experimental steps, and some new directions for which the experiment is poised.

Chapter 2

RF-dressed adiabatic potentials

Following laser cooling, the ability to magnetically confine a dense cloud of atoms for extended periods of time in high vacuum enabled the development of evaporative cooling techniques that reduce the temperature by many orders of magnitude to create a BEC [1, 3]. Optical traps are now commonly used in conjunction with magnetic potentials to give more opportunities for potential shaping [36, 37], with some all-optical techniques now demonstrated [38]. However, a theme that has been followed in the Oxford group is that of extending the range of magnetic trapping geometries by the application of time-varying fields. Most relevant in the context of this work is the application of radiofrequency (RF) radiation to introduce features to the potential landscape on a far smaller scale than the field-producing components [39, 40]. These potentials offer an alternative set of properties to optical traps, including an intrinsic state- and species-selectivity based on the g_F value of trapped atoms [33, 41, 42], an inherently robust construction in comparison to optical traps, and smooth trapping potentials with demonstrably low heating rates when created using macroscopic coils located a few cm from the atoms [31]. The trap properties can be easily tuned using the RF amplitude and polarisation, and properties of the underlying static field, allowing dynamic control over a trapped atomic gas in one, two or three dimensions [31, 43–45].

This chapter serves as a short introduction to RF-dressed magnetic traps, beginning with a discussion of the basic concepts of magnetic trapping in Sec. 2.1. The dressed atom formalism used to describe the interaction with the applied RF field is introduced in Sec. 2.2, before a description of RF dressed magnetic traps and the associated experimental developments in Sections 2.3 and 2.4. Finally, the multiple-RF dressed potentials that are the subject of this work are introduced in Sec. 2.5.2, laying the groundwork for their first experimental implementation as described in Chapters 6 and 7.

2.1 Magnetic trapping

The interaction between a weak magnetic field $\mathbf{B_0}$ and a neutral atom with magnetic moment $\boldsymbol{\mu}$ causes a Zeeman splitting with energy:

$$V_{\text{mag}} = -\boldsymbol{\mu} \cdot \mathbf{B_0} = g_F \mu_B m_F |\mathbf{B_0}| \tag{2.1}$$

with μ_B the Bohr magneton and the atom in a hyperfine level denoted by total angular momentum F with Landé g-factor g_F and $(2F+1)$-fold degeneracy in m_F, the projection of $\mathbf{F}$ onto the quantisation axis defined by the external field. For the weak field limit to apply, the interaction with the external field must be much less than the hyperfine interaction between total electronic and nuclear angular momenta $\mathbf{J}$ and $\mathbf{I}$ respectively, such that we can consider only the projections of $\mathbf{I}$ and $\mathbf{J}$ onto $\mathbf{F}$.

An atom in a magnetic field gradient experiences a spatial variation of its Zeeman energy levels, resulting in a magnetic force

$$\mathbf{F}_{\text{mag}} = -g_F \mu_B m_F \boldsymbol{\nabla}(|\mathbf{B_0}|). \tag{2.2}$$

States with $m_F g_F > 0$, or equivalently $\boldsymbol{\mu}$ antiparallel to $\mathbf{B}$, are attracted to a region of lower field and can be confined at a local field minimum. Conversely, the high-field-seeking states with $m_F g_F < 0$ are repelled from low-field regions and are therefore the anti-trapped states[1]. For the low fields considered in this work, states with $m_F g_F = 0$ are unaffected by the linear variation of the field.

Magnetic trapping of neutral particles was first demonstrated in 1977 using a cold neutron beam [47] in an experiment using a superconducting magnetic storage ring with a $35\,\mathrm{T}$ field to compensate for the small magnetic moment of the neutron. Magnetic trapping of atoms was first achieved in 1985 [48] using laser-cooled sodium atoms in a magnetic quadrupole potential. A simple trapping configuration formed using two coaxial coils carrying currents that flow in opposite directions, the quadrupole field takes the form

$$\mathbf{B}_0(\mathbf{r}) = B'_q(x\hat{\mathbf{e}}_x + y\hat{\mathbf{e}}_y - 2z\hat{\mathbf{e}}_z) \tag{2.3}$$

with B'_q the radial quadrupole gradient[2]. The resultant field takes a value of zero in the centre of the coil configuration and increases linearly in amplitude with distance from this point, and the iso-magnetic field surface takes the form of an oblate spheroid.

The trap strength and parameters depend on the coil dimensions, current, and magnetic moment of the confined atomic state/species. However, prior to the advent of laser cooling techniques to sufficiently pre-cool atoms below the trap depth, it was possible only to mildly deflect atoms rather than trap them with typical laboratory fields [49]. For the example of sodium described in [48], the trap depth was on the order of $17\,\mathrm{mK}$, requiring substantial pre-cooling of the trapped atoms. The trap lifetime in this example was 8.3 s, limited by collisions with the background gas. However, the quadrupole trap has an intrinsic limitation from Majorana transitions which occur when the magnetic moment of a trapped atom cannot adjust adiabatically to

---

[1] High-field-seeking states could in principle be trapped at a maximum in the magnetic field. However, given Maxwell's equations $\nabla \cdot \mathbf{B} = 0$ and $\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}\right) = 0$ for static fields in free space, static magnetic field maxima are precluded by Wing's theorem [46].

[2] The factor of 2 in the z field is necessary to satisfy $\nabla \cdot \mathbf{B} = 0$.

local variations in the orientation of the magnetic field, causing a transition into an untrapped state. For an atom to follow the field adiabatically, the speed of its motion must be slow in comparison to the rate at which the magnetic moment precesses about the field, i.e. $\omega_L \gg \omega_{\text{trap}}$. The Larmor frequency $\omega_L = g_F \mu_B B/\hbar$ determines the rate of precession about the magnetic field with amplitude $B = |\mathbf{B_0}|$, describing the timescales over which the atom can adapt to a changing field, and ω_{trap} is the orbital angular frequency of the trapped atom. The likelihood of a non-adiabatic transition is therefore greater the higher the frequency of atom motion, and the smaller the length scale of magnetic field variation. This becomes a limitation in the sharply-varying region near the quadrupole field zero, forming an increasingly significant loss channel as atoms are cooled towards quantum degeneracy due to the increased atomic density and the increased sampling of regions close to the field zero.

The non-adiabatic region of the potential must therefore be modified to address this issue. One approach is the application of a blue-detuned laser beam to repel atoms from this region, as used in the second realisation of BEC (the first realisation in sodium) [3]. Another is to modify the magnetic field itself, for example by applying a time-varying bias field as employed in the first experimental demonstration of BEC [1]. This creates a time-orbiting potential (TOP) [50] of the form

$$\mathbf{B}_T(\mathbf{r}) = B'_q(B_{\text{TOP}}\cos(\omega_{\text{TOP}}t)x\hat{\mathbf{e}}_x + B_{\text{TOP}}\sin(\omega_{\text{TOP}}t)y\hat{\mathbf{e}}_y - 2z\hat{\mathbf{e}}_z), \tag{2.4}$$

where B_{TOP} is the applied bias field amplitude. The angular frequency of rotation ω_{TOP} is chosen such that $\omega_{\text{trap}} \ll \omega_{\text{TOP}} \ll \omega_L$ and the atoms experience an effective harmonic time-averaged potential, constrained spatially by Majorana losses at the zero-field-crossing which is swept in a circular orbit with radius $r_0 = B_{\text{TOP}}/B'_q$. This limitation to the temperature of trapped atoms means that the TOP field is applied after an initial stage of evaporative cooling in the quadrupole.

Another alternative is to remove the field zero altogether by implementing some variant on the Ioffe-Pritchard configuration [51], which uses a two-dimensional quadrupole generated by four parallel wires and supplementary circular pinch coils to provide confinement in the axial direction. This potential takes the form of an anisotropic harmonic trap for atoms close to the trap centre with $k_B T \ll g_F \mu_B B_0$, and varies on length scales sufficient to maintain adiabaticity as the atoms traverse the trap centre. Many variants on this trapping scheme have been implemented, most notably the cloverleaf [52] and baseball [53] configurations[3], and QUIC trap [54].

Increasingly complex magnetic trap structures are possible using the trapping wire geometry to shape the potential, with time-varying fields allowing the dynamic manipulation and waveguiding of trapped atom clouds. These techniques have been employed on a scale of a few hundred $\mu\mathrm{m}$ by placing trapping wires on chips [55]. The length-scales of these traps lead to extremely high quadrupole gradients which characterise the atom chip RF-dressing schemes discussed in Sec. 2.4. A significant challenge of atom chips has been the fragmentation of trapped clouds due to imperfections in the current-carrying wires: any variation in current translates directly to a variation in the potential due to the close proximity of the atoms to the trapping wires [56]. Time-varying fields have been applied to time-average these corrugations to zero [57].

Interesting magnetic trapping configurations have also been realised on a macroscopic scale, either by shaping the field-generating wires to create, for example, ring traps for interferometry [58–60], or by applying time-averaging techniques to create configurations including a magnetic double well [61]. However, dramatic changes to the form of a magnetic potential are possible by the application of RF fields to a static magnetic trap. Just as the quadrupole can be modified by both static and slowly varying magnetic fields to maintain the adiabaticity of atom motion near the trap centre, RF fields can be used to couple the Zeeman-modified hyperfine states of a trapped

---

[3] This baseball trap was used to demonstrate the state-selectivity of magnetic trapping, exploited for sympathetic cooling.

atom to similarly smooth the transition between m_F states and allow adiabatic transfer of atoms between them. Referred to more generally as ‘adiabatic potentials’ (APs) or ‘RF-dressed adiabatic potentials’, these RF-modified magnetic traps are most easily explored within the framework of the dressed atom picture of atom-radiation interaction.

2.2 A primer on the dressed atom picture

Although now perhaps better known for its description of the optical dipole force, the dressed atom picture was first developed in 1969 as an intuitive quantum mechanical explanation of the resonances observed in light scattered by atoms interacting with an RF field in optical pumping experiments [15]. Its application to the interaction of atoms with intense laser beams was appreciated in 1977 [62], and it has since become a standard description of the interaction between atoms and both RF and optical fields. The formalism has been applied to describe statistical behaviour such as spontaneous emission and even to describe the process of Sisyphus cooling [63]. While it is was developed as a descriptor of strong field interactions, the dressed atom picture can also be applied to weak fields; in this case the system reduces effectively to a description of the light shift [63, 64].

The central concept of the dressed atom formalism is that the atom and radiation field, and the interaction between them, can be treated as a combined system described by the dressed atom Hamiltonian

$$H_{\text{total}} = H_A + H_L + H_R + V_{AL} + V_{AR} \tag{2.5}$$

where H_A, H_L and H_R are, respectively, the Hamiltonians for the isolated atom, laser field, and an initially-empty reservoir, and V_{AL}, V_{AR} describe the interaction between the atom and laser field[4], and atom and reservoir. Strictly speaking, this models the atom interacting with the laser mode inside a cavity, such that the only variation in photon number can be caused by absorption or emission by the atom [64]. The extension of this to more general experimental situations depends on the assumptions that the experimental laser field is locally the same as that experienced by an atom in the cavity, and that the cavity is not small enough to modify the spontaneous emission. This corresponds to requiring that the energy density can be described by $\langle N \rangle / V$ where $\langle N \rangle$ is the mean photon number and V the cavity volume, that this energy density is fixed, and both $\langle N \rangle$ and V tend to infinity [63, 64]. The laser intensity is assumed constant throughout the interaction, with the number of photons lost from the field negligible relative to the photon number distribution but this distribution narrow relative to the large mean photon number $\langle N \rangle$. The reservoir provides a means of damping for the atom-light system, so in the case of negligible spontaneous emission as in the RF-dressed system, it is possible to consider just the atom, radiation field, and interaction between them. The dressed atom picture for a three-level system interacting with an RF field is illustrated in Fig. 2.1.

The ^{87}Rb $|F = 1\rangle$ level comprises three states $m_F = \pm 1, 0$, their degeneracy lifted by an external magnetic field[5]. The atomic Hamiltonian is given by $H_A = \hbar\omega_0 |{+1}\rangle \langle{+1}| - \hbar\omega_0 |{-1}\rangle \langle{-1}|$. Here, $\pm\hbar\omega_0$ give the Zeeman splitting of the $m_F = \pm 1$ levels, with $m_F = 0$ experiencing no magnetic interaction. These states are illustrated in panel (a) of Fig. 2.1. The Hamiltonian of the photon field is $H_L = \hbar\omega_L(a^\dagger a + \frac{1}{2})$ where $a^\dagger$ and a indicate photon creation and annihilation operators and the $\hbar\omega_L/2$ term corresponds to the zero-point energy.

In the non-interacting system $V_{AL} = 0$ the atom and light field are completely independent but described as a joint system by the dressed atom formalism. Illustrated in Fig. 2.1 (b), these uncoupled, or 'bare' states are described by the Hamiltonian

---

[4] While the case of RF dressing does not involve a laser field, this terminology is applied in the more conventional description of the dressed atom picture. This can be generalised to the case of RF dressing.

[5] Given that $g_F = -1/2$ in this state, $m_F = -1$ will lie highest in energy and $m_F = +1$ lowest.

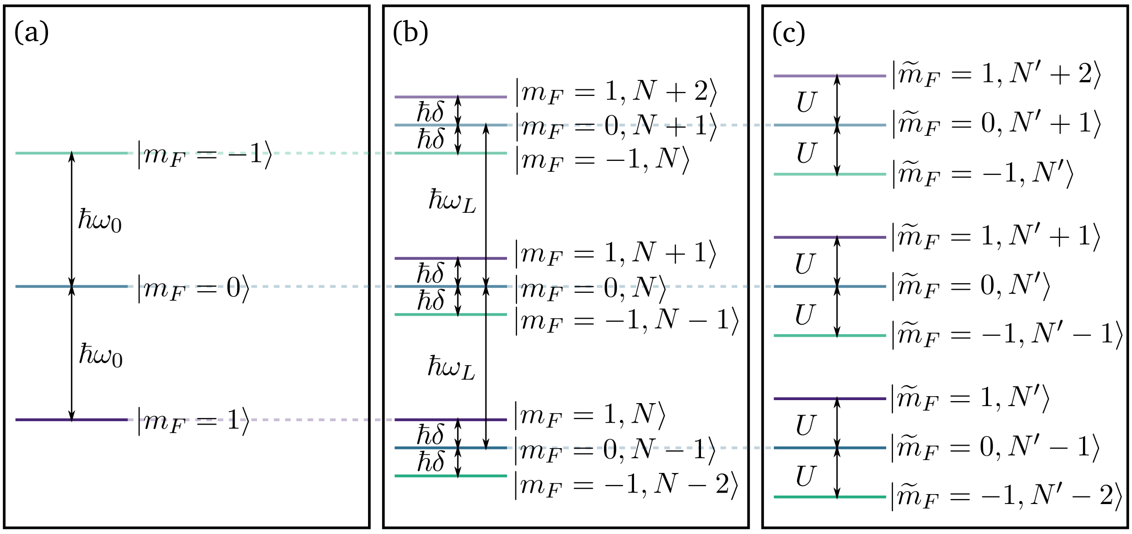

Figure 2.1: Transforming Zeeman-separated atomic energy levels into dressed states interacting with an optical or RF field. (a) Separation of the three m_F components of the $|F = 1\rangle$ level due to a static external magnetic field. (b) Bare states of the atom/photon-field system, neglecting interactions. An infinite ladder of states are formed, labelled by the original m_F state and the photon number N. In the case shown here, the energy splitting $\hbar\omega_L$ between identical m_F states with consecutive photon numbers exceeds the Zeeman separation between m_F states by a small detuning $\hbar\delta$. The original m_F states are unchanged, but is now considered in conjunction with the photon field. (c) Taking into account the interaction between atoms and photon field, the bare states are separated by energy U according to the interaction strength, and relabelled in terms of new quantum numbers $\widetilde{m}_F$ and N'. In the context of AP formation as described here, the interaction strength associated with a given field i is parameterised on-resonance by the Rabi frequency Ω_i, with $U = \hbar\Omega_i$.

$H_{\text{bare}} = H_A + H_L$ and labelled by $|m_F, N\rangle$ with N the photon number. The corresponding eigenergies are then given by

$$H_{\text{bare}}\,|{+1}, N+1\rangle = \hbar(\omega_A + (N+1)\omega_L)\,|{+1}, N+1\rangle \tag{2.6}$$

$$H_{\text{bare}}\,|0, N\rangle = \hbar(N\omega_L)\,|0, N\rangle \tag{2.7}$$

$$H_{\text{bare}}\,|{-1}, N-1\rangle = \hbar(-\omega_A + (N-1)\omega_L)\,|{-1}, N-1\rangle \tag{2.8}$$

where the zero-point energy is neglected. The states are grouped in this manner, rather than simply as $|m_F, N\rangle$ due to the assumption of a small detuning $|\delta_L| = \omega_L - \omega_0 \ll \omega_0$, as is the case for RF dressing. This means that the photon energy $\hbar\omega_L$ is close to resonance with the Zeeman splitting $\hbar\omega_L$ such that, for example, the state $|0, N+1\rangle$, lying at $\hbar\omega_L$ above $|0, N\rangle$, ends up much closer to $|1, N\rangle$ than $|0, N\rangle$ with a separation given by $\hbar\delta_L$. Applying this to arbitrary numbers of photons added and subtracted from the system, this gives rise to the groupings of states illustrated in Fig. 2.1 (b)[6]. The combined system can therefore be viewed as a grouping of the three states $m_F = -1, 0, 1$ with corresponding photon numbers $N-1, N, N+1$, into manifolds separated by $\hbar\delta$ and repeated at intervals of $\hbar\omega_L$. Each manifold has constant $(\mathrm{sgn}\,(g_F) \cdot m_F + N)$ between its constituent states. Isolated states at the extremes of this ladder can be neglected: its extent is determined by the number of photons available in the system which can be considered effectively infinite for all applications in this work.

While the presence of the radiation field affects the way the atomic energy levels are represented, no modification has yet been made to the levels themselves. The interaction V_{AL} is proportional to both the atomic dipole moment and the time-varying magnetic field amplitude, and corresponds to the absorption of a photon resonant with the spacing between m_F states to mediate the transition between them at the expense of a photon from the external field. The dipole moment provides the usual selection rule $\Delta m_F = \pm 1$, subject to the polarisation of the applied field[7], while the nonzero matrix elements of the electric field component allow a change in photon number $\Delta N = \pm 1$. Where such a resonant transition occurs, the strongest coupling is therefore between allowed states within each manifold, with a matrix element

$$\langle 0, N|V_{AL}|1, N+1\rangle = \frac{\hbar\Omega_{N+1}}{2} \tag{2.9}$$

where the coupling strength is described by the Rabi frequency $\Omega_{N+1} = \Omega_0\sqrt{N+1}$[8].

The vacuum Rabi frequency Ω_0 describes the transition in the presence of one

---

[6] Fig. 2.1 assumes $\delta > 0$ such that upon addition of a photon, a state with a given m_F will overshoot the next state as dictated by $\hbar\omega_0$.

[7] While the general case is considered here, the specific cases of different RF polarisations are considered in Sec. 2.3.1.

[8] The Rabi frequency, denoted by Ω, describes an angular frequency whenever mentioned in this work.

photon. Likewise, the matrix element is proportional to N for the transition $|-1, N-1\rangle \to |0, N\rangle$. The new eigenstates are a linear combination of the bare states that are coupled to create them, and are labelled as $|m_F', N'\rangle$ with an energy splitting $\hbar\widetilde{\Omega}_N$ where $\widetilde{\Omega}_N = \sqrt{\Omega_N^2 + \delta^2}$. The Rabi frequency of the coupling determines the separation between these dressed state levels, and is defined by the matrix element for the interaction. These eigenstates are illustrated in Fig. 2.1(c).

Where permitted by the dipole moment selection rules, nonresonant couplings are also possible between states in different manifolds. These can be neglected in an initial treatment in an approach equivalent to making the rotating wave approximation. However, as discussed in Sec. 2.3.3, these effects cannot be neglected entirely.

2.3 Theory of rf-dressed adiabatic potentials

The extension of the dressed atom picture to rf-dressed trapping potentials relies simply on the inclusion of a spatial dependence in the magnetic field, and thus also a reliance on the polarisation of the applied RF. A spatially-varying field introduces a spatial dependence to the Zeeman splitting of atomic hyperfine levels, and thus to the detuning of the applied RF from resonance. The adiabatic potential forms at the spatial location of this resonance, where the interaction between the atom and RF field forces the energy levels apart to form an avoided crossing on which atoms can be confined. The form of this trapping potential is modified by factors including polarisation, gravity and nonresonant processes.

2.3.1 AP formation

For the experiment discussed in this work, the static field is the magnetic quadrupole described by Eq. 2.3. We apply a spatially homogeneous RF field

$$B_{\text{RF}}(t) = B_1[\cos(\omega_1 t)\hat{\mathbf{e}}_x - \sin(\omega_1 t + \varphi_1)\hat{\mathbf{e}}_y] \tag{2.10}$$

where B_1 denotes the RF field amplitude, ω_1 its angular frequency, and φ_1 controls the phase difference between x and y components and thus the field polarisation.

Given the energy scales of RF radiation, the relevant Zeeman splitting is that between the atomic hyperfine levels, and described by Eq. 2.1. For the quadrupole field, these levels vary linearly in space relative to the quadrupole centre. With typical quadrupole gradients accessible in the lab of up to $\sim 300\,\mathrm{G\,cm^{-1}}$ and the ^{87}Rb $|F = 1\rangle$ Zeeman splitting of $0.7\,\mathrm{MHz\,G^{-1}}$, the Zeeman splitting between hyperfine states is on the order of up to a few MHz.

As in Fig. 2.1, the application of RF field in the absence of interactions results in a repeated manifold of bare states $g_F \mu_B m_F |\mathbf{B}_0| + \sum_i n_i \hbar\omega_{\text{rf}}$, each separated by the photon angular frequency $\hbar\omega_{\text{rf}}$ and effectively infinite in extent[9]. As the coupling is between Zeeman-separated components of the $|F = 1\rangle$ hyperfine level of the electronic ground state, spontaneous emission and therefore the reservoir coupling term are neglected. This ladder of bare eigenstates is illustrated in Fig. 2.2(a), this time plotted against the spatially-varying magnetic field strength.

Taking, for simplicity, the two bare states $|1, N\rangle$ and $|0, N+1\rangle$, applying the zero-interaction Hamiltonian yields eigenvalues $\hbar(\omega_A + N\omega_{RF})$ and $(N+1)\hbar\omega_{RF}$. At the point in space in which the hyperfine splitting is equal to the energy of the RF photon, these two states should cross, as illustrated in Fig. 2.2 (a), corresponding to a detuning $\delta = 0$ in Fig. 2.1(b). However, with the interaction taken into account as

---

[9] A semiclassical description would equally well apply to this system, and indeed a semiclassical calculation is ultimately used to perform the final calculation of the RF-dressed states considered. However, the dressed atom picture provides a clear intuitive picture of the RF dressing process.

in Fig. 2.1(c), the states are forced apart by the interaction. Formerly, atoms traversing the level crossing would have maintained the same m_F state and experienced no trapping potential. However, with an interaction mediating the avoided crossing, an atom can instead experience this as a trapping potential, adiabatically following the change in m_F value. Associated with this adiabatic traversal of the avoided crossing is a gradual evolution in the bare-state composition of the overall atomic state. Implied by the colour change of the trapped state in Fig. 2.2(b), this variation is described by the mixing angle $\xi = \arccos\left(\frac{\delta}{\sqrt{\delta^2+\Omega_R^2}}\right)$. This formulates the state composition as an angle between $0°$ and $180°$, with $0°$ a pure $|-1, N-1\rangle$ state, $180°$ a pure $|1, N+1\rangle$ state

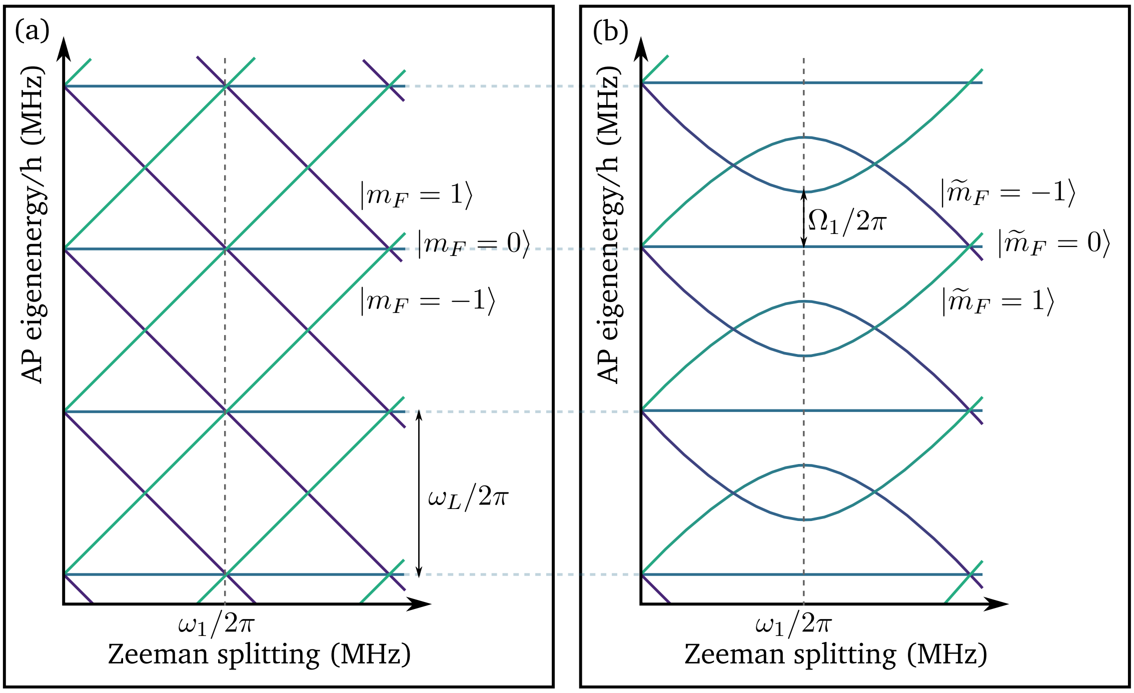

Figure 2.2: Eigenstates of the rf-dressed quadrupole field in the absence of gravity. (a) The ladder of bare states arising from considering the atom and RF field in the absence of interactions: the three m_F states separated by the Zeeman effect are repeated in an infinite ladder at intervals of the RF photon energy $\hbar\omega_{\mathrm{rf}}$. (b) Avoided crossings form at points in space where the resonance condition $\omega_0 = n\omega_{\mathrm{rf}}$ is fulfilled, forming a trapping potential for atoms in $|\widetilde{m}_F = 1\rangle$ and barrier for atoms in $|\widetilde{m}_F = -1\rangle$. The magnitude of these potential features are described in terms of the Rabi frequency Ω_1. With purple, blue and green lines indicating the bare states $m_F = 1, 0, -1$ respectively, the adiabatic transition between states is indicated by the colour change in the avoided crossing.

and 90° corresponding to the resonance and a maximal $|0, N\rangle$ component in the mixed state. The translation of this state decomposition into space therefore depends on both detuning and coupling strength.

If the interaction between the atoms and dressing RF is suffiently strong, and the variation of the static field orientation sufficiently slow with position, an atom traversing the avoided crossing will adiabatically follow the new eigenstate labelled by the quantum number $\widetilde{m}_F$ [65], with atoms in $\widetilde{m}_F = 1$ confined in the AP. The trapping potential takes the form

$$U_{\mathrm{AP}}(\mathbf{r}) = \widetilde{m}_F \hbar \sqrt{\delta^2(\mathbf{r}) + \Omega_1^2(\mathbf{r})} \tag{2.11}$$

and as such depends on both the angular frequency detuning of the RF from resonance, $\delta(\mathbf{r}) = |g_F \mu_B B_0(\mathbf{r})/\hbar| - \omega_1$, and the coupling strength as expressed by the Rabi frequency Ω_1. This coupling strength determines the extent to which the levels are separated, with a stronger interaction corresponding to a greater energy level separation and flatter potential well, and also introduces the vectorial nature of the potential. The Rabi frequency is determined by a combination of the RF field amplitude and its polarisation according to

$$\Omega_1 = \frac{g_F \mu_B}{2\hbar} B_1^{\perp}(\mathbf{r}) \tag{2.12}$$

where $B_1^{\perp}(\mathbf{r})$ is the RF field component perpendicular to the local magnetic field.

To extend this picture to three dimensions, the coupling is considered over the isomagnetic surface of the underlying static field. For the quadrupole field this forms an oblate spheroidal resonant shell. Atoms are confined on the shell surface with a coupling strength that varies according to the dressing RF polarisation as parameterised by $B_1^{\perp}$. The interaction therefore falls from its maximum to zero as the RF varies between perpendicular and parallel.

This variation in coupling strength can be quantified by first generalising 2.10 to a spatially inhomogeneous, three-dimensional field

$$B_{\text{RF}}(\mathbf{r},t) = B_x\cos(\omega_1 t)\hat{\mathbf{e}}_x + B_y\cos(\omega_1 t + a)\hat{\mathbf{e}}_y + B_z\cos(\omega_1 t + b)\hat{\mathbf{e}}_z \tag{2.13}$$

where $B_{x,y,z}$ are the RF field amplitude in the x, y and z dimensions respectively, and a and b describe the phase of the y and z components with respect to the x component of the RF field. Using spherical coordinates, the perpendicular component of the field can be further decomposed [66, 67]:

$$\begin{aligned}(B_1^\perp)^2 =& B_x^2\cos^2(\theta)\cos^2(\phi) + B_x^2\sin^2(\phi) + B_y^2\cos^2(\phi)\\ &+ B_y^2\cos^2(\theta)\sin^2(\phi) + B_z^2\sin^2(\theta) + 2B_xB_y\cos^2(\theta)\cos(a)\sin(\phi)\\ &+ 2B_yB_z\cos(\theta)\sin(\phi)\sin(\theta)\cos(a-b) + 2B_xB_z\cos(\theta)cos(\phi)\sin(\theta)\cos(b)\\ &- 2B_xB_y\sin(\phi)\cos(\phi)\cos(a) + 2B_xB_y\sin^2(\phi)\sin(a)\cos(\theta)\\ &+ 2B_yB_z\sin(\theta)\cos(\phi)\sin(a-b) + 2B_xB_y\sin(a)\cos(\theta)\cos^2(\phi)\\ &+ 2B_xB_z\sin(\phi)\sin(b)\sin(\theta).\end{aligned} \tag{2.14}$$

Here, θ and ϕ are the angles in spherical polar coordinates that define the unit vector of the quadrupole field:

$$\hat{\mathbf{e}}_q = \frac{1}{\sqrt{x^2+y^2+4z^2}}\begin{pmatrix} x \\ y \\ -2z \end{pmatrix} = \begin{pmatrix} \sin(\theta)\cos(\phi) \\ \sin(\theta)\sin(\phi) \\ -\cos(\theta) \end{pmatrix}. \tag{2.15}$$

For the case of a circularly polarised RF field, the Rabi frequency is therefore given by

$$\Omega_1 = \frac{g_F\mu_B B_1}{2\hbar}\left(1 \pm \frac{2z}{\sqrt{x^2+y^2+4z^2}}\right) \tag{2.16}$$

with B_1 the magnetic field amplitude of the ω_1 RF field and x, y, z Cartesian coordinates with an origin at the quadrupole centre. In this work the handedness is chosen such that the coupling is maximised at the south pole of the resonant spheroid, and zero at the north pole. The other case employed in this work is that of an RF field linearly polarised in the xy plane, with a Rabi frequency:

$$\Omega_1 = \frac{g_F \mu_B B_1}{2\hbar} \left(\frac{r_\perp^2 + 4z^2}{r_\perp^2 + r_\parallel^2 + 4z^2} \right)^{1/2} . \tag{2.17}$$

Here, $r_\parallel$ and $r_\perp$ describe the coordinates parallel and perpendicular to the polarisation direction in the xy plane. The resonant ellipsoid therefore has maximum coupling at points for which the parallel component is zero, and zero coupling at those points on the equator for which the perpendicular component is zero.

2.3.2 Adiabaticity and losses

As with the quadrupole-based trapping techniques discussed in Sec. 2.1, maintaining adiabaticity is a key factor in minimising atom loss from RF-dressed APs. Non-adiabatic traversal of the avoided crossing results in atom loss due to transitions into untrapped states. This is by no means an isolated problem: the losses are known as Landau-Zener transitions in analogy with Zener's 1932 treatment of non-adiabatic following of molecular energy levels in response to finite-timescale variation of system parameters [68][10]. Following an extension of this to a three-level system [72], the probability that the atoms follow the dressed state eigenenergy adiabatically is given by [42, 73]

$$P_{\mathrm{LZ}}(\mathbf{r}) = 1 - \exp\left(-\frac{h\Omega(\mathbf{r})^2}{4 g_F \mu_B \partial_t (B_q(vt))} \right) \tag{2.18}$$

where $\Omega(\mathbf{r})$ denotes the spatially-varying coupling strength as described by the Rabi frequency. The quadrupole field amplitude B_q and atom velocity v define the relevant quantity as the field gradient as experienced by the moving atom. The loss rate is there-

---

[10] This work was published separately in 1932 by Zener [68], Landau [69], Majorana [70] and Stueckelberg [71].

fore largely determined by the coupling strength itself, with a strong coupling pushing the states involved in the avoided crossing further apart to flatten the potential and increase the deviation of energy levels from the bare states. The quadrupole gradient determines the gradient of the resulting adiabatic potential, where a sharper change in the gradient at the avoided crossing increases the probability of non-adiabatic transitions [73]. The higher the velocity of the confined atoms traversing the crossing, the less likely they are to adiabatically follow the change in internal state.

2.3.3 Modifications to the potential

The exact form of the potential is modified by gravity, and a gravitational energy mgh term incorporated into the Hamiltonian. This shifts the potential minimum downwards in space, corresponding here to a translation towards higher frequencies. As the trapped atoms are supported against gravity by the quadrupole field, a high field gradient reduces the effect of this shift.

Further considerations include non-resonant couplings and multiple-photon resonances. In an effect known as the Bloch-Siegert shift [63, 74, 75], a spatial translation of the location of the resonant avoided crossing can occur due to non-resonant coupling between manifolds that acts to impose a small shift of the eigenvalues involved. Avoided crossings can also occur at integer values of the dressing RF angular frequency ω_1 subject to conservation of angular momentum. These effects are exacerbated by strong coupling.

2.4 RF dressed adiabatic potential experiments

The dressed atom formalism as applied to both atoms [15] and neutrons [76] has been employed to describe dipole force trapping associated with both optical and microwave fields [77], and off-resonant microwave dressing in a magnetic trap [13]. This latter work underscores the relationship between adiabatic potentials created by radiation in different time-varying regimes, and highlights the benefits of increased atom density and lifetime over static magnetic traps, and the negligible effects of spontaneous emission heating and nonadiabatic transitions in contrast to optical dipole traps, despite the cost of large microwave field amplitudes. Interest then turned to the RF regime, and the dressed atom framework has been used to describe the effects of RF radiation on the eigenenergies of magnetically confined atoms. The first proposal for this as a means to confine ultracold atoms emphasised the application of this method to accessing the 2D regime with a trapped quantum gas [14].

Experimental implementations of RF dressed adiabatic potentials has diverged down two paths that share the underlying physics while each exploiting different features of the method: RF-dressed potentials based on macroscopic coil arrays, in which the first experimental implementation of RF-dressing was realised [78], and those created using atom chips, in which a BEC was first produced [79]. These two methods have very different characteristics, and correspondingly different applications.

2.4.1 Adiabatic potentials on atom chips

Atom chip experiments [80] are performed over length-scales of a few hundred $\mu\mathrm{m}$, with RF and static fields generated using wires tens of $\mu\mathrm{m}$ in width and separation. These often exploit the polarisation-dependence of the RF-dressed potential [79], associating variations in coupling strength with the strong static field gradient to shape the potential to create features on the order of a few $\mu\mathrm{m}$.

It was in a chip-based adiabatic potential that the vectorial nature of the coupling strength was first recognised [79], and this feature was central to subsequent work on shaping the potential and manipulating trapped atoms. The exploitation of this vectorial coupling, combined with the close proximity of the trap to the chip itself, to create a potential with tuneable small-scale features was predominantly motivated by the desire to create a coherent, adiabatic matter-wave beamsplitter that could be applied to interferometry and precision measurements [65, 79]. The difficulty in designing such a process using static fields lies in constraints imposed by Maxwell's equations, requiring higher-order multipole terms to create the fine features required for beamsplitting sequences. The associated reduction in confinement strength results in a non-adiabatic beam splitting process [11]. However, the application of RF dressing techniques can overcome this problem, allowing the introduction of small length-scale features without loss of trap strength or adiabaticity during atom manipulations [11, 65, 81]. Experimental challenges include high loss rates at the atom densities required for precision measurements [82], and the creation of a sufficiently compact, low-power, and high repetition rate apparatus for practical application in the field.

Atom chip construction comprises separate wires to generate the static and RF-dressing fields, tens of microns in width and separation, producing highly asymmetric trapping potentials that can operate within both the 1D [79] and 3D [32] regimes depending on the combination of applied currents. Additional bias fields are also employed to position the potential and assist with its shaping.

The fundamental element of these traps is a double well that can be used as either a temporal [79] or spatial [32] beamsplitter. These double wells are variants on the initial implementation [79], in which a DC magnetic trapping wire supplemented by an external bias field traps atoms 80 $\mu\mathrm{m}$ below the RF wire. The high static field gradient results in a corresponding rapid variation in the angle between the quantisation axis and the polarisation axis of the dressing RF, leading to a spatial variation of the coupling strength. This coupling strength, and the corresponding dressed state level splitting, is maximised directly beneath the RF wire; over short distances the vectorial coupling variation supercedes the static-field Zeeman splitting such that the dressed level separation reduces outwards from this point to create a double well. Potential shaping is achieved by changing the amplitude of the RF current, with the strong confinement aiding adiabaticity. The well separation is controlled using the dressing field frequency, with a range of 500 kHz to 4 MHz varying the well spacing between 3 and 80 μm and corresponding to both tunnelling and fully isolated regimes.

This first experiment demonstrated adiabatic, phase-coherent splitting of a trapped condensate, and subsequent work using the same principles but based on a three-wire RF-dressed Ioffe trap performed interference measurements to distinguish between a split BEC and those formed separately in independent wells [43]. A full characterisation of the effects of polarisation and angular dependence is also presented in this later work. A variety of wire configurations based on these principles have so far been both proposed and demonstrated, each associated with a particular class of potential landscape to produce cold atom analogues to various optical elements, ranging between different schemes for double well configurations and rings for Sagnac interferometry [65]. With radial trap frequencies on the order of a few kHz, these traps are suitable for interferometric characterisation of 1D systems [83].

In a significant step towards the application of these potentials to precision measurements, subsequent work demonstrated that RF dressing allows splitting and interference of a trapped condensate to be performed at sufficiently high atomic densities for interferometric or precision measurement applications [82]; the general suppression of field fluctuations in RF-dressed traps has also been highlighted [32]. The compact nature of atom chips lends them to applications in portable interferometry systems, and recent work has demonstrated the compact, high-repetition rate setups that makes this scaling feasible [84]. While less flexible than optical dipole trapping methods, the ability to create small-scale features without laser beams is a significant advantage to applications in which simplicity and robustness are prized.

Theoretical studies of RF-dressed atom chip potentials, supported by experimental measurements, have studied the significant effect of violating the rotating wave approximation on the eigenenergies and allowed transitions [85]. The effect of nonlinear Zeeman shifts on the trap strength and resulting spatial distribution of atoms has also been theoretically considered [86]. State selective manipulations have been suggested as canditates for quantum information processing procedures [43], with an initial experimental demonstration of this technique using a ^{87}Rb-^{40}K mixture on an atom chip [41]. More recent theoretical work regarding atom chip adiabatic potentials includes the proposals for RF-dressed lattices created by wire arrays [87], and the application of these to quantum information processing [88]. Exotic connected geometries for 1D waveguides have also been suggested by combining RF dressing fields with an inductive magnetic trap [89], while an inductive ring trap design combined with RF dressing has been proposed that would reduce the scale of ring traps to those conventially only associated with diffraction-limited optical dipole traps [90].

2.4.2 RF dressing using macroscopic coils

In contrast, coils several cm in extent can be used to create both the static magnetic and dressing RF fields, trapping atoms on an isomagnetic bubble- or shell-like surface that echoes the static field structure. The properties of these shell traps depend on the underlying static field strength, and the frequency and amplitude of the dressing RF, but are typically hundreds of $\mu\mathrm{m}$ to a few mm in extent, with radial trapping frequencies on the order of a few Hz and axial frequencies tens of Hz, with a static field gradient of tens to hundreds of $\mathrm{G\,cm^{-1}}$. Depending on the combination of trap parameters, atoms either coalesce on the lower surface of the shell under the influence of gravity, or, for a higher coupling strength, can flow to fill the entire surface.

While atom chips produce highly anisotropic traps with high field gradients, well-suited to potential shaping in the 1D regime, macroscopic coil APs are more naturally oriented towards 2D physics on a surface smoothed by distance from the trapping wires. The exact form of the resonant trapping surface depends on the underlying static magnetic trap; APs have been generated using a QUIC [78], Ioffe-Pritchard [91], and quadrupole [44] static field, with corresponding trapping surfaces manifesting as asymmetric cigar-shaped bubbles, curved sheets, and, for the spherically symmetric quadrupole, spheroidal shells. The vectorial nature of the coupling manifests as a variation in coupling strength over the shell surface, and this can be exploited to manipulate the trapping geometry. Using a low RF field amplitude, the first magnetically-trapped 2D gas has been demonstrated and studied [31, 92]. These macroscopically-generated APs are more generally characterised by their smoothness, robust construction, and low heating rates [31].

The potential landscape experienced by trapped atoms can be diversified further by the application of additional fields, namely an optical dipole trap, or time-varying bias field. In the former approach, an optical field is overlapped with the RF-dressed shell. A stacked sequence of rings created by the application of a blue-detuned standing wave was proposed [93] and a single ring sandwiched by two such sheets has been demonstrated [66]. A local sculpted optical potential can also be applied to atoms confined on a shell surface; progress towards combining the RF-dressed quadrupole and a rapidly-scanned painted optical potential is explored in Chapter 8.

The application of a bias field varying at kHz frequencies creates a time-averaged adiabatic potential (TAAP) [94]. This technique has been used to create a double well [95], and variable-radius rings [96] with demonstrated applications in state-selective manipulations that could be useful in Sagnac interferometry [45]. This has also provided the means to flatten and shape the shell surface, with schemes developed for evaporation and rotation of trapped atoms [44]. The vertically-offset double well is used in our shell trap loading scheme as described in Sec. 6.1.2.

2.4.3 Extending the RF dressing technique

A number of aspects of the RF dressing method have also begun to curve towards more applied uses. A direct application is that of using an RF-dressed interferometry setup to probe non-equilibrium coherence dynamics in 1D systems [85]. In more of a tangential step, RF dressing has also been proposed as a means of reducing the dephasing observed in magnetically confined microwave clocks [97]; The strong coupling and precise features in the potential landscape have also lent the RF dressing technique to the production of an atom laser [98].

However, one seemingly obvious extension of RF dressing has been little explored. Multiple-RF dressing was first proposed in the early days of RF dressing experiments [42], with the suggestion of a range of potentials from double wells to lattices and combs using an appropriate combination of frequencies. The application of the MRF method to independent state- or species-selective manipulation of trapped atoms was also alluded to at this time. However, this method was not adopted in earnest until nearly a decade later, when renewed interest inspired a series of more rigorous theoretical treatments [99–102] and a recent proposal for creating multiple-RF dressed lattices for artificial magnetic fields using Raman laser beams [103]. The first experimental demonstration of MRF potentials is discussed in this thesis, where we implement a three-RF double well as a demonstration and exploration of MRF techniques, laying the groundwork for future applications [104].

2.5 Theory of multiple-RF dressed adiabatic potentials

The extension to a multiple-RF dressed system is conceptually straightforward: each frequency is associated with a spatially distinct resonance that forms an avoided crossing. The manifestation of the avoided crossing as either a well or barrier for a given eigenstate is determined by the eigenenergy gradient, and since one avoided crossing reverses this gradient at a given point in space, the application of multiple RF components creates an alternating sequence of wells and barriers.

Experimental and technical considerations are discussed in Chapters 6 and 7, and relevant aspects of the numerical solution in Sections 2.5.1 and 2.5.2. This work focusses on the simplest example of MRF dressing: three frequencies applied to generate two trapping wells with an independently-tuneable barrier to separate them. While this configuration alone has scope for future work in 2D and species-selective double well physics, these principles can be extended to any combination of RF signals subject to adiabaticity and technical constraints to create complex geometries including lattices, asymmetric double wells and flat-bottomed potentials.

This description of the MRF theory follows that presented in our recent publication [104]; a more detailed discussion of the theoretical aspects of MRF dressing is forthcoming in Ref. [102].

2.5.1 Atoms in a multi-component RF field

The homogeneous RF field of Eq. 2.10 extended to multiple RF frequency components takes the form

$$\mathbf{B}_{\mathrm{MRF}}(t) = \sum_i \frac{1}{\sqrt{1+\kappa_i}} B_i[\cos(\omega_i t + \phi_i)\hat{\mathbf{e}}_x - \kappa_i \sin(\omega_i t + \phi_i)\hat{\mathbf{e}}_y] \tag{2.19}$$

where B_i, ω_i, and ϕ_i are the amplitude, angular frequency and relative phase of each frequency component. The phase term ϕ_i is common to both x and y components of each frequency to dictate the phase difference between each frequency component. In keeping with our experimental implementation of either circular or linear RF fields, the polarisation of each field component is reduced to a description by the parameter κ_i, with $\kappa_i = 1$ denoting circularly polarised RF and $\kappa_i = 0$ linear.

For either a circularly or linearly polarised RF field, MRF dressing yields a system again described by the dressed-atom Hamiltonian:

$$V = \sum_i \hbar\omega_i a_i^\dagger a_i + g_F \mu_B \mathbf{F} \cdot [\mathbf{B}_0(\mathbf{r}) + \mathbf{B}_{\mathrm{MRF}}(t)] \tag{2.20}$$

where $\mathbf{B}_{\mathrm{MRF}}(t)$ is the second quantised operator for the MRF field in Eq. 2.19. The first term accounts for the energy of the i^{th} RF field component with angular frequency ω_i and corresponding photon creation and annihilation operators $a_i^\dagger$ and a_i. The second term describes the interaction between the atomic spin operator $\mathbf{F}$, and the total magnetic field comprising static and RF components $\mathbf{B}_0(\mathbf{r})$ and $\mathbf{B}_{\text{MRF}}(t)$ respectively, following the convention in Ref. [105].

In the absence of interactions with the RF field, the system eigenstates are the tensor products of the Fock states of each RF field and the atomic Zeeman substates $|n_1, n_2, ..., m_F\rangle$. These states are labelled by the magnetic quantum number m_F and the n_i photons in the i^{th} frequency component of the RF field. They form a ladder of eigenenergies $g_F \mu_B m_F |\mathbf{B_0}| + \sum_i n_i \hbar \omega_i$ in which the three Zeeman substates are repeated with a spacing of ω_f, the highest common factor of RF angular frequencies ω_i.

While the dressed-atom picture provides an intuitive visualisation of the RF dressing process, the large mean photon number of the RF field allows it to be represented classically by replacing $a_i^\dagger$ and a_i by their mean field values $\bar{n}_i$. This is performed within the context of the interaction picture, in which $V \to U_{\text{RF}}^\dagger V U_{\text{RF}}$ and $|\psi\rangle \to U_{\text{RF}}^\dagger |\psi\rangle$ with $U_{\text{RF}} = \exp\Big(i \sum_i a_i^\dagger a_i \omega_i t\Big)$. The RF field is decomposed into components parallel and perpendicular to a unit vector $\hat{\mathbf{z}}'$ defined in the direction of $\mathbf{B_0}$, where $\mathbf{F} \cdot \mathbf{B}_0 = B_0 F_{z'}$.

The parallel component has amplitude

$$\zeta_i(\mathbf{r}) = \frac{1}{\sqrt{1+\kappa_i}}(\sin\theta\cos\phi + i\kappa_i \sin\phi), \tag{2.21}$$

where $\cos\theta = -2z(x^2+4z^2)^{-1/2}$ and $\cos\phi = [(x^2+4z^2)/(x^2+y^2+4z^2)]^{1/2}$ from the definition of the static quadrupole field. The anticlockwise and clockwise rotating components of the perpendicular field have amplitudes

$$\alpha_i(\mathbf{r}) = \frac{1}{\sqrt{2+2\kappa_i}}(\cos\theta - i\sin\theta\sin\phi - \kappa_i\cos\phi), \tag{2.22}$$

$$\beta_i(\mathbf{r}) = \frac{1}{\sqrt{2+2\kappa_i}}(\cos\theta + i\sin\theta\sin\phi + \kappa_i\cos\phi). \tag{2.23}$$

In this basis the semiclassical version of the Hamiltonian Eq. 2.20 becomes

$$\begin{aligned}V(t) = g_F\mu_B B_0 F_z + \frac{g_F\mu_B}{2}\sum_i \Bigg[&\left(\frac{\alpha_i}{\sqrt{2}}F_- + \frac{\beta_i}{\sqrt{2}}F_+ + \zeta_i F_z\right) B_i e^{i(\omega_i t+\phi_i)} \\ &+ \left(\frac{\alpha_i^*}{\sqrt{2}}F_+ + \frac{\beta_i^*}{\sqrt{2}}F_- + \zeta_i^* F_z\right) B_i e^{-i(\omega_i t+\phi_i)}\Bigg]\end{aligned} \tag{2.24}$$

which is periodic in time with period $T = 2\pi/\omega_f$. The coefficients α_i, β_i and ζ_i give the projection of the field operator in the local circular basis, with $|\alpha_i|^2 + |\beta_i|^2 + |\zeta_i|^2 = 1$.

From Floquet theory, the eigenstates of this time-periodic Hamiltonian can be expressed in the form $|\psi(t)\rangle = \exp(iE't/\hbar)\,|\Psi(t)\rangle$, a product of a phase term and the time-periodic function $|\Psi(t)\rangle$. The time evolution operator $U(t, t+T)$ advances the system between times t and $t+T$, such that $U(0,T)\,|\psi(0)\rangle = |\psi(T)\rangle = \exp(iE'T/\hbar)\,|\psi(0)\rangle$. The phases $E'T/\hbar$ can be associated with the energy of the dressed eigenstates of Eq. 2.20 at time T [75, 102] such that the dressed state eigenenergies modulo $\hbar\omega_f$ are given by the $2F+1$ eigenvalues of $(-i\hbar/T)\log U(T)$. We calculate U through numerical integration of the Schrödinger equation.

2.5.2 The three-RF double well

The MRF field creates a first order avoided crossing at the location of the resonance with each of the three dressing frequencies. For atoms in the $|\widetilde{m}_F = 1\rangle$ state, this corresponds to trapping wells associated with ω_1 and ω_3, and a barrier at ω_2 to form a double well in which field amplitudes are varied to tune the well depths and barrier height. The dressed state ladder now has a spacing of $\hbar\omega_f$, where ω_f is the highest common factor of the constituent dressing angular frequencies. In the experimental work we use an MRF field with $\omega_f = 2\pi \times 0.6$ MHz and angular frequency components $\omega_{1,2,3} = (5, 6, 7) \times \omega_f$, i.e. $\omega_{1,2,3} = 2\pi \times (3.0, 3.6, 4.2)$ MHz. This example is plotted with experimentally realistic parameters in Fig. 2.3, alongside a plot of a single-RF well at the central angular frequency $\omega_2 = 2\pi \times 3.6$ MHz.

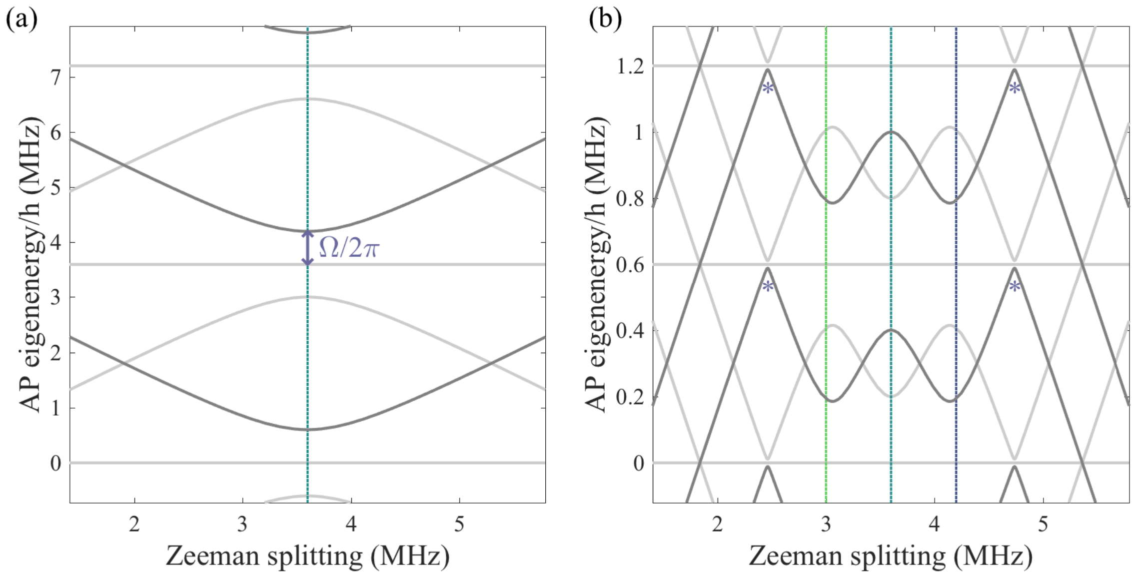

Figure 2.3: Examples of single- and triple-freqency RF dressing, plotted in the absence of gravity. (a) Dressed state eigenstates showing an avoided crossing formed by a single dressing RF at $\omega_1 = 2\pi \times 3.6$ MHz with field amplitude $\Omega_1 = 2\pi \times 200$ kHz. The ladder of dressed sates is repeated at an interval of the $\hbar\omega_1$. The trapped $|\widetilde{m}_F = 1\rangle$ state is emphasised. (b) Upon application of three dressing angular frequencies $\omega_{1,2,3} = 2\pi \times (3.0, 3.6, 4.2)$MHz with amplitudes $\Omega_{1,2,3} = 2\pi \times 200$ kHz, three avoided crossings form to create a double well trapping potential for $|\widetilde{m}_F = 1\rangle$. The energies of the eigenstates at each resonance are shifted by the presence of the other RF field components, translating the well minima in space. Weak avoided crossngs are also formed by multi-photon couplings at integer values of $\hbar\omega_f$, indicated by asterisks. The system periodicity is now defined by $\omega_f = 0.6$ MHz, the highest common factor of the RF components[11].

The potential shaping capabilites of this trapping configuration are explored in Chapter 6. Evident in Fig. 2.3(b) is the shift in the eigenenergy at each resonance due to the other RF field components, which has the effect of squeezing the double well minima closer together. The claim of independent tuneability of the trapping wells is therefore not a statement of the independence of the wells themselves, as there exists the influence of each field component at any given point in space. Rather it concerns the ability to independently tune the signals that predominantly influence each feature of the potential landscape, and use the MRF waveform itself to influence the cross-talk effects. An additional influence is that of gravity, omitted from Fig. 2.3 but considered in the context of potential shaping and atom manipulation in Chapter 6. For clarity, the effect of gravity is excluded from the figures presented in this work unless relevant.

Weak higher order resonances, highlighted in Fig. 2.3, also form at intervals of ω_f, creating a loss channel that influences the choice of dressing field parameters and the temperature at which atoms are loaded into the trap. This is considered in the loading scheme presented in Chapter 6, and explored in Ref. [102].

The experimental angular frequency separation $\omega_f = 2\pi \times 0.6$ MHz was chosen for two main reasons: at the standard[12] quadrupole field gradient $B'_q = 62.45\,\mathrm{G\,cm^{-1}}$, this corresponds to a shell separation of order $100\,\mu\mathrm{m}$, easily imaged with the low-resolution horizontal imaging system described in Sec. 4.5.4.1, and furthermore, this separation is a good match to the few hundred kHz field amplitudes of early operation in terms of maintaining adiabaticity. These adiabaticity constraints motivate the choice of experimental parameters including the frequency separation, RF amplitudes and static field gradient: an atom with constant velocity v moving through the spatially-varying potential is retained with a probability approximately given by the Landau-Zener model of Eq. 2.18. With a view to reducing the well spacing in or-

---

[11] This and similar figures of AP eigenenergies are created using MATLAB code developed by Elliot Bentine [35].

[12] This is the AP quadrupole gradient at which most exploratory and diagnostic work was performed; different values of the gradient have at various times been used.

der to perform double well experiments, one aspect of experimental investigation has therefore been to implement MRF potentials at reduced field amplitudes. This has required systematic removal of RF noise, discussed in Ref. [35] and is an ongoing line of work[13]. This condition also dictates our theoretical approach to the MRF problem: as the piecewise approach presented in Ref. [42] is invalid in this limit [99], Floquet theory is employed to calculate the MRF dressed state eigenenergies. Numerical artefacts are removed by appropriate meshing over the range of magnetic field values considered. An intuitive depiction of MRF dressing that uses the resolvent formalism to discard these artefacts is explored in Ref. [102], and another Floquet theory-based approach to the MRF problem has also been recently been released on ArXiv [101].

Each trapping well in the MRF configuration corresponds to a distinct AP shell with a coupling strength variation determined by the polarisation of the field component in question. In this work, we consider either linear or circular polarisation with a handedness that maximises coupling at the south pole, in both cases with a field amplitude sufficiently low that atoms flow to the south pole under the influence of gravity. The resonant surfaces therefore take the form of concentric spheroidal shells, and the resultant potential for atoms in this experimental regime is a vertically offset double well consisting of the lower surface of each shell, illustrated in Fig. 2.4.

In this work we consider a three dimensional atom cloud on each shell surface. However, a reduction in the field amplitudes Ω_i would allow the manipulation of a 2D gas [31]. This will allow us to exploit the easy tuneability of these double well potentials to investigate double well physics and interferometry in the 2D regime, analogous to the studies of 1D interferometry performed using RF-dressed double wells on atom chips. The three-RF double well is implemented and characterised in Chapters 6 and 7.

---

[13] We have implemented single-RF APs with Rabi frequencies as low as $\Omega \sim 50$ kHz, which should place us in the regime of a 2D gas, although this was not explicitly verified. Work is ongoing to reduce the amplitudes Ω_i and angular frequency spacing ω_f in the MRF potential to access the 2D regime.

[14] 3D sketches of AP surfaces are based on a figure originally calculated by Kathrin Luksch for Ref. [104].

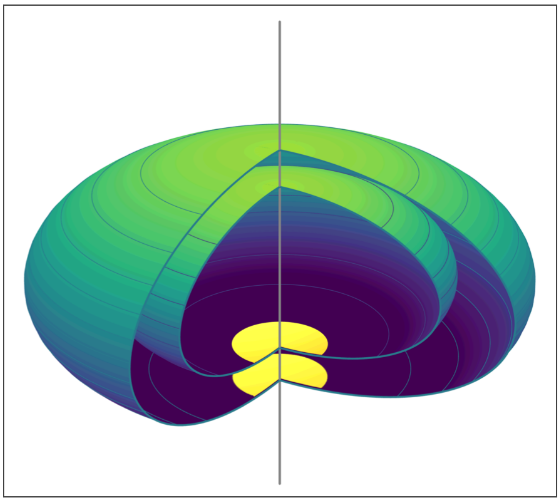

Figure 2.4: Cut-away sketch of the double shell trapping configuration, in which atoms (yellow) are trapped at the south pole of each trapping shell under the influence of gravity. These shells are associated with the resonances at ω_1 and ω_3 with a circularly polarised RF field. The separation between each shell is determined by the frequency separation between RF components and the quadrupole field gradient. The variation in coupling strength associated with the circularly-polarised RF field is indicated by shading, with purple (dark) indicating strong coupling and green (light) weaker coupling[14].

Chapter 3

Optical dipole traps for ultracold atoms

The term ‘adiabatic potential’ could as well be applied to optical dipole traps as RF-dressed potentials. Like the RF-dressed traps, optical potentials are based on the ability of an atom to adiabatically follow a perturbation to its energy levels upon the application of a field, and this interaction can again be intuitively described in terms of the dressed atom picture, though with a reservoir term that accounts for spontaneous emission. This short chapter introduces the interaction between an atom and optical field, in the context of both laser cooling and the creation of optical potentials. Rather than repeat a discussion of the dressed atom picture, focus here is on laser beam shaping techniques, to motivate the choices behind the rapidly-scanned dipole trapping setup discussed in Chapter 8.

3.1 Atom-light interactions

The interaction between an atom and laser field can be described in terms of the spontaneous and stimulated emission processes that define, respectively, the scattering and dipole forces. Viewed in a simplified two-level picture, photon scattering is the isotropic spontaneous emission process that can follow absorption of a resonant photon. Incoherent and dissipative, this is the irreversible process behind laser cooling, with a force given by [106]

$$F_{sp} = \frac{\hbar k S_0 \gamma}{2\left(1 + S_0 + \left(\frac{2\delta}{\gamma}\right)^2\right)} \tag{3.1}$$

where $\hbar k$ is the photon momenum and γ the spontaneous emission rate. The saturation parameter S_0 arises from the competition provided by the non-dissipative relaxation route provided by stimulated emission. Unlike the spontaneous emission rate that depends only on the excited state lifetime, this is proportional to the laser intensity. The scattering force therefore saturates at high intensities. As this force is proportional to the light intensity, it cannot alone create a potential minimum[1].

The dipole force is the conservative force associated with stimulated emission and the coherent redistribution and lensing of the electric field by the atom [107, 108]. Changes to the internal energy structure of the atom arise due to the interaction of an induced dipole $\mathbf{d} = \alpha(\omega)\mathbf{E_0}\cos(\omega t)$ with α the atomic polarisability, and the oscillating electric field $\mathbf{E_0}\cos(\omega t)$. This results in a coupled system described by the dressed-atom Hamiltonian whereby the system eignenstates are equivalent to those of a bare atom moving in an effective interaction potential. For the simplified case of a 2-level atom, the resultant dipole trapping potential is given by

$$U_{dip}(\mathbf{r}) = -\frac{3\pi c^2 \Gamma}{2\omega_0^3} I(\mathbf{r}) \left(\frac{1}{\omega_0 - \omega} + \frac{1}{\omega_0 + \omega}\right) \tag{3.2}$$

where ω is the angular frequency of the trapping light, ω_0 the atomic transition angular frequency, $I(\mathbf{r})$ the trapping intensity and Γ the transition damping rate [109].

---

[1] A potential minimum implicitly requires $\nabla^2 U > 0$, such that a corresponding force $\mathbf{F} = -\nabla U$ would require $\nabla \cdot \mathbf{F} \propto \nabla \cdot \mathbf{E} < 0$, implying an unphysical negative divergence of the electric field at the potential minimum.

With small angular frequency detunings $\delta = \omega - \omega_0$, the rotating wave approximation is valid and the $1/(\omega_0 + \omega)$ term can be neglected. This corresponds to neglecting the rapidly rotating terms in the interaction Hamiltonian. In this approximation the trapping potential reduces to:

$$U_{dip}(\mathbf{r}) = \frac{3\pi c^2}{2\omega_0^3} I(\mathbf{r}) \frac{\Gamma}{\delta} \tag{3.3}$$

This trapping potential varies with I/δ: for a red-detuned beam with $\delta < 0$ the trapping force acts towards high-intensity regions and for blue-detuned $\delta > 0$ the force is repulsive. Furthermore, the scattering force is proportional to I/δ^2 such that detuning the trapping beam far from resonance reduces heating associated with scattering events.

3.2 Laser beam shaping: motivation and techniques

The simplest form of a dipole trap is a single tightly-focussed red-detuned laser beam with trap depth determined by the beam intensity and tightness of confinement by the intensity gradient, analogous to optical tweezers [110]. A crossed-beam geometry provides good all-round confinement and is the standard approach to all-optical evaporative cooling whereby the beam combination facilitates simultaneous compression and evaporation [38]. Enhanced complexity in the dipole trap geometry has further improved control over this technique [111–114].

Optical lattices have allowed single-site control and imaging of a trapped condensate [5, 115, 116]. With a BEC loaded in an optical lattice, the Bose-Hubbard Hamiltonian can be realised, providing a powerful experimental framework for simulating the otherwise inaccessible electron gas in solids. For example, control of individual lattice sites has allowed investigation into disorder, the Mott insulator transition, topological states and quantum registers [117–120]. Lattices are formed from pairs of counterpropagating laser beams, the interference between which creates arrays of intensity maxima. Intensity variations tune interactions between atoms [118] while dynamic manipulation of lattice sites has been realised using acousto-optic techniques [22, 121]. However, the complexity of standing-wave lattices is inherently restricted. Furthermore, a harmonic external trapping potential or the Gaussian envelope of the lattice beams beams introduces a spatial dependence to the system density, blurring phase transitions and reducing simulation accuracy [122–124].

To create a full range of arbitrary lattice structures and continuous potentials, diffractive beam shaping techniques can be employed, incorporating the technologies of spatial light modulation and acousto-optic deflection. Whilst site resolution of interfering-beam lattices is half the laser wavelength, maintaining inter-site coherence in correlated systems [5, 121], that of diffractive techniques is the diffraction limit of the focussing optics. However, the power of diffractive methods lies in the increased range of accessible potentials, for atom confinement in both red- and blue-detuned light patterns. Examples of accessible structures range between bowtie lattices, a circular lattice corresponding to an infinite 1D array, and corral-type lattice potentials [125–130]. Arbitrary dynamic control over individual lattice sites has also been demonstrated [131] whereby a single input beam is manipulated to smoothly deform a potential and a confined condensate.

Of continuous potentials, flat-topped beams are of particular interest as square well potentials or box traps, uniform surfaces in structures such as quantum logic gates [132], and a canvas on which to form other potentials. Box traps have also been implemented using a series of blue-detuned potential barriers [133–135]. Power-law potentials are also extremely useful, extending small-scale compensation of external potential features [136] to broader compensation of an external trapping potential [124]. The ability to rapidly update the potential landscape would allow dynamic control over these potentials.

A spatial light modulator (SLM) is an array of pixels to which a computer-generated hologram is applied. Each pixel imposes a programmable phase retardation or intensity attenuation on an incident laser beam. In phase-modulation mode, the output plane intensity profile is related to the Fourier transform of the input plane phase profile. Binary SLMs such as ferroelectric liquid crystal, or the closely-related technology of digital micromirror devices (DMDs), are computationally simple enough that a direct search or simulated annealing algorithm can calculate the phase profile [131, 137]. DMDs have also been used to directly modify the amplitude of an incident beam [138], and as such can be applied to create large flat-topped potentials with less flexibility, but reduced computational cost than phase-modulation techniques. Able to access a wider range of output intensities, a higher phase resolution SLM requires more efficient techniques such as conjugate gradient minimisation [139, 140] or an iterative Fourier transform algorithm (IFTA) [141]. Maximising IFTA accuracy for large or complicated intensity patterns requires restricted output plane amplitude freedom, introducing noise close to the trapping region. IFTAs also suffer from optical vortices introduced during the calculation process. Difficult to remove once established, these correspond to a 2π phase winding and zero intensity [142]. Methods of limiting vortex formation have been investigated [132, 139], but computational intensity is still a limiting factor for large arbitrary continuous potentials. SLM technology has also been employed to obtain simultaneous control over both the phase and amplitude of the trapping plane optical field, using both a nematic liquid crystal SLM programmed using a conjugate gradient minimisation calculation [140], and using a DMD with superpixel technology [143].

Dynamic manipulation of a trapped condensate requires a sufficient update rate between static frames. Binary SLMs are associated with the highest pixel switching frequencies: the Texas Instruments digital micro-mirror device around 50 kHz and the Boulder Nonlinear Systems (BNS) ferroelectric liquid crystal SLM approximately 1 kHz. However, the BNS nematic liquid crystal SLM offering the more versatile range of accessible potentials is limited to a switching frequency of hundreds of Hz [131, 138, 144]. The power of SLM technology in precisely sculpting many aspects of an optical trapping field is therefore limited in situations where smooth dynamic manipulation of trapped atoms is required.

Acousto-optic deflection is based on the Bragg diffraction of an incident beam by an acoustic wave established in the acousto-optic deflector (AOD) crystal. The first order deflection angle depends on the acoustic wave frequency. Beam shaping can be performed either by rapid scanning of a single beam to paint a continuous potential or series of lattice points [145], or by the superposition of multiple static beams [124]. The pattern quality depends on the input beam profile and AOD response to the pre-programmed input frequency sequence. Although efficiency is reduced by diffraction into other orders, these are sufficiently spatially separated from the trapping potential that they need not disrupt the trapping potential. Furthermore, the output intensity comprises beams diffracted by different AOD frequencies, such that the interference effects outside the trapping plane associated with a single large coherent beam will not arise. An AOD has an acoustic frequency update rate on the order of 10 MHz, although the rate of change of deflection angle is limited by the finite time taken for the acoustic wave to cross the incident beam to a few tens or hundreds of kHz[2]. Despite this limitation, this approach can therefore combine the update rate necessary for smooth dynamic variation of the trapping potential and ability to accurately reproduce a range of target patterns.

Each beam shaping technique is associated with different strengths, and selecting a method depends on the details of the experiment in question. Phase-only SLMs are capable of precisely sculpting a variety of output plane features, including amplitude and phase, in an almost arbitrary manner, although with moderate computational intensity required to do so. They are, however, not ideally suited to creating large continuous intensity distributions or flat-topped features, and their update rate is limited

---

[2] With a sound speed an order of magnitude higher than in an AOD, an AOM could feasibly be used to scan an incident beam with frequencies up to a few MHz.

to a few tens of Hz, which is too slow for some dynamic manipulation applications given the timescales of cold atom systems. DMDs have the benefits of a higher update speed, and the capability to create accurate flat-topped beams by operation in direct amplitude-modulation mode. This is, however, offset against a fairly low power efficiency due to the necessity of discarding portions of the input beam, and a limited range of achievable intensity patterns. The range of painted potentials achieved by scanning an AOM or AOD is somewhat limited in scope, their complexity ultimately restricted by the timescales over which a scan must be completed in order to satisfy the time-averaging criterion. However, the intensity shaping method is implicitly simple and direct, and individual points on the trapping plane can be easily addressed. Painted potentials could be combined with signal multiplexing to increase the versatility of this technique. They are also capable of rapidly updating the effective potential seen by trapped atoms, on timescales up to a few hundred kHz.

In this experiment, we require a smooth trapping beam, with the capability to deform a circular trap into an ellipse to induce rotation in the trapped atoms. Smoothness can be achieved by time-averaging, requiring the beam to be scanned at a frequency much faster than the timescales of motion of the atoms in the trap, in our case with a target frequency on the order of 50 kHz. In light of this requirement on update rate, and the relatively simple and cyclical nature of the potentials of interest, the AOD painted potential method was chosen to create the trapping potential in the rotating dipole trap setup. Methods associated with both time-averaging and signal multiplexing using the AOD are discussed in Chapter 8.

Chapter 4

Experimental apparatus

The core of the experimental apparatus used to produce a Bose-Einstein condensate of ^{87}Rb atoms was constructed as described in Refs. [67, 146, 147]. As well as implementing the hardware necessary to create multiple-RF dressed adiabatic potentials as explored in Sections 5 and 6, and the optical dipole trap described in Sec. 8, substantial modifications and upgrades were made to this core apparatus in order to ensure its continued reliable operation and suitability for future applications.

This chapter explores these modifications and describes the process by which a BEC is produced. First, an overview of the experiment sequence is outlined in Sec. 4.1. Aspects of the apparatus are then considered in turn in Sections 4.2 to 4.4 with a focus on the recent changes. These are then considered in the context of a detailed explanation of BEC production in Sec. 4.5.

4.1 Experiment overview

An overview schematic of the apparatus is shown in Fig. 4.1, and a typical experimental sequence is illustrated in Fig. 4.2.

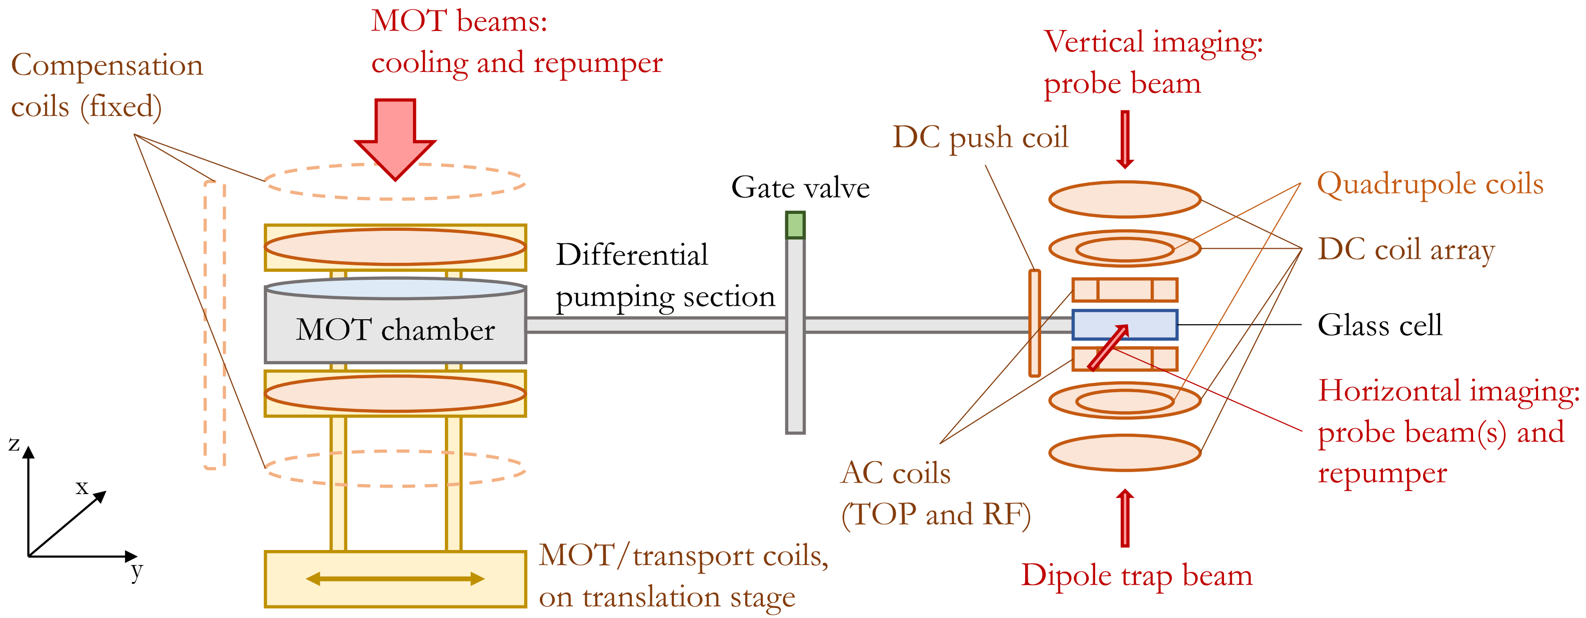

Figure 4.1: Schematic of the experiment apparatus. Atoms are dispensed and laser cooled in the pyramid MOT chamber, before being transported through a differential pumping tube to the glass cell by mechanical translation of the quadrupole MOT coils. A stage of magnetic transport completes their transfer to a magnetic quadrupole potential in the glass cell. Forced RF evaporative cooling takes place in the quadrupole, and continues in the time-orbiting potential (TOP) created by applying a bias field rotating at 7 kHz. Following evaporative cooling, atoms can be transferred into the RF dressed shell trap created by applying a field oscillating at a few MHz. To characterise the resultant distribution of atoms, absorption imaging can be performed in either the horizontal or vertical directions. The vertical imaging system is coupled with a red-detuned beam that will in future be used to create a painted dipole potential on the lower surface of the RF-dressed trapping shell.

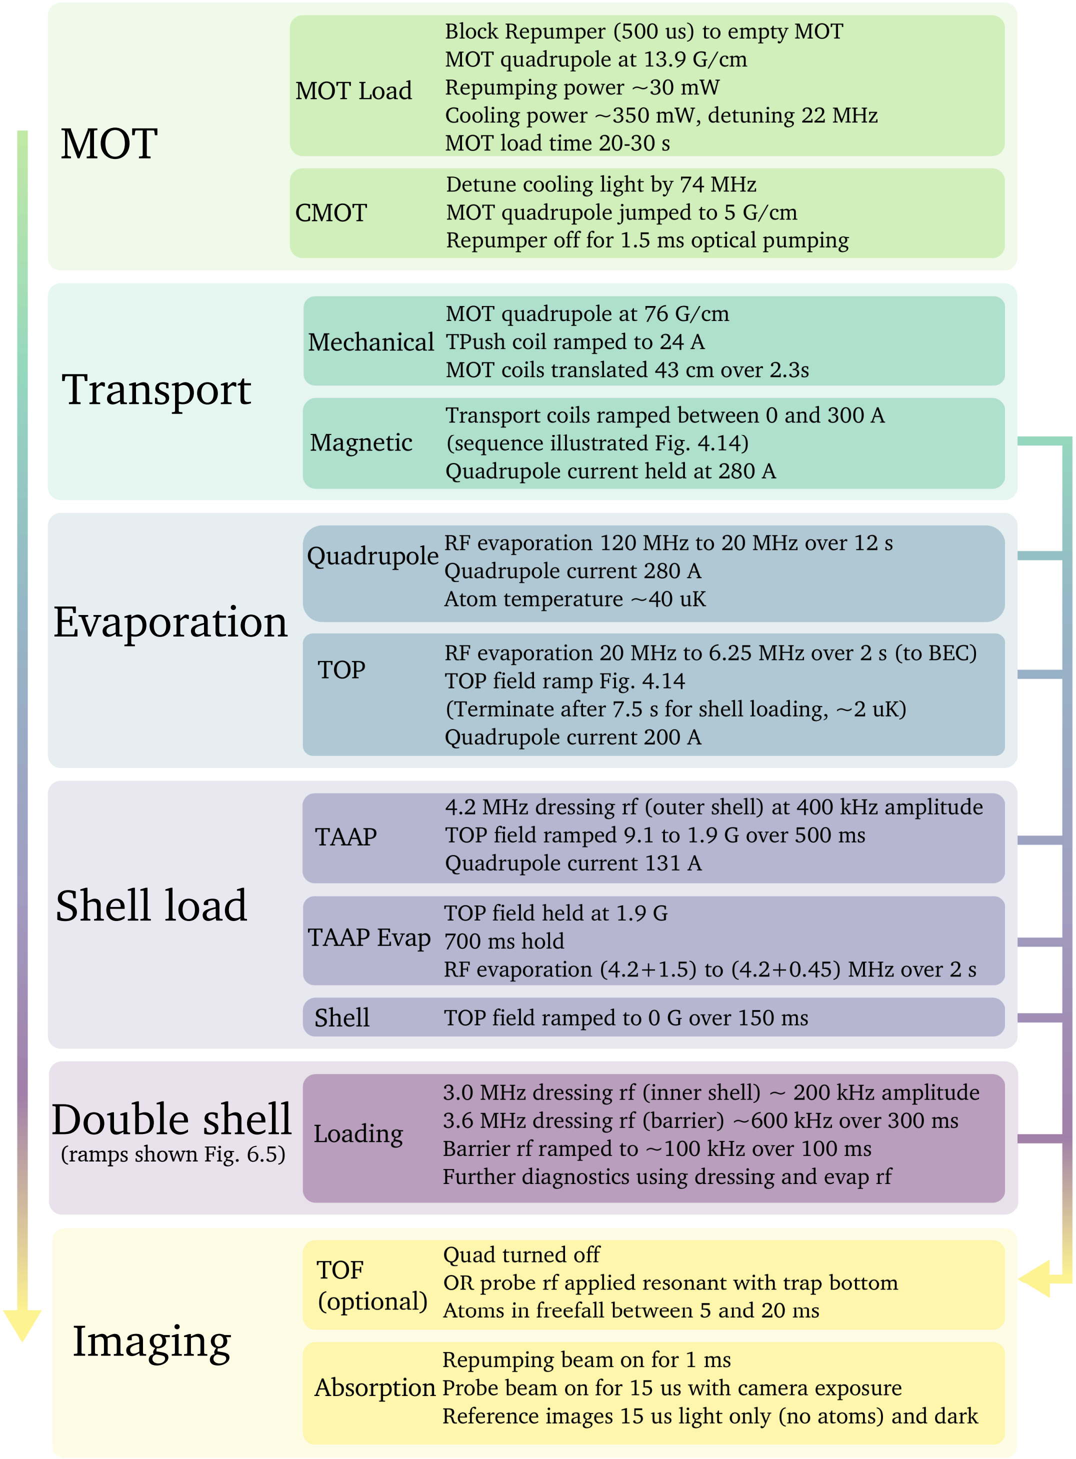

Figure 4.2: Flowchart summarising the stages of a typical experimental sequence. Atoms can be released for time of flight imaging at any point after transport to the glass cell.

This procedure can be broken into several component stages: the MOT; hybrid mechanical/magnetic transport; evaporative cooling in the quadrupole and TOP traps, either to BEC, or to a thermal cloud that can be loaded into the RF-dressed shell potential; and finally absorption imaging of the atomic cloud, generally following some time-of-flight expansion. The computer system used to control these sequences is described in Ref. [147].

The success of these sequences requires the backdrop of a reliably functioning experimental apparatus. This apparatus has several fundamental components: an ultra-high vacuum with the capacity for transporting atoms between chambers; DC coils to provide magnetic confinement; AC coils for evaporation from a laser-cooled cloud to BEC, and to provide the RF fields needed to create the AP trapping surface; and finally a laser system for both laser cooling and imaging. For a visual overview of the system, a series of photographs of the apparatus (excluding the laser systems) is shown in Fig. 4.3.

These system components will be described in detail in the following sections, culminating in a discussion of how these are brought together to create a BEC according to the first few steps of the experimental sequence shown in Fig. 4.2. This chapter therefore lays the foundation for discussions of the RF-dressed potentials in Chapters 6 and 7, and the dipole trapping system addressed in Chapter 8.

4.2 The laser system

A healthy laser system is the bedrock of a cold atom experiment, and is imperative to both the initial step of laser cooling and the final imaging of the atom cloud in this apparatus. Discussion of the laser system has been split across two sections. This first section focusses on the structural aspects of the system: the requirements on the laser light, and how it is generated and stabilised in preparation for its application to the atoms. The subsequent discussion in Sec. 4.5 will describe the experimental implementation of this light to the relevant stages of BEC production, namely its role in the MOT and the atom imaging systems. Description of the dipole trapping laser

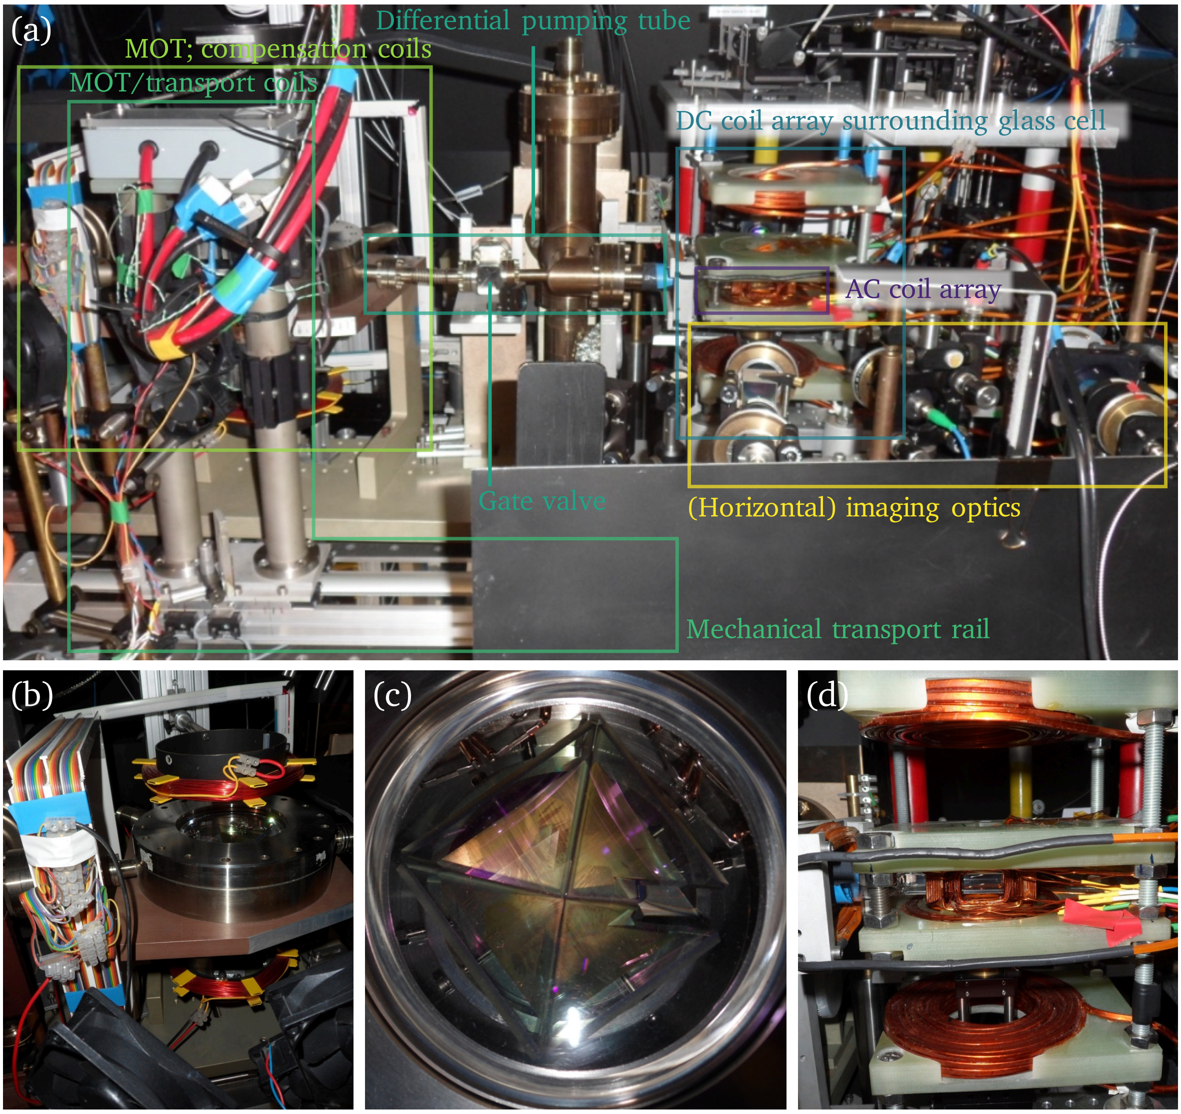

Figure 4.3: Photograph of the experiment apparatus. (a) Overview of the vacuum system, magnetic trapping coils, and imaging optics. Atoms are laser cooled in the MOT chamber (far left) and transported to the glass cell using a combination of mechanical transport of the MOT quadrupole trapping coils, and a final magnetic transfer stage. The pressure differential between MOT chamber and glass cell is maintained using a differential pumping tube. The glass cell is surrounded by a DC coil array to facilitate magnetic transport and subsequent confinement in a quadrupole trap, and an AC coil array that provides the 7 kHz TOP bias field, and both dressing and evaporative RF fields. Both horizontal and vertical imaging, and imaging lattice optics, populate the tiered optical breadboards surrounding the glass cell. (b) The pyramid MOT chamber without the quadrupole/transport coils. The coils provide a controllable bias field to tune the atom cloud position. Cooling optics are mounted vertically. (c) The pyramid MOT mirrors inside the chamber. Rubidium dispensers are visible above the top left and top right mirrors. Atom transport is conducted in the direction indicated by the gaps in the mirrors on the right hand side. (d) The glass cell, surrounded by DC and AC coil arrays. The AC coils fit snugly around the glass cell.

setup is omitted from this section as it is difficult to disentangle the system from its specific design requirements: this is discussed in Chapter 8.

4.2.1 ^{87}Rb transitions

The ^{87}Rb D2 line ($5^2\mathrm{S}_{1/2}$ to $5^2\mathrm{P}_{3/2}$) with wavelength 780 nm provides the atomic structure exploited during the cooling and imaging processes. The hyperfine structure associated with these levels is illustrated in Fig. 4.4, with the cooling and repumping transitions indicated.

Laser cooling is based on the dissipative scattering force as discussed in Sec. 3, whereby photon absorption and isotropic re-emission results in a net momentum transfer that opposes the direction of atom motion [2, 149]. The necessary directional momentum transfer requires an imbalance in absorption probability according to the motion of the atom relative to the direction of beam propagation. This is achieved using the Doppler shift: by detuning to the red of the cyclic cooling transition we ensure that the absorption probability is significant only for atoms moving counter to the direction of beam propagation. Efficient cooling requires a cyclic transition, with a minimal probability of decay to a level inaccessible to the cooling laser. The best candidate for this is the $|\mathrm{F}=2\rangle \to |\mathrm{F}'=3\rangle$ transition. We have empirically chosen to detune the cooling beam from this transition by a value of 22 MHz, or approximately 4Γ, where $\Gamma = 2\pi \times 6.1$ MHz is the natural linewidth of the transition [148].

However, a finite probability exists for the cooling beam to excite a transition into $|\mathrm{F}'=2\rangle$, from where it has a 50% chance of decay into the $|\mathrm{F}=1\rangle$ ground state, where the atom is several GHz from resonance and thus blind to the cooling beam. Over several cooling cycles, this optical pumping effect would result in the loss of a significant number of atoms from the cycle, severely restricting the laser cooling process. A repumping beam is therefore employed to retrieve atoms from the dark state, returning them to the $5^2P_{3/2}\,|F=2\rangle$ excited state to give them a chance of de-excitation back

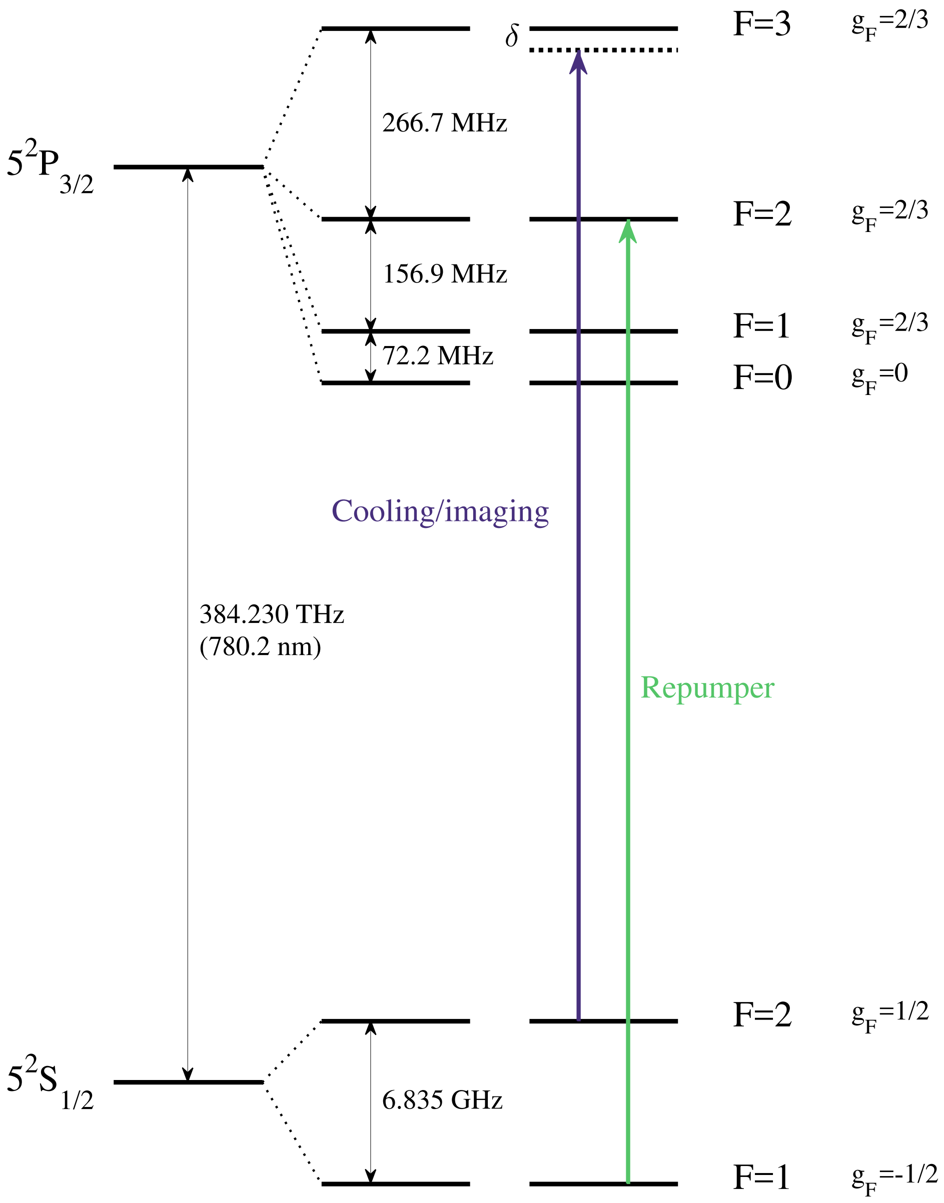

Figure 4.4: Hyperfine structure of the ^{87}Rb D2 line. The magnitude of the g_F factors $g_F = 2/3$ and $g_F = 1/2$ correspond to Zeeman splittings between adjacent sublevels of 0.93 MHz/G and 0.7 MHz/G respectively. Energy level data was taken from Ref. [148].

into the $|F = 2\rangle$ ground and back into the cooling cycle.

4.2.2 Laser frequency stabilisation

Diode lasers are exclusively used on this experiment, due to a combination of factors including low cost, ease of use, straightforward tuneability of the output wavelength and a narrow linewidth once locked to an appropriate signal. The free-running linewidth of these laser diodes is typically on the order of a few tens of MHz. For applications in both cooling and imaging, this must be narrowed to less than the natural linewidth of a typical ^{87}Rb excited state transition, corresponding to approximately $2\pi \times 6$ MHz for the D_2 line [148].

For each of the cooling and repumping lines, a 'master' laser, its frequency controlled by an external cavity to form an external cavity diode laser (ECDL), is locked to a signal derived from ^{87}Rb spectroscopy. Meanwhile, injection-locked [150] 'slave' lasers provide a power boost for the cooling light. In our setup, acousto-optic modulators (AOMs) are used to vary the input beam frequency to control the detuning from atomic resonance.

4.2.2.1 Laser diode control

The broad linewidth of a free-running laser diode is artificially narrowed by the injection of a narrow, stabilised beam derived either from some portion of the output light stabilised to a cavity (the master ECDLs) or a separate stabilised laser source (the injection-locked slave lasers). This externally-derived light then encourages single-mode, narrow-linewidth operation by the process of gain saturation: amplification of light by the gain medium over this small frequency range deprives alternative modes of this gain, thus preventing them from lasing, and itself asserting control of the laser output.

In the Littrow ECDL configuration used in this experiment, the feedback frequency is set by a diffraction grating, aligned such that its first diffracted order is reflected directly back into the diode, with the grating angle tuneable either by piezo or manually (for the initial coarse alignment stage). Since the grating acts to spatially separate wavelength components of a polychromatic source according to the relation $n\lambda = 2d\sin(\theta)$ with n the diffraction order, d the grating constant and θ the diffraction angle, this determines which spectral component is aligned back into the diode to force single-mode operation. The laser frequency also has a dependence on the external cavity length L, still determined by the piezo controller, with the mth resonant frequency of the cavity given by $\nu_m = mc/2L$ and arising from the condition for standing wave formation in the cavity. The output facet of the diode is anti-reflection coated to minimise competition between the mode associated with the internal cavity established by reflections and that enforced by the wavelength-selective external cavity. With careful grating alignment, an ECDL setup can produce linewidths of around 100 kHz [151], well below the MHz atomic transition linewidth, as averaged over the tens of ms associated with the bandwidth of the laser control servo.

The injection-locked slave laser setup is far simpler: a low-power, carefully mode-matched frequency-stabilised beam derived from the ECDL is aligned into the anti-reflection coated output facet of a laser diode. This 'seed' beam saturates the gain medium at its own operating frequency, encouraging single-mode operation of the slave laser that follows any variation in the seed beam wavelength within the constraints of the diode output spectrum.

In both ECDL and injection-locked laser systems, the careful positioning of mode-hops is necessary to ensure single-mode operation throughout the frequency sweeps implemented during an experimental sequence. Various lasing modes can exist within a cavity, all with some degree of spatial overlap and competition for the available gain. The gain associated with each mode is determined by a combination of factors including the intrinsic gain lineshape of the semiconductor gain medium, the properties of the cavities formed between front and back facets of the diode and the grating, and the frequency spectrum back-diffracted into the diode from an external cavity grating. Competition between these effects gives rise to mode-hops: the net gain associated with a number of modes can be very similar in value, such that a small perturbation can diminish the gain of the original dominant mode and empower another, causing a jump in the lasing mode exhibited. Once a given mode has won the competition for gain, it is able to seize control of single-mode operation by the same gain saturation process exploited to lock the laser. This effect can be diminished by an incomplete spatial overlap of the modes, and even a well-designed system can be sensitive to differential response of the various cavities to temperature drifts. It is therefore possible that due to imperfect alignment or other drifts in the system a different mode becomes able to dominate the gain profile, resulting in a mode-hop.

The location of mode-hops can be adjusted using the diode temperature and current, in addition to realigning the external cavity or seed beam if necessary. The temperature affects the length of both internal and external cavities due to thermal expansion, affecting both the mode frequency and spacing, while the current affects the refractive index of the semiconducting gain medium in the diode and thus its effective length and corresponding frequency [152]. The influence of the laser current over the output frequency can be exploited for frequency stabilisation, by splitting the current signal used to control the piezo and diverting some proportion of this to directly modulate the diode current in a feed-forward configuration for additional fast control that reduces the resultant linewidth.

4.2.2.2 Laser locking to an atomic spectrum

Tuning the frequency of a laser diode is therefore fairly straightforward. However, this frequency must also be stabilised for the duration of the experimental sequence. Locking a laser to a given frequency requires a discriminator signal that is related to the difference between the set-point and measured frequency. This discriminator signal can be derived from a variety of sources, including spectroscopic measurements, cavities and electronically via a frequency offset locking system [153].

To probe an atomic transition using spectrosocpy, the angular frequency ω of a laser beam propagating through some atomic sample is scanned around the atomic resonance ω_0 to map the absorption. However, in this simplest configuration, the transition linewidth is broadened by the Doppler effect associated with thermal motion of atoms in a room temperature vapour cell, obscuring features smaller than the Doppler width of around $1\,\mathrm{GHz}$. The saturated absorption technique [154] is a popular approach to obtaining the few hundred kHz laser lock required. In this technique, the probe beam that propagates through the vapour cell, and whose absorption profile is detected on the photodiode, is joined by a counterpropagating pump beam strong enough to saturate the atomic transition for those atoms in the velocity class it interacts with. This saturation is still experienced by the atoms that are also resonant with the weaker probe beam, i.e. those with a Doppler shift of zero within the natural linewidth Γ, such that their ability to interact with this beam is reduced. By thus limiting the interaction to those atoms with close to zero velocity, the effects of Doppler broadening are eliminated and narrow peaks of reduced absorption, width Γ, are created on the broad Doppler intensity profile [63].

The Doppler-free absorption spectrum for the ^{85}Rb and ^{87}Rb D2 line is shown in Fig. 4.5. The four broad dips in the spectrum correspond to the Doppler-broadened transitions from the $5^2\mathrm{S}_{1/2}|F=1\rangle$ and $|F=2\rangle$ states in each species, and the small peaks superimposed on this profile the isolated transitions between these states and specific components of $5^2\mathrm{P}_{3/2}$.

However, more sub-Doppler peaks are evident than transitions exist. This is due to the crossover peaks that fall midway between each pair of transitions. These peaks arise due to the sensitivity of the probe resonant with atoms on one transition to those same atoms' saturation by the pump on another. If the pump is Doppler-shifted into resonance with atoms with angular frequency $\omega_1 = \omega + kv$, and the probe with $\omega_2 = \omega - kv$, then this crossover saturation will occur at $\omega = (\omega_1 + \omega_2)/2$ where the atoms in this velocity class are able to experience the effects of saturation, although exhibiting them on the unintended transition when probed [63]. Crossover resonances can be comparable to or larger than the transition peaks, and as such can either provide a useful alternative reference point, as in our repumping system, or a problematic

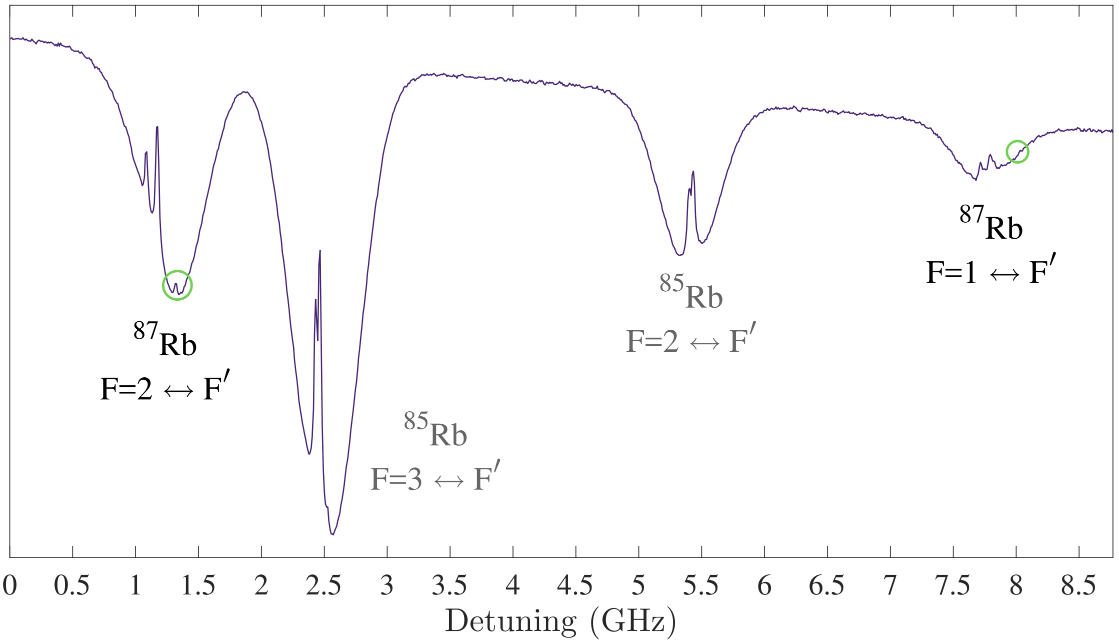

Figure 4.5: Doppler-free absorption spectrum over a 9 GHz frequency scan, showing the D2 lines of ^{87}Rb and ^{85}Rb. The Doppler peaks for the ^{87}Rb $|F=1\rangle \leftrightarrow |F'\rangle$ and $|F=2\rangle \leftrightarrow |F'\rangle$, and the ^{85}Rb $|F=2\rangle \leftrightarrow |F'\rangle$ and $|F=3\rangle \leftrightarrow |F'\rangle$ are visible. Sub-Doppler features are visible within these broad peaks, including both real transitions and crossover peaks. The sub-Doppler peaks for the cooling and repumping lines are circled in green.

source of background noise diminishing the stability of an attempted lock to a nearby peak.

Modulation transfer spectroscopy (MTS) is a simple extension of the saturated absorption method that creates a low-noise, zero background discriminator signal on which to lock the laser frequency. Already an established technique for high-resolution spectroscopy [155, 156], this was first applied to laser locking using the caesium D_2 transition [157]. In addition to a broad application in many cold atom experiments, MTS has since been characterised for rubidium [158], and the potassium D2 line [159]. As in saturated absorption, the MTS setup is based on a pump and probe beam counterpropagating through a vapour cell. However, in this case both beams are of approximately equal intensities, and the pump is modulated to produce sidebands at some frequency comparable to Γ. This modulation is typically applied using an electro-optic modulator (EOM), although acousto-optic modulation has also been successfully demonstrated in Rb spectroscopy [160]. When both pump and probe beams interact sufficiently strongly with the atoms in the vapour cell, four-wave mixing can occur, which transfers the modulation onto the probe beam. The conditions for this nonlinear interaction include sufficient intensity in both beams, and a closed transition that ensures sufficient efficiency of the modulation process [63].

The modulation transfer process occurs only at frequencies where the four-wave mixing criterion is fulfilled, and a high-gradient discriminator signal can be generated by detecting those frequencies at which the sidebands are present on the probe beam. The creation of this discriminator signal is described in the context of this experimental setup in Sec. 4.2.4. MTS signals for the Rb cooling transitions are plotted in Fig. 4.6.

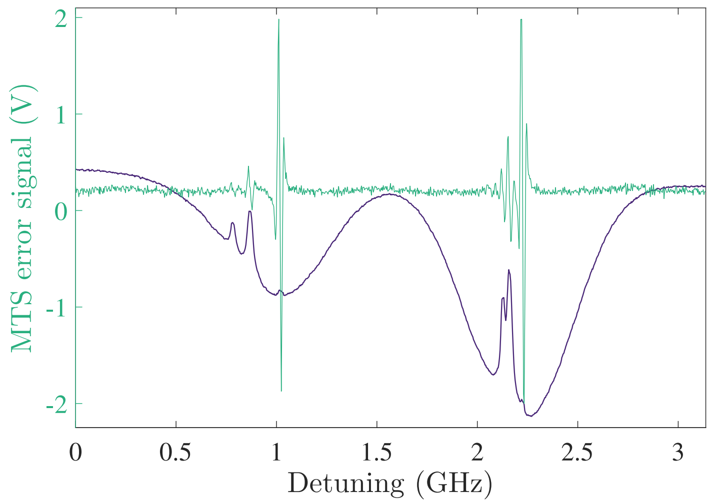

Figure 4.6: Modulation transfer spectroscopy on the ^{87}Rb and ^{87}Rb cooling lines. Purple traces show the spectra and green the corresponding MTS discriminator signal. The key MTS characteristics are apparent here: large-amplitude, high-gradient error peaks associated with the closed transitions $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ (^{87}Rb) and $|F = 3\rangle \leftrightarrow |F' = 4\rangle$ (^{85}Rb), set against a uniform background and significantly lower-amplitude peaks for the other large spectroscopic features. This peak provides a precise, stable lock point.

Unlike the Pound-Drever-Hall method [161] used with the saturated absorption repumper signal, modulation is restricted to the spectroscopy beams and so the main laser output remains free from these effects. This is particularly important for the cooling light, where the beams are detuned from resonance by 22 MHz in the standard case, and up to 74 MHz during the CMOT stage; sidebands with tens of MHz spacing would risk resonance with the atoms during these sweeps. Another prominent characteristic of MTS is the presence of error signal peaks only for closed transitions. This reduces the likelihood that the lock will jump to a neighbouring transition, or that a preferable nearby lock point exists. The drawback of this is that not all desirable lock points possess a corresponding MTS peak. For example, too many de-excitation routes inherently exist for the ^{87}Rb repumping transition. However, a ground state crossover can act as a suitable closed transition if the ground state energies are close enough together to be thermally occupied. This ground state spacing determines the velocity class that is dually selected in the crossover interaction: this velocity is given by $(\omega_1 - \omega_2)/2k$, and for the ^{87}Rb ground state is on the order of $2500\,\mathrm{m\,s^{-1}}$. In contrast, with a smaller ground state splitting and corresponding crossover velocity of just $190\,\mathrm{m\,s^{-1}}$, an MTS signal has been observed on the potassium ground state crossover [100].

4.2.3 Laser system overview

We have constructed parallel laser systems that are frequency-stabilised to the cooling and repumping transitions, with variable detuning achieved by acousto-optic modulation. These systems were designed based on two paradigms: stability and modularity.

In aid of stability, we favoured methods with minimal degrees of freedom. A noteable example of this is the tapered amplifier (TA) outlined in Sec. 4.2.6, where, despite the short-term alignment challenges, this approach has proved worthwhile in terms of the significant decrease in ongoing realignment and maintenance[1].

The modular approach is one in which each system is split into subsystems that each perform a specific task. Each subsystem is contained within its own breadboard, coupled using optical fibres, to minimise the impact of any upstream maintenance and

---

[1] A new TA chip housing style and robust alignment optics has reduced maintenance from 2-3 times a week with the old design, to approximately once every 6 months.

allow both offline system development and easy transport between locations[2]. In its previous incarnation, the laser system had been constructed entirely on the same optical table as the rest of the experiment apparatus. Aside from spatial constraints, this caused stability issues such as frequent unlocking of the repumper during the mechanical transport stage, and of the cooling MTS due to MHz noise during RF evaporation. The flexibility of the modular approach outweighs the drawbacks of the losses and polarisation drifts associated with fibre coupling. To combat the former, an additional slave laser was added to the cooling system, granting the benefit of an extra injection-locked laser to share the frequency shift during the CMOT stage. Polarisation drifts due to imperfect matching of the polarisation axis of the fibre with linearly polarised incident light can be minimised by careful alignment, and their effect mitigated by careful choice of the break points between subsystems. If necessary, active power stabilisation could be introduced.

The layout of the laser systems are outlined in the following sections and series of figures. This has recently been developed further to incorporate frequency offset locks for both cooling and repumping light from a single ECDL locked to the ^{87}Rb cooling transition, and will be discussed in Ref. [35].

4.2.4 Cooling light

The cooling system is contained within boards C1-C5, illustrated in Fig. 4.7, with amplification implemented by a tapered amplifier (TA) as shown in Fig. 4.9 and the distribution into the imaging system shown in Fig. 4.16.

Cooling light is ultimately derived from a Toptica DL PRO 780, providing a maximum output power of 80 mW between 765 and 800 nm with current and temperature control provided by the DCC 110 and DTC 110 modules respectively. Following a stage of isolation and beam shaping (C1), this is immediately fibre coupled onto a dis-

---

[2] This flexibility has been a particularly useful defence against nearby building works.

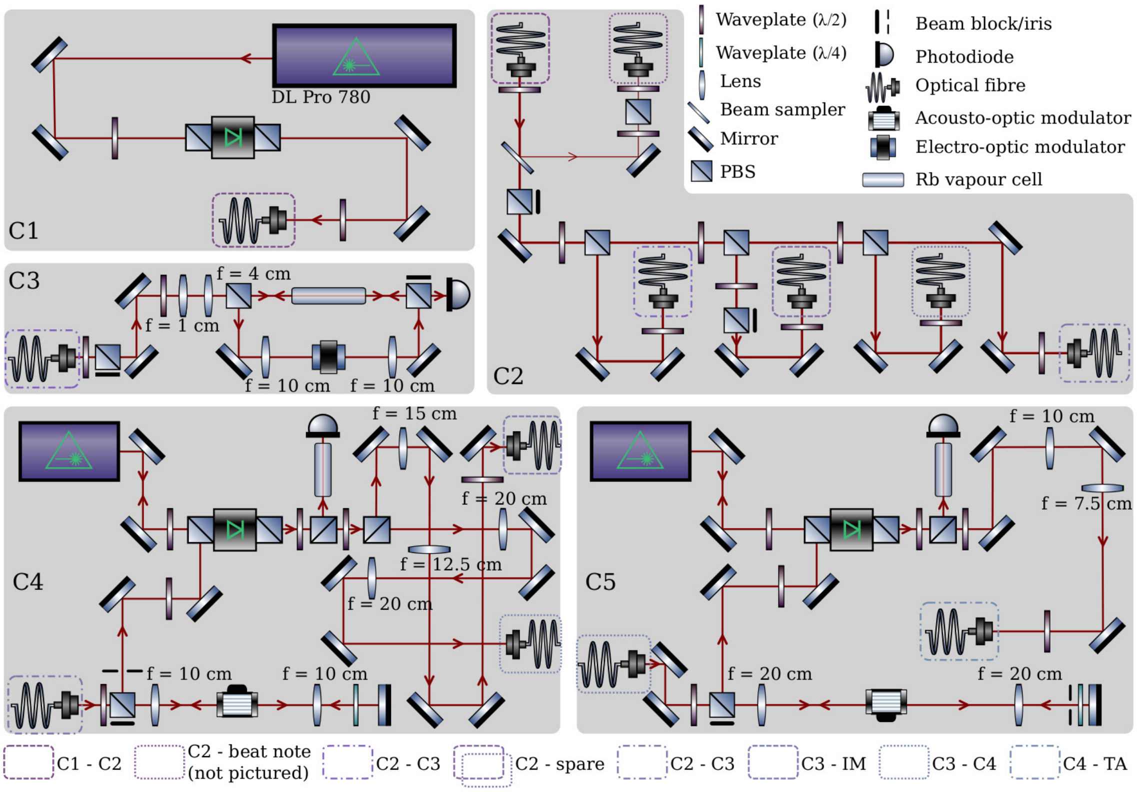

Figure 4.7: Schematic of the cooling laser setup. Laser light is derived from a DL PRO (C1), and immediately transferred to a distribution board (C2). Fibres then transfer light to the spectroscopy setup (C3) and onwards for amplification. Additional fibre outputs can provide laser power for other experiments and for beat-note detection in an offset-locking scheme [35]. A first stage of amplification is performed using two injection-locked lasers (C4, C5). The seed beam frequency in each case is controlled using an acousto-optic modulator (AOM), with a frequency shift of order 100 MHz imparted in opposite directions on each board to give a variable cooling beam detuning tens of MHz from the cooling transition. Outputs from board C4 lead to the second injection-locked laser on C5, and directly to the imaging system shown in Fig. 4.16 where another AOM allows independent detuning of cooling and imaging light. The output from board C5 is amplified by the TA in Fig. 4.9. The component key carries over to subsequent diagrams, and colour-coded boxes around the optical fibres match in pairs to illustrate fibre destinations. This and subsequent optics diagrams are created using components adapted from the Component Library [162].

tribution board (C2) which was constructed to allow for the extension of the system to accommodate the frequency offset locking system and allow the DL PRO to act as a stable, narrow-linewidth laser source for multiple experiments. The laser is locked to the cycling transition using an error signal derived from modulation transfer spectroscopy (C3) as shown in Fig. 4.6, facilitated by the Digilock 110 in which PID values can be digitally set, controlling the cavity via the SC 110 piezo scan module.

To create the MTS discriminator signal, the modulated probe beam is incident on a photodiode with a sufficiently fast response to detect the beat pattern between probe and sidebands. This signal is then amplified (*Mini-Circuits ZFL-500HLN*) and mixed with a phase reference derived from a second function generator to perform phase sensitive demodulation. By low-pass filtering (*Mini-Circuits BLP-1.9+*) the mixer output, the DC component is selected to form the PID input that determines the piezo signal and thus cavity length and grating angle, locking the signal to the zero-point of the sharp edge of the MTS error signal.

Stabilised resonant light is then taken from the low-noise laser lab and transported by a 21 m fibre (*PMJ-3AF3AF-850-5/125-3AS-21-1*) to the main experiment[3]. Polarisation variations were minimised to within 5% by careful alignment and fibre insulation.

From here, the light is split into cooling and imaging components, each with precise frequency control granted by two complementary double-pass AOMs (*Crystal Technology 3110-140*), acting to give a controllable frequency detuning of up to 90 MHz. The first amplification stage affects both cooling and imaging branches, and consists of an injection locked laser offset from resonance (C4). The original diodes used for the injection-locked lasers were *Sanyo DL7140-201S*, housed in a home-built laser head driven using an in-house current control unit and a temperature controller (*Newport 325*). With seed light of a few hundred μW and good input alignment and mode-matching, these diodes could output around 60 mW with a mode-hop free tuning range of around 2 GHz. After the discontinuation of this line, we have used the 200 mW *Panasonic LNC728PS01WW*[4], outputting 50 mW with a seed power around

---

[3] A dedicated laser lab protects the laser system from both traffic and environmental variations, and also systematic experimental perturbations. When the ECDLs were set up on the same table as the main apparatus, evaporative RF at a few discrete frequencies around 30 MHz would couple into the MTS locking electronics and intermittently kick the cooling laser out of lock; until the new system was built we developed modified ramps that jumped over the problematic frequencies. Vibrations caused by the mechanical transport would routinely destroy the lock of the repumping ECDL.

1.5 mW.

One branch of the light is diverted into the imaging system (IM), and the second, cooling, branch provides the seed beam for a second injection-locked slave laser, having been shifted back to $\sim 4\Gamma$ below resonance (C5). This light is combined with the repumper and directly provides the seed for the TA discussed in Sec. 4.2.6.

4.2.5 Repumping light

As illustrated in Fig. 4.8, the repumping light for both MOT and imaging is derived from a single ECDL (R1), using an *Eagleyard EYP-RWE-0790* diode in a *Toptica DL100* housing[5]. Current and temperature control are provided by the *DCC 110* and *DTC 110* systems respectively, and the piezo scan controlled using the *PID 110* The laser frequency is locked to a vapour cell spectroscopy signal (R2).

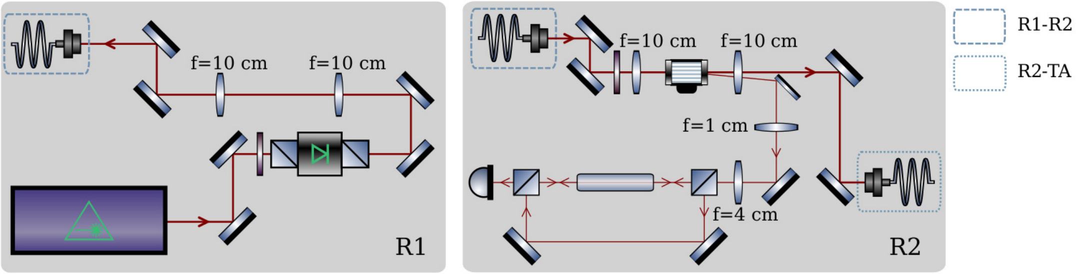

Figure 4.8: Schematic of the repumping laser setup. The *DL 100* laser source (R1) is stabilised by a Pound-Drever-Hall lock using saturated absorption spectroscopy (R2), placing 20 MHz sidebands on the beam. This light is transmitted to the TA board for amplification.

An error signal is derived from saturated absorption spectroscopy using the Pound-Drever-Hall technique implemented using the *Toptica PDD 110* module to modulate the laser driving current at 20 MHz. As the Doppler-free peak associated with the

---

[4] This followed an extended period of trialling a range of alternatives, none of which offered the performance of the original Sanyo diodes.

[5] We have also used EYP-RWE-0780 diodes in the DL100 housing. Anecdotally, while initiating laser feedback using these diodes seemed easier, it was more difficult to subsequently lower the lasing threshold to the same extent as the RWE-0790 diodes. Furthermore, the emission plane is closer to the front surface of the diode. A spacer was therefore required to make these compatible with our housing and collimating optics.

repumping transition is the smallest of the $|F=1\rangle \leftrightarrow |F'\rangle$ transtions, with a correspondingly small error signal, this does not provide an ideal lock point. To maximise lock stability and prevent it jumping to a larger signal, we instead lock on the $|F=1\rangle \leftrightarrow |F'=1/2\rangle$ crossover, using the AOM on board R1 to detune the 78.45 MHz back to resonance. The trimpots that set the PID parameters lack the precision and repeatability of digital control, and eventually required replacement after becoming increasingly unreliable. This rempumping system has recently been modified to derive its locking signal from a frequency offset lock [35].

4.2.6 The tapered amplifier

The tapered amplifier (TA) amplifies both the cooling and repumping beams to provide sufficient power for the large-diameter MOT beam. The associated breadboard is illustrated in Fig. 4.9. The TA chip contains a tapered region of gain medium, anti-reflection coated at both ends, seeded at the narrow input to filter the spatial mode. The tapered structure of the gain medium allows significant amplification of the seed

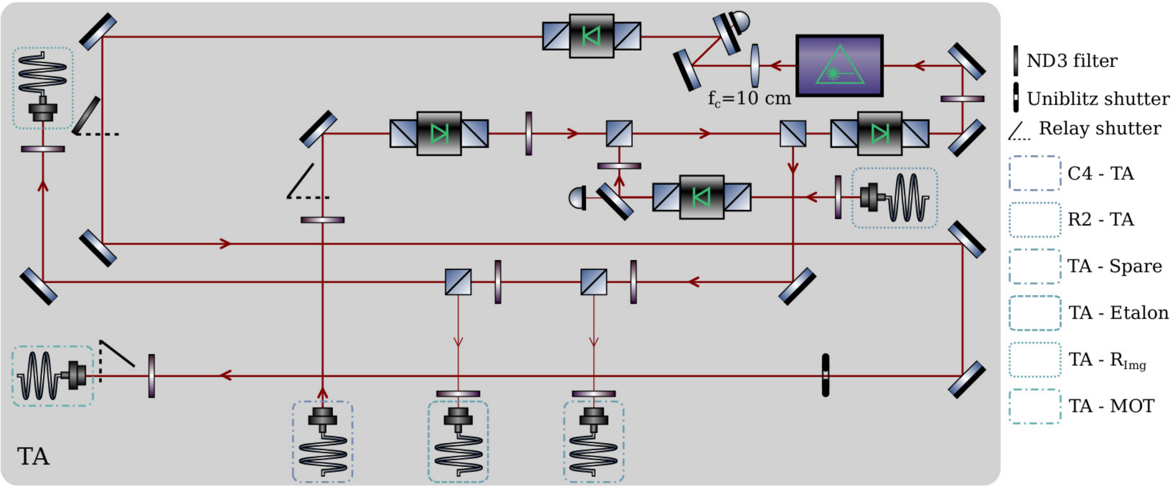

Figure 4.9: Schematic of the tapered amplifier optics. Cooling and repumping beams are combined to seed the TA, and a few mW of repumping light diverted to the imaging system (cooling light is independently shuttered). Additional outputs are used for diagnostic purposes including monitoring with a Fabry-Pérot Etalon. The TA output mode is shaped using a cylindrical lens and free-space propagation to maximise fibre coupling efficiency. This provides the ~ 400 mW of light required for the pyramid MOT.

light without endangering the diode output facet, although this results in a strongly astigmatic output beam.

We use an *Eagleyard TPA-0780-01000-3006-CMT03-0000* chip with a maximum 1 W output, in a home-built monolithic housing. While the previous housing design contained adjustable mounts for the aspheric lenses used to collimate the output beam and back-ejected spontaneous emission to ensure correct focussing of the collimated seed beam onto the TA chip, the new housing approaches a monolithic design in which the lens positions are fixed within the cylindrical aperture around the chip. Initial alignment was painstaking, since the lenses had to be adjusted within a small aperture without risking back-reflections that would damage the TA chip. However, this system has since been completely stable, in contrast to the previous design which required realignment of both input and output beams on a regular basis due to drifts in the lens positions. A Peltier cooler (*RS 490-1424* with maximum 36.3 W of cooling power) and temperature control unit (*Thorlabs PRO 8000-4*) are used to stabilise the TA diode temperature, while current is regulated at 1.9 A using this same control module.

Cooling and repumping light are combined as shown in Fig. 4.9 and aligned to the TA chip using the reverse-emitted spontaneous emission as a guide. The output beam is both highly astigmatic and non-Gaussian; alignment into an optical fibre requires collimation of the horizontal profile using a cylindrical lens, and a ~ 1.5 m propagation path to allow for a degree of ‘self-cleaning’ of the transverse mode. Even with careful mode-matching the coupling efficiency peaked at 67% using a standard polarisation-maintaining fibre. Additional outputs on this board collect stray light for use in diagnostics, and to provide the repumping light used in imaging. Photodiodes are used to monitor any power variations and could be used in active power stabilisation if necessary.

4.3 The vacuum

The vacuum consists of three main sections: the MOT chamber, differential pumping tube and glass science cell. Details of its dimensions and construction are discussed in [67, 146, 147] and a schematic is shown in Fig. 4.10.

The MOT chamber is the starting point for any given experimental sequence, and, housing the rubidium dispensers, is the highest-pressure and most variable region of the vacuum system, with a typical pressure around 10^{-9} mbar. Meanwhile, with good optical access, the glass science cell is maintained at a vacuum of approximately 10^{-11} mbar, protected from the higher pressure in the MOT chamber by the differential pumping tube. Vacuum is maintained on both sides of the vacuum apparatus by the two ion pumps, one attached to the MOT chamber and one located between the gate valve and glass cell, and a non-evaporable getter (NEG) pump located between gate valve and glass cell that removes atoms by adsorption to the getting surface. A gate valve protects the glass cell against changes to the MOT chamber.

4.3.1 The rubidium dispensers

Rubidium is dispensed into the experiment by means of a pair of dispensers (*SAES Rb/NF/3.4/12 FT10+10*) driven in series by a constant current source. These Rb dispensers are located in the MOT chamber, behind the pyramid mirrors, with two such pairs installed in 2008. Until November 2014, the original dispenser pair was driven by a constant current of $2.9\,\mathrm{A}$ throughout a full day of experiment operation. Within 30 minutes to an hour, the vapour pressure in the MOT chamber, as monitored by the MOT size itself, would stabilise to give a loading time of approximately $20\,\mathrm{s}$ with a final atom number of 3.9×10^9.

The optimal vapour pressure is a compromise between loading a sufficiently large MOT in a reasonable time, and the transport efficiency allowed due to collisions with background atoms. Assuming no contamination in the chamber, a good work-

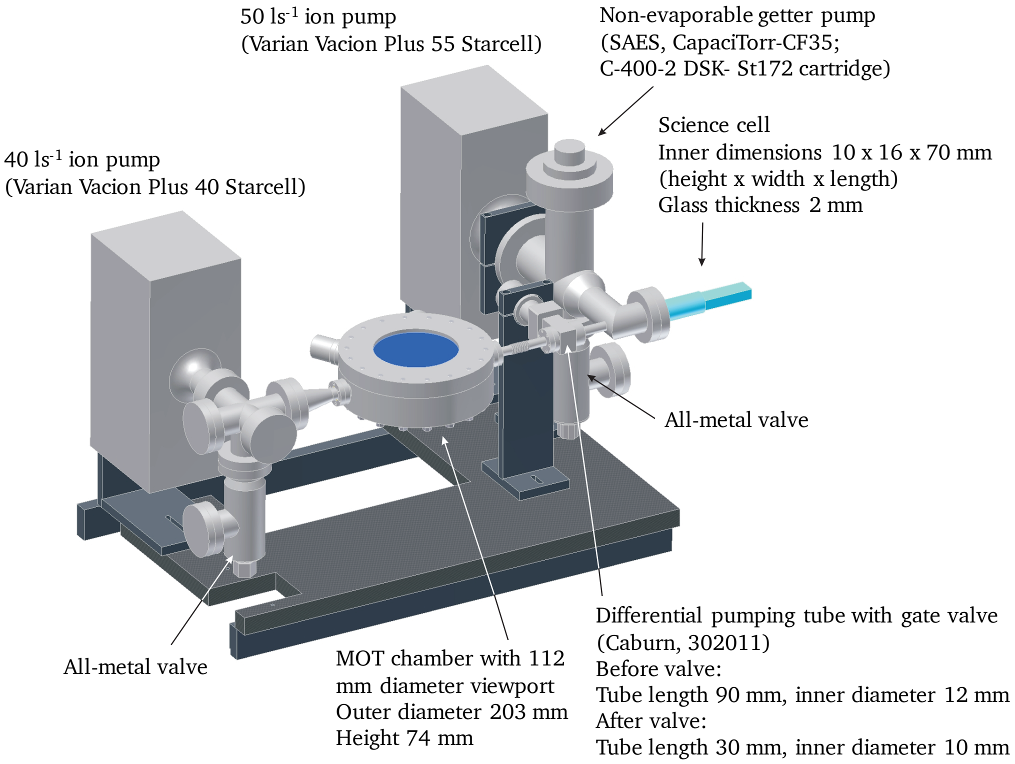

Figure 4.10: Principal components of the vacuum system: the MOT chamber, differential pumping tube, and glass science cell. Image reproduced from [146].

ing vapour pressure range has been found empirically to correspond to loading rates between 15 and 30 s to a maximum MOT size of around 4×10^9 atoms. The transported cloud is prohibitively small with loading less than 10 s.

However, by 2014 there was a serious unreliability in experiment operation that was, in part, traced back to these dispensers. The atom transport efficiency was unreliable, and would degrade significantly throughout the day. Furthermore, the MOT size itself would fluctuate in a surprising manner: within the 30 minutes of turning on the dispensers, the MOT number would gradually rise as the rubidium pressure in the chamber increased, but would then fall off to approximately 80% of this maximum value over the course of a few minutes. A reasonable hypothesis is that this fall in atom number was due to the emission of contaminant atoms or molecules, raising the background pressure in the MOT chamber and causing heating. As there was no associated fall in loading rate, and due to the short timescales of the response, this was not attributed to a reduction in Rb release as the dispensers heated. Different modes of operation were trialled, including application of ultraviolet light in a light-induced atomic desorption scheme [163, 164], and pulsed operation of the dispensers with various currents and pulse periods. However, the fundamental unreliability remained, and we switched to the pre-installed secondary dispenser pair in November 2014.

These were initialised by running at $8\,\mathrm{A}$ for $10\,\mathrm{s}$ to remove the oxidised layer that would have been a remnant of the manufacturing process as well as any contaminant residue. To prevent an accumulation of atoms in the differential pumping tube or high-vacuum glass cell during this period of significant rubidium emission and corresponding rise in vapour pressure, the gate valve was closed and a partial bake performed following the method described in Sec. 4.3.2. This bake was maintained for several days until the vapour pressure had reduced to a reasonable level[6].

The mode of operation used for the previous dispensers was untenable for the new ones; even a current of less than $1\,\mathrm{A}$ drives the vapour pressure too high if run continuously throughout experiment operation, and these low currents are associated with a higher contaminant fraction in the emitted vapour. After experimenting with a few different methods of running the dispensers, the optimum for stable operation was found to be to run them at $3.8\,\mathrm{A}$ for approximately 10 minutes each evening, adjusting on a daily basis to any observable build-up or decline in pressure. This causes an initial spike in the vapour pressure, which settles overnight to provide a stable loading pressure from which to work throughout the day. Responding to the observable pressure on a daily basis, this has been gradually changed over time, such that by November 2016 the run time had been increased to 14 minutes each evening to prevent a decline in pressure over the afternoon. Towards the end of November 2016, this was changed to the empirically equivalent approach of running at $4.2\,\mathrm{A}$ for

---

[6] The rate of vapour pressure falloff in the MOT chamber was so slow that this provided perfect experimental conditions for a little over a week following the initial firing.

5 minutes.

4.3.2 System maintenance: a partial bake

As a precaution when firing the new dispensers, a selective bake of the differential pumping tube was performed to prevent the high vapour pressure extending into the tube and posing a heating risk to atoms during transport. This baking method has been repeated on several other occasions, as both a maintenance and precautionary procedure, and is illustrated in Fig. 4.11. To minimise disruption to the experiment, the bake was localised to the region of the pumping tube between the gate valve and MOT chamber. The gate valve was closed to avoid atoms being driven into the glass cell, and a temperature gradient established outwards from this point, with the pumping tube to the MOT side of this point being heated to at least 75°C. Care was taken to ensure that there was no significant temperature gradient across the MOT chamber viewport, and that the pumping tube was no higher than 30°C prior to its connection to the glass cell.

Spatial control over the baking process was achieved using two rolls of heater tape, powered by separate variacs and wrapped in foil to control the regions heated. The first tape was used exclusively to wrap the gate valve, which was heated separately for 2 hours to ensure that the interior was completely heated through and establish the temperature gradient, preventing atom accumulation in this region. The second roll of heater tape was used to cover the remainder of the differential pumping tube, with multiple layers necessary to sufficiently heat the connection to the MOT chamber. This region was heated over the course of several hours, with the temperature gradient maintained by careful layering of the tape and occasional adjustments to the variac output current. During dispenser firing, the bake was maintained for the 3 days it took for the MOT chamber vapour pressure to approach our standard operating pressure, while for general maintenance purposes it was implemented for just 36 hours at higher

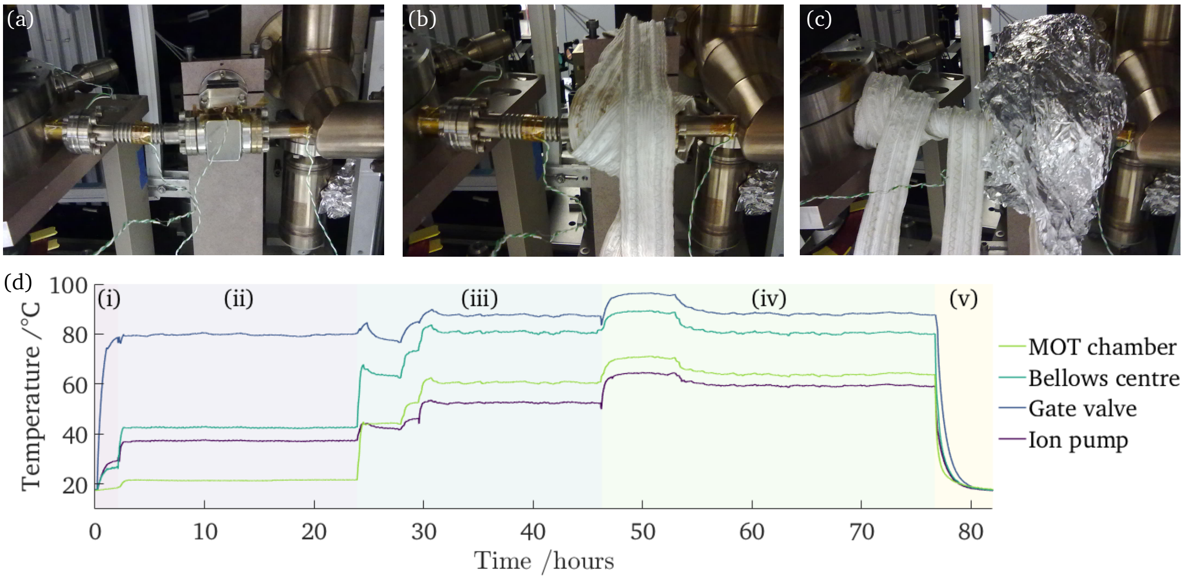

Figure 4.11: The differential pumping tube baking method. (a) Thermocouple positions: close to the junction between MOT chamber and differential pumping tube; 5 cm to the left of the gate valve, in the centre of the bellows section of the tube; at the gate valve itself; and 5 cm to the right of the gate valve, just before the differential pumping tube widens at the position of the ion pump. (b) The coverage of the gate valve using the first roll of heater tape. (c) Full coverage of the differential pumping tube in the second roll of tape. This was subsequently completely covered in foil. (d) Temperature variation at each monitored location during the bake performed for the dispenser change, with key stages of this heating process highlighted. In (i), only the gate valve is actively heated for a period of 2 hours, allowing this to completely heat through and ensure that the necessary temperature gradient is established. In (ii), the second variac is switched on and temperatures allowed to stabilise, and in (iii) the variacs are altered in tandem, and insulation of the system adjusted to ensure adequate heating of the junction with the MOT chamber. The temperatures are unilaterally increased for the final period of the bake in (iv) before both variacs are switched off in (v) and the system allowed to gradually cool with the insulation still in place. Short-timescale variations in temperature are due to adjustment of the foil or tape positions during heating. Subsequent bakes were identical in character, but shorter in duration (up to about 36 hours) and with a higher maximum temperature of around 120 °C.

temperatures of up to 120 °C. The heater tape and foil were left in place to maintain the correct temperature gradient during cooling.

4.4 Magnetic trapping

The magnetic trapping system comprises both static and AC fields. DC fields are used for the MOT, magnetic transport, and quadrupole trap. Once confined in the magnetic quadrupole, an AC coil array is employed to generate the TOP field, evaporative RF, and dressing RF fields.

4.4.1 MOT/transport coils

The MOT field is created using a coil pair connected in a quadrupole configuration, i.e. with current flowing in opposite directions in each coil. These are driven in series by a constant-current power supply (Xantrex XFR 20V- 60A) and their operation is described in previous work [67, 146, 147]. Supplementary bias fields are applied to tune the position of the MOT and compensate for gravity and stray magnetic fields. A short-circuited coil is also used to balance the effects of eddy-currents generated in the MOT chamber. These coils are described in more detail in [67, 146, 147].

The *MOT/Transport coils* are mounted on a mechanical translation stage to confine the atoms during transport to the glass cell in a process described in Sec. 4.5.2. These are illustrated alongside the other DC coils in Fig. 4.12. As they do not require water cooling, they are constructed from solid wire with a 1.8×3.6 mm cross-section.

A supplementary lateral steering bias coil (*TPush*) is mounted in the same array and applied during transport to steer the cloud through the differential pumping tube. This coil is constructed from the same wire as the *MOT* coils, and consists of 3×3 turns with inner dimensions $25 \times 100\ \mathrm{mm}^2$. The coil is fan-cooled, but would benefit from water cooling if a higher repetition rate is sought for the experiment.

4.4.2 DC coils

Three pairs of water-cooled coils surround the science cell to generate magnetic fields acting in the z (vertical) direction, with a single coil mounted perpendicular to the direction of atom transport to provide a bias field in y. These coils are constructed from hollow copper wire with a square profile, and are water cooled at a flow rate of approximately $0.05\,\mathrm{L\,s^{-1}}$ with water temperature between 10 and $12\,^{\circ}\mathrm{C}$[7]. A sketch of this array is shown in Fig. 4.12. These coils are nominally the same as those described in [67, 146, 147][8].

The *Quadrupole* coil pair, closest to the glass cell and connected in a nominally anti-Helmholtz configuration with current flowing in opposite directions in each coil, generates the quadrupole potential used to confine the atoms throughout evaporation and subsequent loading into the adiabatic potential. *Push* and *Auxiliary* coils are used during transport: *Push* to translate atoms in y to aid the transfer between the longitudinally-offset centres of transport and quadrupole trapping fields, and the counter-flowing *Auxiliary* pair to supplement the quadrupole field during magnetic transport. The *Helmholtz* coils are connected such that the current in each coil flows in the same direction. This creates a vertical bias field to displace the trap centre. This is a useful tool for high-resolution imaging given the restricted depth of focus of the objective, and aids the objective alignment process as described in Sec. 4.5.4.2. This bias field is also used to ensure the overlap of atoms confined on the AP shell surface with the location of the lock-in imaging lattice and dipole trap.

These coils are powered by a high-current power supply (Magna-Power Electronics SQA50-330) capable of outputting a maximum current of $330\,\mathrm{A}$ at a voltage of $50\,\mathrm{V}$. This current is controllably divided between the coils using the MOSFET bank described in [147, 166]. Gate voltages are controlled by a driver box and set by the computer control system.

---

[7] As discovered during separate incidents involving limescale buildup inside the DC coil array and algae growth in the water filter, a reduction in water pressure significantly reduces the cooling power available: with a ‘good’ flow rate of around $0.05\,\mathrm{L\,s^{-1}}$ though the coils, the *Quadrupole* coils reach around $53\,^{\circ}\mathrm{C}$ during a standard BEC sequence, measured using a thermocouple attached to the outside of the output wire, while a constricted flow rate of even $0.04\,\mathrm{L\,s^{-1}}$ resulted in a peak coil temperature of up to $70\,^{\circ}\mathrm{C}$. With these reduced flow rates, any variation in cooling water temperature can regularly push this temperature above the $70\,^{\circ}\mathrm{C}$ interlock point set to protect the coils from damage due to excessive heat, requiring a reduction in the quadrupole current during evaporation. Regular descaling of the coils helps to mitigate these problems, although an independent water supply would allow us to regulate our own water use.

[8] The quadrupole coils were replaced by Edward Owen in 2012 to the same design [29].

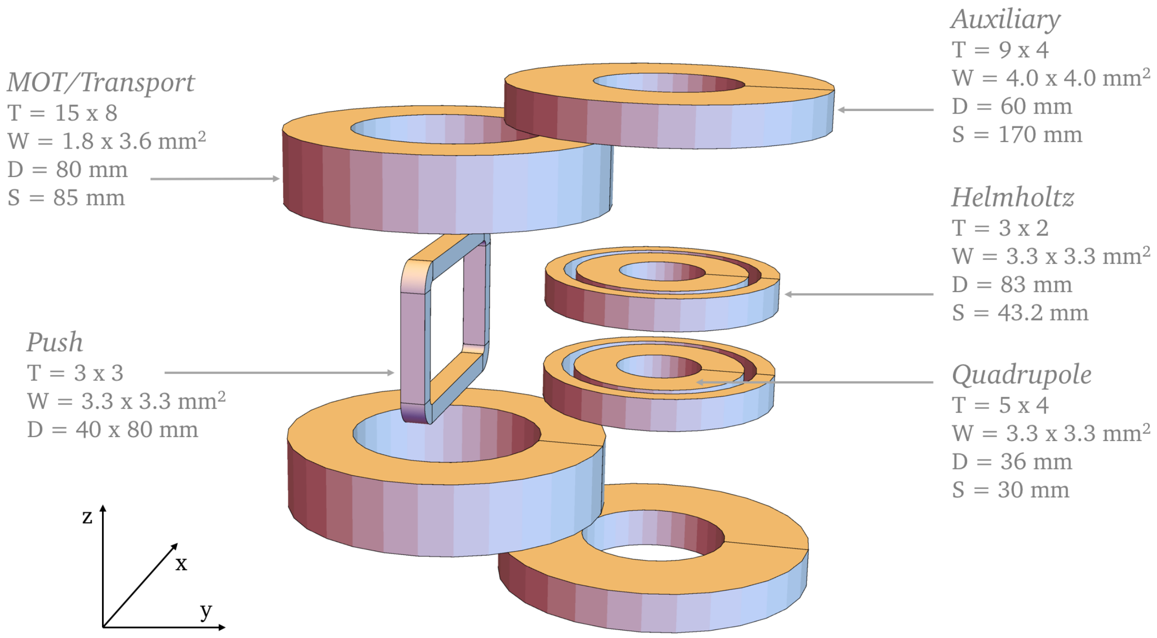

Figure 4.12: The DC coil array properties (*TPush* omitted for clarity). Atoms are transported along y, and the *Transport*, *Auxiliary*, *Helmholtz* and *Quadrupole* coils are arranged in vertically symmetric pairs about the atom transport plane The listed quantities are: T = wire turns (horizontal by vertical); W = wire dimensions; D = inner diameter of coil; S = vertical offset between coil pair centres. The AC coil array and glass cell sit between the planes formed by the *Quadrupole* and *Helmholtz* coils. The vertical axis of the experiment is defined by the centre of these coils. Coil properties obtained from Refs. [67, 146]. Figure created using Radia [165].

4.4.3 AC coils

The AC coil array is used to generate the fields associated with RF dressing and evaporation, and the 7 kHz rotating bias field that transforms the quadrupole potential into a TOP trap. This array comprises the TOP coils in the x and y directions, a separate zTOP (unused in the present work), RF dressing coils in x and y, and a single evaporative RF coil.

The TOP coils, connected in pairs along x and y respectively, are used to apply a 7 kHz rotating bias field to the quadrupole or RF dressed potential. Constructed from copper wire with diameter 1.5 mm, these coils have a rectangular profile, with outer dimensions 26×48.7 mm, and a depth 11.1 mm corresponding to a single layer of 7 turns, slightly separated by the glue used to set them. Their internal separation is

44.3 mm in y (parallel to the transport direction) and 46.7 mm in x. The 7 kHz signal is generated using an *Agilent 33220A* function generator operated in burst mode, outputting a signal of amplitude 460 mV. This signal is split in two, and the resulting components passed through a series of circuitry [167] designed to independently and controllably attenuate each signal and impose a $\pi/2$ phase shift between the two resultant TOP components. The attenuated signals are then amplified using a *Crown XLS 5000* audio amplifier. A transformer is used to impedance match the Crown output to the coils at 7 kHz, ensuring sufficient transmission of the signal [67]. The attenuation settings are coarsely set by balancing the currents in each TOP coil pair, and fine-tuned to allow for differences between the pairs by RF spectroscopy measurements of atoms in the TOP field, measuring the effect of varying the amplitude of signal emitted by each pair until they match. The TOP field calibration measurements are detailed in Appendix A.1.

The evaporative RF coil is a single 3.3 cm diameter wire loop. The evaporation signal is procured using direct digital synthesis (DDS) with maximum output voltage 0.4 V peak-to-peak, as described in Ref. [147], and amplified by a 1 W amplifier. The RF powers required for evaporation are so low that no impedance matching of this antenna to the amplifier output is necessary, and it is testament to the resilience of the evaporation process that the variation in the transmitted power during the frequency ramp does not adversely affect cooling.

The dressing RF coils are nested inside the TOP array with rectangular cross-section 24×29.2 mm. They are constructed from 1.3 mm diameter wire and assembled as a 3-layer coil with $(5, 4, 5)$ windings in each layer to give an outer-layer depth of 6.9 mm. Dressing RF signals are generated using DDS and implemented as described in Chapters 5 and 6.

4.5 BEC production

A streamlined BEC production process is derived from these constituent systems, from which the shell loading procedure described in Chapter 6 can be initiated. This process begins in the MOT (Sec. 4.5.1), followed by mechanical and magnetic transport to the glass cell (Sec. 4.5.2) and then magnetic confinement and RF evaporation. At whichever stage the process is terminated, be it thermal cloud, BEC or atoms confined to the RF-dressed shell, destructive absorption imaging is the final diagnostic step (Sec. 4.5.4).

4.5.1 The MOT

The first stage of laser cooling takes place in a pyramid MOT [168]. In this configuration, sketched in Fig. 4.13 and discussed in detail in Ref. [146],the six beams required to provide three-dimensional cooling are generated by the reflection of a single large incident beam from the four faces of an inverted pyramidal mirror array. This geometry therefore requires cooling and repumping beams with a diameter approaching the size of the pyramid structure, and correspondingly a high beam power.

Both cooling and repumping light are incident from a common optical fibre derived from the TA output. This beam is collimated to a $1/e^2$ diameter of 80 mm when incident on the MOT mirrors. This is constrained by the inner diameter of the transport coils through which it must pass, but ideally fills the effective diameter presented by the MOT mirrors. Beam powers typically take the values 350 mW for cooling, and 40 mW for repumping light. Two monitoring schemes are in place: a small amount of light is picked off from the main beam using a prism and incident on a photodiode to track any variation of incident power, while the MOT fluorescence as detected by a *Foculus FO442B* camera is used to monitor the size and shape of the MOT cloud. Sensitive to slightly different factors, these methods in conjunction are generally sufficient

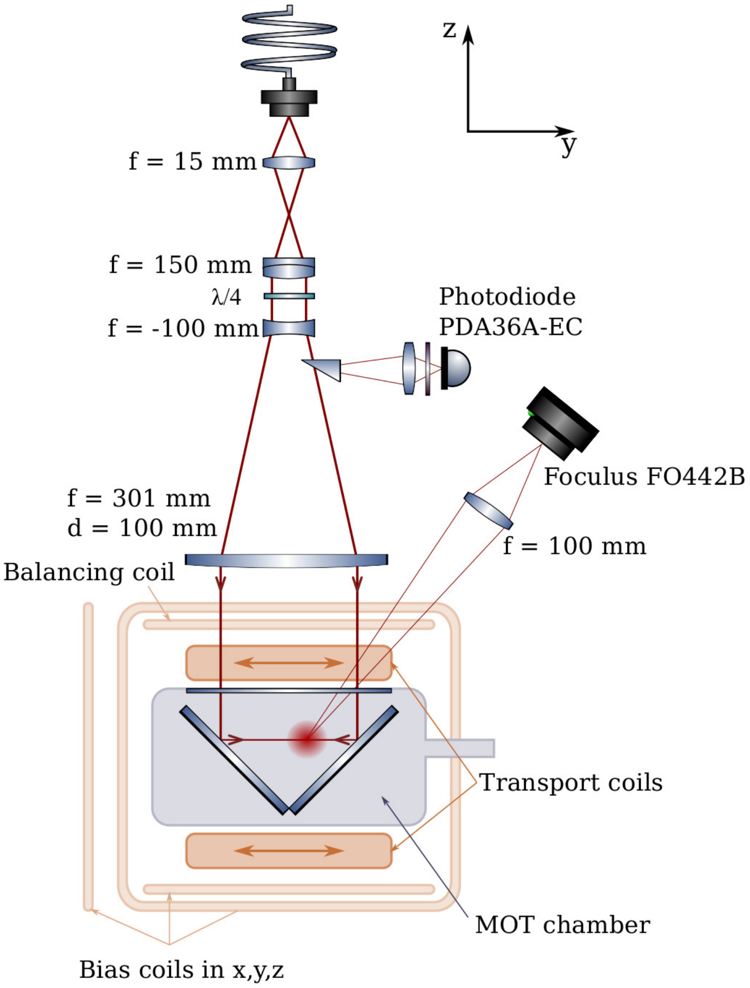

Figure 4.13: Schematic of the pyramid MOT optics, with transport and bias coils. Vertical optics are secured to a vertical support mounted on a translation stage. The atom location is offset from the intersection of the pyramid mirrors, due to gaps in reflection.

to diagnose any problems with the MOT, including drifts in laser power or frequency, and any misalignment of the MOT optics or fields. However, the photodiode is itself sensitive to the MOT fluorescence to an extent, which can slightly obscure the diagnostic procedure.

The quadrupole field that provides the imbalance in scattering rates necessary for atom confinement is generated using the *MOT/Transport* coils located above and below the MOT chamber. An additional set of three coils provides a bias field to offset the cloud from the centre of the mirror array, where the mirrors intersect and there is a gap in reflection. These also provide compensation against gravity and stray fields.

Transport to the glass science cell directly follows Doppler cooling in the MOT. A compressed MOT (CMOT) stage acts to increase the density of the MOT prior to loading into the transport quadrupole. There are two main mechanisms limiting the density of atoms confined in a MOT: for small clouds, collisions associated with the thermal motion of individual atoms, and for the larger clouds dealt with as the precursor to this experimental work, the re-radiation of resonant photons causing internal heating and expansion [169, 170]. Reducing the density to optimise the quadrupole loading must therefore address not only the gradient of the quadrupole field itself but also these rescattering processes. Several different compression methods exist, including the reduction in the outward pressure of rescattered radiation and losses from excited-state collisions by pumping the atoms into a dark state, either over a small spatial region in the centre of the trap as in the dark SPOT trap by witholding the repumping light [170], or by varying the laser detuning and intensity parameters in time, accompanied by an increase in quadrupole field gradient to create the temporal dark SPOT [169, 171]. The approach employed in this work, discussed in [67, 146], is to similarly increase detuning to reduce the effects of photon rescattering and increase the cloud density, while decreasing the MOT quadrupole field to increase the volume over which sub-Doppler cooling mechanisms apply and yield an overall colder transported cloud. Over 50 ms the MOT quadrupole field is ramped between 13.9 and 5 $\mathrm{G\,cm^{-1}}$, while the cooling beam detuning is increased from 22 MHz to 34 MHz over 40 ms, and further to 74 MHz over the next 10 ms. The repumping light is also blocked for 1.5 ms to pump atoms into the $|F=1\rangle$ state. The cooling light is turned off and the cloud loaded into the transport quadrupole potential with a field gradient of 76 $\mathrm{G\,cm^{-1}}$ for transport. A rapid current ramp in this coil is facilitated by the quick-start circuitry discussed in Ref. [146].

4.5.2 Atom transport

As discussed in Sec. 4.3, the MOT chamber and glass science cell are separated by a differential pumping tube, along which the atoms must be transported. This transport is conducted in two stages: mechanical and magnetic, and follows the method developed in Ref. [146].

The transport coils are mounted on a *Parker Hannifin 404 XR* translation stage, with a translation profile programmable using the *ViX 250ie* drive unit. The motor (*Parker SMB60 30 1,4 892VB64 0*) gives a translation precision of a few tens of $\mu\mathrm{m}$. Following a period of unreliable operation during which the end-point of the translation stage would vary by up to $2\,\mathrm{mm}$[9], the final position of the translation stage is monitored using a light gate created using a razor blade attached to the translation stage, positioned such that it blocks approximately half of a laser beam incident on a photodiode. Following maintenance, a precsion of a few $\mu\mathrm{m}$ was restored.

The lateral steering (*TPush*) coil mounted perpendicular to the transport and is powered by a *Lambda ESS* current source. To create the necessary smoothly controlled current ramp profile, the current through this coil is controlled by a MOSFET bank driven by a control voltage, with unwanted current diverted through a similar coil mounted on water cooling plates.

This translation stage carries the MOT/transport coils and TPUSH coils, and therefore the confined atoms, a distance of 430 mm along the differential pumping tube over 2.3 s, ending in line with the perpendicular Push coil as illustrated in Fig. 4.12. The vertical imaging optics prevent mechanical transport to the glass cell quadrupole centre. Instead, the translation is terminated with the atoms 90 mm from their ultimate destination and the remainder of the transport carried out by a purely magnetic transport sequence. This sequence is instigated using the dc coil array shown in

---

[9] This was remedied by two maintenance procedures: cleaning and re-lubricating the ballscrew, which fixed an audible scraping of the translation stage during transport, and tightening the coupler screw connecting the motor block and ballscrew. Following the initial period of unreliability in Spring 2015, this maintenance has been performed approximately once a year.

Fig. 4.12, and the corresponding current ramps plotted in Fig. 4.14.

The first field to be established is that of the *Push* coil, which provides a bias field to push the atoms the last few cm into the glass cell. While it does so, the *Auxiliary* current rises to catch the atoms in a quadrupole potential. The interplay between *Auxiliary* and *Quadrupole* coils then mediates the transfer of atoms into the smaller-volume, tighter quadrupole potential. These ramps are based on those presented in [146].

Following this transport sequence, approximately 5×10^8 atoms are deposited in the glass cell quadrupole trap[10]. With some heating unavoidable during the extended transport sequence, these atoms have a temperature of a few hundred $\mu\mathrm{K}$ and are ready to begin evaporative cooling.

4.5.3 Evaporative cooling

Forced RF evaporation in the glass cell is carried out in two sequential magnetic potentials: a quadrupole potential and subsequent TOP formed by applying a rotating bias field. During this process a weak 'RF knife' resonant with the Zeeman splitting between trapped and untrapped states is applied in a sweep from high to low frequencies. Since trapped atoms with higher energies oscillate with larger amplitudes, removal of atoms by transitions to untrapped states is performed for atoms with successively lower energies.

The initial quadrupole evaporation stage is performed at a quadrupole gradient of $330\,\mathrm{G\,cm^{-1}}$, with an evaporative RF sweep between 120 MHz and 20 MHz over 12 s to give a cloud of approximately 8×10^7 atoms at a temperature of around $20\,\mu\mathrm{K}$.

The TOP bias field is then ramped to 19.4 G over 1 s to begin TOP evaporation to BEC. The TOP trap evaporation ramps are illustrated in Fig. 4.15, which shows the change in TOP field as the driving current is varied, and the RF ramps applied. The corresponding progression in the phase-space density of the atoms is plotted alongside

---

[10] This value is inferred from a fit to an absorption image of the transported cloud, which forms an image too large to be seen on the CCD in its entirety.

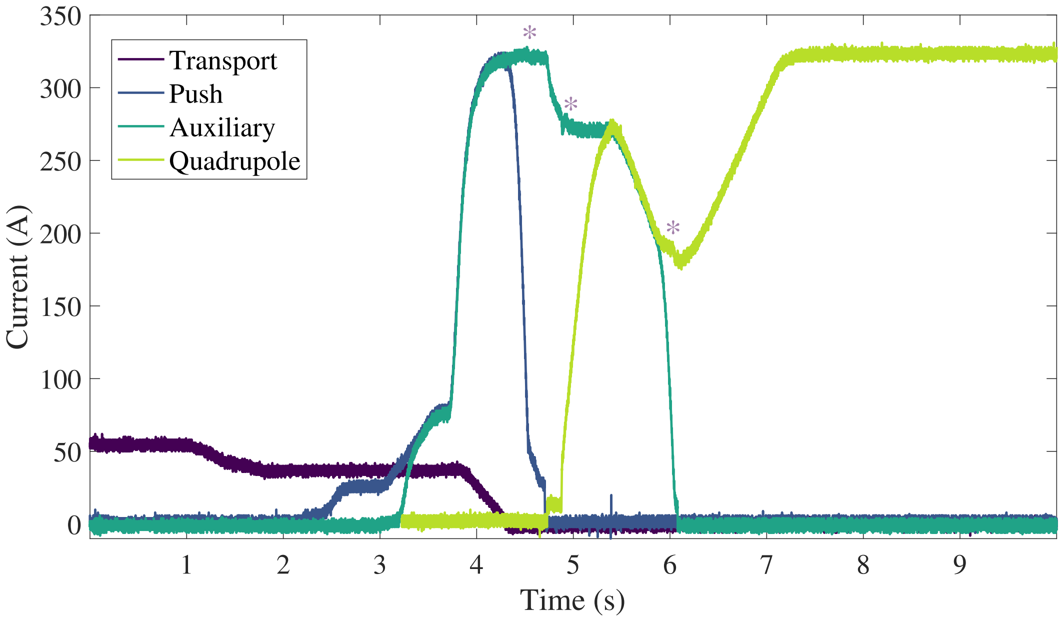

Figure 4.14: Magnetic transport current ramps through the DC coil array. Asterisks indicate locations where current spikes most often appear if there are imperfections in the MOSFET bank circuitry; depending on the severity these can significantly affect the characteristics of the transported atom cloud. A minimal effect has been observed on the temperature and number of atoms transferred to the quadrupole with a variation in peak amplitude of the *Push* coil.

this. The phase-space density plot can be used to guide and optimise the evaporation process, aiming to maximise the atom number and phase space density at each step.

The TOP itself introduces an additional mechanism for selective atom loss: the *circle of death*, or circling field zero that can eject atoms on contact. The circle of death radius is given by $r_0 = 2B_T/B'_q$, where B_T is the TOP field amplitude and B'_q the radial quadrupole gradient. The amplitude of the TOP field therefore sets an upper limit to the spatial extent, thus temperature, of the trapped atoms, in addition to affecting the trap frequency. A lower TOP amplitude B_T corresponds to a higher trap frequency and rethermalisation rate, with radial and axial trapping frequencies given by

$$\omega_x = \omega_y = B'_q\sqrt{\frac{g_F m_F \mu_B}{2m_{\text{Rb}}B_T}}, \tag{4.1}$$

$$\omega_z = B'_q\sqrt{\frac{4g_F m_F \mu_B}{m_{\text{Rb}}B_T}}. \tag{4.2}$$

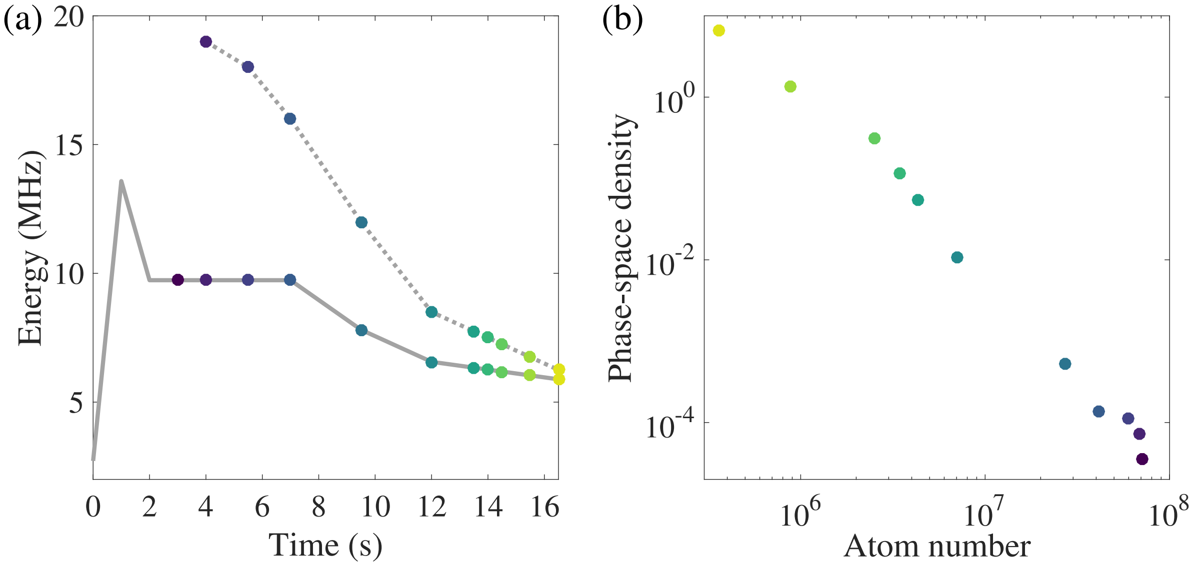

Figure 4.15: (a) TOP field (solid line) and evaporative RF (dotted line) ramps throughout the evaporative cooling process. For the first 4 s in the TOP trap, no evaporative RF is applied. Coloured points indicate times at which the atom number and phase-space density were measured. (b) The corresponding change in atom number and phase-space density during evaporation, charting the formation of BEC. The colour of the points indicate the corresponding stage in the TOP and RF field ramps.

The interplay between RF and TOP ramps take this into account. The quadrupole gradient climbs to $235\,\mathrm{G\,cm^{-1}}$ over the first three seconds of TOP confinement, following a drop to $20\,\mathrm{G\,cm^{-1}}$ over ~ 500 ms at the end of quadrupole evaporation.

Implicit from Fig. 4.2, TOP evaporation need not always carry the cloud to BEC; heating during adiabatic potential loading means that we choose to load a thermal cloud at around $1.7\mu\mathrm{K}$ to avoid unnecessary atom loss, completing the forced evaporation process in the time-averaged adiabatic potential as described in Chapter 6. Whether a BEC or thermal cloud, in the TOP or shell trap, the imaging process is a necessary final step that allows us to measure the atomic distribution, inferring trap properties and atom behaviour.

4.5.4 Imaging systems

The atom cloud is characterised by absorption imaging: a resonant probe beam illuminates the atoms and the shadow cast by absorption is imaged to infer the atom cloud properties. Three pictures are taken at an interval of 1 s and exposure time of $15\,\mu\mathrm{s}$: C_1 describes the measured counts for the probe beam with its absorption shadow; C_2, the probe beam alone after the irradiated atoms have fallen from the trap; and C_3, a dark image taken with all probe beam shutters closed to give the intrinsic camera dark counts for removal. By subtracting the dark image from each of the others to give 'atom' and 'light' images $C_A = C_1 - C_3$ and $C_L = C_2 - C_3$ respectively, the atom distribution can be reconstructed.

Experimentally, absorption imaging is performed in three separate directions: vertical (z), transverse horizontal (x) and longitudinal horizontal (approximately along y). The following discussion relates to imaging along x, which forms our primary horizontal imaging direction, but can be easily extended to the other imaging axes by substituting the appropriate variables. The intensity $I(y, z)$ of the probe beam profile following absorption is described in terms of the undisturbed beam intensity $I_0(y, z)$ by the Beer-Lambert law:

$$I(y, z) = I_0(y, z)e^{-D(y,z)}. \tag{4.3}$$

This depends on the optical column density $D(y, z)$ of the atom cloud, calculated by integrating the effective atom number visible to the beam along the imaging direction, in this case x:

$$D(y, z) = \int \sigma(I, \delta)n(x, y, z)\,\mathrm{d}x \tag{4.4}$$

where $n(x, y, z)$ is the atom number at a given point in space, and the effective atom number is calculated by taking into account the absorption cross-section

$$\sigma(I,\delta) = \frac{\Gamma}{2} \frac{2\frac{I(y,z)}{I_{\text{sat}}}}{1+2\frac{I(y,z)}{I_{\text{sat}}}+4\frac{\delta^2}{\Gamma^2}} \frac{\hbar\omega}{I(y,z)} \tag{4.5}$$

$$= \frac{\sigma_0}{1+\frac{I(y,z)}{I_{\text{sat}}}+4\frac{\delta^2}{\Gamma^2}} \tag{4.6}$$

where $I(y,z)$ is the spatially-varying probe beam intensity and δ the variable detuning of the probe beam from resonance. $I_{\text{sat}} = 1.67\mathrm{mW}$ is the saturation intensity for σ-polarised light, $\Gamma = 2\pi \times 6$ MHz the natural linewidth of the excited state and $\sigma_0 = 2.9 \times 10^{-9}\mathrm{cm}^2$ the resonant absorption cross section for σ-polarised light [148].

To calculate the atom number associated with the absorption imaged by each pixel (i,j), the atom number in the associated column $n(x,y,z)$ can be integrated along the imaging direction. Using Eq. 4.4, this can be rearranged in terms of column density D, and via Eq. 4.3 the measurable quantities $I(y,z)$ and $I_0(y,z)$ described by C_A and C_L respectively:

$$N(i,j) = \int_x \int_A n(x,y,z)\,\mathrm{d}A\,\mathrm{d}x \tag{4.7}$$

$$= -\frac{A}{\sigma(I,\delta)} \ln\left(\frac{C_A(i,j)}{C_L(i,j)}\right) \tag{4.8}$$

From this pixel-by-pixel atom number reconstruction, the spatial distribution of the trapped cloud can therefore be extracted, characterised by fitting an appropriate function, and total atom number calculated by summing over all pixels on the CCD.

However, the Beer-Lambert law is only valid for an optically thin cloud: one for which the column density $D < 1$. This is untrue of most of the cold atom clouds considered in this work. In accordance with Eq. 4.4 increasing the probe beam detuning is one option to decrease this column density. However, the associated increase in dispersion also causes lensing effects and distortion of the image if taken to too great an extreme, which is itself exploited in non-destructive imaging techniques [172].

Instead, atoms are released from the trap and imaged after free expansion. More crucially than just reducing the column density, this time of flight expansion exposes the energy distribution of the atom cloud: upon abruptly removing the trapping potential, the interaction energy is converted into kinetic energy and the cloud expands as it enters freefall [173]. During the ballistic expansion period, atoms with more kinetic energy travel further, resulting in a spatial separation of the atoms according to their energy in-trap and a resulting image of the momentum distribution. Three distinct distributions can be fitted during image analysis: a Gaussian fit to a thermal cloud, Thomas-Fermi distribution (inverted parabola) for a BEC, and bimodal function to chart the emergence of the BEC peak from a background thermal cloud.

Our imaging procedure begins with repumping light applied for 160 ms to repump the atoms from the $|F=1\rangle$ to $|F=2\rangle$ hyperfine level. Resonant light is derived from the cooling laser as shown in Fig. 4.7 after the first AOM detuning stage on board C4. This light is shifted back towards resonance by a second AOM, to give a detuning from the cooling transition by 6 MHz to bring the light into resonance with atoms confined in the TOP. This probe beam is applied to the atoms for a pulse of $15\,\mu\mathrm{s}$ duration, and is σ^--polarised to maximise the absorption cross-section [148], with the atoms aligned to a quantisation axis provided by the TOP field. This is applied with an amplitude of 3.2 G and timed such that the field is aligned with the probe beam during the $15\,\mu\mathrm{s}$ probe duration. Atom, light and dark images are taken as described above. These are spaced as closely as possible to minimise the emergence of fringes during the processing stage, due to slight differences between the beam appearance in each image resulting from slight drifts in the position of imaging optics.

The optical breadboard that controls the probe light distribution between horizontal and vertical imaging systems is illustrated in Fig. 4.16. The AOM is used to control the detuning and act as a fast shutter, with a response time on the order of ns, while mechanical relay shutters provide a slower secondary stage of shuttering. These prevent any disruption to the cloud from residual light and allow the AOM to be operated in its ‘on’ state throughout the duration of the experiment, allowing temperature stabilisation.

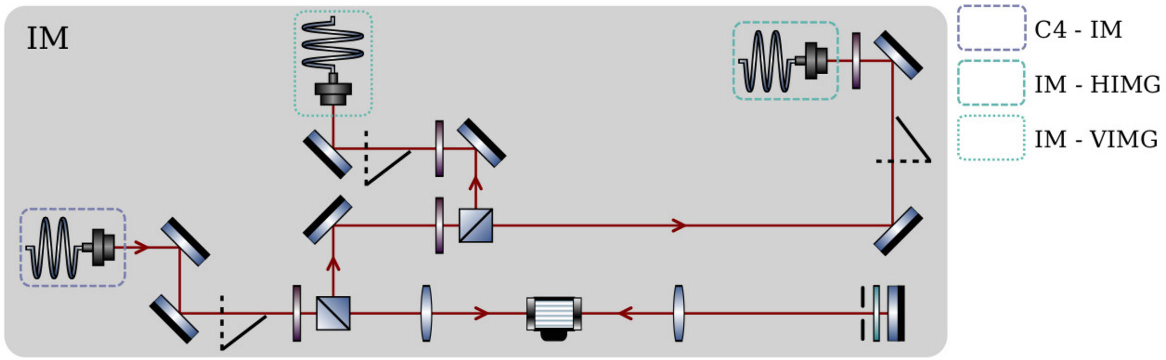

Figure 4.16: Imaging light distribution board. A $\lambda/2$ waveplate controls the proportion of light between the vertical and horizontal directions. Each beam path contains independent mechanical shutters, while the AOM acts as a fast shutter and controls the beam detuning.

4.5.4.1 Horizontal imaging

Absorption imaging is carried out in the horizontal plane in both the lateral direction and an approximation to the longitudinal directions (x and y respectively). A schematic of the optical setup is shown in Fig. 4.17.

The imaging repumper light is salvaged from the TA setup shown in Fig. 4.9, with shuttering controlled prior to this by the mechanical Uniblitz shutter on the repumping breadboard shown in Fig. 4.8. The polarising beamsplitter cube ensures a fixed polarisation of repumping light, and allows continuous monitoring of the beam power. The repumping beam is expanded to a beam with diameter 1 cm to ensure uniform illumination of atoms prior to imaging. Efficient repumping requires beam powers above 0.8 mW after the beamsplitter; a uniform response in atom number is observed for beam powers up to 2 mW.

The lateral imaging direction is the one in primary use, for day-to-day cloud diagnostics and the production of high-quality images. This setup has a measured magnification of 1.028 (calibration methods discussed in Appendix A.3), and is a simple

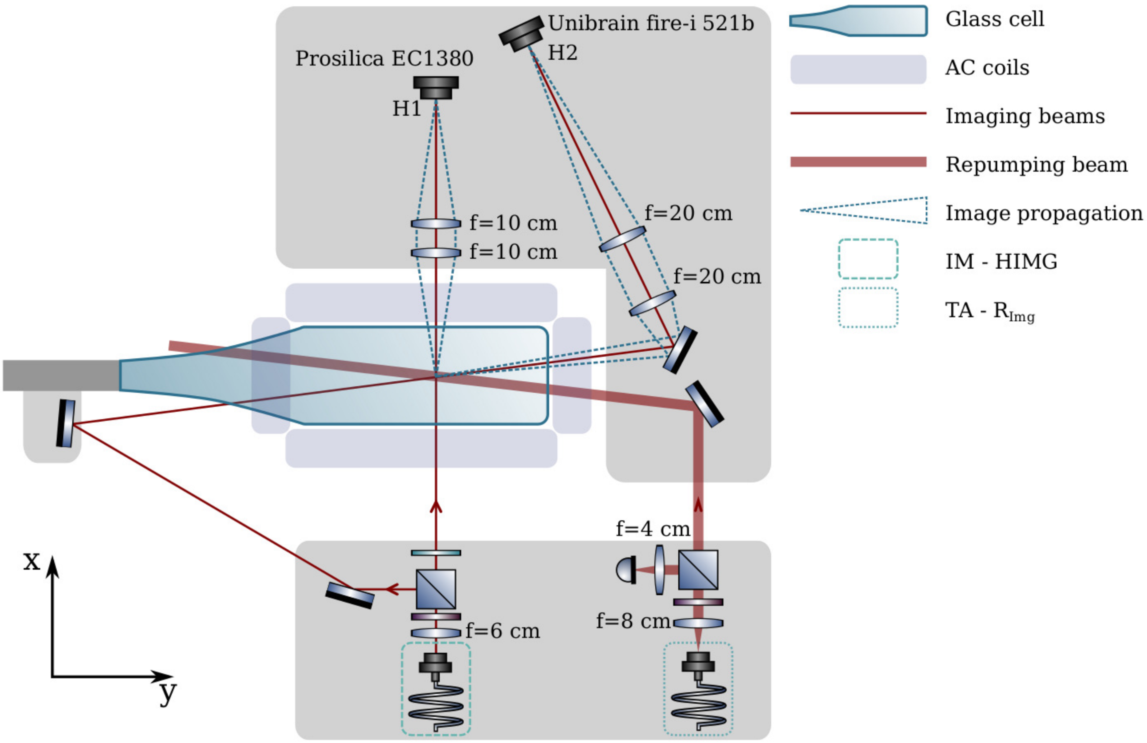

Figure 4.17: Horizontal imaging optics layout, showing both transverse (primary imaging, H1) and longitudinal (diagnostic only, H2) imaging arrangements. Beams are illustrated as solid red lines, and the image propagation as blue dotted lines. The optics for horizontal imaging are contained within two separate breadboards raised from the optical table. Optical lattice optics (not pictured) are also contained on the camera breadboard on this plane. The probe beam (optical fibre on the left) is controllably split into two mechanically-shuttered independent imaging paths, both consisting of a 1:1 telescope. Optical access for the diagnostic beam is highly constrained due to lattice optics and the AC coil array; the beam enters through the curved region of the glass cell, introducing severe distortions to the beam. The photodiode on the repumping beam (optical fibre on the right) is used to monitor any amplitude or polarisation fluctuations.

1:1 imaging system consisting of two $f = 10$ cm lenses placed a focal length from the atom position and CCD respectively. The *Prosilica EC1380* camera has pixel side length $6.45\,\mu\mathrm{m}$. With freedom in the spacing between these lenses, they are simply placed as close together as is practical, with a 1 cm spacing[11].

The longitudinal imaging system was implemented for purely diagnostic purposes, facilitating comparable alignment of the dipole trapping beam in both *x* and *y* direc-

---

[11] The inter-lens spacing is not truly arbitrary due to the finite Gaussian beam divergence. However, flexibility in lens position is reasonable for the beams and length-scales considered here.

tions. This is again a 1:1 imaging system, consisting of two $f = 20$ cm lenses due to space constraints imposed by the lattice optics. This system has a measured magnification o 0.997, imaged onto a Unibrain Fire-i 521b CCD camera, with pixel side length $9.9\,\mu\mathrm{m}$. As can be seen in Fig. 4.17 this imaging beam is not perfectly perpendicular to the primary horizontal imaging direction, as dictated by optical access constraints, and is for the same reason directed through the curved region of the glass cell where the rectangular profile is moulded into a cylinder to match the vacuum chamber fittings. This causes visible distortion to the image. Instabilities also arise due to the mirror mounted on a small breadboard attached to a vacuum system support. As such, the images are far from publication quality, and the coupling between horizontal imaging directions must be taken into account in analysing the resultant images. A further important note regarding this setup is that the camera is mounted upside down. However, this system is extremely useful for diagnostic purposes. The alignment procedure for these and the vertical imaging system is described in Appendix B.

4.5.4.2 Vertical imaging

The vertical imaging system was designed to accommodate three different goals, with sometimes conflicting requirements: vertical absorption imaging, fluorescence imaging, and the optical dipole trap. The common feature between these is the four-lens objective designed to capture fluorescence from the imaging lattice described in Ref. [29], which therefore necessarily forms part of the high-magnification vertical absorption imaging setup and is used in reverse to focus the dipole trapping beam to a $1.6\,\mu\mathrm{m}$ diffraction-limited focus at the TOP centre. The imaging system must also allow for the 830 nm dipole trapping light to be coupled in at a location in keeping with the design for trap rotation discussed in Chapter 8.

This four-lens objective illustrated in Fig. 4.18 was constructed for the work on rotating optical lattices described in Refs. [174, 175], characterised in Ref. [176]

and based on the design in Ref. [177]. Constructed from standard catalogue anti-reflection-coated lenses, this four-lens system was designed to account for the 2 mm glass cell thickness, compensate for the spherical aberration associated with the constituent lenses, and provide a numerical aperture of 0.27.

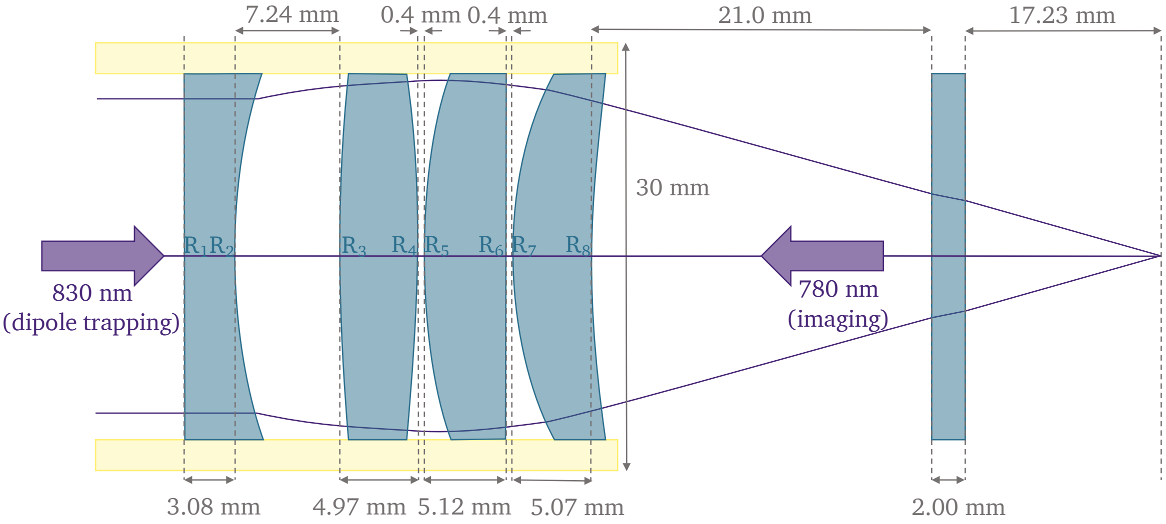

Figure 4.18: Schematic of the four-lens objective, with ray propagation shown in purple. This consists of catalogue lenses mounted in a brass tube, and was designed to compensate for the 2 mm glass cell thickness also shown. The constituent surface radii of curvature are: $R_1 = \infty$, $R_2 = 39.08$ mm, $R_3 = 103.29$ mm, $R_4 = -103.29$ mm, $R_5 = 39.08$ mm, $R_6 = \infty$ mm, $R_7 = 26.00$ mm, $R_8 = 78.16$ mm.

The compound focal length of this system is 3.76 cm, with a front focal plane 1.04 cm and back focal plane 0.34 cm within the lens system. The corresponding diffraction-limited resolution of this lens system is given by Rayleigh's criterion

$$\Delta l = 1.22\frac{f\lambda}{D}. \tag{4.9}$$

With a lens diameter $D = 24.5$ mm and considering a wavelength $\lambda = 780$ nm, this resolution is given by $\Delta l = 1.5\,\mu\text{m}$. Its depth of field as parameterised by the Rayleigh length of the beam $z_R = \pi\omega_0^2/\lambda$ is therefore $9\,\mu\text{m}$, with the focal tolerance of the objective $\Delta z = \left(\frac{f}{r}\right)^2\lambda = 8\,\mu\text{m}$ with r the lens radius [178].

Illustrated in Fig. 4.19, the imaging system itself is conceptually identical to the horizontal systems discussed above, its main challenges arising from the rapidly diverging beam following the objective and the difficulty in optical access to the vertical arrangement.

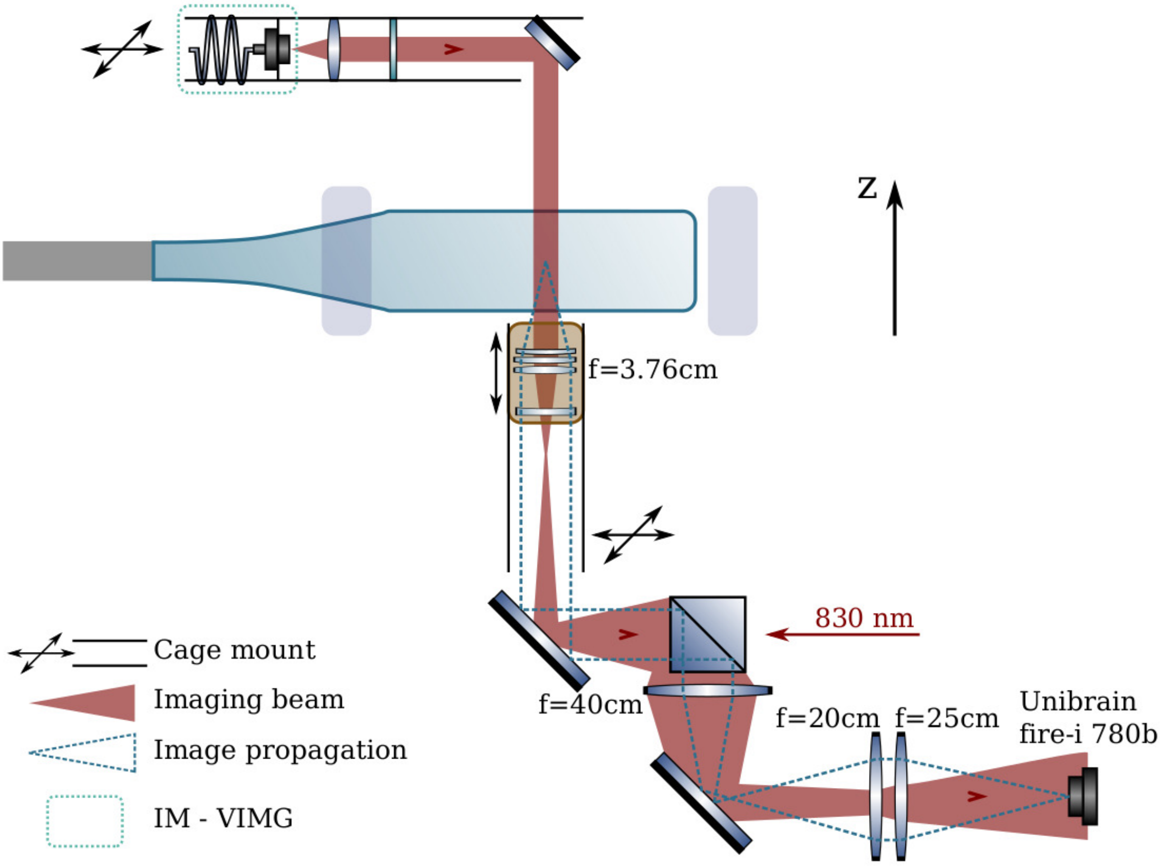

Figure 4.19: Vertical imaging optics layout, showing the vertical propagation of probe light through the glass cell. The probe beam is shown in red, with illustrative beam width variation, and the atom image as the blue dotted lines. The probe light is derived from an upper-tier optical bread board and held within a cage-mount construction mounted on an *x*-*y* translation stage to steer the imaging beam. Telescope lens separations are constrained by the vertical construction and necessity for the beamsplitter at which the 830 nm dipole trapping light is applied. The imaging telescopes incorporate the four-lens objective, which is held in a vertical cage-mount construction with three-dimensional translation. The glass cell and y AC coils are shown for orientation purposes.

The imaging system comprises two separate telescopes, the first including the four-lens objective and a $f = 40$ cm lens, and the second lenses with focal lengths $f = 20$ cm and $f = 25$ cm, to give a magnification of approximately 14. As with the horizontal telescopes, the four-lens objective is placed at its focal length from the atoms and the final imaging lens at its focal length from the CCD, with confocal spacing between telescopes. The spacing within each telescope is again flexible, and is kept to a minimum to maximise the amount of light captured by the imaging system. To the same end, the focal lengths of the constituent lenses were carefully chosen, and lenses with a 50.8 mm diameter used. Precluded for now by the surrounding system and requirements placed by the dipole trapping optics, the signal-to-noise ratio of the imaging system could be improved by using a shorter first telescope to capture a larger fraction of the rapidly-diverging imaging beam. This could in future be incorporated into the design of both vertical imaging and dipole trapping systems, though would require a substantial redesign of both.

The input beam is launched from a cage-mount construction clamped securely in place on a breadboard at a height of several cm above the glass cell. The imaging beam propagates vertically downwards through the centre of the glass cell. The four-lens objective acting to collect this signal is housed in a cage-mount construction whereby the objective is supported by cage-mount rods attached to an x-y translation stage, offset along x for optical access. A z translation stage holding the cage-mount system provides additional fine axial adjustment. This is an extension of a previous design that suffered from stability problems due to large-amplitude vibrations of the objective as a result of its mounting on a springboard construction. Due to constraints of space and optical access the new construction followed the same principles but minimised the longitudinal distance between the cage mount and base, and maximised the thickness of all supporting components.

Vertical imaging alignment is more challenging than in the horizontal direction, due to the short depth of field of the four-lens objective and spatial constraints around it. However, it is also more critical, since impeccable objective alignment is necessary for both fluorescence imaging and the dipole trap The alignment process for this system is described in detail in Appendix B.

Chapter 5

Impedance matching for multiple-RF dressing

Significant amongst the challenges of adapting the experimental apparatus to incorporate multiple dressing frequencies is that of generating sufficiently strong RF fields at each frequency of interest. This is achieved by impedance matching the RF coils. This chapter introduces the principles of impedance matching in Sec. 5.1 and the necessary tools for matching network calculations in Sec. 5.2. The characteristics and requirements of the dressing RF network are discussed in Sections 5.3 and 5.4 and the details of narrow and wideband impedance matches in Sections 5.5 and 5.6. Improvements to the matching networks are considered in Sec. 5.7.

5.1 Impedance matching

A change in the impedance experienced by an AC signal as it travels along a transmission line causes reflection of some fraction of the signal power. Not only do these reflections restrict the total power transmission to the load and inflict damage on the outputs of upstream components, but multiple reflections can result in the formation of standing waves, degrading the signal quality [179]. The impedance must therefore be matched at the end of transmission lines and between components. To account for the frequency dependence of both the load and any matching circuitry, circuit analysis is performed across the full range of frequencies of interest, and a compromise made between the precision and bandwidth of the resulting match.

Commercial RF components and devices are typically manufactured with an input impedance at the $50\,\Omega$ industry standard and can therefore be connected without additional matching required[1]. The circuit elements of concern are therefore the home-built AC coils. The following discussion focusses on the dressing RF coils matched at a few MHz, but the TOP coils are also matched to the 7 kHz signal using a transformer. The evaporative RF coil is unmatched, and so the 100 MHz range of evaporation frequencies applied would be difficult to accommodate without significant variation in transmission. While there is some variation in the frequency response of the evaporative RF amplitude, the evaporation process is sufficiently resilient to amplitude variations that this is not problematic; neither has any signal distortion or damage to the amplifier been observed.

While the principles remain the same, the details of the matching process vary significantly according to the frequency range in question, such as between microwave and long-wavelength RF regimes. The dressing fields of interest in this experiment have angular frequencies in the range $1\mathrm{MHz} < \omega/2\pi < 5\mathrm{MHz}$, a band officially designated for amateur communications and some maritime/aeronautical mobile [180]. The impedance matching techniques are therefore quite different to the standard microwave methods often presented. However, this frequency regime has a good response to standard laboratory components, allowing precision impedance matching using just ceramic capacitors and transmission lines of reasonable length.

---

[1] Custom matching may be desirable for certain components or applications. For example, it could be beneficial to apply a narrow-band match to AOMs used as single input frequency optical shutters, reducing or eliminating the signal amplification required. The measured impedance characteristics of the AOD used in the dipole trapping setup described in Chapter 8 could also be improved.

Mathematically, reflections are treated in an identical way to those that occur for any electromagnetic wave at the interface between two media [181], by considering the total voltage $U(l)$ at distance l along a transmission line as the sum of forward-propagating and reflected waves:

$$U(l) = U_f e^{i\beta l} + U_r e^{-i\beta l} \tag{5.1}$$

where U_f and U_r are the forward-propagating and reflected wave amplitudes $U_{f,r} = \frac{1}{2}(U_0 \pm I_0 Z_0)$ with phase constant β. Z_0 describes the chosen reference impedance, taken to be $50\,\Omega$ for subsequent discussions.

The reflection coefficient $\Gamma(l)$ is simply the ratio of these reflected and forward-propagating wave components and, at the position of the load with $l = 0$, is given by

$$\Gamma = \Gamma(l = 0) = \frac{U_0/I_0 - Z_0}{U_0/I_0 + Z_0} = \frac{Z_L - Z_0}{Z_L + Z_0} = \frac{z - 1}{z + 1} \tag{5.2}$$

with Z_L the load impedance and $z = Z_L/Z_0$ the normalised impedance[2]. It is this reflection coefficient that the impedance matching process seeks to minimise. Some definitions useful in the discussion of impedance matching are as follows:

Impedance: $Z = R + iX$ is a measure of the opposition to the flow of current in an AC circuit.

Resistance: The real part R of the complex impedance, associated with an ideal resistor.

Reactance: The imaginary part X of the complex impedance, associated with an ideal capacitor or inductor.

Admittance: The reciprocal of complex impedance, $Y = 1/Z = G + iB$, measuring the tolerance of a circuit to the flow of AC current.

---

[2] In subsequent discussions, the subscript L refers to the property of the load, and subscript 0 the reference value.

Conductance: The real part G of the complex admittance.

Susceptance: The imaginary part B of the complex admittance.

5.2 Impedance matching toolbox

Working out an impedance match requires the analysis of the reflection characteristics of the system over a range of frequencies. A vector network analyser (VNA) is a reflectometer that measures both the magnitude and phase of the reflection coefficient over a frequency sweep [182]. We use the *Omicron Bode 100* VNA, with a frequency range of 10 Hz to 40 MHz.

To calculate a match, the (angular) frequency-dependent impedance $Z(\omega)$ of each dressing RF coil pair was measured and used as an input to an RF circuit simulation created using a MATLAB toolbox [183]. A matching network with the correct characteristics was constructed using this simulation, using known values of accessible components. These simulations provide a good starting point to the experimental match with final adjustments made using live analysis of the circuit impedance due to imperfections in real components, and the stray capacitance and inductance associated with circuit construction The following sections detail the main tools and concepts used to perform RF circuit analysis.

5.2.1 The Smith Chart

The Smith chart is a logarithmically-scaled polar plot of the normalised complex impedance plane, that simultaneously represents the complex impedance and associated complex reflection coefficient of a network via the relation given by Eq. 5.2 [179, 181]. A sketch of the transformation of the complex impedance plane into a Smith chart is shown in Fig. 5.1. This starts with the complex impedance plane, running from 0 to ∞ on the real (resistance) axis and between $\pm\infty$ on the imaginary (reactance) axis. The reactance axis is deformed to form a polar plot, creating circles of constant resistance anchored at $Z = \infty + 0i$. Logarithmic scaling is applied such that the entire infinite half-plane is enclosed within this plot, with a real axis bounded by 0 and ∞ and the reference impedance located at the origin.

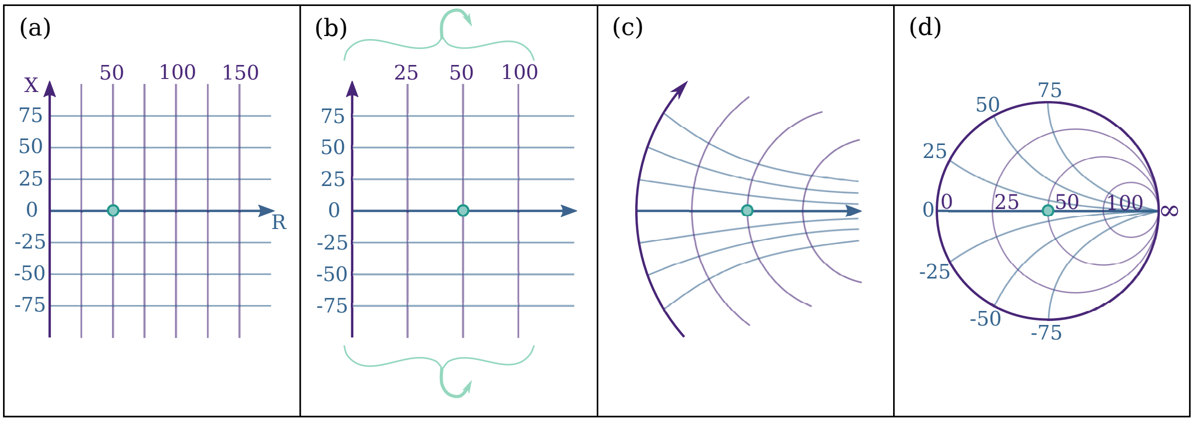

Figure 5.1: The Smith chart as a transformation of the complex impedance plane. (a) The complex impedance plane, with vertical (purple) lines of constant resistance, and horizontal (blue) lines of constant reactance. (b) Logarithmic scaling is applied, orienting the reference impedance (in this case $50\,\Omega$) at the centre of the plot. (c) The imaginary (reactance) axis is deformed to create a polar plot. (d) The Smith chart: the real (resistance) axis runs from 0 to ∞, with circles of constant resistance (purple) and arcs of constant reactance (blue). The logarithmic scaling means that the entire impedance plane, between $\pm\infty$ in the imaginary axis and 0 to ∞ in the real axis, is enclosed within the outer circle.

The mapping of this impedance plot to the complex reflection coefficient is most clear from considering Eq. 5.2. In terms of the normalised impedance $z = r + ix$, this can be rearranged to give normalised resistance and reactance values

$$r = \frac{-(\Gamma_R^2 + \Gamma_I^2 - 1)}{(\Gamma_R - 1)^2 + \Gamma_I^2} \tag{5.3}$$

and

$$x = \frac{2\Gamma_I}{(\Gamma_R - 1)^2 + \Gamma_I^2} \tag{5.4}$$

where Γ_R and Γ_I are the real and imaginary parts of the reflection coefficient respectively. Rearranged, Eq. 5.3 gives

$$\left(\Gamma_R - \frac{r}{1+r}\right)^2 + \Gamma_I^2 = \left(\frac{1}{1+r}\right)^2. \tag{5.5}$$

This describes a circle centred at $\left(\frac{r}{1+r}, 0\right)$ with radius $\frac{1}{1+r}$: the circles of constant resistance shown in the Smith chart. Likewise, Eq. 5.4 can be rearranged to give

$$(\Gamma_R - 1)^2 + \left(\Gamma_I - \frac{1}{x}\right)^2 = \left(\frac{1}{x}\right)^2 \tag{5.6}$$

which describes a circle centred at $\left(1, \frac{1}{x}\right)$ with radius $\frac{1}{x}$. These circles correspond to the arcs of constant reactance.

The complex Γ plane can therefore be visualised as having its origin at the centre of the real axis, corresponding to zero reflection at Z_0. This point therefore becomes an easily-identifiable target for the matching process: the closer the network impedance to the origin, the closer it is to the reference impedance, and a reflection coefficient of zero. Γ increases isotropically from this point, reaching its maximum value of 1 at the normalised bounding circle corresponding to $Z = 0 + iX$. Circles of constant Γ have been added to the Smith chart illustrated in Fig. 5.2a.

The Smith chart can also be plotted in the admittance form, by a rotation of $180°$ about the origin, allowing a direct visual mapping between impedance, admittance and reflection properties of the network. The circles of constant conductance thus emanate from the point of $Z = 0$, $Y = \infty$, with arcs of constant susceptance perpendicular to them. This full plot is shown in Fig. 5.2b where, in the impedance frame, the upper half of the plot represents the capacitive region and the lower half the inductance region. In the admittance frame the reverse is true.

The reflection coefficient associated with a frequency sweep can also be represented on the Smith chart, producing a line on which each point represents the network impedance or reflection coefficient at a given frequency. The length of this line

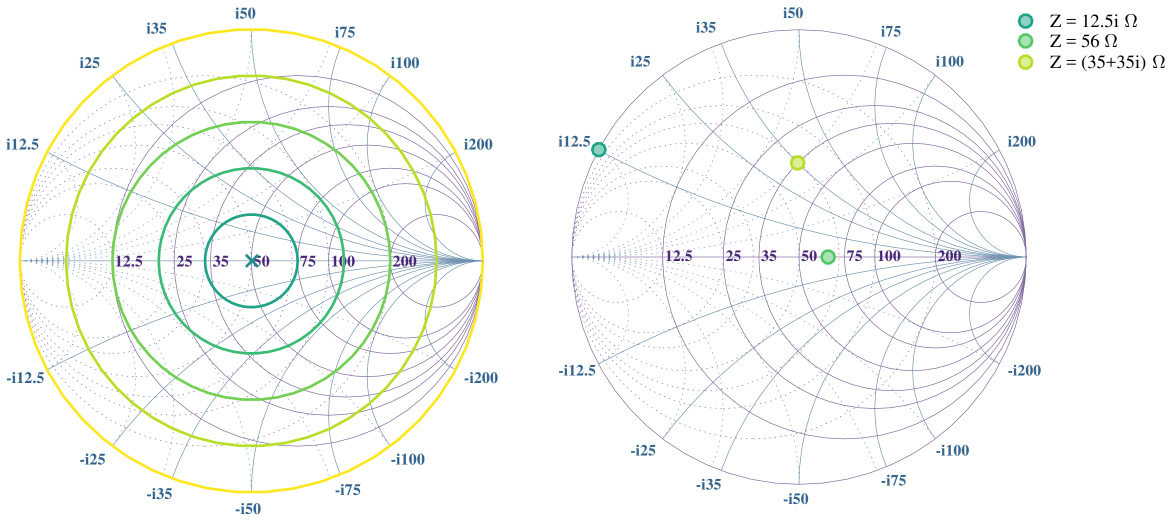

(a) Reflection coefficient

(b) Complex impedance

Figure 5.2: Representing physical quantities on the Smith chart. (a) Circles of constant Γ, at intervals of 0.2 between 0 and 1 extending from the centre of the plot (green to yellow). (b) Representing the complex impedance. Three values of impedance are plotted on the Smith chart, each point associated with a reflection determined by its proximity to the $50\,\Omega$ reference impedance.

represents the variation of the frequency response of the network over the range of interest, such that a short line corresponds to a wideband response clustered around a particular impedance or reflection coefficient, and a longer line a narrow-band response. The Smith chart can also be used to provide an intuitive representation of the effect of various circuit components on the network impedance, and is therefore extremely useful in guiding the modelling process and assessing the final match.

5.2.2 Circuit components

The components that we use are either dissipative (resistors), reactive (capacitors and inductors) or some combination of the two (transmission lines). Each such component corresponds to a particular mapping on the Smith chart, and an associated change in the complex impedance. An ideal capacitor with capacitance C is purely reactive, with a complex impedance

$$Z_C(\omega) = \frac{1}{i\omega C}. \tag{5.7}$$

Adding a capacitor in series, at a fixed value of ω, therefore corresponds to subtracting a reactance of $1/\omega C$ from the initial impedance, and thus a transformation anticlockwise along the iso-resistance circle. However, the frequency-dependence means that the reflectance measured over a frequency range forms an arc in which the lowest frequencies have been transformed by the greatest amount. To add a capacitor in parallel, the impedances of the load and capacitor are added in reciprocal and it is therefore easiest to consider a transformation in the admittance chart. In this case, a capacitor added in parallel corresponds to adding an admittance of ωC, and a transformation clockwise along the iso-conductance curve, in which a higher value of C or ω is transformed with a greater magnitude.

Likewise, an ideal inductor with inductance L has complex impedance

$$Z_L(\omega) = i\omega L. \tag{5.8}$$

An inductor in series maps the impedance clockwise along the iso-resistance curve, with a greater magnitude for large values of L or ω, and an inductor in parallel correspondingly maps the impedance anti-clockwise along the iso-conductance curve, with a reduced magnitude for larger values of L or ω. The effect of these reactive components on an arbitrary initial load Z_L is illustrated in Fig. 5.3.

Propagation through a transmission line imposes a phase shift on the signal, corresponding to a rotation of the complex impedance about the Smith chart origin[3]. This means that the effect of a transmission line on the reflection characteristics of the

---

[3] This effect makes it extremely important to calibrate the VNA reflectance with any diagnostic connections or transmission lines in place, so that their impact on the measured reflectance characteristics of the network can be removed.

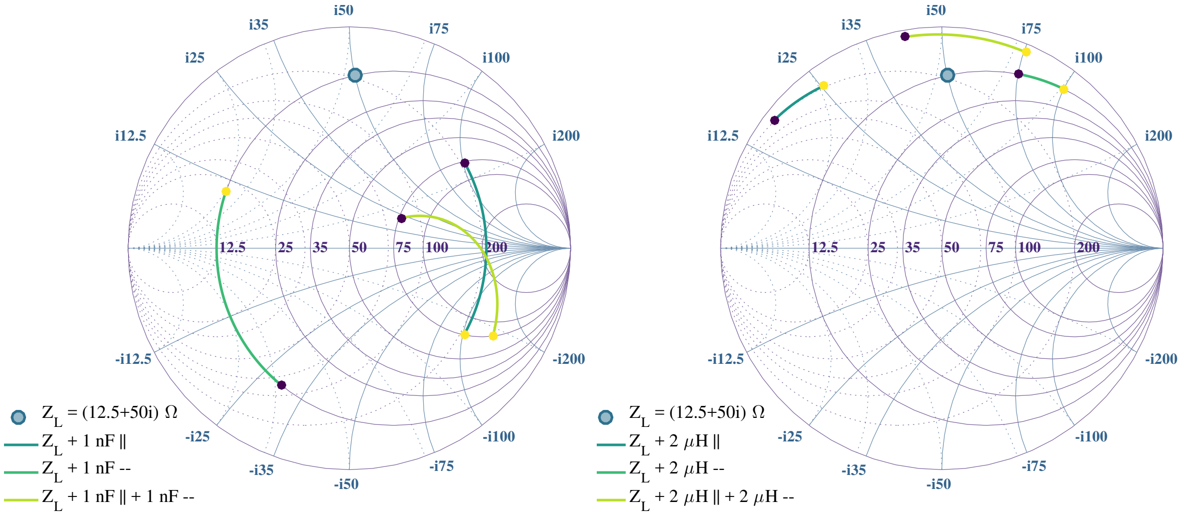

(a) Capacitative networks **(b)** Inductive networks

Figure 5.3: Transformation of a single impedance point due to reactive circuit components over a frequency sweep between 2 MHz (purple dot) and 4 MHz (yellow dot). || indicates a component added in parallel, and – in series. (a) Response to capacitive networks. (b) Response to inductive networks. The frequency dependence of each type of component acts to spread the impedance response over the Γ plane, with a change in the trace curvature in a concatenated network.

network at a given frequency will be diminished for a frequency that more closely corresponds to perfect impedance matching at the centre of the Smith chart. Thus nominally, a sharply well-matched system can be connected to an upstream standardised RF element, such as an amplifier that expects an output at $50\,\Omega$, by a transmission line of arbitrary length with minimal disruption to the matching characteristics. However, the wideband match ultimately used for the multiple-RF dressed potentials possesses an inherently less localised impedance response and is therefore more sensitive to post-match transmission lines. A short transmission line was therefore incorporated into the experimental matching network tested by the VNA and thus accounted for in the experimental fine-tuning stage. Transmission lines can also be connected in parallel to create a stub tuning network [181].

5.3 Characteristics of an impedance match

A number of metrics can be used to quantify an impedance matching network, and the choice of metric depends on the chosen application.

5.3.1 Quality factor

The loaded quality factor (Q, or Q-factor) of a circuit[4] is a measure of the bandwidth of the impedance match, and is described by the usual definition [182]:

$$Q = \frac{\omega_R}{\Delta\omega} \tag{5.9}$$

with ω_R the angular frequency at which reflections are minimised and $\Delta\omega$ the bandwidth as defined by -3 dB transmission. This can be equivalently defined according to

$$Q = \frac{|X|}{R} = \frac{|B|}{G}. \tag{5.10}$$

The appropriate Q-factor depends on the desired frequency response. In this experimental work, a large bandwidth, and associated high tolerance to variations in component values, were desirable, and as such the Q-factor kept deliberately low.

Again, the Smith chart provides an intuitive visualisation of the Q-factor to help guide the matching procedure, by relating Q to the reflection coefficient Γ [179]. Given the definition of Q in 5.10 and the definition of Γ in Eq. 5.2, the real and imaginary parts of the reflection coefficient can be expressed in terms of Q and normalised resistance $r = R_L/R_0$:

$$\Gamma_R = \frac{\frac{1}{r^2} - 1 - Q^2}{\left(\frac{1}{r} + 1\right)^2 + Q^2} \tag{5.11}$$

---

[4] The loaded Q-factor is distinct from the unloaded Q-factor: the latter relates only to the Q-factor of individual components. This quantifies losses in components and is given by the ratio of reactance to resistance [182].

and

$$\Gamma_I = \frac{\mp 2Q\frac{1}{r}}{\left(\frac{1}{r}+1\right)^2 + Q^2}. \tag{5.12}$$

These expressions describe the response of the reflection coefficient to a variation in the real part of the complex impedance for a given value of Q, and thus form of lines of constant Q as displayed on the Smith chart. Equations 5.11 and 5.12 can be rearranged to give:

$$\Gamma_R^2 + \Gamma_I^2 \mp \frac{2}{Q}\Gamma_I - 1 = 0. \tag{5.13}$$

Equation 5.13 describes two circles, both with radius $r_Q = \sqrt{1+\frac{1}{Q^2}}$, centred at $(0, \pm i/Q)$. As can be seen from eqs. 5.11 and 5.12, and the definition of Q in Eq. 5.10, the iso-Q lines are symmetric both with respect to reflection in the real axis, and the $180°$ rotation that maps to the admittance chart. These lines therefore take the form illustrated in Fig. 5.4, where the real axis corresponds to $Q = 0$, and the bounding circle of unity gain to $Q = \infty$.

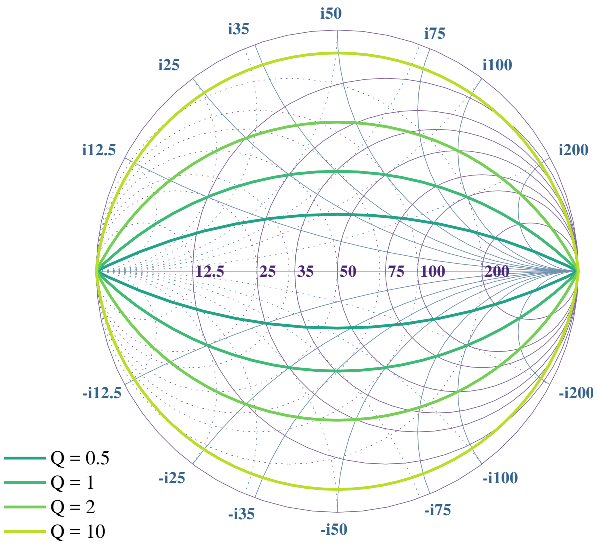

Figure 5.4: Arcs of constant Q-factor represented on the Smith chart.

The Q-factor of the final match is determined by the largest value of Q within which the $\Gamma(\omega)$ trace of the network lies. A high-Q match therefore requires a trace that traverses the outer regions of the Smith chart, and conversely a low-Q match needs a path constrained to the region close to its centre. This informs the choice of matching network components. For example, to keep Q as low as possible it can be preferable to create a matching network from a larger number of components that each correspond to a small translation on the Smith chart.

5.3.2 Voltage standing wave ratio

Interference between the forward-propagating and reflected waves establishes a standing wave in the transmission line, with the imbalance in the amplitude of the two signals manifesting as a remnant travelling wave. The voltage standing wave ratio (VSWR) [181, 182] is defined as the ratio of the maximum to minimum voltage on this line:

$$\text{VSWR} = \frac{|V_{\text{max}}|}{|V_{\text{min}}|} = \frac{|V_f| + |V_r|}{|V_f| - |V_r|} \tag{5.14}$$

where $|V_f|$ and $|V_r|$ are the amplitudes of the forward-propagating and reflected voltages. Using the definition of Γ as the fraction of an incident wave reflected at a boundary, this can also be expressed in terms of the reflection coefficient:

$$\text{VSWR} = \frac{1 + |\Gamma|}{1 - |\Gamma|}. \tag{5.15}$$

This connection to the reflection coefficient thus allows the VSWR to be used to describe the quality of a match at a given frequency. VSWR $= 1$ for $\Gamma = 0$ and thus a perfect impedance match, which tends to ∞ for $\Gamma = 1$ where the forward and reflected waves have amplitudes that exactly cancel.

5.3.3 Return loss

The return loss is a description of the power lost due to reflections at a boundary. This is simply given by the reflection coefficient Γ expressed in dB according to

$$\mathrm{RL}(\omega) = -20\ln(\Gamma(\omega)) \tag{5.16}$$

where $|\Gamma|^2$ gives the fraction of the incident power reflected at the boundary. A low value of the return loss thus corresponds to a small reflection coefficient and greater power transfer. This can be plotted as a function of frequency to complement the Smith chart analysis and easily visualise the depth and width of the matching resonances. In the analysis of the experimental system that follows, both the Smith chart and a plot of the return loss are plotted to characterise the impedance match.

5.4 The experimental network

The dressing RF is applied to the atoms using the orthogonal pairs of racetrack coils described in Sec. 4.4.3. The x and y components are generated independently by direct digital synthesis (DDS), to form independent but analogous networks that differ only in the details of component choices to compensate for slightly different coil characteristics. A schematic of the RF network is illustrated in Fig. 5.5. The RF source takes the form of a series of DDS boards (*Analog Devices AD9854*), each generating one frequency component and sharing a common clock[5]. The maximum peak-to-peak output voltage of each channel is 0.4 V. These independently-generated frequency components are combined at a bank of two-way splitters (*Mini-Circuits ZSC-2-2*), before being passed through a 1.2 MHz high-pass filter (*Mini-Circuits ZFHP-1R2-S+*) to remove any stray low-frequency components, and into a $25\,\mathrm{W}$ amplifier (*Mini-*

---

[5] The details of the clocking procedure for the different measurements presented in this work are discussed in Sec. 6.1.1.

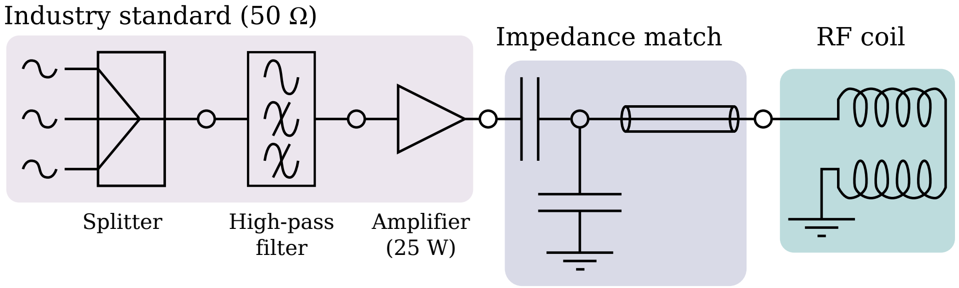

Figure 5.5: Schematic of the RF network. The x and y networks differ only in small differences in coil construction and cable lengths, with correspondingly tuned impedance matching parameters. The splitter is used here to combine the three RF signals used to form the multiple-RF field, each of which are created by independent DDS channels with a maximum signal amplitude of 0.4 V peak-to-peak. Industry standard components are impedance matched to $50\,\Omega$, and so a custom match is required to match this output to the home-wound RF coils.

Circuits ZLY-22+). These components all comply with the $50\,\Omega$ industry standard, and as such can be connected without additional impedance matching. However, a match is required to mediate between the amplifier output and the home-wound RF coil pair. This match is calculated based on the coil impedance characteristics as described in the following sections.

Matches were developed according to both their reflection characteristics and the current that flows through the RF coils, that in turn determines the RF field strength. The precise current required has varied substantially over time and between different applications. However, as an initial estimate that would allow a sufficient margin for early investigations, a target current of 150 mA amplitude was chosen, corresponding to an RF field amplitude around $0.5\,\mathrm{G}$.

A Smith chart and return loss plot characterising the bare dressing RF coils are shown in Fig. 5.5. In these and subsequent analyses, the measured frequency sweep covers the range 100 kHz to 10 MHz.

The frequency range of the VNA sweeps out a significant variation in both impedance and the corresponding reflection coefficient, although at no point does it naturally offer a good match with low reflection. The self-resonance (SR) is the frequency at which the reflection coefficient trace crosses the real axis, its nature

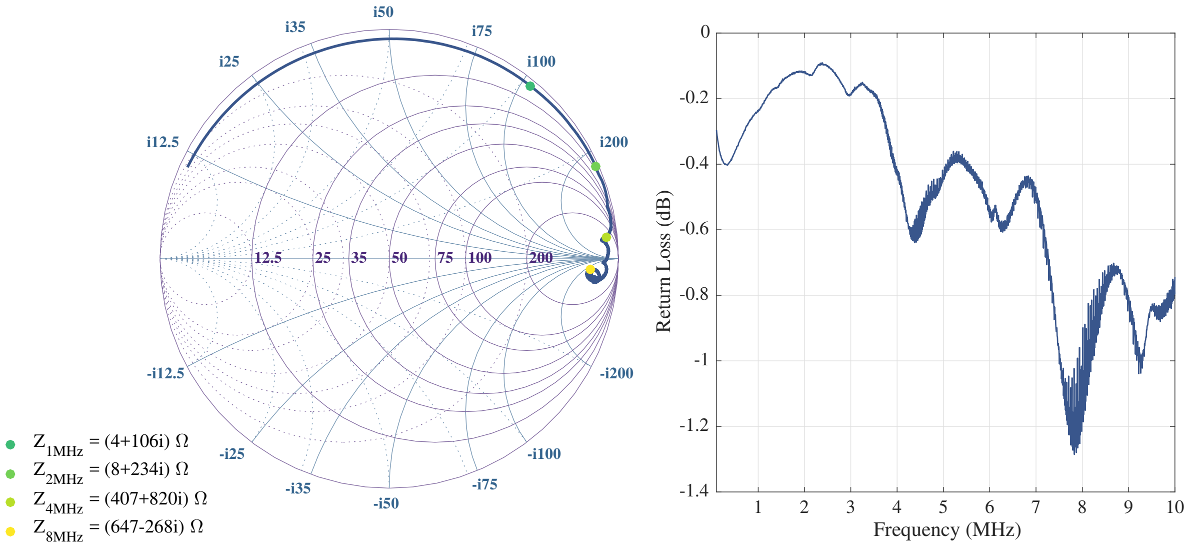

(a) Unmatched *x* coil Smith chart **(b)** Unmatched *x* coil return loss

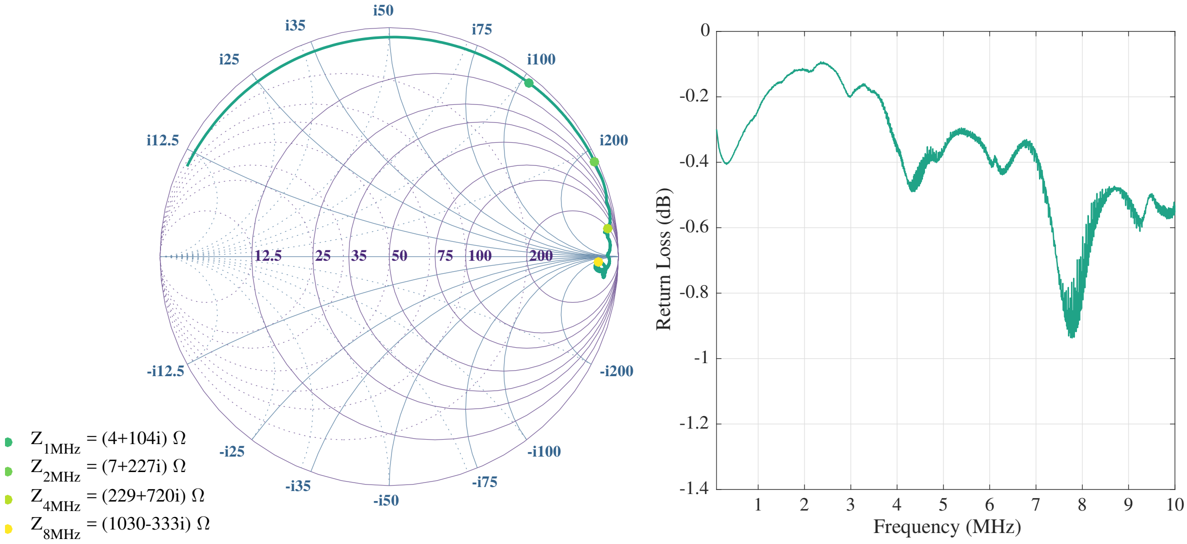

(c) Unmatched *y* coil Smith chart **(d)** Unmatched *y* coil return loss

Figure 5.6: Unmatched RF coil impedance characteristics. (a) The reflection coefficient of the *x* coils as measured using a VNA, plotted on the Smith chart. The network impedance is labelled at a series of frequencies to illustrate the velocity of the trace as the frequency is varied. (b) The corresponding return loss over the measured frequency range. The rapid variations on the return loss response can also be distinguished in the variation in the Smith chart trace in the spiralling region towards the end of the trace. (c) and (d) present the same information for the *y* coil, which is superficially similar, though with differences that become important during the matching process. The impedance is poorly matched over this full range and reflections are prohibitively large without impedance matching.

transforming from capacitive to inductive at this point; this provides an upper limit to the range of useable frequencies. The self-resonances of the two RF coils extracted from these VNA measurements are given by $\mathrm{SR}_x = (5.77 \pm 0.06)$ MHz and $\mathrm{SR}_y = (6.51 \pm 0.06)$ MHz. The uncertainty in each measurement reflects the noise in the measured impedance trace, and was estimated by considering the range of frequencies that lie within the $\pm 75\,\Omega$ reactance arcs that converge at the open-circuit side of the Smith chart.

The two traces are superficially similar, but their differences make a surprisingly substantial difference in the details of implementing the impedance match. The following sections discuss the experimental approaches to impedance matching these RF coils for both narrow-band and wideband operation, and some additional approaches are shown in Appendix C. These matches were calculated and implemented by a combination of logic and intuition, based on an understanding on the impedance transformations associated with different components. As such, they are almost certainly not a uniquely perfect match in each case, but I will discuss some of the benefits and drawbacks of each method used in relation to our experimental requirements.

5.5 Narrow-band impedance matching

To create a single-RF shell with dressing angular frequency $\omega = 2\pi \times 2$ MHz, a narrow-band impedance match was centred on this frequency, with the goal of maximising the quality factor and narrowing the match to minimise the power requirements and exclude spurious signals. A purely capacitive network was used because of the availability of a wide range of commercial component values.

5.5.1 Simulated matching network

The matching circuits designed, and their simulated reflection characteristics, are shown in Fig. 5.7. The circuits are extremely similar in design, each consisting of a shunt capacitance to transform the 2 MHz point $\Gamma(\omega)$ along the constant reactance circle to reach the circle of constant conductance $Y = 1/(50\,\Omega)$. A series capacitance is then employed to traverse this iso-conductance circle to reach the $50\,\Omega$ origin. By ensuring that the $\Gamma(\omega)$ trace lies predominantly on the outer edge of the Smith chart, making only a brief diversion to the central region, a high Q and small bandwidth is maintained.

5.5.2 Experimental matching network

The simulated matching circuit provides a good starting point to the experimental match. However, discrepancies arise between the simulated and experimental systems due to a number of factors including imperfections or inaccuracies in component values, and stray impedances arising from construction and soldering. Following construction of the modelled matching network, the impedance match is therefore fine-tuned using a VNA to provide live diagnostics. The final matching networks and associated reflection characteristics for the 2 MHz narrowband match are shown in Fig. 5.8.

The impedance matches are built in shielded boxes, using ceramic through-hole capacitors soldered together in free space. Adjustments can therefore be made easily, although since even small stray impedances can make a significant difference to the resultant match, particularly in a frequency range with a sharp response, the accuracy of the final match is limited by the tolerance to slight changes in soldering or the orientation of components. While the match is sensitive to these effects during construction, the effect of this is predominantly limited to the details of the resonance peak, and furthermore the match is stable once the box is sealed and undisturbed.

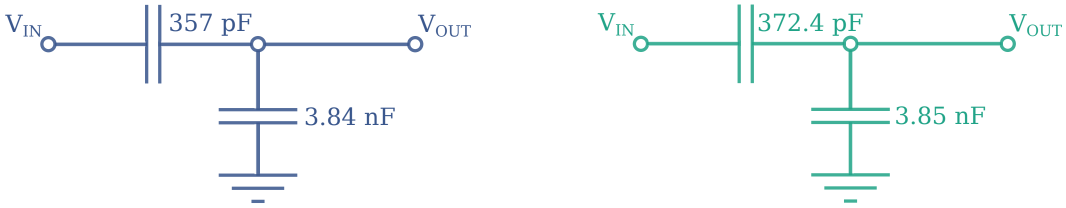

(a) 2 MHz *x* matching network **(b)** 2 MHz *y* matching network

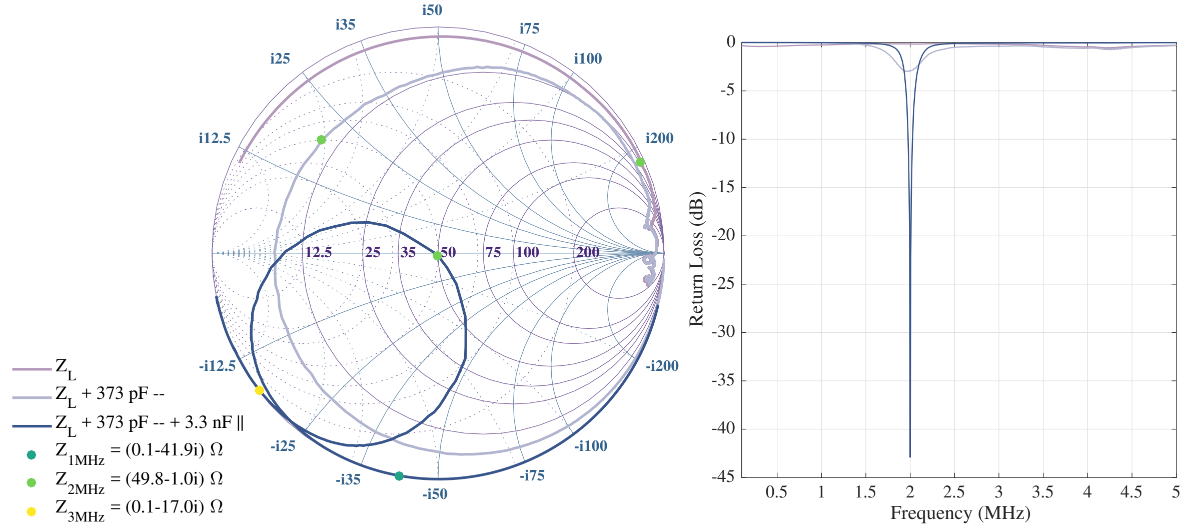

(c) Modelled Smith chart *x* **(d)** Modelled return loss plot *x*

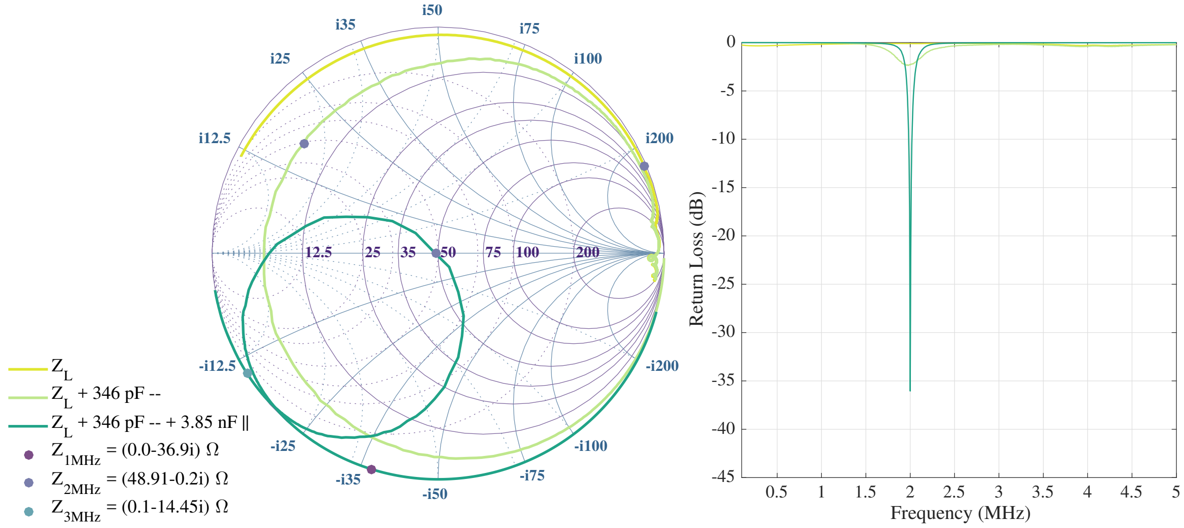

(e) Modelled Smith chart *y* **(f)** Modelled return loss plot *y*

Figure 5.7: Modelled 2 MHz narrowband match. (a),(b) A sketch of the 2 MHz narrowband matching networks for the *x* and *y* coils respectively. These both consist of a shunt and series capacitor. (c) Smith chart showing the reflectance trace $\Gamma(\omega)$ after the addition of each component in the matching network. A narrow match is ensured by the high Q-factor and smooth sweep. The impedance values at 1, 2 and 3 MHz have been labelled to illustrate the scaling of the sweep in the vicinity of the matching frequency. Markers have been placed at the 2 MHz point on each trace to track the impedance transformation associated with each component added (impedance value is only listed for the final trace). (d) Return loss as a function of the frequency sweep plotted on the Smith chart, after the addition of each component in the matching network. (e), (f) show the same information for the *y* network.

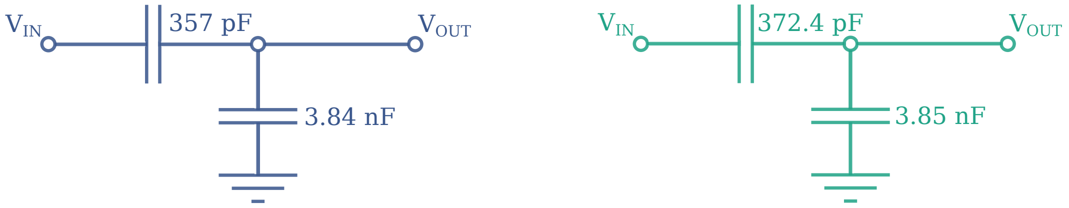

(a) 2 MHz *x* matching network **(b)** 2 MHz *y* matching network

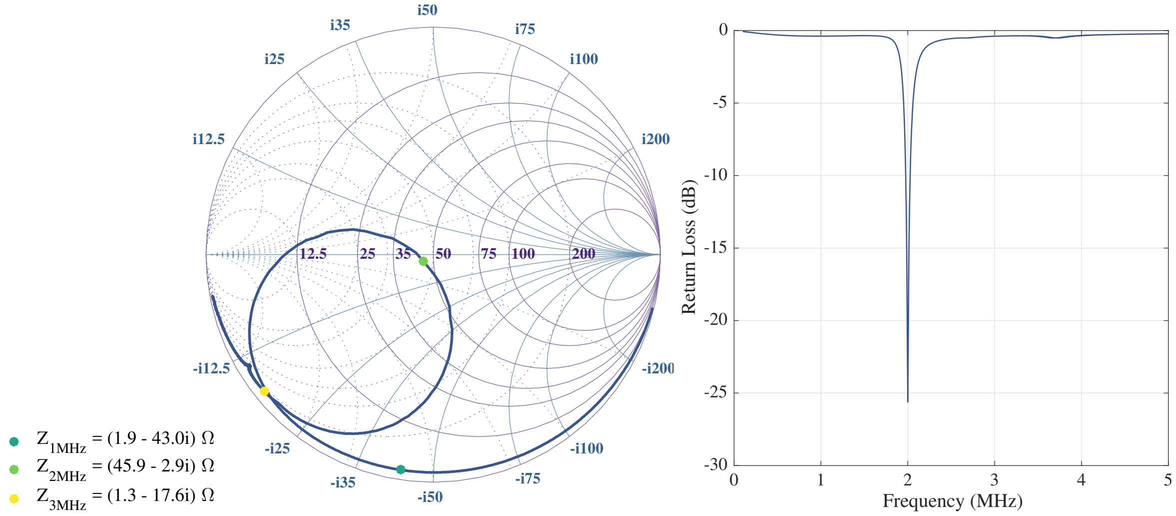

(c) Measured Smith chart *x* **(d)** Measured return loss plot *x*

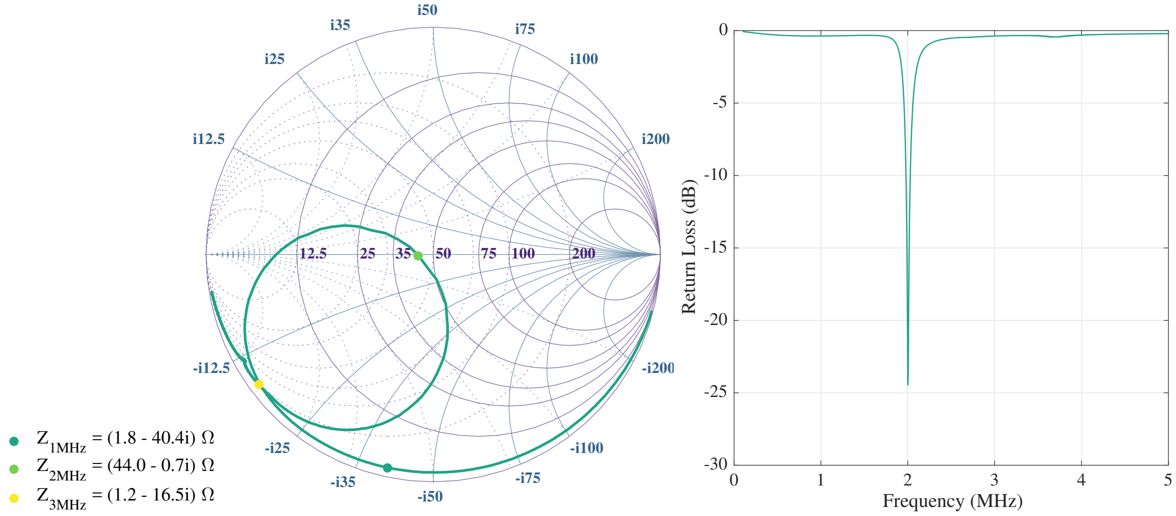

(e) Measured Smith chart *y* **(f)** Measured return loss plot *y*

Figure 5.8: Experimental 2 MHz narrowband match. (a),(b) A sketch of the 2 MHz narrowband match for the *x* and *y* coils respectively. The stated capacitor values indicate the combined capacitance of several components and do not take into account component variations, or stray resistance or capacitance arising from circuit construction. These both consist of a shunt followed by a series capacitor. (c) Smith chart showing the measured reflectance trace $\Gamma(\omega)$ after matching. The impedance values at 1, 2 and 3 MHz have been labelled for comparison of the sweep scaling with the simulated circuit. (d) Return loss as a function of the frequency sweep plotted on the Smith chart. (e), (f) show the same information for the *y* network. Live adjustments were made during circuit construction, based on both Smith chart and return loss plots, but experimental influences made the sharp resonance difficult to match perfectly.

Stability is therefore not a concern as long as there is no significant disturbance to the matching circuit. Both reflection and current response has been stable when tested at intervals ranging between hours and weeks, and within temperature variations of a few degrees as would be typical of standard operation in the lab. However, the sensitivity of the match, translating directly to the stability of the dressing RF field, while not problematic in standard operation, is a potential weakness of the current design. Methods necessary to stabilise the wideband match are discussed in Sec. 5.6, and an alternative approach to circuit construction in Sec. 5.7. The capacitances quoted are obtained from the component values and as such do not account for stray effects.

Taking into account dissipation in the circuit, the reflectance does not map precisely to the current transmitted through the coils, as measured using AC current probes (*Tektronix 6019* and *6021* with adaptors *011-0089-00* and *011-0155-00*) and illustrated

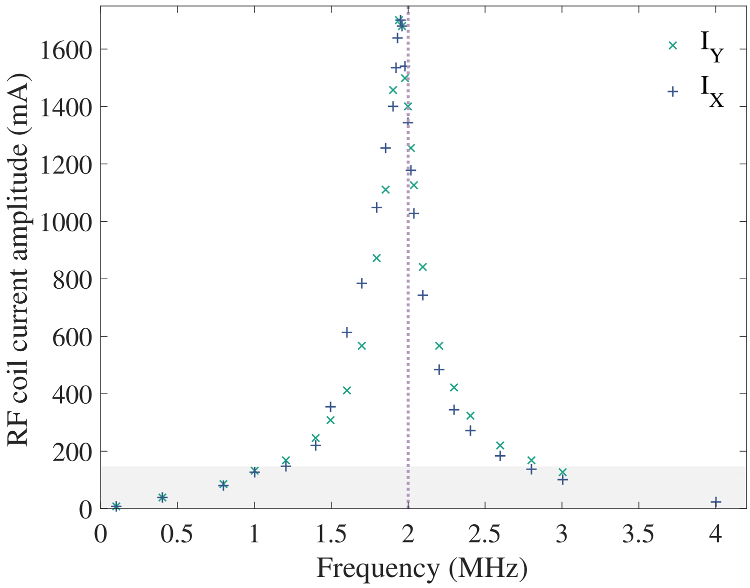

Figure 5.9: Measured current $I(\omega)$ through the narrowband-matched coils with the maximum DDS output of 0.4 V peak-to-peak. The vertical dotted line indicates the target frequency at 2 MHz. The resonance peak is slightly offset from this point since the sharpness of the match made precise tuning about this point very challenging. The sensitivity of the network to minor construction adjustments limited the precision of placing the peak at around 100 kHz. The shaded region indicates the 147 mA target current amplitude.

in Fig. 5.9[6]. The slight offset in the current peak from the 2 MHz target dressing frequency reflects the difficulty in precisely tuning the match beyond the response of the network to slight variations in circuit construction. The discrepancy between the frequency response of the transmitted current and measured reflection can be accounted for by the effects of power dissipation. However, this network is sufficient: current transmission remains above the estimated diagnostic requirements over a range of approximately 1.5 MHz. A narrower match was not sought as a higher gradient has the potential to transform between acoustic and phase noise, and the match is narrow enough to suppress the majority of noise in the system.

The lifetime of thermal atoms with temperature approximately 10 $\mu\mathrm{K}$ in the $\omega_1 = 2\pi \times 2$ MHz shell with $\Omega_1 \approx 300$ kHz was on the order of several minutes, measured as the time at which the atom number falls to $1/e$ of its original value[7].

5.6 Wideband impedance matching

In order to implement the MRF dressed potential, an impedance match was required that would allow sufficient current to flow in the dressing RF coils to produce the desired Rabi frequency for atoms confined in shells at each dressing frequency of interest. These parameters are not fixed, but were chosen for the exploratory work to be a current of 147 mA as above, and dressing angular frequencies $\omega_{1,2,3} = 2\pi \times (3, 3.6, 4.2)$ MHz. An angular frequency separation of $2\pi \times 0.6$ MHz was initially chosen in part because the resultant shell separation at a quadrupole current

---

[6] These current probe readings were sufficiently accurate for diagnostic purposes but experienced a variation on the order of a few percent when the BNC connection to the oscilloscope was adjusted. As a result, these current measurements are used for guidance rather than calibration. Subsequent detailed work with the MRF potentials is monitored using calibrated pickup coils co-wound with the dressing RF coils, with details of RF monitoring and calibration discussed in Sec. 6.2.2.

[7] A more precise value of the lifetime was difficult to obtain as we were limited by the maximum length of the experiment sequence. This data was obtained in early 2015, at a time of experiment upheaval and prior to the systematic noise source elimination process. As focus had moved towards multiple field components by this time, the $\omega_1 = 2\pi \times 2$ MHz lifetime was not measured again. However, we would expect this to increase slightly as damaging noise sources were removed.

of 20 A is on the order of $100\,\mu\mathrm{m}$ and thus easily resolvable with our low-resolution horizontal imaging system. Another compelling reason for this combination of parameters was that Rabi frequencies of a few hundred kHz were the most familiar and easiest to work with in our single-shell experiments. With this combination of parameters, the MRF double well has distinct wells with a good tuning range and higher-order effects that are not too intrusive[8]. The match was required to extend to as high a frequency as possible to facilitate easier access to the 2D regime. Furthermore, we aimed to retain as much flexibility as possible to vary trapping parameters during the dianostic and exploratory process. A variety of strategies could be employed here, including no matching, a wideband match, or a multiple-band match.

A simple solution to this could be to just leave the coils unmatched and work in the fairly uniform high-reflection regime, using robust RF amplifiers. However, this strategy requires the flexibility to waste large amounts of RF power. The typical dressing RF amplitude used for investigative and diagnostic work on the MRF potential was on the order of $\Omega_i = 2\pi \times 300$ kHz. In order to access these field amplitudes with sufficient clearance to probe the parameter space and without additional preamplification, a current with approximate amplitude 120 mA was required. However, reflections from the unmatched coils would limit this current to the region of 8 mA. The use of preamplifiers to boost this signal is constrained by the 20 dBm (3.17 V amplitude) input limit.

A wideband match could instead be used to reduce the reflections around the frequency range of interest; for flexibility this would span approximately a 2 MHz range centred at 3.6 MHz. However, this match would ideally not span too wide a frequency range to minimise transmission of unwanted noise arising from mixing processes in the amplifiers, or low-frequency noise from external sources such as ground loops. The challenge of this method is in maintaining a sufficiently low *Q*, which is made more difficult by the extended frequency response of the reflection trace $\Gamma(\omega)$ of the

---

[8] The form of the MRF potentials is discussed in detail in Sec. 6.3.

load. A multiple-stage network is generally better suited to low-Q networks, while careful component selection could also compress the frequency response. A complicated capacitor network is, however, challenging to implement and tune in practice.

A narrow multiple-band match eliminates the problem of noise transmission, and one example of this is shown in Appendix C.3. The general multiple-band method tested here was to use a long transmission line to extend the frequency response of the $\Gamma(\omega)$ trace such that it repeatedly crosses the real axis in the vicinity of the $50\,\Omega$ origin. Each such crossing corresponds to a transmission resonance. However, these multiple narrow peaks are even more difficult to tune in coordination than the single narrow match described above. A wide match also allows more flexibility during the diagnostic stages in which a variety of frequency combinations were tested. A multiple narrowband match might therefore be an appropriate method only once a final collection of frequencies has been chosen for subsequent investigations.

The final wideband matching network is outlined below, and was used to obtain all the MRF results in the following chapters. For completeness, some other illustrative approaches to this problem are discussed in Appendix C. This includes an early approach to wideband matching that was unsuccessful due to high transmission at low frequencies, severely limiting the atom lifetime. The final method presented here is a combination of the multiple-band and wideband approaches: multiple resonance peaks are used to widen the match, but the accessible frequency range is not constrained to these resonant values.

5.6.1 Simulated wideband matching network

The fundamental components of the wideband impedance match include a series capacitor to block low frequencies, possibly some shunt capacitance to retrospectively adjust the frequency response, a long transmission line to extend the frequency response and induce multiple real-axis crossings, and a shunt capacitance to tune the crossing points. As it is easier to experimentally adjust capacitor values than transmission line lengths, the tuning of the match is performed using the two sets of capacitances, and the transmission line held fixed at an appropriate value based on measured cable lengths. This is slightly restrictive, but reasonable within the flexibility of the wideband match, and could be remedied simply by trimming cables if necessary. Given the width of the match, the 0.5 m BNC cable used to connect the amplifier output to matching boxes was also incorporated into the model. The modelled matching networks for each of the x and y coils is shown in Fig. 5.10 alongside the modelled reflection characteristics.

A network with multiple stages could decrease the Q-factor of the match, but calculations during the design stage suggested that the gains from this do not offset the increase in experimental complexity. Furthermore, the broad frequency response of the load makes a low Q difficult to achieve. Instead, the response is extended using the long transmission line to create multiple real-axis crossings, the details of which can be tuned by the surrounding capacitive stages. The two transmission bands correspond to the real-axis crossings near the $50\,\Omega$ Smith chart origin. The frequency separation of these bands is small enough, and the bands themselves sufficiently wide, that the reflection coefficient is also reduced in the intermediate regime. As is apparent from both the Smith chart traces and return loss plots, the two bands are significantly more closely matched for the y network than x. However, the simulated networks were not fully optimised in anticipation of the experimental fine-tuning necessary due to discrepancies in component values.

5.6.2 Experimental wideband matching network

The increased complexity of the wideband matching network relative to the narrowband match means that there are more soldered connections and uncertainty in component values, all of which affect the impedance of the network. The long trans-

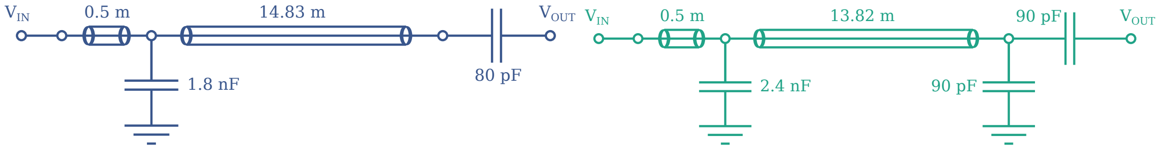

(a) Wideband x matching network **(b)** Wideband y matching network

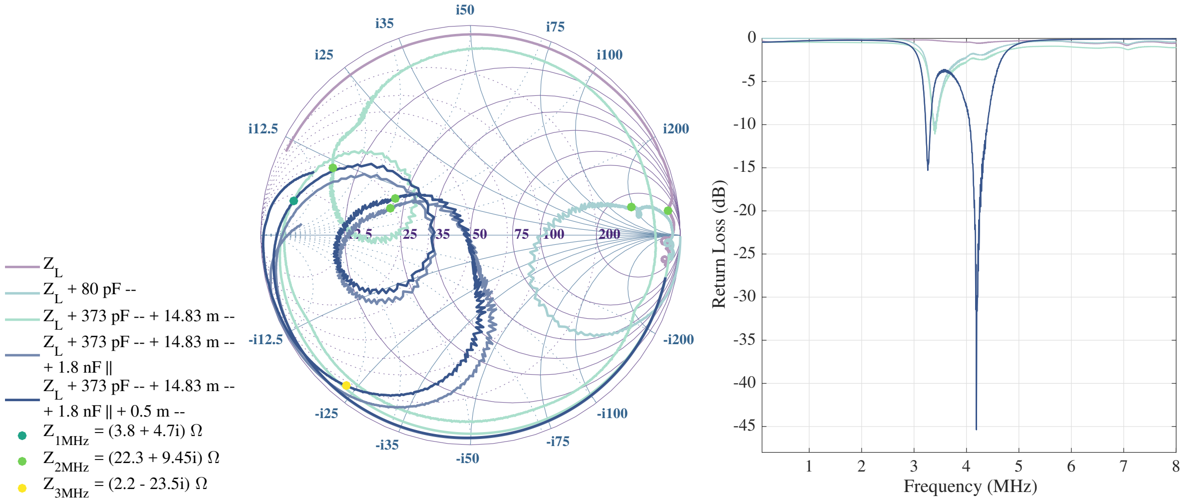

(c) Modelled Smith chart x **(d)** Modelled return loss plot x

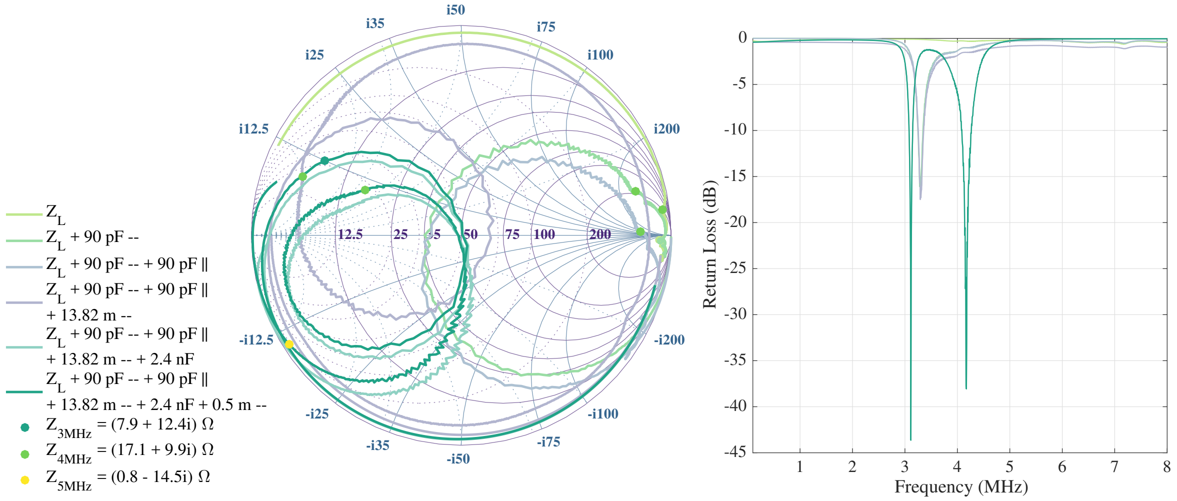

(e) Modelled Smith chart y **(f)** Modelled return loss plot y

Figure 5.10: Modelled 4 MHz wideband match. (a),(b) A sketch of the wideband matching networks for the x and y coils respectively. A series capacitor next to the load blocks low-frequency signals, followed by a combination of long transmission lines and shunt capacitors to tune the match. (c) Smith chart showing $\Gamma(\omega)$ after the addition of each component. Each real-axis crossing in the vicinity of $50\,\Omega$ corresponds to a transmission peak. The Q-factor is as low as possible while minimising the number of network components. The impedance values at 3, 4 and 5 MHz have been labelled to illustrate the scaling of the sweep. Markers have been placed at 4 MHz on each trace to track the impedance transformation as each component is added (impedance value is only listed for the final trace). (d) Return loss as a function of the frequency sweep plotted on the Smith chart, after the addition of each component in the matching network. Two distinct peaks are apparent, corresponding to the real axis crossings near the origin. (e), (f) show the same information for the y network. Dual-band matching has been more successful in y than x.

mission lines are constructed from lengths of BNC cable and have an uncertainty of approximately $\pm 10\,\mathrm{cm}$, although the effect of this can be mitigated using the capacitive stages. Experimental live adjustment was therefore more significant for the wideband than narrow networks. The final experimental networks, where capacitors are labelled by the sum of constituent component values and do not take stray construction effects into account, are shown in Fig. 5.11 alongside the corresponding measured Smith and reflectance plots.

The peaks of greatest transmission are centred around 3.4 MHz in both cases, with a secondary peak at 4.5 MHz (x) and 4.3 MHz (y). The x network also experiences a small peak in transmission at 1.3 MHz. By positioning a peak at the highest end of the frequency range of interest rather than relying on the falloff in transmission of a peak closer to the centre of this range, the current suppression associated with increasing proximity to the self-resonance is compensated. As can be seen in both the Smith and return loss plots, balanced $50\,\Omega$ dual-band matching has been more successful in the y network than x. This is not important for the application considered here, but is a useful demonstration that this method is possible and could be applied to a true multiple-band match in future. The close spacing of the transmission peaks means that a region of reduced reflection also exists in the intermediate frequencies. The sharp features in the reflection plot are smoothed further in the measurement of the transmitted current $I(\omega)$, shown in Fig. 5.12.

The current response is noticeably smoothed in comparison to the return loss plots, forming a single broad peak with a 2.5 MHz range above the target current, comfortably encompassing the three target dressing frequencies. This could arise due to a combination of the proximity to the self-resonant frequency of the RF dressing coils and dissipation in the network. The absence of a second distinct peak above 4 MHz can be attributed to the current suppression associated with proximity to the self-resonance, although the reduction in reflections at these higher frequencies works to

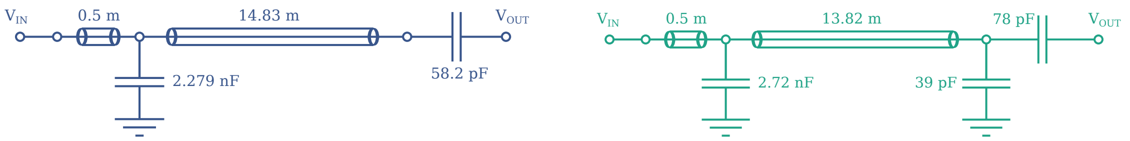

(a) Wideband *x* network **(b)** Wideband *y* network

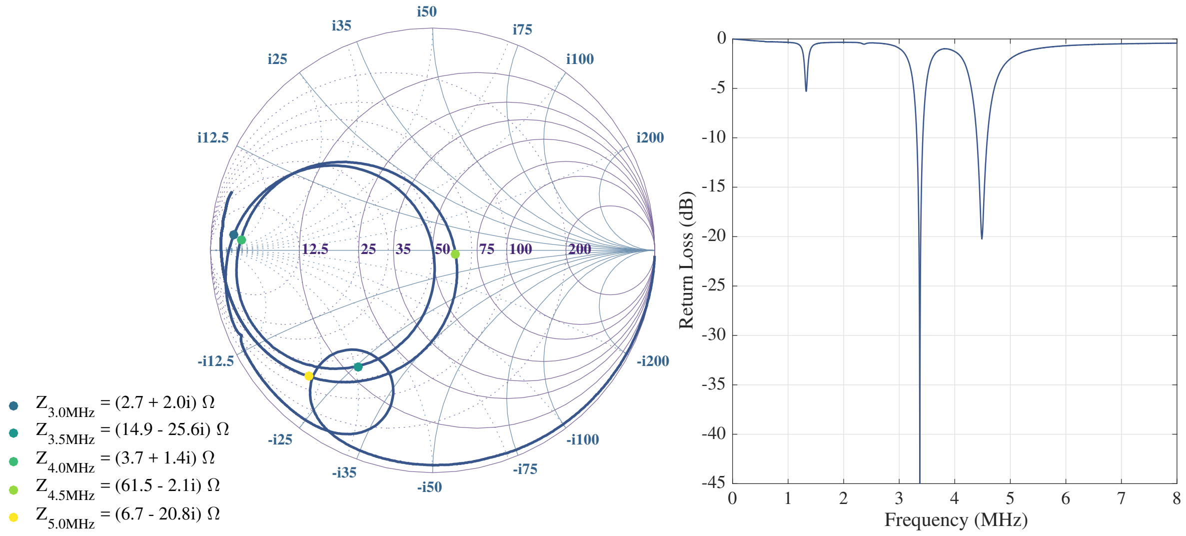

(c) Measured Smith chart *x* **(d)** Measured return loss plot *x*

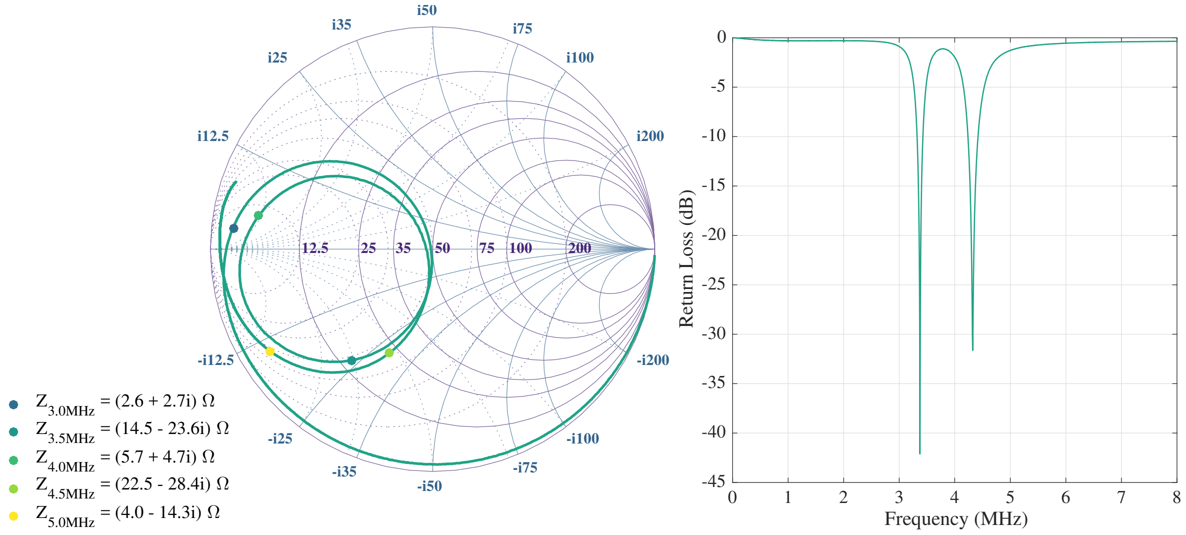

(e) Measured Smith chart *y* **(f)** Measured return loss plot *y*

Figure 5.11: Experimental 4 MHz wideband match. (a),(b) The wideband matching networks for the *x* and *y* coils. The stated capacitor values indicate the combined capacitance of several components and does not take into account variation or the effects of circuit construction. The capacitive stages before and after the transmission line are coordinated to tune the match. (c) Smith chart showing the measured reflectance trace $\Gamma(\omega)$. The impedance values at 500 kHz intervals between 3 MHz and 5 MHz are labelled to indicate the scaling of $\Gamma(\omega)$ and allow comparison with the return loss plot. (d) Return loss, with transmission peaks corresponding to the Smith chart real-axis crossings. The series capacitance next to the load blocks low-frequency transmission. (e), (f) show the same information for the *y* network. Live adjustments were made during circuit construction, based on both Smith chart and return loss plots.

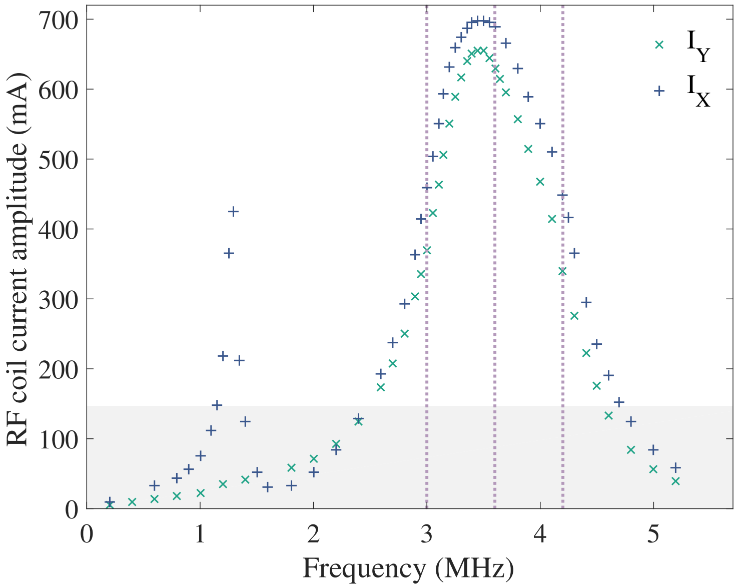

Figure 5.12: Measured $I(\omega)$ through the wideband-matched coils with the maximum DDS output of 0.4 V peak-to-peak. The vertical dotted lines indicates the target RF dressing angular frequencies $\omega_{1,2,3} = 2\pi \times (3, 3.6, 4.2)\text{MHz}$. The shaded region indicates the 147 mA target current amplitude. The effects of current dissipation have caused a significant discrepancy compared to the reflection plots, smoothing the current in the intermediate frequency regime to create a wideband response. The 1.3 MHz transmission peak in the *x* network is evident. The effects of current suppression due to the coil self-resonances are apparent at the higher frequencies, despite the transmission peaks above 4 MHz.

broaden the current response[9]. The transmission peak in the *x* network at 1.3 MHz is apparent in the measured current, but is neither exploited for, nor damaging to, the work considered here. This demonstrates the potential for multiple-band matching at frequencies that are sufficiently separated relative to the individual match bandwidths.

Within the input power limits of the amplifiers, the accessible frequency range can be increased to approximately 5 MHz in each coil. Even for a low input power, distortion of the signal is apparent in the *y* coil above 5 MHz, and in the range 500 kHz to 1.6 MHz. Distortion was seen in the *x* coils around 1.8 MHz.

To ensure the stability of the match, the long BNC cables were coiled and set in glue, and capacitor network boxes secured in place. Prior to being set, the impedance of the network varied slightly (the current response at the transmission peak changing

---

[9] A slight bump in the right hand edge of the *x* current plot could be attributed to the higher-frequency resonance.

on the order of a few percent) according to the orientation of the coiled transmission line, and separation between turns. While the dielectric properties of the setting glue affected the matching characteristics, the cable layout was subtly adjusted while it set to recover the original network properties. Relative to the previous current response, the amplitude of the primary current peaks differ by 0.1% for the x network and 2.2% for y but the overall shape of the peak underwent a negligible change. The only significant difference was in the x network 1.3 MHz peak, as a result of the associated small loop in the reflection coefficient as seen on the Smith chart. Current and reflection measurements have both been repeated at various intervals ranging between hours and weeks, with varying ambient temperatures and with artificially long applied RF durations, with a negligible difference detected. This match is therefore sufficiently stable for standard operation, and coarse calibration of experimental field amplitudes. For later work and measurements in which an accurate knowledge of field amplitude is necessary, calibrated co-wound pickup coils are employed.

The lifetime of thermal atoms confined in a single-frequency shell created using this matching network is on the order of a few seconds. This places no restrictions on the experiments considered in this work, in contrast to the earlier version of the match discussed in Appendix C.1, but is significantly shorter than the corresponding lifetime in the narrowband match discussed in Sec. 5.5. This can be attributed to the suppression of noise by the narrowband matching network. Since these observations were made, a large number of RF noise sources have been eliminated, and it is likely from observations of the evaporation process, for example, that the trap lifetime has increased as a result. However, lifetime measurements have not been formally repeated since lifetime is not a significant concern in the current work. A longer lifetime could be obtained for future experimental work by reducing the width of the match to the minimum range necessary to block extraneous frequencies, especially in the case of a reduced MRF frequency separation.

5.7 Developing the impedance match

The wideband impedance match is adequate for current experimental requirements, and was particularly well suited to the exploratory and diagnostic stages. However, improvements could be made to both the characteristics of the match and stability of the design. Firstly, the disparity in lifetimes between atoms confined in shells created using the narrowband and wideband matches indicates that it could be useful to tailor an individual matching network to each experiment. Examples of this include narrow multiple-band matches for the well-separated frequencies used in state- or species-selective manipulation, or simply matches that are tuned more sharply to the frequency range under consideration. The precision of the network design could be improved by actively trimming the transmission lines to a specified length rather than using capacitors to tune the match. Tuneable matching would also be possible, for example by incorporating a tuneable element such as an air-spaced capacitor, which would be useful to improve the flexibility of a narrow match but might result in instabilities problematic for a narrow bandwidth.

A tailored approach to matching would require an improved design stability: the existing setup is robust against day-to-day disturbances and moderate adjustments to the positions of the matching boxes, but would be unlikely to withstand regular adjustments. An alternative approach trialled was to construct the match on PCBs housed within shielded boxes. Unfortunately, the match properties deviated substantially from expectations, most likely because standard PCBs were used, for example with no ground plane. Although this method was laid aside, it remains an option for the future. An RF PCB could also offer a more stable approach to transmission lines, which could also be easier to manage if created using smaller-radius cables.

Despite the potential for improvements, the methods used to create the impedance matches discussed in this chapter have the benefits of convenience, flexibility, and ease of adjustment. The matching networks created have proved stable over the lifetime of the experimental work on MRF dressing, producing shell traps with atom lifetime and heating rates more than adequate for the experiments of interest. Crucially, these have also provided us with the experimental flexibility to extensively explore the dressing RF parameter space, and develop methods to load, manipulate, and characterise the resultant potentials.

Chapter 6

Experimental implementation of multiple-RF dressed adiabatic potentials

This work describes the first experimental implementation of multiple-RF dressed adiabatic potentials. We have explored and characterised this method using the example of three independently-controlled dressing frequencies to create a highly configurable double well potential. With the experimental parameters used here, the resonant trapping surfaces take the form of two concentric spheroidal shells as introduced in Sec. 2.5.2, and atoms fall to the lowest points on each surface under the influence of gravity to create the double well. A high level of intuitive control is afforded over the potential landscape by independent variation of the relative RF amplitudes and polarisations. Here, a double well with large well spacing is created as a proof-of-principle testing ground to explore and characterise the MRF potential. Following on from this work, the well spacing can be reduced to investigate double well physics and interferometry in the 2D regime afforded by the AP shell structure, switching between different regimes using the RF amplitudes to vary the coupling strength. Other immediate extensions of the method include independent control of trapped species with different values of g_F, and the creation of further trapping wells and potential landscape structures using additional RF components.

This chapter explores the key experimental methods underpinning the creation of these potentials: methods to create and load the single-RF AP are explored in Sec. 6.1, characterisation of the dressing RF using atoms confined in these single-frequency traps presented in Sec. 6.2, and finally schemes for controlled atom loading into the double well presented in Sec. 6.3. The development of these techniques has taken place concurrently with significant improvements to the experimental apparatus. As such the methods presented here represent particular snaphots in time, but are still representative of the techniques currently implemented and used to obtain the results in Chapter 7.

6.1 Trapping atoms in the single well AP

While not our final objective, a single shell provides the means of characterising the dressing RF properties and resultant potentials, and is a necessary stage in loading the multiple-RF dressed traps. Atoms are loaded and cooled using the single-RF TAAP, and a BEC is subsequently loaded into the single-RF shell. This is used to explore experimental parameters and the atom response, and more formally to characterise the RF field amplitude and polarisation for every dressing frequency of interest. Finally, this potential can be deformed by the additional dressing fields to load atoms into the combined MRF potential.

6.1.1 Creating the AP

The RF signals are derived from independent DDS sources, and then combined and filtered as discussed in Chapter 5. Each field is associated with two DDS sources: one for each of the x and y components of each dressing frequency. The eight available DDS channels (one reserved for evaporative RF) are divided between two crates, each with its own back-plane and system clock; the DDS sources and the adaptations required to ensure stable MRF operation are discussed in forthcoming work [35]. The DDS channels associated with each frequency component share a common clock to maintain precise control of each field polarisation, and a single clock extended to all dressing RF channels for measurements in which the relative phase between each field component is important for experiment stability. The importance of this clocking process on the resulting MRF potential is addressed briefly in [104] and will be described in more detail in [35][1]. MRF dressing experiments commenced without a common clock for the two separate DDS crates. However, a shared clock for all dressing RF channels was subsequently enforced; the clock configuration is specified alongside experimental results where relevant.

We originally developed our understanding of AP loading and confinement using a dressing angular frequency of $\omega = 2\pi \times 2$ MHz and a narrow-band match, in order to work with a set of parameters fairly similar to the $\omega = 2\pi \times 1.4$ MHz APs in previous experiments on this apparatus [67, 147]. With a coupling strength of order $\Omega = 2\pi \times 400$ kHz, the lifetime of thermal atoms loaded at $\sim 10\ \mu\mathrm{K}$ was of order several minutes as described in Chapter 5. Loading techniques were developed using this AP, and subsequently transferred to those created at dressing angular frequencies

---

[1] For some diagnostic applications, including double shell loading, a common reference clock between the different field components is not imperative, although the changing relative phase between these components causes some experimental unreliability. However, for any application where precise knowledge of the energy level structure is important, these channels must share a common clock. The data in figures 6.6 and 7.1 were obtained using a common clock for the $\omega_{1,2} = 2\pi \times (3.0, 3.6)$ MHz dressing fields, and a separate clock for the $\omega_3 = 2\pi \times 4.2$ MHz signals; the RF spectroscopy measurements in Fig. 7.3 use a common clock for all RF field components.

between $\omega = 2\pi \times 3$ MHz and $\omega = 2\pi \times 4.5$ MHz, implemented using the wideband impedance match and with a comparable thermal atom cloud lifetime of a few seconds.

We predominantly used field amplitudes corresponding to coupling strengths of a few hundred kHz which give a three-dimensional gas of trapped atoms in a regime largely clear of RF noise. This is also a good match for the predominant angular frequency spacing of $\Delta\omega = 2\pi \times 600$ kHz. After careful reduction of RF noise, single-RF APs with a coupling strength of $\Omega \approx 2\pi \times 50$ kHz have also been realised, paving the way for MRF potentials with a reduced coupling strength and frequency separation, and for operation in the 2D regime [31]. We typically load the AP with a quadrupole gradient of 60 $\mathrm{G\,cm^{-1}}$, and operate shell manipulations with quadrupole gradients between this value and 151 $\mathrm{G\,cm^{-1}}$.

6.1.2 Loading atoms from the TOP

The atom cloud, whether thermal or condensate, is loaded from the TOP onto the AP shell surface via the vertically-offset double well TAAP [44, 95, 96] as illustrated in Fig. 6.1. The 7 kHz rotating TOP bias field serves a dual purpose here, in providing the bias field to time-average both the quadrupole and the position of the resonant spheroid, sweeping the quadrupole field in a horizontal, circular orbit with a rotation radius given by $B_T/B_q^{'}$ where B_T is the amplitude of the TOP field. Possible difficulties with the method and an alternative approach attempted are mentioned at the end of this section.

We begin with a thermal cloud of $\sim 1.2 \times 10^7$ atoms cooled to $\sim 1.7\,\mu\mathrm{K}$ in the TOP, using the standard BEC TOP sequence terminated at a value $B_T = 9.4\,\mathrm{G}$, 4.5 s before the end of the ramps shown in Fig. 4.15. This value is optimised for a combination of atom number and temperature upon loading the shell, taking account of finite heating upon the initial point of TAAP loading. The dressing RF is switched on while the

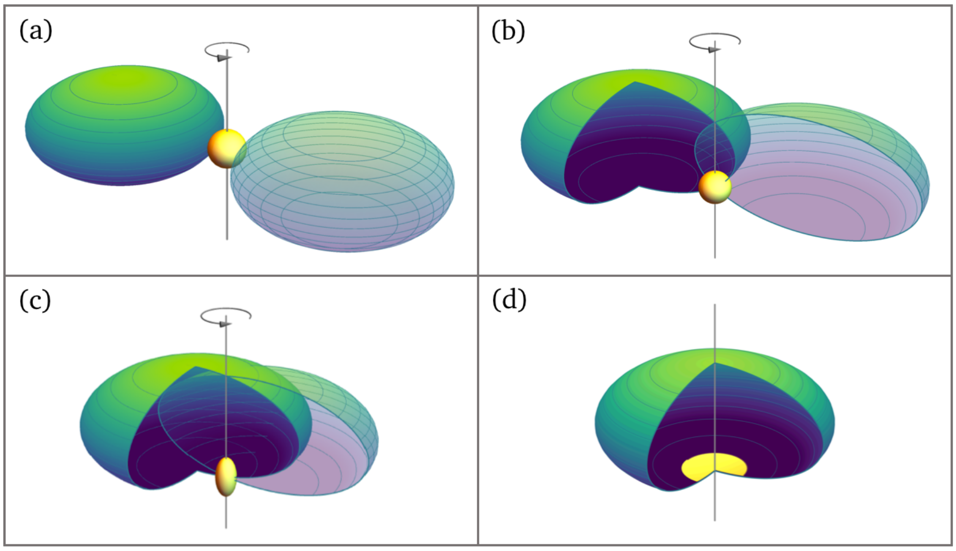

Figure 6.1: Single-shell AP loading scheme for circularly polarised RF. The two resonant spheroids are separated in time under the influence of the bias field rotating with a frequency of 7 kHz about the vertical axis. Shading indicates the relative coupling strength, with darker (purple) shading representing strong coupling at the south pole, and lighter (green) shading the weaker coupling towards the north pole. (a) Before AP loading, the atoms are confined in the TOP. Upon applying the dressing RF, the rotating bias field of amplitude B_T moves the resonant spheroid in a circular orbit. For $B_T > \hbar\omega_1/\mu_B g_F$ the shell orbits outside the thermal atom cloud (yellow). (b) Lowering B_T causes an intersection of spheroid and rotation axis, creating the vertically-offset TAAP double well. The atom cloud is loaded into the lower well. (c) B_T and the corresponding rotation radius are reduced further. Forced RF evaporation to BEC is performed at an intermediate stage of the TAAP. (d) $B_T \rightarrow 0$, loading the BEC onto the static shell surface.

TOP field satisfies $2\hbar\omega_1/g_F\mu_B > B_T > \hbar\omega_1/g_F\mu_B$ such that the resonant spheroid is swept in an orbit outside the location of the atom cloud with minimal perturbation to it. The upper bound to the TOP amplitude B_T accounts for the first harmonic of the resonant spheroid. For example, for the three angular frequencies used in the prototype MRF scheme, $\omega_{1,2,3} = 2\pi \times (3.0, 3.6, 4.2)$ MHz, the respective TOP loading values are $B_T = (6.5, 7.8, 9.1)$ G following a 50 ms TOP field ramp from $B_T = 9.4\,\mathrm{G}$ at the termination of the truncated TOP evaporation sequence, at a quadrupole gradient $B'_q = 154.1\,\mathrm{G\,cm^{-1}}$. The RF field is circularly polarised in the laboratory frame, with a handedness that maximises the interaction strength at the south pole of the resonant spheroid, and has an amplitude such that the coupling strength is on the order of $\Omega_i = g_F\mu_B B_i = 2\pi \times 400\,\mathrm{kHz}$ at this point. Decreasing B_T such that the intersection of the resonant spheroid and rotation axis separates into a double well along the vertical axis, we load the atoms into the TAAP formed at the lower of these intersections under the influence of gravity. This is associated with an atom number of $\sim 1.1 \times 10^7$, corresponding to a transfer efficiency of approximately 92%, and a small amount of heating to a temperature of $\sim 3\,\mu\mathrm{K}$. An optional stage of forced RF evaporation, described in Sec. 6.1.3, can be implemented at this stage to cool atoms to BEC. This is implemented at $B_T = 1.9\,\mathrm{G}$, after 700 ms hold in this TAAP well. A BEC of atom number $N \gtrsim 1 \times 10^5$ with no discernible thermal component and negligible atom loss is loaded onto the lower surface of the shell by reducing B_T to zero over 150 ms. Once shell loading has been performed using circularly polarised dressing RF, the polarisation can be altered to suit the subsequent work, for example by ramping one component to zero to achieve linear polarisation for amplitude calibration measurements. Such ramps are typically performed over 50 ms.

As a comment on issues of which to be aware, early implementations of this loading method were unsuccessful due to system noise that became resonant with the atoms during the TAAP ramp; this sweeps the resonance of trapped atoms through a range of values, making them particularly susceptible to noise sources at inopportune frequencies. However, once the noise spectrum had been cleaned, there were no problems with this method. As an intermediate stage, an alternative loading scheme was devised, inspired by the notion of the ‘shortcut to adiabaticity’ [184]. In this method, the thermal atom cloud was dropped from the TOP by turning off the trapping fields for the duration required for the atoms to fall under gravity to the location of the resonant spheroid surface. The fields and dressing RF were then applied to catch the atoms. This method was tolerably successful, although there was significant heating and motion of the atoms confined in the shell. This informed our early work on the trapping shells, and is a method that could most likely be improved with further work; the TAAP ramp, however, was a far more promising route to pursue. Additional factors that influence the success of the AP loading process include the RF field amplitude and polarisation; the means of characterising and correcting these are discussed in Sec. 6.2.

6.1.3 Cooling atoms in the TAAP

Forced RF evaporation is performed over 2 s in a TAAP formed at a quadrupole field of $B'_q = 154.1\,\mathrm{G\,cm^{-1}}$, circularly polarised RF with amplitude $\Omega_1 = 2\pi \times 400$ kHz, and a constant TOP field of $B_T = 1.9\,\mathrm{G}$. For an example plotted in Fig. 6.2, this transforms a cloud of 2.6×10^6 atoms at $3\,\mu\mathrm{K}$ to a condensate of 1×10^5 atoms with no discernible thermal component. The RF ramp and a corresponding progression of phase-space density are plotted for the example of a $\omega_1 = 2\pi \times 4.2$ MHz dressing angular frequency. The evaporative RF corresponds to a coupling between dressed states with $|\Delta\widetilde{m}_F| = 1$, and we have chosen a ramp corresponding to the transition $\omega_1 + \Omega_1 + \Delta(t)$, namely the sum of the dressing frequency, its amplitude as defined by the Rabi frequency, and such a detuning Δ as to facilitate the appropriate cut from high- to low-energy atoms. Accounting for the differences in dressing frequency, this same ramp can be

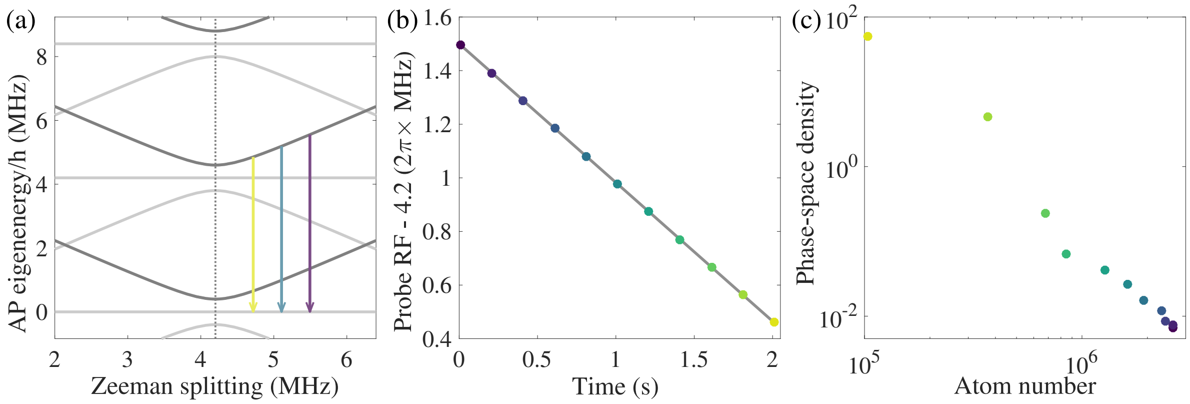

Figure 6.2: Evaporative cooling in the TAAP. (a) Dressed state eigenenergies for the $2\pi \times 4.2$ MHz AP. We evaporate in the TAAP trap by applying an RF knife $\omega_1 + \Omega_1 + \Delta(t)$, illustrated by the coloured arrows, sweeping downwards in energy towards the trap bottom over time. (b) The evaporative RF sweep in the TAAP; the same sweep can be applied to a TAAP with any dressing frequency, by adding the appropriate ω. Coloured dots indicate intervals of 0.2 s, and map to the points shown on the PSD plot. (c) Phase-space density plotted against atom number at 0.2 s intervals along the evaporative RF sweep. Atom number and temperature are obtained for each point by analysis of absorption images obtained after 15 ms time of flight expansion. The final data point here is fitted as a pure BEC (Thomas-Fermi distribution) and the preceding points as thermal Gaussian clouds.

applied to any of the shells so far trialled, with dressing angular frequencies in the range $2.0\,\mathrm{MHz} \leq \omega/2\pi \leq 4.5\,\mathrm{MHz}$.

6.2 Refining and characterising the potential

6.2.1 RF spectroscopy

RF spectroscopy is a technique employed to characterise trapping potentials, including RF-dressed APs, by probing the energy of atomic transitions [85, 185]. In this method, a weak probe RF at a fixed frequency is applied to atoms held in the AP, expelling atoms from the trap when the probe RF is resonant with a transition between trapped and untrapped states. As the probe RF is varied, these resonances appear as dips in the measured atom number. This is similar to the process of RF evaporation, but with the probe RF detuning varied between experimental runs, held constant during each, rather than applied as a time-dependent ramp. The probe RF is applied using the evaporation coil and its same preceding electronics, and we likewise again choose to probe the $\omega_1 + \Omega_1 + \Delta$ transition.

The structure of the transition lineshape is exposed over several experimental runs, its exact non-analytical form dependent on both the details of the potential and the energy distribution of the trapped atoms. To minimise this latter effect, we perform spectroscopy with a BEC such that the resonance has a width on the order of the chemical potential (typically kHz) and we can precisely pinpoint the energy of each transition. The resultant lineshape is markedly asymmetric, with a steeper gradient on the low frequency approach to the resonance than from the high frequency side. This can be explained by considering the energy structure: efficient coupling of the RF probe ensures that all atoms with energy greater than that of the coupling point are expelled from the trap [185], and so with reference to Fig. 6.2(a) it is expected that cutting from below will be associated with abrupt atom loss as the probe RF varies from no resonance at all to resonance with the trap bottom, and from above with a more gradual decline in atom number as the probe cuts through the residual energy distribution of the atom cloud. We extract the resonance peak by fitting an appropriate approximate function; the fitting process is discussed in the MRF context in Sec. 7.2. The probe RF field must be sufficiently weak that it does not itself shift the transition. For the APs used here the Rabi frequencies of the dressing RFs are hundreds of kHz while that of the probe is below $100\,\mathrm{Hz}$. Selected RF spectroscopy measurements were repeated with probe amplitudes spanning one third to three times its standard value, with no measurable shift of the resonance.

RF spectroscopy has been applied in this manner to characterise both single- and multiple-RF dressed potential eigenenergy structure, and the fields critical to the implementation of these traps as discussed below. This method was also implemented to characterise the TOP and quadrupole traps as described in Appendix A.

Resonant RF pulses can also be used to release atoms for time of flight imaging, elminating any effects on the cloud of the magnetic trapping fields switching off. With the quadrupole field remaining on, this has the additional benefit of separating the m_F components for further analysis. This method can also be used in pulsed operation, where short pulses of resonant probe RF are applied to the potential, releasing the atoms in periodic bursts of time of flight. A close spacing between each pulse ensured that several such bursts can be imaged simultaneously to reconstruct the in-trap radial motion perpendicular to the imaging direction. This method of 'RF time of flight', implemented using a signal generator in burst mode, has been used to detect lateral movement of the atoms on the shell surface and to image atoms where the action of the static quadrupole field on different m_F components is important. An example of this is in the imaging magnification calibration discussed in Appendix A.3.

6.2.2 RF amplitude calibration

The amplitudes of the applied RF signals are monitored by pickup coils co-wound with the TOP coils[2]. For each dressing RF used, these measured voltages are calibrated using RF spectroscopy to associate with each measured pickup amplitude the corresponding transition frequency. Linearly polarised RF is used to measure the x and y directions independently and eliminate the effects of cross-coupling between the pickup directions. These measurements are then associated with the corresponding Rabi frequencies Ω_i using Floquet theory, in a calculation that incorporates the effect of gravity and Bloch-Siegert shifts, to give the RF field amplitudes. RF amplitudes are therefore expressed in terms of these single-RF Rabi frequencies Ω_i.

The linearity of the pickup coil response was verified by repeating the single-RF spectroscopy measurements for a variation in RF amplitude of up to $50\,\%$. In another experimental concern, the combined MRF input amplitude for the initial tests presented here approaches a value close to the saturation of the amplifier, which results in up to $4\,\%$ compression of the amplitudes of each RF component for the highest RF powers applied. This saturation is accounted for by the RF pickup measurement in all comparisons between experimental and theoretical work. The field amplitudes are monitored for each experimental run, giving us the means of not only monitoring any drifts over time, but also of accurately correcting for any small amplitude variations in calculations of predicted results. Amplitude drifts on the order of a few kHz have been observed over a period of several weeks, though with variation in short-term severity and the associated experimental impact.

The RF source stability can also be monitored by mixing a small proportion of it with a reference oscillator and extracting the DC component, although this scheme becomes more challenging in the multiple frequency implementation. Using this method, active PID stabilisation of the dressing RF amplitudes was also trialled in the early

---

[2] Early measurements of RF current were made directly with Tektronix AC current probes as addressed in Sec. 5.5.2

stages of this work [35]. However, more significant improvements were made to RF stability by tackling the root causes in terms of the DDS source and peripheral noise, and the feedback scheme was laid aside until such a time as experimental stability requirements become more stringent.

6.2.3 Tuning the RF polarisation

The dressing RF polarisation is determined by the phase offset between the x and y RF components. This in turn dictates the variation in coupling strength over the shell surface. For atoms confined at the south pole of the resonant spheroid by gravity, the trap lifetime is in principle maximised for perfectly circular polarisation where the coupling strength is strongest at the position of the atoms.

It seems a reasonable assumption that the polarisation could be calibrated by varying the phase difference between RF components to maximise the lifetime. However, local maxima exist, caused by noise sources that have at various times plagued the experimental system[3]. As the RF polarisation is varied, the energy of trapped atoms was swept through resonance with certain of these noise sources, causing atom loss in either narrow or broad dips according to the source. This was to become a rare instance in which the system noise could become an asset, and was exploited to improve the accuracy of the polarisation calibration. Since the energy response of trapped atoms is symmetric about the point of circular polarisation with respect to ellipticity in either x or y directions, there exists reflection symmetry in the noise response to phase variation. Finding the central point in this distribution with reference to the narrowest noise peaks is therefore a more accurate means of identifying the balanced phase corresponding to perfect circular polarisation[4].

An example of this method as applied to the $2\pi \times 4.2$ MHz shell is illustrated in

---

[3] A full discussion of the RF noise is presented in [35].

[4] This is not to say that the clearance of RF noise need hinder such measurements; a controllable low-power RF probe could instead be applied to create the necessary features.

Fig. 6.3. Atoms at $\sim 3\mu$K are held in the shell for 8 s, sufficient for significant atom loss, and the atom number measured for a range of values of the phase offset. This allowed a determination of the correct polarisation to within a phase resolution of $\pm 0.01\pi$ radians. The relative phase between the pickup signals in x and y can then be calibrated to provide a polarisation reference.

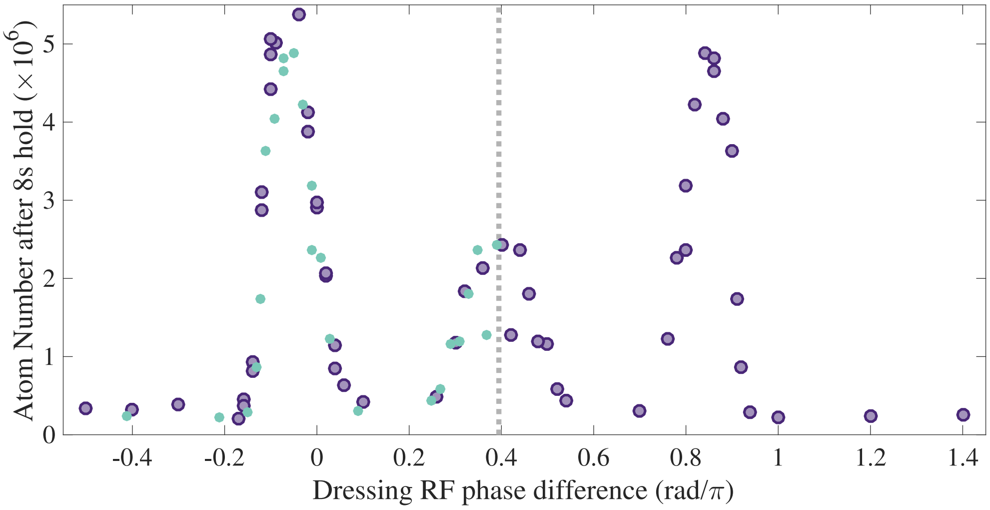

Figure 6.3: RF polarisation calibration in the $\omega_1 = 2\pi \times 4.2$ MHz shell: measured atom number after 8 s hold as the phase between x and y RF components is varied (purple circles), the identified symmetry point (grey dashed line) and the reflected points used to match features (green dots). With a perfectly clean RF noise spectrum, the longest trap lifetime corresponds to perfect circular polarisation; with noise at a frequency resonant with the trapped atoms, atoms are lost and local minima arise. In this case, the symmetry of the variation in resonance frequency with respect to deviation from circular polarisation is a useful reference for locating this point. By observing the overlap between the measured points on one half of the phase sweep with the reflection of those in the other half, the symmetry point can be determined to within $\pm 0.01\pi$ radians. Note: this measurement was chosen to illustrate the method due to the exaggerated peaks. The noise sources were subsequently eliminated to give a single symmetric peak with the same reflection point as above and the longest lifetime corresponding to circular polarisation as expected.

6.3 Loading atoms into the MRF double well

As described in Sec. 2.5.2, the MRF double well is created by three dressing RFs, each corresponding to a separately-controlled single first-order avoided crossing. This results in two trapping wells separated by an anti-trapping barrier for $\widetilde{m}_F = 1$, where trapping is realised on two concentric spheroids. In the present work we consider a parameter regime in which atoms pool on the lower surface of each spheroid under the influence of gravity, to create a configuration akin to two vertically-offset sheets. By a change of parameters to operate the shell trap in the 2D regime, this would correspond to a 2D double well.

The relative heights of the barrier and both wells are controlled by the three separate input RFs, although multi-photon interactions lead to cross-talk between these features and an associated shift in the details of the overall potential landscape with a variation in amplitude of any of the dressing RF components. The main features of the MRF double well are sketched in Fig. 6.4. The primary manipulations of the potential in this work are performed by increasing and decreasing the amplitude Ω_2 of the barrier RF to respectively lower and raise the barrier. This includes a consideration of the effect of gravity on the overall potential landscape.

Within the constraints of the coil array self-resonance and the wideband impedance match, we can confine atoms in APs with dressing angular frequencies in the range $2.7\,\mathrm{MHz} < \omega/2\pi < 4.4\,\mathrm{MHz}$ without additional amplification. Mixing processes in the amplifiers constrain us to use only combinations of dressing angular frequencies with a common fundamental $\omega_f \gtrsim \Omega_i$, which ensures that the resulting intermodulation products are far detuned from transitions between dressed states such that we avoid losses. For this proof-of-principle implementation, we confine atoms in the double well created by dressing RF components $\omega_{1,2,3} = 2\pi \times (3.0, 3.6, 4.2)$ MHz. In our apparatus the $\omega_f = 2\pi \times 0.6$ MHz angular frequency difference between RF components maps to a spatial well separation of $140\mu\mathrm{m}$ at a quadrupole gradient $B_q^{'} = 62.45$ G/cm, allowing the trapping wells to be clearly resolved with our low-resolution imaging system. A smaller angular frequency separation and quadrupole gradient would reduce the spatial separation between the wells. With the existing power supply and water cooling system, we can access quadrupole field gradients of up to $468\,\mathrm{G\,cm^{-1}}$. A smaller frequency separation would require operation at lower field amplitudes. We have suc-

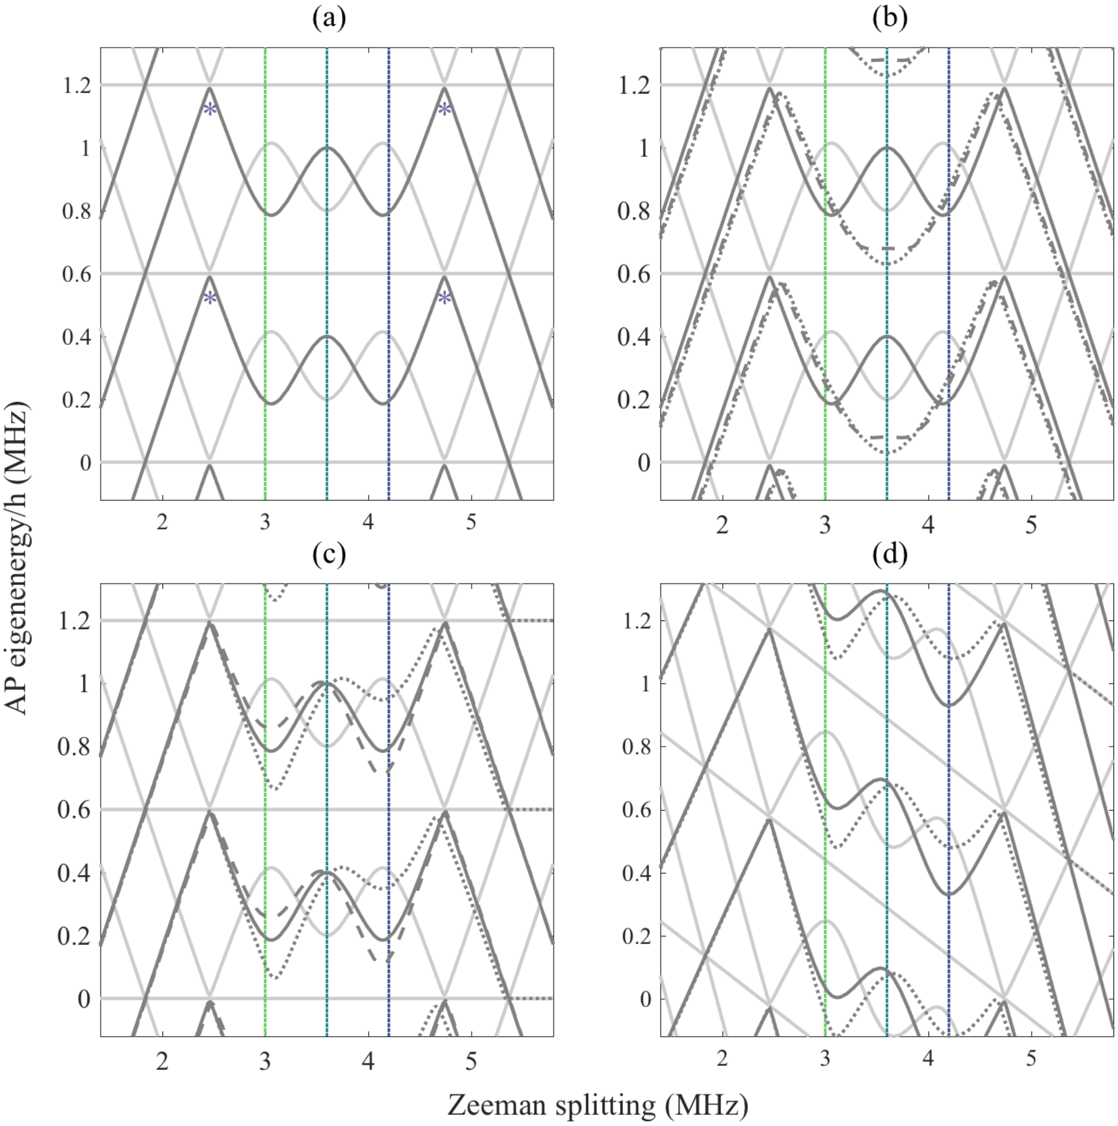

Figure 6.4: Tuning the three-RF double well by varying the relative amplitudes of the constituent dressing RFs $\omega_{1,2,3} = 2\pi \times (3.0, 3.6, 4.2)$ MHz. (a) The three-RF double well plotted in the absence of gravity with equal amplitudes $\Omega_{1,2,3} = 200$ kHz to create a balanced double well. Asterisks indicate the weak multi-photon avoided crossings. (b) The barrier amplitude is adjusted to transform the potential between a double well (solid line), flat-bottomed potential (dashed line) and broad single well (dotted line). With amplitudes $\Omega_{1,3} = 200$ kHz fixed for each plot, the barrier amplitudes are $\Omega_2 = 2\pi \times (200, 520, 570)$ for the solid, dashed and dotted lines respectively. (c) The relative well depths can also be varied to shape the potential. With the $\Omega_{1,2,3} = 200$ kHz reference plotted as a solid line, the dashed and dotted lines show amplitudes $\Omega_{1,2,3} = (270, 200, 120)$ kHz and $\Omega_{1,2,3} = (80, 200, 370)$ kHz. (d) Incorporating the effects of gravity dramatically alters the appearance of the double well, depending on the quadrupole field gradient. The eigenenergies of the double well with amplitudes $\Omega_{1,2,3} = 200$ kHz are again plotted here as the solid lines with gravity incorporated and a quadrupole field gradient 60.4 G/cm. To balance the energy of each well, the field amplitudes are adjusted to $\Omega_{1,2,3} = (80, 200, 360)$ kHz (dotted line). This act of balancing the potential reduces the depth of the well corresponding to trapping angular frequency ω_3, limiting the temperature of confined atoms.

cessfully confined atoms on single-frequency shells with field amplitudes in the region of $\Omega = 2\pi \times 50$ kHz, indicating that this will not be a constraint in the MRF case. The frequency constraints and variation in RF transmission as a function of frequency mean that frequency ramps would be challenging to implement. However, sufficient control over the potential is granted over more easily tuneable parameters that this need not be inhibitive.

The MRF potential takes as its starting point a single shell; whether this is the shell with dressing angular frequency ω_1 or ω_3, each corresponding to the trapping components of the MRF configuration, is a matter of choice, and the principles remain the same. A typical procedure by which to load a BEC from a single shell at ω_3 into the double well is shown in Fig. 6.5. To begin, we ramp up Ω_1, which has a minimally perturbative effect on the potential near the atoms but establishes this resonance in preparation for the subsequent barrier manipulations. This is accompanied by a slight increase in the field amplitude Ω_3, acting to raise the corresponding trapping well in anticipation of ultimately balancing the energy of the two wells. The avoided crossing formed by ω_2 takes the form of an anti-trapping barrier, and as Ω_2 increases, this barrier is lowered and the combined three-frequency potential is flattened, rounded out, or tilted slightly according to the relative values of Ω_1, Ω_2 and Ω_3. For balanced shell loading, we increase Ω_2 to round out the combined potential, with Ω_1 held at an artificially high value to minimise any sudden changes in the width of the potential well as the barrier is lowered. Only once the barrier has been ramped down fully is Ω_1 lowered to the value at which atoms can be transferred. Once atoms equilibrate within this new three-RF single well, we raise the barrier to separate the wells and complete the loading process. Placing limits on this process, and marked with asterisks in the panels of Fig. 6.5 are the second-order resonances that can act as loss channels. The combination of RF amplitudes and frequency separation are therefore chosen to complement the atom temperature.

Figure 6.5 also shows the calculated shape of the potential at key stages of the

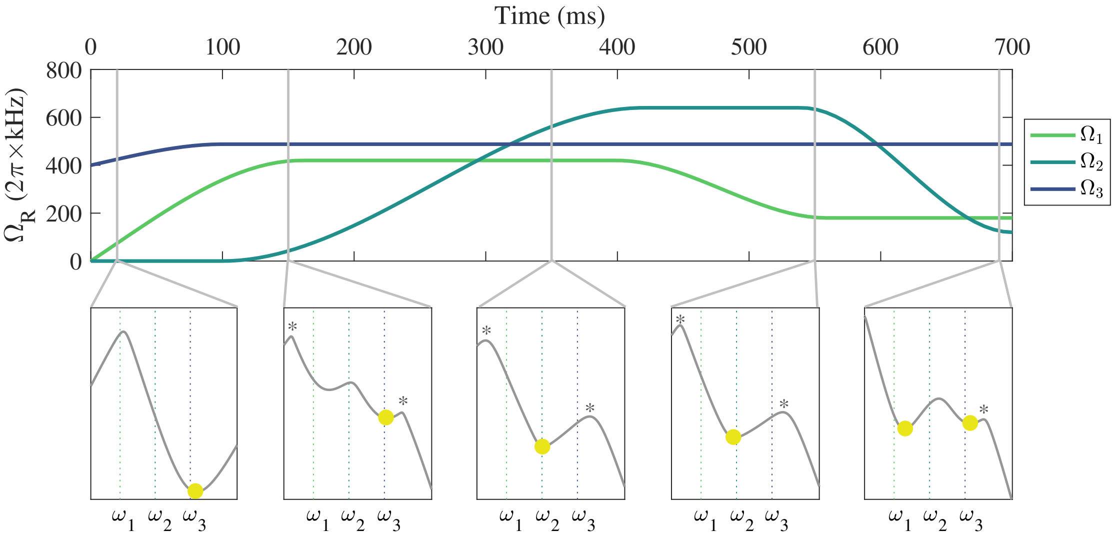

Figure 6.5: A typical time sequence of the dressing RF amplitudes used to load a BEC into a double well from a single shell with dressing angular frequency $\omega_3 = 2\pi \times 4.2$ MHz. $\Omega_i = g_F \mu_B B_i/\hbar$ denote the constituent field amplitudes of the three dressing fields. The lower panels show the three-frequency potential (incorporating the effects of gravity) at key times during this loading sequence, with dotted lines indicating the locations associated with the first-order resonances of the dressing angular frequencies ω_1, ω_2 and ω_3. As this sequence progresses, a single-frequency single well is transformed into a three-frequency single well with a flattened barrier, associated with high Ω_2, and an artificially high value of Ω_1 to restrict the width of the MRF single well. Ω_1 is then reduced to establish the associated trapping well, and the barrier raised to separate the wells and split the cloud between them. The experimental parameters chosen avoid losses due to the second-order resonances, indicated by asterisks in the panels. The relative amplitudes of the RF components determine the final distribution of the atomic population between each well. The offset of each avoided crossing from the guiding lines at each dressing RF results from a combination of the effects of multi-photon interactions and the associated cross-talk between potential features, and the effect of gravity in shifting the potential minimum downwards in space.

loading process. These plots incorporate the effects of gravity, which acts along the central vertical axis of the resonant surface and thus increases the potential depth towards higher resonant frequencies. This effect must be compensated to balance the two trapping wells, explaining the discrepancy between the field amplitudes Ω_1 and Ω_3.

An alternative approach to loading is an extremely similar scheme to that presented above, but with Ω_1 ramped to its final value straight away, so skipping the intermediate high value; this method was used to obtain the data presented in Fig. 6.6 and in a slightly variant form that of Fig. 7.1 in Sec. 7. As a method this was almost equal to that presented above, but with a slightly higher tendancy towards sloshing of the atom cloud during transfer, caused by the sudden exposure of the atoms to a broad well once the barrier was lowered sufficiently that they could flow across. Another loading method trialled was that of loading and evaporating directly in the three-frequency single-well TAAP, eliminating the need for the initial transfer from the single-RF well. However, the restricted lifetime of atoms in this potential at that time prevented this. While it may be a direction explored in future, the success of the method shown above in loading a BEC into the MRF double well without appreciable heating, and more recent work with a smaller well spacing yielding preliminary results suggesting that coherent splitting of a condensate is possible with this method[5], mean that small variants upon the loading scheme explored in Fig. 6.5 will be sufficient for future work.

This loading method allows stable, repeatable loading into the three-RF double well, with precise control over the distribution of the atomic population between the wells granted by the relative amplitudes of each RF component during the ramp. The effect of barrier height is illustrated in Fig. 6.6, where we vary the maximum value of Ω_2 to load a controllable proportion of atoms between the lower and upper wells formed by ω_3 and ω_1 respectively. Starting from a cloud of thermal atoms in the spatially lowest shell at ω_3, the RF components ω_1 and ω_2 are turned on adiabatically following a similar procedure to that described in Fig. 6.5 in which Ω_1 is ramped directly to its final value. Initially, few atoms possess sufficient energy to cross the high barrier that results from a small Ω_2, and minimal population redistribution between the wells occurs. However, upon increasing Ω_2, the barrier is lowered and more atoms are able to populate the second well. Around $\omega_2 = 2\pi \times 400$ kHz the avoided crossing at ω_2 no longer presents a barrier and the atoms distribute themselves across the broad single well formed by the three RF dressing fields. Finally, ω_2 is decreased to

---

[5] The work on coherent splitting has been performed during the preparation of this thesis and so results will be presented in future work [35].

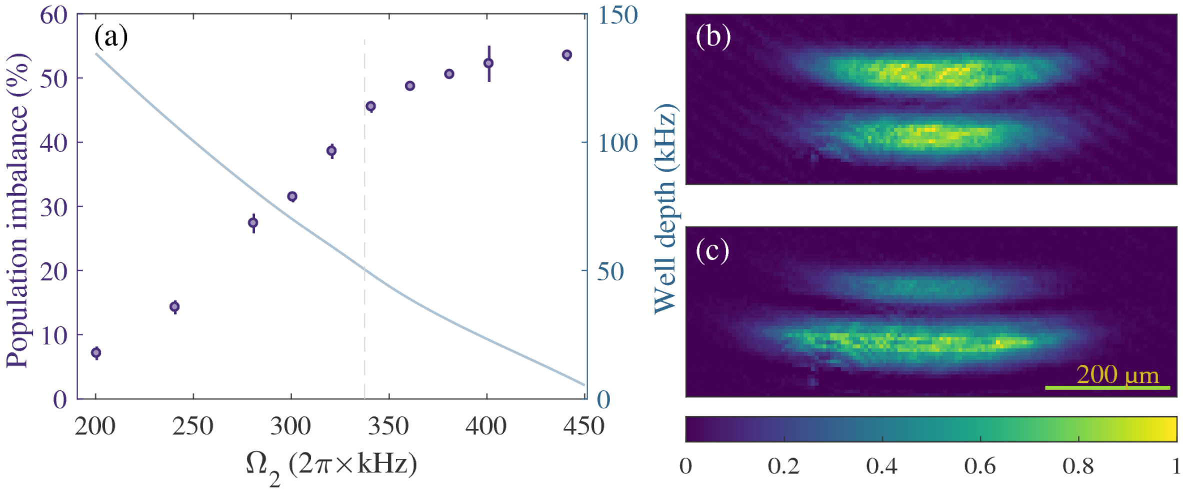

Figure 6.6: Controlled double well loading. (a) The percentage of atoms loaded from the well formed by ω_3 to that corresponding to ω_1 for a given maximum amplitude of the ω_2 field, expressed in terms of $\Omega_2 = g_F \mu_B B_2/\hbar$. The RF amplitude ramps are similar to those plotted in Fig. 6.5, with $\Omega_1 = 2\pi \times 192\,\mathrm{kHz}$ and $\Omega_3 = 2\pi \times 442\,\mathrm{kHz}$. This amplitude disparity compensates the effects of gravity, with a quadrupole gradient of $154\,\mathrm{G\,cm^{-1}}$. The barrier was ramped to its maximum value over $400\,\mathrm{ms}$, then reduced to $2\pi \times 90\,\mathrm{kHz}$ over $100\,\mathrm{ms}$. The blue line shows the effective well depth (right hand scale) seen by atoms in the well at ω_3 for each final value of Ω_2, and the dashed vertical line indicates the barrier height for which a separate well at ω_1 can no longer be resolved. (b), (c) Absorption images of thermal atoms confined on the double shell surfaces at a quadrupole gradient of $60\,\mathrm{G\,cm^{-1}}$ after 1 ms time of flight, with (b) an approximately balanced configuration with $52\,\%$ of atoms in the upper shell and (c) $75\,\%$ of the population in the lower shell. The absorption images are separately normalised to their maximum value and shown in false colour.

raise the barrier and split the population distribution into two distinct wells, with the final proportion reflecting any imbalance between the lowest energy of each well. Figure 6.6(a) illustrates a loading process that transfers 52% of the atoms into the well defined by ω_1. This could be corrected, or exacerbated, by adjusting either Ω_1 or Ω_3 to raise or lower the potential energy minimum of each well to preferentially tip or retain the atoms in a particular direction.

The following chapter presents a further example of precise population control in the context of potential shaping, and characterisation of the energy eigenvalues as the barrier amplitude is changed to shape the potential.

Chapter 7

Ultracold atoms in multiple-RF potentials

Three separate RF dressing fields are used to create two trapping wells separated by a potential barrier, each independently controlled by a separate DDS channel. While each dressing field influences the potential landscape more broadly than just at its own specific resonance, this still provides a high level of control over the potential landscape and interplay between the wells. This chapter begins by exploring the potential shaping capabilities of the MRF double well in Sec. 7.1. RF spectroscopy is then used to probe the potential landscape in Sec. 7.2, to both verify the theoretical model and quantify the cross-talk between RF field components. Finally, the experimental outlook of the multiple-RF technique is discussed in Sec. 7.3.

7.1 Potential shaping and the double well

We have explored the range of potential landscapes accessible with the three-RF prototype dressing field, and have demonstrated the ease of transferring atoms between these potential configurations. Within this system, the landscape can take various forms: a single well corresponding to resonance with the dressing fields at ω_1 or ω_3, a double well with controllable barrier height, or a broad three-frequency well that can approach a square-well potential.

Figure 7.1 shows some different potential configurations accessible by varying the relative amplitudes of the constituent RF fields, illustrating the atom density that arises from two possible transport sequences from the $\omega_3 = 2\pi \times 4.2$ MHz single-frequency shell. These potentials were loaded using a sequence in which Ω_1 is ramped immediately to its final value. As in the scheme presented in Fig. 6.5, Ω_2 is gradually increased to lower the barrier and load atoms into a broadened MRF well; here, however, we show the trapped atom density as a strip along the vertical axis of the resonant spheroid for several steps in this process. From the broad MRF well, two transport sequences are shown: Ω_1 can be increased to tip atoms into the trapping well formed by ω_1 as the barrier is raised by decreasing Ω_2, or the two trapping wells can be kept at approximately equal energies to load a balanced double shell. Any intermediate population mismatch is also possible and determined by the relative well depths, and a single well comprising features of all three dressing field components can also be created. This also demonstrates the effect of the barrier amplitude on the positions of the two trapping wells that is shown in the calculated energy levels in Fig. 6.4: the potential minima for ω_1 and ω_3 become closer together as the barrier is lowered to form the broad single well. This is illustrated here using in-trap images of thermal atoms at a few $\mu\mathrm{K}$, but MRF loading has also been demonstrated using a condensate.

7.2 Probing the potential landscape

RF spectroscopy was performed to verify our theoretical model of the MRF potential. For this purpose, we used a BEC confined in the well near resonance with the $\omega_1 = 2\pi \times 3.0$ MHz dressing field and observed the variation in transition energy with a change in barrier amplitude. After loading into the circularly polarised single-RF

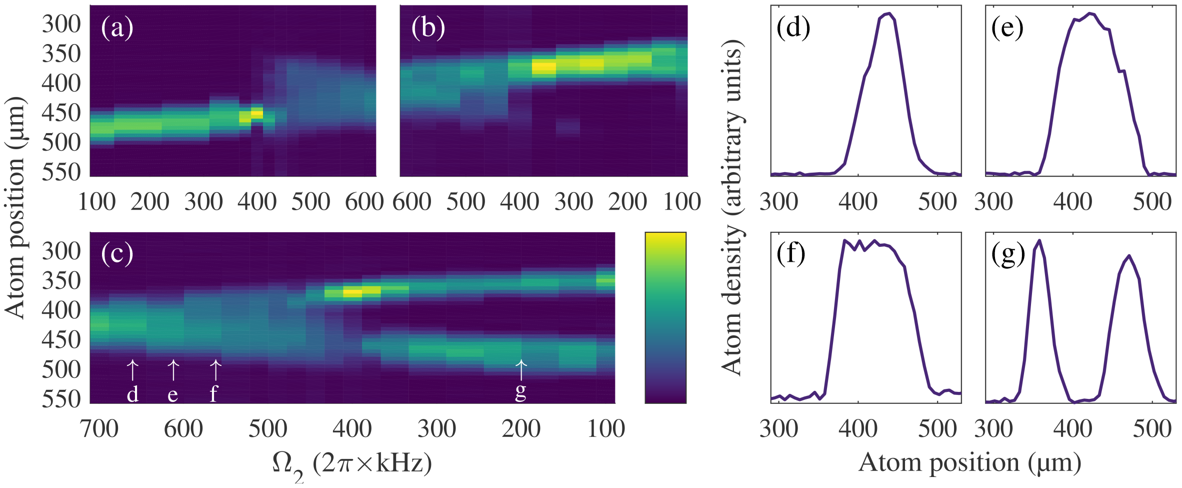

Figure 7.1: Atom density in the 3-RF potential with varying barrier height. (a)-(c) Vertical slices through in-trap absorption images of the MRF potential plotted against barrier amplitude ω_2 and averaged over several experimental runs. Displacement is measured from the centre of the quadrupole trap, and each slice scaled to the same total atom number. (a) All atoms begin in the shell at ω_3. With $\Omega_1 = 192\,\mathrm{kHz}$ and $\Omega_3 = 446\,\mathrm{kHz}$, Ω_2 is ramped up (lowering the barrier) to flatten the potential and load the atoms into a broad single well. This is the starting point for the sequences shown in either (b) or (c) depending on subsequent parameter choice. (b) A transport sequence with $\Omega_1 = 2\pi \times 192$ kHz and $\Omega_3 = 2\pi \times 511$ kHz. This tips atoms across the flattened 3-RF potential to load atoms into the upper shell upon reducing Ω_2 to raise the barrier. (c) Loading a double-shell configuration from the flattened 3-RF potential, with $\Omega_1 = 2\pi \times 192$ kHz and $\Omega_3 = 2\pi \times 446$ kHz to maintain approximately equal atom populations in each well. The highest values of Ω_2 correspond to a single well, with the two distinct wells forming and separating as Ω_2 is reduced to raise the barrier. The apparent transfer of atoms into the shell at ω_1 around $\Omega_2 = 2\pi \times 400$ kHz is a normalisation artefact, resulting from atom loss from the lower well due to technical noise in the apparatus (this noise source has since been eliminated). All absorption images are shown in false colour, normalised to their maximum value; the colorbar displayed scales linearly from 0 to 1. (d)-(g) show the line plots of atom density for the snapshots in the double shell load sequence at barrier amplitudes marked in (c) and corresponding to $\Omega_2 = 2\pi \times 660, 622, 577, 266$ kHz for (d)-(g) respectively. This shows the progression from 3RF single well (d) and (e) to flat-bottomed 'box trap' (f) and double-shell potential (g).

shell, the polarisation was ramped to linear over 500 ms for the application of the MRF potential. Linear polarisation minimised the number of experimental variables to control, eliminating three RF field sources and the polarisation of each field component, and thus the uncertainty of the measurement. The single component of each of the ω_2 and ω_3 fields are ramped directly to their final amplitudes with a fixed relative phase on each experimental run. RF spectroscopy is then performed to probe the transitions available to the trapped BEC.

As in the single-RF case, a resonant probe RF field can drive transitions between different Floquet manifolds, although the reduced spacing $\hbar\omega_f$ of the ladder of dressed eigenenergies means that a large number of transitions can be driven within a given probe RF range. However, many of these transitions correspond to higher-order multiple-photon processes with low transition rates, and so do not have a significant effect in practice. Theoretical transition frequencies are obtained by calculating the the AP eigenenergies using Floquet theory as introduced in Sec. 2.5 [35], from which a double-well trapping level is selected. From the range of available transitions to untrapped states, the energy separation between the atoms at their location in the well and a chosen untrapped state is then selected. As for a single dressing RF, the probed transition is the one that would, in the absence of the additional dressing fields, correspond to $\omega_1 + \Omega_1$; this is illustrated in Fig. 7.2.

The experimental variation in this transition energy with barrier amplitude is shown in Fig. 7.3. The constituent linear RF fields take amplitudes $\Omega_1 = 2\pi \times 187\,\mathrm{kHz}$ and $\Omega_3 = 2\pi \times 248\,\mathrm{kHz}$, while ω_2 takes values between 0 and $2\pi \times 332\,\mathrm{kHz}$ with a quadrupole gradient $B'_q = 82.5\,\mathrm{G\,cm^{-1}}$. Over the course of the Ω_2 amplitude ramp, Ω_1 falls by 5% and Ω_3 rises by 1% due to amplifier saturation and nonlinearities. This is reproducible and is accounted for in the calculated transition frequencies. The field amplitudes are set such that the condensate remains confined to the initial well for the spectroscopy measurements, during which the weak RF probe is applied for a duration of $40\,\mathrm{ms}$. The potential is deformed slowly to avoid sloshing of the condensate; ramps occur over an $800\,\mathrm{ms}$ duration that is slow compared to the inverse of the 200 to 400 Hz axial trap frequencies. The probe duration is sufficiently long that any residual sloshing in the wells would only manifest as a broadening of the measured RF spectroscopy resonances.

The resonance point is extracted from the asymmetric spectroscopy profile [185] by fitting a function of the form $a(x-b) + c/\sqrt{x-d}$. This function provides a good approximation to the asymmetric lineshape of the resonance profile from which the

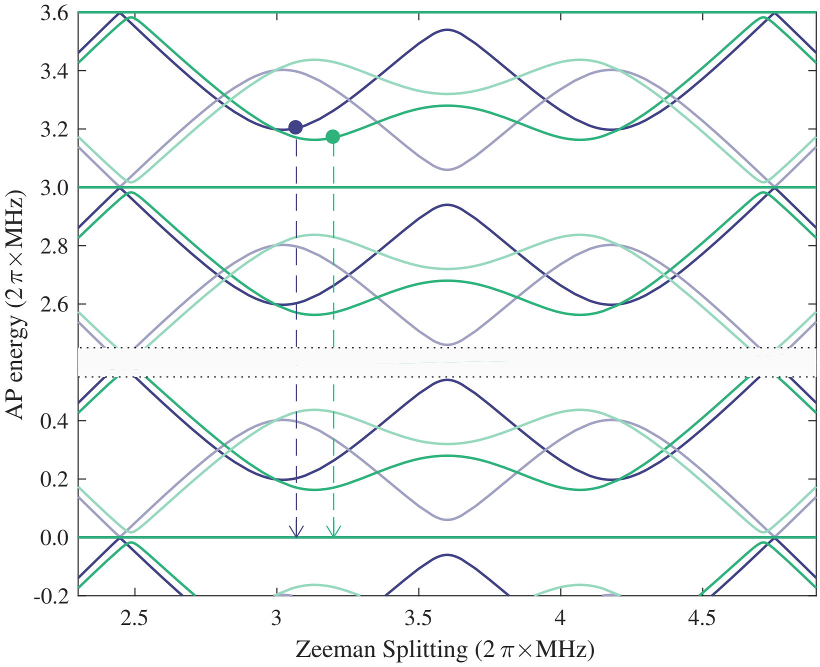

Figure 7.2: Spectroscopy transitions in the MRF potential. Dressed state eigenenergies are shown at two different barrier heights, plotted in the absence of gravity. A BEC is confined in the well near ω_1, its location illustrated here with the effects of gravity incorporated. We apply a probe RF resonant with the transition that would be labelled $\omega_1 + \Omega_1$ in the limit $\omega_2, \omega_3 \to 0$, to track the variation in the MRF eigenenergies at the location of the atoms as the barrier amplitude Ω_2 is varied while keeping Ω_1 and Ω_3 constant.

resonant probe frequency that minimises the atom number can be extracted. Only the data points lying within the range of the resonance were included in the fit, such that the asymmetric parabola captures the centre of the resonance with minimal free parameters. The uncertainty in the fitted resonance location for both single-RF calibration and MRF potentials is estimated from the $99\,\%$ confidence interval of the fitted minimum, and is of order 1 to $3\,\mathrm{kHz}$, although with a statistical accuracy limited by the sample size. This forms the dominant source of uncertainty in the measured transition frequencies, with a smaller influence from uncertainty in measuring dressing RF amplitudes with the pickup coils.

An alternative method based on asymmetric Gaussian fitting to incorporate the background data points was also implemented for comparison; while this incorporates the full atom number distribution, successful fitting requires severe parameter

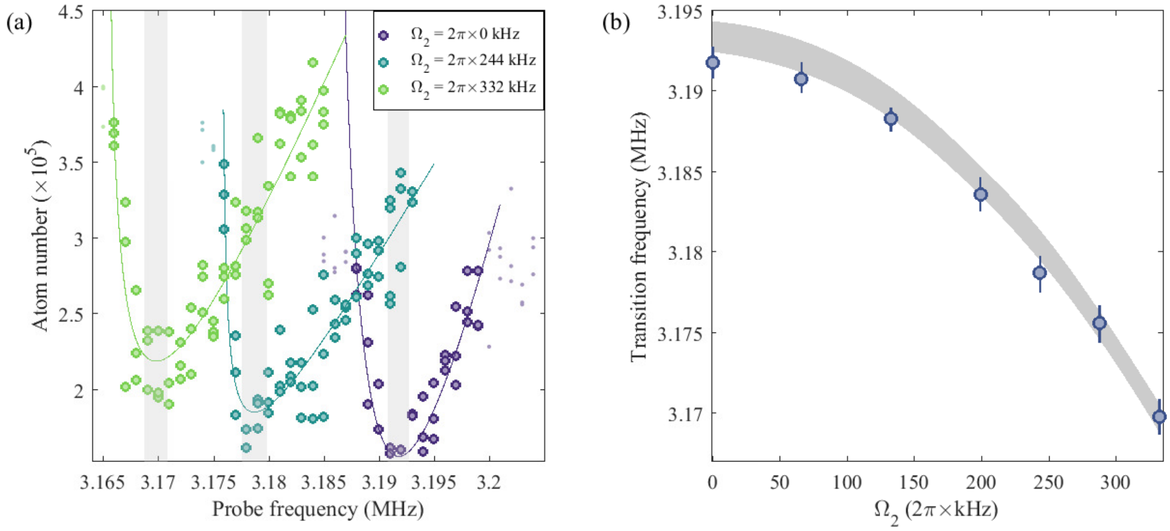

Figure 7.3: Measured variation in the MRF eigenenergies with barrier amplitude. Atoms are confined in the well corresponding to $\omega_1 = 2\pi \times 3.0$ MHz and the $\omega_1 + \Omega_1$ transition measured using a probe RF field. (a) Measured spectroscopy resonances at $\Omega_2 = 2\pi \times 0, 244, 332\,\mathrm{kHz}$. Data points shown in bold are included in the fit used to extract the minimum of the resonance (solid lines, see text), with grey sections indicating the 99% confidence interval for each minimum. (b) Change in measured (points) and theoretical (line) resonances in the MRF potential for a range of values of Ω_2. Error bars are calculated using the 99% confidence interval in the spectroscopy resonance fit in combination with uncertainty in the RF amplitude and its calibration. The theory line was obtained with no free parameters by calculating the transition energy for each value of Ω_2 probed experimentally, with an interpolation between these values. Its finite width corresponds to the experimental uncertainty in the three measured RF amplitudes Ω_i at each value of Ω_2.

constraints in certain cases and it is unclear which method is more accurate. However, for the majority of data points, the agreement between the two fitting methods was within 0.8 kHz; for the measurement at barrier amplitude $\Omega_2 = 2\pi \times 332$ kHz the values differ by 1.1 kHz. Given this close agreement, within the quoted fit uncertainty, either fitting method seems reasonable.

The actual lineshape can be simulated numerically [185], and is influenced by the amplitudes of both dressing and probe RF fields, and the chemical potential of the trapped condensate. With these factors, a separate fit for each spectroscopy data set is impractical and at risk of overfitting. Qualitative comparisons between the simulated lineshape and chosen fit function suggest that the systematic uncertainty arising from a discrepancy between these models would be smaller than a kHz [35].

As shown in Fig. 7.3, increasing Ω_2 to lower the barrier reduces the energy separation between trapped and untrapped states for the measured transition. A shift in the measured RF spectroscopy resonance on the order of tens of kHz is observed as Ω_2 is varied, in good agreement with the theory.

7.3 Outlook

We have demonstrated controllable loading of the three-RF double well, and used this example to explore the capability for potential shaping by varying the dressing RF amplitudes. Trapped atoms have been transferred between a variety of potential configurations. We have also verified our theoretical understanding of the MRF potential by RF spectroscopy of a BEC confined in a single well as the barrier amplitude was changed. In doing so, we have both characterised the three-RF system and developed robust and reliable methods by which to load and manipulate trapped atoms.

This lays the groundwork for a variety of extensions and applications of the MRF dressing technique. These include reducing the well separation sufficiently to study double well physics, including in the 2D regime [35], exploiting the g_F-dependence of the resonant surface to perform independent species-selective manipulations [33, 35, 41], arbitrary potential shaping by the application of additional frequency components, and extension of the accessible trapping geometries by applying techniques such as time averaging, the addition of an optical dipole trap, or the exploitation of other trap parameters such as RF polarisation.

Progress towards the addition of an optical dipole trap to the shell surface is detailed in Chapter 8, but the main topics of current experimental work are those of species-selectivity and well separation. In progress towards the scheme outlined in Ref. [33] in which a single well containing a ^{87}Rb condensate is overlapped with a double well containing ^{85}Rb impurity atoms to act as a probe of the bulk condensate properties, both atomic species have been simultaneously confined at their respective resonance locations by the application of a single-frequency RF field [35]. The well spacing has also been significantly reduced, by a combination of increased quadrupole gradient and reduced frequency spacing. During the preparation of this thesis, a double well with spacing $7.5\,\mu\mathrm{m}$ was achieved in the lab, with the MRF field $\omega_{1,2,3} = 2\pi \times (3.6, 3.8, 4.0)$ MHz and $\Omega_{1,2,3} = 2\pi \times (102, 128, 120)$ kHz at a quadrupole gradient of $211\,\mathrm{G\,cm^{-1}}$. There have also been preliminary indications of coherent splitting of a BEC between the MRF wells. Having separately achieved a single-RF with field amplitude $\Omega_1 \sim 2\pi \times 50$ kHz, it should be possible to operate this double well in the regime of a 2D gas.

The application of additional dressing fields will allow the creation of lattice structures, sculpted single wells, and asymmetric connected wells. The current experiment setup, with seven DDS channels available to generate dressing RF signals, would allow up to six separate frequencies to be applied, with circular polarisation used to load the initial well and subsequently transformed to linear polarisation for the MRF potential. A new set of DDS sources or use of an arbitrary waveform generator would allow this to be extended, subject to the 20 dBm input power limit of the amplifiers. To load these potentials, an extension of the three-RF loading scheme could be employed, with the atoms initially loaded into an AP in the middle of the frequency range and successive trapping wells added in turn using the similar amplitude ramps to those shown in Fig. 6.5. Experimental challenges will include the sensitivity of the MRF system to external noise sources, and choosing parameters that both maintain adiabaticity and minimise higher-order loss channels.

Time averaging techniques can be employed as in the single-RF AP [94, 95] to create trapping geometries including multiple wells and multiple ring structures. More speculatively, additional tuning of the potential could also be explored by working in a regime where the atoms fill a larger fraction of the shell surface, and manipulating the polarisation of each RF field component independently to make full use of the vectorial coupling properties and work with interesting surface geometries. An example of this could be to lower the barrier between the shells at points on the equator to create localised connections between the trapping surfaces.

The understanding of MRF potentials, their challenges, and the associated experimental methods that were developed in this work therefore provide a framework upon which this method can be extended in a variety of directions. The straightforward tuneability of these potentials using the RF field structure means that the trapping geometry can be significantly altered without changing the underlying structure of the trapping coils. In combination with the existing benefits of macroscopically generated adiabatic potentials, these features offer a range of new opportunities beyond the existing single-RF adiabatic potential experiments.

Chapter 8

The hybrid optical-magnetic trap

The optical dipole trap was designed and installed to sculpt the radial confinement on the adiabatic potential shell surface. The 830 nm red-detuned trapping beam is shaped using an acousto-optic deflector (AOD). While the dipole trap was originally designed to create a time-averaged Gaussian to increase the radial trapping frequency and improve the conditions for accessing strongly correlated states in rotation experiments, this configuration is also capable of almost arbitrary potential shaping, using the shell surface as a two-dimensional canvas. This can also be extended to the double-shell configuration to achieve three-dimensional potential shaping as illustrated schematically in Fig. 8.1. Beam shaping techniques using acousto-optic deflection have therefore been explored more broadly to incorporate both time-averaging and the superposition of multiple static beams.

This chapter covers the experimental methods underpinning both time-averaged and composite beam approaches in the context of the system designed for time-averaged potential painting on the AP surface. An overview of the theory of acousto-optic deflection and its influence over experimental choices is presented first in

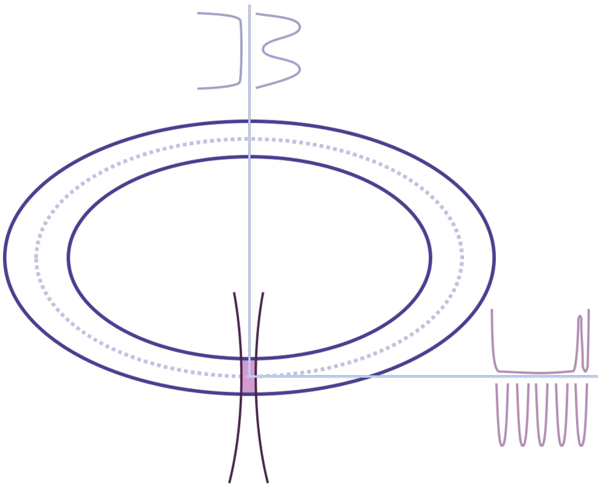

Figure 8.1: Dipole trap schematic, illustrating the intersection of the dipole trap focus with atoms confined on the shell surface to sculpt the radial potential. This can be used either independently from or in conjunction with the MRF axial potential shaping technique. The potentials sketched in the vertical and horizontal directions are representative of the type of shapes produced. Blue lines represent the shell trap surfaces, and the dotted line the MRF barrier; red curves show the dipole trap focus.

Sec. 8.1.1, followed by the results of the composite beam shaping method in Sec. 8.1.2. Section 8.1.3 then presents the rapidly-rotated dipole trap: its design requirements, experimental setup and details of the AOD programming. Progress towards a fully hybrid trapping potential is described in Sections 8.2.1 and 8.2.2, while the next steps and some longer-term applications are mentioned in Sec. 8.2.3.

8.1 Acousto-Optic Deflection: Experimental methods

An acousto-optic deflector is a diffractive optical element that uses an acoustic wave driven by an RF control signal to establish a diffraction grating for an incident laser beam. The resulting beam deflection can be used as a means of shaping the effective potential experienced by atoms confined in an optical dipole trap. Two such approaches to beam shaping are presented here: the superposition of multiple static beams, and a setup for rapidly-scanned potential painting.

8.1.1 Acousto-optic deflection

The acoustic wave in the AOD crystal is generated by radiofrequency modulation of a piezoelectric transducer, producing a diffraction grating acting in the Bragg regime. The AOD is analogous in action to that of an AOM; the differences lie in the larger, more circular, aperture (9.3 mm compared to $\sim 2.5\text{mm} \times 1\text{mm}$), the slightly slower speed of sound in the AOD crystal ($\sim 600\,\mathrm{m\,s^{-1}}$ as opposed to $\sim 4000\,\mathrm{m\,s^{-1}}$ for devices in the 780-850 nm range with a central frequency of a few tens of MHz [186–188]), and the correspondingly wider range of deflection angles (a deflection bandwidth of $\sim 1.9°$ in contrast to $\sim 0.23°$). As a result, AOMs predominantly act as switches or for frequency modulation of a laser beam, and AODs to steer beams over large deflection angles[1].

The angular deviation of the incident laser beam about the central frequency to which the AOD is aligned is given by

$$\Delta\theta = \frac{\Delta f \cdot \lambda}{u_s} \tag{8.1}$$

where u_s is the speed of sound in the AOD crystal, Δf the frequency deviation and $\lambda = 830$ nm the incident laser wavelength. To achieve beam deflection in both x and y radial directions on the AP shell surface, we use an Isomet LSA110A-830XY dual-axis AOD with a 50 MHz central frequency, about which the intensity varies by approximately 10% in a 10 MHz interval. The AOD is optimised to an 830 nm incident laser beam and its $\mathrm{TeO_2}$ crystal has a sound speed $u_s = 610\,\mathrm{m\,s^{-1}}$, inferred from measurements of angular separations for known driving frequencies. This AOD has already been successfully implemented in a pair creating a dynamic and variable standing-wave optical lattice [22, 121, 192]. As the grating is a travelling wave, each deflected beam component acquires a frequency shift at the value of the driving frequency. This

---

[1] Having said this, AOMs have also been successfully applied to optical trapping and potential shaping techniques involving both scanning and signal multiplexing [189–191].

reduces out-of-plane interference effects between components with different deflection frequencies.

An AOD was chosen over SLM or DMD techniques for this hybrid trap due to the faster associated update rate: a scan frequency on the order of 50 kHz is necessary to achieve time-averaging for atoms confined with a trap frequency of a few kHz. This AOD is limited by its sound speed of $610\,\mathrm{m\,s^{-1}}$, which corresponds to an effective rate of change of deflection angle of 140 kHz accounting for the time taken for the acoustic wave to traverse an incident beam with width 2.2 mm (experimental details are discussed in Sec. 8.1.3.1). If necessary, a higher scanning frequency could be achieved by reducing the beam diameter, but this is easily sufficient to satisfy our experimental requirements.

8.1.2 Applying composite beams

A multiplexed input RF signal can be generated using DDS or an arbitrary waveform synthesiser (AWS) to create an array of static beams. These can create distinct dipole trapping sites, applied, for example, to create an array of Josephson junctions with tunnelling across the MRF double well, or be spaced more closely to create a large continuous trapping pattern on the shell surface. The relative amplitudes within the multiplexed input determine the proportion of the total light diverted into each first diffracted order, such that the overall intensity pattern is the sum of the constituent diffracted beams. We thereby achieve precise control over both the position and amplitude of each diffracted beam, and the overall trapping pattern. This work has been published in Ref. [124].

This approach was tested independently from the main hybrid trapping setup, using a 1.3 mm radius laser beam aligned to the AOD at a driving frequency of 45 MHz, the resultant deflection pattern captured by a single-lens imaging system comprising a lens ($f = 50$ cm) placed at its focal length from the AOD, and separated by another focal length from the *Unibrain fire-i 521b* CCD camera (pixel size 9.9 $\mu\mathrm{m}$). The beam is focussed to a 0.104 mm waist at the camera, and the separation between deflected components fixed by the lens. The AOD was driven using a *Hewlett Packard 8770A* AWS, the 1 dBm AWS output amplified by a 1.6 W amplifier.

8.1.2.1 Controlling deflection using the arbitrary waveform synthesiser

For this proof-of-principle system, the calculation process for the required RF signals works backwards from the known target distribution at the CCD. The deflection frequencies are then transmitted to the AWS internal memory using a LabView GPIB interface. The AWS input parameter *n* is calculated for each frequency according to $n_i = f_i N / f_c$ where $f_c = 125$ MHz is the internal clock frequency and N the number of points per waveform period. This is ideally a multiple of both 8 and the time period associated with the lowest beating frequency in the signal as a fraction of the internal clock period, to minimise flicker during the looping of the wave segment. Furthermore, the packet length L must be less than 5320, the number of periods per waveform segment multiplied by the number of points in each period. After choosing an appropriate n_i and number of elements to satisfy these conditions and produce the correct output beam separations, the multiplexed GPIB signal for frequency components i transmitted to the AWS is of the Fourier form

$$S_{in} = \sum_i a_i \sin\left(\frac{2\pi n_i}{L}\right). \tag{8.2}$$

The values a_i indicate the amplitude associated with each frequency component f_i. The AWS then produces a multiplexed output signal f_{out} allowing independent control of frequency components and their amplitudes:

$$f_{out} = \sum_{i=1}^{N} a_i f_i \tag{8.3}$$

8.1.2.2 Creating an asymmetric double well

The multiplexed beam shaping example presented here is an asymmetric double well consisting of a single lattice point connected to a larger flat-topped reservoir [193], of interest for investigating decoherence processes. A smooth flat-topped intensity profile is difficult to obtain by laser beam shaping, as large continuous patterns are challenging for spatial light modulators to create due to optical vortex formation during calculation. However, with the AOD a scalable series of independently-deflected beams can be superimposed to directly construct the target distribution. The AOD frequencies required to generate a flat-topped potential, whether composite or time-averaged, are calculated using the Sparrow resolution criterion [194, 195]. This refinement of the Rayleigh criterion states that the composite intensity distribution of a sequence of beams is perfectly flat, and the constituent beams indistinguishable, if the second derivative is 0. The resultant separation a between adjacent beams is calculated by

$$\frac{d}{dx}\big\{f(x)+f(x+a)\big\}=0 \quad \text{and} \quad \frac{d^2}{dx^2}\big\{f(x)+f(x+a)\big\}=0 \tag{8.4}$$

The resulting two-dimensional intensity distribution $f(x,y)$ is the sum of N constituent Gaussian beams with $1/e^2$ waist w and relative amplitudes A_n and positions (x_n, y_n):

$$f(x,y)=\sum_{n}^{N} A_n e^{-2((x-x_n)^2+(y-y_n)^2)/w^2}. \tag{8.5}$$

Given the optical system and number of beams in the desired distribution, the necessary deflection angles can easily be calculated. For example, the Sparrow spacing for a 9-beam reservoir is $0.552w$, corresponding in the case of the test setup described above to deflection angles differing by 115 $\mu\mathrm{rad}$ between adjacent beams and frequency separations of 84 kHz. Following this criterion does not guarantee a perfect flat-top since the accuracy diminishes as more beams are added. However, it provides a good starting point for iterative optimisation.

The asymmetric double well shown in Fig. 8.2 has a root-mean-square relative error of 2.1% over the entire pattern region. To generate this pattern, the reservoir region was first optimised, and then two additional points added: the single well and a smaller spot controlling the barrier height. Gaussian fitting was used to calculate the initial relative intensities for this prior to an iterative optimisation procedure inspired by Ref. [196] in which a fraction of the difference between original target and measured intensities is added to the previous target on each iteration[2]. A sequence of frames would allow dynamic variation of the trapping potential to optimise loading, vary the relative widths of the wells and investigate the effects of the changing barrier height, making the most of the high AOD update rate. This method has also been applied to various other beam profiles including simple flat-topped beams, and a power-law profile designed to compensate external potentials to improve the accuracy of quantum simulation experiments [124]. These were associated respectively with RMS errors of 1.4% over the flat trapping region and 1.8% over the entire pattern.

Although this example demonstrates a single line of beams, the extension in size or dimension simply requires additional deflection frequencies up to the limits imposed by the maximum deflection angle and the $1\,\mathrm{W}$ maximum AOD input power. This method can also realise sharper edges to flat-topped beams than SLM methods, which struggle to attain a sharp edge due to the high Fourier-space frequencies required. A super-Lorentzian target profile can be used instead, with an order that compromises between flatness and calculation accuracy, but this contrasts unfavourably with the falloff gradient from the peak intensity in the composite beam approach which is limited only by the beam waist w. For example, a composite flat-top of 10 beams has an intensity drop from 95% of its maximum value to 5% over $1.6w$ whereas the eighth-order super-Lorentzian used to generate the highest accuracy SLM flat-topped

---

[2] Iterative improvement of this intensity profile was carried out manually, but the procedure has since been automated during the MPhys project of Dawei Fu.

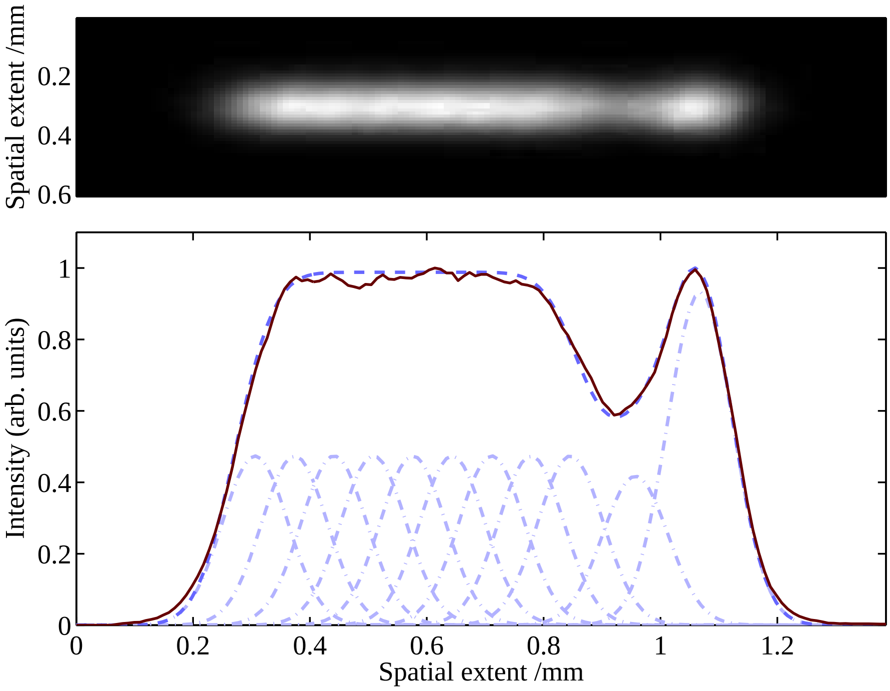

Figure 8.2: Intensity distribution of the asymmetric double well composite beam potential measured using a CCD camera (top) and corresponding line profile (bottom, solid line). The target intensity (dashed line) is the sum of the dash-dot Gaussians.

beam [197] has the same intensity falloff over $2.0w$.

One possible difficulty with using the composite beam method as a trapping potential for cold atoms is beating between adjacent spots that can cause intensity fluctuations. This is exacerbated for two dimensional deflection. A further consideration is interference that between beams deflected in x- and y-directions that acquire the same frequency shift, i.e. those that lie diagonally adjacent. This can be avoided using different frequency spacings in each direction and exploiting the naturally elliptical foci of linearly polarised beams in a high numerical aperture system [198] to maintain Sparrow spacings. Alternatively, an elliptical beam could be used or each constituent spot time-averaged to create a more exaggerated elliptical shape. The range of two-dimensional intensity patterns is also constrained by the deflection mechanism: applying an arbitrary set of deflection frequencies in the second direction reproduces the initial pattern at these intervals, but cannot create, for example, a gap within a square distribution. There are routes around this, such as creating a pattern in two or more spatially-separated segments and combining these in the trapping region using a birefringent material or separately controlled input laser beams, but this is still restricted to fairly simple patterns.

The composite beam method is a valid alternative to SLM methods of creating large continuous trapping patterns, albeit restricted to the deflection symmetry discussed above. Dynamic manipulation of these trapping potentials is straightforward given the hundreds of kHz update rate of the AOD. This method could also trivially create arrays of isolated trapping points. Furthermore, the multiplexed method combined with time-averaging would provide the means to overcome any disruption to trapped particles caused by beating of neighbouring points, by time-averaging fully separated beams to create a continuous potential.

8.1.3 The time-averaged dipole trap

The 830 nm dipole focussed on the lower AP shell surface has been designed to increase the radial trapping frequency. This increase in radial confinement will allow fast rotation of small numbers of atoms and access to strongly correlated states. The target trapping frequency is on the order of 5 kHz with a rotation frequency of trapped atoms approaching this value. To allow unhindered rotation the trapping potential must be extremely smooth, and so the trapping beam is scanned in a circular orbit to smooth out beam imperfections. Scanning must be performed at a sufficiently faster rate than the trapping frequency that determines the atom's motion [145], e.g. with a target rotation rate of 50 kHz. Again, the Sparrow criterion described in Sec. 8.1.2.2 is employed to determine the spacing between time-averaged rotation components. For a two-beam configuration this corresponds to a rotation diameter of half the focussed beam waist.

The dipole trap is implemented through the four-lens objective of the vertical imaging system described in Sec. 4.5.4.2. At 830 nm, the diffraction-limited focus has radius $1.62\,\mu\text{m}$; time-averaging as described extends the effective radius to $2.08\,\mu\text{m}$ and reduces the effective power by a factor of 0.6.

8.1.3.1 Experimental implementation

The optical system was designed such that the beam forms a focus on the lower surface of the shell, but also that the time-separated rotation components are as close as possible to parallel at this point[3]. As the four-lens objective that focusses the beam is aligned vertically to the reference point provided by the TOP centre, the shell surface is shifted to this point using a vertical bias field. The chromatic shift associated with 830 nm light is $77.5\,\mu\text{m}$ beyond the 780 nm focus, and so the capability to correct for this, and make any other live adjustments of the focus location, was incorporated into the optical system design. The laser beam should almost fill the objective to achieve a diffraction-limited focus, and a large beam also improves the AOD diffraction efficiency. Critically, lens focal lengths must also be such that they fit in around the existing vertical imaging system.

A dual-telescope system was constructed after the AOD. A schematic of this optical setup is shown in Fig. 8.3, indicating the calculated position and width of the deflected beam. The constituent lenses have focal lengths $f_1 = 15$ cm, $f_2 = 20$ cm, $f_3 = 25$ cm, and $f_4 = 50$ cm. These lenses, labelled L_1 to L_4, are placed with confocal spacing, the first lens its focal length from the AOD, and the last confocal with the back focal plane of the objective. This ensures that wherever the beam forms a focus (after the first, third, and objective lenses), a deflected component propagates parallel to the optical axis, and that it is collimated in the intervening regions in which it crosses the axis. The position of the third lens is adjusted to correct for the chromatic shift; the necessary $77.5\,\mu\text{m}$ backward shift can be achieved with a translation of 1.373 mm. The deflection angle required for half beam waist rotation in this system is $0.001\,65^\circ$, corresponding

---

[3] Time-averaging would be ineffective if the post-objective focus coincides with the optical axis as this would only produce rotation about a point.

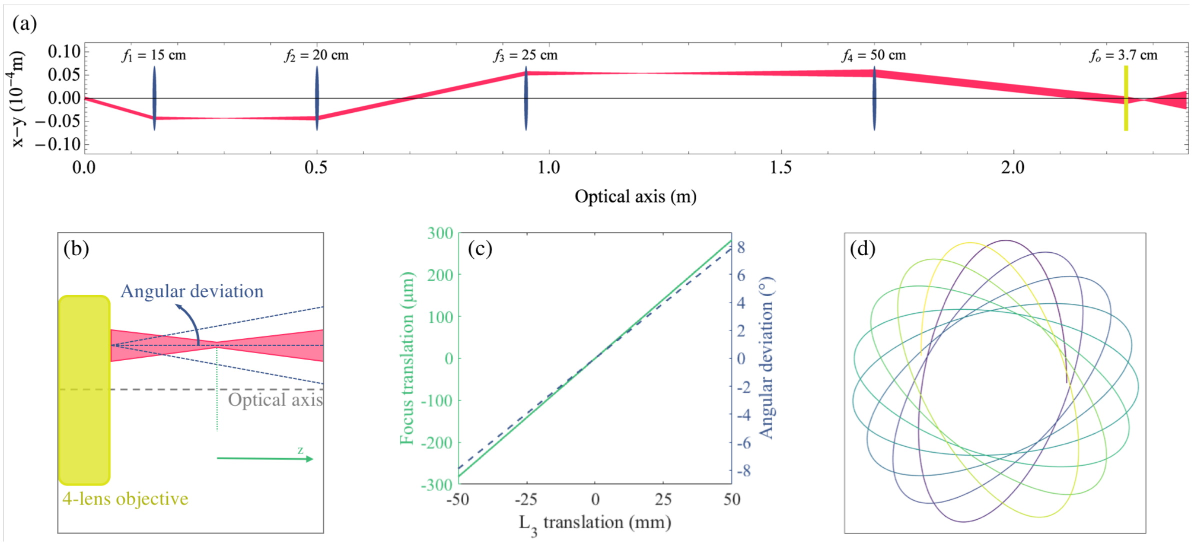

Figure 8.3: Rotating dipole trap beam propagation. (a) Calculated propagation of the deflected beam, showing deflection in the x-y plane plotted against the optical axis, without chromatic shift compensation. The dual telescope sequence enlarges the beam to a 5.8 mm radius to obtain a diffraction-limited trap focus while maintaining freedom of construction. The axis of rotation of the focus is as close as possible to the optical axis, at an angle of $-2.14°$ following the 13.73 mm translation of L_3. To aid visualisation, the displacement of the beam from the optical axis is magnified by a factor of 10^4; beam widths have not been magnified. Lenses are not to scale. (b) Sketch of the key characteristics of the beam focus following the objective, showing the z-axis along which the focus can be translated by moving L_3, and the corresponding change in orientation of the propagation axis. (c) Calculated changes in the location of the focus and angle of beam propagation, with a translation of L_3 by ± 5 cm. Translation of this lens gives a level of control over the location of the focus well-suited to the translation stage precision of $\pm 500\,\mu\text{m}$, with a full extent of travel of several cm. (d) Scanned trajectory required to create a time-averaged rotating ellipse to induce rotation in trapped atoms, as calculated in Sec. 8.1.3.2. The colour change along the trajectory indicates the evolution in time, from purple (dark) to yellow (light).

to a 21.0976 kHz frequency shift. Chromatic shift correction results in a deviation of $-2.14°$ of the post-objective beam from parallel propagation, which should correspond to corrected deflection frequency of 21.0964 kHz to ensure the correct rotation radius. This is so small as to be within alignment uncertainty and so will be accounted for in the last stages of alignment.

The full optical system is shown in Fig. 8.4. The 830 nm source is a free-running 50 mW laser diode (*Thorlabs HL8338MG*), controlled using a home-built current driver, and temperature-stabilised (*Newport 325*)[4]. A filter at 830 nm (*Semrock FF01-*

---

[4] This temperature controller was later found to produce noise between 3.4 and 3.6 MHz in pulses

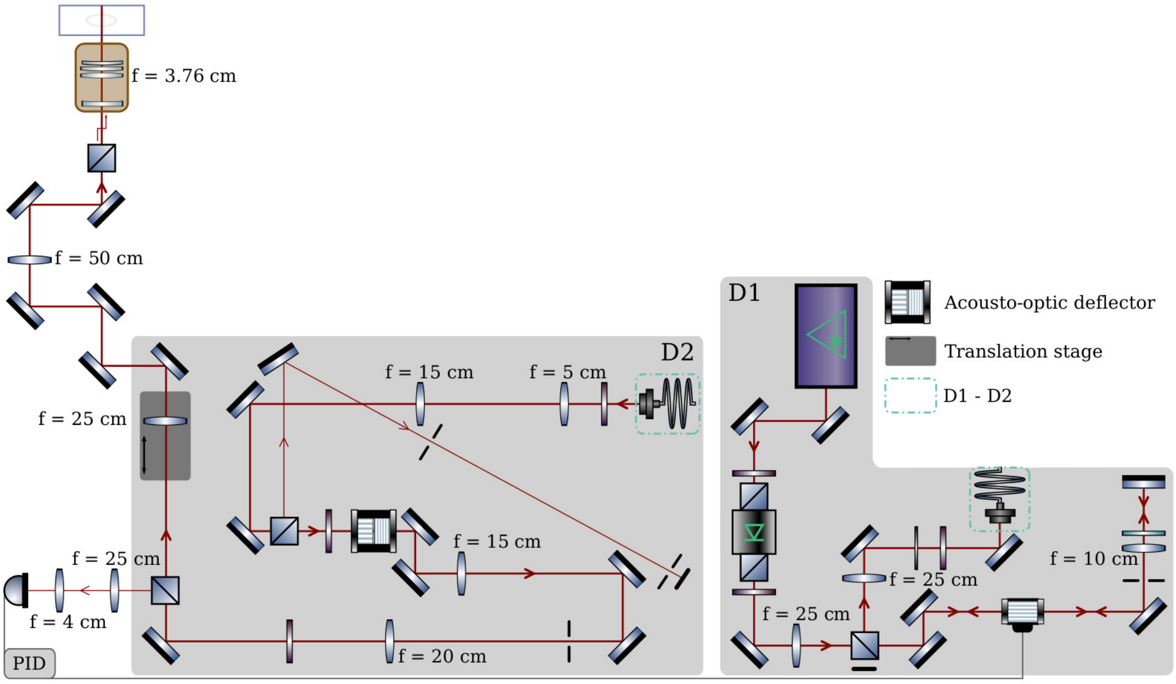

Figure 8.4: The optics layout for the rotating dipole trap. This is constructed over two breadboards for ease of alignment to the shell trap, as discussed in detail in Sec. 8.2.1. The later stages of the rotating beam alignment move from breadboard D2 to the optical table in order to fit in with the existing system. Final alignment to the shell trap is performed using the last pair of steering mirrors before the polarising beamsplitter cube that combines the vertical imaging and dipole trapping systems. Following this cube, a mirror diverts the beam onto a vertical path through the four-lens objective. The rotating beam power is stabilised using the AOM on board D1, with PID control taking as an input the beam power measured by the photodiode after the first deflection-shaping telescope.

830/2-25) narrows the linewidth. Control of the source light is performed on a separate optical breadboard from the deflection, both to maintain independence and preserve beam quality. This source control board, labelled D1 in Fig. 8.4, contains a double-pass AOM that is used to stabilise the beam intensity during its subsequent rotation via PID control (*SRS SIM960*)[5].

Once transported to the deflection breadboard, labelled D2, the beam is expanded to a 2.2 mm beam width to maximise AOD efficiency while maintaining beam quality,

---

of a few second duration (presumably when active temperature control was occurring). This caused significant atom loss from the $\omega = 2\pi \times 3$ MHz shell and so was removed.

[5] This setup could also allow for multiple input laser frequencies for a Raman dressing scheme [103].

circularly polarised, and then aligned to the AOD such that the 1-1 order emerges horizontal to the board at the central deflection frequency of 50 MHz. The AOD control frequencies are provided by two DDS channels. Prior to this, the beam polarisation is cleaned using a polarising beamsplitter cube (PBS) directly before the AOD input. The leakage beam from the PBS is aligned through two irises to provide an alignment reference to aid recovery in case of AOD misalignment. Following the AOD, the beam path severely constrained by the available space, is the dual telescope discussed above. These telescopes are aligned along the optical axis by ensuring that off-axis beams deflected at different frequencies are parallel in the required regions directly after L1 and L3, and that they are collimated after L2 and L4; final alignment to the objective is described in Sec. 8.2.1. Uncertainty in lens position on the order of ± 0.5 cm is introduced by the separation between the AOD crystals, which prevents perfect alignment to the beams deflected in both x and y. Lens alignment was therefore performed to the centre point between these. Care was taken to align the translation axis of L3 to the optical axis. Due to a height disparity between the breadboard-mounted AOD and the vertical imaging mirror to which the beam must be aligned, a gradual beam altitude decline totalling 2 cm is mediated at each uninterrupted mirror pair. The telescopes transform the beam to a collimated measured width of 5.8 mm prior to the four-lens objective (with L3 in its neutral position without chromatic shift correction), sufficient to achieve a diffraction-limited focus. An iris blocks unwanted diffraction orders.

The dipole trapping beam is combined with the vertical imaging system at a polarising beamsplitter cube with favourable dichroic properties. This cube is associated with a loss of $\sim 20\%$ as opposed to $\sim 5\%$ using a dichroic filter (*Semrock FF801-Di02-25x36*). However, the Semrock filter introduced astigmatism to the reflected vertical imaging beam, with the specified curvature tolerance associated with the coating process sufficient to explain the astigmatic difference observed (the distance between the meridional and sagital foci). Another attempt using a dichroic mirror caused ghosting in the imaging system due to reflections from both mirror surfaces, and so the beamsplitter was used.

A few μW are split from the main beam between L2 and L3. A lens with $f = 25$ cm placed at the sum of focal lengths from L2 causes beam propagation parallel to the optical axis; a 4 cm focal length lens placed 8 cm from this then focusses the beam to coincide with the point of intersection with the optical axis. A photodiode at this point provides the input to the PID system that controls the AOM transmission and acts to stabilise the beam power during AOD rotation.

Initial testing of this system was performed using a 50 cm focal length imaging lens, and later using a replica objective[6] and the focus imaged with a microscope objective (*Nikon Plan Fluor 100x*). Beam deflection and scanning worked as expected; minimal corrections should therefore be necessary to the final system, including adjustments to the focus using L3, and beam deflection frequencies.

8.1.3.2 A rotating dipole trap for atoms

To induce rotation of small numbers of atoms confined in the time-averaged dipole trap, the beam can be scanned to create an elliptical potential, and the symmetry axis of the painted ellipse then itself rotated. This can then be smoothly deformed back into a circular potential once rotation has been initiated. This can be achieved in a straightforward manner using the DDS.

The following parameters define the properties of the 'static' circular and elliptical time-averaged beams, and the subsequent rotating ellipse:

- ω_C: Central (offset) AOD driving angular frequency ($2\pi \times 50$ MHz) - defines the static location of the deflected beam
- ω_M: Driving angular frequency modulation ($2\pi\times \sim 20$ kHz) - deviation from the central frequency that determines the amplitude of the beam position modula-

---

[6] We have two objectives with nominally the same optical components and properties. One was assembled in the lab, and the other by the lens manufacturer.

tion

- ω_{TA}: Time-averaging angular frequency ($2\pi \times 50$ to $2\pi \times 100$ kHz) - rate of modulation of the central beam position
- ω_R: Rotation angular frequency ($< 2\pi \times 6$ kHz) - angular frequency of rotation for the time-averaged ellipse
- ω_{off}: Offset angular frequency ($\sim 2\pi \times 5$ MHz) - angular frequency temporarily appearing prior to the mixing stage to offset the sidebands
- E: Ellipticity of the modulation ($E = 1$ denotes circular, $E = 0$ a 1D line)

The AOD input required to produce a static deflected beam is given by

$$\begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} \cos(\omega_C t) \\ \cos(\omega_C t) \end{pmatrix}. \tag{8.6}$$

Modulation of this central frequency oscillates the deflected beam with angular frequency ω_{TA}. If the x and y axes are modulated with a $\pi/2$ phase shift then the resultant motion is circular, with an amplitude determined by ω_M. An ellipse is generated by introducing an ellipticity factor E into the modulation amplitude of one axis:

$$\begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} \cos([\omega_C + E\omega_M \cos(\omega_{TA} t)]t) \\ \cos([\omega_C + \omega_M \sin(\omega_{TA} t)]t) \end{pmatrix} \tag{8.7}$$

where the actual modulation is determined by the signal:

$$\begin{pmatrix} X_{\text{mod}} \\ Y_{\text{mod}} \end{pmatrix} = \begin{pmatrix} E\omega_M \cos(\omega_{TA} t) \\ \omega_M \sin(\omega_{TA} t) \end{pmatrix} \tag{8.8}$$

To rotate this ellipse at an angular frequency ω_R, a rotation matrix is applied to give

$$\begin{pmatrix} X_{mod} \\ Y_{mod} \end{pmatrix} = \begin{pmatrix} \cos(\omega_R t) & \sin(\omega_R t) \\ -\sin(\omega_R t) & \cos(\omega_C t) \end{pmatrix} \begin{pmatrix} E\omega_M \cos(\omega_{TA} t) \\ \omega_M \sin(\omega_{TA} t) \end{pmatrix} \tag{8.9}$$

$$= \frac{\omega_M}{2} \begin{pmatrix} (1+E)\cos([\omega_{TA} - \omega_R]t) - (1-E)\cos([\omega_{TA} + \omega_R]t) \\ (1+E)\sin([\omega_{TA} - \omega_R]t) + (1-E)\sin([\omega_{TA} + \omega_R]t) \end{pmatrix}. \tag{8.10}$$

With this modulation the AOD input becomes

$$\begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} \cos\big(\omega_C + \frac{\omega_M}{2}((1+E)\cos([\omega_{TA} - \omega_R]t) - (1-E)\cos([\omega_{TA} + \omega_R]t)))\big) \\ \cos\big(\omega_C + \frac{\omega_M}{2}((1+E)\sin([\omega_{TA} - \omega_R]t) + (1-E)\sin([\omega_{TA} + \omega_R]t)))\big) \end{pmatrix}. \tag{8.11}$$

By programming the DDS to produce the cosine of an arbitrary function, this can be directly implemented using two DDS channels. However, while the DDS can output single frequencies indefinitely by looping, frequency ramps require a sequence of updates, and stored sequences are limited to 2.1 MHz-seconds by the 64 MB memory. The ramp update frequency is determined by ω_{TA}. For example, at the maximum frequency of 100 kHz, an update rate of at least 500 kHz is necessary. With two channels, this restricts us to 2.1 s of rotation.

To produce multiple identical rotating dipole traps, multiple constant frequencies can be mixed into both channels. The effect of sidebands can be avoided by substituting $(\omega_C - \omega_{\text{off}})$ for ω_C in Eq. 8.11, with $\omega_{\text{off}} \gtrsim 5$ MHz. This would mean that the frequencies mixed into the signal are sufficiently large that the deflection efficiency at the unwanted sideband frequency would be low enough, and the deflected beam sufficiently separated from the trapping beam that it could be blocked, that this would not adversely affect the trapping potential. Single-sideband mixing was considered, but the complexity this would introduce to the system due to the requirement for additional mixers and phase shifters was not considered worthwhile. Alternatively, the additional mixing frequencies could be carefully chosen such that the sidebands themselves can be used as components of the trapping array.

The parameters E, ω_C, ω_R and ω_M can be functions of time without imposing further restrictions on rotation time, except in the unlikely case that their required update frequency exceeds that of the time-averaging. The overall signal amplitude that determines the beam intensity can also be a function of time. Some amplitude ramps will be necessary to run concurrently with the frequency modulation to stabilise the output amplitude, though a lower update rate than for the frequency would be sufficient. However, intensity ramps over a large range such as during dipole trap loading would probably be best performed externally using a VCA so as not to limit the rotation time. This sequence has been used to demonstrate rotation of a time-averaged ellipse in a test system, and imaged on a CCD. The path associated with Eq. 8.11 is plotted in Fig. 8.3(d) for realistic parameters $\omega_m = 20$ kHz, $\omega_{\text{TA}} = 100$ kHz, $\omega_R = 6$ kHz, and $E = 0.5$.

8.2 Creating the hybrid potential

The most challenging stage of setting up the dipole trap is the alignment of the beam to the atoms themselves, due to a combination of the necessary precision and a lack of access to the surrounding area. The following section gives details of the best procedure found for this, and covers the alignmnent process, some initial characterisation of the setup, and the next steps to develop the dipole trap.

8.2.1 Alignment

As shown in Fig. 8.3, the dipole beam control optics were predominantly constructed on a breadboard for modularity. However, the final integration of this setup with the vertical imaging system through which it is algined is necessarily performed on the experiment optical table and therefore takes into account the considerable apparatus and optics already present.

This process follows the accurate alignment of the vertical imaging system as described in Appendix B, which ensures that the 4-lens objective is aligned to the TOP centre, and provides a reference point on the CCD. The dipole trap location is set by the lattice alignment to be the centre of the TOP trap, and the relevant shell surface can be shifted to this point using a bias field. The 830 nm dipole trapping beam is therefore also aligned laterally through the four-lens objective and focussed at this reference point. The integration with the vertical imaging system is illustrated in Fig. 8.5. This is experimentally challenging due to the difficulty in making precise distance measurements in the vertical region, obstructed access, and challeges of making measurements above the objective which has a short focal length and is close to the glass cell. Set against this, it is necessary to place the freestanding lens L4 at the sum of focal lengths from the objective, and the preceding breadboard such that lenses L3 and L4 are separated by their sum of focal lengths. Fine corrections are possible by translating L3. This problem is tackled from both directions, and a step-by-step guide to the alignment procedure is as follows:

1. **Identify the centre of the TOP trap** This is the target for dipole trap alignment, both as a general reference point and as defined by the lattice alignment. The corresponding pixel location on each horizontal imaging camera can be identified by in-trap imaging of a TOP BEC.

2. **Align the 4-lens objective to the TOP centre** This was achieved as part of the vertical imaging alignment, described in Sec. 4.5.4.2 and Appendix B.

3. **Install supporting optics** Align the PBS (part of vertical imaging setup) and steering mirrors to the beam. To allow scope for subsequent adjustments to lens positions, these must be as close to the objective as possible.

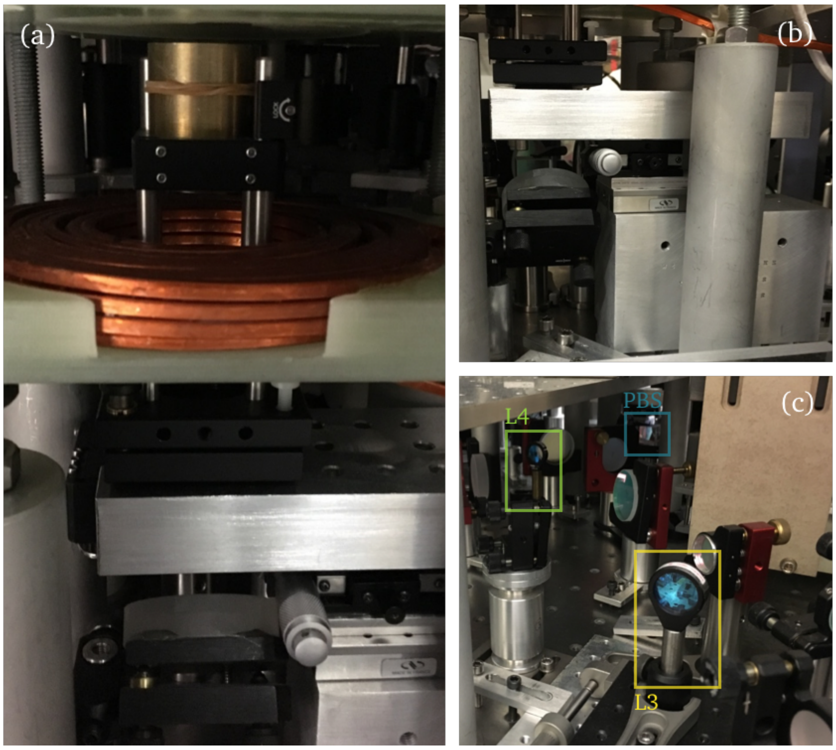

Figure 8.5: Photographs showing the integration between vertical imaging and dipole trap. (a) Vertical beam propagation through the four-lens objective (the brass tube above the copper magnetic trapping coils). (b) The objective is supported and adjusted using a translatable composite mount. This was constructed to be as stable as possible given the spatial constraints imposed by the mirror beneath. (c) The reverse side of the apparatus, showing the final telescope of the dipole trap optics. The dipole trapping beam is introduced to the vertical imaging system at the dichroic beam-splitting cube (PBS) pictured.

4. **Align L4 to the vertical imaging beam (both radial and axial positions)** The vertical imaging beam rapidly diverges following the objective and so overfills the lens after propagating the sum of their focal lengths[7]. Nonetheless, align the lens as usual, collimating the beam (assume that the incident imaging beam is collimated as established during vertical imaging alignment).

5. **Install supporting optics** Align the steering lenses between L4 and the breadboard.

---

[7] As mentioned in Sec. 4.5.4.2, a redesign of the system could be implemented in which L4 has as short a focal length as possible and sits after the PBS to form the first telescope in the vertical imaging system. This could be used to improve the signal-to-noise ratio of the vertical imaging, but would complicate the dipole trapping system due to space constraints and the requirement for beam size incident at the objective.

6. **Align dipole trapping beam to L4 and imaging system** The next step is to ensure that the breadboard is positioned such that L3 and L4 are separated by the sum of their focal lengths, and that the dipole trapping beam is correctly aligned through the objective. Ensure that L3 is at its neutral position, i.e. the sum of focal lengths from L2, and that the 830 nm beam undergoes 50 MHz deflection through the AOD.

    (a) **Approximate alignment of dipole trapping beam** Perform extremely coarse (to nearest few cm) breadboard and beam positioning to align the dipole trapping beam through L4 and the objective.

    (b) **Place guiding rails around breadboard** Achieving the correct separation between L3 and L4 requires translating the entire breadboard by several cm; this is easiest if not accompanied by lateral motion. Intercept the beam between L3 and L4 (after the breadboard) and collimate it; allow the beam to propagate several m. Clamp guiding rails either side of the breadboard such that its translation over tens of cm precisely matches the optical axis and the beam path is undeviated. Ensure that the same is true for movement of the L3 translation stage over its full range of motion; iterate between translation stage and breadboard translation until neither affects the beam path[8].

    (c) **Collimate dipole beam after L4** Translate the breadboard with L3 in its neutral position to collimate the dipole trapping beam after L4.

    (d) **Overlap trapping and imaging beams** Use the steering mirror pair immediately before f_4 to overlap the dipole trapping and vertical imaging beams. Iterate between this and step 6c to ensure that both collimation and accurate alignment have been achieved.

---

[8] This was tested after aligning to the atom cloud as described in step 7, and moving the breadboard over approximately 10 cm made no noticeable difference to the beam alignment to the atoms.

7. **Selectively repump to identify beam location** Steps 4 to 6d constitute coarse alignment of the dipole trapping beam. Finer alignment can be achieved by passing repumping light through the dipole trapping optics, to selectively repump only the narrow slice of the atom cloud illuminated by the dipole beam prior to absorption imaging in both horizontal directions. To compensate for the change in wavelength, the AOD driving frequency is changed from 50 MHz to 53.2 MHz; the location of the focus is adjusted with L3 in its neutral position, allowing for chromatic shift correction using the translation stage. The process outlined in steps 7b to 7d is iterated with progressively lower repumping powers and smaller clouds. In each of these steps, compare the response of the beam in each imaging direction to monitor any possible astigmatism. Absorption images illustrating this process are shown in Fig. 8.6.

    (a) **Define axes on the horizontal imaging CCDs** Absorption imaging in each horizontal direction is used to ensure that the beam is centred and perfectly vertical. Take a series of images of a BEC falling under different times of flight to define the vertical axis.

    (b) **Locate the beam** Use absorption imaging to locate the beam. As a general guide, a repumper power of 0.1 mW after the AOD applied to a thermal cloud directly after transport for a duration of $40\,\mu\mathrm{s}$, is useful as a starting point, brightly illuminating a broad swathe of atoms such that the beam can be detected even if slightly outside the imaging field of view. Towards the final alignment stages, a cold thermal cloud in the quadrupole or TOP traps with a repumping power of around $12\,\mu\mathrm{W}$ allows a precise image of the beam focus. The beam might be slightly out of shot in one direction but visible, fainter than expected, in the other: use this to steer it until clearly visible in both. Any strong gradient in apparent atom density, particularly if a high density is seen above the focus, can also be used to infer a tilt along the viewing axis.

(c) **Ensure that the dipole beam is both vertical and transversely aligned to the TOP centre.** Use the final steering mirror pair to walk the beam to the trap centre. Once approximately centred, continue adjusting the mirrors until it is both centred and vertical.

(d) **Translate the beam focus** Adjust the spacing between L3 and L4 to move the focus to the TOP centre. As far as possible, translate the entire breadboard to achieve this; only once the precision of this motion limits the alignment accuracy, employ the L3 translation stage. Try to restrict independent L3 translation to ~ 2 mm so as not to unduly disrupt beam rotation once implemented; save this for chromatic shift correction.

As it stands, the absorption images resulting from both the aligned vertical imaging setup and selective repumper dipole alignment provide a useful reference point that can be used to recover any drifts in the two optical systems, but most crucially in the alignment of the four-lens objective. This method also provides an early illustration of the effects of the additional diffracted orders, one of which appeared as a blurring shadow around the main repumping beam, motivating the iris placement after the AOD.

Selective repumping is therefore an efficient method to assess dipole beam alignment. However, the response of the optical system to 780 nm differs to 830 nm in both angular deflection by the AOD, and the axial location of the focus; while these effects have been calculated to mitigate their effects it is necessary to verify this experimentally and to implement a fine alignment stage using the 830 nm beam. Due to the pursuit of other experimental directions, this last alignment step has not yet been performed. This would consist of trapping atoms in the 830 nm dipole trap to verify its transverse position, and performing trap frequency measurements as the position of L3 is scanned to locate the 830 nm focus.

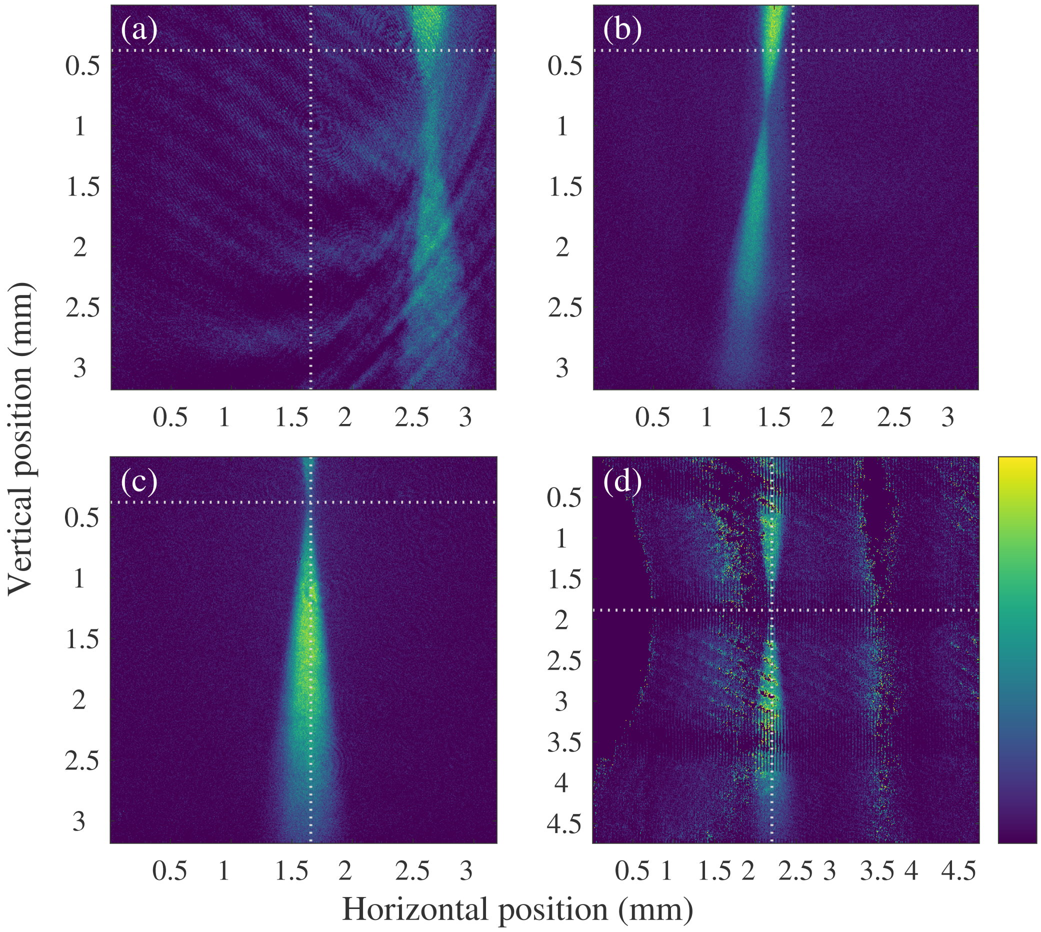

Figure 8.6: Absorption images showing the selective repumping alignment method as imaged with a thermal atom cloud directly after transport. Repumping light is diverted through the dipole trapping optics prior to imaging, illuminating only those atoms in its path. This allows efficient and intuitive alignment of the beam through the four-lens objective. (a-c) were taken using the standard horizontal absorption imaging setup, and (d) the auxiliary system in which distortion due to the curved surface of the glass cell are apparent. Dashed lines guide the eye to the measured TOP centre. The axes of each image indicate the in-trap location relative to the edge of the field of view of each CCD; the colormap has been scaled to take values between 0 and the maximum atom density for each image (except for (d), see below) and does not indicate relative atom numbers between each subfigure. (a) Initial repumping alignment absorption image, showing the first stages of coarse alignment (step 6 in the alignment guide above). This results in the correct alignment within a few hundred μm in each direction. (b) Repumping beam during alignment, as the beam is walked closer to lateral alignment with the TOP centre. (c) Final repumping dipole beam alignment after correction to beam angle and lateral position, and the axial position of the beam focus. Alignment is within $\pm 10\,\mu\mathrm{m}$ of the TOP centre, the precision limited by the imaging system resolution and identification of the focus and TOP centre. (d) Final repumping dipole beam alignment as imaged in the auxiliary horizontal imaging system; imaging iterates between the two horizontal directions during the alignment process. This image has been rescaled to reduce the effects of fringes on beam visibility, with the highest-intensity regions set artifically to zero. The focus in this direction is aligned to within $\pm 20\,\mu\mathrm{m}$ of the TOP centre.

8.2.2 Characterising the optical setup

Selective repumping can also be used to verify that the beam reacts in the anticipated manner to any changes in the optical system; this is useful both to characterise the setup and to inform subsequent experimental decisions. As an example of this, Fig. 8.7 shows the shift in the measured location of the beam focus with translation of L3, in comparison with the calculated effect. This measurement can be used to both assess the alignment of the dipole trapping system, and later to inform the ultimate translation of L_3 to compensate for the 830 nm chromatic shift.

The measurements displayed in Fig. 8.7 provide a preliminary characterisation of the system response. As the position of L_3 was varied, the focus position was extracted from each absorption image by compressing the image and fitting a 1D Gaussian, and a linear relationship between focus and L_3 positions obtained by fitting to this data. This was compared to the results of a beam propagation calculation, with an error calculated from the recognised uncertainty of ± 5 mm in the location of each telescope lens due to the separation between AOD crystals and limits on the precision of manual alignment. The variation in measured position for each value of L_3 translation warrants further investigation prior to the final alignment stages due to the spread in focus positions measured for a given L_3 location; any backlash or variation in the translation stage response should be fully characterised before this step. However, these results are promising, suggesting that the response of the focus to L_3 translation is within the combined uncertainty of the measurement, and known alignment limitations. The general alignment therefore has reasonable agreement with the modelled optical system, and will require only slight adjustments in the final fine-tuning stages.

The selective repumping images can also be used to characterise the objective alignment. For example, the position of the focus in each imaging direction, and its response to motion of either the breadboard or L3, can be used to determine any astigmatism present. Based on these measurements at the end of the repumping alignment

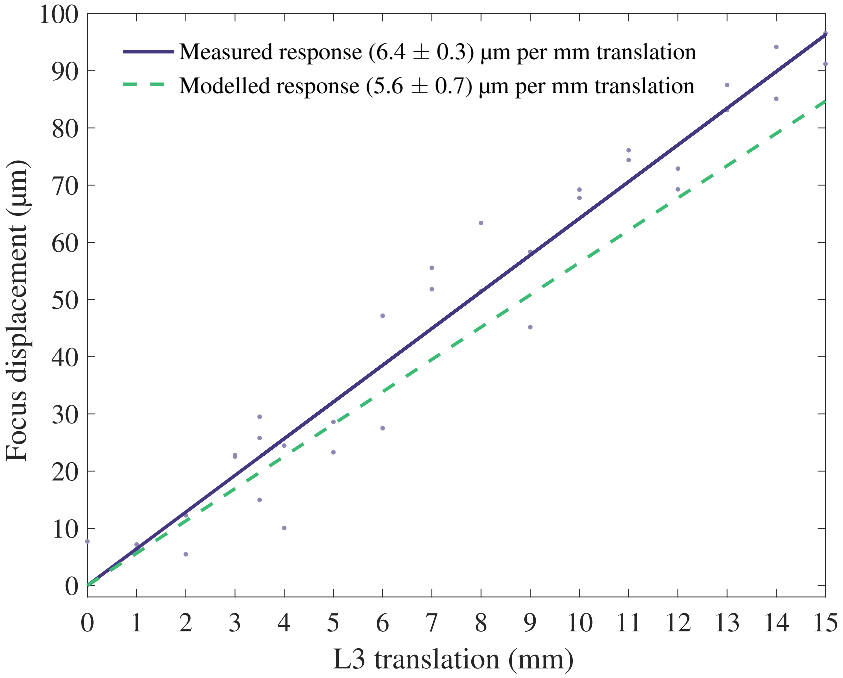

Figure 8.7: Dipole trapping optics characterisation using selective repumping: measured shift in the 780 nm selective repumping dipole beam focus with translation of L3 (data points with a linear fit, purple), plotted alongside the calculated response (green, dashed). The uncertainty in the gradient of the measured response is extracted from the 95% confidence interval of the fit, and that of the model from the estimated ±5 mm uncertainty in the alignment of each lens following the AOD resulting from the finite spacing between the two crystals. The measured and calculated responses agree within these uncertainties, suggesting that the alignment procedure was accurate within the recognised constraints.

process, an upper bound to the astigmatic difference was deduced to be $9\,\mu\mathrm{m}$. This is within the uncertainty of the position measurements in terms of both the resolution limits imposed by pixel size, and the uncertainty of the gaussian fit to the focus profile, and it is therefore reasonable to conclude that no significant astigmatism is apparent in the trapping beam.

8.2.3 Outlook

Due to conflicting experimental pressures, work on the dipole trap was unfortunately halted at this point: the laser beam focus is aligned to the TOP trap centre, and the preceding system set up to scan the dipole trap amplitude and lateral position.

Manual adjustments to the axial location of the focus are possible using the translatable lens, with a system response in reasonable agreement with the beam propagation calculation, while the modular setup allows for interchangeable laser sources. The selective repumping method provides a straightforward means of aligning and characterising the beam profile, and also a useful reference point for alignment, particularly in combination with the vertical imaging system. In light of this, the next step is to develop a beam power amplitude ramp to allow atom loading from either the TOP or shell into a hybrid trap using the 830 nm beam. Rotation experiments can then be performed with trapped atoms in both the single- and multiple-RF wells.

A tangential application of the established apparatus would be to extend the method of selective repumping used for alignment to a full imaging technique for atoms in the magnetic potentials: the beam could be scanned across the distribution, repumping only the narrow slice illuminated by the beam focus, to build up a tomographic image.

Chapter 9

Conclusion

This thesis has primarily concerned the development of an experimental apparatus to manipulate ultracold atoms confined in multiple-RF-dressed adiabatic potentials and the experimental progress towards a rapidly-scanned optical dipole trap to enhance the radial confinement on the adiabatic potential surface. Modifications to the experimental apparatus necessary to ensure its reliable operation and extension to MRF dressing have also been detailed, with particular focus on the implementation of the wideband RF impedance matching networks required to investigate the frequency regime of interest. Relevant aspects of the theory of atom confinement have also been discussed.

In this final chapter, some aspects of the current experimental work and future directions are considered, before a final reflection on the experimental achievements made so far.

9.1 Experimental outlook

This work has demonstrated a proof-of-principle experimental implementation of the MRF dressing method, using a three-component dressing field to create a double well potential. Future work can build upon this stable, highly-configurable double well implementation either by modifying the double well properties to enter a regime in which interesting phenomena can be studied, or extending the method to access a range of different trapping potentials. Some examples of immediate future work are outlined below.

9.1.1 Species-selective confinement

From Eq. 2.1 we can see that the response of an atom to a given magnetic field varies according to the g_F value of an atomic state or species confined. This change in the Zeeman splitting for a given external field strength translates into a change in the location of the resonance upon application of a dressing frequency, to create spatially distinct trapping potentials for the respective components of a mixture of atomic species. With multiple dressing frequencies applied, each produces a distinct potential for each trapped species, and with careful selection of the frequencies used we can produce a system where we consider just one potential for each species, each independently controlled by a different frequency. In a scheme described in Refs. [33] and [35], we will combine this inherent species-selectivity of the RF-dressed APs with the MRF double shell developed in this work by applying four dressing frequencies such that a double well containing ^{85}Rb impurity atoms is overlapped with a single well containing a ^{87}Rb condensate. The tunnelling properties of the ^{85}Rb impurities can then be monitored to probe the bulk condensate.

In work completed during my experimental tenure, we separately confined thermal atoms of ^{87}Rb and ^{85}Rb in a $\omega_1 = 2\pi \times 4.2$ MHz AP. During the preparation of this thesis, both species were simultaneously confined in a single-RF AP and work begun on studying the properties of both trapped species [35]. The final implementation of the impurity probe scheme requires the well spacing to be reduced to a few $\mu\mathrm{m}$ in order to access the tunnelling regime for the double well. An experimental exploration of the lifetime and properties of both species must also be performed.

In addition to the ^{85}Rb-^{87}Rb implementation, the g_F difference between $|F = 1\rangle$ and $|F = 2\rangle$ states of ^{87}Rb offer these as independently manipulable species within our current experimental setup. The method could also be extended to two or more completely different atomic species using a different apparatus.

9.1.2 Tunnelling between wells

Double well physics can be explored if the well spacing is reduced to a few $\mu\mathrm{m}$. This can be achieved by a combination of increasing the quadrupole field gradient (currently limited by the water-cooled quadrupole coils) and reducing the frequency spacing between RF components. This requires a corresponding reduction in the field amplitudes applied to create the adiabatic potentials. The main experimental challenges in achieving this have been in eliminating all RF noise sources that can cause atom loss at any stage during the necessary amplitude ramps of the RF field and quadrupole gradient. This is an ongoing line of experimental work, with extremely positive results so far.

9.1.3 A 2D gas in the MRF potential

Using the precise tuning properties of the three-RF double well, interferometry experiments could be performed using a 2D gas confined on the AP shell surfaces in analogy with the 1D interferometry experiments carried out on atom chips. In addition to the reduced well separation necessary to achieve this, a 2D condensate must first be achieved in the MRF potential. Given that single-RF field amplitudes of order $\Omega_1 \sim 50$ kHz have been previously demonstrated, the extension of this to multiple RF fields should be possible.

9.1.4 Extended potential shaping

Subject to maintaining adiabaticity, additional dressing RF components can be employed to sculpt the MRF potential, creating both lattices and continuous structures. Within the example of the three-RF double well, the ability to shape the potential using the RF field amplitudes was studied, but future work could include an exploration of additional tuning parameters such as the RF polarisation. While truly arbitrary potential shaping would not be possible in the manner of optical dipole traps, a diverse range of potentials could be explored for systems in which the properties of magnetic trapping techniques are desirable.

9.1.5 Evolution of the RF dressing hardware

While the existing RF coil array has been adequate for our needs, a subsequent evolution of the experiment could benefit from a PCB-based antenna design. These are compact, potentially improving optical access to the glass cell, and are associated with a self-resonant frequency on the order of tens of MHz which would significantly extend the range of accessible dressing frequencies. However, the associated RF field amplitude would be significantly lower than the existing dressing coils for a given input power. The existing dressing coils are sufficient for the current experiment goals.

A move towards a specialised RF PCB construction for the impedance matching networks could provide a neater alternative to the existing free-space approach, though contained in similar shielded boxes, and could be more repeatable and robust against variation in network characteristics during circuit construction. Once final experimental parameters have been chosen for MRF investigations, the match could also be narrowed around the frequency range of interest to minimise the transmission of RF noise. Again, however, the match implemented as described in Chapter 5 has been shown to provide stable operation over a frequency range wide enough to explore the full frequency range of interest for exploratory and diagnostic work without unduly limiting the lifetime of trapped atoms.

9.1.6 The hybrid optical-magnetic trap

The rapidly-scanned laser beam has been focussed onto the glass cell reference point provided by the centre of the TOP trap, to which the fluorescence imaging lattice has also been aligned [29]. A vertical bias field can be employed to translate the AP trapping surface to this point to provide axial confinement. Trapping of atoms in the 830 nm beam must be demonstrated, and the dipole trap characterised. Beam scanning to create a time-averaged rotating ellipse has been verified using a test setup, and the necessary deflection parameters for the experimental implementation calculated. This time-averaged trap must be implemented, any necessary adjustments made, and the resulting potential characterised.

This system could then be used to study the rotation of trapped atoms, ultimately working towards rapid rotation of small atom numbers. The MRF potential could also be useful here, for example by using one shell surface to impart momentum to trapped atoms [44] and the other to act as a phase reference.

The multiplexed static beam deflection can also be used to create a lattice structure on the surface of either the single-RF shell or MRF potential. Implemented with, for example, the MRF double well in a tunnelling regime, this could be used to create an array of Josephson junctions.

9.2 Conclusion

The first experimental implementation of a multiple-RF-dressed adiabatic potential has been presented, demonstrating their application to shaping the potential landscape, characterising their energy levels and showing good agreement with theoretical predications. Most importantly, techniques have been developed by which to load and manipulate these potentials. Important experimental considerations have been identified, including the effect of higher order resonances and the influence of technical features such as RF mixing and noise. As a result of this work, we can make use of the high level of configurability to confidently pursue a variety of more intricate trapping configurations. These potentials retain the traditional benefits of macroscopic RF dressing including the creation of smooth, robust trapping surfaces with an inherent state- or species-selectivity and the ability to transform between the 2D and 3D regimes of a trapped atomic gas. MRF techniques could also be employed to extend the potential shaping capabilities of any RF dressing experiment, including on atom chips.

Extensive modifications and upgrades have also been performed to an existing apparatus, ensuring the reliable and stable operation that will allow the pursuit of the experimental work identified above. As part of this work, the optical system for a rapidly-scanned optical dipole trap has been designed and constructed, and the trapping beam aligned to the atoms. The combination of dipole trap and multiple-RF dressed adiabatic potential will allow precise sculpting of the trapping potential in both axial and radial directions, and access to a vast range of trapping geometries.

The imminent application of this work to species-selective double well potentials makes use of the precise and intuitive tuning that characterises the multiple-RF dressed potential. This work lays the foundation for an extremely versatile extension of the RF dressing technique that could, in light of the increased understanding provided by this experimental characterisation, be implemented in a straightforward manner on any existing RF dressing experiment.

Appendix A

Calibrations

A.1 TOP field

TOP calibration takes two stages: converting between the computer control output voltage and resultant current through the coils, and between this current and the field experienced by the atoms. These steps are both illustrated in Fig. A.1. This first step is a simple measurement of current over a range of control voltage values using an AC current probe (*Tektronix 6019* and *6021*). A linear fit to these data points gives a current calibration of $I_x = (3.95 \pm 0.02)$ A/V and $I_y = (3.93 \pm 0.02)$ A/V with uncertainty obtained from the 95% confidence interval in each fit. These values were tuned using the attenuation stage prior to the Crown amplifier to produce a balanced TOP current in each direction for a given control voltage. These calibration values also serve as a target current response after any adjustments to the control electronics. A prefactor of 1.04 is subsequently applied to the y-TOP control voltage to balance the fields produced by the coils in each direction. This prefactor value is verified using two different methods: varying the amplitude of each field component in turn during the evaporative RF sweep to produce an equal effect on condensate production, and minimising the width of the RF spectroscopy resonance.

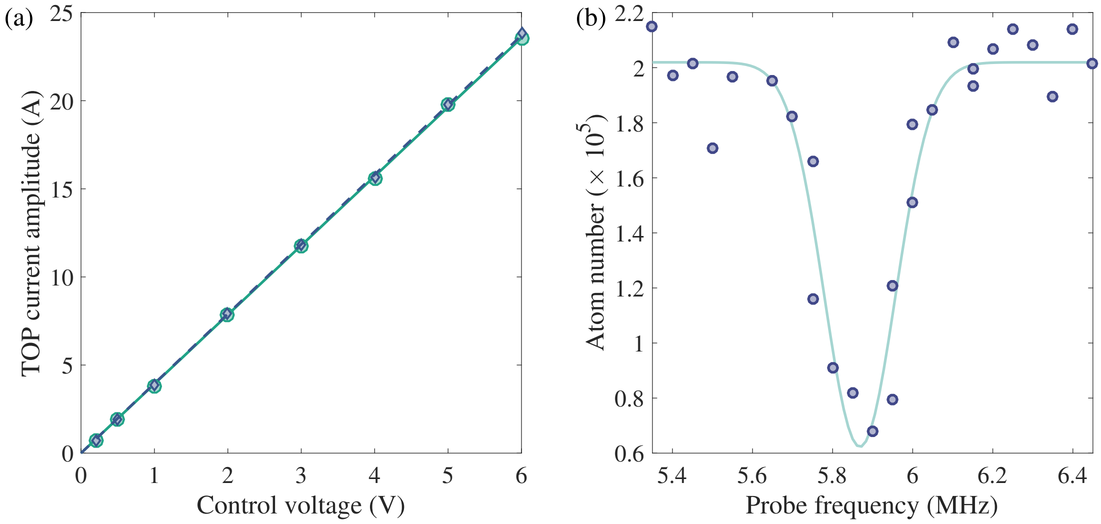

Figure A.1: Calibrating the TOP field. (a) Current response to the control voltage for x (blue) and y (green) TOP coils, balanced using the variable attenuation stage of the TOP control electronics. A linear fit to each data set yields $I_x = (3.95 \pm 0.02)$ A/V and $I_y = (3.93 \pm 0.02)$ A/V with uncertainty obtained from the 95% confidence interval in each fit. (b) An example of RF spectroscopy in the TOP trap at a measured y current of (10.23 ± 0.05) A. The fitted Gaussian has a minimum at (5.86 ± 0.01) MHz.

RF spectroscopy also provides the means of calibrating the TOP field for a given measured current; an example of this is shown in Fig. A.1(b). This particular measurement was made at a current $I_y = (10.23 \pm 0.05)$ A, with the uncertainty obtained from the variation in measured current values over several experimental runs. Spectroscopy revealed a resonance at (5.86 ± 0.01) MHz with uncertainty representing the 95% confidence interval in the Gaussian fit. With the $5^2\mathrm{S}_{1/2}$ Zeeman splitting of 0.7 MHz/G, this corresponds to a field calibration $B_{\mathrm{TOP}} = (0.820 \pm 0.004)$ G/A for the y coils, and x values differing from this by the factor $1/1.04$. For any measurements for which precise knowledge of the TOP field amplitude is required, the co-wound pickup coils are used to reliably monitor the field amplitude, calibrated using RF spectroscopy immediately prior to the measurement in question.

A.2 Quadrupole gradient

The quadrupole field gradient is a critical parameter in characterising RF-dressed potentials, and is controlled by the current flowing through the quadrupole coil. This is determined by a combination of the Magna Power output current and the MOSFET bank control voltages, and is continuously monitored using a DC current probe (*Fluke i1010*). The calibration between this current and the resulting field gradient can be obtained by measuring the radial trapping angular frequency ω_x of an atom cloud confined in the TOP according to [199]

$$B'_q = \left(\frac{2 B_T m \omega_x^2}{g_F m_F \mu_B} \right) \tag{A.1}$$

where B_T is a known TOP field amplitude and $m = 1.443 \times 10^{-25}$ the atomic mass [148].

The TOP field is measured using RF spectroscopy as described above and the resonance plotted in Fig. A.2(a) with a parabolic fit of the resonance minimum. This resonance is converted into a field amplitude using the known Zeeman shift of 0.7 MHz/G in the $5^2\mathrm{S}_{1/2}$ state. TOP trap frequency measurements for this known TOP field of (8.17 ± 0.02) G, at a measured current of (224.4 ± 0.2) A are shown in Fig. A.2(b), with radial motion of the atom cloud perpendicular to the imaging direction instigated using a kick from the Push coil.

This yields a quadrupole calibration of $B'_q = (1.173 \pm 0.001)$ G/cm/A. The uncertainty is calculated by a propagation of the errors in fitting the TOP resonance and trap frequencies, both obtained from the 99% confidence interval of the fit, and the 0.2 A uncertainty estimated from observation of the DC probe variation.

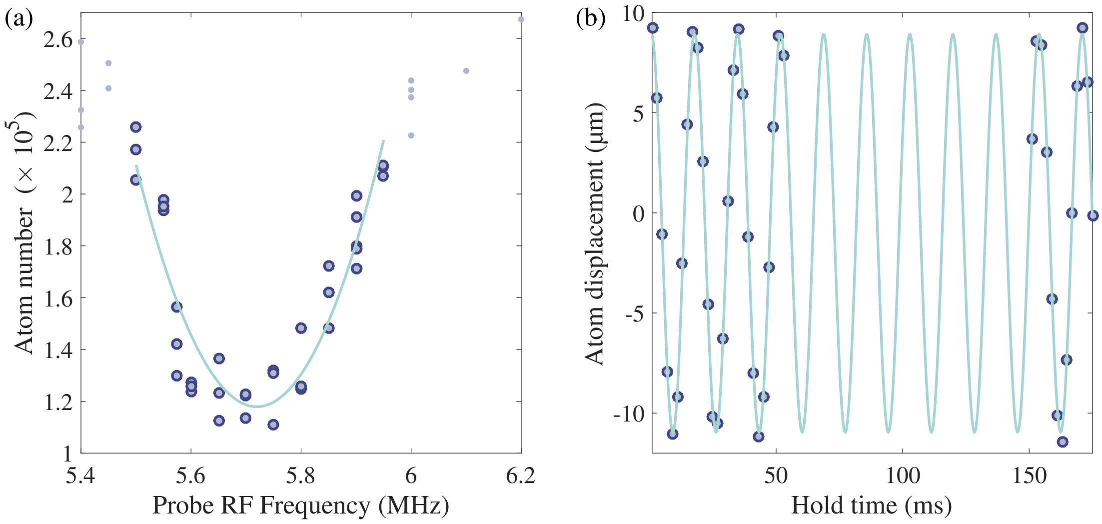

Figure A.2: Calibrating the TOP and quadrupole fields. (a) RF spectroscopy on a BEC confined in the TOP, with a minimum (5.72 ± 0.01) MHz extracted from a parabolic fit of the minimum. (b) TOP trap frequency measurement following a kick in the y direction using the Push coil. A sinusoid is fitted to this data and a trap angular frequency of $\omega_x = 2\pi \times (58.73 \pm 0.03)$ Hz extracted from the fit. These values are used to calculate the quadrupole gradient in G/cm at this current.

A.3 Imaging system magnification

Precise knowledge of the imaging system magnification is imperative to our reliance on atom imaging as a source of system diagnostics. The magnification was calibrated using the separation of m_F components under RF time of flight. Upon application of a resonant RF pulse to a condensate confined in the AP, the three m_F components are separated, with $m_F = 0$ undergoing standard freefall under the influence of gravity g, $m_F = -1$ travelling towards the quadrupole centre with acceleration $a_q - g$, and $m_F = +1$ accelerating away from the quadrupole centre with $a_q + g$. The acceleration due to the quadrupole field in the vertical direction is described by $a_q = g_F \mu_B m_F B'_z/m = 2 g_F \mu_B m_F B'_q/m$. The separation of the $m_F = 0$ and $m_F = -1$ components with time is therefore given by

$$\Delta d = \frac{g_F \mu_B B'_q}{m} t^2. \tag{A.2}$$

By comparing this with the measured acceleration between the m_F components, the system magnification is obtained as 1.028. The component separation is plotted in Fig. A.3.

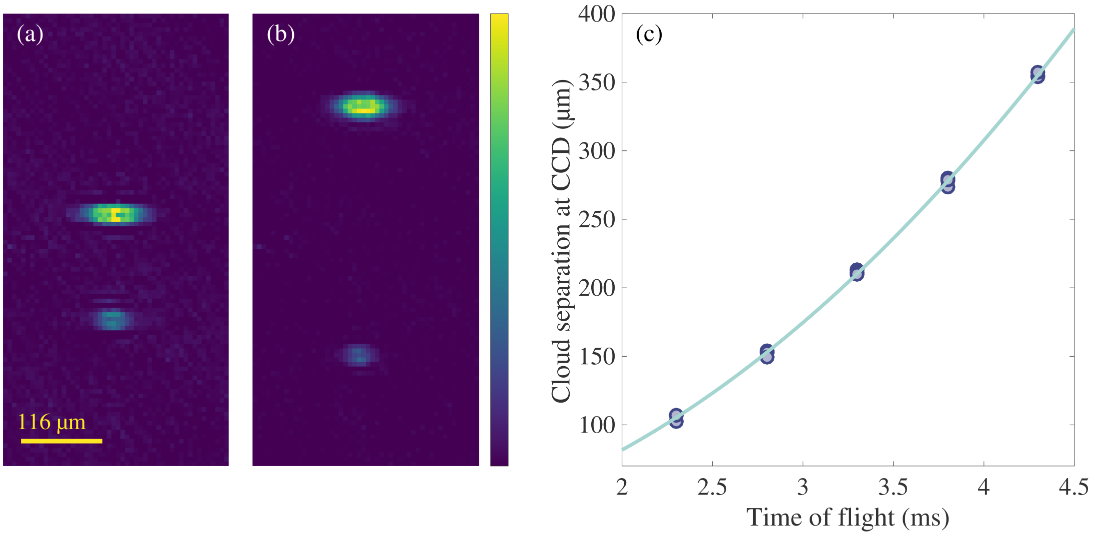

Figure A.3: Calibrating the horizontal imaging magnification by measuring the distance between m_F components at varying times of flight. (a),(b) Absorption images after 2.8 ms and 4.3 ms time of flight respectively, showing $|m_F = 0, -1\rangle$ components. These are shown in false colour, and each image normalised to its maximum value. The colourbar indicates linear scaling between zero and maximum absorption. (c) Variation in m_F component separation with time of flight. We fit the function $\Delta d = at^2 + bt + c$ to the data, comparing the coefficient $a = 20.010$ to the calculated value of 19.467 to obtain a magnification of 1.028.

Alternative means of calibrating the horizontal imaging includes simple measurements of atoms in time of flight and comparison between the measured acceleration and gravity; this method was used to calibrate the secondary horizontal imaging direction.

A.4 Vertical bias field

The vertical bias field is applied using the *Helmholtz* DC coils. Its calibration consists simply of varying the MOSFET gate voltages that control the current through the DC coils and observing the variation in the position of a BEC confined in the TOP. The condensate is released in 15 ms time of flight to allow imaging over the full range of translation. The results are shown in Fig. A.4. As this illustrates, the relationship

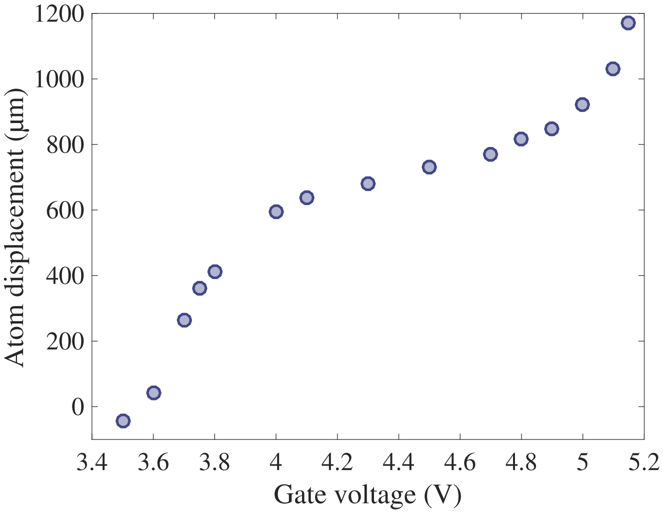

Figure A.4: Calibrating the vertical bias field: mean displacement from the zero-bias position of a TOP BEC after 15 ms time of flight for a varying bias field gate voltage.

between gate voltage and current, hence atom displacement, is highly nonlinear; interpolation between data points provides a good approximation to a desired cloud displacement, which can then be fine-tuned.

Appendix B

Guide to aligning an imaging system

This chapter briefly outlines methods used to align the imaging systems in this experiment. Commencing with the simple setup of the horizontal imaging, the bulk of the discussion focusses on the vertical imaging system, the multipurpose nature of which introduces a number of additional constraints and whose narrow depth of field makes accurate alignment more critical.

B.1 Horizontal imaging

Horizontal imaging, illustrated in Fig. 4.17, consists of two independent 1:1 telescopes in almost perpendicular configurations. For the purposes of the following description, the two lenses will be referred to as *imaging* and *focussing*, respectively referring to the lens directly following the atoms and that focussing the image onto the CCD. In terms of alignment procedure, these systems are equivalent. A step-by-step approach is outlined below:

1. **Align probe beam to atoms.** This step can be carried out simply by careful

measurement of the system geometry, ensuring that the collimated beam passes through the coil array centre, propagating parallel to the table (this is trickier than it sounds in the longitudinal configuration).

2. **Align the final lens to the probe beam.** The two-lens telescope configuration grants a reasonable amount of freedom in the space between the two lenses: the image itself is collimated in this region and the only concern is therefore in having an appropriate probe beam profile on the CCD. We make the most of this freedom in the longitudinal imaging setup (and vertical imaging system described below) to circumvent constraints imposed by a crowded optical table. Whether this is exploited or not (as in the lateral horizontal direction where the two lenses are placed as close together as possible) aligning the second lens first allows the collimated probe beam to mimic the image of the atoms in the completed setup.

3. **Align the camera to the probe beam focus.** Using the probe beam now as a reproduction of the final image behaviour, we can place the CCD at the focus of the beam with confidence that this will ultimately be the focus of the completed imaging system. This will remove any uncertainty when focussing the final image using the first lens. Make a note of the location of the beam focus.

4. **Align the first lens to the probe beam and atom location.** Align the first lens to the probe beam at approximately its focal length from the centre of the coil array. Depending on the alignment precision of the beam to the trap centre, this should be approximately aligned to the atoms also. Image the atoms.

5. **Adjust lateral alignments.** This is a good point at which to iterate the process to whatever degree is necessary to optimise alignment. Start by adjusting beam orientation such that the atom image falls in the centre, and repeat steps 1 to 4. Minimal changes should be necessary, with no adjustment to the spacing between

the second lens and CCD. By the end of this process, the image of the atoms should fall at the same position on the CCD as the beam focus in step 3.

6. **Adjust imaging lens focussing.** Align the imaging lens along the optical axis, adjusting its position until the image of the atoms is in focus.

B.2 Vertical imaging

In terms of basic procedure, the same steps apply to the vertical imaging system as to horizontal. However, this is complicated by a number of factors: lack of access for components, short depth of focus of the four-lens objective, and the rapidly diverging beam that emerges from this objective. The precision of the objective alignment is imperative not only for accurate high-resolution imaging of the atom cloud, but also to ensure that the dipole trapping beam, when perfectly aligned through the objective to minimise distortion of the trapping potential, will form its focus at the centre of the adiabatic potential shell surface.

As a result, the initial alignment procedure used a ‘dummy’ objective, namely a lens with focal length 10 cm, mounted in the same cage-mount arrangement as the real four-lens objective. A single pass through the alignment procedure with this dummy lens thus fixed the optical axis, and the real objective was inserted for the final stages of alignment.

1. **Align probe beam to the atoms.** Carefully measure the system geometry and align the beam vertically through the centre of the coil array, with the array centre defining the vertical axis. A good way to ensure this is to lay lens tissue on each accessible layer of the coil array former, marking the centre on each layer and viewing the propagating beam with an IR viewer. Accurate completion of this step is imperative for the objective alignment in step 4b.

2. **Align supporting optics.** As before, use the collimated probe beam to simulate

the image of the atoms after collimation by the objective. This setup contains two separate telescopes. As with the horizontal telescope, the inter-lens spacing is unconstrained within each telescope, while the separation between atoms and objective, and focussing lens and CCD must again be their respective focal lengths. Likewise the spacing between the two telescopes is the sum of their focal lengths. Due to the high magnification of this system and strong divergence of the beam following to the objective, the unconstrained spacings should be as small as possible to minimise the size of the imaging beam falling on the CCD, to maintain the signal-to-noise ratio in the absorption image. Ensuring that the beam is collimated after the first lens of the second telescope, the three supporting lenses, excluding the objective, can be aligned to the probe beam.

3. **Define optical axis.** Place irises along the beam path to aid subsequent objective alignment.

4. **Insert (dummy) objective.** Align the objective to the imaging beam, placing it approximately its focal length from the trap centre. This process is complicated by the vertical system:

   (a) Insert lens into cage mount and ensure that the cage mount rods are raised to an equal height above the base. With the lens approximately horizontal, align it to the imaging beam, using the axis defined by the irises. Perform this coarse alignment by moving the entire mounting block.

   (b) Place a mirror either on the upper surface of the objective mount or in an additional cage-mount lens holder. Adjust the tilt of the the reflected beam back through the fibre to ensure that the lens is normal to the imaging beam. Remove the mirror and double-check that the back-reflections from the lens surfaces follow the same path (this will be difficult with the true objective).

   (c) Place lens at approximately its focal length from the atom position and image the atoms, starting with a large cloud and adjusting the alignment until a BEC can be seen .

5. **Adjust lateral alignment.** As in the horizontal system, iterate over steps 1 to 4 to centre the probe beam on the atoms and correspondingly adjust the supporting optics. This process is easiest with the dummy objective for which the beam size is reasonably constrained. Ensure that the (dummy) objective is centred on the atoms by the end of this process. Use the fine alignment of the translation stage for this step.

6. **Adjust imaging lens focussing.** The (dummy) objective must be aligned to a reference point in the glass cell. Defined by the position of the fluorescence imaging lattice, this is the centre of the unbiased TOP trap. Therefore, bias the location of the trap centre vertically upwards such that the atoms will fall to this position over a given time of flight. Align the vertical position of the objective, making sure that in doing so the image of the atoms on the CCD does not change. If it does then repeat steps 1 to 4 again (sorry!). Make sure that the image forms a clear focus - it is far easier to adjust the supporting optics with the dummy lens than the real thing.

7. **Insert four-lens objective.** Remove the dummy objective. Maintaining the cage mount position and tilt, insert the four-lens objective and place it approximately at its focal length from the TOP centre. Use the translation stages to compensate for any change to the position of the atom cloud image on the CCD; if any such changes are more than a few pixels then check that the objective angle is correct by looking again at back-reflections, and adjust the lateral alignment. If any of these changes seem significant then return to the dummy objective and ensure that this alignment is as accurate as possible.

8. **Adjust objective focussing.** Once satisfied that the alignment is as accurate as possible, adjust the vertical position of the objective, first by hand and then fine-tuning using the z translation stage adjustment. In the last stages this, it is easier and more reliable to vary the time of flight than move the objective position.

Appendix C

Alternative approaches to wideband impedance matching

Numerous options for different approaches to wideband or multiple-band impedance matching arose from the initial calculations performed. Those discussed here are of particular interest, raising issues of expeirmental subtlety or concepts possibly worth pursuing in future. In each of these examples, only the x networks and results are shown; as in Sec. 5.6 the corresponding networks for the y coils differ only in the details, and as such do not provide any additional insight.

C.1 Original wideband impedance match

The wideband impedance match that was originally experimentally implemented is extremely similar in design to the final match discussed in Sec. 5.6, and is illustrated in Fig. C.1. As in the final network, a long transmission line is used to create a ripple on the frequency response and spread this response in a full loop around the Smith chart with multiple real-axis crossings. A shunt capacitance then transforms the region around $4\,\mathrm{MHz}$ to the $50\,\Omega$ chart origin. The corresponding reflection coefficient and

Figure C.1: Sketch of the initial wideband impedance matching network, consisting of a 14.83 m transmission line and 3.13 nF shunt capacitance separating the amplifier output and RF coil. The original calculation consisted of a 14.74 m transmission line and 3.2 nF shunt capacitor.

return loss are plotted for both the model and experimental systems in Fig. C.2.

As can be seen in the return loss plots in both cases, a small transmission peak exists in the region of 400 kHz. In practice, this opened a large window to the transmission of current at low frequencies, as shown in the RF coil current plots in Fig. C.3. Noise at a few hundred kHz is particularly problematic due to its proximity to the Rabi frequencies used in the early tests and characterisation of the single and multiple-RF dressed potentials; in the early stages of this work, noise sources ranging from hundreds of Hz to a few hundred kHz were also particularly prevalent. The transmitted noise therefore severely inhibited atom loading and the trap lifetime in the shell (on the order of hundreds of ms for thermal atoms, rather than seconds as in the final match implemented). Furthermore, the restricted power transmission around the frequencies of interest did not leave much room for varying the trap parameters during the exploratory stages. In order to increase the barrier RF amplitude Ω_2 sufficiently to flatten the barrier, preamps (*Mini-Circuits ZFL-500HLN*) were used. These did not noticeably add noise to the system and as such pose no experimental problems in principle, save for issues of scalability regarding the maximum combined input power of the amplifiers, and an additional stage of system complexity.

Due to the lifetime problems arising from noise transmission, this design was modified to incorporate the series capacitor next to the load, acting as a high-pass filter in the matching circuit, that was experimentally successful and discussed in Sec. 5.6.

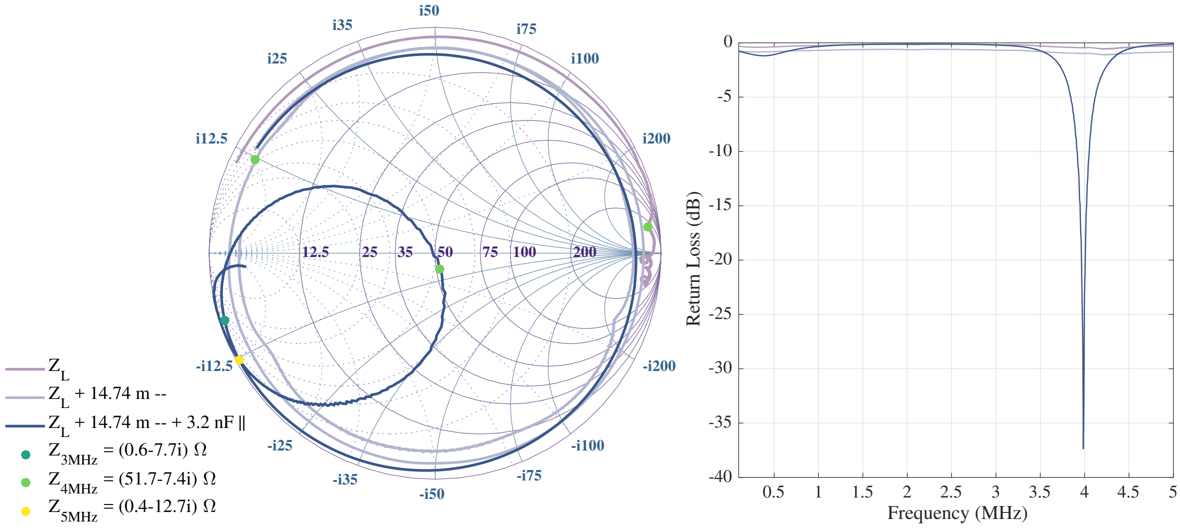

(a) Modelled matching network Smith chart **(b)** Modelled return loss plot

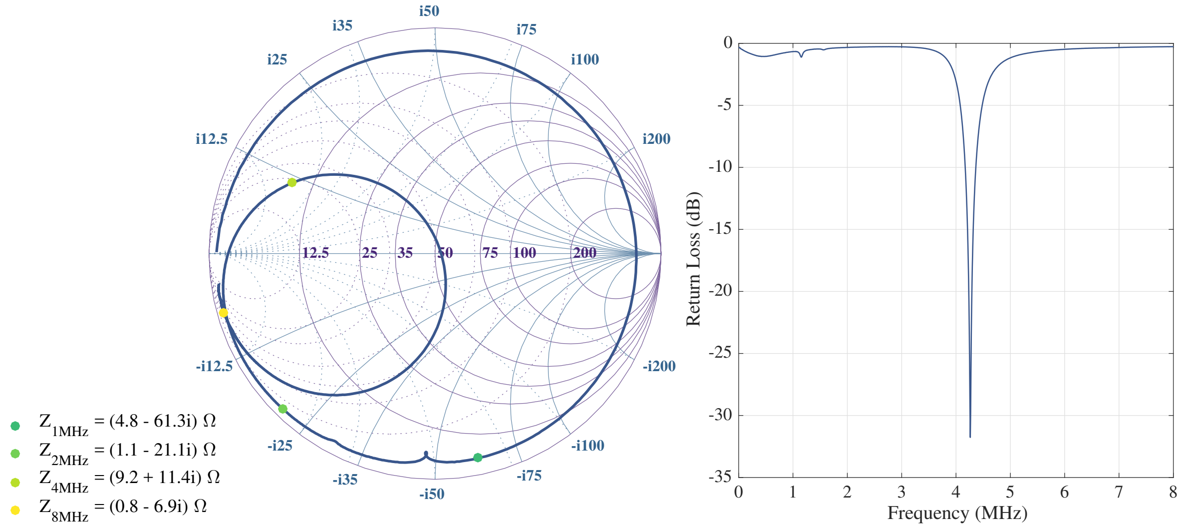

(c) Experimental matching network Smith chart **(d)** Experimental return loss plot

Figure C.2: Smith charts and return loss plots for the initial experimental wideband match. (a) and (b) show the modelled results after adding each component to the matching network; (c) and (d) show the corresponding experimental measurements. Markers have been placed on both Smith charts to illustrate the value of the impedance at certain frequencies. The 4 MHz impedance has been marked on each trace in the modelled Smith chart (a) to illustrate the effect of each successive component on the impedance transformation.

C.2 Resistive matching

Resistive matches use series or parallel resistors to traverse the lines of constant reactance and susceptance. The frequency-independence of the resistance means that series resistors, corresponding to a transformation anticlockwise along an arc of con-

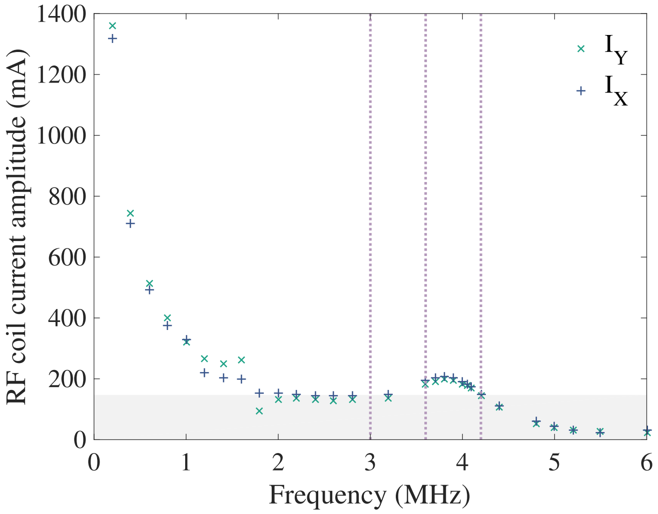

Figure C.3: Transmitted current in both x and y coils with the initial impedance matching network. Vertical dotted lines indicate the angular frequencies of the three dressing fields $\omega_{1,2,3} = 3, 3.6, 4.2$ MHz used to create and characterise the original MRF double well. The shaded region denotes the 147 mA target current amplitude. This network offers no suppression of low-frequency noise arising either from nonlinear mixing processes in the amplifiers or external sources such as ground loops.

stant reactance, effectively compress the frequency response of the network, thus reducing the Q-factor and widening the match. Benefits of the resistive approach also include the reliability of high-power commercial resistors and the relative resilience of the resulting network to stray capacitances and inductances in their experimental implementation. One example of a wideband resistive network experimentally trialled is illustrated in Fig. C.4, alongside the corresponding Smith chart and return loss of the modelled system.

The width of the match is related to its low Q, corresponding to the compression of the frequency response upon adding the final resistor. The small circle of the $\Gamma(\omega)$ trace in the stage immediately after adding the $100\,\Omega$ resistor, prior to the addition of the second transmission line and $25\,\Omega$ resistor, can also be clearly seen to lie within a low-Q arc and corresponds to a far wider transmission peak, albeit far shallower and further from resonance. The two stages of resistor and transmission line result in a lower Q than would be achieved with a single stage, due to the distortion of the circular trace by the second transmission line.

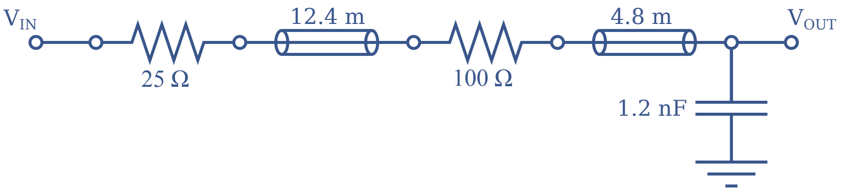

(a) Resistive matching network

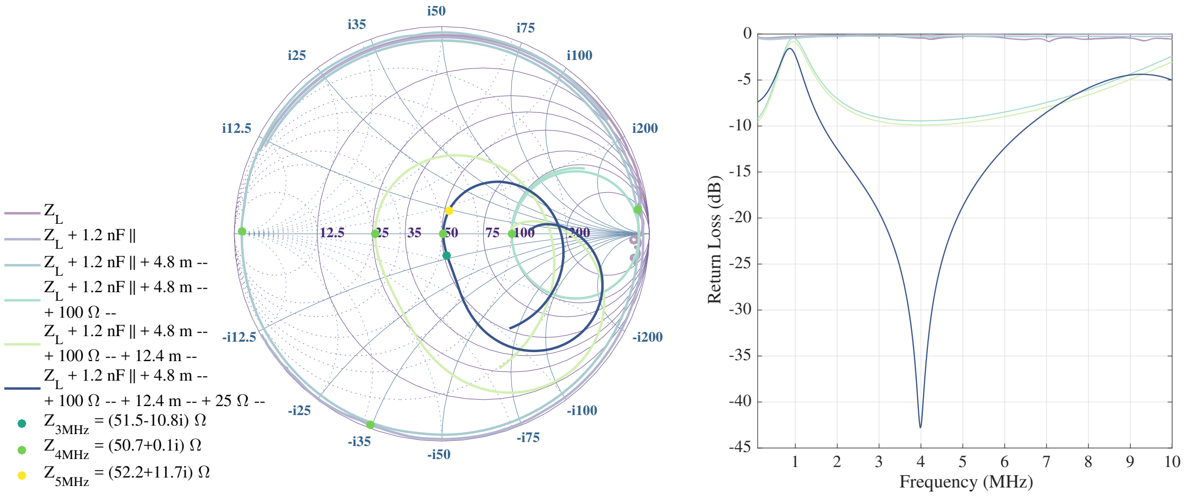

(b) Modelled Smith chart **(c)** Modelled return loss plot

Figure C.4: Example of a resistive wideband impedance match. (a) The matching network, consisting of a shunt capacitance, and two stages of transmission line and series resistance. (b) Smith chart showing the frequency-dependent reflection coefficient $\Gamma(\omega)$ after the addition of each component in the network. By transforming along the arcs of constant reactance, the series resistors act to compress the frequency response, reducing the value of Q for the final match. (c) This low Q is apparent in the calculated return loss from this network. Traces are shown for each component added to the network.

The experimental challenges of this approach relate to the dissipation of power in the series resistors; in the circuits tested, this did not outweigh the gains of reduced reflectance. Furthermore, the high-power resistors used experienced significant heating, and would be challenging to adequately heat-sink in the shielded impedance matching boxes used. Having been used to create a shell trap, this particular network would also have suffered from the low-current transmission problems of the match described in Sec. C.1, although this could have been remedied in a similar manner. Although promising in terms of the match width, this method was abandoned; however, it may be possible that a slightly different network could be found that would offset dissipation and work reliably with careful heatsinking.

C.3 Butterfly stub tuning

An example of a triple-band modelled stub matching circuit inspired by a butterfly network [181] is illustrated in Fig. C.5. This example illustrates the creation of a multiple-band match by enforcing multiple real-axis crossings. The tuning here is far from perfect: the axis crossings could be more closely clustered around the $50\,\Omega$ orgin to maximise the depth of each resonance, and for most applications would benefit

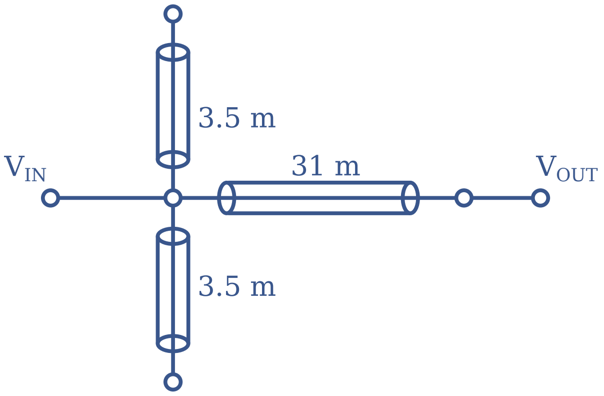

(a) Butterfly stub matching network

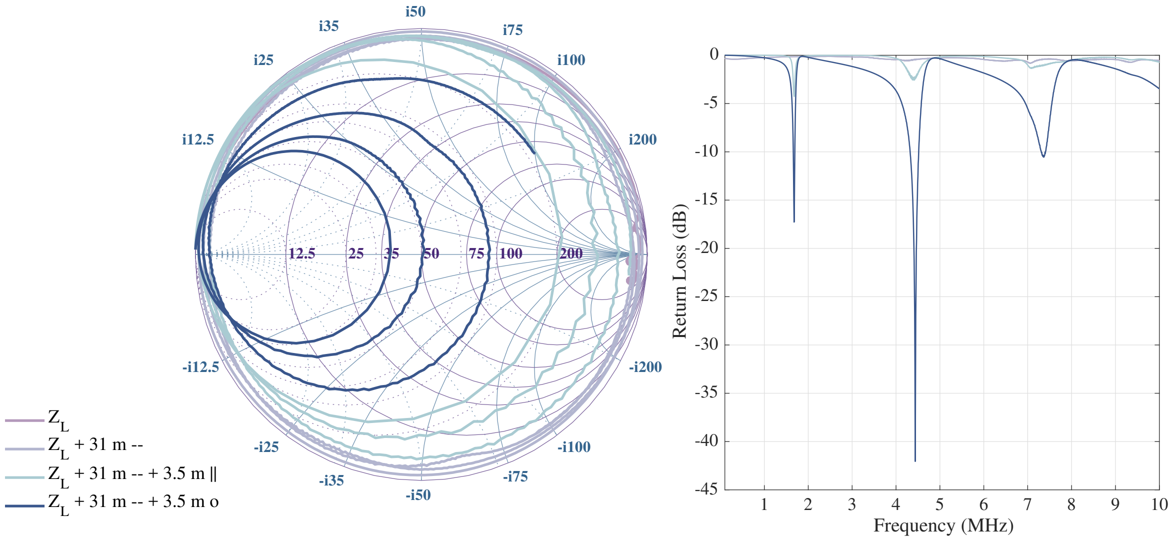

(b) Modelled Smith chart **(c)** Modelled return loss plot

Figure C.5: Example of a butterfly stub-tuned impedance match. (a) The matching network, consisting of a series and a butterfly stub transmission line. (b) Smith chart illustrating $\Gamma(\omega)$ after the addition of each component in the network. The penultimate trace corresponds to a single 3.5 m open stub, and the final trace a double stub. (c) Return loss plot following the addition of each network component. The three real axis crossings correspond to the three resonances; additional resonances could be introduced by lengthening the frequency response of the reflection trace.

from having an approximately equal width. Fine-tuning would be possible by different combinations of transmission line stages, possibly in combination with a capacitative stage.

This method was neither optimised nor implemented due in part to the impracticality of the cable lengths, in addition to the relative ease of live adjustment of capacitor networks in contrast to fine adjustment of cable lengths. Furthermore, for the reasons discussed in Sec. 5.6, a wide-band match was favoured over multiple bands. However, this example demonstrates the use of the dispersion of the reflection response by a transmission line to create multiple real-axis crossings and a multiple-band match if desired. This method, in conjunction with fine-tuning using capacitors, could feasibly be employed to create multiple narrow resonances close to the $50\,\Omega$ point should such a match be required, for example to suppress noise outside a specified set of dressing fields. However, achieving a narrow match at each of a specific set of frequencies would be challenging to achieve.

References

[1] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. “Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor”. *Science (80-.)*. 269(5221) (1995), pp. 198–201 (cited on pp. 1, 10, 13).

[2] S. Chu, L. Hollberg, J. E. Bjorkholm, A. Cable, and A. Ashkin. “Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure”. *Phys. Rev. Lett.* 55(1) (1985), pp. 48–51 (cited on pp. 1, 51).

[3] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. Van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. “Bose-Einstein condensation in a gas of sodium atoms”. *Phys. Rev. Lett.* 75(22) (1995), pp. 3969–3973 (cited on pp. 1, 10, 13).

[4] I. Bloch, J. Dalibard, and S. Nascimbène. “Quantum simulations with ultracold quantum gases”. *Nat. Phys.* 8(4) (2012), pp. 267–276 (cited on p. 1).

[5] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner. “A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice.” *Nature*, 462(7269) (2009), pp. 74–77 (cited on pp. 1, 41, 42).

[6] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard. “Optics and interferometry with atoms and molecules”. *Rev. Mod. Phys.* 81(3) (2009), pp. 1051–1129 (cited on p. 1).

[7] B. T. Seaman, M. Krämer, D. Z. Anderson, and M. J. Holland. “Atomtronics: Ultracold-atom analogs of electronic devices”. *Phys. Rev. A - At. Mol. Opt. Phys.* 75(2) (2007), pp. 1–12 (cited on p. 1).

[8] C. Ryu and M. G. Boshier. “Integrated coherent matter wave circuits”. *New J. Phys.* 17(9) (2015) (cited on p. 1).

[9] Y. H. Wang, A. Kumar, F. Jendrzejewski, R. M. Wilson, M. Edwards, S. Eckel, G. K. Campbell, and C. W. Clark. “Resonant wavepackets and shock waves in an atomtronic SQUID”. *New J. Phys.* 17(12) (2015) (cited on p. 1).

[10] T. J. Davis. “2D magnetic traps for ultra-cold atoms: A simple theory using complex numbers”. *Eur. Phys. J. D*, 18(1) (2002), pp. 27–36 (cited on p. 2).

[11] I. Lesanovsky, T. Schumm, S. Hofferberth, L. M. Andersson, P. Krüger, and J. Schmiedmayer. “Adiabatic radio-frequency potentials for the coherent manipulation of matter waves”. *Phys. Rev. A*, 73(3) (2006), pp. 1–5 (cited on pp. 2, 27).

[12] R. Grimm, M. Weidemüller, and Y. Ovchinnikov. “Optical dipole trap for neutral atoms”. *Adv. At. Mol. Opt. Phys.* 42 (2000), p. 95 (cited on p. 3).

[13] R. J. C. Spreeuw, C. Gerz, L. S. Goldner, W. D. Phillips, S. L. Rolston, C. I. Westbrook, M. W. Reynolds, and I. F. Silvera. “Demonstration of neutral atom trapping with microwaves”. *Phys. Rev. Lett.* 72(20) (1994), pp. 3162–3165 (cited on pp. 3, 25).

[14] O. Zobay and B. M. Garraway. “Two-dimensional atom trapping in field-induced adiabatic potentials”. *Phys. Rev. Lett.* 86(7) (2001), pp. 1195–1198 (cited on pp. 3, 26).

[15] C. Cohen-Tannoudji and S. Haroche. “Absorption et diffusion de photons optiques par un atome en interaction avec des photons de radiofrequence”. *J. Phys.* 30 (1969), pp. 153–168 (cited on pp. 3, 15, 25).

[16] G. Li and E. Y. Andrei. “Observation of Landau levels of Dirac fermions in graphite”. *Nat. Phys.* 3 (2007), p. 623 (cited on p. 3).

[17] T. Taychatanapat, K. Watanabe, and P. Jarillo-Herrero. “Quantum Hall effect and Landau-level crossing of Dirac fermions in trilayer graphene”. *Nat. Phys.* 7 (2011), p. 621 (cited on p. 3).

[18] S. Viefers. “Quantum Hall physics in rotating Bose–Einstein condensates”. *J. Phys. Condens. Matter*, 20 (2008), p. 123202 (cited on pp. 3, 4).

[19] H. L. Stormer. “Nobel Lecture: The fractional quantum Hall effect”. *Rev. Mod. Phys.* 71 (1999), p. 875 (cited on p. 4).

[20] D. N. Sheng, Z.-C. Gu, K. Sun, and L. Sheng. “Fractional quantum Hall effect in the absence of Landau levels”. *Nat. Commun.* 2 (2011), p. 389 (cited on p. 4).

[21] N. R. Cooper. “Rapidly rotating atomic gases”. *Adv. Phys.* 57(6) (2008), pp. 539–616 (cited on pp. 4, 5).

[22] R. A. Williams, S. Al-Assam, and C. J. Foot. “Observation of Vortex Nucleation in a Rotating Two-Dimensional Lattice of Bose-Einstein Condensates”. *Phys. Rev. Lett.* 104 (2010), p. 50404 (cited on pp. 4, 42, 151).

[23] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. “Vortex Formation in a Stirred Bose-Einstein Condensate”. *Phys. Rev. Lett.* 84 (2000), pp. 806–809 (cited on p. 4).

[24] M. R. Matthews, B. P. Anderson, P. C. Haljan, D. S. Hall, C. E. Wieman, and E. A. Cornell. “Vortices in a Bose-Einstein Condensate”. *Phys. Rev. Lett.* 83 (1999), pp. 2498–2501 (cited on p. 4).

[25] N. R. Cooper and J. Dalibard. “Reaching fractional quantum hall states with optical flux lattices”. *Phys. Rev. Lett.* 110(18) (2013), pp. 1–5 (cited on p. 4).

[26] M. Daghofer and M. Haque. “Viewpoint: Toward Fractional Quantum Hall physics with cold atoms”. 49 (2013) (cited on p. 5).

[27] N. Goldman, J. C. Budich, and P. Zoller. “Topological quantum matter with ultracold gases in optical lattices”. *Nat. Phys.* 12(7) (2016), pp. 639–645 (cited on p. 5).

[28] N. Gemelke, E. Sarajlic, and S. Chu. “Rotating Few-body Atomic Systems in the Fractional Quantum Hall Regime” (2010), pp. 1–22 (cited on p. 5).

[29] E. H. T. Owen. “Towards an Imaging Lattice for Magnetically Trapped Atoms”. DPhil Thesis. University of Oxford, 2016 (cited on pp. 5, 72, 87, 179).

[30] N. Read and N. R. Cooper. “Free expansion of lowest Landau level states of trapped atoms: a wavefunction microscope”. *Phys. Rev. A*, 68 (2003), p. 035601 (cited on p. 6).

[31] K. Merloti, R. Dubessy, L. Longchambon, A. Perrin, P. E. Pottie, V. Lorent, and H. Perrin. “A two-dimensional quantum gas in a magnetic trap”. *New J. Phys.* 15 (2013) (cited on pp. 6, 10, 30, 37, 125).

[32] J. J. P. Van Es, S. Whitlock, T. Fernholz, A. H. Van Amerongen, and N. J. Van Druten. “Longitudinal character of atom-chip-based rf-dressed potentials”. *Phys. Rev. A*, 77(6) (2008), pp. 1–9 (cited on pp. 6, 27, 28).

[33] E. Bentine, T. L. Harte, K. Luksch, A. Barker, J. Mur-Petit, B. Yuen, and C. J. Foot. “Species-selective confinement of atoms dressed with multiple radiofrequencies”. *J. Phys. B*, 50 (2017), p. 094002 (cited on pp. 7, 10, 146, 176).

[34] T. Fukuhara, A. Kantian, M. Endres, M. Cheneau, P. Schauß, S. Hild, D. Bellem, U. Schollwöck, T. Giamarchi, C. Gross, I. Bloch, and S. Kuhr. “Quantum dynamics of a mobile spin impurity”. *Nat. Phys.* 9(4) (2013), pp. 235–241 (cited on p. 7).

[35] E. Bentine. “DPhil Thesis, in preparation”. PhD thesis. University of Oxford (cited on pp. 7, 36, 37, 60, 61, 64, 124, 132, 138, 143, 145–147, 176).

[36] D. M. Stamper-Kurn, H. J. Miesner, A. P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle. “Reversible formation of a Bose-Einstein condensate”. *Phys. Rev. Lett.* 81(11) (1998), pp. 2194–2197 (cited on p. 10).

[37] Y. J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto. “Rapid production of R 87 b Bose-Einstein condensates in a combined magnetic and optical potential”. *Phys. Rev. A*, 79(6) (2009), pp. 1–8 (cited on p. 10).

[38] M. D. Barrett, J. A. Sauer, and M. S. Chapman. “All-optical formation of an atomic Bose-Einstein condensate.” *Phys. Rev. Lett.* 87(1) (2001), p. 010404 (cited on pp. 10, 41).

[39] B. M. Garraway and H. Perrin. “Recent developments in trapping and manipulation of atoms with adiabatic potentials”. *J. Phys. B At. Mol. Opt. Phys.* 49(17) (2016), pp. 1–21 (cited on p. 10).

[40] H. Perrin and B. M. Garraway. “Trapping atoms with radio-frequency adiabatic potentials”. *Adv. At. Mol. Opt. Phys.* 66 (2017), pp. 181–262 (cited on p. 10).

[41] M. H. T. Extavour, L. J. LeBlanc, T. Schumm, B. Cieslak, S. Myrskog, A. Stummer, S. Aubin, and J. H. Thywissen. “Dual-species quantum degeneracy of 40K and 87Rb on an atom chip”. *AIP Conf. Proc.* 869(2006) (2006), pp. 241–249 (cited on pp. 10, 29, 146).

[42] P. W. Courteille, B. Deh, J. Fortágh, A. Gunther, S. Kraft, C. Marzok, S. Slama, and C. Zimmermann. “Highly versatile atomic micro traps generated by multifrequency magnetic field modulation”. *J. Phys. B*, 39(5) (2006), pp. 1055–1064 (cited on pp. 10, 24, 31, 37).

[43] S. Hofferberth, I. Lesanovsky, B. Fischer, J. Verdu, and J. Schmiedmayer. “Radiofrequency-dressed-state potentials for neutral atoms”. *Nat. Phys.* 2(10) (2006), pp. 710–716 (cited on pp. 10, 28, 29).

[44] M. Gildemeister, B. E. Sherlock, and C. J. Foot. “Techniques to cool and rotate Bose-Einstein condensates in time-averaged adiabatic potentials”. *Phys. Rev. A*, 85(5) (2012), pp. 1–6 (cited on pp. 11, 29, 30, 125, 179).

[45] P. Navez, S. Pandey, H. Mas, K. Poulios, T. Fernholz, and W. von Klitzing. “Matter-wave interferometers using TAAP rings”. *New J. Phys.* 18(7) (2016), pp. 1–17 (cited on pp. 11, 30).

[46] W. Wing. “On Neutral Particle Trapping in Quasielectrostatic Electromagnetic Fields”. *Prog. Quant. Elect.* 8 (1984), p. 181 (cited on p. 12).

[47] K. J. Kügler, W. Paul, and U. Trinks. “A magnetic storage ring for neutrons”. *Phys. Lett. B*, 72(3) (1978), pp. 422–424 (cited on p. 12).

[48] A. L. Migdall, J. V. Prodan, W. D. Phillips, T. H. Bergeman, and H. J. Metcalf. “First observation of magnetically trapped neutral atoms”. *Phys. Rev. Lett.* 54(24) (1985), pp. 2596–2599 (cited on p. 12).

[49] E. A. Hinds and I. G. Hughes. “Magnetic atom optics: mirrors, guides, traps, and chips for atoms”. *J. Phys. D. Appl. Phys.* 32(18) (1999), R119–R146 (cited on p. 12).

[50] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell. “Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms”. *Phys. Rev. Lett.* 74(17) (1995), pp. 3352–3355 (cited on p. 13).

[51] D. E. Pritchard. “Cooling neutral atoms in a magnetic trap for precision spectroscopy”. *Phys. Rev. Lett.* 51(15) (1983), pp. 1336–1339 (cited on p. 14).

[52] M.-O. Mewes, M. R. Andrews, N. van Druten, D. Kurn, D. Durfee, and W. Ketterle. “Bose-Einstein Condensation in a Tightly Confining dc Magnetic Trap”. *Phys. Rev. Lett.* 77(3) (1996), pp. 416–419 (cited on p. 14).

[53] C. J. Myatt, E. A. Burt, R. W. Ghrist, E. A. Cornell, and C. E. Wieman. “Production of Two Overlapping Bose-Einstein Condensates by Sympathetic Cooling”. *Phys. Rev. Lett.* 78(4) (1997), pp. 586–589 (cited on p. 14).

[54] T. Esslinger, I. Bloch, and T. W. Ha. “Bose-Einstein condensation in a quadrupole-Ioffe-configuration trap”. *Phys. Rev. A*, 58(4) (1998), pp. 2664–2667 (cited on p. 14).

[55] J. Fortágh and C. Zimmermann. “Magnetic microtraps for ultracold atoms”. *Rev. Mod. Phys.* 79(1) (2007), pp. 235–289 (cited on p. 14).

[56] S. Kraft, A. Günther, H. Ott, D. Wharam, C. Zimmermann, and J. Fortágh. "Anomalous longitudinal magnetic field near the surface of copper conductors". *J. Phys. B*, 35 (2002), p. L469 (cited on p. 14).

[57] J. B. Trebbia, C. L. Garrido Alzar, R. Cornelussen, C. I. Westbrook, and I. Bouchoule. "Roughness suppression via rapid current modulation on an atom chip". *Phys. Rev. Lett.* 98(26) (2007), pp. 2–5 (cited on p. 14).

[58] J. A. Sauer, M. D. Barrett, and M. S. Chapman. "Storage Ring for Neutral Atoms". *Phys. Rev. Lett.* 87(27) (2001), p. 270401 (cited on p. 14).

[59] A. S. Arnold, C. Garvie, and E. Riis. "Large magnetic storage ring for Bose-Einstein condensates". *Phys. Rev. A*, 73(4) (2006), p. 041606 (cited on p. 14).

[60] P. F. Griffin, E. Riis, and A. S. Arnold. "Smooth inductively coupled ring trap for atoms". *Phys. Rev. A*, 77(5) (2008), p. 051402 (cited on p. 14).

[61] T. G. Tiecke, M. Kemmann, C. Buggle, I. Shvarchuck, W. Von Klitzing, and J. T. M. Walraven. "Bose-Einstein condensation in a magnetic double-well potential". *J. Opt. B Quantum Semiclass. Opt.* 5(1) (2003), pp. 119–123 (cited on p. 14).

[62] C. Cohen-Tannoudji and S. Reynaud. "Dressed-Atom Description of Resonance Fluorescence and Absorption-Spectra of a Multilevel Atom in an Intense Laser-Beam". *J. Phys. B*, 10(3) (1977), pp. 345–363 (cited on p. 15).

[63] C. Cohen-Tannoudji and D. Guéry-Odelin. *Advances in Atomic Physics: An Overview*. 1st. Singapore: World Scientific Publishing, 2011 (cited on pp. 15, 16, 25, 56, 58).

[64] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg. *Atom-photon interactions: basic processes and applications*. Wiley-Interscience publication. J. Wiley, 1992 (cited on pp. 15, 16).

[65] I. Lesanovsky, S. Hofferberth, J. Schmiedmayer, and P. Schmelcher. "Manipulation of ultracold atoms in dressed adiabatic radio-frequency potentials". *Phys. Rev. A*, 74(3) (2006), pp. 1–10 (cited on pp. 22, 27, 28).

[66] W. H. Heathcote, E. Nugent, B. T. Sheard, and C. J. Foot. "A ring trap for ultracold atoms in an RF-dressed state". *New J. Phys.* 10 (2008) (cited on pp. 23, 30).

[67] B. Sherlock. “Ultracold quantum gases in time-averaged adiabatic potentials”. DPhil Thesis. University of Oxford, 2011, p. 159 (cited on pp. 23, 46, 66, 71–74, 77, 124).

[68] C. Zener. “Non-adiabatic crossing of energy levels”. *Proc. R. Soc. London. Ser. A, Contain. Pap. a Math. Phys. Character*, 137 (1932), pp. 696–702 (cited on p. 24).

[69] L. Landau. “Zur Theorie der Energieubertragung. II”. *Phys. Zeitschrift der Sowjetunion*, 2 (1932), pp. 46–51 (cited on p. 24).

[70] E. Majorana. “Atomi orientati in campo magnetico variabile”. *Nuovo Cim.* 9(2) (1932), pp. 43–50 (cited on p. 24).

[71] E. C. G. Stueckelberg. “Theorie der unelastischen Stösse zwischen Atomen”. *Helv. Phys. Acta*, 5 (1933), p. 369 (cited on p. 24).

[72] C. Carroll and F. Hioe. “Generalisation of the Landau-Zener calculation to three levels”. *J. Phys. A Math. Gen.* 19 (1986), pp. 1151–1161 (cited on p. 24).

[73] K. A. Burrows, B. M. Garraway, and H. Perrin. “Non-adiabatic losses from radio-frequency dressed cold atom traps: beyond the Landau-Zener model”. *ArXiv:1705.00681* (2017) (cited on pp. 24, 25).

[74] F. Bloch and A. Siegert. “Magnetic resonance for nonrotating fields”. *Phys. Rev.* 57(6) (1940), pp. 522–527 (cited on p. 25).

[75] J. H. Shirley. “Solution of the Schrödinger equation with a Hamiltonian periodic in time”. *Phys. Rev.* 138(4B) (1965), B979 (cited on pp. 25, 34).

[76] E. Muskat, D. Dubbers, and O. Schärpf. “Dressed neutrons”. *Phys. Rev. Lett.* 58(20) (1987), pp. 2047–2050 (cited on p. 25).

[77] C. C. Agosta, I. F. Silvera, H. T. C. Stoof, and B. J. Verhaar. “Trapping of neutral atoms with resonant microwave radiation”. *Phys. Rev. Lett.* 62(20) (1989), pp. 2361–2364 (cited on p. 25).

[78] Y. Colombe, E. Knyazchyan, O. Morizot, B. Mercier, V. Lorent, and H. Perrin. “Ultracold atoms confined in rf-induced two-dimensional trapping potentials”. *Eur. Lett.* 67(4) (2004), pp. 593–599 (cited on pp. 26, 29).

[79] T. Schumm, S. Hofferberth, L. M. Anderson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Krüger. “Matter-wave interferometry in a double well on an atom chip”. *Nat. Phys.* 1 (2005), p. 57 (cited on pp. 26, 27).

[80] M. Streif, A. Buchleitner, D. Jaksch, and J. Mur-Petit. “Measuring correlations of cold-atom systems using multiple quantum probes”. *Phys. Rev. A - At. Mol. Opt. Phys.* 94(5) (2016), pp. 1–10 (cited on p. 26).

[81] J. Estève, T. Schumm, J. B. Trebbia, I. Bouchoule, A. Aspect, and C. I. Westbrook. “Realizing a stable magnetic double-well potential on an atom chip”. *Eur. Phys. J. D*, 35(1) (2005), pp. 141–146 (cited on p. 27).

[82] G. B. Jo, Y. Shin, S. Will, T. A. Pasquini, M. Saba, W. Ketterle, D. E. Pritchard, M. Vengalattore, and M. Prentiss. “Long phase coherence time and number squeezing of two Bose-Einstein condensates on an atom chip”. *Phys. Rev. Lett.* 98(3) (2007), pp. 98–101 (cited on pp. 27, 28).

[83] S. Hofferberth, I. Lesanovsky, B. Fischer, T. Schumm, and J. Schmiedmayer. “Non-equilibrium coherence dynamics in one-dimensional Bose gases”. *Nature*, 449(7160) (2007), pp. 324–327 (cited on p. 28).

[84] S. J. Kim, H. Yu, S. T. Gang, and J. B. Kim. “Matter-wave beam splitter on an atom chip for a portable atom interferometer”. *Appl. Phys. B*, 123(5) (2017) (cited on p. 28).

[85] S. Hofferberth, B. Fischer, T. Schumm, J. Schmiedmayer, and I. Lesanovsky. “Ultracold atoms in radio-frequency dressed potentials beyond the rotating-wave approximation”. *Phys. Rev. A*, 76(1) (2007), pp. 1–5 (cited on pp. 28, 30, 129).

[86] G. Sinuco-León and B. M. Garraway. “Radio-frequency dressed atoms beyond the linear Zeeman effect”. *New J. Phys.* 14 (2012) (cited on p. 28).

[87] G. A. Sinuco-León and B. M. Garraway. “Radio-frequency dressed lattices for ultracold alkali atoms”. *New J. Phys.* 17(5) (2015), p. 053037 (cited on p. 29).

[88] G. A. Sinuco-León and B. M. Garraway. “Addressed qubit manipulation in radio-frequency dressed lattices”. *New J. Phys.* 18(3) (2016), p. 35009 (cited on p. 29).

[89] G. A. Sinuco-León, K. a. Burrows, A. S. Arnold, and B. M. Garraway. “Inductively guided circuits for ultracold dressed atoms.” *Nat. Commun.* 5(May) (2014), p. 5289 (cited on p. 29).

[90] M. Vangeleyn, B. M. Garraway, H. Perrin, and A. S. Arnold. “Inductive dressed ring traps for ultracold atoms”. *J. Phys. B At. Mol. Opt. Phys.* 47(7) (2014), p. 071001 (cited on p. 29).

[91] M. White, H. Gao, M. Pasienski, and B. DeMarco. “Bose-Einstein condensates in rf-dressed adiabatic potentials”. *Phys. Rev. A*, 74(2) (2006), pp. 1–4 (cited on p. 29).

[92] K. Merloti, R. Dubessy, L. Longchambon, M. Olshanii, and H. Perrin. “Breakdown of scale invariance in a quasi-two-dimensional Bose gas due to the presence of the third dimension”. *Phys. Rev. A*, 88(6) (2013), pp. 4–7 (cited on p. 30).

[93] O. Morizot, Y. Colombe, V. Lorent, H. Perrin, and B. M. Garraway. “Ring trap for ultracold atoms”. *Phys. Rev. A*, 74(2) (2006), pp. 1–10 (cited on p. 30).

[94] I. Lesanovsky and W. Von Klitzing. “Time-averaged adiabatic potentials: Versatile matter-wave guides and atom traps”. *Phys. Rev. Lett.* 99(8) (2007), pp. 1–4 (cited on pp. 30, 147).

[95] M. Gildemeister, E. Nugent, B. E. Sherlock, M. Kubasik, B. T. Sheard, and C. J. Foot. “Trapping ultracold atoms in a time-averaged adiabatic potential”. *Phys. Rev. A - At. Mol. Opt. Phys.* 81(3) (2010), pp. 3–6 (cited on pp. 30, 125, 147).

[96] B. E. Sherlock, M. Gildemeister, E. Owen, E. Nugent, and C. J. Foot. “Time-averaged adiabatic ring potential for ultracold atoms”. *Phys. Rev. A*, 83(4) (2011), pp. 1–5 (cited on pp. 30, 125).

[97] G. A. Kazakov and T. Schumm. “Magic radio-frequency dressing for trapped atomic microwave clocks”. *Phys. Rev. A*, 91(2) (2015), pp. 1–9 (cited on p. 31).

[98] V. Bolpasi, N. K. Efremidis, M. J. Morrissey, P. C. Condylis, D. Sahagun, M. Baker, and W. Von Klitzing. “An ultra-bright atom laser”. *New J. Phys.* 16 (2014) (cited on p. 31).

[99] T. Morgan, T. Busch, and T. Fernholz. “Adiabatic potentials using multiple radio frequencies”. *arXiv:1405.2534* (2014) (cited on pp. 31, 37).

[100] D. Trypogeorgeos. “Periodically driven atomic systems”. DPhil Thesis. University of Oxford, 2014 (cited on pp. 31, 59).

[101] A. Chakraborty and S. R. Mishra. “A Floquet theory approach to study rf-dressed magnetic trap for cold atoms”. *ArXiv:1703.03552* (2017) (cited on pp. 31, 37).

[102] B. Yuen. *Multiple-radiofrequency dressed adiabatic potentials (in preparation)*. 2017 (cited on pp. 31, 32, 34, 36, 37).

[103] T. Andrijauskas, I. B. Spielman, and G. Juzeliunas. “Topological lattice using multi-frequency radiation” (2017), pp. 1–10 (cited on pp. 31, 160).

[104] T. L. Harte, E. Bentine, K. Luksch, A. J. Barker, D. Trypogeorgos, B. Yuen, and C. J. Foot. “Ultracold atoms in multiple-radiofrequency dressed adiabatic potentials”. *ArXiv:1706.01491* (2017) (cited on pp. 31, 32, 37, 124).

[105] C. J. Foot. *Atomic physics*. Oxford master series in physics. Oxford: Oxford University Press, 2005 (cited on p. 33).

[106] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. 1st ed. New York: Springer-Verlag, 1999 (cited on p. 40).

[107] S. Chu. “Laser manipulation of atoms and particles”. *Science (80-.).* 253 (), pp. 861–866 (cited on p. 40).

[108] C. Adams and E. Riis. “Laser cooling and trapping of neutral atoms”. *Prog. Quant. Elect.* 21(1) (1997), pp. 1–79 (cited on p. 40).

[109] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov. “Optical dipole traps for neutral atoms”. *Adv. At. Mol. Opt. Phys.* 42 (2000), pp. 95–170 (cited on p. 40).

[110] T. L. Gustavson, A. P. Chikkatur, A. E. Leanhardt, A. Görlitz, S. Gupta, D. E. Pritchard, and W. Ketterle. “Transport of Bose-Einstein Condensates with Optical Tweezers”. *Phys. Rev. Lett.* 88 (2001), p. 20401 (cited on p. 41).

[111] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl, R. Grimm, H.-C. Nägerl, and R. Grimm. “Bose-Einstein Condensation of Cesium”. *Science (80-.).* 299 (2003), pp. 232–235 (cited on p. 41).

[112] T. Kinoshita, T. Wenger, and D. S. Weiss. “All-optical Bose-Einstein condensation using a compressible crossed dipole trap”. *Phys. Rev. A - At. Mol. Opt. Phys.* 71(1) (2005), pp. 1–4 (cited on p. 41).

[113] J.-F. Clément, J.-P. Brantut, M. Robert-de-Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel, and P. Bouyer. “All-optical runaway evaporation to Bose-Einstein condensation”. *Phys. Rev. A*, 79 (2009), p. 61406 (cited on p. 41).

[114] D. Jacob, E. Mimoun, L. D. Sarlo, M. Weitz, J. Dalibard, and F. Gerbier. “Production of sodium Bose–Einstein condensates in an optical dimple trap”. *New J. Phys.* 13(6) (2011), p. 65022 (cited on p. 41).

[115] I. Bloch. “Ultracold quantum gases in optical lattices”. *Nat. Phys.* 1 (2005), pp. 23–30 (cited on p. 41).

[116] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr. “Single-atom-resolved fluorescence imaging of an atomic Mott insulator.” *Nature*, 467(7311) (2010), pp. 68–72 (cited on p. 41).

[117] J. Billy, V. Josse, Z. Zuo, A. Bernard, B. Hambrecht, P. Lugan, D. Clément, L. Sanchez-Palencia, P. Bouyer, and A. Aspect. “Direct observation of Anderson localization of matter waves in a controlled disorder”. *Nature*, 453 (2008), pp. 891–894 (cited on p. 41).

[118] M. Greiner, O. Mandel, T. Esslinger, T. W. Hansch, and I. Bloch. “Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms”. *Nature*, 415(6867) (2002), pp. 39–44 (cited on pp. 41, 42).

[119] R. N. Palmer and D. Jaksch. “High-field fractional quantum hall effect in optical Lattices”. *Phys. Rev. Lett.* 96(18) (2006), pp. 1–4 (cited on p. 41).

[120] C. Weitenberg, S. Kuhr, K. Mølmer, and J. F. Sherson. “Quantum computation architecture using optical tweezers”. *Phys. Rev. A*, 84 (2011), p. 32322 (cited on p. 41).

[121] A. Al-Assam, R. A. Williams, and C. J. Foot. “Ultracold atom in an optical lattice with dynamically variable periodicity”. *Phys. Rev. A*, 82 (2010), p. 21604 (cited on pp. 42, 151).

[122] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller. “Cold Bosonic Atoms in Optical Lattices”. *Phys. Rev. Lett.* 81(15) (1998), p. 3108 (cited on p. 42).

[123] V. L. Campo, K. Capelle, J. Quintanilla, and C. Hooley. “Quantitative Determination of the Hubbard Model Phase Diagram from Optical Lattice Experiments by two-parameter Scaling”. *Phys. Rev. Lett.* 99(24) (2007), p. 240403 (cited on p. 42).

[124] D. Trypogeorgos, T. Harte, A. Bonnin, and C. Foot. “Precise shaping of laser light by an acousto-optic deflector”. *Opt. Express*, 032322(21) (2013), p. 10 (cited on pp. 42, 44, 152, 155).

[125] W. Zhang, L. Li, and W. Guo. “Hard core bosons on the dual of the bowtie lattice”. *Phys. Rev. B*, 82 (2010), p. 134536 (cited on p. 42).

[126] B. Olmos and I. Lesanovsky. “Rydberg Rings”. *Phys. Chem. Chem. Phys.* 13 (2011), pp. 4208–4219 (cited on p. 42).

[127] L. Amico, A. Osterloh, and F. Cataliotti. “Quantum Many Particle Systems in Ring-Shaped Optical Lattices”. *Phys. Rev. Lett.* 95 (2005), p. 63201 (cited on p. 42).

[128] S. Franke-Arnold, J. Leach, M. J. Padgett, V. E. Lembessis, D. Ellinas, A. J. Wright, J. M. Girkin, P. Ohberg, and A. S. Arnold. “Optical ferris wheel for ultracold atoms”. *Opt. Express*, 15(14) (2007), pp. 8619–8625 (cited on p. 42).

[129] H. Xiong and B. Wu. “Atomic quantum corrals for Bose-Einstein condensates”. *Phys. Rev. A*, 82 (2010), p. 53634 (cited on p. 42).

[130] A. S. Arnold. “Extending dark optical trapping geometries”. *Opt. Lett.* 37(13) (2012), p. 2505 (cited on p. 42).

[131] V. Boyer, R. M. Godun, G. Smirne, D. Cassettari, C. M. Chandrashekar, A. B. Deb, Z. J. Laczik, and C. J. Foot. “Dynamic manipulation of {Bose-Einstein} condensates with a spatial light modulator”. *Phys. Rev. A*, 73(3) (2006), p. 31402 (cited on pp. 42, 43).

[132] A. L. Gaunt and Z. Hadzibabic. “Robust Digital Holography For Ultracold Atom Trapping”. *Sci. Rep.* 2 (2012), p. 721 (cited on pp. 42, 43).

[133] A. L. Gaunt, T. F. Schmidutz, I. Gotlibovych, R. P. Smith, and Z. Hadzibabic. “Bose-Einstein condensation of atoms in a uniform potential”. *Phys. Rev. Lett.* 110(20) (2013), pp. 1–5 (cited on p. 42).

[134] T. F. Schmidutz, I. Gotlibovych, A. L. Gaunt, R. P. Smith, N. Navon, and Z. Hadzibabic. “Quantum Joule-Thomson effect in a saturated homogeneous bose gas”. *Phys. Rev. Lett.* 112(4) (2014), pp. 1–5 (cited on p. 42).

[135] C. Eigen, A. L. Gaunt, A. Suleymanzade, N. Navon, Z. Hadzibabic, and R. P. Smith. “Observation of Weak Collapse in a Bose-Einstein Condensate”. *arXiv*, 041058(ii) (2016), pp. 1–8 (cited on p. 42).

[136] T. P. Meyrath, F. Schreck, J. L. Hanssen, C.-S. Chuu, and M. G. Raizen. “Bose-Einstein condensate in a box”. *Phys. Rev. A*, 71(4) (2005), p. 41604 (cited on p. 42).

[137] A. G. Kirk and T. J. Hall. “Design of binary computer generated holograms by simulated annealing”. *J. Mod. Opt.* 39 (1992), pp. 2531–2539 (cited on p. 43).

[138] L. Brandt, C. Muldoon, T. Thiele, J. Dong, E. Brainis, and A. Kuhn. “Spatial light modulators for the manipulation of individual atoms”. *Appl. Phys. B*, 102 (2011), pp. 443–450 (cited on p. 43).

[139] T. Harte, G. D. Bruce, J. Keeling, and D. Cassettari. “Conjugate gradient minimisation approach to generating holographic traps for ultracold atoms”. *Opt. Express*, 22(22) (2014), p. 26548 (cited on p. 43).

[140] D. Bowman, T. Harte, V. Chardonnet, C. de Groot, S. Denny, G. Le Goc, M. Anderson, P. Ireland, D. Cassettari, and G. Bruce. “High-fidelity phase and amplitude control of phase-only computer generated holograms using conjugate gradient minimisation”. *Opt. Express*, 25(10) (2017), p. 11692 (cited on p. 43).

[141] M. Pasienski and B. DeMarco. “A high-accuracy algorithm for designing arbitrary holographic atom traps”. *Opt. Express*, 16(3) (2008), pp. 2176–2190 (cited on p. 43).

[142] P. Senthilkumaran, F. Wyrowski, and H. Schimmel. “Vortex Stagnation problem in iterative {F}ourier transform algorithms”. *Opt. Lasers Eng.* 43(1) (2005), pp. 43–56 (cited on p. 43).

[143] S. A. Goorden, J. Bertolotti, and A. P. Mosk. “Superpixel-based spatial amplitude and phase modulation using a digital micromirror device.” *Opt. Express*, 22(15) (2014), pp. 17999–8009 (cited on p. 43).

[144] V. Boyer, C. M. Chandrashekar, C. J. Foot, and Z. J. Laczik. “Dynamic optical trap generation using FLC SLMs for the manipulation of cold atoms”. *J. Mod. Opt.* 51(14) (2004), pp. 2235–2240 (cited on p. 43).

[145] K. Henderson, C. Ryu, C. MacCormick, and M. G. Boshier. “Experimental demonstration of painting arbitrary and dynamic potentials for {Bose-Einstein} condensates”. *New J. Phys.* 11 (2009), p. 43030 (cited on pp. 44, 157).

[146] B. T. Sheard. “Magnetic Transport and Bose-Einstein Condensation of Rubidium Atoms”. DPhil Thesis. University of Oxford, 2010, pp. 1–209 (cited on pp. 46, 66, 67, 71–73, 75, 77–79).

[147] M. Gildemeister. “Trapping Ultracold Atoms in Time-Averaged Adiabatic Potentials”. DPhil Thesis. University of Oxford, 2010, pp. 1–146 (cited on pp. 46, 49, 66, 71, 72, 74, 124).

[148] D. A. Steck. “Rubidium 87 D Line Data” (2003), pp. 1–29 (cited on pp. 51–53, 83, 84, 183).

[149] A. Ashkin. “Trapping of atoms by resonance radiation pressure”. *Phys. Rev. Lett.* 40(12) (1978), pp. 729–732 (cited on p. 51).

[150] R. Lang. “Injection Locking Properties of a Semiconductor Laser”. *IEEE J. Quantum Electron.* 18(6) (1982), pp. 976–983 (cited on p. 53).

[151] S. D. Saliba and R. E. Scholten. “Linewidths below 100 kHz with external cavity diode lasers”. *Appl. Opt.* 48(36) (2009), p. 6961 (cited on p. 54).

[152] S. D. Saliba, M. Junker, L. D. Turner, and R. E. Scholten. “Mode stability of external cavity diode lasers.” *Appl. Opt.* 48(35) (2009), pp. 6692–700 (cited on p. 55).

[153] J. Appel, A. MacRae, and Lvo. “A versatile digital GHz phase lock for external cavity diode lasers”. *Meas. Sci. Technol.* 20(20) (2009), p. 055302 (cited on p. 55).

[154] T. W. Hänsch, M. D. Levenson, and A. L. Schawlow. “Complete Hyperfine Structure of a Molecular Iodine Line”. *Phys. Rev. Lett.* 26(16) (1971), pp. 946–949 (cited on p. 56).

[155] R. K. Raj, D. Bloch, J. J. Snyder, G. Camy, and M. Ducloy. “High-Frequency Optically Heterodyned Saturation Spectroscopy Via Resonant Degenerate Four-Wave Mixing”. *Phys. Rev. Lett.* 44(19) (1980), pp. 1251–1254 (cited on p. 57).

[156] J. H. Shirley. “Modulation transfer processes in optical heterodyne saturation spectroscopy.” *Opt. Lett.* 7(11) (1982), pp. 537–9 (cited on p. 57).

[157] F. Bertinetto, P. Cordiale, G. Galzerano, and E. Bava. “Frequency Stabilization of DBR Diode Laser Against Cs Absorption Lines at 852 nm Using the Modulation Transfer Method”. *IEEE Trans. Instrum. Meas.* 50(2) (2001), p. 490 (cited on p. 57).

[158] D. J. McCarron, S. A. King, and S. L. Cornish. “Modulation transfer spectroscopy in atomic rubidium”. *Meas. Sci. Technol.* 19(10) (2008), p. 105601 (cited on p. 57).

[159] L. Mudarikwa, K. Pahwa, and J. Goldwin. “Sub-Doppler modulation spectroscopy of potassium for laser stabilization”. *J. Phys. B At. Mol. Opt. Phys.* 45(6) (2012), p. 065002 (cited on p. 57).

[160] V. Negnevitsky and L. D. Turner. “Wideband laser locking to an atomic reference with modulation transfer spectroscopy.” *Opt. Express*, 21(3) (2013), pp. 3103–13 (cited on p. 57).

[161] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. “Laser phase and frequency stabilization using an optical resonator”. *Appl. Phys. B*, 31(2) (1983), pp. 97–105 (cited on p. 58).

[162] A. Franzen. *Component library*. URL: `http://www.gwoptics.org/ComponentLibrary/` (visited on 01/01/2013) (cited on p. 61).

[163] M. Meucci, E. Mariotti, P. Bicchi, C. Marinelli, and L. Moi. “Light-Induced Atom Desorption”. *Eur. Lett.* 25(9) (2007), pp. 639–643 (cited on p. 68).

[164] B. P. Anderson and M. A. Kasevich. “Loading a vapor-cell magneto-optic trap using light-induced atom desorption”. *Phys. Rev. A*, 63(2) (2001), pp. 1–6 (cited on p. 68).

[165] ESRF. *Radia*. URL: `http://www.esrf.eu/Accelerators/Groups/InsertionDevices/Software/Radia` (visited on 01/01/2014) (cited on p. 73).

[166] G. Zürn. “High current control for magnetic transport of cold atoms”. Project report. University of Oxford, 2008 (cited on p. 72).

[167] W. Heathcote. “A toroidal trap for ultracold atoms in a rf-dressed state A toroidal trap for ultracold atoms in a RF-dressed state”. DPhil Thesis. University of Oxford, 2007 (cited on p. 74).

[168] K. I. Lee, J. a. Kim, H. R. Noh, and W. Jhe. “Single-beam atom trap in a pyramidal and conical hollow mirror.” *Opt. Lett.* 21(15) (1996), pp. 1177–1179 (cited on p. 75).

[169] W. Petrich, M. H. Anderson, J. R. Ensher, and E. a. Cornell. “Behavior of atoms in a compressed magneto-optical trap”. *J. Opt. Soc. Am. B*, 11(8) (1994), p. 1332 (cited on p. 77).

[170] W. Ketterle, K. B. Davis, M. A. Joff, A. Martin, and D. E. Pritchard. “High densities of cold atoms in a dark spontaneous-force optical trap”. 70(15) (1993), pp. 2253–2256 (cited on p. 77).

[171] C. G. Townsend, N. H. Edwards, K. P. Zetie, C. J. Cooper, J. Rink, and C. J. Foot. “High-density trapping of cesium atoms in a dark magneto-optical trap”. 53(3) (1996), pp. 1702–1714 (cited on p. 77).

[172] M. Gajdacz, P. L. Pedersen, T. Mørch, A. J. Hilliard, J. Arlt, and J. F. Sherson. “Non-destructive Faraday imaging of dynamically controlled ultracold atoms”. *Rev. Sci. Instrum.* 84(8) (2013), p. 083105 (cited on p. 83).

[173] Y. Castin and R. Dum. “Bose-Einstein Condensates in Time Dependent Traps”. *Phys. Rev. Lett.* 77(27) (1996), pp. 5315–5319 (cited on p. 84).

[174] R. A. Williams. “Vortex nucleation in a rotating optical lattice of ultracold atoms”. DPhil Thesis. University of Oxford, 2009 (cited on p. 87).

[175] S. Al-Assam. “Quantum Simulation Using Ultracold Atoms in Two-Dimensional Optical Lattices”. DPhil Thesis. University of Oxford, 2010 (cited on p. 87).

[176] P. Baranowski. “Towards Number Correlated States of a Bose-Einstein Condensed Gas”. Master’s Thesis. University of Oxford/Freie Universitat Berlin, 2005, pp. 1–77 (cited on p. 87).

[177] W. Alt. “An objective lens for efficient fluorescence detection of single atoms”. *Optik (Stuttg).* 113(3) (2001), p. 3 (cited on p. 88).

[178] M. Born and E. Wolf. *Principles of Optics*. London: Pergamon Press, 1959, Chapter 8.8.2 (cited on p. 88).

[179] R. Sorrentino and G. Bianchi. *Microwave and RF Engineering*. Hoboken, GB: Wiley, 2010 (cited on pp. 92, 94, 100).

[180] Committee on UK Spectrum Strategy. *United Kingdom Frequency Allocation Table*. 2013. URL: `https://www.ofcom.org.uk/%7B%5C_%7D%7B%5C_%7Ddata/assets/pdf%7B%5C_%7Dfile/0016/103309/uk-fat-2017.pdf` (visited on 01/01/2017) (cited on p. 92).

[181] F. Gustrau. *RF and Microwave Engineering*. Somerset, GB: Chichester: Wiley, 2012, pp. 236–240 (cited on pp. 93, 94, 99, 102, 198).

[182] J. B. Hagen. *Radio-frequency electronics: Circuits and applications*. 2nd. Cambridge, GB: Cambridge University Press, 2009 (cited on pp. 94, 100, 102).

[183] N. Tucker. *RF Utilities V1.2*. 2010. URL: `http://uk.mathworks.com/matlabcentral/fileexchange/22996-rf-utilities-v1-2` (visited on 01/01/2013) (cited on p. 94).

[184] X. Chen, A. Ruschhaupt, S. Schmidt, A. Del Campo, D. Guéry-Odelin, and J. G. Muga. “Fast optimal frictionless atom cooling in harmonic traps: Shortcut to adiabaticity”. *Phys. Rev. Lett.* 104(6) (2010), pp. 1–4 (cited on p. 127).

[185] R. K. Easwaran, L. Longchambon, P.-E. Pottie, V. Lorent, H. Perrin, and B. M. Garraway. “RF spectroscopy in a resonant RF-dressed trap”. *J. Phys. B At. Mol. Opt. Phys*, 43(6) (2010), p. 065302 (cited on pp. 129, 143, 145).

[186] ISOMET. *Datasheet: LS110A-NIR-XY Dual Axis AO Deflector*. URL: `http://www.isomet.com/AO%7B%5C_%7DPdf/LS110A-NIR-XY.pdf` (cited on p. 151).

[187] *Application Note – Modulator Model 3000 Series*. URL: `http://sites.fas.harvard.edu/%7B~%7Dphys191r/Bench%7B%5C_%7DNotes/D4/AO%7B%5C_%7DModulator3000%7B%5C_%7Dappnote.pdf` (cited on p. 151).

[188] *Datasheet: Gooch and Housego*. 2003. URL: `https://goochandhousego.com/wp-content/pdfs/3080%7B%5C_%7D122%7B%5C_%7D97%7B%5C_%7D01280%7B%5C_%7D01rD.pdf` (cited on p. 151).

[189] N. Houston, E. Riis, and A. S. Arnold. "Dark dynamic acousto-optic ring lattices for ultracold atoms". *J. Phys. B At. Mol. Opt. Phys.* 41(21) (2008), p. 6 (cited on p. 151).

[190] T. A. Bell, J. A. P. Glidden, L. Humbert, M. W. J. Bromley, S. A. Haine, M. J. Davis, T. W. Neely, M. A. Baker, and H. Rubinsztein-Dunlop. "Bose-Einstein condensation in large time-averaged optical ring potentials". *New J. Phys.* 18 (2016), p. 035003 (cited on p. 151).

[191] Y. Shin, M. Saba, T. A. Pasquini, W. Ketterle, D. E. Pritchard, and A. E. Leanhardt. "Atom interferometry with Bose-Einstein condensates in a double-well potential." *Phys. Rev. Lett.* 92(5) (2003), p. 4 (cited on p. 151).

[192] R. A. Williams, J. D. Pillet, S. Al-Assam, B. Fletcher, M. Shotter, and C. J. Foot. "Dynamic optical lattices: two-dimensional rotating and accordion lattices for ultracold atoms". *Opt. Express*, 16(21) (2008), pp. 16977–16983 (cited on p. 151).

[193] S. Hunn, K. Zimmermann, M. Hiller, and A. Buchleitner. "Tunneling decay of two interacting bosons in an asymmetric double-well potential: A spectral approach". *Phys. Rev. A*, 87 (2013), p. 43626 (cited on p. 154).

[194] C. M. Sparrow. "On spectroscopic resolving power". *Astrophys. J.* 44 (1916), pp. 76–86 (cited on p. 154).

[195] A. J. den Dekker and A. van den Bos. "Resolution: a survey". *J. Opt. Soc. Am. A*, 14(3) (1997), pp. 547–557 (cited on p. 154).

[196] G. D. Bruce, M. Y. H. Johnson, E. Cormack, D. a. W. Richards, J. Mayoh, and D. Cassettari. "Feedback-enhanced algorithm for aberration correction of holographic atom traps". *J. Phys. B*, 48(11) (2015), p. 115303 (cited on p. 155).

[197] J. Liang, J. Kohn, M. F. Becker, and D. J. Heinzen. “High-precision laser beam shaping using a binary-amplitude spatial light modulator”. *Appl. Opt.* 49(8) (2010), pp. 1323–1330 (cited on p. 156).

[198] R. Dorn, S. Quabis, and G. Leuchs. “The focus of light - linear polarization breaks the rotational symmetry of the focal spot”. *J. Mod. Opt.* 50(12) (2003), pp. 1917–1926 (cited on p. 156).

[199] C. Pethick and H. Smith. *Bose-Einstein condensation in dilute gases*. 2nd. Cambridge: Cambridge University Press, 2008 (cited on p. 183).