Lisa Marie Starkey née Bennie

Precise engineering of the Bose–Einstein condensate wavefunction using magnetic resonance control

Submitted in total fulfilment of the requirements
of the degree of Doctor of Philosophy

Supervisory committee:
Dr. Lincoln Turner
Dr. Russell Anderson
Prof. Kristian Helmerson

School of Physics & Astronomy
Monash University

July, 2016

Copyright Notices

Abstract

To *engineer the wavefunction* of a Bose–Einstein condensate is to exert control over both the density and phase of the Bose–Einstein condensate order parameter. Having the ability to engineer the condensate order parameter down to the smallest length scale relevant to condensate dynamics—the healing length scale—would enable the study of new combinations of topological defects and may pave the way to using Bose–Einstein condensates as versatile, precise quantum simulators.

In this thesis, we present a new wavefunction engineering technique which reaches the sub-optical healing length scale. Influenced by magnetic resonance imaging, we name this technique *magnetic resonance control*. This technique uses time-varying coupling between internal spin states of a spinor Bose–Einstein condensate within a magnetic field gradient to address spatial regions of the condensate, enabling control over both the phase and the density of the condensate order parameter down to the healing length scale.

Techniques already exist to engineer condensate wavefunctions, but not with such a fine degree of control. These techniques primarily rely on either the intensity variation of a laser beam, which limits the resolution to the diffraction limit (larger than the typically sub-optical healing length scale), or the adiabatic inversion of a magnetic trapping potential, which can not be easily changed to produce a variety of structures in the condensate wavefunction.

To develop our magnetic resonance control technique, we simulate a spinor condensate in one dimension with time-dependent coupling between spin states and time-dependent external magnetic field gradients using the Gross–Pitaevskii equation. We show that magnetic resonance control can engineer a single black soliton using experimentally feasible parameters. A black soliton is an ideal target state to select for this demonstration because engineering such a state requires control over both the phase and density of the condensate with healing-length resolution. We demonstrate that magnetic resonance control can be extended to engineer more complicated target states by simulating the creation of multiple solitons in a condensate, with control over the initial positions and trajectories of the solitons.

When magnetic resonance control is applied to Bose–Einstein condensates in the laboratory, it will be necessary to have an imaging system capable of resolving

the fine structures created. As an alternative to high-cost, custom-manufactured lenses, and in-vacuum optical systems, I have designed and bench-tested an objective lens with a high numerical aperture (0.36) and a long working distance (35 mm) consisting of standard catalogue lenses. Using 780 nm light, suitable for imaging rubidium condensates, this objective can achieve a resolution of 1.3 μm across a diffraction-limited field of view of 360 μm through a 5 mm thick glass window of a science cell. By changing the spacing between the lens elements, this objective lens can compensate for the aberrations produced by a glass window up to 15 mm thick, and by changing the aperture size the objective becomes suitable for diffraction-limited monochromatic imaging on the D2 line of all the alkalis.

Before performing proof-of-principle magnetic resonance control experiments on real Bose–Einstein condensates, we needed to construct an experimental apparatus capable of producing spinor Bose–Einstein condensates. In this thesis I summarise my main contributions to this group endeavour, including: constructing the ultra-high vacuum system; supervising the bakeout of our vacuum system; designing and aligning the optical systems to produce laser beams of different, tunable frequencies to trap, cool, and image rubidium gas; trapping a cloud of rubidium atoms in a magneto-optical trap; constructing laser beam shutters, photodetectors, and photodetector signal filters; and designing and constructing our side imaging and top imaging systems.

Using our spinor Bose–Einstein condensate apparatus, I performed the first proof-of-principle magnetic resonance control experiments. With a pair of condensates side-by-side, separated by 200 μm, we can use magnetic resonance control to invert the spin of one condensate only, while leaving the other condensate unaffected.

We anticipate magnetic resonance control being used in the laboratory to engineer the first black soliton in a Bose–Einstein condensate. Looking beyond solitons, magnetic resonance control could have applications to the field of magnon spintronics, and extending the technique to higher dimensions could enable the study of exotic topological defects such as spin knots in a quantum fluid.

Declaration

This thesis contains no material that has been accepted for the award of any other degree or diploma in any university or other institution. To the best of my knowledge the thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis. For parts of this thesis that are based on joint research or publications, the relative contributions of the respective authors are detailed appropriately.

Lisa Starkey

Dr Lincoln Turner

Dr Russell Anderson

Professor Kristian Helmerson

Acknowledgements

I have been fortunate to be a member of the Physics and Astronomy family at Monash University, and I would like to take this opportunity to thank them all for supporting me through my PhD. In particular, I would like to thank the following people below, who have played an integral role in my PhD journey.

Lincoln Turner first sparked my fascination with Bose–Einstein condensates during the first colloquium he presented at the School. His passion and excitement for research inspired me to do my PhD project with him, and without Lincoln this thesis would not exist.

Russell Anderson, our Mathematica magician, spent countless hours in paper-torture editing sessions with me, teaching me how to write with precision. Russ knows everything there is to know about our Bose–Einstein condensate 'rattus', and if not for him the laboratory might fall apart.

Kris Helmerson, my associate supervisor, has always been available for a chat about my project, carrying a repository of useful references around in his memory.

Tapio Simula, the group's cheerful theorist, has provided me with alternative ways of thinking about challenging problems.

Michael Morgan, our Head of School, proved a loyal supporter of our research group as the inevitable delays and problems arose. He provided financial support when needed, and always encouraged us to publish in Nature.

David Paganin and Alexis Bishop have both been postgraduate coordinators for the School during my PhD, forming an excellent base of support, and advocating for any extensions needed.

Jean Pettigrew is the amazing mother of the School, tirelessly battling bureaucracy and always making sure we are all taken care of.

Considering that our research group started with an empty laboratory, we were so fortunate to have an electronic and mechanical workshop staffed with such talented people. In particular, Nino Benci and Brett Williams both patiently worked with me to translate my vague requirements into specific, achievable designs.

Martijn Jasperse and Alex Wood were my comrades-in-arms in the spinor lab, and together we achieved milestones which sometimes seemed impossible on the

journey towards condensation. I want to thank them both, for making our time together in the lab a memorable experience.

Our research group is lucky to have not one, but two Bose–Einstein condensate machines. Thank you to the members of the dual-species laboratory, Phil Starkey, Shaun Johnstone, and Chris Billington, for all your contributions to the group, especially not minding (too much) when we stole all your Thorlabs components. In particular, I want to thank Phil and Chris for developing our awesome laboratory control system software, which allowed me to record all my experimental data ridiculously efficiently.

Over the years our research group has grown, with new members joining us while others left for greener pastures. I would like to thank all past and current members of the group, for creating a vibrant research community. Beyond our research group, I would like to thank all the other PhD students in the School, who celebrated when things went well, commiserated when they didn't, and made the tea room a welcoming respite from the laboratory.

Thank you to my parents, who always believe I can do the impossible.

Finally, I would like to thank Phil (again), for being my partner both in work and life. Thank you for everything, always and forever.

Contents

Chapter 1

Introduction

Welcome to my thesis! I hope you find the contents interesting and useful for your own research endeavours. I have designed each chapter to be largely self-contained, making it easier to skip to areas you are interested in; unless, of course, you are assessing this thesis, in which case I hope you enjoy the cover-to-cover read. In this chapter I place this thesis in the broader context of quantum simulation, explain why we are interested in devising better ways to engineer the Bose–Einstein condensate (BEC) wavefunction, and review the current techniques for engineering solitons in Bose–Einstein condensates. I summarise the main scientific contributions of this thesis, discuss the evolution of my research project, and provide an outline of the structure of this thesis.

1.1 *Engineering the Bose–Einstein condensate wavefunction for quantum simulation*

An outstanding challenge in the physical sciences is the simulation of quantum mechanical systems. During the surging popularity of computers in the early 1980s people realised that simulating quantum mechanics with classical computers was going to be a very difficult problem. In his 1982 lecture, *'Simulating Quantum Mechanics'*, Richard Feynman discussed the situation as follows:

> "And therefore, the problem is, how can we simulate the quantum mechanics? There are two ways we can go about it. We can give up on our rule about what the computer was, we can say: Let the computer itself be built out of quantum mechanical elements which obey quantum mechanical laws. Or we can turn the other way and say: Let the computer still be the same kind that we thought of before—a logical, universal automaton; can we imitate this situation?" [1]

These two options have since formed two active fields of research; the former being *quantum simulation*, the latter *quantum computation*.

While much progress has been made in the field of quantum computation [2–4], we are still far from the realisation of a quantum computer with sufficient qubits

to simulate a condensed matter system of moderate size. A different approach is to use a quantum simulator [5] instead of a quantum computer. While a quantum computer uses bits and gates to simulate a quantum system, a quantum simulator uses one complex quantum system to mimic another completed quantum system without the need for programmable logic [6–11].

For a quantum system to be a suitable quantum simulator, we need to be able to do the following three things [12]:

1. Set up a mapping between the Hamiltonian of the system we seek to simulate and the Hamiltonian of the quantum simulator.
2. Prepare the initial quantum state of the system for the start of the simulation with high fidelity.
3. Make high-precision measurements of the quantum simulator.

Systems proposed for quantum simulators include neutral atoms in optical lattices, trapped ions, polar molecules, cavity arrays, quantum dots, superconducting circuits, photons, and nuclear spins. A summary of the potential of these systems for quantum simulation, and experimental realisations, can be found in Reference 5. Examples of experimental implementations include simulating the Bardeen–Cooper–Schrieffer (BCS) model with liquid state nuclear magnetic resonance [13], observing the BCS-BEC crossover in Fermi gases [14,15] and simulating the Dirac equation using a single trapped ion [10].

For this thesis, we are particularly interested in the potential of Bose–Einstein condensates as quantum simulators. Several proposals have been made to use Bose–Einstein condensates to simulate aspects of cosmology [16–18], and experimentally Bose–Einstein condensates have been used to study Hubbard models and quantum phase transitions [6], and black hole analogues [19]; do Bose–Einstein condensates have the potential to simulate other systems?

For the purposes of most experiments, a Bose–Einstein condensate is well described by the mean-field Gross–Pitaevskii model, in which a multi-component condensate is viewed as a macroscopically occupied coherent quantum state, which is then described by an order parameter in which a density and a phase is associated with each spatial location within the condensate, for each internal component of the condensate [20]. In this thesis, we follow the convention in our field and refer to such an order parameter as a *macroscopic wavefunction*. Being macroscopic wavefunctions, Bose–Einstein condensates provide an isolated and highly controllable platform which may prove ideal for quantum simulation of a variety of physical systems [21].

At the start of my research project, we asked ourselves: can the Bose–Einstein condensate wavefunction be controlled rapidly, with arbitrarily high precision, to create the desired initial conditions for quantum simulation of a variety of physical systems? This is a very ambitious question to say the least, encompassing an entire sub-field of investigation rather than a single doctoral project! Consequently, we reduced the scope of this question to the following:

> *Can the Bose–Einstein condensate wavefunction be controlled rapidly, with resolution at the smallest relevant length scale of the condensate, to engineer a stable black soliton in the condensate?*

One PhD project later, we can state that the answer to this research question is *yes* in theory, and therefore *probably* for the purposes of experimental implementation.

We note that in practice it is unlikely that a *perfect* black soliton can be engineered. Even aside from experimental imperfections which would cause deviations from the black soliton mode (for example, drift of the dipole trap position relative to the magnetic gradient coils), the rate of quantum depletion of the condensed fraction into the non-condensed fraction is higher in a condensate with a black soliton than one without; this quantum depletion perturbs the soliton so it is no longer stationary [22]. Regardless, a black soliton is an appropriate target state for testing a new wavefunction engineering technique, for the following reasons. First, a soliton is a one-dimensional object. This allowed us to develop our proof-of-principle technique in only one dimension, which is conceptually easier to consider, before considering how to extend the technique to higher dimensions. Second, solitons are the smallest one-dimensional structure that can exist with stability in a quantum fluid. Engineering a stationary soliton in a Bose–Einstein condensate demonstrates the ability to control the wavefunction down to the smallest relevant length scale. Third, to engineer a black (stationary) condensate rather than simply a grey (moving) soliton requires control over both the density *and* the phase of the condensate.

Techniques exist to engineer wavefunctions beyond the context of solitons and vortices; for example, engineering the spin state of atoms in optical lattices [23], or spatially controlling the condensate scattering length to form spin interfaces [24]. For the remainder of this chapter, we focus our discussion to existing engineering techniques which aim to engineer topological defects in superfluids.

1.2 Dark solitons in Bose–Einstein condensates

Non-linear systems support solitary waves, or *solitons*: excitations stable against perturbations that propagate without changing shape. The field of soliton research emerged in 1834 when Scottish naval architect John Scott Russell observed a shallow water wave created by the sudden halt of a boat in a narrow channel propagating without dispersion [25]. Since then, solitons have been observed in other non-linear systems including optical fields [26] and, of particular interest for this thesis, Bose–Einstein condensates [27, 28].

Bose–Einstein condensates support either *bright* solitons or *dark* solitons, depending on the sign of the s-wave scattering length a_s. When $a_s < 0$ the particle self-interaction is attractive and bright solitons can form. These bright solitons are localised density packets that propagate without dispersion, reminiscent of solitons formed in water, and were first engineered at Rice University in 2002 [29]. When $a_s > 0$ the particle self-interaction is repulsive and dark solitons can form; these are the type considered in this thesis. Dark solitons are stable notches in

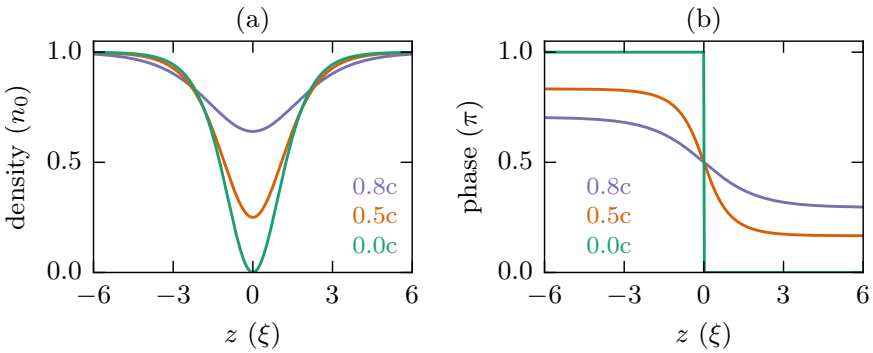

Figure 1.1: The (a) density and (b) phase of a soliton with different velocities v_{sol} in a homogeneous condensate, where c is the speed of sound in the condensate. These curves were made using Equation 1.3.

the condensate density, with an associated maximum phase shift of π, formed by the balance of two opposing tendencies: the repulsive self-interaction seeking to shrink the density notch, and the chemical potential limiting the sharpness of density gradients [30, p. 86]. Dark solitons in condensates were first engineered in Hannover [27] and the National Institute of Standards and Technology (NIST) [28] in 1999 using a phase-imprinting technique discussed in Section 1.3.1.

In Bose–Einstein condensate literature, the term 'dark soliton' is used in two contexts. The first is to refer to all solitons formed in a repulsive condensate, in contrast to bright solitons (as used in the previous paragraph). The second is as one of three categories of solitons formed in a repulsive condensate, qualitatively determined by the depth (or 'darkness') of the soliton: a *grey* soliton has a shallow density dip with a broad, shallow phase gradient, a *dark* soliton has a deeper density dip with a steeper phase gradient, while a *black* soliton has a density dip that reaches zero with a π phase discontinuity. This is demonstrated in Figure 1.1. The fact that the phase only exhibits a discontinuity when $v_{\text{sol}} = 0$ leads to the following observation: although solitons are sometimes referred to as topological defects [31], they do not have true topological protection since they can be smoothly deformed away. Some believe solitons should be classed as pseudo-topological defects [32].

The *healing length* ξ of the condensate,

$$\xi = \frac{1}{\sqrt{8\,\pi\,n_0\,a_s}} = \frac{\hbar}{\sqrt{2\,m\,\mu_{\text{cp}}}}\,, \tag{1.1}$$

determines the width of the soliton in the condensate. Here n_0 is the peak condensate density, a_s is the s-wave scattering length, m is the atomic mass and μ_{cp} is the chemical potential;

$$\mu_{\text{cp}} = 4\pi\,\hbar^2\,\frac{n_0\,a_s}{m}\,. \tag{1.2}$$

The healing length is the shortest distance over which density perturbations exist in a stationary state of a condensate; consequently, it is the smallest length scale relevant to wavefunction engineering. To give an example, rubidium condensates typically have a peak density of 10^{14} atoms/cm^3 and an s-wave scattering length of

5.3 nm. For such a condensate the healing length is 270 nm, which is smaller than the visible transitions in rubidium: typically 780 nm, with a more technically challenging transition at 420 nm. This sub-optical healing length scale precludes the use of optical potentials alone for wavefunction engineering at the healing length scale; this issue is discussed further in Section 1.3. While lasers do exist which produce such short wavelengths, using them is an impractical solution because producing a diffraction-limited spot size requires a very high numerical aperture lens, which extremely limits the optical access around the condensate.

The wavefunction for a soliton in a one-dimensional homogeneous condensate is

$$\psi_{\mathrm{sol}}(z,\,t) = \sqrt{n_0}\left(i\,\frac{v_{\mathrm{sol}}}{c} + \sqrt{1-\left(\frac{v_{\mathrm{sol}}}{c}\right)^2}\,\tanh\left(\frac{5\,(z-v_{\mathrm{sol}}\,t)}{2\sqrt{2}\,w_{\mathrm{sol},\,v}}\right)\right) e^{-i\,\frac{\mu_{\mathrm{cp}}\,t}{\hbar}}\,, \tag{1.3}$$

where v_{sol} is the velocity of the soliton, $c = \sqrt{\mu_{\mathrm{cp}}/m}$ is the speed of sound in the condensate[1], and $w_{\mathrm{sol},\,v}$ is the velocity-dependent width of the soliton,

$$w_{\mathrm{sol},\,v} = \frac{5\,\xi}{2\sqrt{1-\left(\frac{v_{\mathrm{sol}}}{c}\right)^2}}\,. \tag{1.4}$$

[1] An excellent derivation of the soliton wavefunction can be found in "Bose–Einstein Condensation of Dilute Gases" by C. J. Pethick and H. Smith [33, p. 216–218].

The density at the base of a soliton is

$$n_{\mathrm{min}} = n_0\left(\frac{v_{\mathrm{sol}}}{c}\right)^2\,. \tag{1.5}$$

Equation 1.5 shows that the soliton depth depends on its velocity; if $v_{\mathrm{sol}} = 0$, we form a black soliton with width $w_{\mathrm{sol}} = 5\,\xi/2$. For the case of grey or dark solitons in a trapped condensate, the soliton oscillates within the condensate as though it were a particle of mass $2\,m$ [33, 34], turning at the point when the density dip reaches zero; consequently, grey solitons have a larger amplitude of oscillation than dark solitons.

While dark (moving) solitons have been engineered in condensates before, a stable, single black soliton has to date not been experimentally realised in a condensate. The black soliton is an ideal target state for the demonstration of any new wavefunction engineering protocol, because it requires both density control and phase control at the healing length scale. If the engineering protocol forms a density notch wider than the healing length, the notch will decay into multiple grey solitons (Figure 1.2(a)). Even if the density notch is engineered to the correct width, the protocol must also engineer the requisite π phase discontinuity for the soliton to be stable (Figure 1.2(b)).

1.3 Engineering solitons in Bose–Einstein condensates

1.3.1 Phase-imprinting

The first dark solitons in condensates were engineered using a technique called *phase-imprinting*, where an off-resonant laser beam is masked to produce a 'sharp' intensity profile, then imaged onto the condensate [27, 28], as shown in Figure 1.3.

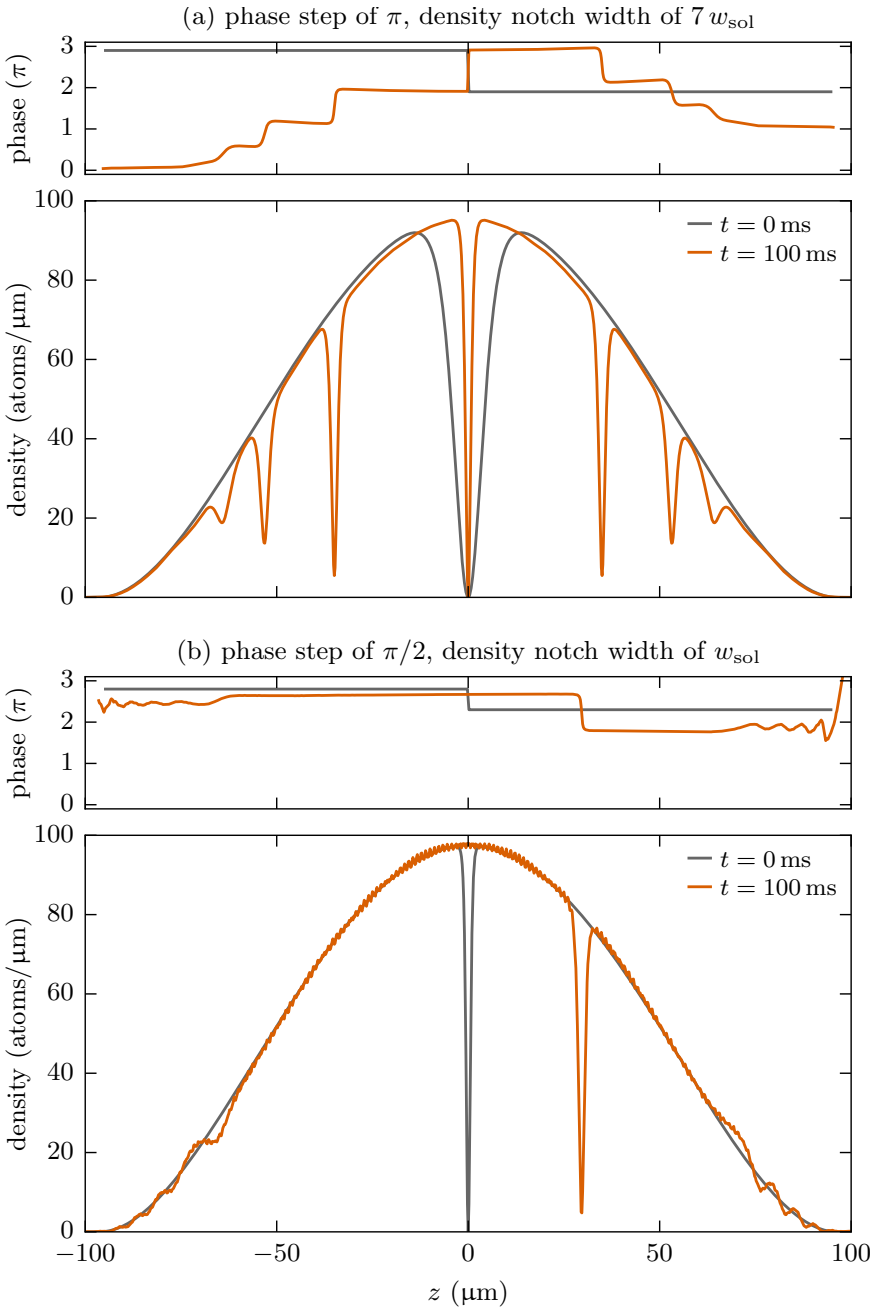

Figure 1.2: Examples of poor soliton engineering. (a) If the density notch is wider than the healing length (in this example, seven times wider) the feature is unstable and decays into multiple grey solitons, even if the initial phase step is π. (b) If the density notch is initially the correct width (w_{sol}) but the phase step is not equal to π, the feature will decay into one grey soliton by shedding sound waves. This data is from simulations of the Gross–Pitaevskii equation, discussed in Chapter 4. This simulated condensate consists of 10^4 rubidium atoms (^{87}Rb) in a harmonic trap with an axial frequency of 2.4 Hz and a radial frequency of 158.4 Hz, but the general principles illustrated apply to all condensates.

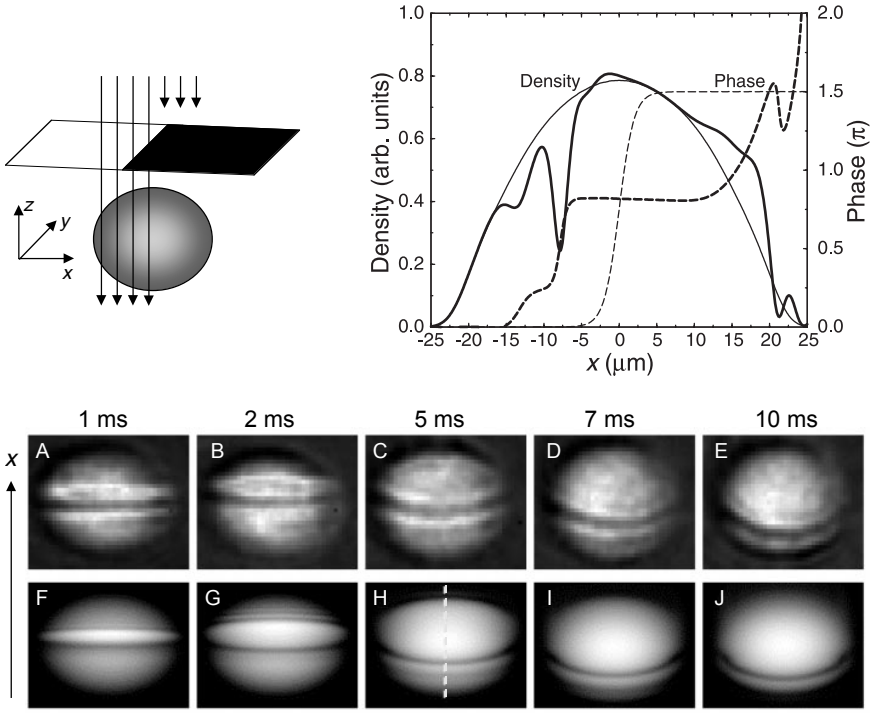

Figure 1.3: From J. Denschlag et. al. 2000 [28]. Reprinted with permission from AAAS. A schematic of the phase imprinting technique (upper left) shows a mask being used to shine light on half the condensate, to imprint a π phase step. Since this technique does not also engineer the condensate density, the engineered feature evolves into a grey soliton by shedding sound waves (upper right, lineout from 'H'). Experimental data (A–E) shows the grey soliton moving downwards through the condensate, while the higher density sound waves move upwards. This behaviour is also apparent in the simulated data (F–J).

The beam induces a light shift (ac Stark shift), causing the phase of the condensate to accumulate faster in the exposed region. If the condensate is exposed to this beam for an appropriate length of time a phase shift of π is produced. But phase-imprinting is fundamentally limited by diffraction at the edge of the beam; the resulting phase profile is not a discontinuity, instead the phase changes across a distance greater than the (sub-optical) healing length.[2] This means that, at best, this technique can only produce moving solitons. A further problem is that phase-imprinting does not allow direct control over the condensate density. The induced phase gradient causes density to be ejected from the soliton site in the form of either sound waves or supersonic shock waves [28], further perturbing the condensate from the target state.

[2] Phase-imprinting was first proposed by Dobrek et al. [35] as a way of producing a vortex in a condensate, using a circular phase mask. However this was never successfully implemented due to the deleterious diffraction effects at the 2 π boundary of the phase mask (K. Helmerson, private communication).

1.3.2 Matter-wave interference

Dark solitons can be engineered using a matter-wave interference technique to control the density profile of the condensate. In this technique, first proposed by Scott et al. in 1998 [36], two condensates initially held in a double-well potential expand into each other with low collisional energy as the potential is slowly removed.

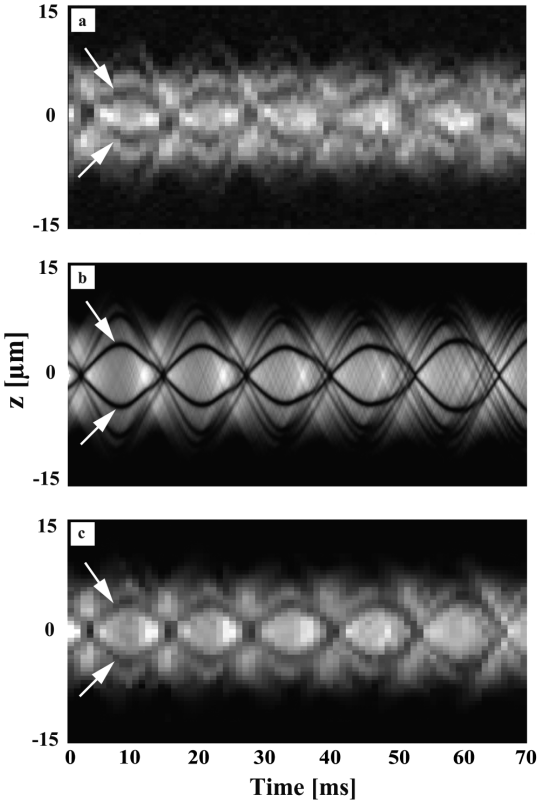

Figure 1.4: Reprinted with permission from A. Weller et. al. 2008 [39]. Copyright 2016 by the American Physical Society. Two condensates are transferred from the double-well potential of an optical lattice to a single harmonic potential. This is an example of matter-wave interference being used to engineer solitons. As the condensates interfere with low collisional energy an array of grey solitons form (the two deepest solitons are marked with arrows). Experimental data is shown in (a), with simulations shown in (b) and (c); the latter has been convolved with the resolution of the experimental data. Since only the density of the condensate has been directly controlled, not both the density and phase, this technique can only form even numbers of solitons.

Normally the collision of two condensates would result in a cosine standing-wave interference pattern, as first observed by W. Ketterle in 1996 [37]. However, if the collisional energy is less than the repulsive interaction energy [38] the atomic interactions modify the shape of the interference pattern into an array of grey solitons [36]. Viewed another way, as the collisional energy decreases the wavelength of the interference pattern increases. If the wavelength of the interference pattern equals the width of a soliton, a soliton array is formed [38]. In experiments, the requisite double well potential has been formed by either an optical lattice (Figure 1.4) [39] or a blue-detuned laser which forms a repulsive barrier in a single condensate [38].

Matter-wave interference alone does not add an extra π phase step across the condensate, which means that it will form only *even* numbers of solitons moving symmetrically apart (such that the array of phase steps across the condensate sum

to zero). But matter-wave interference combined with phase-imprinting can produce odd-numbered soliton arrays. This technique was proposed by Carr et al. in 2001 [40], and realised experimentally by Theocrasis et al. in 2010 [41]. In principle this combination of phase and density engineering could be used to make a single black soliton, but at best has matched the achievements of phase-imprinting by making a single grey soliton [41]. This is because matter-wave interference shares the same fundamental limitation as phase-imprinting: it is diffraction-limited and as such cannot engineer sub-optical features.

1.4 Other wavefunction engineering techniques

This literature review would not be complete without discussing wavefunction engineering techniques that aim to form structures other than solitons.

1.4.1 Laguerre–Gaussian beams

Laser beams with Laguerre–Gaussian spatial modes carry quantised angular momentum, which can be coupled into two internal states of a Bose–Einstein condensate to form a quantised vortex in one internal state [42–44]. Intricate ring lattices suitable for trapping cold quantum gases can be formed by interfering two Laguerre–Gaussian beams [45]. These techniques have a similar limitation to phase imprinting, in that they are optical techniques limited in resolution by diffraction effects. It is not apparent that these techniques could be used to form non-radially symmetric density and phase structures.

1.4.2 Topological phase imprinting

By adiabatically inverting a magnetic trapping potential, a relative geometric phase shift can form across a condensate [46]. This technique has been used to engineer a vortex [47, 48] and a two-dimensional skyrmion [49, 50] in Bose–Einstein condensates. It is possible that this technique could be used to form a spin knot in a condensate [51], or a vortex with very high vorticity [52–54]. While it is apparent that this is a powerful technique, it is not very flexible because it is reliant on the configuration of the magnetic trapping potential on the experimental apparatus, so it would not be easy to create a situation where a variety of structures could be formed in succession using the same condensate machine.

1.5 What's new in this thesis

The primary scientific result in this thesis is the development of a new wavefunction engineering technique, which we name *magnetic resonance control* (MRC). The optical diffraction limit does not apply to MRC because it is a non-optical technique; instead the resolution of MRC is provided by a linear magnetic field gradient that induces a spatially-varying level splitting between two internal states. In the presence of this gradient a strong time-dependent coupling between the states is

used to control both the phase and the density of the quantum fluid, more rapidly than the fastest system dynamics of the uncontrolled fluid.

In this thesis our exemplar quantum fluid is a spinor Bose–Einstein condensate with three Zeeman sublevels coupled by radiofrequency-driven magnetic dipole transitions. More generally, MRC is applicable to any spatially-extended quantum system with internal states as long as the splitting between these states can be made spatially-dependent and the states can be coupled in a time-dependent way. Suitable quantum systems include Fermi gases [55], atoms in optical lattices [23,56,57], and ^{3}He films [58].

I demonstrate in simulations that magnetic resonance control can engineer a black soliton in a Bose–Einstein condensate using experimentally realistic parameters, and perform the first proof-of-principle experiments using this technique in our laboratory. Our vision is that magnetic resonance control will one day enable Bose–Einstein condensates to be used as quantum simulators for a variety of quantum systems, and enable the controlled creation and study of exotic topological defects.

An instrumentation result presented in this thesis that the quantum gas community may find interesting is the high-resolution objective lens I have designed and bench-tested. This objective lens should enable diffraction-limited imaging of topological defects in a Bose–Einstein condensate through the thick glass wall of a vacuum cell using optical elements available in standard lens catalogues. Capable of achieving resolution of around one micrometer with a wide field of view, for the imaging wavelengths of all the alkalis, it provides a versatile, low-cost option for high-resolution imaging.

1.6 Chronology of this research project

We planned for my research project to be primarily experimental in nature. We started a collaboration with theorist Joseph Hope at the Australian National University, with the intent that his group would provide the theoretical expertise and we would provide the experimental expertise. As such, for the first two years of my project I devoted my time primarily to the construction of our spinor Bose–Einstein condensate apparatus, a monumental task shared between the other members of my research group. As collaborations are wont to do, ours evolved into a joint theoretical collaboration, with both parties performing theoretical investigations for the remainder of my research project. This evolution is reflected in the structure of this thesis: first focusing on building a spinor Bose–Einstein apparatus and associated hardware, then presenting my own extensive simulations developing our magnetic resonance control protocol, and finally exploring in the laboratory what interesting things can be achieved by applying magnetic resonance control to our condensates within the capabilities of our apparatus.

1.7 Outline of this thesis

Chapter 2 presents my contributions to building our spinor Bose–Einstein condensate apparatus: constructing the ultra-high vacuum system; baking the vacuum system to achieve the vacuum pressure required for long-lived Bose–Einstein condensates; installing optics and electronics needed to trap and cool atoms; designing our 'workhorse' side-imaging system and creating our low-resolution top-imaging system.

Chapter 3 presents the design, construction and performance testing of our high-resolution objective lens. This work was published in Optics Express [59]. I then discuss the design of a complete imaging system to be used with the objective lens.

In Chapter 4 I move away from the laboratory and delve into simulations, developing our new technique for engineering condensate wavefunctions. I simulate a spin-1 condensate in one dimension and demonstrate how magnetic resonance control can engineer a stable black soliton in the condensate. I extend this work by simulating the creation of multiple solitons, with control over their positions and initial trajectories. Sections of this chapter form the basis of our manuscript published on the ArXiv [60], currently in peer-review.

Chapter 5 explores the range of condensate parameters in which our magnetic resonance control technique can successfully engineer a stable black soliton with experimentally realistic parameters. This work will be valuable for implementing the protocol in laboratories.

In Chapter 6 I apply magnetic resonance control to the Bose–Einstein condensates in our laboratory, displaying individual control over the spin projection of two spatially separated condensates.

Chapter 7 concludes this thesis by speculating on the future directions of magnetic resonance control.

Chapter 2

Building a spinor Bose–Einstein condensate machine

When I joined the spinor Bose–Einstein condensate research group at Monash University in 2009, the laboratory contained two empty optical benches and dreams for the future. Transforming the space into a functional research laboratory was a monumental task and a team effort from the whole research group. In this chapter I present the aspects of the spinor Bose–Einstein Condensate apparatus that I was primarily involved with. Further information about other aspects of our apparatus is presented in the PhD theses of my colleagues Alex Wood [61] and Martijn Jasperse [62].

2.1 Overview of the apparatus

Our spinor Bose–Einstein machine follows the design of J. Porto and I. Spielman at NIST [63]. Construction of the vacuum system is discussed in Section 2.2, with the process to obtain ultra-high vacuum in Section 2.3. An atomic beam is formed using a collimated effusive oven, and slowed using a single-layer, variable pitch Zeeman slower [64]. This atomic beam loads a 3D magneto-optical trap (MOT) in the science cell, an optically flat glass cell chosen to maximise optical access, described further in Section 2.5. A photograph of the apparatus, taken in June 2014, is shown in Figure 2.1.

I was heavily involved with constructing the apparatus up to and including loading atoms into the MOT, and with building the associated electronics (Section 2.6) and optics required (Section 2.4 and Section 2.7). I then developed the high resolution imaging system presented in Chapter 3 while Alex Wood and Russell Anderson continued down the temperature scale to make BEC. Full details of how we create BECs are in Alex's thesis [61]. Briefly, we use a molasses cooling stage and magnetic trapping stage, then load into a hybrid magnetic and optical dipole trap, before finally reaching condensation in a purely optical cross-beam dipole trap. The absence of magnetic trapping in the final stage creates a *spinor* Bose–Einstein condensate, in which there remains a spin degree of freedom. This

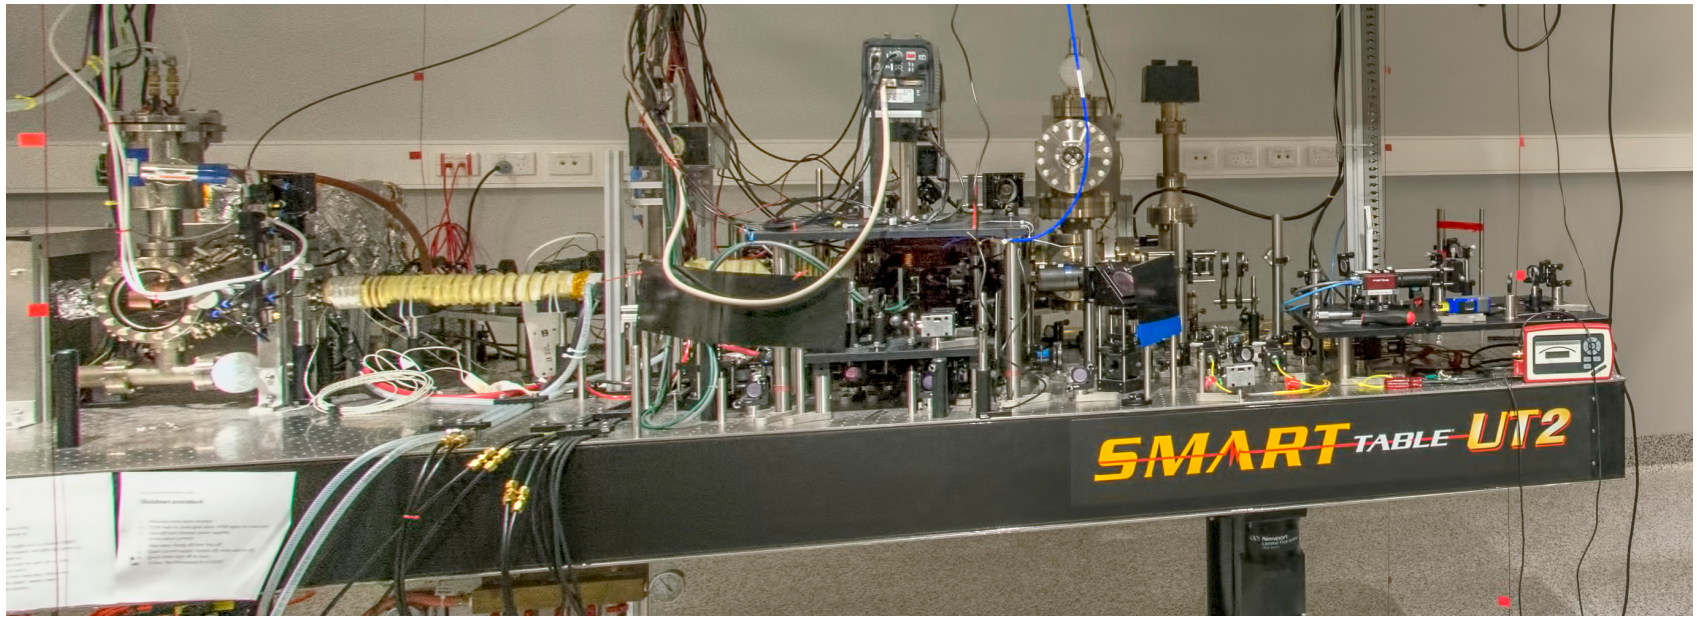

Figure 2.1: Over the course of five years, with one relocation to a new building, the empty vacuum table evolved into this spinor Bose–Einstein condensate machine. This was a result of the efforts of myself, Alex Wood, Martijn Jasperse, Russell Anderson and Lincoln Turner, with help from our colleagues Philip Starkey, Shaun Johnstone, Chris Billington and Kristian Helmerson. Not shown in this photograph is the second table enveloped by optical components and lasers, and the abundance of electronic components, interlock systems, and control software we built to make and perform experiments on BECs.

is essential, not only for my work on magnetic resonance control, but also for my colleagues' work on studying the effects of magnetic gradients on spinor condensates [61] and on Faraday imaging of condensates [62].

2.2 Constructing the vacuum system

At the start of my PhD project in March 2011, the vacuum system was still in two pieces. One section contained the rubidium oven, a 75 L/s ion pump, atomic beam collimator, and Zeeman slower,[1] and was constructed by Alex Wood as part of his Honours project in 2010 [66]. This vacuum section is shown in Figure 2.2. During my Honours project in 2010 [67] I constructed the ultra-high-vacuum (UHV) section, which houses a 75 L/s ion pump (Gamma Vacuum 75S), an UHV-grade ion gauge, an Extorr residual gas analyser (RGA) and the titanium sublimation pump required to maintain ultra-high vacuum pressure of less than 10^{-11} Torr. To these vacuum components I attached a temporary (not optically-flat) glass cell fed by a rubidium getter (Alvatec AS-3-Rb-10-V). This vacuum section is shown in Figure 2.3. The reason for this separate construction was that our optically flat glass cell, custom-made by Hellma Optik GmbH Jena, suffered extended delays. While waiting for it to arrive we wanted to characterize the Zeeman slower and practice making a magneto-optical trap (MOT).

[1] I constructed the Zeeman slower in a 3rd year research project in 2010 [65].

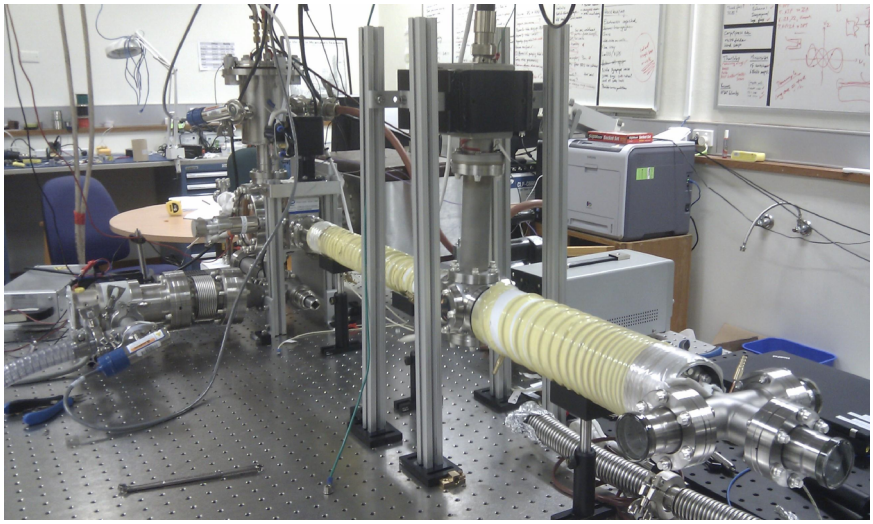

Figure 2.2: The state of the apparatus at the start of my research project. This is one half of the vacuum system, containing the rubidium oven, ion pump, atomic beam collimator, and Zeeman slower. Notice the junction of viewports in the lower right corner, which was later removed to join the two haves together. The other half of the vacuum system is shown in Figure 2.3.

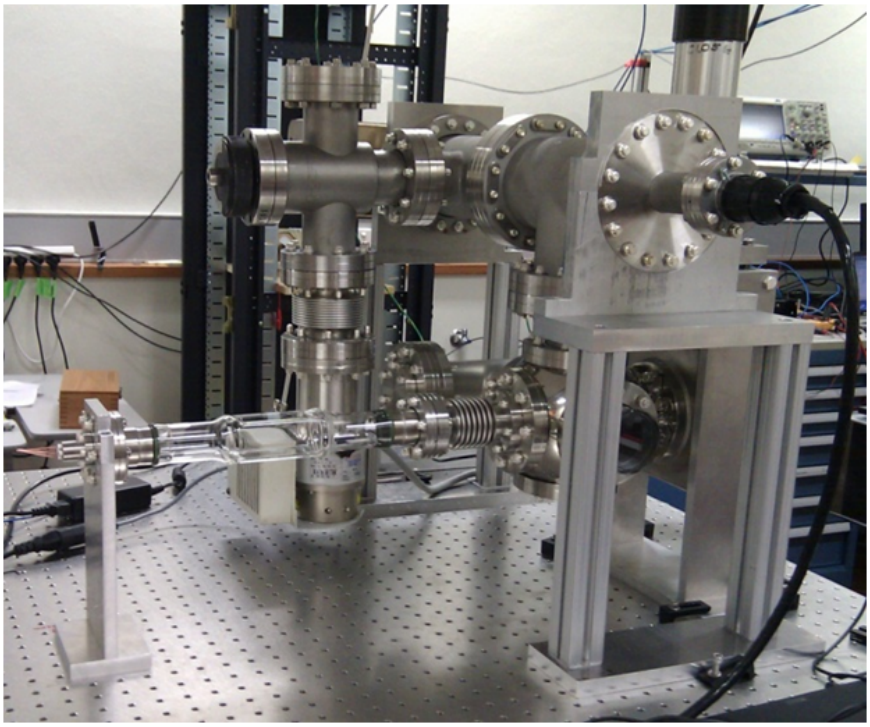

Figure 2.3: The state of the apparatus at the start of my research project. This is the UHV system containing the ion pump, ion gauge, residual gas analyser (RGA) and titanium sublimation pump. This photograph shows the temporary glass cell, fed by a getter, used to make a 'practice' magneto-optical trap (MOT) during my Honours project as we awaited the arrival of our optically-flat glass cell. Once the cell arrived, we mated this half of the vacuum system to the other half (shown in Figure 2.2).

2.2.1 *Installing the glass cell in our vacuum system*

In December 2010 we had a MOT in the temporary glass cell, and by August 2011 the Zeeman slower had been extensively characterized by Alex Wood [61]. By this time the custom glass cell had finally arrived, enabling the completion of our vacuum system. The mating of the two sections was the result of a heroic day's work, where we broke vacuum in the UHV section, removed the temporary glass cell, pushed the UHV section into position next to the Zeeman slower, and installed the glass cell. Some action shots of the UHV migration are shown in Figure 2.4.

Attaching the glass cell to the vacuum system was not easy. We first attached the glass cell to the Zeeman slower vacuum flange via a flexible metal bellows.[2] We then carefully positioned the very heavy (around 100 kg) UHV system by hand to sit less than 1 mm from the open end of the glass cell with an accuracy of better than 100 μm, such that once the glass cell was bolted in place the cell remained axial with the atomic beam. I monitored the strain on the glass cell while the final bolts were tightened, to make sure we didn't break the glass cell by over-tightening the bolts, by viewing the glass cell through a pair of crossed-polarisers and making sure that there was no significant change to the strain pattern as each bolt was

[2] Positioning a copper gasket in mid-air for a horizontal vacuum join is an art in itself. Kudos to Alex Wood for perfecting this technique.

Figure 2.4: Positioning the UHV system, in preparation for installing the glass cell, by Alex Wood and Russell Anderson. The School's talented workshop technician Brett Williams was on hand to provide professional advice. Shown here are photographs automatically taken by a webcam as they (a) were getting ready to push the UHV chamber towards the Zeeman slower, (b) were halfway through the migration, and (c) after the UHV chamber was carefully positioned, leaving the precise gap needed for the glass cell.

Figure 2.5: Lincoln Turner had the honour of tightening the final bolt to join the glass cell to the vacuum system, while I monitored the strain in the glass cell using crossed-polarisers.

gradually tightened. A photograph of this event is shown in Figure 2.5.

2.3 *Bakeout*

In order for a BEC to have long trap lifetime of several seconds it must exist within ultra-high vacuum, where the pressure is on the order of 10^{-11} Torr. Vacuum pumps cannot reach this pressure until the water monolayer coating the inside of the steel walls, and the hydrogen embedded in the steel, is removed. To do this we performed a *bakeout*, where the UHV chamber was heated to an average temperature of $250\,^{\circ}\mathrm{C}$ for several days.

Although our bakeout ultimately had a successful conclusion, it was an extended lesson in how *not* to do a bakeout. What should have taken two weeks to complete took close to two months due to a succession of vacuum leaks. Here are the two main lessons we learned:

- *Support the vacuum system at a minimum number of points*, and if possible don't bolt the system down to the table, so that it is free to expand and contract during the bake. Our leaks were caused by our vacuum system straining against supports and opening joins while it was heating.
- *It is preferable to over-tighten, rather than under-tighten, bolts on vacuum flanges.* We initially followed the torque guidelines recommended by Reference 68. When our vacuum system kept springing leaks, we tightened the bolts further, with no ill-effects.

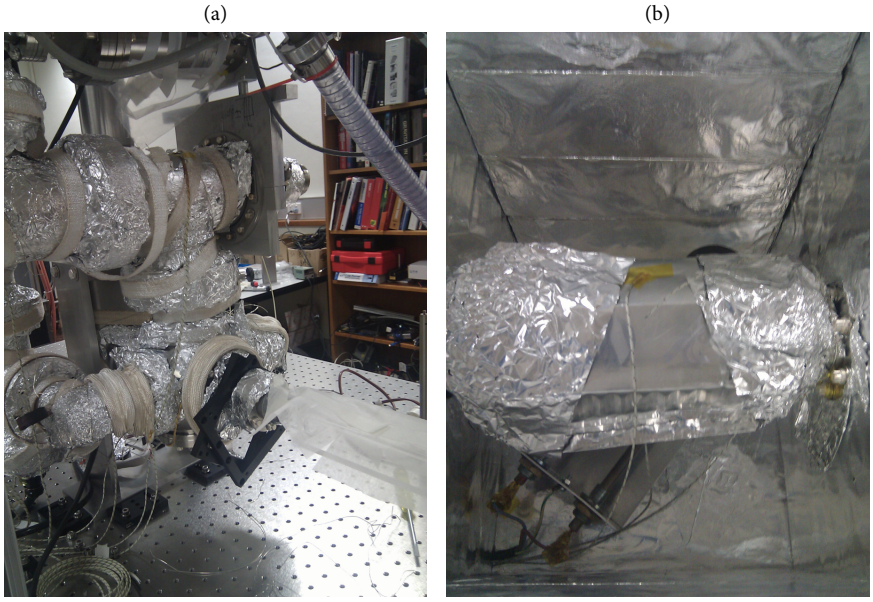

Figure 2.6: Preparing for bakeout. (a) The UHV chamber wrapped in heater tapes, with a tight base layer of aluminium foil for increased conduction, before the addition of a loose top layer of aluminium foil to trap air for insulation. The glass cell is shown here wrapped in lens tissue for protection. Thermocouples for temperature monitoring can be seen leaving the vacuum system in the lower left corner. The black rectangles near the glass cell are some of the bias coil supports. (b) The view inside the oven constructed around the glass cell using aluminium foil-wrapped fire bricks. The glass cell was protected from baking residues by an aluminium sheet. At the base of the oven is the heating element used to heat the cell via convection rather than conduction, to ensure even heating.

At first our bakeout went relatively to plan. I wrapped the UHV system in fibreglass heater tape and insulated it with several layers of aluminium foil (Figure 2.6(a)). We constructed an oven around the glass cell made of fire bricks and ceramic fibre board, so the cell was heated by convection from a heating filament instead of heater tape (Figure 2.6(b)). This prevented the glass cell developing hot spots, which could have otherwise broken the cell.

The heating elements were manually controlled using Variacs (variable autotransformers), with thermocouples monitoring temperatures at 16 locations.[3] To prevent leaks occurring during the bakeout, the vacuum system had to be heated slowly (at a rate of less than 20 °C per hour) and evenly (notably, keeping the temperature difference across any glass-to-metal seals smaller than 10 °C). To achieve this, I monitored the reports from the thermocouple reader using a LabVIEW program (developed by Martijn Jasperse) and manually adjusted the Variacs accordingly. Example screenshots of this software are shown in Figure 2.7 and Figure 2.8. So that we could sleep (mildly) at ease, the Variacs were connected to an interlock system that would notify us through emails and text messages if any temperatures deviated beyond set ranges. I highly recommend using an interlock system during bakeout; further details of the interlock system used is in Alex Wood's thesis [61].

[3] Limited by the 16 channels of our SRS SR630 thermocouple reader.

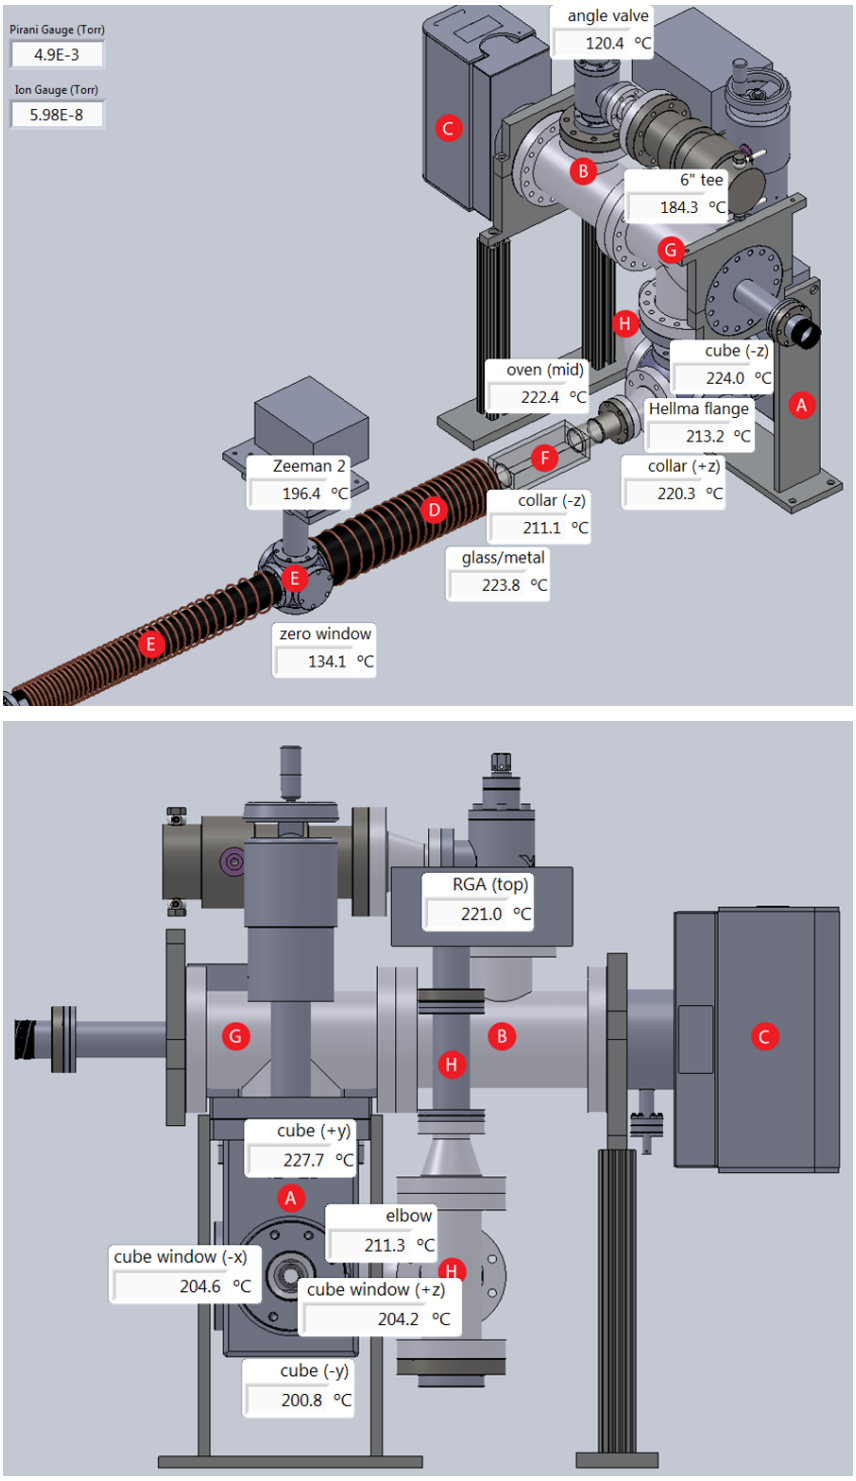

Figure 2.7: Screenshots of part of the temperature monitoring software used during the bake-out. Red letters indicated which Variacs controlled the temperature of different regions of the vacuum system. Temperatures were shown in real time, reported from the thermocouple monitor.

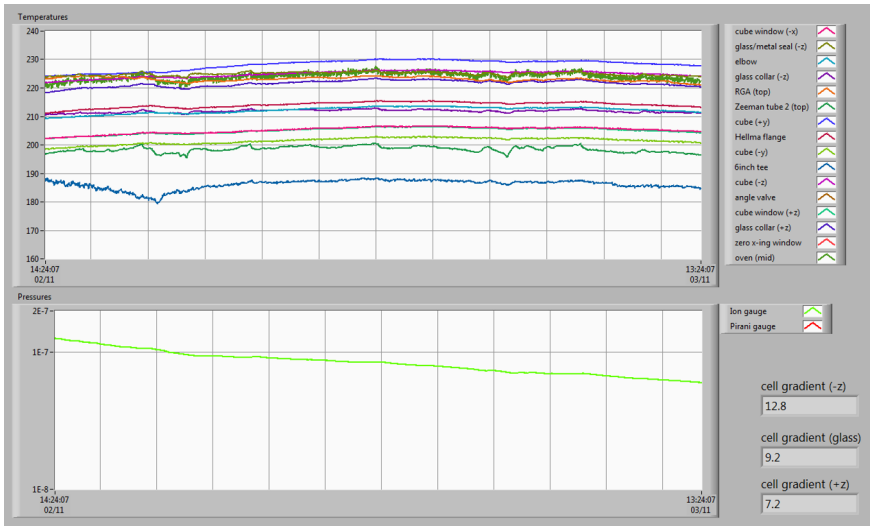

Figure 2.8: A screenshot of part of the temperature monitoring software use during the bakeout. This showed real-time graphs of the temperature at different regions of the vacuum system, allowing us to become 'human PID controllers' by adjusting the Variacs to ensure even, slow heating of the vacuum system. In this screenshot, the temperatures have stabilised. This software also displayed the pressure in the vacuum system, which on good days slowly decreased (like shown here) and on bad days suddenly spiked as leaks opened up in the vacuum system.

Our first significant bakeout misfortune occurred while we were cooling the vacuum system back to room temperature after baking at an average temperature of 250 °C for six days. When the average temperature was 200 °C, with the bulk metal of the UHV end still above 250 °C, a leak suddenly appeared which raised the pressure from 10^{-7} Torr to 10^{-4} Torr (the pressure stabilised here due to the efforts of the roughing pump). We initially thought this leak was due to either cooling too rapidly, or not enough tension in the bolts around the leak point. However after several smaller leaks appeared over the next several weeks[4] we realised the problem lay in the vacuum support structures. The UHV system was held by three supports, which seemed a sensible choice given the weight of these components. Unfortunately, because the supports were necessarily at different heights, they expanded by different amounts during heating; as the system heated and expanded, the supports stressed the core of the vacuum system, opening up joins. To solve this problem we removed the least-essential third support and removed the clamps fixing the supports to the table so that the system could expand more freely.[5] After doing so, we were finally able to bake the vacuum system for four consecutive days, and cool back to room temperature without any leaks occurring, thus concluding two months of bakeouts. A picture of the vacuum system post-bakeout is shown in Figure 2.9.

The post-bake pressure in the vacuum system was 10^{-10} Torr. To reduce this pressure by a final order of magnitude, we fired the titanium sublimation pump filaments to coat the interior surfaces of the vacuum system with titanium. We

[4] Yes, our bakeout was a pretty trying experience. If you are about to embark on a bakeout, I hope your experience is better.

[5] By removing the third support we had to decide to remove the gate valve at the end of the atomic beamline, originally included to allow future expansion of the vacuum system. Hopefully this decision will not be cursed by future members of the research group.

happily achieved a final pressure of 4×10^{-11} Torr. This value approaches the floor of the ion gauge, which has an x-ray limit of 2×10^{-11} Torr. This pressure is low enough to enable BEC trap lifetimes of several minutes. Our pressure may have been limited in the end to the effectiveness of our titanium sublimation pump, which was exposed to atmosphere while hot during the first vacuum leak. As a result, we believe that it did not coat the steel surfaces with as much titanium as it should have, to absorb the remaining hydrogen outgassing from the steel.

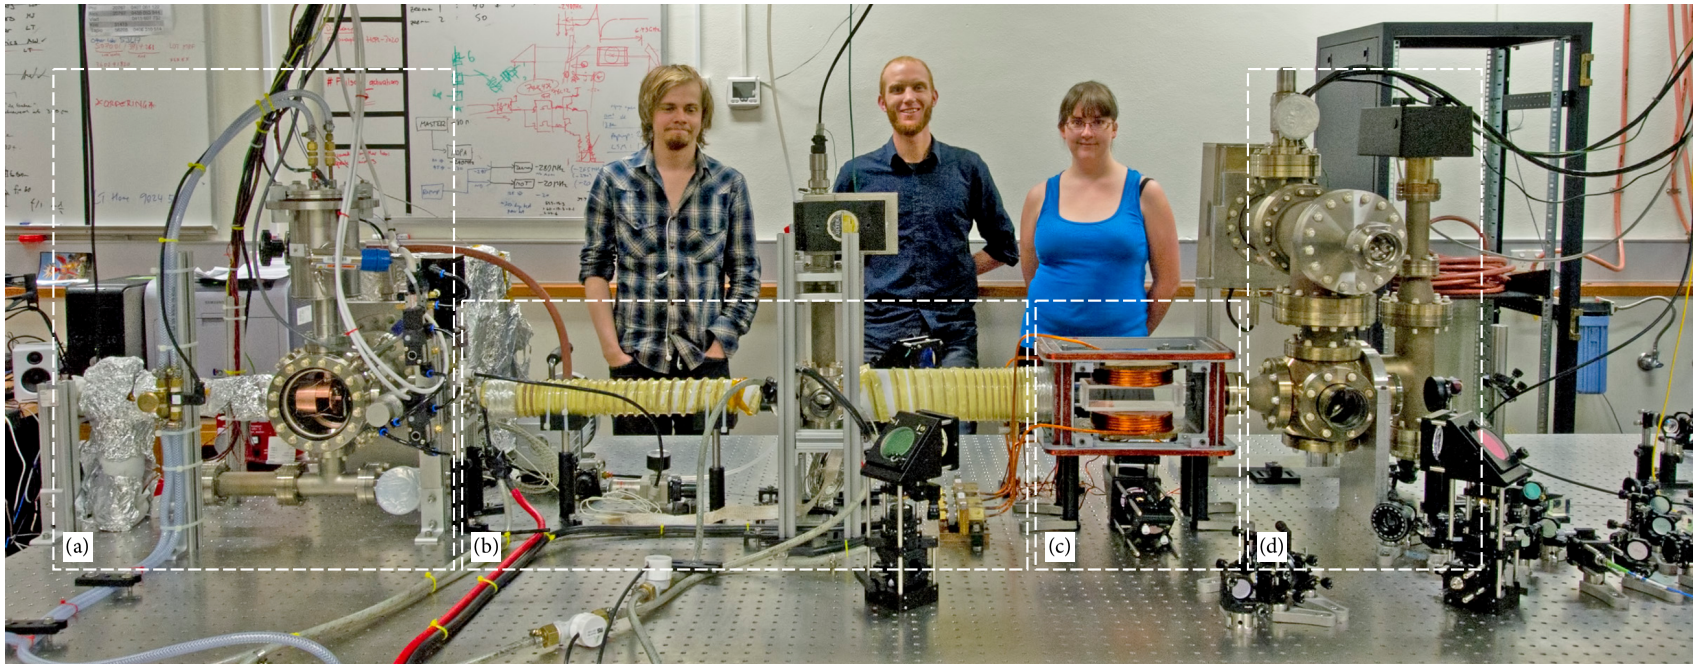

Figure 2.9: The completed vacuum system after bakeout, showing (a) the rubidium oven and collimated atomic beam source, (b) the Zeeman slower, (c) the glass cell science chamber surrounded by our magnetic quadrupole coils and bias field coils, and (d) the UHV end housing the ion pump, ion gauge, RGA and titanium sublimation pump. This photograph was taken in November 2011, before any optics around the glass cell were installed. For comparison, in the photograph of the apparatus shown in Figure 2.1, optomechanics now completely engulf the glass cell. I am standing in the background with Alex Wood and Russell Anderson, after spending the day carefully installing the magnetic coils around the glass cell.

2.4 Optics after the tapered amplifier

Trapping atoms in the MOT requires laser beams of four different frequencies, around the cooling and repump transitions. For ^{87}Rb, the cooling and repump transitions are between $F = 2 \rightarrow F' = 3$ and $F = 1 \rightarrow F' = 2$ respectively. The MOT trapping light is detuned 16 MHz below the cooling transition, and requires the ability to sweep the frequency 60 MHz below resonance during the molasses stage. The MOT repump light is resonant with the repump transition. The Zeeman slower uses slowing light detuned 160 MHz below the cooling transition, with repump light detuned 160 MHz below the repump transition. Imaging requires yet another frequency; the imaging beam is resonant with the cooling transition.

To produce these beams we designed an optical layout for the laser table that uses only three lasers, which are all MOGlabs external cavity diode lasers (ECDLs). The master laser provides the input for a tapered amplifier (TA-0780-1500, m2k-Laser GmbH) housed in the master oscillator power amplifier (MOPA), and is locked to a single frequency using a modulation transfer spectroscopy (MTS) lock developed by my colleague Vlad Negnevitsky [69]. The MOPA produces up to 800 mW of light, which is used to form the MOT trapping beams, Zeeman slower beam and imaging beam. The Zeeman repump laser provides the repump light for the Zeeman slower, and is locked referenced to the master laser with a beat-note lock [70]. The MOT repump laser provides the repump light for the MOT, and is locked using a current-modulated saturated absorption spectroscopy lock. Schematics of the three laser locks are shown in Figure 2.10. Ideally a single laser would provide all the repump light needed, but our repump lasers have both suffered from poor performance, producing less than 50 mW each.[6]

[6] This issue is actually being remedied by my colleagues as I write this thesis, demonstrating how quickly apparatus chapters can become dated.

The lasers and MOPA are located on one optical bench (the 'laser table') while the vacuum system is located on a second optical bench (the 'vacuum table'). This is to minimise electrical equipment and stray fields near the science cell. We use polarisation-maintaining (PM) optical fibres to transfer light from the laser table to the vacuum table. I was responsible for designing and building the optics after the MOPA which split the amplified light into three paths, or *beamlines*[7], for the MOT trapping beams, Zeeman slower beam and imaging beam, shifting the frequency of each beam as required with acousto-optic modulators (AOMs). A schematic of the beamline optics is shown in Figure 2.11

[7] We were perhaps influenced in this naming by the Australian Synchrotron across the road from our laboratory.

The beamlines experienced several iterations as our requirements changed. For instance, our original intention was use a six-way fibre splitter to form the six MOT beams, which required the MOT repump light and the MOT trapping light to be coupled into the same fibre. This is very difficult and wasted weeks of time, avoid it if you can. We abandoned this approach for other reasons, discussed in Section 2.5.

AOMs change the pointing direction of the output beam, depending on the frequency shift required. If we do not need to change the frequency of a particular beamline during an experiment, or from one experiment to the next, this is not a problem because the static pointing direction does not reduce the efficiency of the

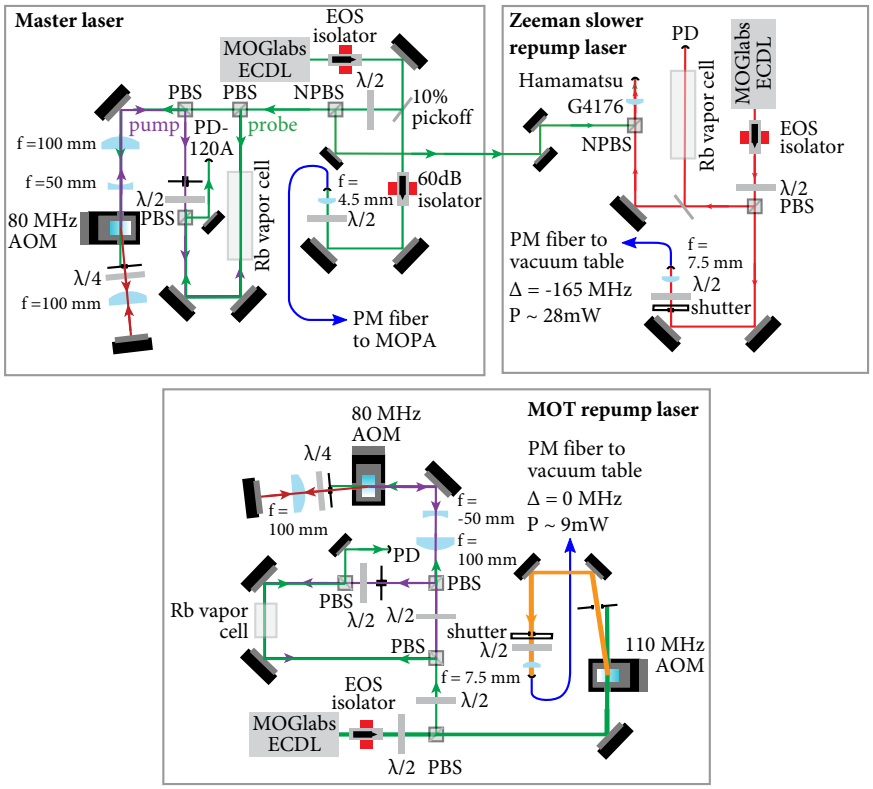

Figure 2.10: Layout of the locks for the master laser, Zeeman slower repump laser and MOT repump laser. All three lasers are MOGlabs external cavity diode lasers. The master laser uses a modulation transfer spectroscopy lock. The Zeeman repump laser is referenced to the master laser using a beatnote lock. The MOT repump laser uses a standard current-modulated saturated absorption spectroscopy lock. If the Zeeman repump laser could produce 100 mW of light we would retire the MOT repump laser, which has an unreliable lock, and produce all our repump light with the one laser. Image courtesy of Alex Wood, modified with permission [61].

fibre couple. But in some cases we need to sweep the frequency over a large range during an experiment, like the MOT beamline during the molasses stage. This would ordinarily drastically reduce the fibre coupling efficiency. To combat this, we set up these AOMs in a double-pass configuration [71], shown in Figure 2.12. A perfectly aligned double-passed AOM does not allow the pointing direction of the beam to change as the frequency changes. In practice, some deviation in pointing direction still occurs but it is greatly reduced. For example, as the frequency of the AOM in the MOT beamline is swept, the fibre coupling efficiency maintains a FWHM of 85 MHz.

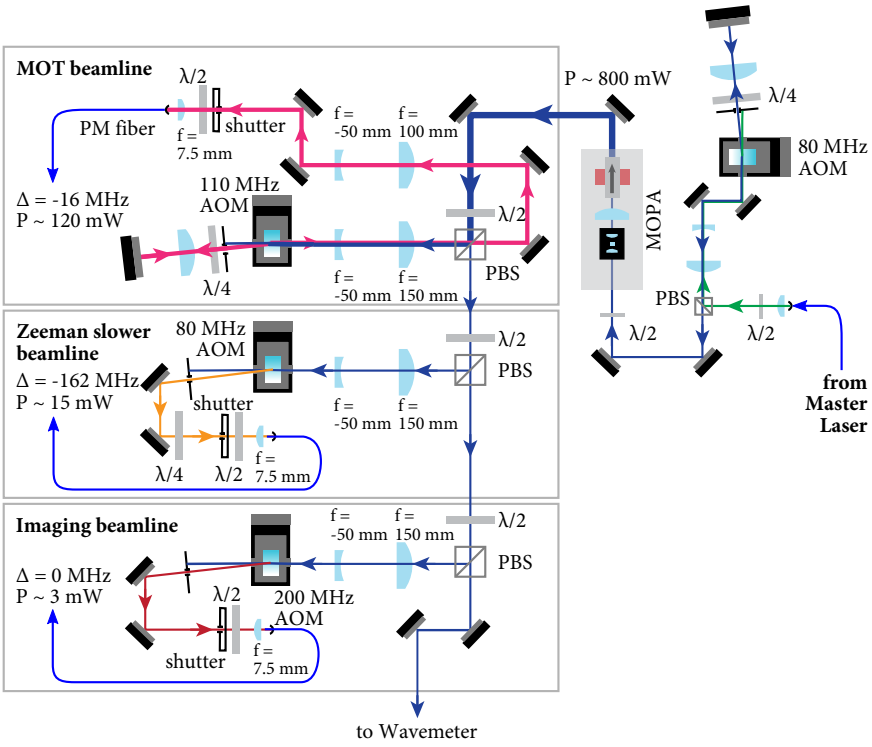

Figure 2.11: Layout of the optics on the laser table, using light fibre-coupled from the master laser. The master oscillator power amplifier (MOPA) uses a tapered amplifier to create 800 mW of light, which is split into the MOT beamline, Zeeman slower beamline, and imaging beamline. AOMs change each beamline to the required frequency (detunings shown are relative to the cooling transition). The AOM in the MOT beamline uses a double-pass configuration to enable frequency sweeps without degradation of the fibre couple efficiency (see Figure 2.12 for more detail). The MOT beamline is closest to the MOPA output to minimise its propagation distance, as we require it to have the greatest power stability. Image credit to Alex Wood, modified with permission [61].

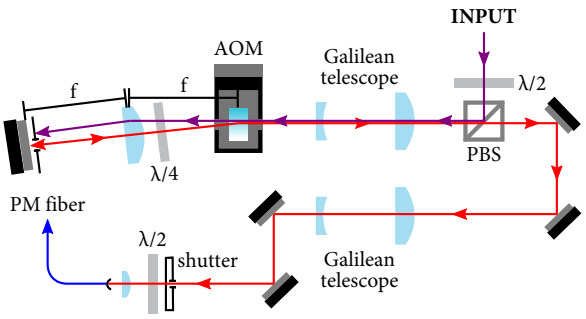

Figure 2.12: Typical layout of a double-passed AOM and fibre couple. In our lab the input beam has a large diameter of around 4 mm to allow a long propagation distance after the MOPA. We use a telescope to shrink the beam before the AOM for maximum diffraction efficiency. To minimise changes in the pointing direction of the double-passed beam as the AOM frequency changes, the lens after the AOM is positioned one focal length from both the AOM and the mirror.

2.5 Trapping atoms in the magneto-optical trap

With an operational vacuum system and light for the MOT and Zeeman slower piped over to the vacuum table, we could then build the optics on the vacuum table to start trapping atoms in the MOT. While Alex Wood took charge of the Zeeman slower optics, I created the six MOT beams.

The first stage was to combine the MOT cooling light and repump light, and then split this beam into six. To do this, I designed a system using minimal optical components where the power balance of each pair of counter-propagating beams is independently adjustable, making the task of balancing the powers in each MOT beam relatively straight forward. This layout is shown in Figure 2.13. These optical components were mounted onto fixed-height 1 inch diameter posts, ranging between 17.5 mm and 25 mm tall, to set the beam height at only 45 mm above the vacuum table. This was the height required for direct input into the MOT beam expanders.

The MOT cooling and repump light is collimated after the optical fibre using a Thorlabs F810APC-780 collimator to form a $1/e^2$ beam diameter of 7.5 mm. After splitting, we expand the six beams to a diameter of 16 mm using beam expanders I designed and constructed during my Honours project [67]. Within each beam expander (not shown in Figure 2.13) the MOT beam is circularly polarised using a quarter-wave plate.

Our original plan when designing these beam expanders was that the MOT cooling and repump light would be combined into one fibre, and then this fibre would be split into six fibres which could then be attached directly to each beam expander; removing the need for a lot of optics on the vacuum table. However, Philip Starkey determined during his Honours project that it is very difficult to build your own fibre splitter that both splits the power evenly and maintains polarisation stability [72], and unfortunately a commercial fibre splitter was beyond our budget.

Once the MOT beams were aligned and the powers in each beam were balanced, and the Zeeman slower beam was aligned and slowing atoms, we were able to start trapping atoms. A photograph of the MOT when it is well aligned is shown in Figure 2.14.

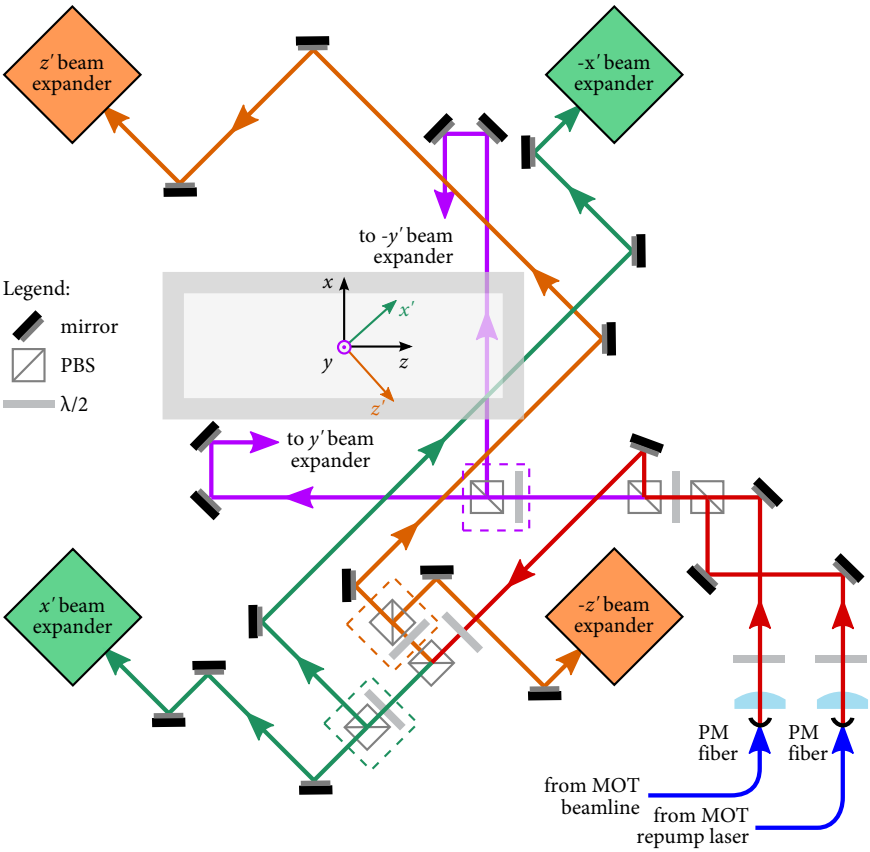

Figure 2.13: Layout of the optics around the glass cell used to form the six MOT beams. For simplicity, only optics on the ground floor layer of the vacuum table are shown; the six MOT beam expanders are not included in full detail. This diagram shows that the power balance in each pair of counter-propagating beams (shown in a different colour) can be independently adjusted using the half-wave plates and polarising beam splitter (PBS) cubes indicated by a dashed outline. To maintain polarisation purity, the design enables all but one mirror to be aligned at 45 degrees to the beam path.

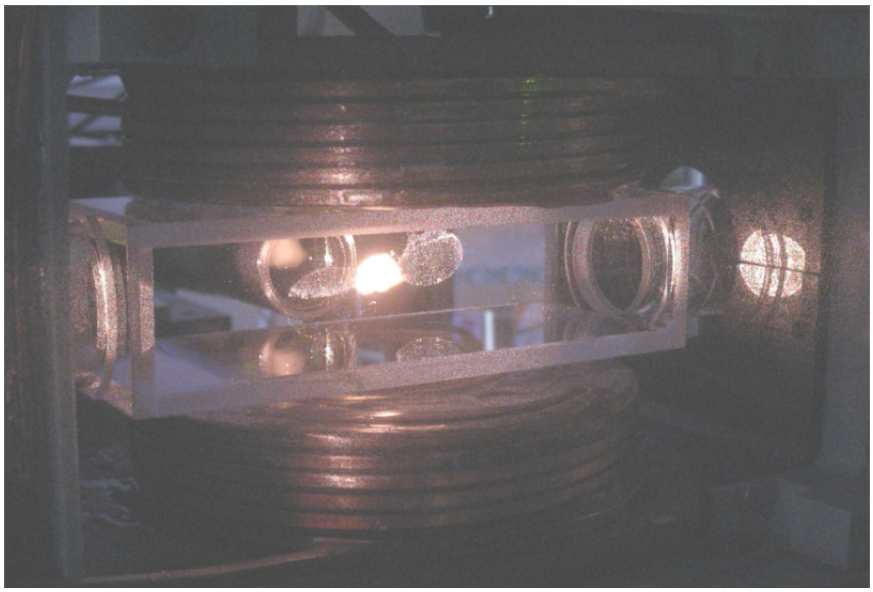

Figure 2.14: A photograph of atoms trapped in the MOT, taken with a camera with sensitivity extending into the near IR. For scale, the glass cell is 40 mm tall.

2.6 Laser shutters, photodetectors and signal filters

While building our spinor BEC apparatus, the entire research group constructed many of the electronic devices needed to make and perform experiments on BECs. As my contribution to this effort, I built several laser shutters out of recycled hard disc drives (HDDs), and designed and constructed photodetectors and signal filters for our laser beam power monitoring system.

2.6.1 Laser shutters

During an experimental run we need to automatically block different laser beams. We use acousto-optic modulators (AOMs) not only to change the frequency of beams, but also to rapidly switch beams on and off (with typical switching times $< 1\,\mu s$). But sometimes during an experiment we require beams to be switched off for several seconds. If we switched the AOM off for this long, it would cool down, changing the diffraction efficiency (and hence the beam power) during an experiment. To avoid this problem, we first block the beam rapidly using an AOM, then use a (comparatively slower) mechanical laser shutter to block the beam, letting us switch the AOM back on before it cools.

On some beamlines we use commercial laser shutters (SRS SR475); for others I constructed custom shutters made from the voice-coil actuators (or read-write arm) from recycled HDDs [73], and assembled electronic drivers for each using printed circuit boards purchased from MOGlabs (SD-002). Figure 2.15 shows an assembled shutter and driver, while Figure 2.16 shows the driver circuit diagram.

The main feature of the driver is the 1 mF capacitor that discharges into the voice-coil of the HDD read-write arm, moving the arm at 17 m/s (the exact speed being dependent on the model of HDD used), which blocks a 1 mm $1/e^2$ diameter beam in 100 µs. The HDD shutter has three main advantages over commercial shutters: it is faster than the SR475 shutter, which moves at only 10 m/s, blocking the same beam in 190 µs (Figure 2.17); it is much cheaper, at a quarter of the cost of the commercial shutter; and it has a more compact design which requires less free beam propagation distance. The disadvantage of our HDD shutters is that they are less robust than commercial shutters. In our laboratory we use the HDD shutters where speed is essential, access is tight, or wherever our budget requires them.

2.6.2 Photodetectors

I designed and constructed several photodetectors with the intent of installing a laser power monitoring system at locations throughout our optical system, to monitor the stability of alignment into fibre couples and AOMs. In practice we decided that a complete laser power monitoring system was not necessary; instead we have used them for other applications. We use them to continuously monitor the 'seed' input beam to the tapered amplifier as part of an interlock system that turns off the tapered amplifier if the seed power drops below a threshold value, to prevent

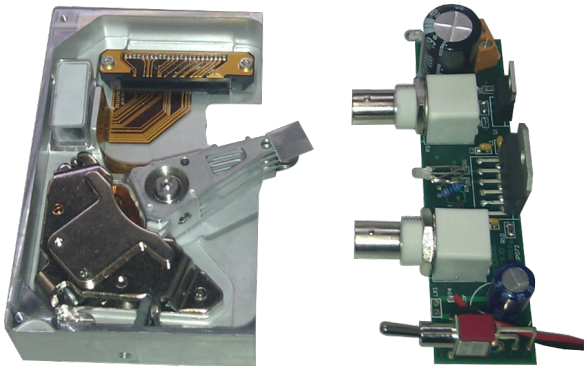

Figure 2.15: A photograph of the voice-coil actuator cut from a HDD, with an aluminium flag attached used to block the laser beam, and the circuit used to control the shutter (either with a manual switch or a transistor-transistor logic (TTL) signal).

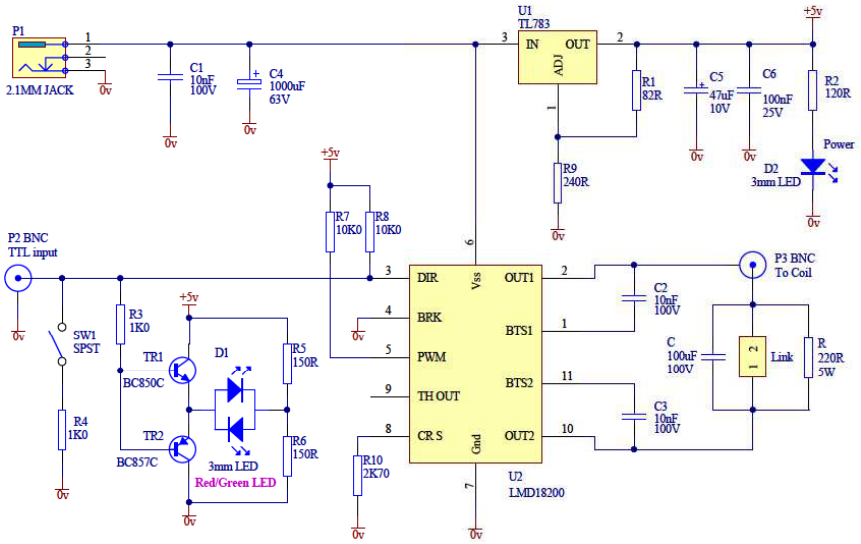

Figure 2.16: Circuit diagram for the HDD shutter driver, from the MOGlabs SD-002 manual.

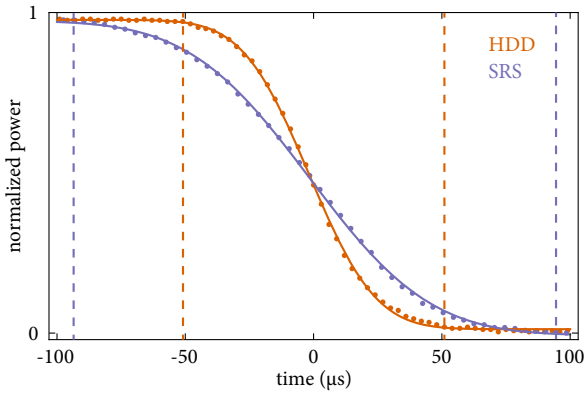

Figure 2.17: The HDD shutter blocks a beam with $1/e^2$ diameter of 1 mm in 100 μs, almost twice as fast as the commercial SRS shutter which blocks the same beam in 190 μs.

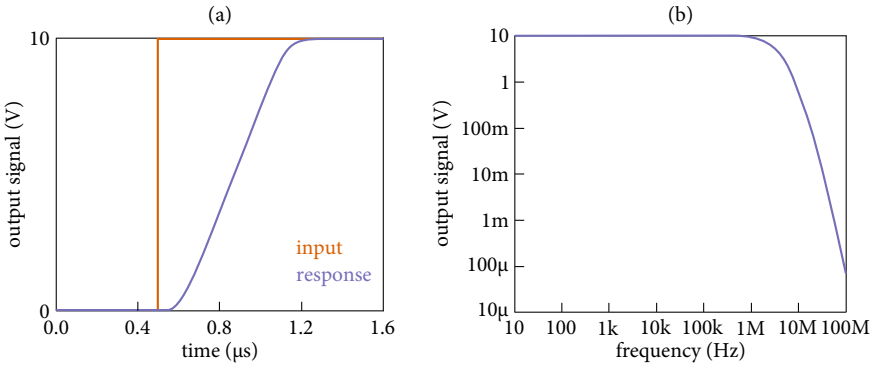

Figure 2.18: SIMetrix simulation of the photodetector circuit. (a) Response to a step function input signal showing a rise time of 0.8 μs. (b) Frequency response curve showing a −3 dB point, or bandwidth, of 5 MHz.

damage to the tapered amplifier. We use them to measure delay times for the HDD shutters, which are dependent on beam alignment.

The photodetectors convert a 0–500 μW input signal incident on a silicon photodiode to a 0–10 V output signal. To be sensitive to weak pick-off beams, they are shot noise limited at input powers above 60 μW. To achieve this performance I optimised resistor and capacitor values by simulating the circuit using SIMetrix, and selected two op-amps with stringent requirements on voltage offsets, voltage and current noise densities, bias currents and gain-bandwidth products [74, ch. 18]. SIMetrix simulations indicate the photodetector has a rise time (10% − 90%) of 400 ns and a bandwidth of 5 MHz (Figure 2.18). This is faster than the AOM switching time, which means they can resolve the salient features of an experimental run. The photodetector circuit diagram is shown in Figure 2.19. A photograph of a constructed photodetector is shown in Figure 2.20

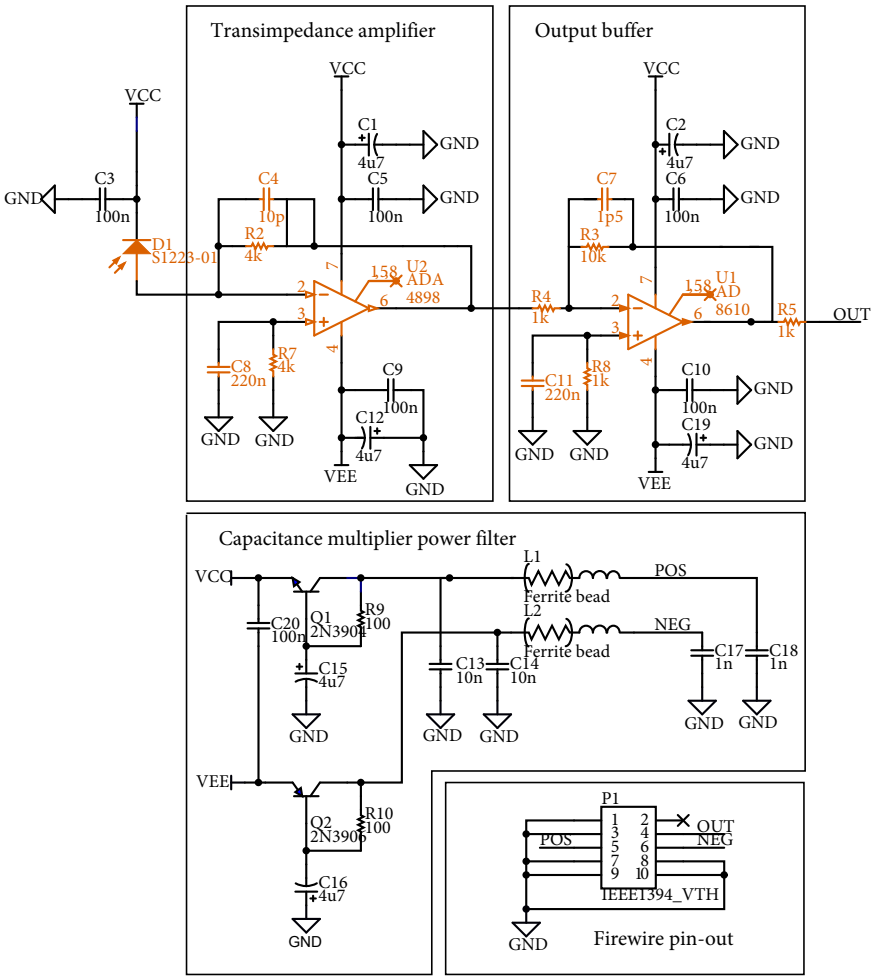

Figure 2.19: Circuit diagram of the photodetectors which convert an input light signal of 0–500 µW (D1) to an output signal of 0–10 V (OUT). Black components are suitable for most applications to filter noise from the power supply, while orange components are specifically chosen to produce a shot-noise limited signal between 60–500 µW with minimal voltage offset. The transimpedance amplifier converts the input current to a 0–1 V signal and the output buffer multiplies this signal by 10. The output buffer adds capacitive load-driving ability so that the photodetector can be connected to the filter box by a long firewire cable. The capacitance-multiplier power filter reduces the power supplied by a separate 15 V power supply (POS, NEG) to 11 V (VCC, VEE) while removing AC noise from the power supply.

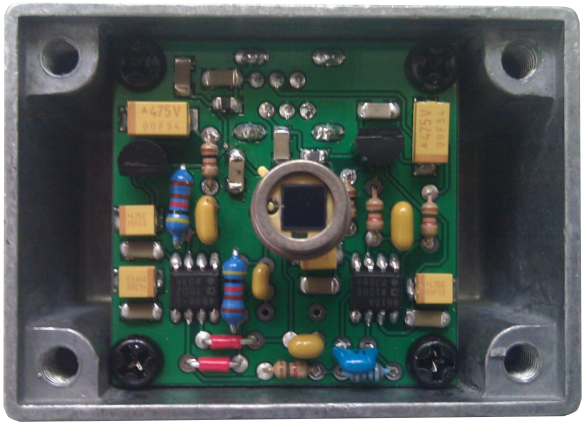

Figure 2.20: One of the photodetectors housed in a small box (with lid removed for photograph). For scale, the circuit board is 25 mm across.

2.6.3 *Signal filters*

Our data acquisition system (National Instruments PCIe-6363) can monitor one photodetector at 2 MSPS, corresponding to a bandwidth of 700 kHz, or monitor 16 photodetectors at the slower rate of 60 kSPS. Since the photodetectors designed in Section 2.6.2 have a bandwidth of 5 MHz, monitoring multiple photodetectors simultaneously would result in a noisy signal. To allow us to simultaneously monitor powers in different locations, I designed and constructed anti-alias filters for each photodetector to optionally reduce the bandwidth to 100 kHz. The filter circuit diagram is shown in Figure 2.21. Results of simulations using the filter design software LTspice are shown in Figure 2.22. Using a Bessel-type filter, in contrast to Butterworth or Chebyshev filters, maintains a good response in the time domain, with minimal overshoot or 'ringing' [75, p 272]. We can use the filters when monitoring many photodetectors, or bypass the filter when we wish to perform a wider bandwidth measurement of only one photodetector.

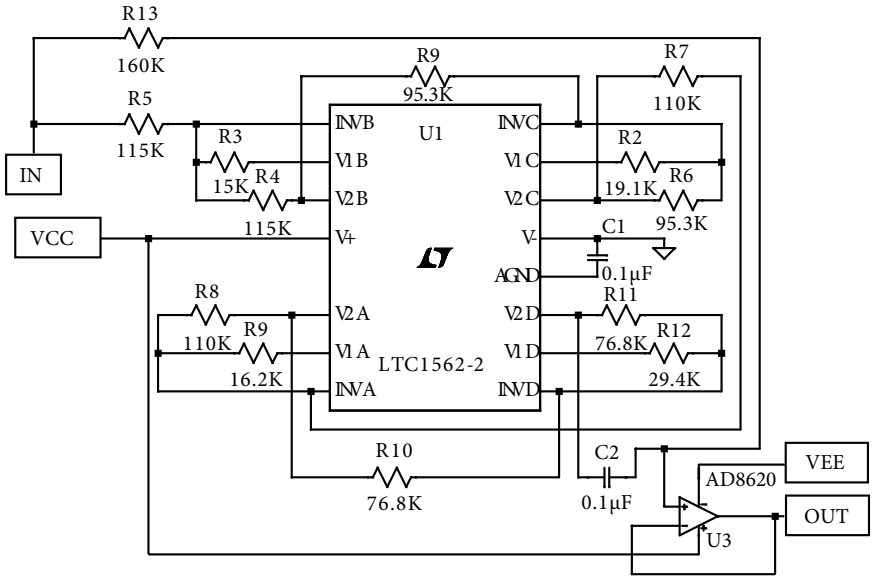

Figure 2.21: Circuit diagram of the Bessel filter used to reduce the bandwidth of the photodetectors when monitoring multiple locations in the optics setup. The values of the resistors surrounding the universal filter chip were set given requirements for the -3 dB point and sharpness of drop in the frequency response curve. The output op-amp topology restores DC accuracy, removing the offset introduced by the filter chip. The circuit is connected to a ± 11 V power supply (VCC, VEE).

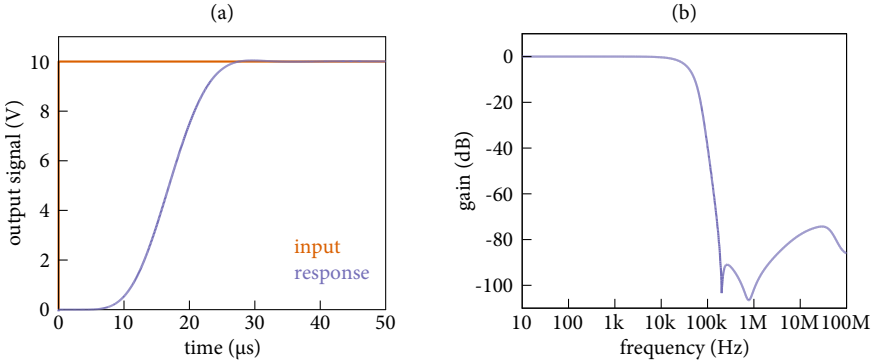

Figure 2.22: LTspice simulation of the photodetector filter circuit. (a) Response to a step function input signal showing the smooth rise with minimal overshoot characteristic of Bessel filters. (b) Frequency response curve showing a -3 dB point of 100 kHz.

2.7 Our imaging systems

The primary measurement technique used in Bose–Einstein condensate physics is absorption imaging [76], where resonant light is absorbed by the condensate and a shadow is cast onto a camera.[8] In our laboratory we have the ability to image condensates from the side or from the top as required. I designed the image-forming optics for the side imaging system, and was in charge of the design and alignment of the image-forming optics for the top imaging system. The experimental details of our typical absorption imaging sequence is discussed in Alex Wood's thesis [61].

As discussed below, both of these imaging system have relatively low resolution; suitable for observing the bulk structure of condensates, but not fine structures like solitons. In Chapter 3 I present a high resolution imaging system, to be installed when we require the ability to image fine structures within condensates.

[8] An example of an alternative measurement technique is Faraday imaging of condensates, developed by my colleague Martijn Jasperse [62].

2.7.1 Side imaging

The simplest of imaging systems is a spherical singlet lens and a camera; outweighing its simplicity is the high degree of spherical aberration such a system exhibits, particularly with a thick glass cell in front of the object. We reduce spherical aberration in our side imaging system by using two 'back-to-back' achromats, with a small distance between the lenses to accommodate an aperture. The objective lens and image-forming lens have a focal length of $f_1 = 50\,\text{mm}$ and $f_2 = 100\,\text{mm}$ respectively, and are both 1 inch in diameter. This small diameter allowed close placement of the optics to the glass cell without obstructing the MOT beams. The lenses and aperture are housed inside a Thorlabs 1 inch lens tube attached to the camera to reduce stray light.

The achromats are infinity-corrected, meaning that aberrations are minimised when the rays from the object are collimated between the two lenses. To achieve this the lenses are positioned as shown in Figure 2.23. The aperture between the lenses has a diameter of 15 mm to further reduce aberration. The size of this aperture was chosen by simulating the imaging system in the imaging design software Zemax. In this simulation, we sought an aperture size which produced a ray trace spot size equal to the diffraction-limited Airy disc; a smaller aperture would cause the spot size to be dominated by diffraction, while a larger aperture would cause the spot size to be dominated by aberration.

If the lenses were positioned *exactly* as shown in Figure 2.23, the magnification of the system would be 2; by tracking the centroid position of the condensate for different drop times, we measured the magnification of the side imaging system to be 2.1. For the AVTGX1920 camera CCD, with a pixel size of 4.5 μm, this magnification gives an effective pixel size at the object plane of 2.1 μm.[9]

[9] This isn't a typographical error, the effective pixel size in μm is coincidentally the same number as the magnification.

2.7.2 Top imaging

Currently our top imaging system uses the y-directed MOT beam[10] as the imaging light; since the MOT beams are detuned 16 MHz below resonance, and are highly

[10] This coordinate system is defined in Figure 2.13.

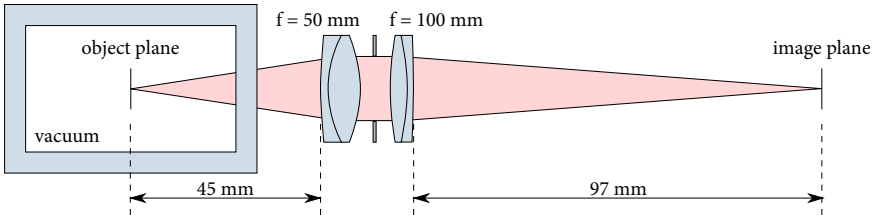

Figure 2.23: Schematic of our side imaging system (to scale). Two infinity-corrected back-to-back achromats magnify our condensates by two, sufficient for imaging the bulk structure of our condensates. A 15 mm diameter aperture between the achromats is used to reduce spherical aberration. This is our primary imaging system, which we use for measurements of condensate atom number and spin projection.

aberrated by the MOT beam expanders, this is far from ideal. Since we to date only use the top imaging system for alignment of the cross-beam dipole trap, this probe beam is sufficient. A 'proper' resonant probe beam will be installed when this top imaging system is replaced with the high-resolution system described in Chapter 3. The $-y$-directed MOT beam reflects off an electromechanically-activated flipper mirror into the glass cell during the MOT trapping and molasses stages. For top imaging, this flipper mirror is then moved out of the MOT beam path to allow the y-directed MOT beam to travel into the top imaging optics.

This top imaging system is similar to the side imaging system, in that it uses two back-to-back infinity corrected achromats. The key difference is that we require space for the flipper mirror in the beam path, preventing us from placing the objective achromat close to the glass cell. As such, to obtain a useful magnification of greater than two, the total track length of the imaging system must be much longer than the side imaging system. The length of the top imaging system is an issue because of space limitations around the apparatus, but this problem can be reduced by using a telephoto lens design.

I chose achromats 2 inches in diameter with focal lengths of $f_1 = 200$ mm and $f_2 = 400$ mm respectively, with the addition of a 1 inch diameter diverging singlet with focal length $f = -100$ mm.[11] This diverging singlet in combination with the achromat increases the effective focal length of the image-forming lens, increasing the magnification of the system without a significant increase in the track length.

[11] This choice of lens combination was influenced by the availability of lenses in the laboratory at the time.

From a simulation using Zemax this choice of lenses could give a magnification of 2.9, while partially cancelling the spherical aberration introduced by the glass cell, if the lenses are positioned as shown in Figure 2.24. In practice, the actual placement of lenses resulted in a magnification of 2.0, measured by comparing the size of the condensate viewed with the top imaging system to the calibrated side imaging system. With this magnification and the camera used (Andor Neo sCMOS), the effective pixel size in the object plane is 3.2 μm.

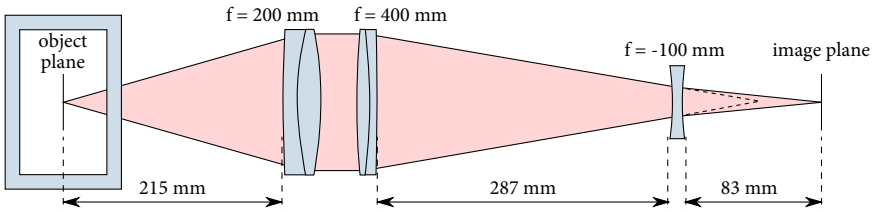

Figure 2.24: Schematic of our top imaging system (not to scale). Two back-to-back achromats and a diverging spherical singlet produce a magnification of 2.9. This telephoto lens design enables the imaging system to fit the space requirements around the apparatus while still producing a sufficient magnification. We use this imaging system primarily for alignment of the cross-beam optical dipole trap.

2.8 Summary

Building a spinor Bose–Einstein condensate apparatus 'from scratch' was an experience I feel proud to have been involved in. My main contributions to this endeavour were: constructing the ultra-high vacuum system and science cell; leading the saga of the vacuum system bakeout to its successful conclusion; assembling the optics to form the different laser beams needed to cool, trap, and image rubidium; making laser shutters, photodetectors, and photodetector signal filters; creating the six intersecting beams for the magneto-optical trap and capturing our first cloud of cooled atoms in this trap; and designing our absorption imaging systems. For the full details of all the elements of our spinor Bose–Einstein condensate apparatus, the interested reader can combine this chapter will associated chapters from the theses of my colleagues Alex Wood [61] and Martijn Jasperse [62].

Chapter 3

A high resolution imaging system

In this chapter I present a high resolution imaging system which we plan to use to resolve small features within condensates, such as solitons. I have bench-tested the performance of the objective lens of this imaging system, but we have not yet used it to take images of our condensates. Once we require the ability to take high-resolution images of condensates, this imaging system will be installed onto our apparatus.

Section 3.1 outlines our requirements for this imaging system, and the main challenges in design. Section 3.2 presents our published paper [59] which documents the design, construction and performance testing of the objective lens. Section 3.3 describes a suggested set of image-forming optics which could be combined with the objective lens when installing the imaging system onto our apparatus.

3.1 Our requirements for the imaging system

In Section 2.7 I described the imaging systems that are currently installed on our apparatus. While these imaging systems are capable of producing absorption images to measure the condensate atom number and spin projection, they do not have a high enough resolution to image small spatial features within the condensate *in situ*. Since one of our aims is to use magnetic resonance control to engineer a black soliton in a condensate, we will need the ability to resolve roughly 1 μm features in our absorption images to verify the success of our new wavefunction engineering technique experimentally.

The main obstacle to achieving this resolution with our apparatus is the thick (5 mm) glass walls of our science cell. The cell was engineered to be optically flat (less than 200 nm of curvature across the surface), which necessitated thick glass to avoid deformation under vacuum pressure. Unfortunately, rays which diverge through a plane-parallel glass plate become spherically aberrated by an amount proportional to the thickness of the glass plate [77, p.85]. This means that in order to achieve a diffraction-limited resolution the imaging system needs to be able to correct for this spherical aberration introduced by our glass cell.

One approach to solving this problem is to place a lens, or system of lenses, inside the glass cell, to collimate the rays before they pass through the glass cell. Such collimated rays would not develop any aberration as they pass through the glass cell. This approach has been used to create some beautiful high-resolution imaging systems [78, 79], but was unappealing to us because it requires in-vacuum optics; since in our laboratory we would prefer to never break vacuum ever again[1], such a design would prevent changes to our imaging system if our requirements changed over time. For us, the best approach was to have all our optics outside the glass cell, and correct for the introduced spherical aberration.

[1] You might understand this preference if you have read Section 2.3 of this thesis.

Our high-resolution imaging system is split into two stages. The first stage, called the *objective lens*, collimates the Abbe-rays[2] from the object plane while correcting the spherical aberration from the glass cell. The second stage, called the *imaging lens*, focuses the rays to form the image. Such a system is common in microscopes, and enables us to replace the (simpler, cheaper) imaging lens with moderate ease if we decide to change the magnification of the imaging system, without modifying the objective lens.

[2] If we were working with fluorescent images, it would suffice to consider the classical rays emanating from the object. For absorption images, we must instead consider the Abbe-rays. An excellent discussion of this subject can be found in "An introduction to Fourier optics" by J. W. Goodman [80].

The objective lens needs to have a long working distance to be positioned outside the glass cell, 35 mm from the object plane. This rules out the majority of commercially-available microscope objectives, designed to be placed close to a microscope sample slide. Even if a microscope objective had a long working distance, they are designed to correct only the aberration introduced by a thin glass cover slip (typically 170 μm thick) rather than a thick glass wall.

In addition to our resolution and working distance requirements, the imaging system needs to produce a wide diffraction-limited field of view of a few hundred micrometers, so that the entire condensate can be imaged. We could have decided to pay several thousand dollars for a custom-manufactured objective lens, where the individual lens elements are custom-ground to meet, or even exceed, our particular requirements [81–86]. Instead, with patience and the lens design software Zemax, I was able to design an objective lens which uses a combination of commercially-available spherical singlets which is both perfectly adequate for our needs and better suited to our research budget, costing less than one thousand dollars to manufacture.

3.2 *Designing and testing the objective lens*

In this section I present our published paper on the objective lens, which details the design, construction and performance testing of the objective lens. I should first briefly outline the contributions of the authors of this paper. I am responsible for the design and assembly of the objective lens, with the advice of my supervisors Lincoln Turner and Russell Anderson. Martijn Jasperse wrote the camera control software which we continue to use in our research lab. Philip Starkey and Christopher Billington developed our lab control system [87] which I used to automate the data collection for mapping the field of view of the objective lens. Philip Starkey earned the place of second author by also assisting me with interfacing the elec-

tronic translation stages with the control system software, and teaching me how to use Python to analyse my data.

I have reproduced this paper here with permission of the publishers, with the agreement that the following text is included:

> *"This paper was published in Optics Express and is made available as an electronic reprint with the permission of OSA. The paper can be found at the following URL on the OSA website:* www.opticsinfobase.org/oe/abstract.cfm?uri=oe-21-7-9011. *Systematic or multiple reproduction or distribution to multiple locations via electronic or other means is prohibited and is subject to penalties under law."*

A versatile high resolution objective for imaging quantum gases

L. M. Bennie,* P. T. Starkey, M. Jasperse, C. J. Billington, R. P. Anderson, and L. D. Turner

School of Physics, Monash University, Victoria 3800, Australia

Abstract: We present a high resolution objective lens made entirely from catalog singlets that has a numerical aperture of 0.36. It corrects for aberrations introduced by a glass window and has a long working distance of 35 mm, making it suitable for imaging objects within a vacuum system. This offers simple high resolution imaging for many in the quantum gas community. The objective achieves a resolution of 1.3 µm at the design wavelength of 780 nm, and a diffraction-limited field of view of 360 µm when imaging through a 5 mm thick window. Images of a resolution target and a pinhole show quantitative agreement with the simulated lens performance. The objective is suitable for diffraction-limited monochromatic imaging on the D2 line of all the alkalis by changing only the aperture diameter, retaining numerical apertures above 0.32. The design corrects for window thicknesses of up to 15 mm if the singlet spacings are modified.

OCIS codes: (110.0180) Microscopy; (120.3620) Lens system design; (020.1475) Bose–Einstein condensates.

References and links

1. Y. Kawaguchi and M. Ueda, "Spinor Bose–Einstein condensates," Phys. Rep. **520**, 253–381 (2012).
2. E. W. Streed, A. Jechow, B. G. Norton, and D. Kielpinski, "Absorption imaging of a single atom," Nat. Commun. **3**, 933 (2012).
3. H. Gross, *Handbook of Optical Systems* (Wiley-VCH, 2005) vol. 1.
4. Y. R. P. Sortais, H. Marion, C. Tuchendler, A. M. Lance, M. Lamare, P. Fournet, C. Armellin, R. Mercier, G. Messin, A. Browaeys, and P. Grangier, "Diffraction-limited optics for single-atom manipulation," Phys. Rev. A **75**, 013406 (2007).
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**, 74–77 (2009).
6. E. A. Salim, S. C. Caliga, J. B. Pfeiffer, and D. Z. Anderson, "High-resolution imaging and optical control of Bose–Einstein condensates in an atom chip magnetic trap," arXiv:1208.4897 (2012).
7. Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, "Rapid production of ^{87}Rb Bose–Einstein condensates in a combined magnetic and optical potential," Phys. Rev. A **79**, 063631 (2009).
8. D. Douillet, E. Rolley, C. Guthmann, and A. Prevost, "An easy-to-build long working distance microscope," Physica B **284–288**, 2059–2060 (2000).
9. W. Alt, "An objective lens for efficient fluorescence detection of single atoms," Optik **113**, 142–144 (2002).
10. T. B. Ottenstein, "A new objective for high-resolution imaging of Bose–Einstein condensates," Diploma thesis, University of Heidelberg (2006).
11. K. D. Nelson, X. Li, and D. S. Weiss, "Imaging single atoms in a three-dimensional array," Nat. Phys. **3**, 556–560 (2007).
12. R. Bücker, A. Perrin, S. Manz, T. Betz, C. Koller, T. Plisson, J. Rottmann, T. Schumm, and J. Schmiedmayer, "Single-particle-sensitive imaging of freely propagating ultracold atoms," New J. Phys. **11**, 103039 (2009).
13. 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**, 68–72 (2010).

#185535 - $15.00 USD Received 19 Feb 2013; revised 25 Mar 2013; accepted 27 Mar 2013; published 4 Apr 2013

14. B. Zimmermann, T. Müller, J. Meineke, T. Esslinger, and H. Moritz, "High-resolution imaging of ultracold fermions in microscopically tailored optical potentials," New J. Phys. **13**, 043007 (2011).
15. B. P. Anderson and M. A. Kasevich, "Spatial observation of Bose–Einstein condensation of ^{87}Rb in a confining potential," Phys. Rev. A **59**, 938–941 (1999).
16. The objective can also be modeled using the free software OSLO EDU, which is limited to ten surfaces, because the first surface is flat and need not be included in the model.
17. H. Gross, *Handbook of Optical Systems* (Wiley-VCH, 2007) vol. 3.
18. W. J. Smith, *Modern Lens Design*, 2nd ed. (McGraw-Hill, 2005).
19. R. E. Fischer and K. L. Mason, "Spherical aberration – some fascinating observations," Proc. SPIE **0766**, 53–60 (1987).
20. P. T. Starkey, C. J. Billington, S. P. Johnstone, M. Jasperse, K. Helmerson, L. D. Turner, and R. P. Anderson, "A scripted control system for autonomous hardware-timed experiments," arXiv:1303.0080 [cond-mat.quant-gas].

1. Introduction

Imaging systems are an essential component of quantum gas experiments. There is an increasing need for high resolution imaging ($< 3\,\mu$m) to extract fine detail, such as topological defects in a Bose–Einstein condensate [1], or a single trapped ion [2]. One difficulty faced is that the object of interest is held within a vacuum chamber and must be viewed through a thick glass window. Even a perfect optical flat introduces significant spherical aberration [3, p. 572] in a large numerical aperture (NA) system.

One approach is to house an infinity-corrected objective within the vacuum chamber [4, 5]. Collimated rays from such objectives are not significantly aberrated because they pass through an optically flat window at near-normal incidence. With this approach the objective can have a short working distance, allowing for large NAs and permitting the use of commercial microscope objectives [6]. However *in vacuo* optics must be vacuum-compatible and bakeable, and are fixed in positions that limit optical access to the experiment.

Optics housed external to the vacuum chamber are particularly desirable, for example, in a Bose-Einstein condensate apparatus that shares a common source chamber and science chamber [7], or in the study of liquid helium using optical cryostats [8]. The *ex vacuo* objective must have a long working distance if the object is located far ($\gtrsim 5\,$mm) within a vacuum chamber and must also correct for the spherical aberrations introduced by the vacuum window. Fortunately in quantum gas experiments the illumination is typically monochromatic with a small field of view ($\sim 1\,$mm), nevertheless long working distance objectives that correct for large spherical aberrations are not available commercially and remain a design challenge. Published designs that meet these criteria [9–14] require manufacturing at least one custom singlet, which is a slow and costly process. Because of this difficulty, many experiments employ a single lens, or a single light-gathering objective lens followed by a single image forming lens [15]. Such imaging systems are not corrected for aberrations and consequently have low resolution.

This article provides an alternative solution to the problem of high resolution imaging with *ex vacuo* optics: an objective consisting of entirely catalog singlets that achieves diffraction-limited imaging with an NA of 0.36. The objective achieves a resolution of 1.3 µm with a field of view (FOV) of 360 µm for a 5 mm thick window using 780 nm illumination. The objective is suitable for diffraction-limited imaging of all the alkalis, retaining NAs above 0.32 by changing only the aperture diameter and accounting for the chromatic focal shift. The objective remains well-corrected into the near ultraviolet; we predict sub-micrometer resolution when imaging Yb^+ at 370 nm in ion trapping experiments. A useful diffraction-limited FOV can be retained for glass windows up to 15 mm thick by changing the singlet spacings. The required spacings can be found through re-optimization of the design in ray-tracing software.

#185535 - $15.00 USD Received 19 Feb 2013; revised 25 Mar 2013; accepted 27 Mar 2013; published 4 Apr 2013

Table 1. (a) Lens prescription of the objective suitable for entry into a ray-tracing system ($\lambda = 780$ nm, 5 mm silica window, 35 mm working distance, $f = 47$ mm). (b) Changing only the aperture diameter enables diffraction-limited imaging of all the alkalis and Yb^+.

(a) Objective prescription

Surface	Curvature rad. (mm)	Thickness (mm)	Material
1	∞	4.00	BK7
2	51.50	10.50	air
3	89.47	10.36	BK7
4	-89.47	0.76	air
5	47.90	7.30	BK7
6	119.30	1.12	air
7	30.30	9.70	BK7
8	65.80	7.30	air
9	∞	14.94	air
10	∞	5.00	silica
11	∞	15.00	vacuum

(b) Objective performance

	λ (nm)	Ap (mm)	Res (μm)	NA	FOV (μm)
Li	671	24.0	1.20	0.34	400
Na	589	22.0	1.13	0.32	560
K	767	25.0	1.33	0.35	400
Rb	780	25.4	1.31	0.36	360
Cs	852	26.0	1.44	0.36	340
Fr	718	24.0	1.29	0.34	470
Yb^+	370	16.0	0.90	0.25	520

Ap = aperture diameter.
Res = resolution
FOV = diffraction-limited FOV diameter

2. Design and construction

Our Bose–Einstein condensate apparatus is similar to [7], but with an optically flat glass cell as our science chamber to maximize optical access. The magnetic coils surrounding the glass cell accommodate 50.8 mm diameter lenses, while the distance between the 5 mm thick glass window and the condensate demands a working distance greater than 25 mm. We require a diffraction-limited FOV of at least 300 μm for imaging condensates *in situ*.

The objective was designed using Zemax 12 SE [16], a ray-tracing program which varies parameters to optimize the lens performance. The design of Alt [9] was chosen as a starting point because it achieved a diffraction-limited performance using BK7 glass—a material widely available in standard lens catalogs—with only four singlets. Our design is distinguished from Alt's by a larger NA and the use of only catalog lenses (Thorlabs and Newport). After scaling the design to 50.8 mm diameter singlets, we optimized both on- and 200 μm off-axis for maximal NA with the required working distance. The optimization routine chosen in Zemax minimized the squared sum of all wavefront aberrations up to 7th order. The default merit function was modified to seek a diffraction-limited modulation transfer function (MTF). This ensures high image contrast for objects with spatial frequencies up to the resolution determined by the NA [17, p. 132]. Initially all singlet thicknesses, spacings and curvatures were varied until the diffraction limit was reached for the desired working distance and NA. We then replaced the singlet most similar to a catalog lens, fixing its curvature and thickness to the catalog values, and repeated the optimization routine for the remaining lenses. This process was repeated until all singlets were catalog lenses. During this process we occasionally made dramatic changes to lens curvatures and spacings to ensure that the optimization routine did not stagnate in a local minimum [18].

The final design comprises four catalog singlets: two positive meniscus lenses to enable a high NA with minimal addition of spherical aberration, a bi-convex lens, and a rear plano-concave lens to cancel aberrations introduced by the other lenses and by the window [19]. This design is shown in Fig. 1, with the singlet curvatures and spacings listed in Table 1(a). Our objective has an NA of 0.36 yielding a resolution of 1.3 μm at the design wavelength of 780 nm. It has a long working distance of 35 mm measured from the front aperture (surface 9, Fig. 1) and an effective focal length of 47 mm. The objective is infinity-corrected, allowing the image-

#185535 - $15.00 USD Received 19 Feb 2013; revised 25 Mar 2013; accepted 27 Mar 2013; published 4 Apr 2013

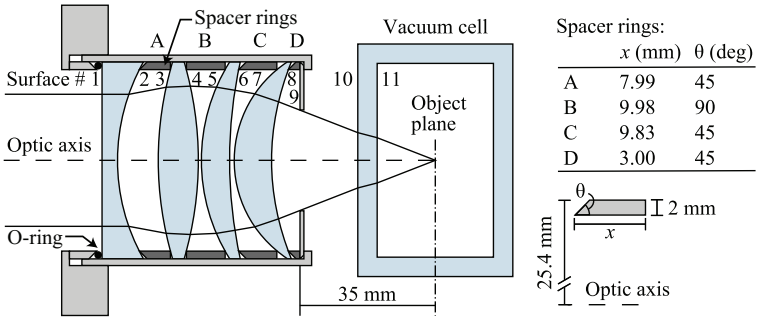

Fig. 1. Cross-section of the objective in its aluminium housing assembly. The four catalog lenses, from left to right, are Thorlabs LC1093-B, Newport KBX151AR.16, Thorlabs LE1418-B and Thorlabs LE1076-B. The cell is fused silica and is optically flat to less than $\lambda/4$. The rays shown propagate from the object plane through the maximum aperture of the objective. The geometry of the spacer rings A, B, C and D are shown.

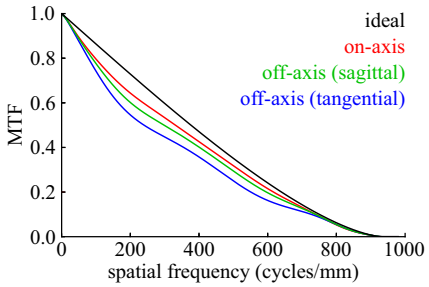

Fig. 2. The calculated MTF of the objective. On-axis and 180 µm off-axis, both tangential and sagittal, exhibit marginally lower contrast than an aberration-free lens with the same NA. Residual aberrations are a consequence of limiting the design to catalog singlets.

forming optics to be chosen separately to suit various applications. The aberration produced by the 5 mm thick window is corrected over a diffraction-limited FOV of 360 µm. Here we consider the FOV to be 'diffraction-limited' where the Strehl ratio is greater than 0.8 [17, p. 90]. The calculated MTF, both on-axis and at the edge of the FOV, is compared to the diffraction-limited MTF in Fig. 2. The similarities between the curves indicate that the residual aberrations will not significantly degrade image contrast across the 360 µm diameter FOV.

Our objective is applicable to imaging all the alkalis and Yb^+ at commonly used wavelengths. Retaining a diffraction-limited FOV of a few hundred micrometers only requires changing the aperture diameter (Table (1b)). The objective is not corrected for chromatic aberration, so the effective focal length changes with wavelength and hence the objective is only suitable for monochromatic light.

While a full tolerance analysis was not performed, the objective appears robust to changes in singlet spacing: rounding the spacings (thickness of surfaces 2, 4 and 6 in Table 1) to the nearest 0.5 mm had negligible effect on the predicted performance. Similarly, the performance and focus are robust to thermal expansion of the spacers across ± 10 K. This allowed the housing assembly to be the simple design shown in Fig. 1. The singlets are held on-axis inside a smooth aluminium tube and are separated by aluminium spacer rings. These rings contact each singlet

#185535 - $15.00 USD Received 19 Feb 2013; revised 25 Mar 2013; accepted 27 Mar 2013; published 4 Apr 2013

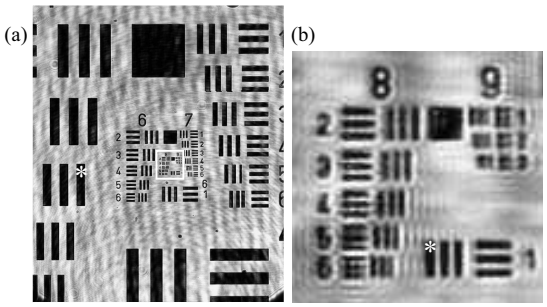

Fig. 3. Image of a USAF 1951 resolution target, illuminated with coherent light, formed by the objective and an $f = 1000$ mm achromat. The white rectangle in (a) marks the region shown in (b). The bars marked with asterisks are (a) 22.1 µm and (b) 1.95 µm wide in the object plane.

2 mm from its outer edge, avoiding chamfers which can result in large spacing errors. The tube screws into a mounting bracket to hold the singlets in place and an o-ring moderates the pressure applied to the singlets. Aluminium components are anodized black to reduce reflections.

3. Experimental performance

We measured the FOV and point spread function of the objective using a USAF 1951 resolution target (Edmund Optics 58-198) and a 1 µm pinhole (Edmund Optics 39-878), both illuminated by collimated 780 nm laser light. A 5 mm fused silica optical flat was used to mimic a vacuum window. An $f = 1000$ mm achromat placed immediately after the objective formed images of the test objects at a magnification of -21.4 on an sCMOS camera (Andor Neo). This camera has a pixel size of 6.50 µm, corresponding to an effective pixel size of 304 nm in the object plane. Including the achromat in our simulation further compensated aberrations, permitting an increase of both the FOV to 400 µm and the aperture diameter to 26 mm. Accordingly, this aperture diameter was used in the constructed objective.

The image of the resolution target (Fig. 3) is undistorted beyond the diffraction-limited FOV. The line pairs in element 6 of group 8 are clearly resolved, corresponding to a resolution of ≤ 2.20 µm. The predicted resolution of 1.3 µm would also resolve the elements in group 9, but the coherent illumination produced diffraction fringes which degraded the image quality.

The point spread function—the image of a point source—provides another measure of the lens performance. The measured image of a 1 µm pinhole is similar to the convolution of a 1 µm top-hat function with the simulated point spread function (Fig. 4(d), inset). The azimuthal averages of these images about the pinhole center (Fig. 4) are in agreement, affirming the objective performed as expected. Asymmetry in the on-axis pinhole image is likely due to tilt of the optical flat or objective relative to the camera [10].

The pinhole was also used to measure the diffraction-limited FOV and to locate the position of the optic axis relative to the camera chip. We raster-scanned the pinhole across the object plane using two motorized translation stages under automated control [20], taking 3600 images on a 10 µm grid in one hour. At each point on the grid we fitted a two dimensional Gaussian to the pinhole image; the geometric mean of the fitted rms diameters (spot size) are shown in Fig. 4(a). Row- and column-wise averages of the spot size revealed distinct minima, which we took to be the position of the optic axis (cross in Fig. 4(a)). The diffraction-limited FOV was measured to be 350 µm in diameter, commensurate with a circle inside which the spot size is less than $\sqrt{2}$ times the on-axis spot size. The measured FOV is smaller than the simulated 400 µm,

#185535 - $15.00 USD Received 19 Feb 2013; revised 25 Mar 2013; accepted 27 Mar 2013; published 4 Apr 2013

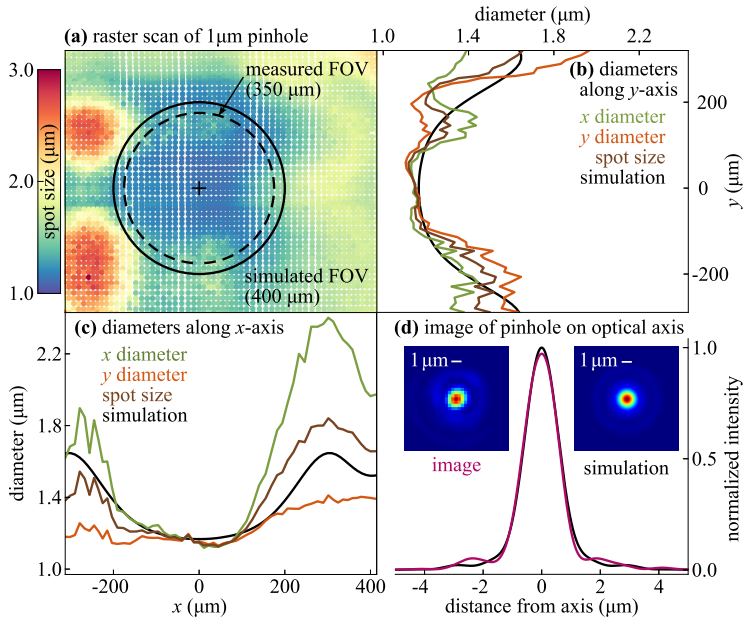

Fig. 4. (a) The spot size of a 1 μm pinhole measured at 3600 positions across the object plane. (b, c) The spot size and constituent diameters across the *y*- and *x*-axis respectively through the optic axis (cross in (a)). The measured spot size is in good agreement with our simulation within the FOV. (d) The on-axis pinhole image compared to the simulated point spread function convolved with a 1 μm pinhole. The curves are the azimuthal averages of the inset images.

owing to the asymmetry in spot size about the optic axis. This asymmetry was likely caused by tilt in the translation stages that moved the pinhole through either side of focus across the object plane. Figures. 4(b) and 4(c) show the diameters of the pinhole image along the column and row that intersect the optic axis respectively. In addition to the spot size, the constituent *x* and *y* Gaussian rms diameters are shown. The variation between the *x* and *y* diameters reveals small astigmatism, suggesting tilt of the optical flat and objective relative to the camera. The measured spot size compares favorably to simulation both on- and off-axis.

4. Conclusion

The combination of catalog singlets presented in this paper form a high resolution, long working distance objective suitable for imaging objects far within a vacuum chamber through a thick window. Optical tests confirm the performance predicted using ray-tracing software. Remarkably this objective enables diffraction-limited imaging across a wide FOV of all the alkalis and ytterbium ions, by changing only the aperture diameter. Spherical aberration from a range of window thicknesses can be corrected by changing the singlet spacings. The versatility and simplicity of this design makes it applicable to many experiments in the field of quantum gases.

Acknowledgments

This work was supported by Australian Research Council grants DP1094399 and DP1096830, and by a Monash University Faculty of Science equipment grant.

#185535 - $15.00 USD Received 19 Feb 2013; revised 25 Mar 2013; accepted 27 Mar 2013; published 4 Apr 2013

3.3 Design of the full imaging system

Figure 3.1 shows a schematic of the high-resolution objective lens positioned in the apparatus, along with suggested image-forming optics. The bore of our quadrupole coils is (deliberately) large enough to accommodate 2-inch lenses housed in a lens tube, so that the high-resolution objective lens can be inserted into this hole and positioned 15 mm from the surface of the glass cell. The vertical MOT trapping beams pass through the quadrupole coils and reflect off an electromechanically-actuated flipper mirror; During the imaging stage this flipper mirror moves aside, so that there are no moving elements used in the imaging system itself.

When the objective lens is installed, the MOT beam travelling downwards into the cell will be focused by the objective lens. To correct for this and ensure that the MOT beam is collimated within the glass cell, the MOT beam will need to be brought to a focus 22 mm before the back element of the objective lens.

When choosing the image-forming optics to accompany the objective lens, we need to satisfy three criteria. First, the magnification must be high enough so that the effective pixel size of the camera CCD in the object plane is smaller than the point-spread function of the objective lens; otherwise we would be underutilising our high resolution objective. Second, the image-forming optics should introduce only minimal aberration, so that the imaging system remains diffraction-limited. Third, the total length of the imaging system should not become too long, otherwise the imaging system may not fit in the space available around the apparatus without several folding mirrors, and we wish to use minimal optical elements to reduce etaloning in the images from mechanical vibrations. The first two criteria could be satisfied by using a single achromat with a long focal length. But to also satisfy the third criteria, we must use two achromats to form a telephoto lens, instead of a single achromat, which decreases the length of the optical system while maintaining the level of magnification.

One possible design of the image-forming optics is to use two achromats with focal lengths $f = 250\,\text{mm}$ and $f = -100\,\text{mm}$ positioned as shown in Figure 3.1. With this design the object plane and imaging plane are separated by $400.4\,\text{mm}$, and the magnification of the image is 10. If we use our Andor Neo sCMOS camera, this gives an effective pixel size of $650\,\text{nm}$ in the object plane, which is smaller than the resolution of the objective lens. Figure 3.2 shows that these lenses maintain a diffraction-limited performance. The point-spread function (PSF) of the full imaging system at the image plane is shown in Figure 3.3.

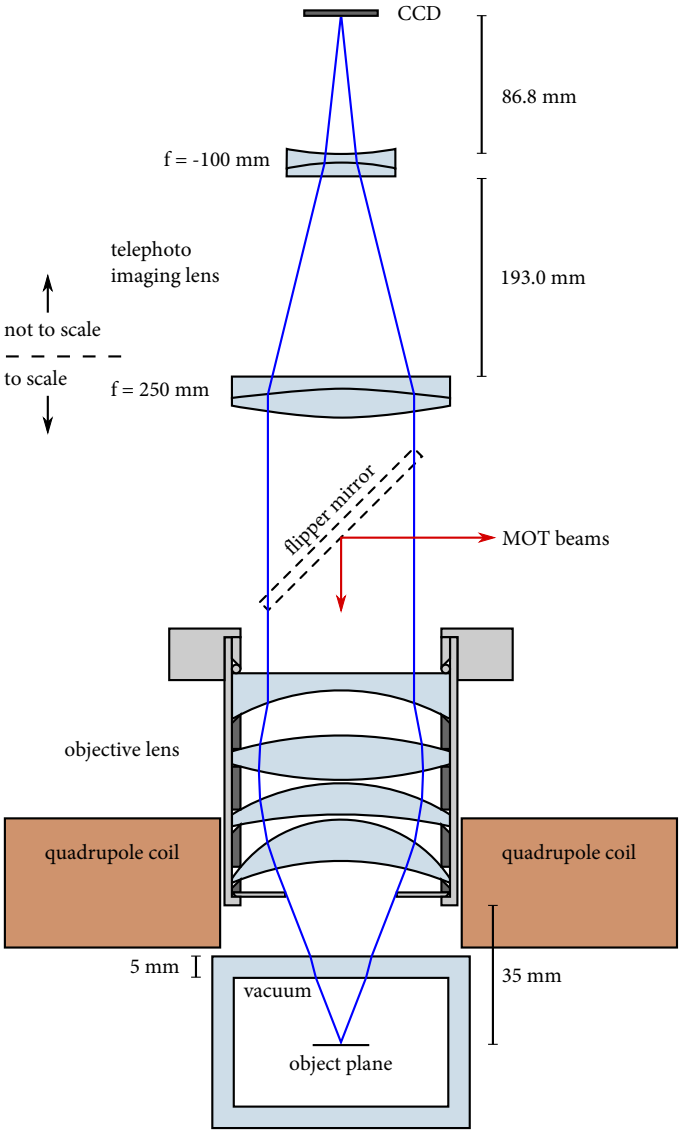

Figure 3.1: A cross-section schematic showing the objective lens positioned inside the quadrupole coil above the glass science cell (to scale), with the imaging lens positioned above (not to scale). Dark blue lines indicate the maximal Abbe-rays emanating from the object plane. These rays are collimated by the objective lens, then brought to a focus on the CCD by the telephoto imaging lens (consisting of two achromats). An electronically-actuated flipper mirror is located between the objective lens and the imaging lens. When moved into position, this directs the vertical MOT beam down into the glass cell.

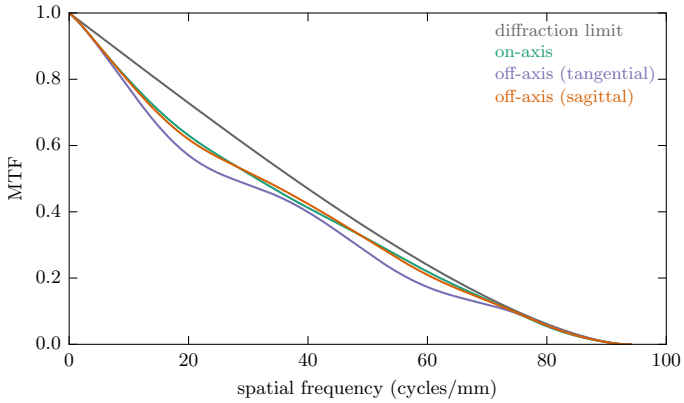

Figure 3.2: The simulated modulation transfer function (MTF) of the full imaging system. By comparing this to the MTF of the objective lens alone, in Figure 2 of Section 3.2, we can see that the choice of the imaging lens has not introduced significant additional aberrations to the imaging system.

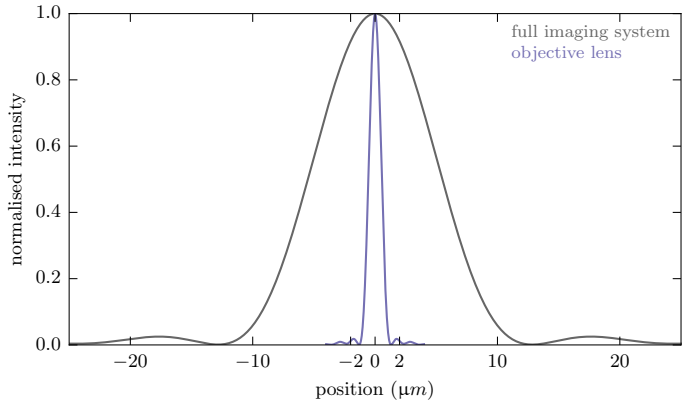

Figure 3.3: The point spread function of the full imaging system (grey) compared to the point spread function of the objective lens. Remember that the imaging lens has deliberately produced a magnification of 10.

3.4 Summary

This objective lens enables high resolution imaging through the thick glass wall of a science cell without the expense of custom-manufactured lens elements. It is suitable for absorption imaging of all the alkali elements, making it widely applicable to the quantum gas research community. In our laboratory, it will be essential when we wish to image topological defects in Bose-Einstein condensates *in situ*.

Chapter 4

Simulating magnetic resonance control

In this chapter I simulate engineering a soliton in a Bose–Einstein condensate using our proposed wavefunction engineering protocol, which we call magnetic resonance control (MRC). The key results of this chapter form the basis of my first-author paper published on the ArXiv [60], currently in peer-review. Aspects of this work were developed in collaboration with PhD student Paul Wigley, Dr Stuart Szigeti and Associate Professor Joseph Hope.

As discussed in Section 1.2, a black (stationary) soliton is the sharpest stable structure supported by a single component condensate. A black soliton exhibits a π phase step across a density zero with a width on the order of the healing length of the condensate. The healing length is $\xi = (8\,\pi\,n_0\,a_s)^{-1/2}$, where n_0 is the peak atomic density of the condensate and a_s is the s-wave scattering length. In rubidium condensates the healing length is typically smaller than optical wavelengths. For example, a condensate with a peak atom density of $n_0 = 10^{14}$ atoms/cm^3 and a s-wave scattering length of $a_s = 5.3$ nm has a healing length of $\xi = 270$ nm.

When magnetic resonance control was still at the whiteboard stage of design, we had to decide on an example target state we would aim to engineer. We decided to engineer a black soliton to demonstrate the precision of MRC. The ability to engineer a black soliton is a stringent test of any wavefunction engineering protocol, because a black soliton can only be engineered using a protocol that manipulates *both* the phase and density of the macroscopic wavefunction with *healing length resolution*. Developing an experimentally feasible technique for engineering a black soliton has been a compelling topic, because a truly black soliton has never been experimentally achieved in a condensate; existing wavefunction engineering techniques, discussed in Section 1.3, have lacked either the required precision, simultaneous phase and density control, or both.

Sections 4.1 and 4.1.1 provide a conceptual overview of the protocol used to engineer a black soliton in a condensate, while Section 4.2 outlines how we simulate one-dimensional condensates. Section 4.3 presents the key result of this chapter: using MRC to engineer a black soliton in a condensate with a sub-optical healing length using experimentally feasible parameters. Section 4.4 discusses the details that must be considered when selecting the parameters of the protocol. We then

extend the base protocol to produce multiple solitons in Section 4.5, including a multiple-soliton state that has yet to be experimentally realised.

4.1 The protocol to engineer a soliton

In magnetic resonance imaging (MRI), a magnetic field gradient combined with coupling between two internal states provides spatial resolution of spin [88]. We propose a 'reversal' of this technique; instead of *imaging* spins we aim to *control* them.

We start by describing magnetic resonance control in the context of a two-level system, then from Section 4.2 onwards we change to a three-level system. We take this approach despite the integer spin of bosons because the essential physics of our technique is the spin degree-of-freedom of the condensate, rather than the *dimension* of this spin degree-of-freedom. In any case, pseudospin-1/2 systems are readily realised in bosonic quantum fluids [89]. After describing this protocol for a pseudospin-1/2 condensate, we move to the full three-level system for our simulations of a condensate in the $F = 1$ state of ^{87}Rb.

Consider a one-dimensional, spin-1/2 condensate described by the spinor

$$\mathbf{\Psi}(z, t) = \begin{pmatrix} \psi_{\downarrow}(z, t) \\ \psi_{\uparrow}(z, t) \end{pmatrix} . \tag{4.1}$$

The spatial population distribution and phase of this condensate is shown in Figure 4.1 during the stages of MRC.

The condensate begins in the spin-down state: $\psi_{\uparrow}(z, 0) = 0$. The condensate is at rest within a harmonic potential, meaning that the condensate phase is uniform; without loss of generality, we can set $\arg(\psi_{\downarrow}(z, 0)) = 0$. This initial state is shown in Figure 4.1(a).

During the first stage of MRC, population in one side of the condensate is transferred from the spin-down state $|\downarrow\rangle$ to the spin-up state $|\uparrow\rangle$. We arbitrarily choose the 'left' side, $z < 0$, as shown in Figure 4.1(b). Ordinarily, resonant coupling would transfer the population of the whole condensate. To obtain spatial selectivity we apply a linear magnetic field $B(z) = B_0 + \frac{dB}{dz}z$ which spatially varies the energy splitting between state $|\downarrow\rangle$ and state $|\uparrow\rangle$. This enables the use of an adiabatic coupling pulse to address a spatial subset of the condensate. The nature of this adiabatic coupling pulse is discussed in Section 4.1.1. For the moment, it suffices to imagine a coupling pulse that transfers the $z < 0$ region of the condensate from state $|\downarrow\rangle$ to state $|\uparrow\rangle$.

The magnetic field gradient causes a spatially-dependent phase across the condensate, thereby exerting a force on the condensate. Since the direction of the force is state-dependent, the phase variation is much smaller on the left side because the left side of the condensate spends roughly half the time in each state during the pulse. As a result, the two sides of the condensate move in different directions after the first coupling pulse. This effect is discussed further in Section 4.4.

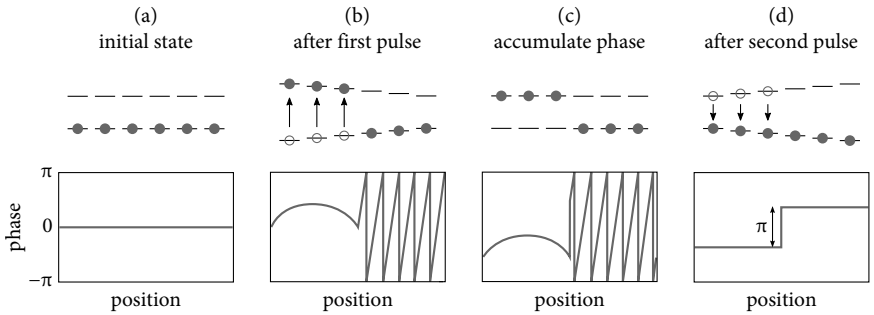

Figure 4.1: The stages of the MRC protocol used to engineer a soliton in a condensate. This figure explains how the phase of the condensate is engineered. Energy level diagrams indicate the population in the two energy levels across the condensate. (a) Initially the condensate is entirely in $|\downarrow\rangle$ with uniform phase of $\arg(\psi_\downarrow(z, 0)) = 0$. (b) A magnetic field gradient shifts the energy level splitting across the condensate, so that an adiabatic coupling pulse can excite the left side of the condensate to $|\uparrow\rangle$. (c) The magnetic field gradient is removed, and we wait until a phase difference of π accumulates between the two sides. (d) The magnetic field gradient is reversed, and a second coupling pulse returns the left side to $|\downarrow\rangle$. This removes the phase gradient introduced by the first magnetic field gradient, revealing the accumulated π phase step.

After the first pulse, once the left side of the condensate is in state $|\uparrow\rangle$ (with the right side still in state $|\downarrow\rangle$), we remove the magnetic field gradient, returning to a uniform energy level splitting across the condensate. Because the left side is in a higher energy state compared to the right side, the left side accumulates phase at a faster rate than the right side. The two states of the condensate will accumulate a phase difference of π after an interval of $t_\phi = \pi/\Delta_\phi$, where Δ_ϕ is the detuning during the interval between the two pulses. This phase-accumulation stage is shown in Figure 4.1(c).

Next we reintroduce the magnetic field gradient, with opposite sign, and apply a second adiabatic coupling pulse to return the left side of the condensate from state $|\uparrow\rangle$ to state $|\downarrow\rangle$. Both sides of the condensate are now in state $|\downarrow\rangle$, with a π phase step at $z = 0$, as shown in Figure 4.1(d).

Note that in addition to the phase difference which accumulates due to one side of the condensate being in a higher energy state, there will also be some phase shift due to mean-field interactions. Such phase shifts depend on the s-wave scattering length and the atomic density. For the $F = 1$ ^{87}Rb condensate we simulate later in this chapter, the singlet and triplet scattering lengths are very similar, and as a result the mean-field phase shift will be very small over the time scale of the protocol compared to the effect of the accumulated phase difference. By simulating the condensate using the Gross–Pitaevskii equation, we include the mean-field phase shift and can adjust the protocol to counter this effect if necessary.

At this point, you should be asking yourself the following:

1. Complete wavefunction engineering requires control over both the phase and density of the condensate; how does MRC engineer the condensate density?

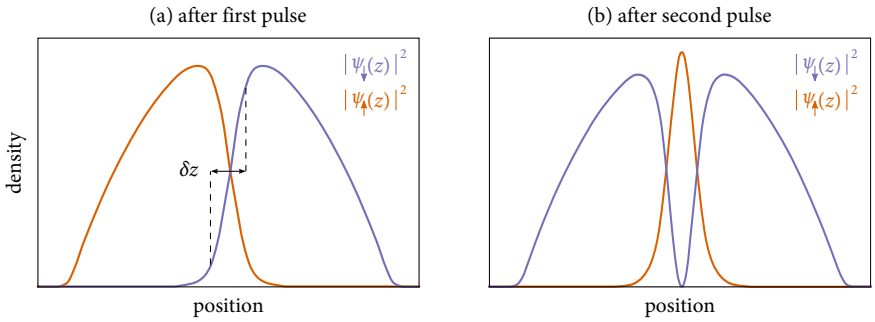

Figure 4.2: The finite pulse duration provides the ability to engineer the condensate density at the phase step. (a) A finite duration pulse has a finite slice sharpness of δz. (b) After the second pulse, this finite slice sharpness leaves some population remaining in $|\uparrow\rangle$ near $z = 0$. This population is then removed to form a notch in the condensate density. If we design the pulse to make this notch have the same width as a black soliton, and at the same time engineer the requisite phase step of π (Figure 4.1), we will engineer a black soliton in the condensate.

2. Wouldn't the magnetic field gradient adversely perturb the condensate?

To answer the first question, MRC engineers the condensate density using an intrinsic feature of the coupling pulse: the finite pulse duration. Only an infinitely long pulse can transfer population in the region $z < 0$ with infinite sharpness at the boundary ($z = 0$). When a finite duration pulse is applied the boundary of the transferred region will be blurred over some distance δz, as shown in Figure 4.2(a). In MRI literature this distance is called the *slice sharpness* [90]. The finite slice sharpness of each coupling pulse results in some population remaining in state $|\uparrow\rangle$ in the region $z = \pm\delta z/2$ after the second pulse (Figure 4.2(b)).

To create an unfilled soliton in state $|\downarrow\rangle$, we remove this 'residual' population from the condensate.[1] This produces a notch in the condensate density at the location of the phase step. We must select the slice sharpness of the coupling pulse so that the resulting density dip has the same width as a black soliton. When the protocol also generates the requisite π phase step, a black soliton is engineered.

[1] A suggested technique for doing this is discussed in Section 4.3.5.

Returning to the second question, if the protocol is too slow then the force from the magnetic field gradient will cause bulk movement of the condensate; an issue discussed further in Sections 4.1.6 and 4.4.2. If this movement is too large then the slice sharpness is broadened, and taken to the extreme the region addressed by the second pulse may not coincide with the region addressed by the first pulse.

We minimise this movement in two ways. The primary way is to make the protocol fast: as discussed in Section 4.4.2, a good 'rule of thumb' is to ensure the protocol duration is less than the healing time $t_\xi = \xi/c$, where c is the speed of sound in the condensate.[2] This still results in some movement, so we further reduce the effect by ensuring that the second magnetic field gradient has the reverse sign to the first gradient. This means that the magnetic field gradient imparts no net impulse to the condensate over the duration of the protocol.

[2] While we didn't invent the term *healing time*, it is seldom used in BEC literature. A similar definition, $\sqrt{2}$ smaller than t_ξ here, is given in Reference 91

4.1.1 Hyperbolic secant pulses

When designing the coupling pulse, we must consider both the pulse duration and the resulting spin projection $F_z(z) \equiv \langle \hat{F}_z(z) \rangle$ of the transferred region. In MRI literature, the resulting spin projection is called the *slice profile* [90]. Since the spatial selectivity of the coupling pulse is provided by a magnetic field gradient, we can describe the slice profile in either the spatial or frequency domain. We work in the frame rotating at the instantaneous radiation frequency ω, for which the detuning $\Delta(z, t)$ is given spatial dependence by the magnetic field strength $B(z)$,

$$\Delta(z, t) = \omega(t) - \gamma B(z) \tag{4.2}$$

$$= \omega(t) - \gamma \left(B_0 + \frac{\mathrm{d}B}{\mathrm{d}z} z \right) , \tag{4.3}$$

where $\gamma = \mu_B |g_F| / \hbar$ is the gyromagnetic ratio and $B_0 = B(z = 0)$ is the magnetic field offset.

It may seem as though the ideal coupling pulse for magnetic resonance control would be one that produces a slice profile as close to a top-hat as possible, to maximise spatial resolution by making the slice edge very sharp. The catch is that this would require a long pulse, and we need to keep the duration of the protocol below the healing time. In addition, we don't want the slice sharpness to be infinitely sharp; we require the slice sharpness to be on the order of the healing length.[3] Moving to the opposite extreme, the fastest pulse that inverts the condensate spin is a π-pulse: a single frequency, resonant, and constant amplitude pulse with duration $t_p = \pi/\Omega$, where Ω is the Rabi frequency. Unfortunately, a π-pulse produces an oscillatory slice profile (close to a sinc function) which does not allow precise spatial selectivity.[4]

[3] We specify this requirement in Section 4.1.2.

[4] Nevertheless, the applications of π-pulses in magnetic resonance control are explored experimentally in Chapter 6.

A good compromise between having a short pulse duration and a sharp slice profile can be achieved using a *hyperbolic secant pulse*; an adiabatic coupling pulse used extensively in MRI. In a hyperbolic secant pulse, the amplitude and detuning of the coupling are both time-dependent. The Rabi frequency and detuning are commonly expressed in the form

$$\Omega(t) = \Omega_0 \operatorname{sech}\left(\beta \left(t - \frac{t_p}{2} \right) \right) , \tag{4.4a}$$

$$\Delta(t, z) = \Delta_0 \tanh\left(\beta \left(t - \frac{t_p}{2} \right) \right) + \Delta_1(z) , \tag{4.4b}$$

$$t \in \left[\frac{-t_p}{2}, \frac{t_p}{2} \right] , \tag{4.4c}$$

where Ω_0 is the peak Rabi frequency, Δ_0 is the peak detuning, β is the sweep rate of the detuning, t_p is the pulse duration and Δ_1 is the offset of the detuning from resonance;

$$\Delta_1(z) = -\gamma \frac{\mathrm{d}B}{\mathrm{d}z} (z - z_0) . \tag{4.5}$$

By comparing Equations 4.2, 4.4b and 4.5, we can see that the modulation in detuning can be affected by changing either the radiation frequency ω, or the magnetic

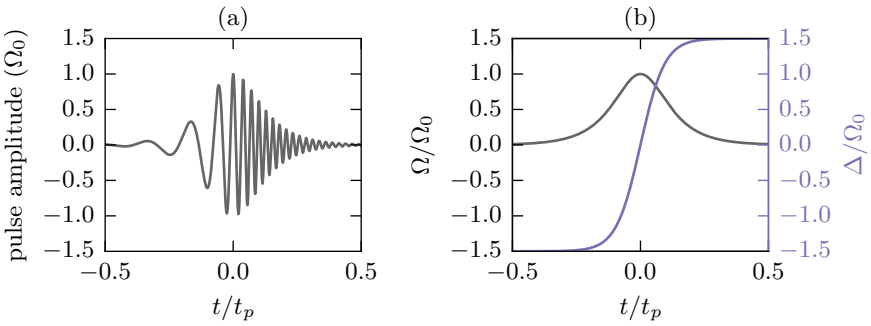

Figure 4.3: An example of a hyperbolic secant pulse. (a) The amplitude of a hyperbolic secant pulse over time. The envelope of the coupling oscillation smoothly increases, then decreases, in the form of a sech function, while the frequency sweeps through resonance. (b) The same hyperbolic secant pulse, with the Rabi frequency $\Omega(t)$ and detuning from resonance $\Delta(t)$ shown separately.

field offset B_0, or both. An example of a hyperbolic secant pulse is shown in Figure 4.3.

4.1.2 *Pulse requirements to engineer a black soliton*

We define the resolution of the spatially-dependent transfer effected by a single pulse to be

$$R = \frac{\Delta z}{\delta z}\,, \tag{4.6}$$

where Δz is the slice thickness and δz is the slice sharpness. For consistency with MRI literature, we define the slice thickness to be the full width at half maximum (FWHM) and the slice sharpness to be the 10% – 90% rise distance of the population transfer $P_{\uparrow,\,\text{final}}$ (see Equation 4.18).

To engineer a single soliton (by addressing one half of the condensate) the slice thickness needs to span at least one Thomas–Fermi radius z_{TF} (see section 4.2.1),

$$\Delta z = \zeta\, z_{\text{TF}}\,, \tag{4.7}$$

where ζ is a constant which adjusts the thickness of the slice. For the condensates we simulate later in this chapter, it is sufficient to set $\zeta = 1.2$, such that when one side of the slice is at the centre of the condensate, the other side of the slice is beyond the extent of the condensate.

To carve the shape of a black soliton into the condensate density, the slice sharpness must approximately equal the FWHM of the black soliton. To clarify this statement, a black soliton in an otherwise homogeneous condensate of density n_0 has a density profile [33] of

$$n(z) = n_0 \tanh^2\left(\frac{z}{\sqrt{2}\,\xi}\right)\,, \tag{4.8}$$

with a corresponding FWHM of

$$w_{\text{sol}} = 2\sqrt{2}\,\tanh^{-1}\left(\frac{1}{\sqrt{2}}\right)\xi \simeq \frac{5\,\xi}{2}\,. \tag{4.9}$$

By observing the shape of the density notch formed after two hyperbolic secant pulses[5] we can observe that the requisite slice sharpness is

$$\delta z \simeq \frac{5\, w_{\mathrm{sol}}}{4} \simeq 3\,\xi\,. \tag{4.10}$$

[5] From analytic expressions given in Reference 88.

Equations 4.7 and 4.10 allow us to express the required resolution of the pulse to be

$$R \simeq \frac{\zeta\, z_{\mathrm{TF}}}{3\,\xi}\,. \tag{4.11}$$

Additionally, we require the pulse duration to be

$$t_p \leq \frac{t_\xi}{4}\,, \tag{4.12}$$

to avoid detrimental movement of the condensate during the pulse. This requirement is explored in more detail in Section 4.4.2.

4.1.3 Reparametrising the hyperbolic secant pulse

When the hyperbolic secant pulse is expressed in the form given in Equations 4.4a and 4.4b, it is not easy to see how to design the pulse to satisfy the pulse resolution requirement (Equation 4.11) and the the pulse duration requirement (Equation 4.12). To develop a method of designing the pulse, we first reparametrise Equations 4.4a and 4.4b using the dimensionless parameters μ, Γ and α, which each have a distinct effect on the shape of the pulse, as defined below.

The range of the frequency sweep, $2\,\Delta_0$, determines the *slice thickness* Δz: the full width at half maximum (FWHM) of the transferred slice. From Equation 4.5, the slice thickness can be expressed as

$$\Delta z = \frac{2\,\Delta_0}{\gamma\,\frac{\mathrm{d}B}{\mathrm{d}z}}\,. \tag{4.13}$$

We can replace Δ_0 with the *normalised pulse bandwidth*[6] μ, where

$$\mu = \frac{\Delta_0}{\Omega_0}\,. \tag{4.14}$$

[6] Using the symbol μ for the normalised pulse bandwidth—the convention of much magnetic resonance literature—is unfortunate given that this symbol is also used to represent the chemical potential of the condensate. To avoid confusion, in this thesis I use μ_{cp} for the chemical potential.

We must ensure that the pulse is adiabatic, otherwise the spin projection at the centre of the slice will be $F_z < 1$ and there will be some population left behind in the original spin state. The criterion for a pulse to be adiabatic [92, p. 195] is $\Omega_0 \gg \sqrt{\Delta_0\,\beta}$, which using Equation 4.14 can be expressed as $\Omega_0 \gg \mu\,\beta$. We thus introduce the *adiabaticity parameter* Γ, where

$$\Gamma = \frac{\Omega_0}{\mu\,\beta}\,, \tag{4.15}$$

and know that the pulse will be adiabatic for $\Gamma \gg 1$ (this statement is refined in Section 4.1.4). The effect of Γ on the slice profile is shown in Figure 4.4(b). The value of Γ determines the peak spin projection at the centre of the slice, and the resolution (ratio of slice thickness to slice sharpness) of the slice.

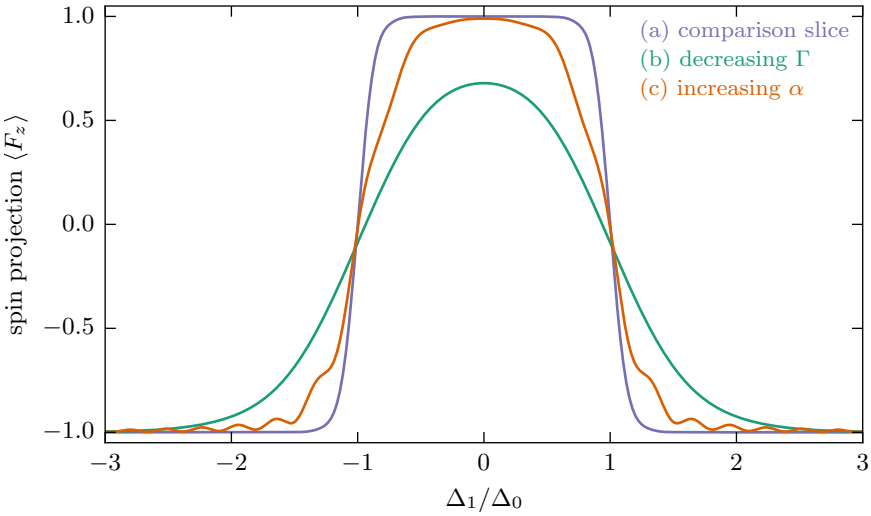

Figure 4.4: The slice profile in frequency space for different hyperbolic secant pulse parameters. For all three curves, the normalised pulse bandwidth was set to $\mu = 1$. (a) Here the adiabaticity parameter is $\Gamma = 5$ and the pulse truncation is $\alpha = 0.003$. This produces a smooth, moderately rectangular slice profile; a shape suitable for magnetic resonance control. (b) Here we reduce the adiabaticity of the pulse by setting $\Gamma = 1$ while keeping all other parameters the same. Even though the slice profile remains smooth, the spin projection at the centre of the slice no longer reaches 1. (c) Here we reduce the smoothness of the slice profile by setting $\alpha = 0.2$. While the spin projection at the centre of the slice is still high, detrimental 'wriggles' appear in the slice profile which reduce the spatial selectivity of the slice.

The *pulse truncation* α is the ratio between the initial Rabi frequency (from Equation 4.4a, $\Omega(-t_p/2) = \Omega_0 \operatorname{sech}(\beta t_p/2)$) and the peak Rabi frequency Ω_0, and can be expressed in the form

$$\alpha = \operatorname{sech}\left(\frac{\beta t_p}{2}\right). \tag{4.16}$$

The pulse truncation determines the smoothness of the slice profile, as shown in Figure 4.4(c).

The spin projections shown in Figure 4.4 were calculated by numerically solving the time-dependent Schrödinger equation for a two-state system using the Rabi Hamiltonian

$$H = \frac{\hbar}{2}\begin{pmatrix} \Delta(z,\, t) & \Omega(t) \\ \Omega(t) & -\Delta(z,\, t) \end{pmatrix}. \tag{4.17}$$

This is equivalent to considering a line of spins fixed in position along the z-axis. In Section 4.3 we observe that this simple calculation shows excellent agreement with mean-field simulations of a three-level condensate provided that the pulse duration is sufficiently short, such that the condensate does not move significantly during the pulse.

4.1.4 Selecting the correct set of pulse parameters

Now that we have defined the dimensionless parameters μ, Γ, and α, we can discuss how to select these pulse parameters to satisfy the pulse resolution and pulse duration requirements (Equations 4.11 and 4.12, respectively). In the limit of an infinitely long hyperbolic secant pulse applied to a line of stationary spins (rather than a fluid condensate), there is an analytic expression for the fractional population transferred from the spin-down state $|\downarrow\rangle$ to the spin-up state $|\uparrow\rangle$ [88]:

$$P_{\uparrow,\,\text{final}} = \frac{\cosh\left(\pi\,\Gamma\,\mu^2\right) - \cosh\left(\pi\,\Gamma\,\mu\,\sqrt{\mu^2-1}\right)}{\cosh\left(\pi\,\Gamma\,\mu^2\right) + \cosh\left(\pi\,\Gamma\,d\,\mu^2\right)}\,, \tag{4.18}$$

where $d = \Delta_1/\Delta_0$ is the normalised detuning offset. We define the pulse fidelity as the maximum fractional population transferred to the final Zeeman state at the centre of the slice profile after application of the pulse, which we estimate by setting $d = 0$ in Equation 4.18. In other words, the hyperbolic secant pulse transfers population most efficiently when the frequency sweep crosses resonance simultaneously with the maximum coupling amplitude;

$$\begin{aligned} P_{\uparrow,\,\text{adiabatic}} &\equiv \max{}_d\left(P_{\uparrow,\,\text{final}}\right) && \text{(4.19a)}\\ &= 1 - \cosh^2\left(\frac{\pi}{2}\,\Gamma\,\mu\,\sqrt{\mu^2-1}\right)\operatorname{sech}^2\left(\frac{\pi}{2}\,\Gamma\,\mu^2\right) && \text{(4.19b)}\\ &\approx 1 - e^{-\pi\,\Gamma/2}\,. && \text{(4.19c)} \end{aligned}$$

The approximation in Eq. 4.19c is accurate to within 1% for $\mu > 2$ and $\Gamma > 2$ (and is exact in the limit of $\mu \to \infty$), and thus Γ serves as a good adiabaticity parameter, analogous to the adiabaticity parameter in the Landau–Zener probability [93, 94].

We can estimate the pulse resolution R by rearranging Equation 4.18 to find the normalised detuning offset for a given fractional population transferred to the spin-up state:

$$\begin{aligned} d(P_{\uparrow,\,\text{final}}) = \frac{1}{\pi\,\Gamma\,\mu^2}\cosh^{-1}\Bigg(&\left(P_{\uparrow,\,\text{final}}^{-1} - 1\right)\cosh\left(\pi\,\Gamma\,\mu^2\right)\\ &- P_{\uparrow,\,\text{final}}^{-1}\cosh\left(\pi\,\Gamma\,\mu\,\sqrt{\mu^2-1}\right)\Bigg)\,. \end{aligned} \tag{4.20}$$

The single-pulse resolution in the limit of an infinitely long pulse is then

$$R(\mu,\,\Gamma) = \frac{\Delta z}{\delta z} = \frac{2\,d(0.5)}{d(0.9) - d(0.1)}\,. \tag{4.21}$$

In this form, Equation 4.21 does not enable immediate insight into how the pulse parameters determine the resolution of the pulse. However, when $\Gamma \geq 3$ and $\mu \geq 1$, Equation 4.21 is well approximated by

$$R \approx \sqrt{2}\,\Gamma\,\mu^2\,, \tag{4.22}$$

agreeing with Equation 4.21 to within 1%.

Truncating the hyperbolic secant pulse to have a finite duration t_p results in non-adiabatic, off-resonant Rabi oscillations at the beginning and end of the pulse,

rendering Equations 4.18–4.19b inexact. Approximating the coupling at the boundaries of the sweep to be that of an unmodulated off-resonant pulse (which is reasonable as both the frequency and amplitude of the hyperbolic secant pulse change most slowly at the boundaries), the off-resonant oscillations in $P_{\downarrow,\text{final}}$ are given by

$$P_{\downarrow,\text{asymptotic}} \equiv \frac{\Omega(t=t_p)^2}{\Omega(t=t_p)^2+\Delta(t=t_p)^2} = \frac{\alpha^2}{\alpha^2+\mu^2\left(d+\sqrt{1-\alpha^2}\right)} . \quad (4.23)$$

These off-resonant Rabi oscillations can not only perturb the resonant pulse fidelity, but also lower the resolution of the pulse if α is not chosen to be sufficiently low (as demonstrated in Figure 4.4). A lower value of α reduces the amplitude of these oscillations, but also increases the duration of the pulse. As such, the trick is to predict how low α needs to be to not significantly degrade the adiabatic-limited resolution given in Equation 4.21. We have found that the resolution limit imparted by the finite pulse duration is related to the width of the curve given by Equation 4.23 as a function of normalised detuning d. This is a Lorentzian curve with a FWHM of $2\,\alpha/\mu$ centred at $d_0 = -\sqrt{1-\alpha^2}$. The resolution of a finite-duration pulse agrees to within 1% of an infinite pulse if

$$\alpha < \frac{1}{\sqrt{2}\,\eta\,\mu\,\Gamma} , \quad (4.24)$$

where $\eta = 30$ typically characterises this resolution limit.

The above considerations motivate the following method for selecting the pulse parameters to engineer a black soliton in a condensate:

1. Given a desired resonant fidelity $P_{\uparrow,\,\text{adiabatic}} = 1 - P_{\downarrow,\,\text{adiabatic}}$, the adiabaticity parameter is well-approximated by $\Gamma \approx -2/\pi \ln P_{\downarrow,\,\text{adiabatic}}$. For example, to obtain a fidelity of 99% (adequate for many experiments), choosing $\Gamma = 3$ is more than sufficient.
2. With Γ chosen, the normalised detuning μ can be obtained using Equation 4.22.
3. With μ selected, we can calculate α using Equation 4.24.
4. The last step is to set the peak Rabi frequency Ω_0. From Equations 4.15 and 4.16, the minimum permissible pulse area is

$$(\Omega_0\,t_p)_{\text{min}} = 2\,\Gamma\,\mu\,\cosh^{-1}\left(\alpha^{-1}\right) . \quad (4.25)$$

Recall that for a given condensate we require $t_p < t_\xi/4$ for the quantum fluid to remain stationary during the pulse sequence. Substituting $t_p = t_\xi/4$ into Equation 4.25 sets the required peak Rabi frequency Ω_0:

$$\Omega_0 = \frac{8\,\Gamma\,\mu\,\cosh^{-1}\left(\alpha^{-1}\right)}{t_\xi} . \quad (4.26)$$

It is important to note that the method outlined in this section is not the only way to find an adequate set of pulse parameters for a given condensate. Equation 4.22 shows that you could set Γ to be larger than necessary (thereby giving your

pulse a higher fidelity than required) and compensate by reducing μ to achieve the same pulse resolution. This increases the required effective pulse area $\Omega_0\, t_p \propto \Gamma\, \mu$, but otherwise results in the same slice profile. Indeed the results presented in the rest of this chapter use $\Gamma > 3$, as the above method was established after I had already simulated MRC to engineer a black soliton, and prepared the main results of this chapter. That a different set of pulse parameters can also produce a black soliton speaks to the robustness of this technique.

4.1.5 Required magnetic field gradient

The equations derived in this chapter lead to an alternate expression for the required magnetic field gradient in terms of the healing length. By combining Equations 1.1, 4.13, 4.14, 4.11, 4.22, and 4.26, we arrive at the equation

$$\left|\frac{\mathrm{d}B}{\mathrm{d}z}\right| \approx \frac{8\,\hbar\,\cosh^{-1}\left(\alpha^{-1}\right)}{3\,\gamma\,m\,\xi^3}\,. \tag{4.27}$$

This shows that the magnetic field required by MRC to engineer a black soliton in a condensate depends only on the healing length of the condensate, without dependence on the Thomas–Fermi radius of the condensate, which may seem counter-intuitive. This equation also allows us to estimate the required magnetic field gradient for a particular condensate before finding appropriate pulse parameters; the 3rd order scaling with healing length shows that the magnetic field gradient needed will increase dramatically for very small healing lengths. For example, setting $\alpha = 0.003$, if a ^{87}Rb condensate has a healing length of $500\,\mathrm{nm}$, the required magnetic field gradient will be on the order of $250\,\mathrm{G/cm}$. If the healing length is decreased to $150\,\mathrm{nm}$, the required magnetic field gradient dramatically increases to the order of $10\,\mathrm{kG/cm}$

4.1.6 Gradient-induced motion

An aspect of condensate motion we have not considered in detail in the above analysis is the Stern–Gerlach effect; the fact that the magnetic field gradient will accelerate the different spin components of the condensate in different directions. This acceleration is given by

$$a_{\mathrm{SG}} = -\frac{\hbar\,\gamma\,m_F}{m}\,\frac{\mathrm{d}B}{\mathrm{d}z}\,, \tag{4.28}$$

and during each pulse the $|{+1}\rangle$ and $|{-1}\rangle$ components will move in opposite directions by the amount

$$\delta z_{\mathrm{SG}} = \frac{1}{2} a_{\mathrm{SG}}\, t_p^2\,. \tag{4.29}$$

To keep this movement negligible during the pulse, we require $|\delta z_{\mathrm{SG}}| < \xi$. Using Equation 4.28, setting $t_p = t_\xi/4$, this requirement can be expressed in the form

$$\left|\frac{\mathrm{d}B}{\mathrm{d}z}\right| < \frac{16\,\hbar}{\gamma\,m\,\xi^3}\,. \tag{4.30}$$

This cubic scaling between magnetic field gradient and length-scale of the achieved resolution is the same relationship previously found in the resolution limits of atomic position measurements of laser-cooled gases and atomic beams [95].

We have found that the upper bound for the magnetic field gradient in Equation 4.30 can be close to the gradient required in Equation 4.27. This means that to be certain of engineering a black soliton, we need to simulate the condensate via the Gross–Pitaevskii equation (Section 4.2) which includes this gradient-induced motion. We discuss in Section 4.4.3 that if this effect is small, it can be mitigated through appropriate choice of the sign of the magnetic field gradient. As the healing length of the condensate decreases, and hence the required magnetic field gradient increases, the gradient-induced motion will become more significant. In these cases, it might be necessary to use a pulse duration shorter than $t_p = t_\xi/4$.

4.1.7 *Relationship to the proposal by Williams and Holland*

After developing our magnetic resonance control technique, we realised that it seemed reminiscent of a wavefunction engineering proposal made by J. E. Williams and M. J. Holland in 1999 [96], which was used to create the first vortex in a Bose-Einstein condensate [97], and later an unstable soliton [98]. After these examples this technique did not become commonly used, perhaps because simpler methods were soon discovered for creating vortices in condensates; for example, stirring the condensate with a repulsive potential [99], or coupling angular momentum from a LagurreGaussian beam [42, 44]. As a result it did not develop a short, descriptive name; in order to compare and contrast this technique with magnetic resonance control, we will refer to it as the Williams–Holland technique.

Both magnetic resonance control and the Williams-Holland technique establish an effective two-level system and create a spatially-varying splitting between these two levels to control the density and phase of the condensate with spatial resolution. However, this is the extent of the similarity between these two techniques.

The Williams-Holland technique varies the level splitting over time in an oscillatory manner, and uses constant coupling between the levels. In contrast, magnetic resonance control modulates the coupling while using a static level splitting.

Unlike magnetic resonance control, the William-Holland technique does not achieve healing-length resolution, instead creating larger structures which then relax into the target state at a rate governed by the underlying dynamics of the uncoupled system. Achieving a finer spatial resolution using the Williams-Holland technique would require weaker coupling between the two levels, necessarily resulting in a slower protocol compared to magnetic resonance control; one advantage of the sub-healing-time speed of magnetic resonance control is that it enables us to engineer the condensate wavefunctions multiple times during the course of an experiment, near-instantaneously.

Williams and Holland proposed that their technique could be used to create wavefunctions less symmetric than a single vortex or soliton by applying time- and spatially-varying level splittings with a profile more complicated than a linear gradient. Achieving this experimentally would usually require either microstructured current elements to produce such magnetic field gradients, or alternatively a spatial light modulator to produce an effective magnetic field gradient using a light-shift.

Both these options add to the experimental complexity of the technique, and potentially impose constraints on the resolution of the technique. We show in Section 4.5 that we can use magnetic resonance control to engineer the less symmetric wavefunctions of multiple off-centre solitons by retaining the simple linear gradient and instead varying the profile of the time-dependent coupling; such pulses are readily achievable with agile radiofrequency synthesisers.

4.2 Simulating a one-dimensional spin-1 condensate

The Gross–Pitaevskii model [100, 101] is used ubiquitously to study the evolution of Bose–Einstein condensates in the mean-field approximation [102,103]. A derivation of this model applied to a two-level system can be found in Reference 104, with the extension to spin-1 systems in Reference 20.

To demonstrate our magnetic resonance control protocol, we simulate a spin-1 Bose–Einstein condensate with three Zeeman sublevels which are coupled by magnetic dipole transitions in a time-dependent magnetic field gradient. The condensate is elongated, held in a cylindrically symmetric potential with radial and axial trap frequencies $\omega_r \gg \omega_z$. In such a condensate, the relevant dynamics occur in the direction of weakest confinement; the axial direction z. This allows us to use a quasi-1D Gross–Pitaevskii equation (GPE) to reduce the number of simulated spatial dimensions, applying a Thomas–Fermi ansatz to the density along the radial direction [105].

This quasi-1D simulation better represents experimental reality than a strictly one-dimensional GPE because it considers the finite radial extent of the condensate. For example, a 1D GPE simulation predicts a parabolic single-component ground state, which differs to the profile observed in an elongated condensate; the lower density at the axial extent of the condensate rounds the sharp parabolic boundary. Such a profile is better predicted by the quasi-1D simulation described below.

In these simulations we model a spin-1 system instead of the two-level system described previously in this chapter. We choose a spin-1 system because it enables a higher rf Rabi frequency than a pseudospin-1/2 condensate with two-photon coupling. For a spin-F system starting in $|F,\, m_F = -F\rangle$, all analytics discussed previously in this chapter map through the relation

$$\frac{\langle F_z \rangle}{F} = 2\, P_{\uparrow,\,\text{final}} - 1 \tag{4.31}$$

since spin-rotations of the polarised state $|F,\, m_F = \pm F\rangle$ map to the Bloch-sphere representation of spin-1 systems [94].

4.2.1 Spin-1 quasi-1D Gross–Pitaevskii equation with magnetic resonance

A cylindrically symmetric spin-1 BEC has a spinor order parameter

$$\mathbf{\Psi}(r,\, z) = \begin{pmatrix} \psi_{-1}\,(r,\, z) \\ \psi_0\,(r,\, z) \\ \psi_{+1}\,(r,\, z) \end{pmatrix}. \tag{4.32}$$

Applying the Thomas–Fermi ansatz described in Reference 105, the wavefunction of the Zeeman component is

$$\psi_{m_F}(r,\, z) = \Phi_\perp(r,\, \chi(z)) f_{m_F}(z)\,, \tag{4.33}$$

where $\Phi_\perp(r,\, \chi(z))$ represents the radial dependence of the condensate density, common to all Zeeman components. This radial density is a Thomas–Fermi profile, with an axially-varying radius $\chi(z)$;

$$\Phi_\perp(r,\, \chi(z)) = \sqrt{\frac{2}{\pi}}\, \frac{1}{\chi(z)} \max\left(1 - \frac{r^2}{\chi(z)^2},\, 0\right)^{1/2}. \tag{4.34}$$

The fields are normalised such that

$$\int |\psi_{m_F}(r,\, z)|^2\, dV = \int_{-\infty}^{\infty} |f_{m_F}(z)|^2\, dz = N_{m_F}\,, \tag{4.35}$$

where N_{m_F} is the number of atoms in the Zeeman state m_F. The 3D atomic density (with units of atoms/m^3) of the m_F component is

$$n_{m_F}(r,\, z) \equiv |\psi_{m_F}(r,\, z)|^2 \tag{4.36}$$

while the *linear* density (with units of atoms/m) of the m_F component is

$$\rho_{m_F}(z) \equiv |f_{m_F}(z)|^2\,. \tag{4.37}$$

From the fields $f_{m_F}(z)$ and axially-dependent radial width $\chi(z)$ we can calculate experimentally relevant quantities such as the peak atomic density n_0, the average density $\langle n \rangle$, and the column density $\tilde{n}(y,\, z) = \int_{-\infty}^{\infty} n(r = \sqrt{x^2 + y^2},\, z)\, dx$. For example, the peak atomic density is

$$n_0 = \max{}_z \left(\frac{2\,\rho(z)}{\pi\,\chi(z)^2}\right), \tag{4.38}$$

where $\rho(z) \equiv \sum_{m_F} \rho_{m_F}(z)$ is the total linear density.

The spinor $\mathbf{f}(z,\, t) = (f_{-1},\, f_0,\, f_{+1})^T$ obeys a coupled quasi-1D GPE

$$i\,\hbar\,\frac{\partial \mathbf{f}}{\partial t} = \left(-\frac{\hbar^2}{2m}\,\frac{\partial^2}{\partial z^2} + V + E_\perp + c_0\,\eta\,\rho + S + H_C\right)\mathbf{f}\,, \tag{4.39}$$

where m is the atomic mass, $V(z) = m\,\omega_z^2\, z^2/2$ is the spin-independent external potential along the axial direction, $E_\perp(z) = m\,\omega_r^2\,\chi(z)^2/6$ is the transverse mode energy, $c_0 = 4\,\pi\,\hbar^2\,(2\,a_2 + a_0)\,/3\,m$ is the spin-independent interaction strength, $\eta(z,\, t) = 4/\left(3\,\pi\,\chi(z)^2\right)$ is the scaling factor (with units of m^{-2}), $\rho(z,\, t)$ is the total linear density, $S(z,\, t)$ is the spin-exchange operator and $H_C(z,\, t)$ is the linear-coupling Hamiltonian (both defined below).

The spin-exchange operator which governs the spin-spin interactions is

$$S = c_2\,\eta \begin{pmatrix} \rho_0 + \rho_{-1} - \rho_{+1} & f_{+1}^*\, f_0 & 0 \\ f_{+1}\, f_0^* & \rho_{-1} + \rho_{+1} & f_{-1}\, f_0^* \\ 0 & f_{-1}^*\, f_0 & \rho_0 + \rho_{+1} - \rho_{-1} \end{pmatrix}, \tag{4.40}$$

where $c_2 = 4\,\pi\,\hbar^2\ (a_2 - a_0)\,/3\,m$ is the spin-dependent interaction strength. The z-dependent radial width $\chi(z)$ at any given time is determined by an auxiliary differential equation arising from the Lagrangian formulation of the quasi-one dimensional GPE [105]:

$$\rho\,\frac{\partial E_\perp}{\partial \chi} + \frac{1}{2}\left(c_0\,\rho^2 + c_2\,S_2\right)\frac{\partial \eta}{\partial \chi} = 0\,, \tag{4.41}$$

where

$$S_2 = (\rho_{-1} - \rho_{+1})^2 + 2\,\rho_0\,(\rho_{-1} + \rho_{+1}) + 2\,f_0^{*2}\,f_{+1}\,f_{-1} + 2\,f_{+1}^{*}\,f_{-1}^{*}\,f_0^2\,. \tag{4.42}$$

For the Thomas–Fermi ansatz applied here, this reduces to an algebraic equation for χ in terms of the one-dimensional fields f_{m_F} and the densities ρ_{m_F}:

$$\chi^2 = \frac{2}{\omega_r}\sqrt{\frac{c_0\,\rho^2 + c_2\,S_2}{\pi\,m\,\rho}}\,. \tag{4.43}$$

The Hamiltonian for linear coupling between the states $|m_F = -1,\ 0,\ +1\rangle$ is

$$H_C = \hbar\begin{pmatrix} -\Delta & \frac{\Omega}{\sqrt{2}} & 0 \\ \frac{\Omega}{\sqrt{2}} & q & \frac{\Omega}{\sqrt{2}} \\ 0 & \frac{\Omega}{\sqrt{2}} & \Delta \end{pmatrix}, \tag{4.44}$$

where Ω is the Rabi frequency for transitions between the three spin-1 Zeeman states, and $q = (E_{-1} + E_{+1} - 2\,E_0)/2\,\hbar$ is the quadratic Zeeman shift.

The spin-exchange terms in Equations 4.39 and 4.43 (those proportional to c_2) vanish in the absence of a superposition of Zeeman states, and are negligible during the pulses. For example, $c_2\langle n\rangle$ is around 10 Hz for $F = 1$ ^{87}Rb condensates with typical densities of $\langle n\rangle = 10^{14}$ atoms/cm^3.

To perform our quasi-1D GPE simulations I used our 'in-house' simulation engine *pygpe* developed by Martijn Jasperse, which I extended to include time-dependent magnetic fields and time-dependent coupling. This simulation engine uses a Python front-end combined with C and FORTRAN to increase computational speed. It numerically solves the Gross–Pitaevskii equation on a regular spatial lattice using a split-step Fourier method; the kinetic energy term is integrated in Fourier space first (using the FFTW library), followed by integrating the potential energy term in real space. We use the high-performance numerical FORTRAN library EXPOKIT for matrix exponentiation of the spin-exchange operator at each point in space. Further details of the computational methods used are discussed in Martijn Jasperse's PhD confirmation report [106].

4.2.2 Choosing the space step and time step

Simulating magnetic resonance control needs a spatial grid with small spacing z_{step} and small time steps t_{step}. We seek to resolve features at the healing length scale, requiring $z_{\text{step}} < \xi/2$, and since the hyperbolic secant pulses occur over a time scale smaller than the healing time, we also require $t_{\text{step}} \ll t_\xi$ during the protocol. If

the space steps or time steps (or both) are too large, the simulated wavefunction degenerates into meaningless noise after a few thousand time steps. The space step needs to be sufficiently small to sample the soliton. Once the space step is chosen, the time step needs to be sufficiently small so that the superfluid flow during one time step is much less than one space step, and so that the modulation of the coupling pulse is sufficiently sampled. But we do not want to choose unnecessarily small space steps and time steps, since this increases the computation time.

A good way to find the largest permissible space step and time step is to simulate an analytically-defined black soliton for one trap period, and check that the soliton remained stationary. This analytically-defined black soliton is the numerically-determined ground state wavefunction multiplied by the analytic expression for a soliton in a homogeneous density profile (equation 1.3). I use this term to distinguish this from a soliton engineered into the condensate using our magnetic resonance control protocol.

4.2.3 *Confirming the physicality of our simulations*

Before simulating magnetic resonance control, I simulated condensates under a variety of familiar conditions to ensure that the simulations were producing physical results.

I checked that in the absence of magnetic field gradients and coupling, the ground-state finder was able to converge to the lowest energy state. The ground-state finder in pygpe uses imaginary time propagation [107] until the change in energy after each time step reaches a defined threshold value.

I confirmed that constant resonant radiofrequency coupling produced Rabi oscillations between the Zeeman states, and that a hyperbolic secant pulse (in the absence of a magnetic field gradient) could transfer the condensate population entirely from $|-1\rangle$ to $|+1\rangle$ (and vice versa).

By pulsing a strong gradient on and off, I observed the sloshing mode of a harmonically trapped single-component condensate. I also checked that when a constant magnetic field gradient was applied, a condensate in the $|-1\rangle$ Zeeman state moved towards smaller $|B|$ while a condensate in the $|+1\rangle$ Zeeman state moved towards larger $|B|$. This is consistent with the weak and strong field seeking behaviour of these states, respectively.

In the final stages of confirming the physicality of my simulations, I had the valuable opportunity to compare my results with collaborators Joseph Hope and Paul Wigley, who were using the simulation environment XMDS2 [108] to simulate the same system. The simulations in XMDS2 use an adaptive time step method with the algorithm ARK45, which is a fourth order 'almost Runge–Kutta' integration technique with five stages [109].

Setting the initial state of the condensate to contain a soliton[7] was another way to confirm that the simulation produced expected results. For example, we confirmed that a dark soliton[8] oscillates within the condensate with a frequency of $\sqrt{2}$

[7] By defining the initial density and phase using the equation for a soliton (Equation 1.3) multiplied by the lowest energy state.

[8] The terminology of 'black', 'dark' and 'grey' solitons is defined in Section 1.2.

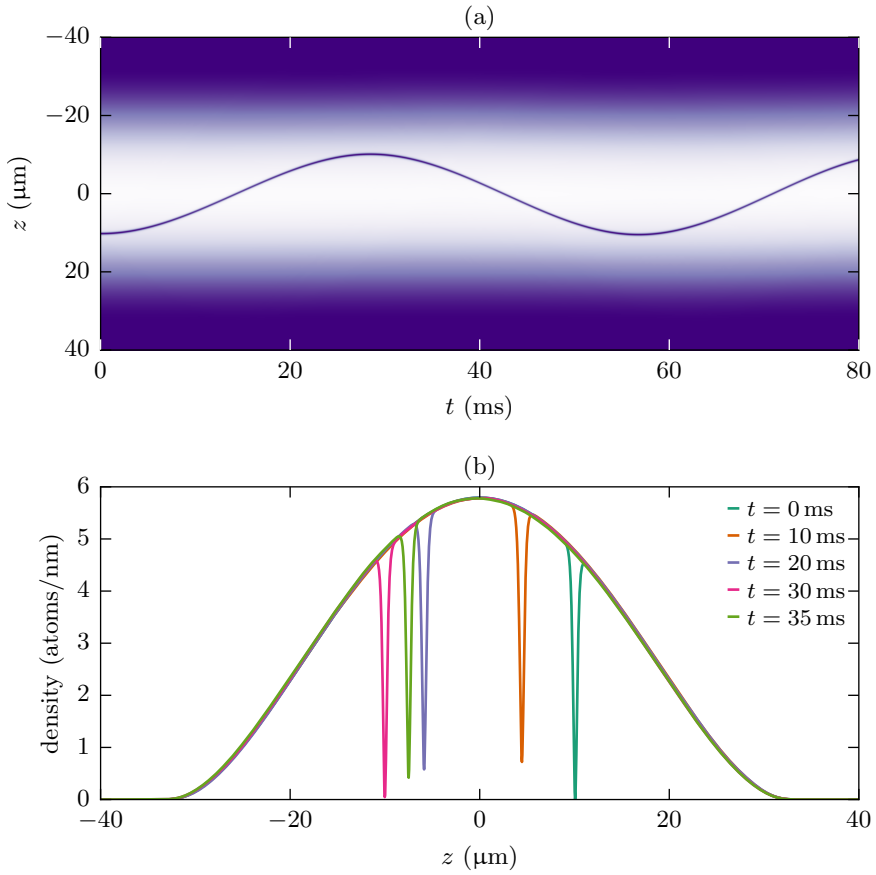

Figure 4.5: Simulating an oscillating soliton, initially located at $z = 10\,\mu\text{m}$. Here the trap frequencies are $f_z = 24\,\text{Hz}$ and $f_\rho = 240\,\text{Hz}$, and the condensate contains 2×10^5 atoms. (a) The soliton oscillates with a frequency of $\sqrt{2}\,f_z$ as expected [33, p. 221]. (b) The linear density of the condensate at different times shows that the turning points of the oscillation occur when the density dip reaches zero.

times the trap frequency [33, p. 221], and saw that the turning point of this oscillation occurs when the soliton density dip reaches zero (Figure 4.5).

We can observe how the stability of the soliton depends on the width of the density dip (while maintaining a π phase step). Figure 4.6 shows that as the width of the density dip is increased beyond $w_{\text{sol}} = 5\,\xi/2$, the feature decays into multiple dark solitons. Interestingly, when the width of the density dip is $w_{\text{sol}} = n\,5\,\xi/2$, where n is an odd integer, the density dip decays into n dark solitons. When n is even, $n + 1$ solitons form. This is because we have defined an initial π phase step; since phase must be conserved across the condensate, such a system can support only odd numbers of solitons.

At this stage of our investigations, we had not yet developed the criteria discussed in Chapter 5 for choosing condensate parameters (atom number and trap frequencies) that are suitable for engineering black solitons. This meant that the choice of condensate and trap parameters for the simulations presented above were

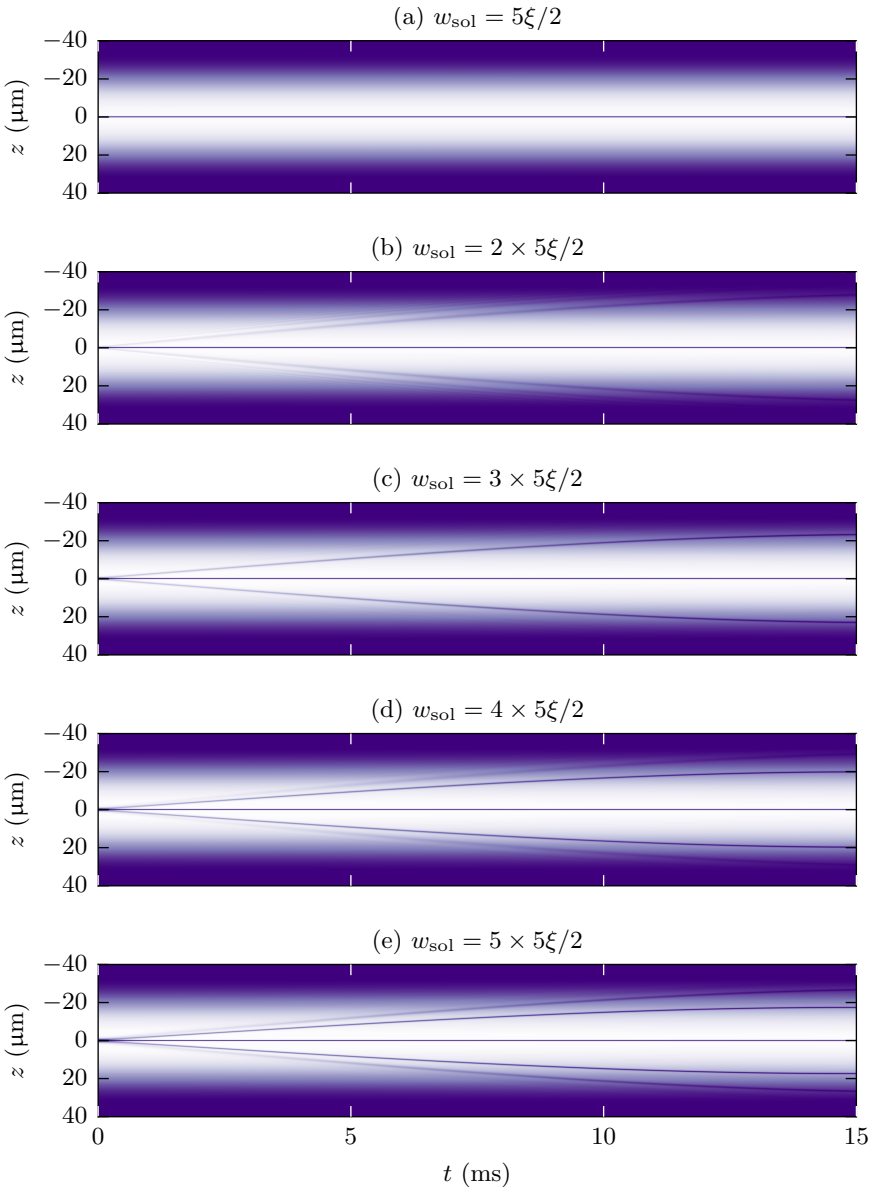

Figure 4.6: Increasing the width of the density dip causes the density dip to decay into multiple solitons. Figure 4.6(a) shows that indeed a black soliton, with width $w_{\mathrm{sol}} = 5\,\xi/2$, remains stationary at the centre of the condensate (Equation 1.3). Figures 4.6(b–e) show that as the width is increased by integer multiples n, the density dip is unstable and decays into n solitons when n is odd, and $n + 1$ solitons when n is even.

fairly arbitrary; here the trap frequencies were $f_z = 24\,\text{Hz}$ and $f_\rho = 240\,\text{Hz}$, with atom number $N = 2 \times 10^5$, similar to the condensates produced in our laboratory. For such a condensate, $\xi = 148\,\text{nm}$ and $z_{\text{TF}} = 33\,\mu\text{m}$. This necessitated a spatial grid of 2048 points, with step size $z_{\text{step}} = 43\,\text{nm}$, and time step $t_{\text{step}} = 1\,\mu\text{s}$.

4.3 Engineering a black soliton using magnetic resonance control

Here I use magnetic resonance control to simulate engineering a single black soliton in a condensate with experimentally feasible parameters. Developing this successful protocol was the culmination of a year of research, including a close collaboration in the last few months with Paul Wigley and Joseph Hope. To keep this section as transparent as possible, I wait until Section 4.4 to motivate some of the choices made in designing the protocol parameters.

4.3.1 The condensate parameters

The condensate considered for the rest of this chapter contains $N = 10^4$ ^{87}Rb atoms in a harmonic trap with axial frequency $f_z = 2.4\,\text{Hz}$ and radial frequency $f_\rho = 158.4\,\text{Hz}$. These condensate parameters satisfy all of the criteria discussed in Chapter 5 necessary for engineering a black soliton using magnetic resonance control. Using equations defined in Section 1.2, this condensate has a Thomas–Fermi radius of $z_{\text{TF}} = 96\,\mu\text{m}$, a healing length of $\xi = 504\,\text{nm}$, a soliton full width at half maximum (FWHM) of $w_{\text{sol}} = 1.3\,\mu\text{m}$, and a healing time of $t_\xi = 492\,\mu\text{s}$.

The spatial grid consists of 2048 points with a step size of $z_{\text{step}} = 195\,\text{nm}$, corresponding to a total spatial extent of $400\,\mu\text{m}$. This space step is fine enough to model a soliton in this condensate, since it corresponds to seven points across the FWHM of the soliton.

The spatial extent is deliberately much larger than the Thomas–Fermi radius of the condensate[9] because the Fourier integration technique used by our simulation engine results in periodic boundary conditions; any portion of the condensate which reaches the boundary would reappear on the other side of the simulation. To manage these non-physical boundary effects, in addition to making the spatial extent much larger that the condensate, we introduce exponential damping of the wavefunction near the edges of space.

[9] The figures in this chapter do not show the full spatial extent, instead truncating just beyond the Thomas–Fermi radius.

The time step was chosen to accommodate this step size without allowing computational error to significantly accumulate. During the protocol, with coupling pulses of shorter duration than the healing time, the time step is $t_{\text{step}} = 100\,\text{ns}$.

To check that a black soliton has been successfully engineered, we need to observe the condensate evolve over several trap periods. With a trap period of 41.7 ms, we choose to propagate up to one second after the protocol. After the second coupling pulse the time step is increased to $t_{\text{step}} = 10\,\mu\text{s}$ with no ill-effects, thereby significantly reducing the computation time (from hours to minutes).

4.3.2 The pulse parameters

For the simulations in this section, we use the following hyperbolic secant pulse parameters:

- the normalised pulse bandwidth $\mu = 3.2$,
- the adiabaticity parameter $\Gamma = 5$,
- the pulse truncation $\alpha = 0.003$,
- the peak Rabi frequency $\Omega_0 = 2\,\pi \times 300\,\mathrm{kHz}$.

These parameters produce a smooth slice profile (shown in Figure 4.7) with a peak spin projection at the centre of the slice of $F_z > 0.999$. Using Equation 4.25, the pulse duration is $t_p = 110.4$ µs, which satisfies the requirement that $t_p < t_\xi/4$. To achieve the necessary slice thickness of $\Delta z = \zeta\, z_{\mathrm{TF}}$, with $\zeta = 1.2$, from Equation 4.13 the magnetic field gradient needs to be $|\mathrm{d}B/\mathrm{d}z| = 237.5\,\mathrm{G/cm}$. In this gradient, the slice sharpness is $\delta z = 1.6$ µm, which satisfies the slice sharpness criterion (Equation 4.10).

A peak Rabi frequency of $\Omega_0 = 2\,\pi \times 300\,\mathrm{kHz}$ is experimentally feasible; because $\gamma = 702\,\mathrm{kHz/G}$ for the $F = 1$ level of ^{87}Rb, such a frequency corresponds to an oscillating magnetic field amplitude of 430 mG, which can be generated by a coil antenna 10 cm from the condensate driven by a 30 W radiofrequency amplifier.

Figure 4.7 shows the spin projection after this hyperbolic secant pulse, both from the stationary spins calculation (Equation 4.17) and from the GPE simulation. The two curves show excellent agreement, validating the simplification of using the stationary spins calculation while selecting the pulse parameters, instead of using the full GPE simulation.

The pulse parameters we have chosen here are not precisely the values obtained by following the procedure outlined in Section 4.1.4; the pulse used here is excessively adiabatic, which is compensated by a lower normalised pulse bandwidth. This is because the findings in Sections 4.1.2 and 4.1.4 were not refined until after these simulations were performed. However, you will see that these pulse parameters do still allow us to engineer a black soliton in this condensate. If we were to follow the procedure outlined in Section 4.1.4 this would allow us to use either a pulse with a shorter duration, or a lower peak Rabi frequency.

4.3.3 Experimental considerations for the magnetic field gradient

For experimentalists accustomed to manipulating spinor condensates with magnetic field gradients generated by bias coils located outside a vacuum chamber, like in our laboratory, a magnetic field gradient of order 200 G/cm seems quite strong. Fortunately, this gradient is comparable in magnitude to those used for magnetic trapping of condensates. This means that an apparatus using a magnetic trapping stage on the way to creating spinor condensates could implement this protocol, by using (comparatively weaker) bias coils to offset the field produced by large quadrupole coils. If the spinor condensate is created on an atom-chip trap, the magnetic field gradient ceases to be an issue. Much larger gradients are routinely

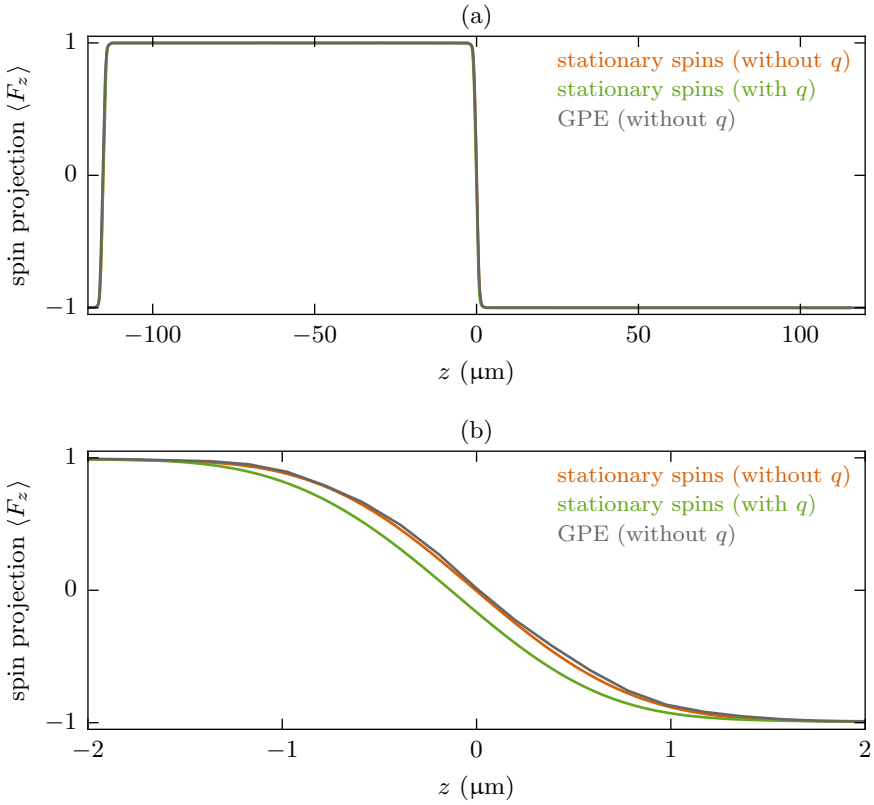

Figure 4.7: The spin projection after a hyperbolic secant pulse with pulse parameters of $\Omega_0 = 2\pi \times 300\,\text{kHz}$, $\mu = 3.2$, $\Gamma = 5$, and $\alpha = 0.003$, corresponding to a pulse duration of $t_p = 110.4\,\mu\text{s}$, in a magnetic field gradient of $|\text{d}B/\text{d}z| = 237.5\,\text{G/cm}$. Results are shown for the stationary spins calculation with no quadratic Zeeman shift (orange), with the quadratic Zeeman shift included (green), and from a simulation of the GPE (grey). (a) The spin projection across the entire condensate ($z_{\text{TF}} = 96\,\mu\text{m}$). The slice thickness has been set to $\Delta z = 1.2 z_{\text{TF}} = 115.3\,\mu\text{m}$ to ensure that the left side of the slice is well outside the condensate. (b) The spin projection near the edge of the slice profile. The slice sharpness is $\delta z = 1.6\,\mu\text{m}$. The curves from the GPE simulation and the stationary spins calculation with no quadratic Zeeman shift show such close agreement that they are difficult to distinguish on this scale. The quadratic Zeeman shift results in a horizontal offset to the slice. This offset can be compensated for by an additional detuning offset to the pulse.

created by atom-chip traps, where current-carrying wires can be tens of microns from the condensate [110].

If due to apparatus limitations the gradient cannot be generated by coils of wire, the necessary spatially-varying energy level splitting could instead be achieved optically. The vector light shift from the side of an off-resonant Gaussian laser beam can produce an effective magnetic field gradient. The deleterious state-*independent* force generated by such an intensity gradient can be made to vanish if the laser is tuned to a magic-zero wavelength [111], while still synthesising a sufficiently large vector light shift [112]. This optical approach has the added advantage of providing the ability to rapidly modulate the effective magnetic field; rapid switching using current-carrying wires can be experimentally more challenging than modulating the polarisation of a laser beam.

4.3.4 *The quadratic Zeeman shift*

We require that the magnitude of the magnetic field across the entire condensate is large enough to allow radiofrequency coupling between Zeeman states (ideally in the regime of the rotating-wave approximation, such that $\omega_L \gg \Omega_0$, where ω_L is the Larmor frequency). To achieve this, it is sufficient to set the minimum magnetic field strength (on one edge of the condensate) to be

$$|B|_{\text{min}} = 10\frac{\Omega_0}{\gamma}\,. \tag{4.45}$$

The pulse we use has $\Omega_0 = 2\,\pi \times 300\,\text{kHz}$, which requires $|B|_{\text{min}} = 4.3\,\text{G}$. For a magnetic field gradient of $237.5\,\text{G/cm}$, the maximum magnetic field strength at the other edge of the condensate (with $z_{\text{TF}} = 96\,\mu\text{m}$) is $|B|_{\text{max}} = 8.8\,\text{G}$.

The quadratic Zeeman shift in an $F = 1$ ^{87}Rb condensate [94, 113] is

$$q\,(z) = 2\,\pi \times \left(B\,(z)\right)^2 \times 71.89\,\text{Hz/G}^2\,, \tag{4.46}$$

which for this maximum magnetic field gives $q_{\text{max}} = 2\,\pi \times 5.6\,\text{kHz}$.

In Figure 4.7 we can compare the slice profile for the same pulse, with and without the presence of this spatially-dependent quadratic Zeeman shift. We see that the quadratic Zeeman shift primarily results in a horizontal offset to the slice on order q (equivalent to an offset in the detuning Δ_1). This is understandable, since the quadratic Zeeman shift changes the resonance of the sweep between $|-1\rangle$ and $|0\rangle$, and the detuning is the difference between the frequency of the pulse and the Larmor frequency. We can therefore compensate for the majority of the effect of the quadratic Zeeman shift with an appropriate offset of the detuning of each pulse.

Since q is not constant across the condensate, the slice offset is not constant; this results in a small distortion of the slice profile (as though the magnetic field gradient has a slight curvature). It would be possible to remove this distortion by using an asymmetric pulse, to deliberately skew the slice profile in the opposite direction. Alternatively, we believe a quadratic Zeeman shift of this magnitude could

be cancelled using a microwave pulse [61]. With these options in mind, for simplicity we choose to neglect the quadratic Zeeman shift term from the Hamiltonian for the simulations presented in this chapter.

4.3.5 The MRC protocol used to engineer a black soliton

The condensate starts at rest in the ground state of the harmonic trap with all population in the $|-1\rangle$ state. We apply a magnetic field gradient of $-237.5\,\mathrm{G/cm}$, and apply the first hyperbolic secant pulse. The first pulse begins with a detuning of $2\pi \times 1.92\,\mathrm{MHz}$, and ends on resonance, with respect to the centre of the condensate. This results in the resonant coupling position starting outside the condensate, at $z = -115.2\,\mu\mathrm{m}$, and sweeping to the condensate centre.[10] At the end of the first pulse, the left side of the condensate is in the $|+1\rangle$ state, while the right side remains in $|-1\rangle$. This is shown in Figure 4.8(a). Due to the finite edge sharpness of the pulse, there is some overlap of all three spin populations at the centre of the condensate, as shown in Figure 4.9(a).

[10] This choice in sign of the gradient is deliberate, and is discussed in Section 4.4.3. The consequences of the direction of the sweep across the condensate are discussed in Section 4.4.1.

At the conclusion of the first pulse, we remove the magnetic field gradient and apply a constant detuning of $\Delta_\phi = 2\pi \times 79.5\,\mathrm{kHz}$ for a duration of $t_\phi = 5\,\mu\mathrm{s}$.[11] During this stage, the Rabi frequency is zero (the radiofrequency oscillator continues to run but no signal is applied to the atoms) and a phase difference is allowed to accumulate between the two sides of the condensate. The density of each spin state at the end of this phase-accumulation stage is shown in Figures 4.8(b) and 4.9(b). Experimentally, this detuning Δ_ϕ can be controlled by adjusting the spatially-uniform Zeeman splitting across the condensate using a magnetic field offset, or by adjusting the frequency of the coupling field between the two pulses.

[11] This choice of Δ_ϕ and t_ϕ is explained in Section 4.4.1.

Then we apply a magnetic field gradient of $+237.5\,\mathrm{G/cm}$, and apply a second hyperbolic secant pulse which starts on resonance and ends with a detuning of $-2\pi \times 1.92\,\mathrm{MHz}$, with respect to the centre of the condensate. The resonant coupling position starts at the centre of the condensate and sweeps back out to $z = -115.2\,\mu\mathrm{m}$. At the end of the second pulse, we remove the magnetic field gradient. The left side of the condensate has been returned to the $|-1\rangle$ state, while the right side (as always) remains in $|-1\rangle$. This is shown in Figure 4.8(c). Due to the finite edge sharpness of the pulse, some residual population remains in states $|0\rangle$ and $|+1\rangle$ at the centre of the condensate, as shown in Figure 4.9(c).

At the conclusion of the second pulse, we have created both a density notch of width $w_{\mathrm{sol}} = 1.6\,\mu\mathrm{m}$ and a π phase step in the $|-1\rangle$ state, filled by population in states $|0\rangle$ and $|+1\rangle$. To create a black soliton this residual population must be removed. Experimentally, this can be achieved by using a microwave coupling pulse to rapidly transfer this population to $|F = 2\rangle$, then an optical pulse can remove this population from the trap entirely [114]. In these simulations, to remove this residual population we set the populations in $|0\rangle$ and $|+1\rangle$ to zero after the second pulse. Figure 4.10 shows the density of each spin state during the entire protocol as a 2D colour-map plot.

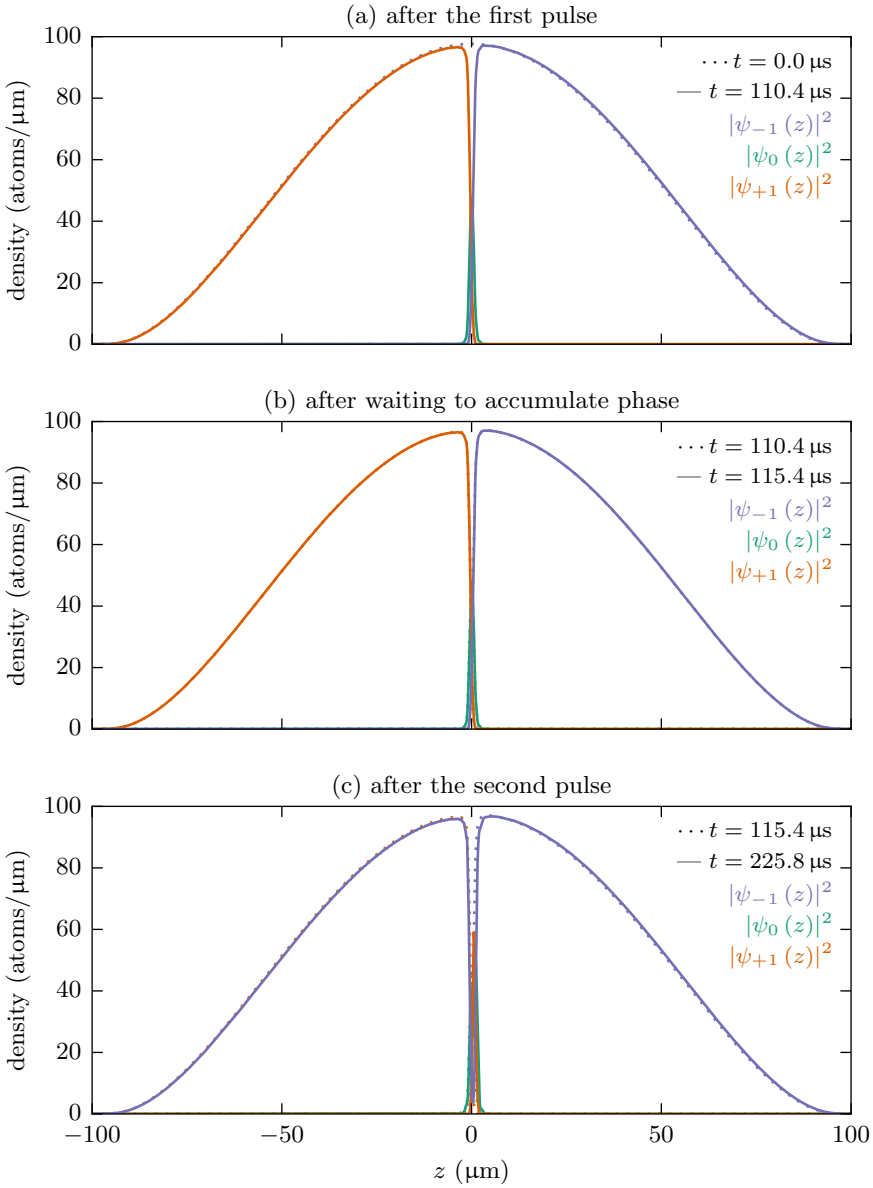

Figure 4.8: The density of each spin state at key times during the protocol. (a) In the magnetic field gradient of $-237.5\,\mathrm{G/cm}$ the first pulse transfers the left side of the condensate to $|{+1}\rangle$. There is a slight bulk movement of the condensate to the right as a result of the magnetic gradient (seen in the offset of dotted lines from solid lines), but this movement is small because the protocol completes faster than the healing time; this is discussed further in Section 4.4.2. (b) After the first coupling pulse, the magnetic field gradient is removed, and a constant detuning is applied for 5 µs to accumulate a phase difference between the two sides of the condensate. (c) In the magnetic field gradient of $-237.5\,\mathrm{G/cm}$ the second pulse returns the left side to $|{-1}\rangle$. Near $z = 0$ some population remains in $|0\rangle$ and $|{+1}\rangle$, creating a density dip in $|{-1}\rangle$. This central region is shown in Figure 4.9.

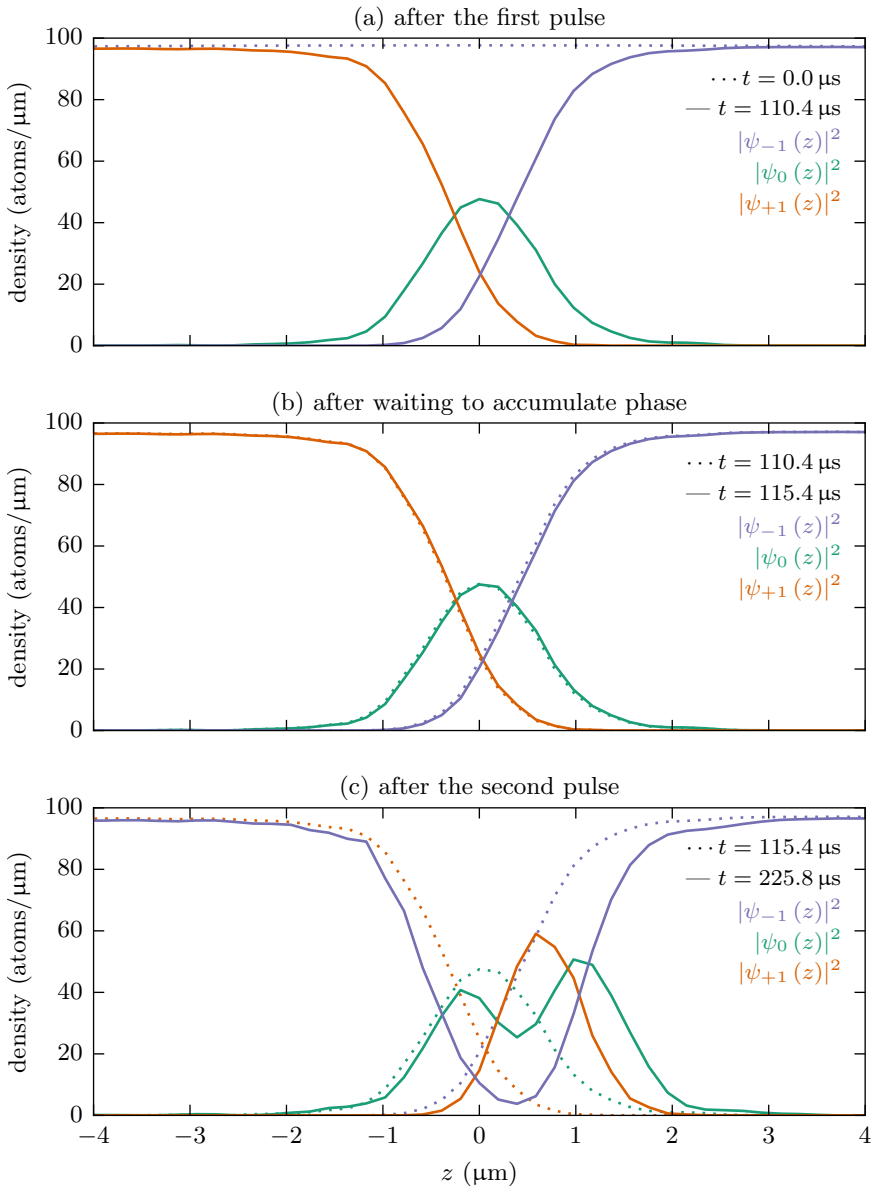

Figure 4.9: The density of each spin state in the central region of the condensate at key times during the protocol. (a) After the first pulse, some population remains in state $|0\rangle$ due to the finite sharpness of the slice profile. (b) Comparison of the dotted and solid lines shows the minimal movement of the condensate during the phase accumulation stage. (c) After the second pulse some population remains in $|0\rangle$ and $|+1\rangle$ at the centre of the condensate, creating a density dip in $|-1\rangle$ with a FWHM of 1.6 µm, matching the width of a black soliton in this condensate. With this density dip and the accumulated phase difference, a black soliton is created after the residual population is removed.

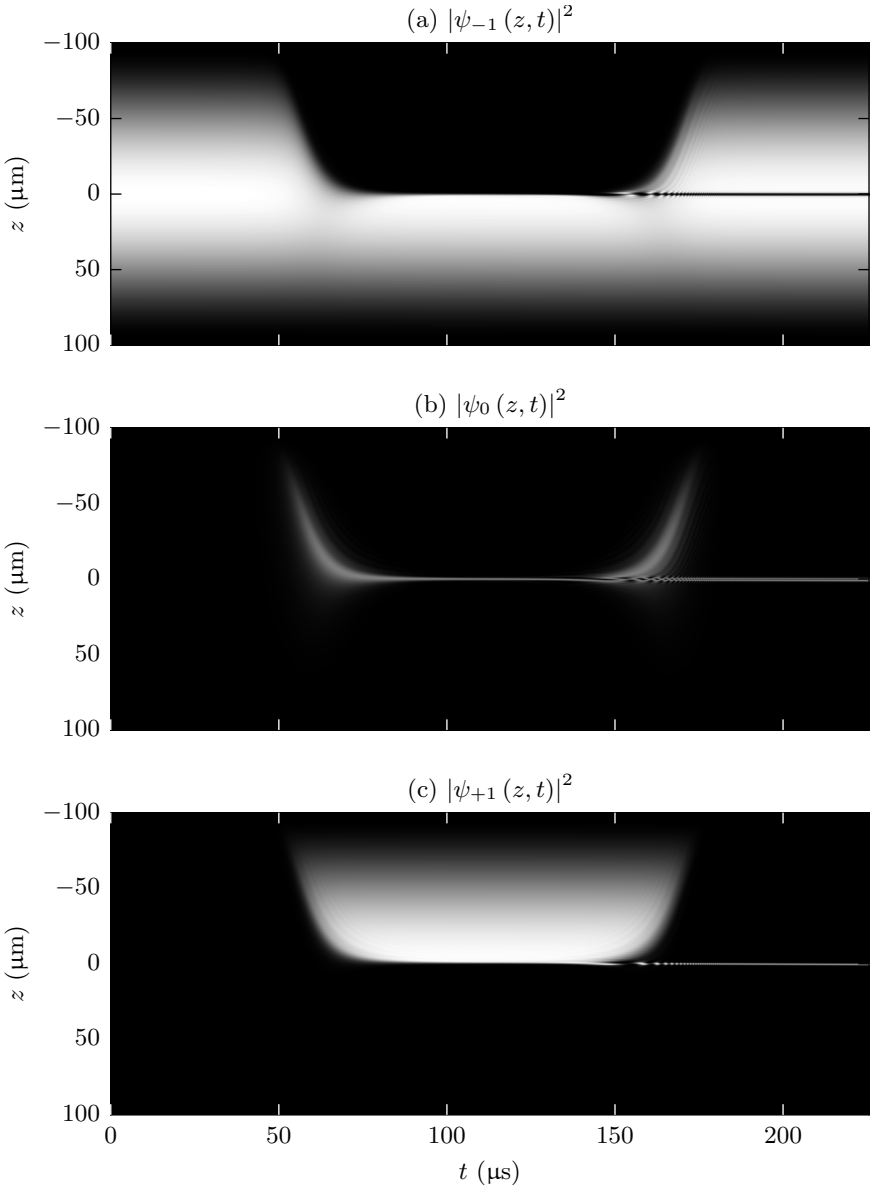

Figure 4.10: The density in each spin state while engineering a black soliton. The first pulse slices in from the left edge of the condensate to the centre, completing the transfer of the left side from $|-1\rangle$ to $|+1\rangle$ at $t = 110.4\,\mu s$. After a $5\,\mu s$ wait time to accumulate a phase difference between the two sides, the second pulse slices out from the centre of the condensate, returning the left side of the condensate to $|-1\rangle$. The finite edge sharpness of the pulse carves a density dip of width $w_{sol} = 1.6\,\mu m$ into the $|-1\rangle$ state, filled by residual population in states $|0\rangle$ and $|+1\rangle$.

4.3.6 Confirming that a black soliton is engineered

To confirm that we have engineered a black soliton, we observe the stability of the soliton for several trap periods. Figure 4.11 shows the total density of the condensate $|\psi_{-1}(z,\, t)|^2$ over one second of evolution. At this scale, the time over which we apply the magnetic resonance control protocol (Figure 4.10) is not visible. This much longer time scale reveals that the protocol did indeed engineer a black soliton; the soliton remains stationary at $z = 0$.

We note that the protocol did not result in an exact black soliton immediately after the protocol. Figure 4.12(a) shows that the engineered density notch only reaches $0.04\, n_0$, not all the way to zero. This 'almost-black' soliton heals to a black soliton by shedding small-amplitude sound waves which oscillate within the condensate as shown in Figure 4.11, with a lineout at $t = 100\,\text{ms}$ shown in Figure 4.12(b). We anticipate that these sound waves could be reduced through further optimisation of the hyperbolic secant pulse parameters. Remember that perfect black soliton creation is not the end goal of magnetic resonance control, but is rather simply an example of wavefunction engineering that requires rapid completion with sub-optical resolution. We have demonstrated here that magnetic resonance control is capable of engineering such a challenging target.

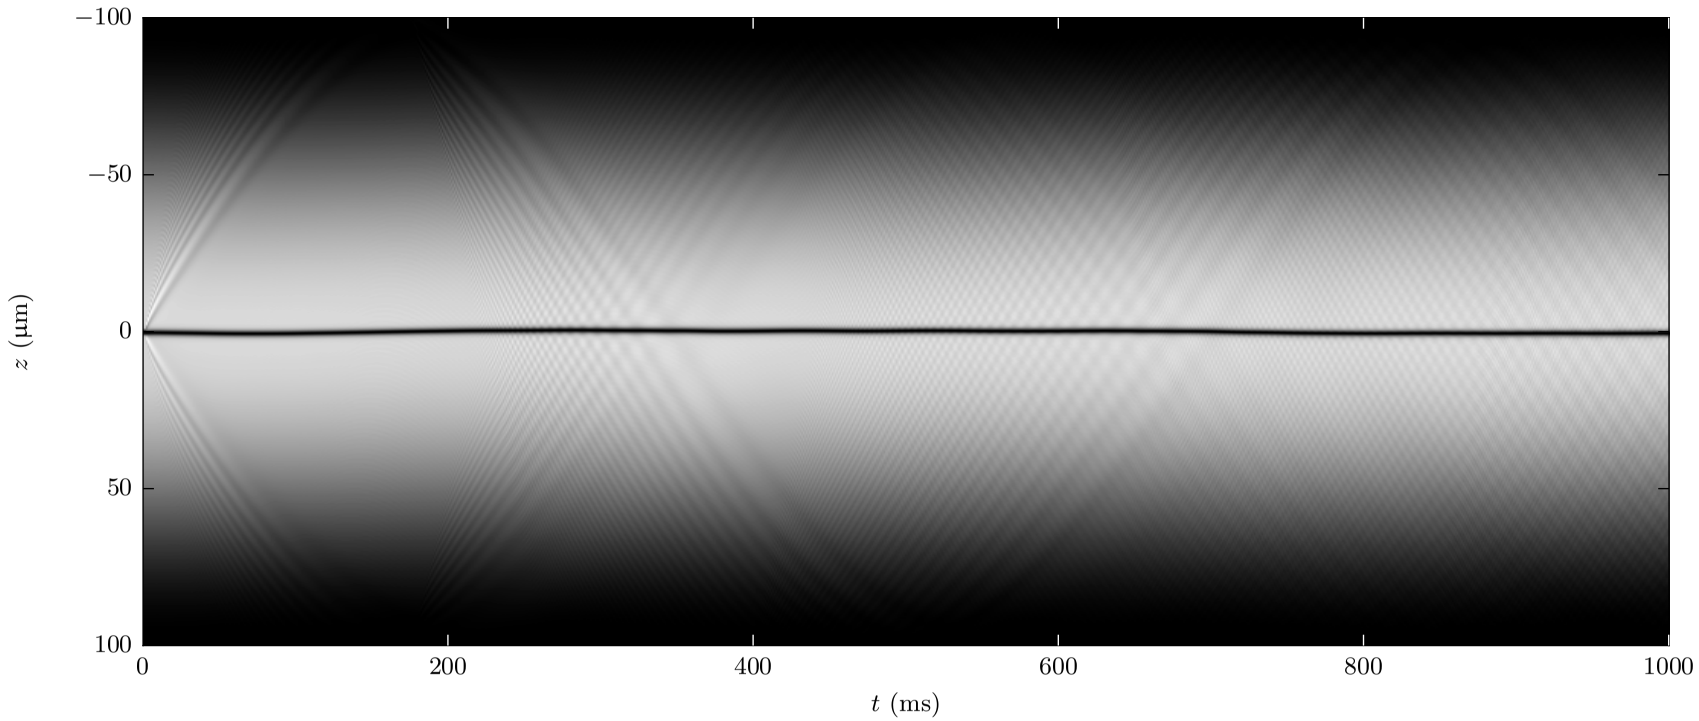

Figure 4.11: The total density of the condensate evolving for one second after applying the MRC protocol, confirming that a *black* soliton was engineered; the soliton remains stationary at $z = 0$ over several trap periods (a rare case of being excited to see 'nothing happen' when evolving a condensate over time). At this time scale, the magnetic resonance control stage is invisible. Some small-amplitude sounds waves are shed from the condensate immediately after the protocol, which can be seen oscillating within the condensate.

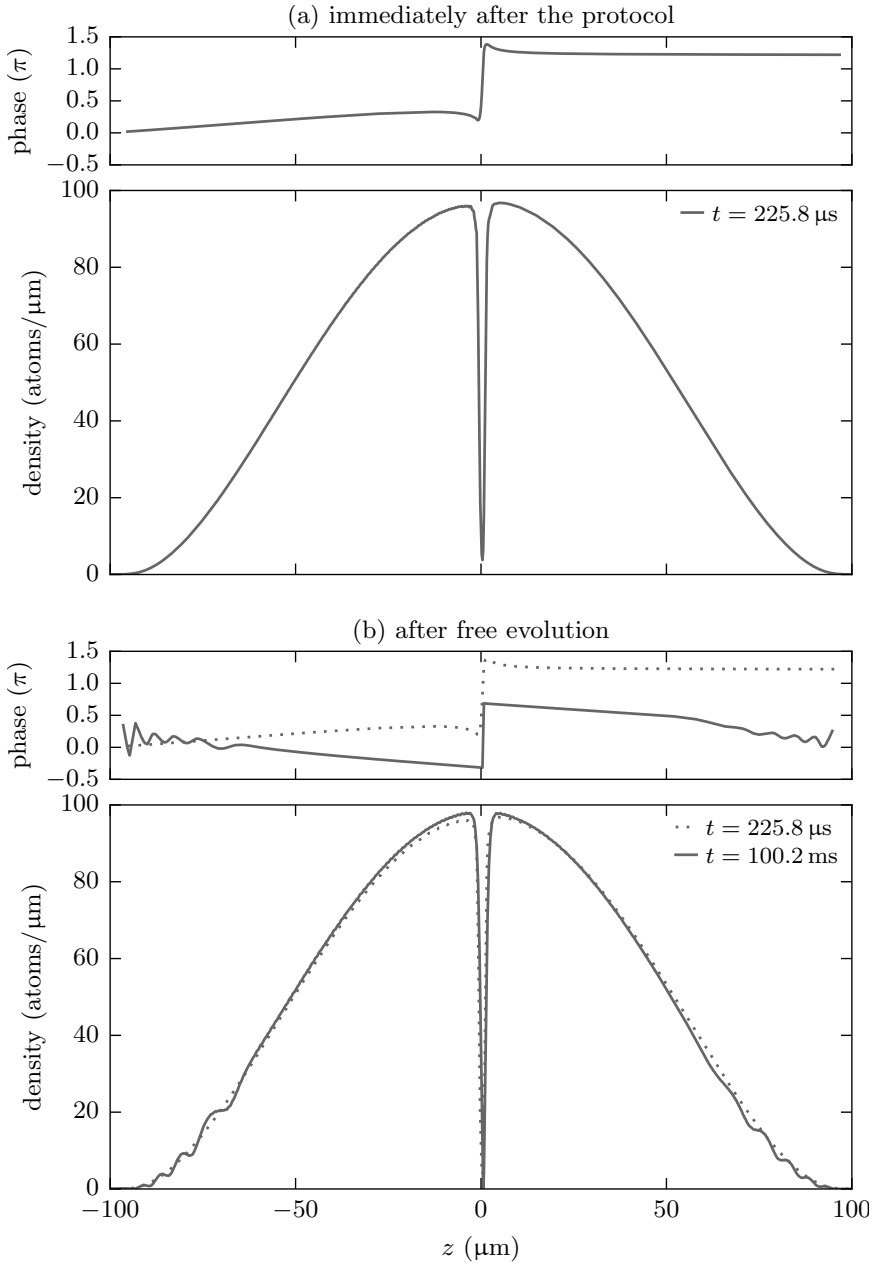

Figure 4.12: (a) The phase and total density of the condensate immediately after the protocol, showing the engineered density dip with FWHM of $w_{sol} = 1.6\,\mu m$ and minimum density of $0.04\,n_0$. (b) The phase and total density of the condensate 100 ms after the protocol, showing that this state has healed into a black soliton by shedding some small-amplitude sound waves. We anticipate that these sound waves could be reduced by further optimisation of the pulse used to engineer the soliton.

4.4 Protocol design details

There were three main factors that needed to be considered when designing the magnetic resonance control protocol used to engineer a black soliton. These factors were the phase accumulated during the time between pulses, the duration of the protocol, and the force applied by the magnetic field gradient (Section 4.4.3), and I discuss these in Sections 4.4.1, 4.4.2, and 4.4.3 respectively.

4.4.1 Ensuring a π phase step

In Section 4.3.5 I stated that after the first pulse we apply a constant detuning of $\Delta_\phi = 2\,\pi \times 79.5\,\text{kHz}$ for a duration of $t_\phi = 5\,\mu\text{s}$ to ensure that at the end of the protocol there is a π phase difference across the condensate. In theory, to accumulate a phase difference of κ between two states $|m_F = \pm F\rangle$, the time-detuning product should be

$$\Delta_\phi\, t_\phi = \frac{\kappa}{2\,F}\,, \tag{4.47}$$

so to accumulate a phase difference of π in our spin-1 system, we should want to set $\Delta_\phi\, t_\phi = \pi/2$. However, we instead found that $\Delta_\phi\, t_\phi = 1.59\,\pi/2$ was the necessary time-detuning product to create a phase difference of π at the end of the protocol. This apparent discrepancy arises from that each pulse contributes a difference in phase across the two sides of the condensate, due to evolution during each pulse. If the pulses used to engineer a black soliton are applied back-to-back, with no phase-accumulation stage in between, the result is a broad phase step of roughly $\pi/2$, decreasing from left to right. The shape of the phase step depends on the pulse duration, and the direction of the detuning sweep (whether the resonance point moves into the condensate, or starts at the centre and moves outwards). This means that there is already a non-uniform phase profile at the beginning of the phase-accumulation stage.

This explains why we need to set the time-detuning product to $\Delta_\phi\, t_\phi \neq \pi/2$, but implies that we should require $\Delta_\phi\, t_\phi < \pi/2$; instead, we find that we require $\Delta_\phi\, t_\phi = 1.59\,\pi/2$ to form a black soliton. This time-detuning product was found empirically by setting $t_\phi = 5\,\mu\text{s}$ and observing the stability of the engineered soliton for various Δ_ϕ until the resulting soliton was black (Figure 4.13). It appears that to create a black soliton, the phase step at the end of the protocol must be a *rising* edge (increasing from left to right) with a height of π; when the phase exhibits a falling edge of π, the resulting density dip is shallow.[12] Since the pulses contribute a phase slope with a falling edge, the time-detuning product needs to be *greater* than $\pi/2$ to result in a rising edge of π at the end of the protocol.

[12] This requirement is presented here as an empirical observation. Future work may investigate the underlying cause of this requirement.

I set $t_\phi = 5\,\mu\text{s}$ to reduce movement of the condensate between the two pulses. The rapid switching of magnetic fields this necessitates is challenging to implement experimentally. This wait time can be increased; a black soliton can still be engineered provided that the detuning Δ_ϕ is adjusted to compensate for the movement of the condensate. For example if $t_\phi = 50\,\mu\text{s}$, then a time-detuning product of $\Delta_\phi\, t_\phi = \pi$ produces a black soliton.

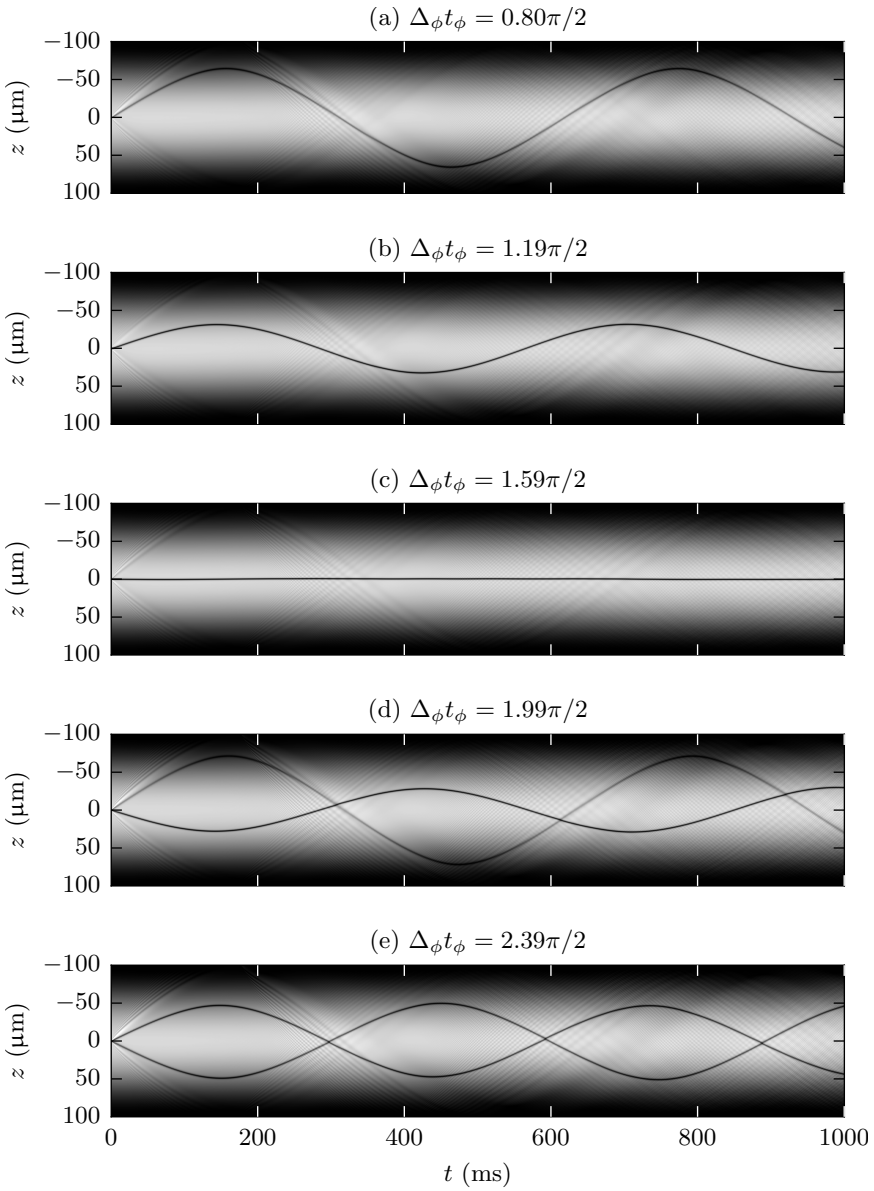

Figure 4.13: The total density of the condensate evolved for one second after the magnetic resonance control protocol. In all cases the hyperbolic secant pulses and the magnetic gradient had the same parameters as used in Section 4.3, and the wait time was $t_\phi = 5$ µs. The simulations shown here have a different detuning Δ_ϕ between the pulses, with the time-detuning product indicated by each label. When $\Delta_\phi\, t_\phi < 1.59\pi/2$, the phase step is less than π and a grey soliton is formed. When $\Delta_\phi\, t_\phi > 1.59\pi/2$ the phase step is greater than π, and the feature decays into two grey solitons.

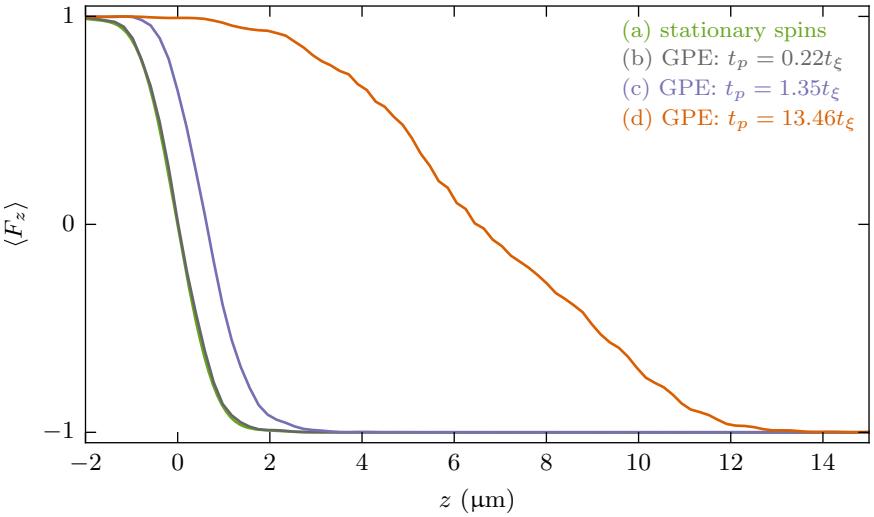

Figure 4.14: The spin projection of the condensate after the first pulse, for different pulse durations. In all cases the pulse parameters are $\Gamma = 5$, $\mu = 3.2$ and $\alpha = 0.003$. (a) The stationary spins prediction is independent of the pulse duration. The slice edge is located at $z = 0$ with a sharpness of 1.6 μm. (b) When $\Omega_0 = 2\pi \times 300$ kHz, the pulse duration is $t_p = 110.4$ μs, and the GPE simulation agrees closely with the stationary spins prediction. (c) When $\Omega_0 = 2\pi \times 50$ kHz, the pulse duration is $t_p = 622.3$ μs. In this case, even though the slice sharpness is only slightly increased (to 1.8 μm), the edge of the slice has noticeably moved compared to the stationary spins prediction (to $z = 0.6$ μm). (d) When $\Omega_0 = 2\pi \times 5$ kHz, the pulse duration is $t_p = 6.6$ ms. Now the effect of the magnetic field gradient is very apparent; the condensate has moved so far to the right during the pulse that the slice edge is now centred at $z = 6.5$ μm, and the sharpness has blurred to $\delta z = 7.6$ μm. In all cases the magnetic field gradient was set using Equation 4.13, with $\Delta z = 1.2\, z_{\mathrm{TF}}$ so that the entire left side of the condensate is coupled. Even though the magnetic field gradient decreases when Ω_0 is decreased, the movement due to the gradient becomes more significant because the duration of the magnetic field gradient is increased.

4.4.2 *The importance of being faster than the healing time*

In Section 4.1 we discussed the necessity of keeping the duration of the protocol below the healing time of the condensate. Figure 4.14 demonstrates the change to the slice profile if the duration of the pulse is increased beyond the healing time. We can see that when the pulse is too slow, the GPE simulations no longer agree with the stationary spins prediction (Section 4.1.1); the edge of the slice profile becomes broader, and moves to the right. This is because, unlike in the single-atom prediction, the condensate in the GPE simulation is free to move. Provided the pulse duration is less than the healing time, this movement does not significantly alter the shape and position of the slice, and the combination of two pulses can be used to engineer a black soliton.

4.4.3 *Tolerating the magnetic gradient 'kick'*

For ^{87}Rb atoms in a magnetic field gradient, atoms in $|-1\rangle$ will experience a force towards smaller $|\mathbf{B}|$ while atoms in $|+1\rangle$ will experience a force towards larger $|\mathbf{B}|$.

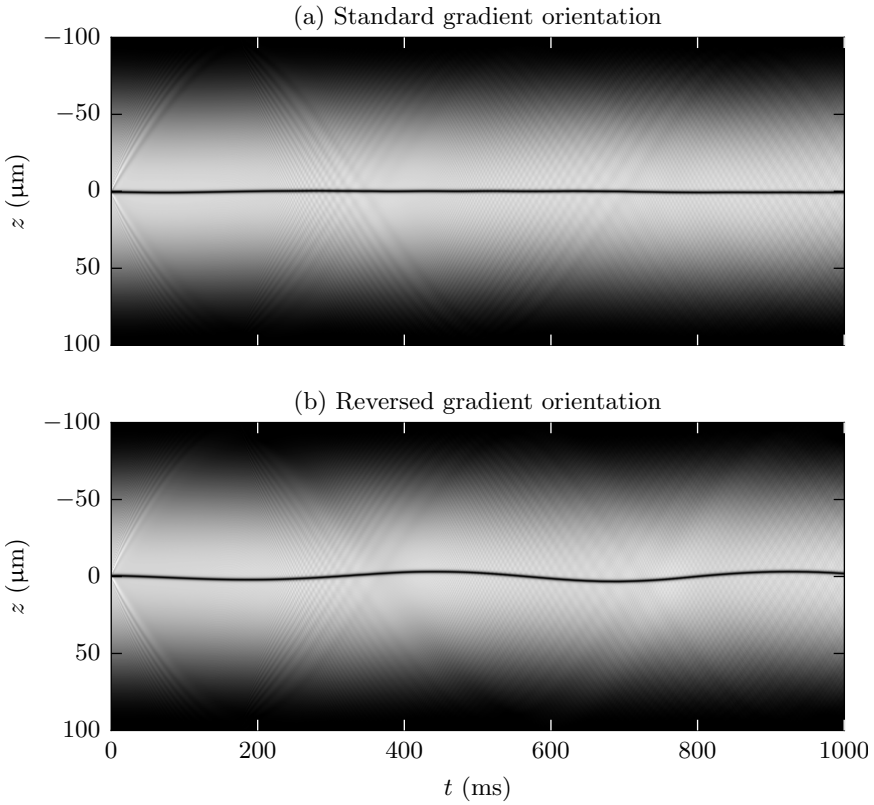

Figure 4.15: (a) The same data presented in Figure 4.11; the black soliton created by the successful protocol, in which we apply a negative gradient during the first pulse, and a positive gradient during the second pulse. (b) Reversing the orientation of the gradient, while still addressing the left side of the condensate, results in larger amplitude sound waves being shed from the soliton, and the soliton oscillates within the condensate.

This means that during the first pulse, the uncoupled side of the condensate will move to one side, the direction depending on both the initial spin state of the condensate and the sign of the magnetic field gradient. The coupled side of the condensate does not move as much as the uncoupled side because, at least on average across that side of the condensate, the population spends equal time in both states.

It is best to ensure that the two sides of the condensate move apart from each other during the first hyperbolic secant pulse, rather than moving towards each other, by choosing the appropriate sign of the gradient. To achieve this, since the condensate starts in $|-1\rangle$ and we choose to couple the left side, we apply a negative gradient. If the two sides are pushed together by the magnetic field gradient, the partial miscibility of the two states causes higher density peaks to form at the boundary. These density peaks decay into sound waves, which perturb the soliton (Figure 4.15).

4.5 Controlled creation of multiple solitons

In the simulations described in this chapter so far, the coupling pulses have only addressed the left side of the condensate; the coupled slice has had a thickness of $\Delta z = 1.2\, z_{\mathrm{TF}}$, and the detuning offset has been set to the amplitude of the detuning sweep ($\Delta_1 = \Delta_0$) to position the right edge of the slice at $z = 0$. This means that the left edge of the slice has been an unused commodity. If both edges are positioned inside the condensate we can engineer two solitons simultaneously, located at each edge of the coupled slice; one of the solitons has a rising phase step, while the other has a falling phase step.

To do this, we need to ensure that as the slice thickness is decreased, the slice sharpness remains unchanged. Otherwise, the resulting density notches will be narrower than a black soliton, and hence produce unstable solitons. To achieve this, we decrease the range of the detuning sweep while maintaining the duration of the pulse. In terms of our dimensionless hyperbolic secant pulse parameters, this corresponds to decreasing μ and increasing Γ in proportion, while keeping all other parameters (including the magnetic field gradient) the same.

4.5.1 A Newton's cradle with two solitons

If we position the left edge of the slice inside the condensate, while keeping the right edge at $z = 0$, we can engineer two solitons that oscillate in the condensate like a two-ball Newton's cradle; initially the left soliton will move towards the (stationary) right soliton, and on collision the momentum will be exchanged so that the left soliton is stationary and the right soliton is moving to the right.

To demonstrate this effect, we set $\mu = 0.5$ and $\Gamma = 32$ while keeping all other parameters unchanged. In the magnetic field gradient of $237.5\,\mathrm{G/cm}$, this results in a slice thickness of $\Delta z = 18.0\,\mu\mathrm{m}$ while the slice sharpness is still $\delta z = 1.6\,\mu\mathrm{m}$[13].

[13] Remember that previously, we used $\mu = 3.2$ and $\Gamma = 5$ to create a single black soliton, with a slice thickness of $\Delta z = 115.3\,\mu\mathrm{m}$ and a slice sharpness of $\delta z = 1.6\,\mu\mathrm{m}$.

The densities of the three spin states during the protocol are shown in Figure 4.16. We can see that the left edge of the coupled slice is now at $z = -18.0\,\mu\mathrm{m}$, and at the end of the protocol a notch in the $|-1\rangle$ density has formed on either side of the slice. Figure 4.17 shows the evolution of the condensate over one second, demonstrating the solitons undergoing momentum-conserving collisions.

4.5.2 Two solitons oscillating in phase

If we move the two solitons closer together and change the detuning offset Δ_1 to position both solitons to one side of the condensate, we can engineer a state in which the two solitons maintain their relative separation and oscillate in phase. This state is of particular interest because such a double soliton state has never been experimentally realised in a quantum fluid. To engineer this state, we set $\mu = 0.14$, $\Gamma = 115.3$ and $\Delta_1 = -2\,\pi \times 166.52\,\mathrm{kHz}$. In the magnetic field gradient of $237.5\,\mathrm{G/cm}$, the slice edges are located at $z = -12.3\,\mu\mathrm{m}$ and $z = -7.3\,\mu\mathrm{m}$. Figure 4.18 shows the evolution of the condensate over one second, demonstrating that the solitons oscillate in phase with each other.

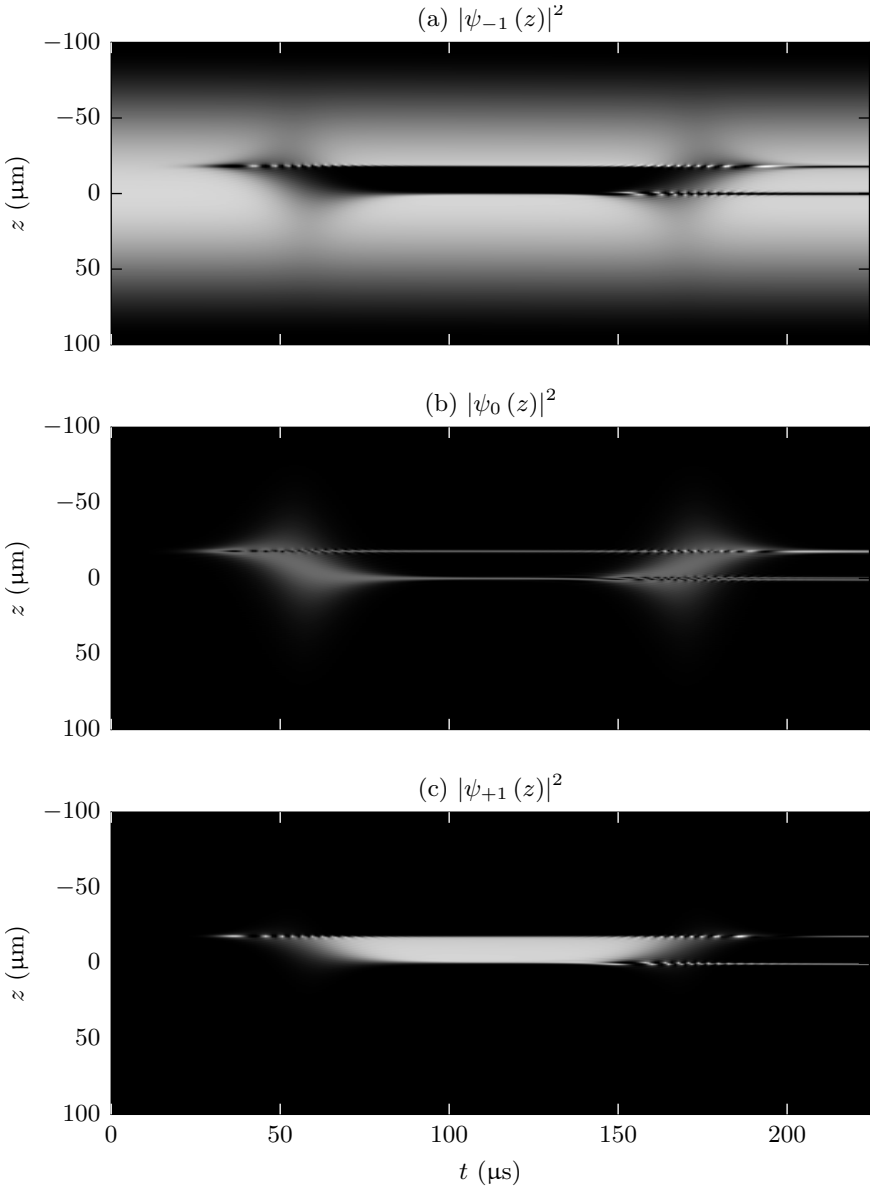

Figure 4.16: The density of each spin state $|\psi_i(z,t)|^2$ across the condensate during the magnetic resonance protocol. Here we have reduced the slice thickness (but maintained the slice sharpness) by setting $\mu = 0.5$ and $\Gamma = 32$, while all other parameters remain the same as those used to engineer a single black soliton (shown in Figure 4.10). Bringing the left edge of the slice within the condensate in this manner results in a soliton forming at each edge of the slice.

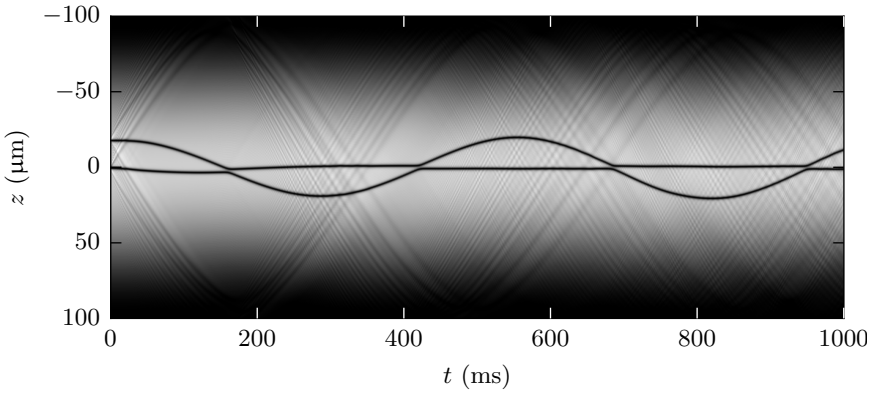

Figure 4.17: The total density of the condensate evolved for one second after the magnetic resonance control protocol. We observe the two solitons oscillating like a two-ball Newton's cradle.

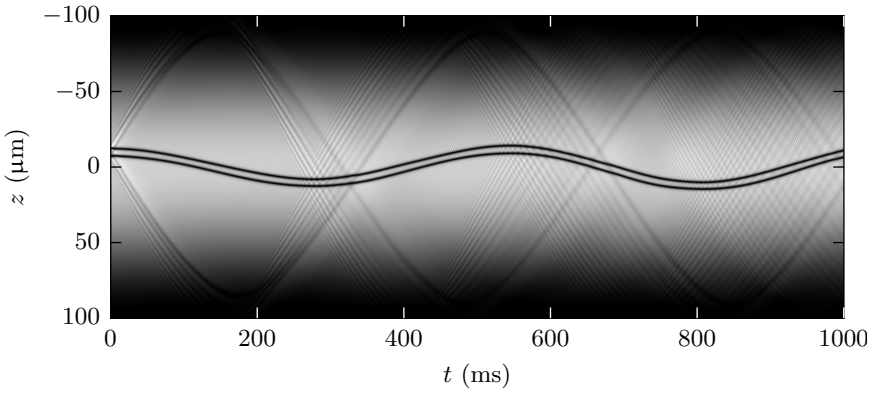

Figure 4.18: The total density of the condensate evolved for one second after engineering two solitons at $z = -12.3\,\mu\text{m}$ and $z = -7.3\,\mu\text{m}$. These two solitons oscillate in phase.

We note that a particular orientation of the magnetic field gradient is not going to be ideal for both solitons. In this case, the right edge of the slice moved first away, then towards, the uncoupled condensate which minimises the amplitude of sound waves produced by the rightmost soliton. Consequently, the left edge of the slice pushed against the uncoupled portion of the condensate, increasing the amplitude of sound waves which shed from the leftmost soliton after the protocol, which then perturb the rightmost soliton. With a small soliton separation of $5\,\mu\text{m}$, the sound waves shed from one soliton will perturb the other soliton. To correct for this so that the two solitons ultimately oscillate in phase, we can adjust the initial trajectory of each soliton by varying the accumulated phase difference. Here, we set $\Delta_\phi\, t_\phi = 1.53\pi/2$, with $t_\phi = 5\,\mu\text{s}$, instead of the previous optimal value of $\Delta_\phi\, t_\phi = 1.59\pi/2$ for creating a single black soliton. We adjust both the time-detuning product and the soliton separation until the solitons oscillate in phase.

4.5.3 Repeating the protocol to make four solitons

We can engineer four solitons by repeating the protocol that makes two solitons, with a different offset to the slice position each time. Because our magnetic resonance control protocol is so fast, it can be successfully applied to the same condensate multiple times to make multiple pairs of solitons.

Figure 4.19 shows the density of the three spin states while engineering four solitons positioned symmetrically about the condensate centre: at $z = -12.3\,\mu\text{m}$, $z = -7.3\,\mu\text{m}$, $z = 7.3\,\mu\text{m}$ and $z = 12.3\,\mu\text{m}$. Each pair alone would oscillate in phase like those shown in Figure 4.18. Together, the four solitons experience symmetric collisions, as shown in Figure 4.20.

Instead of engineering four symmetric solitons, we can offset both slices to one side of the condensate to make four equidistant solitons, initially located to the left of the condensate centre. Figure Figure 4.21 shows four such solitons located at $z = -22.0\,\mu\text{m}$, $z = -17.0\,\mu\text{m}$, $z = -11.5\,\mu\text{m}$ and $z = -6.5\,\mu\text{m}$. The sound waves shed from these solitons perturb the solitons sufficiently to prevent perfect in-phase oscillations of all four solitons; they oscillate roughly in-phase, with the trajectories of two inner solitons being constrained by the trajectories of the two outer solitons.

Here I have demonstrated that our magnetic resonance control protocol can be applied in succession to make multiple pairs of solitons. Since single solitons can be engineered if one side of the slice remains outside the condensate, we could make odd numbers of solitons, not just multiple pairs of solitons. The initial position of each soliton is controlled by the width and offset of the slice, while the initial trajectory of each soliton is controlled by adjusting the phase accumulated between pulses.

Being able to create arbitrary arrangements of solitons, with control over their initial trajectories, could enable the study of soliton interactions with unprecedented control. This could also provide the ability to create soliton arrays, which exhibit fermionisation within the bosonic condensate [115, p. 199]. Further, instead of simply repeating the protocol to make more solitons, we could use a multi-tone rf sweep to simultaneously engineer multiple solitons; I have not attempted this yet, but recommend it as an area to explore.

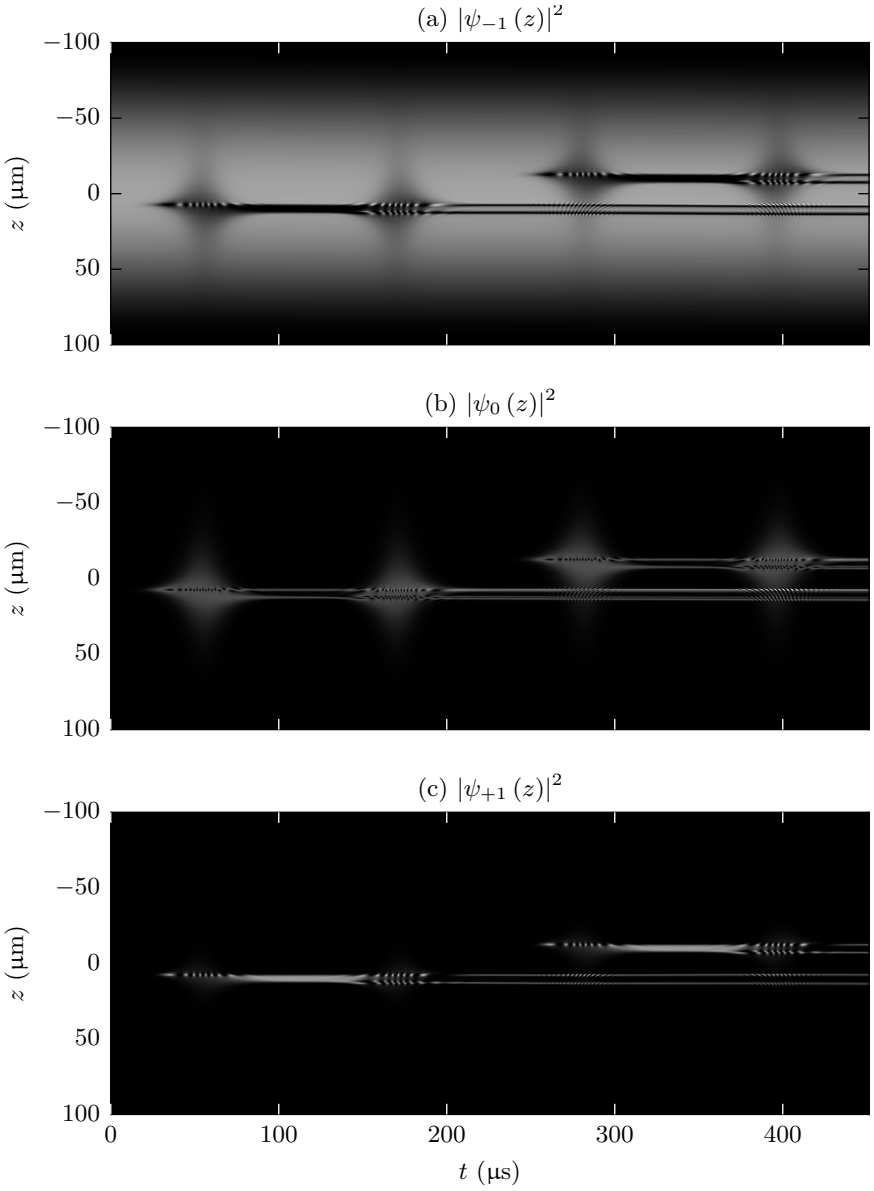

Figure 4.19: The density of the three spin states across the condensate while engineering four solitons. The first iteration of the protocol creates a pair of solitons at $z = 7.3\,\mu\text{m}$ and $z = 12.3\,\mu\text{m}$, while the second iteration produces a soliton pair at $z = -7.3\,\mu\text{m}$ and $z = -12.3\,\mu\text{m}$.

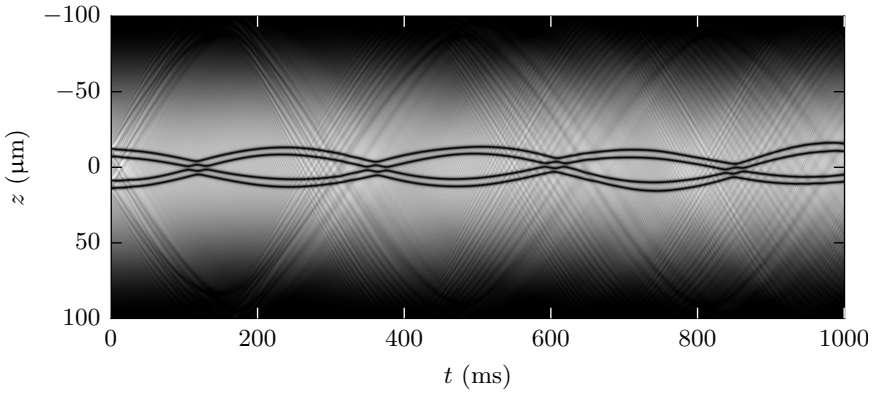

Figure 4.20: The total density of the condensate evolved for one second after engineering four solitons at $z = -12.3\,\mu\text{m}$, $z = -7.3\,\mu\text{m}$, $z = 7.3\,\mu\text{m}$ and $z = 12.3\,\mu\text{m}$. If only one of the two soliton pairs were present, they would oscillate in phase like the pair shown in Figure 4.18.

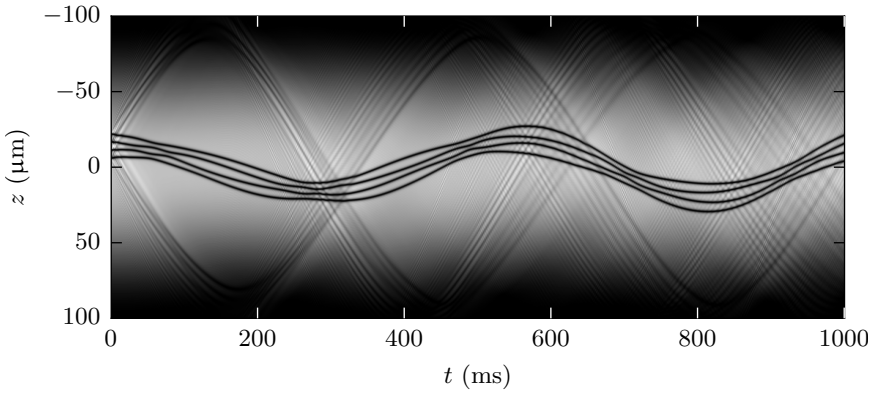

Figure 4.21: The total density of the condensate evolved for one second after engineering four solitons at $z = -22.0\,\mu\text{m}$, $z = -17.0\,\mu\text{m}$, $z = -11.5\,\mu\text{m}$ and $z = -6.5\,\mu\text{m}$. These four solitons almost exhibit in-phase oscillations, but have been perturbed by the sound waves shed by each soliton.

4.6 Summary

The simulations in this chapter demonstrate that for a spin-1 condensate in a linear magnetic field gradient, we can design hyperbolic secant pulses that couple the Zeeman sublevels of the condensate with spatial dependency. We can select the coupled region with healing-length precision, allowing us to control both the phase and the density of the condensate. We use this magnetic resonance control technique to simulate engineering a black soliton in a condensate using experimentally feasible parameters. Further, we have shown this technique to be easily extensible, simulating the creation of multiple soliton states with control over their placement and initial trajectories.

Table 4.1: Record of the datasets (`.py` files) containing the results shown in figures in this chapter. The time stamps are unique identifiers for the files, which have names of the form `<timestamp>_spinor_1D.py`.

Figure	Time stamp
4.5	20140529T094741
4.6	20140527T105810
	20140611T103716
	20140611T103354
	20140611T102931
	20140611T102430
4.7 – 4.12	20150205T112108
4.13	20150205T142530
	20150205T143031
	20150205T112108
	20150205T143055
	20150205T143120
4.14	20150205T112108
	20150207T160257
	20150208T135927
4.15	20150205T112108
	20150207T121704
4.16, 4.17	20150208T153606
4.18	20150209T145358
4.19, 4.20	20150209T165356
4.21	20150209T162519

4.7 Record of simulation files

Every time I ran a simulation, the results were stored in a `.h5` file with a name containing the time and date of the execution and the name of the Python script used to define the simulation parameters. Inspired by a similar system we use in our laboratory for recording and organising our experimental data, for each simulation I also stored a copy of the Python script with the same filename as the `.h5` file. This enabled me to easily keep a log of my simulations, and to revisit a previous simulation to check results or parameters.

If anyone with access to our data repository wishes to investigate any of the results presented in this chapter, Table 4.1 is a record of the `.py` files used to create the results shown in each figure.

Chapter 5

Searching the condensate parameters

In the previous chapter, I presented simulations of using magnetic resonance control to engineer solitons in a ^{87}Rb condensate. This condensate contained 10^4 atoms in a harmonic trap with axial frequency $\omega_z = 2\pi \times 2.4\,\text{Hz}$ and radial frequency $\omega_\rho = 2\pi \times 158.4\,\text{Hz}$. This particular set of condensate parameters was not an arbitrary choice. In this chapter, we explore how these condensate parameters are constrained by four criteria necessary to engineer a black soliton in a condensate using magnetic resonance control. Examining this condensate 'parameter space' informs the design of trapping potentials to create condensates in the laboratory suitable for applying our magnetic resonance control technique.

The condensate must satisfy four criteria to use magnetic resonance control to best demonstrate using magnetic resonance control to engineer a black soliton in the condensate. These criteria are:

1. *The quasi-1D condensate criterion:* To support a stable soliton for several trap periods, long enough to be studied experimentally, the condensate must be quasi-one dimensional.
2. *The sub-optical healing length criterion:* We want the healing length of the condensate to be smaller than $500\,\text{nm}$. This criterion allows us to show that the resolution of magnetic resonance control can surpass the optical diffraction limit.
3. *The reasonable axial frequency criterion:* If the condensate is to be held in an optical dipole trap, the axial frequency of the trap should be greater than 1 Hz, otherwise the stability of the trap becomes extremely sensitive to the tilt of the dipole beam with respect to gravity.
4. *The pulse resolution criterion:* An experimentally feasible hyperbolic secant pulse must achieve the resolution R required to engineer a black soliton.

Here we explore the condensate parameter regions that satisfy each of these criteria, both in terms of the aspect ratio $\eta = \omega_\rho/\omega_z$ of a generic trap, and for the specific case of a single-beam optical dipole trap, which is defined by the power P and the waist size w_0 of the beam. As we see in Section 5.5, happily the intersection

of these four criteria is non-zero, accommodating a range of different condensate shapes and sizes.

5.1 Quasi-1D condensate criterion

To verify that a black soliton has been engineered, we need to be able to observe the soliton motion (or lack thereof) for several trap periods. This requires the condensate to be quasi-one dimensional, otherwise the soliton will rapidly become unstable and decay into vortex rings [31, 98].

A condensate is defined to be quasi-one dimensional [116] if it satisfies the relation

$$\frac{\mu_{\text{cp}}}{\hbar\,\omega_\rho} < \gamma_{\text{crit}}\,, \tag{5.1}$$

with μ_{cp} as the chemical potential,

$$\mu_{\text{cp}} = \left(\frac{N g}{\frac{8\pi}{15}\left(\frac{2}{m\overline{\omega}^2}\right)^{3/2}}\right)^{2/5}, \tag{5.2}$$

where $\overline{\omega} = \left(\omega_\rho^2\,\omega_z\right)^{1/3}$ is the geometric mean of the trapping frequencies, and the mean field interaction strength is $g = 4\,\pi\,\hbar^2\,\left(a_0 + 2\,a_2\right)/3\,m$.

The value of γ_{crit} increases with the aspect ratio of the trap, converging to $\gamma_{\text{crit}} = 2.4$ for very elongated traps ($\omega_\rho \gg \omega_z$). For simplicity, we decide to set $\gamma_{\text{crit}} = 2.4$ always; condensates with a low aspect ratio that satisfy Equation 5.1 with $\gamma_{\text{crit}} = 2.4$ will be more quasi-one dimensional than necessary for stable solitons. This is preferable to using a weaker region that includes some condensates which are not quasi-one dimensional.

We can express Equation 5.1 explicitly in terms of N, η, and ω_z as

$$\eta > \left(\frac{1}{\hbar\,\gamma_{\text{crit}}}\right)^5 \left(\frac{15\,N\,g\,m^{3/2}}{16\,\sqrt{2}\,\pi}\right)^2 \omega_z\,, \tag{5.3}$$

constraining the previously independent parameters N, η, and ω_z. We can now see that there is a linear relationship between the axial trap frequency and the trap aspect ratio for condensates which to satisfy the quasi-one-dimensional criterion. This relationship is shown in Figure 5.1 as a 3D region plot of $N\left(f_z,\,\eta\right)$, and in Figure 5.2 as a stack of 2D region plots of $\eta\left(f_z\right)$ for different values of N. These figures demonstrate that if we seek a quasi-one dimensional condensate with high atom number, we should aim for a condensate with either a high aspect ratio, or a low axial trapping frequency.

We can recast this criterion for the specific case of a single-beam optical dipole trap in terms of the independent parameters N, P, and w_0. For a focussed off-resonant laser beam [117] the radial trap frequency is

$$\omega_\rho = \sqrt{\frac{4\,U_0}{m\,w_0^2}}\,, \tag{5.4}$$

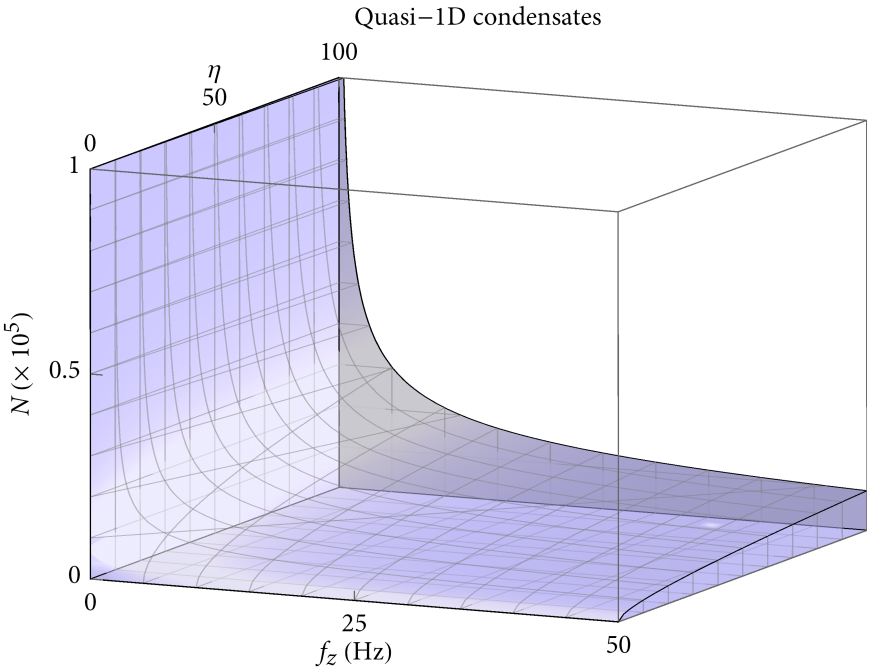

Figure 5.1: This solid region shows the combinations of N, η and f_z for which condensates are quasi-one dimensional (Equation 5.3). Planes of constant atom number from this solid region are shown in Figure 5.2.

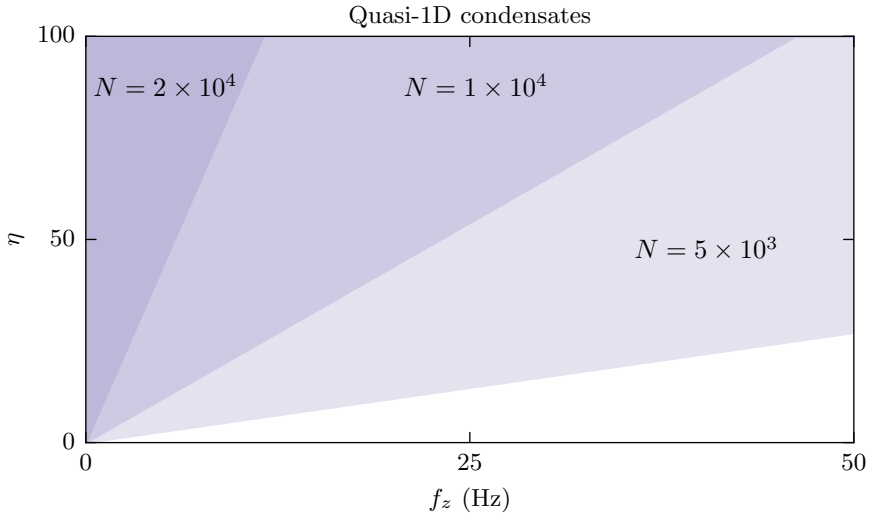

Figure 5.2: These regions show the combinations of η and f_z for which a condensate is quasi-one dimensional (Equation 5.3) with $N = 5 \times 10^3$, $N = 1 \times 10^4$, and $N = 2 \times 10^4$. The regions corresponding to higher atom numbers are contained within the regions for lower atom numbers.

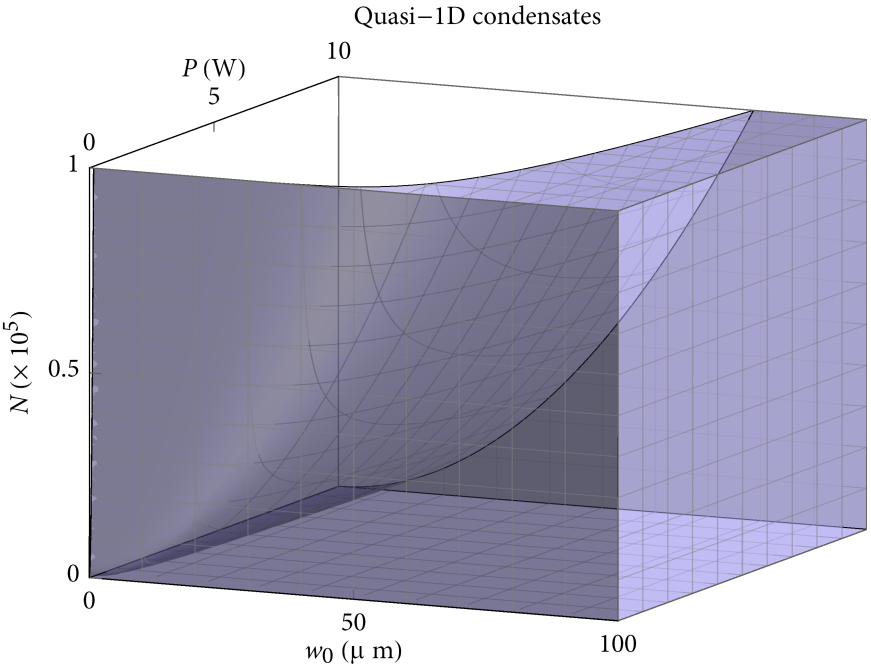

Figure 5.3: This solid region shows the combinations of N, P and w_0 for which a condensate trapped by a single-beam optical dipole trap is quasi-one dimensional (Equation 5.6). Planes of constant atom number from this solid region are shown in Figure 5.4.

and the axial trap frequency is

$$\omega_z = \sqrt{\frac{2\,U_0}{m\,z_\mathrm{R}^2}}\,. \tag{5.5}$$

Here, $U_0 = 2\,\alpha_s\,P/\pi\,w_0^2$ is the trap depth, where α_s is the dynamic scalar polarisability [117] (in units of J/W m²), and $z_\mathrm{R} = \pi\,w_0^2/\lambda$ is the Rayleigh range of the beam with wavelength λ. Using this information, Equation 5.3 can be expressed in terms of the beam power P and waist w_0 as

$$P < \left(\frac{\hbar\,\gamma_\mathrm{crit}}{\sqrt{m}}\right)^{10} \left(\frac{16}{15\,N\,g\,\lambda}\right)^{4} \frac{2\,\pi^9}{\alpha_s}\,w_0^8\,. \tag{5.6}$$

This relationship is shown in Figure 5.3 as a 3D region plot of $N\,(w_0,\,P)$, and in Figure 5.4 as a stack of 2D region plots of $P\,(w_0)$ for different values of N. These figures show that to have a quasi-one dimensional condensate in a single-beam optical dipole trap with high atom number, we should aim for a large beam waist. This 8th order scaling in the beam waist size constrains the waist size more strongly than power for a given atom number.

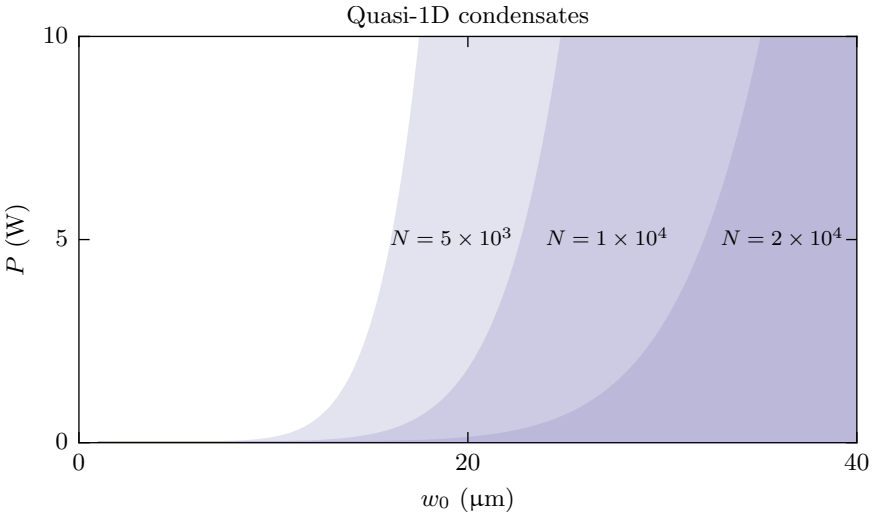

Figure 5.4: These regions show the combinations of P and w_0 for which a condensate trapped by a single-beam optical dipole trap is quasi-one dimensional (Equation 5.6) with $N = 5 \times 10^3$, $N = 1 \times 10^4$, and $N = 2 \times 10^4$. The regions corresponding to higher atom numbers are contained within the regions of lower atom numbers.

5.2 Sub-optical healing length criterion

To showcase that the resolution of magnetic resonance control is not limited to optical wavelengths, we demand the healing length be sub-optical: $\xi < \lambda_{\text{opt}}$, where $\lambda_{\text{opt}} = 500\,\text{nm}$. If we write this succinct condition explicitly in terms of the axial trap frequency ω_z and the trap aspect ratio η, it transforms into the more verbose relation

$$\eta > \left(\frac{\hbar}{\lambda_{\text{opt}}\sqrt{2\,m}}\right)^{5/2} \left(\frac{8\,\pi}{15\,N\,g}\right)^{1/2} \left(\frac{2}{m}\right)^{3/4} \frac{1}{\omega_z^{3/2}}\,. \tag{5.7}$$

This relationship is shown in Figure 5.5 as a 3D region plot of $N\,(f_z,\,\eta)$, and in Figure 5.6 as a stack of 2D region plots of $\eta\,(f_z)$ for different values of N. These figures show that to produce a condensate with a sub-optical healing length and a high atom number, the condensate should have either a high aspect ratio or a high axial trapping frequency. Since we saw previously that a high aspect ratio or *low* axial trapping frequency is required for quasi-one dimensional condensates, we can conclude that having a high aspect ratio would be the best approach for satisfying both these conditions.

We can again express this condition specifically for a single-beam optical dipole trap in the form

$$P > \left(\frac{\hbar}{\lambda_{\text{opt}}\sqrt{m}}\right)^{10/3} \left(\frac{\pi^{7/2}}{15\sqrt{2}\,N\,g\,\lambda}\right)^{2/3} \left(\frac{1}{2\,\alpha_s}\right) w_0^{14/3}\,. \tag{5.8}$$

This relationship is shown in Figure 5.7 as a 3D region plot of $N\,(w_0,\,P)$, and in Figure 5.8 as a stack of 2D region plots of $P\,(w_0)$ for different values of N. These

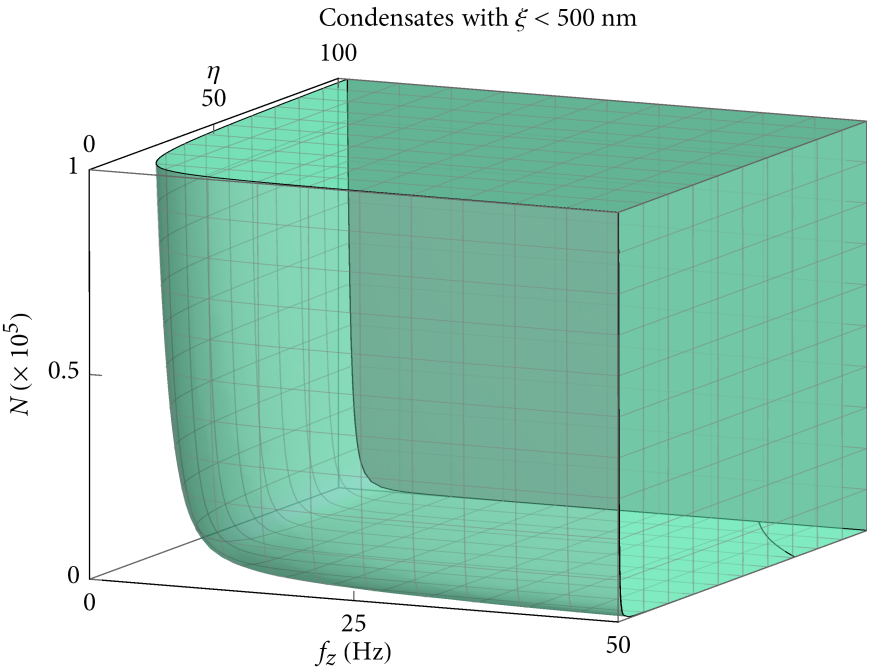

Figure 5.5: This solid region shows the combinations of N, η and f_z for which condensates have a healing length of less than 500 nm (Equation 5.7). This criterion primarily excludes the choice of condensates with both a low aspect ratio and small axial trap frequency. Planes of constant atom number of this solid region are shown in Figure 5.6.

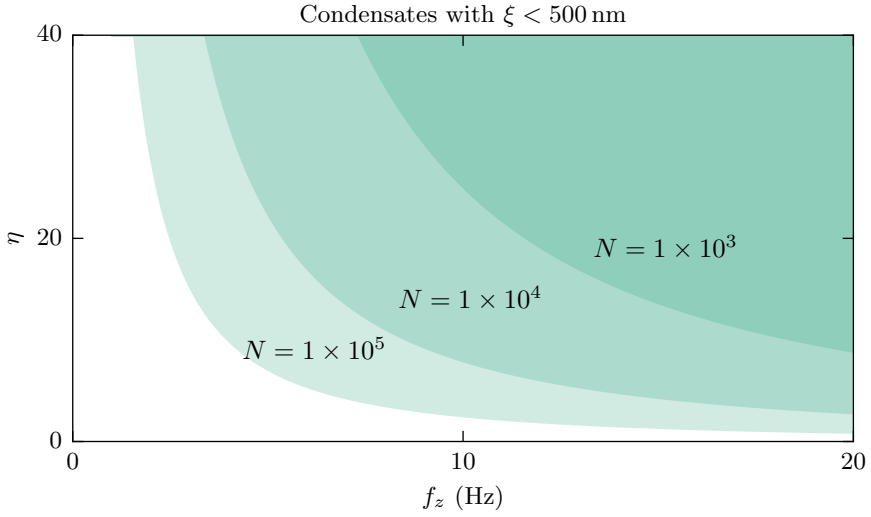

Figure 5.6: These regions show the combinations of η and f_z for which a condensate has a healing length of less than 500 nm (Equation 5.7), with $N = 1 \times 10^3$, $N = 1 \times 10^4$ and $N = 1 \times 10^5$. The regions corresponding to lower atom numbers are contained within the regions of higher atom numbers.

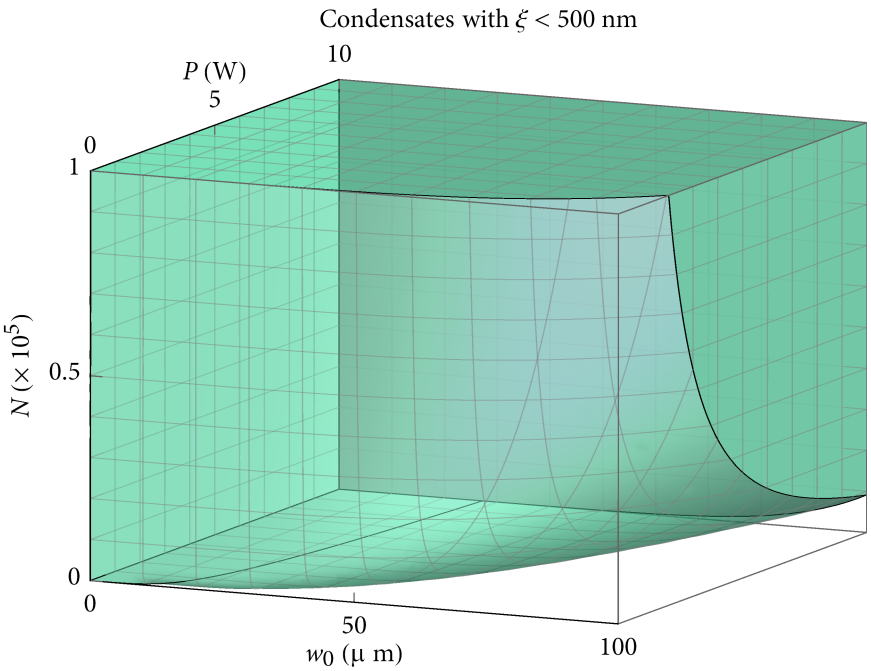

Figure 5.7: This solid region shows the combinations of N, P and w_0 for which a condensate trapped by a single-beam optical dipole trap has a healing length of less than 500 nm (Equation 5.7). This criterion excludes traps with large beam waists for small condensates. Planes of constant atom number of this solid region are shown in Figure 5.8.

figures show that if we need to use a larger beam waist, we will need to increase the number of atoms in the condensate to maintain a sub-optical healing length.

5.3 *Reasonable axial frequency criterion*

This criterion is simple; we ensure that $\omega_z > 2\pi \times 1\,\text{Hz}$, so that the atoms can be adequately trapped in the axial direction. In a single-beam optical dipole trap, the axial trapping frequency is always lower than the radial trapping frequency, and if the axial trapping frequency is too small, then the beam must be aligned precisely horizontally, otherwise the atoms will fall out due to gravity. While this is no issue theoretically, in practice it is much easier to align a dipole trap if some small horizontal tilt is permissible.

For the case of a single-beam optical dipole trap, this criterion becomes

$$P > \frac{m\,\pi^3\,\omega_{z(\text{min})}^2}{4\,\alpha_s\,\lambda^4}\,w_0^6\,, \tag{5.9}$$

which is shown in Figure 5.9 as a 2D region plot of $P\,(w_0)$. This expression is independent of atom number. This relationship shows that for a given power, if the beam waist size becomes too large the condensate will not be sufficiently trapped in the axial direction. The 6th order scaling in waist size demonstrates that, similar to

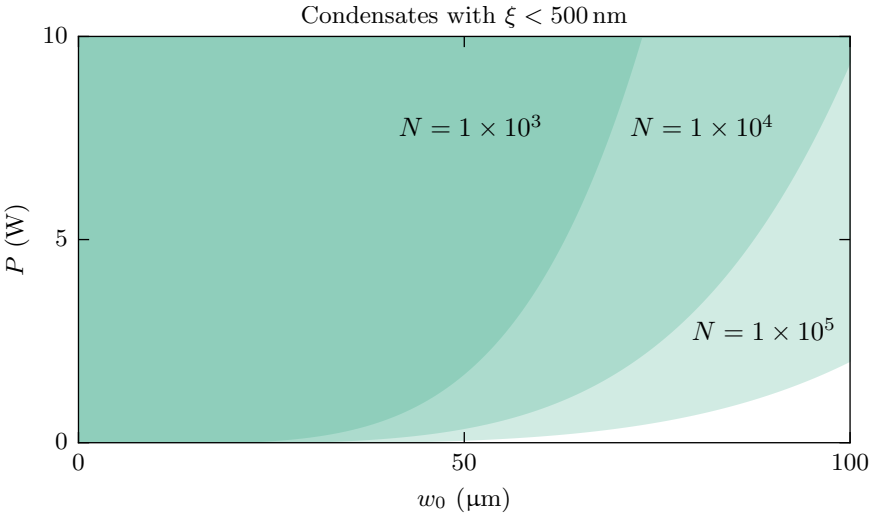

Figure 5.8: These regions show the combinations of P and w_0 for which a condensate trapped by a single-beam optical dipole trap has a healing length of less than 500 nm (Equation 5.8), with $N = 1 \times 10^3$, $N = 1 \times 10^4$ and $N = 1 \times 10^5$. The regions corresponding to lower atom numbers are contained within the regions of higher atom numbers.

the quasi-one dimensional criterion, setting a minimum axial trapping frequency constrains the allowed waist size far more than the beam power.

5.4 Pulse resolution criterion

As discussed in Section 4.1 and Section 4.1.1, to produce a single black soliton using magnetic resonance control we require a hyperbolic secant pulse that achieves a pulse resolution of $R \simeq 2\,\zeta\, z_{\mathrm{TF}}/5\,\xi$ (Equation 4.11), where typically $\zeta = 1.2$, in a time faster than one quarter of the condensate healing time with a pulse fidelity of at least 99.9%. If R is much smaller than this, the engineered density dip will be too wide, and decay into multiple solitons. If R is much larger than this, the engineered density dip will be too narrow, and expand to form a soliton by ejecting sound waves.

Combining Equations 4.11, 4.22 and 4.25, we can express this pulse resolution requirement as

$$\left(\frac{t_\xi}{4}\right)^2 \frac{\sqrt{2}\,\Gamma}{\left(\frac{2\,\Gamma}{\Omega_0}\,\cosh^{-1}\left(\alpha^{-1}\right)\right)^2} \simeq \frac{2\,\zeta\, z_{\mathrm{TF}}}{5\,\xi}\,. \tag{5.10}$$

Making the dependence on N, η and ω_z explicit, the above equation becomes

$$\eta \simeq \left(\frac{\hbar^3\,\pi\,\Omega_0}{2\,\sqrt{2}\,\Gamma\,\left(\cosh^{-1}\left(\alpha^{-1}\right)\right)^2\,360\,m^{5/2}\,N\,g\,\zeta}\right)^{1/2} \frac{1}{\omega_z}\,. \tag{5.11}$$

In Section 4.1.2 we discussed that appropriate values of the adiabaticity parameter Γ and pulse truncation α for engineering a soliton using magnetic resonance

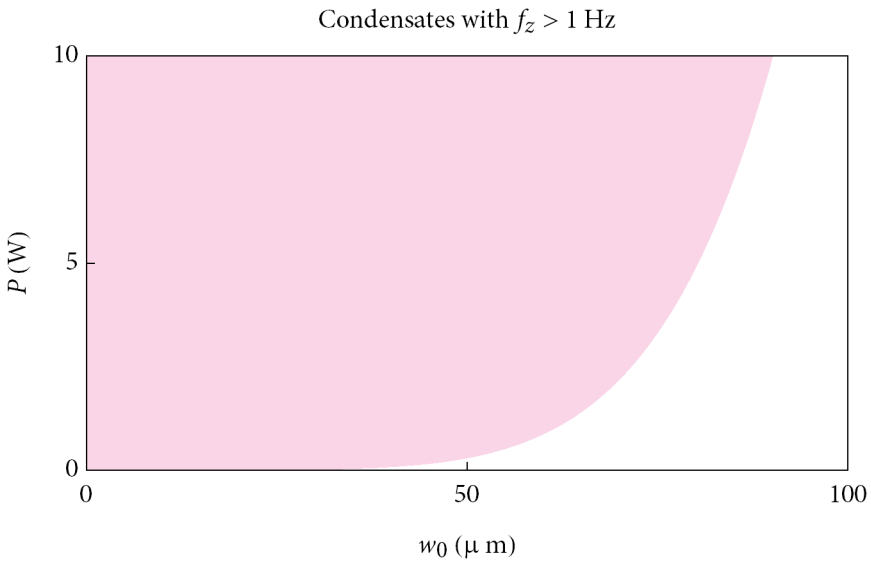

Figure 5.9: This region shows the combinations of P and w_0 for condensates trapped in a single-beam optical dipole trap with an axial frequency of $f_z > 1$ Hz (Equation 5.9). Note that this criterion is independent of atom number.

control are $\Gamma = 3$ and $\alpha = 0.003$. With these values set, the only free parameters in Equation 5.11 are N, η, ω_z, and Ω_0. This allows us to show this criterion in Figure 5.10 as 3D surface plots of $N\,(f_z,\,\eta)$ for example values of Ω_0, and in Figure 5.11 as lines of $\eta\,(f_z)$ for example values of N with the Rabi frequency fixed at $\Omega_0 = 2\pi \times 300\,\text{kHz}$. We select this Rabi frequency in particular because this is the frequency used in the simulations presented in Chapter 4.

Figure 5.10 shows that using a higher Rabi frequency enables the choice of condensates with both higher aspect ratio and higher axial trapping frequency. Figure 5.11 shows that, for a selected Rabi frequency, a higher atom number condensate can only be used if it has a high aspect ratio and low axial frequency, or conversely a low aspect ratio and high axial frequency; resulting from the inverse relation between η and ω_z in Equation 5.11.

For the case of a condensate in a single-beam optical dipole trap, Equation 5.11 can be rewritten in terms of N, P and w_0 as

$$P \simeq \frac{\hbar^3\,\pi^2\,\Omega_0}{5760\,\sqrt{2}\,\Gamma\,\left(\cosh^{-1}\left(\alpha^{-1}\right)\right)^2\,m^{3/2}\,N\,g\,\zeta\,\alpha_s}\,w_0^4\,. \tag{5.12}$$

Again fixing $\Gamma = 3$ and $\alpha = 0.003$, this relation is shown in Figure 5.12 as 3D surface plots of $N\,(w_0,\,P)$ for example values of Ω_0, and in Figure 5.11 as lines of $P\,(w_0)$ for example values of N with fixed Ω_0. Figure 5.12 demonstrates that using a pulse with a higher Rabi frequency requires a higher power beam, while Figure 5.11 shows that for a fixed Rabi frequency, to increase the atom number we require either a lower beam power or larger waist.

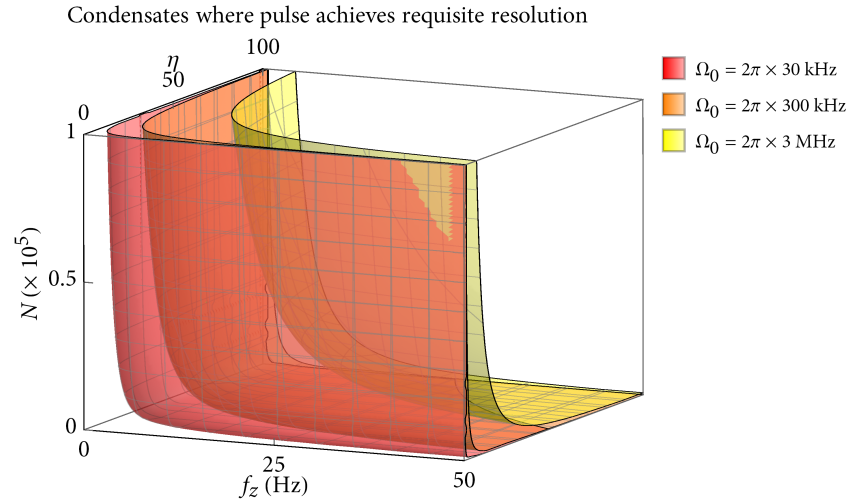

Figure 5.10: These surfaces show the combinations of N, η and f_z for condensates in which a hyperbolic secant pulse can achieve the resolution required to engineer a single soliton, for $\Omega_0 = 2\pi \times 30\,\text{kHz}$, $\Omega_0 = 2\pi \times 300\,\text{kHz}$, and $\Omega_0 = 2\pi \times 3\,\text{MHz}$. Lines of constant atom number from the surface corresponding to $\Omega_0 = 2\pi \times 300\,\text{kHz}$ are shown in Figure 5.11. The yellow triangular region near the upper vertex in the foreground is a rendering artefact, and should be ignored.

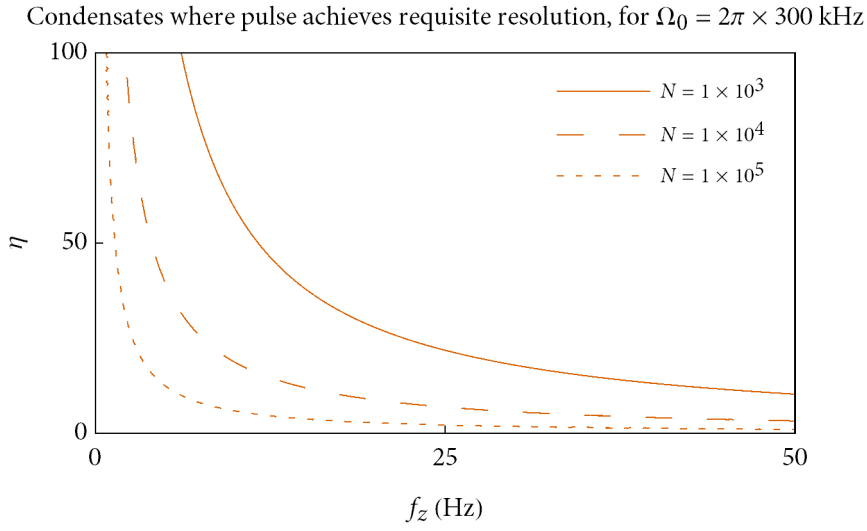

Figure 5.11: These lines show the combinations of η and f_z for condensates in which a hyperbolic secant pulse with Rabi frequency $\Omega_0 = 2\pi \times 300\,\text{kHz}$ can achieve the resolution required to engineer a single soliton, for $N = 1 \times 10^3$, $N = 1 \times 10^4$, and $N = 1 \times 10^5$.

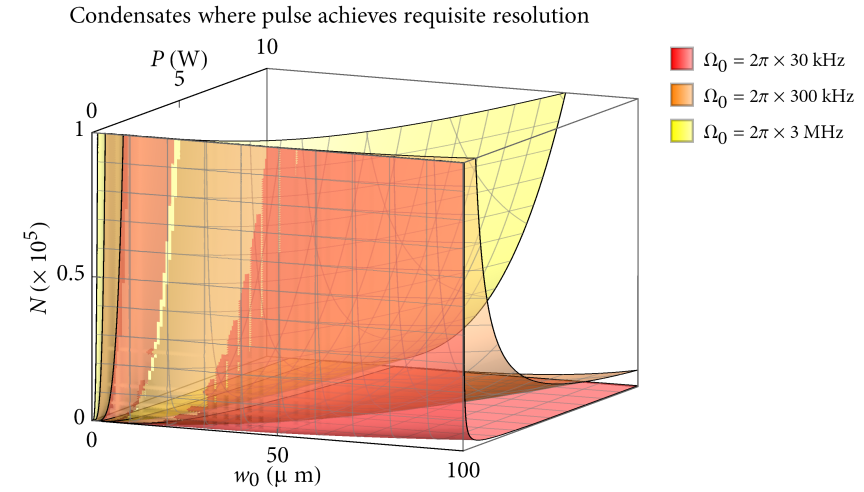

Figure 5.12: These surfaces show the combinations of N, P and w_0 for condensates trapped by a single-beam optical dipole trap in which a hyperbolic secant pulse can achieve the resolution required to engineer a single soliton, for $\Omega_0 = 2\pi \times 30\,\text{kHz}$, $\Omega_0 = 2\pi \times 300\,\text{kHz}$, and $\Omega_0 = 2\pi \times 3\,\text{MHz}$. Lines of constant atom number from the surface corresponding to $\Omega_0 = 2\pi \times 300\,\text{kHz}$ are shown in Figure 5.13. The yellow-orange banding on the left hand side of the figure is a rendering artefact, resulting from the closeness of the three surfaces in this region.

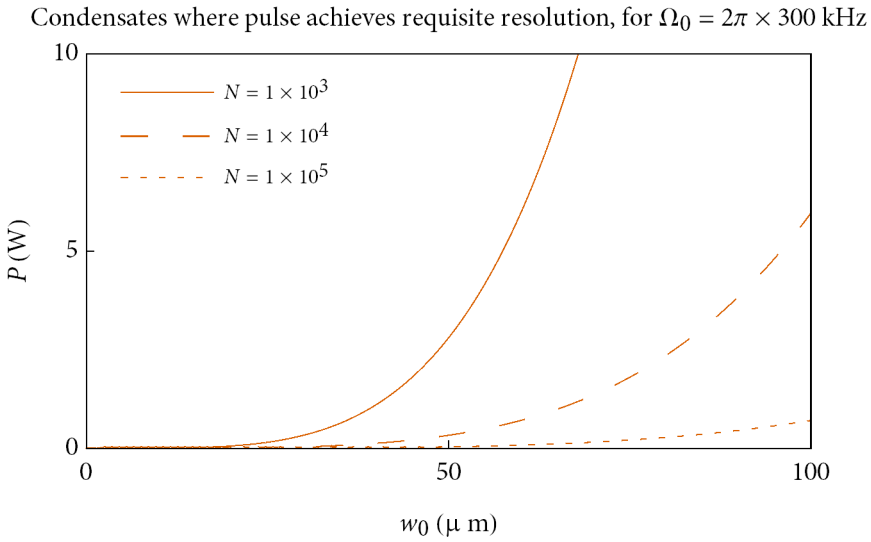

Figure 5.13: These lines show the combinations of P and w_0 for condensates trapped by a single-beam optical dipole trap in which a hyperbolic secant pulse with Rabi frequency $\Omega_0 = 2\pi \times 300\,\text{kHz}$ can achieve the resolution required to engineer a single soliton, for $N = 1 \times 10^3$, $N = 1 \times 10^4$, and $N = 1 \times 10^5$.

5.5 Intersection of all criteria

Having explored the condensate parameters which satisfy our four criteria for engineering a black soliton using magnetic resonance control, we can now observe the intersection of these four criteria; namely, the condensate parameters which will produce a quasi-one dimensional condensate with a sub-optical healing length, and an axial trapping frequency greater than 1 Hz, for which a hyperbolic secant pulse of duration $t_p = t_\xi/4$ can achieve the resolution required to engineer a black soliton using an experimentally feasible Rabi frequency.

The grey 3D solid in Figure 5.14 shows the region of $N\left(f_z, \eta\right)$ for condensates that are both quasi-one dimensional and have a healing length of $\xi < 500\,\text{nm}$. For the ranges of N and η chosen, this solid naturally satisfies the $f_z > 1\,\text{Hz}$ criterion. The orange surface shows the region of $N\left(f_z, \eta\right)$ for condensates in which a hyperbolic secant pulse with a Rabi frequency of $\Omega_0 = 2\pi \times 300\,\text{kHz}$ can achieve the resolution necessary to engineer a black soliton.

Figure 5.15 shows slices of the 3D objects in Figure 5.14, $\eta\left(f_z\right)$, for particular values of N. This figure shows that for the low atom number of $N = 1 \times 10^3$ there is a wide range of η and f_z which satisfy all four criteria; we can choose axial frequencies of $1\,\text{Hz} < f_z < 150\,\text{Hz}$ and aspect ratios of $3 < \eta < 400$, as long as they satisfy Equation 5.11. As the atom number is increased, the ranges of η and f_z become further restricted; at $N = 1 \times 10^4$ we can choose axial frequencies of $1\,\text{Hz} < f_z < 10\,\text{Hz}$ and aspect ratios of $20 < \eta < 150$, while for an atom number of $N = 3 \times 10^4$ we can choose axial frequencies of $1\,\text{Hz} < f_z < 3\,\text{Hz}$ and aspect ratios of $50 < \eta < 70$, provided they satisfy Equation 5.11.

Moving to the case of a condensate in a single-beam optical dipole trap, the grey 3D solid in Figure 5.16 shows the region of $N\left(w_0, P\right)$ for condensates that are quasi-one dimensional, have a healing length $\xi < 500\,\text{nm}$, and an axial trapping frequency $f_z > 1\,\text{Hz}$. The orange surface intersecting the solid shows the region of $N\left(w_0, P\right)$ for condensates in which a hyperbolic secant pulse with a Rabi frequency of $\Omega_0 = 2\pi \times 300\,\text{kHz}$ can achieve the resolution necessary to engineer a black soliton.

Figure 5.17 shows slices of the 3D objects in Figure 5.16, $P\left(w_0\right)$, for particular values of N. This figure shows that for the low atom number of $N = 1 \times 10^3$ there is a wide range of P and w_0 which satisfy all four criteria; we can choose waist sizes of $1\,\mu\text{m} < w_0 < 100\,\mu\text{m}$ and beam powers of $\eta < 50\,\text{W}$, as long as they satisfy Equation 5.12. As the atom number is increased, the ranges of w_0 and P become further restricted; for $N = 1 \times 10^4$ we can choose waist sizes of $10\,\mu\text{m} < w_0 < 30\,\mu\text{m}$ and beam powers of $P < 40\,\text{mW}$. For $N = 1 \times 10^5$, we would need to use a higher Rabi frequency to engineer a black soliton in the condensate, since the pulse resolution criterion using $\Omega_0 = 2\pi \times 300\,\text{kHz}$ does not intersect the quasi-one dimensional, sub-optical healing length region. In general, we can see that for condensates containing a reasonable number of atoms, we are pushed towards fairly tight traps with low trapping powers, which are achievable technical requirements.

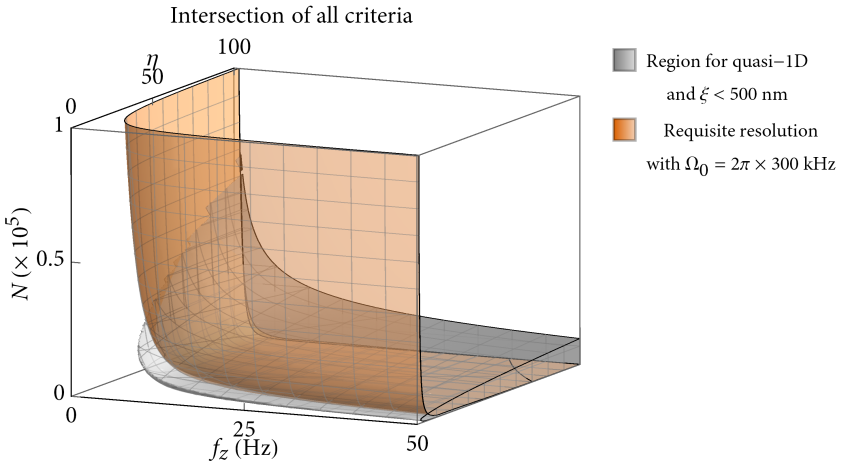

Figure 5.14: The grey solid shows the combinations of N, η and f_z for which a condensate is both quasi-one dimensional and has a sub-optical healing length. For the range shown in this figure, the grey solid naturally satisfies $f_z > 1$ Hz. The orange surface is the same as that shown in Figure 5.10; the surface of condensates for which a hyperbolic secant pulse with $\Omega_0 = 2\pi \times 300$ kHz achieves the resolution required to engineer a black soliton. Planes of constant atom number are shown in Figure 5.15.

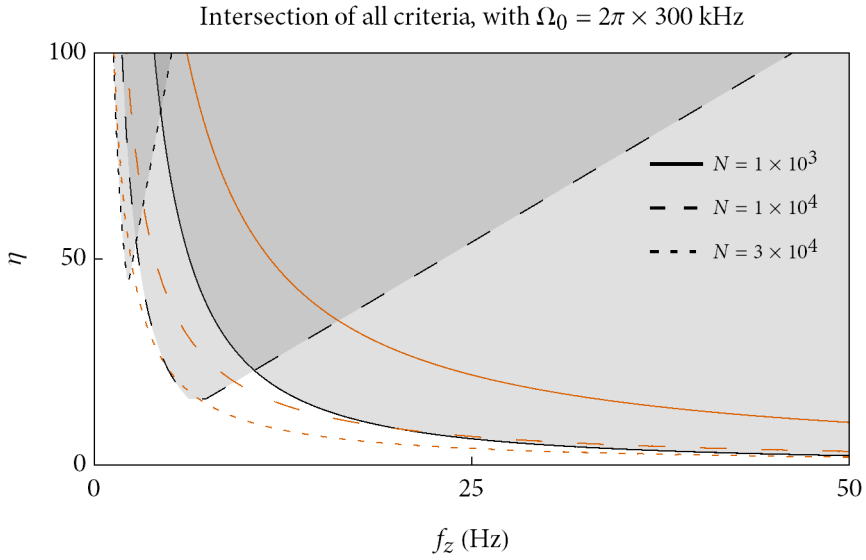

Figure 5.15: The grey regions show the combinations of η and f_z for quasi-one dimensional condensates with sub-optical healing length, for particular values of N (indicated by different outlines). The grey region shrinks with increasing atom number, with the quasi-one dimensional criterion encroaching from the right, and the sub-optical healing length criterion from the left. The orange lines indicate condensates for which a hyperbolic secant pulse with $\Omega_0 = 2\pi \times 300$ kHz achieves the resolution required to engineer a black soliton, again for particular values of N.

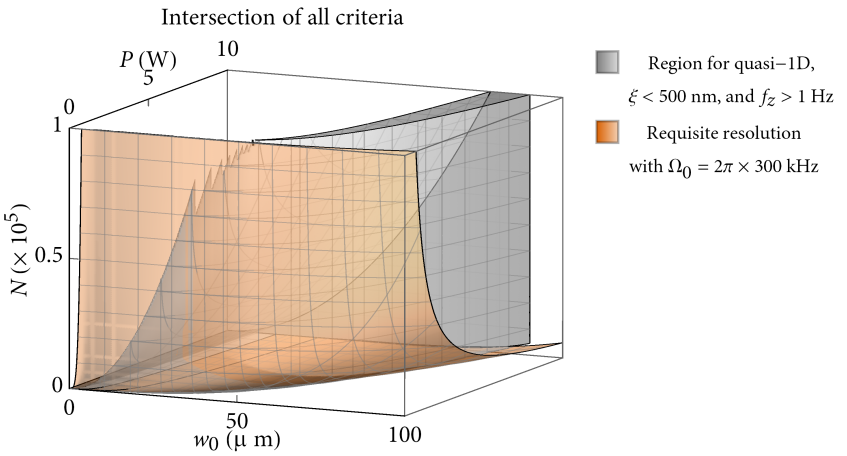

Figure 5.16: The grey solid shows the combinations of N, P and w_0 for which a condensate in a single-beam optical dipole trap is quasi-one dimensional, has a sub-optical healing length, and an axial frequency $f_z > 1$ Hz. Sharp oscillations in the grey solid are rendering artefacts and should be ignored. The orange surface is the same as that shown in Figure 5.12; the surface of condensates for which a hyperbolic secant pulse with $\Omega_0 = 2\pi \times 300$ kHz achieves the resolution required to engineer a black soliton. Planes of constant atom number are shown in Figure 5.17.

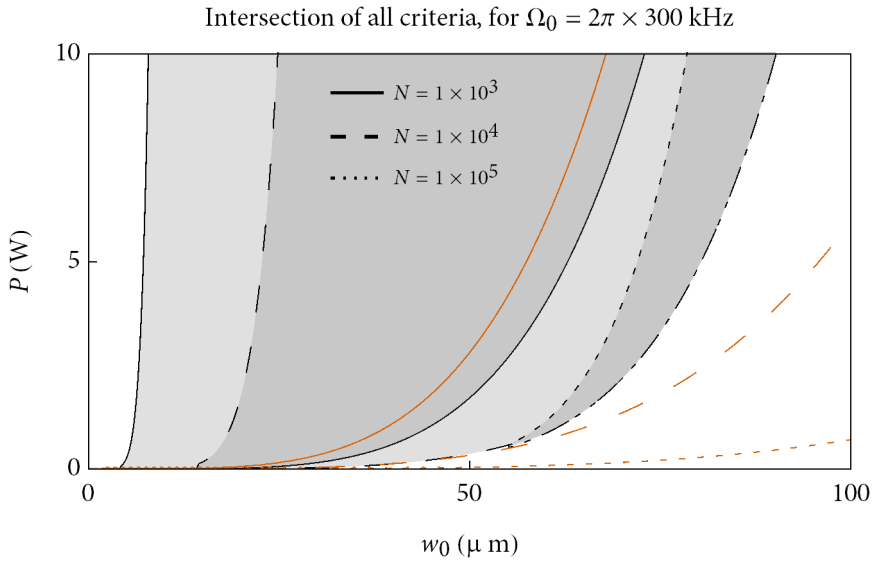

Figure 5.17: The grey regions show the combinations of P and w_0 for quasi-one dimensional condensates with sub-optical healing length and axial frequency $f_z > 1$ Hz, for particular values of N (indicated by different outlines). The grey region shrinks with increasing atom number, with the sub-optical healing length criterion encroaching from the right, and the quasi-one dimensional criterion from the left. The orange lines indicate condensates for which a hyperbolic secant pulse with $\Omega_0 = 2\pi \times 300$ kHz achieves the resolution required to engineer a black soliton, again for particular values of N. For $N = 1 \times 10^5$, the orange line no longer intersects with the grey region; for such large condensates, a higher Rabi frequency would be necessary to engineer a black soliton in the condensate.

5.6 Summary

The results presented in this chapter can be used to guide the choice of condensate parameters for those interested in experimentally implementing magnetic resonance control in order to engineer a black soliton in a condensate. The four criteria we require for the choice of condensate parameters to be able to engineer a black soliton with MRC are: the condensate should be quasi-one dimensional, for soliton stability; the healing length should be sub-optical, otherwise optical wavefunction engineering techniques might be utilised instead; the axial trapping frequency should be greater that 1 Hz, for ease of trap alignment; and a hyperbolic secant pulse of duration $t_p = t_\xi/4$ should be able to achieve the resolution required to engineer a black soliton in the condensate.

Fortunately, we have shown that the intersection of the four criteria forms a comfortably non-vanishing region. This demonstrates that magnetic resonance control should be applicable to a variety of experimental condensate apparatuses, rather than a single 'magic' set of values for atom number, axial frequency and aspect ratio.

Chapter 6

Magnetic resonance control in the laboratory

This chapter presents magnetic resonance control experiments performed using our spinor Bose–Einstein condensate apparatus. At the date of these experiments, the apparatus was incapable of producing magnetic field gradients strong enough to create the solitons described in the previous chapter. Remedying this would require either the installation of additional coils, applying a synthetic gradient using vector light shifts, or rapidly offsetting the quadrupole trap (discussed further in Chapter 7). I did not explore any of these approaches at the time, instead choosing to perform proof-of-principle experiments in the regime of small magnetic field gradients.

First, we explore how to individually control the spin projection of two spatially-separated condensates using a weak magnetic field gradient and a square-profile pulse (a finite-duration coupling pulse with uniform amplitude and frequency, ubiquitous in cold atom physics). Next, we use a hyperbolic secant pulse instead of a square-profile pulse and show that we retain the ability to individually control the two condensates. Finally, we observe the power of hyperbolic secant pulses in magnetic resonance control by changing to a regime where a square-profile pulse cannot address two condensates individually, but a hyperbolic secant pulse still can.

6.1 Square-profile coupling pulses

In Section 4.1.1 we discussed hyperbolic secant coupling pulses in the context of magnetic resonance control, and briefly mentioned the advantages of such a profile over the more common square-profile coupling pulse (in that discussion, specifically a π-pulse, which completely inverts the condensate spin projection). Since in this chapter we will examine the application of square-profile pulses to magnetic resonance control, we first discuss coupling three-level atoms with square-profile pulses.

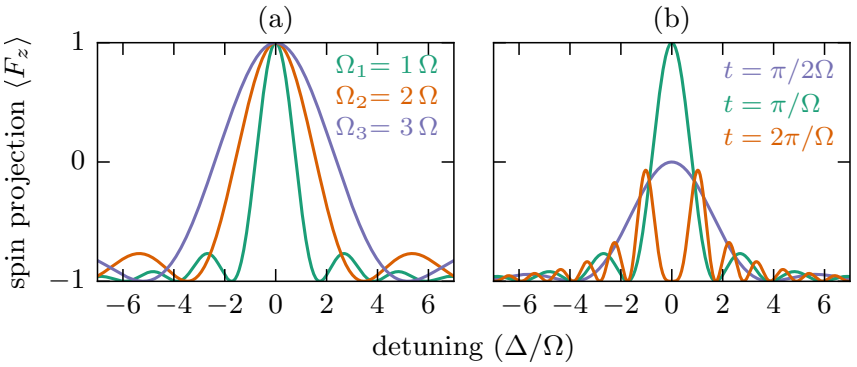

Figure 6.1: Examples of Rabi spectra, formed by plotting the spin projection F_z after a square-profile coupling pulse as a function of the detuning of the pulse from resonance. For the three curves shown in (a), the coupling is always a π-pulse ($t_1 = \pi/\Omega_1$, $t_2 = \pi/\Omega_2$, and $t_3 = \pi/\Omega_3$) so that the spin projection at resonance is $F_z = +1$. We see that the width of a Rabi spectrum increases with greater Rabi frequency Ω. The three curves shown in (b) show the effect of changing the pulse duration while using the same Rabi frequency. If the pulse duration is not $t_\pi = \pi/\Omega$, the spin projection at resonance reduced.

Resonant coupling between two atomic states with single-frequency, constant amplitude radiation causes the spin projection F_z of the condensate to oscillate sinusoidally in time; these are called Rabi oscillations. For example, consider a condensate within the $F = 1$ Zeeman manifold. In a low magnetic field, the splitting between the $m_F = -1$ and $m_F = 0$ states ($|-1\rangle$ and $|0\rangle$) is very similar to the splitting between the $m_F = 0$ and $m_F = +1$ states ($|0\rangle$ and $|+1\rangle$).[1] As such, coupling resonant to these (equal) transitions will cause the spin projection to oscillate between $F_z = \pm 1$. The frequency of this oscillation is called the Rabi frequency Ω. We can invert the condensate spin projection, from $F_z = +1$ to $F_z = -1$ and vice versa, by applying a π-pulse: a resonant coupling pulse with duration $t_\pi = \pi/\Omega$.

[1] In a high magnetic field, the quadratic Zeeman shift causes this splitting to become uneven.

A general expression for the final spin projection of a three-level condensate after a square-profile pulse, where the spin projection was initially $F_z = -1$, is

$$F_z = \frac{2\,\Omega^2}{\Omega^2 + \Delta^2} \sin^2\left(\left(\sqrt{\Omega^2 + \Delta^2}\right)\frac{t}{2}\right) - 1\,, \tag{6.1}$$

where $\Delta = \omega - \omega_0$ is the detuning from resonance [118].

To see how a square-profile coupling pulse can be used in the context of magnetic resonance control, consider the slice profile of such a pulse, or in other words the shape of the graph of spin projection as a function of detuning from resonance. Using Equation 6.1, example slice profiles are shown in Figure 6.1. These slice profiles are called *Rabi spectra*. The frequency corresponding to the maximum spin projection of a Rabi spectrum is called the *Rabi peak*, and the frequencies corresponding to the lowest spin projections are called *Rabi minima*.

We see that the slice profile of a π-pulse contains lobes on either side of the main slice; at certain detunings the spin projection of the condensate will remain in $|-1\rangle$, but at other detunings we see a change in the spin projection. This is not ideal for magnetic resonance control, since it means that we are unable to use this pulse to

address a specific localised region of the condensate. We saw in Section 4.1.1 that a hyperbolic secant pulse produces a slice profile in frequency space more similar to a top-hat function with no side lobes, which allows us to address one specific region of the condensate.

In addition to the issue that the side lobes introduce, a square-profile pulse does not have a resolution R sufficient to address one side of the condensate with healing-length resolution. Recall from Equation 4.6 that the resolution of the slice profile is the ratio of the FWHM of the transferred slice (the slice thickness) to the 10%–90% rise distance of the edge of the slice (the slice sharpness). Using Equation 6.1, the resolution of a π-pulse is $R \approx 1.6$. This is an order of magnitude smaller than the resolution achievable with a hyperbolic secant (remember that a larger value of R is better for magnetic resonance control). This shows that a π-pulse cannot produce a slice profile which both has a large slice thickness and sharp edges.

Nevertheless, one of the simplest experiments demonstrating magnetic resonance control—individually addressing the spin of two spatially-separated condensates—can be achieved using a π-pulse.

In a magnetic field gradient, two condensates A and B separated by a distance Δz will have a different resonance frequency; $\omega_{0\,(A)}$ and $\omega_{0\,(B)}$, respectively. The difference between these two frequencies is

$$\omega_{0\,(B)} - \omega_{0\,(A)} = \frac{\mathrm{d}\,|B|}{\mathrm{d}z}\,\gamma\,\Delta z\,. \tag{6.2}$$

We seek to align the Rabi peak of condensate B with the first Rabi minimum of condensate A. To do this, we need to locate the Rabi minima of condensate A. By inspection, Equation 6.1 satisfies this condition (by reducing to $F_z = -1$) when

$$\Delta_n = \Omega\sqrt{(2n)^2 - 1}\,, \tag{6.3}$$

where $n \in \mathbb{Z}^+$. Hence, the first Rabi minimum corresponds to $\Delta_1 = \sqrt{3}\,\Omega$.

Combining Equations 6.3 and 6.2, the Rabi peak of condensate B will coincide with the first Rabi minima of condensate A if the condensate separation and magnetic field gradient satisfy the condition

$$\frac{\mathrm{d}\,|B|}{\mathrm{d}z}\,\gamma\,\Delta z = \sqrt{3}\,\Omega\,. \tag{6.4}$$

If Equation 6.4 is satisfied, a π-pulse resonant with the centre of condensate A will not tip the spin of the centre of condensate B (Figure 6.2). To avoid also coupling the edges of condensate B, we require the condensate separation to be large enough so that the whole condensate 'fits' within the first Rabi minima; otherwise there would be a significantly non-uniform spin-tip across the width of each condensate. In other words, since the resolution R of a π-pulse is around 1.6, the ratio of the condensate size to condensate separation needs to be much less than 1.6.

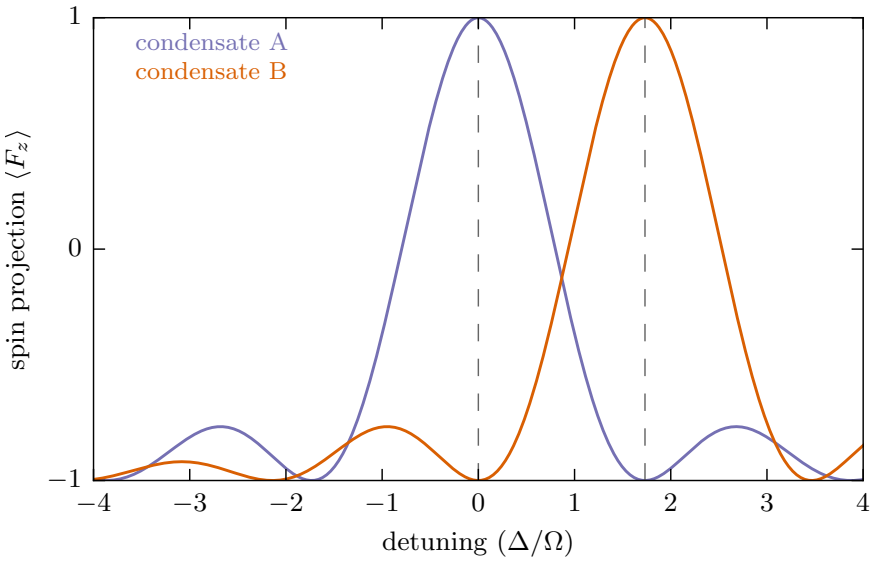

Figure 6.2: Rabi spectra for two spatially-separated condensates, where the detuning Δ is relative to the resonance frequency of condensate A. This system satisfies Equation 6.4 such that a π-pulse resonant with condensate A will not change the spin projection of condensate B.

6.2 *Experimental sequence for proof-of-principle MRC*

Our standard experimental sequence for making a single spinor Bose–Einstein condensate has been described in Alex Wood's thesis [61]. Here I describe the modification to this standard sequence to produce two spatially-separated condensates, and then the procedure used for magnetic resonance control experiments. This section includes hardware settings common to the presented results; values obtained from atomic measurements, such as magnetic field gradients, are presented in later sections alongside the experimental data.

6.2.1 *Making two condensates*

As discussed in Reference 61, to make a single spinor condensate we load atoms from the magnetic trap into two crossed dipole beams ('dipole beam 0' and 'dipole beam 1'), then evaporate to condensation. Dipole beam 0 has a waist size of 70 µm, and dipole beam 1 has a waist size of 100 µm [61]. To make two spatially-separated condensates, we can send a two-tone radiofrequency signal into the AOM controlling either dipole beam 0 or dipole beam 1, splitting the beam to form *two* crossed dipole traps. If the two crossed dipole traps are separated by less than 200 µm, atoms can be loaded directly from the magnetic trap into these two dipole traps, and then evaporated to condensation. If we require a larger separation, after evaporation the condensates can be moved further apart by smoothly increasing the difference between the two-tome radiofrequency signal driving the AOM. Figure 6.3 shows a schematic of our dipole traps and defines the lab coordinates and dipole

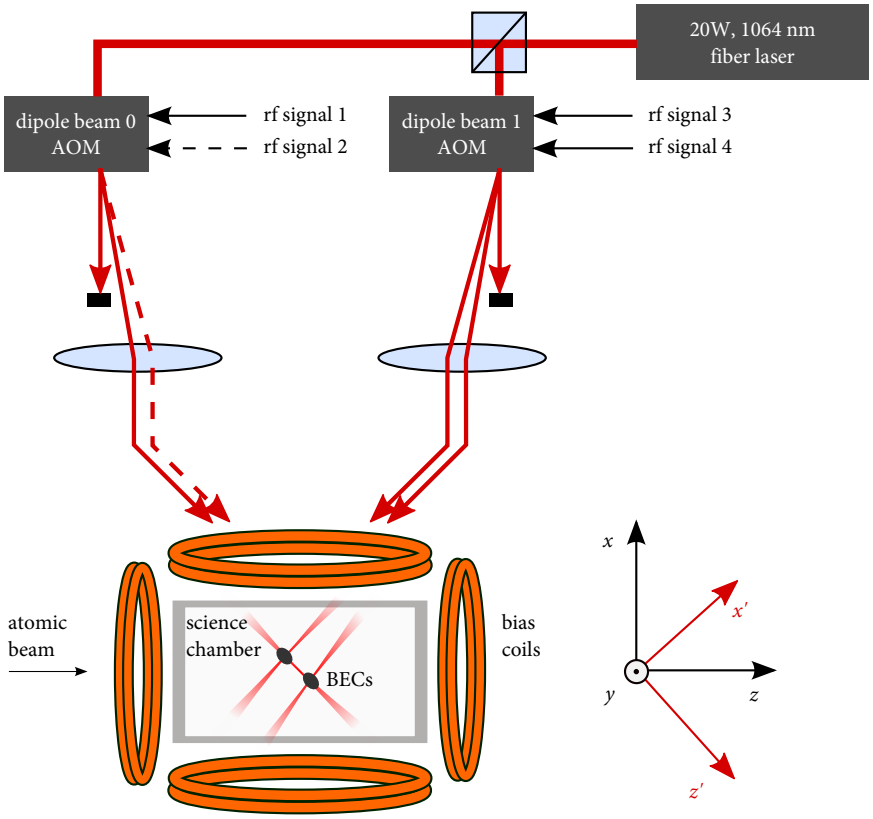

Figure 6.3: Schematic of the crossed dipole trap (not to scale) showing how we split our dipole beams to create two condensates. Only essential optical components are shown; free-space bends in beam path indicate excluded optics. We define a right-handed laboratory coordinate frame, where the origin is at the center of the science chamber, with gravity along $-y$ and z pointing in the direction of the atomic beam. The dipole beams are aligned at 45 degrees to the laboratory frame in the x-z plane. Dipole beam 0 points along z' while dipole beam 1 points along $-x'$. To form two condensates split along z' we drive the dipole beam 1 AOM with a two-tone radiofrequency signal.

beam coordinates used in the rest of this chapter. For the results presented in this chapter the two condensates are always separated along z', by sending a two-tone radiofrequency signal into the AOM controlling dipole beam 1.

We use a similar technique when performing magnetic tensor gradiometry to measure the magnetic gradient tensor at the centre of our science chamber. This technique was developed by my colleague Alex Wood, and I contributed to these efforts by collecting a large portion of the data used to measure the magnetic gradient tensor, earning my place as second author on our paper, Reference 119.

6.2.2 *Applying the magnetic gradient*

To create a magnetic field gradient we apply current through only one of the z-bias coils, while using the x- and y-bias coils to null the background magnetic field along x and y. The currents used to null the background magnetic field are $I_x =$

391.6 mA, $I_y = 311.8$ mA, $I_{z1} = 8.8$ A, and $I_{z0} = 0$. The procedure for finding these nulling currents is described in Reference 119. Applying a current through only one of the pair of z-bias coils produces the largest gradient possible without reversing the current in one coil. We keep the magnetic field offset to below 1 G, so that the quadratic Zeeman shift is negligible. For higher offsets, the population transfer between the three Zeeman sublevels becomes unbalanced, reducing the maximum spin-tip (or pulse fidelity) possible using a π-pulse.

6.2.3 *Measuring the spin projection*

After ramping the magnetic field gradient gradually over 10 ms to prevent exciting the condensates into sloshing modes, and waiting 75 ms for ringing in the coils to fade, we apply either a square-profile pulse or a hyperbolic secant pulse (depending on the experiment performed), which might change the spin projection F_z of either or both condensates. To detect the effect of the coupling pulse, we measure the spin projection of each condensate.

We measure the spin projection by separating the three spin components of each condensate. We release the condensates from the trap, and 11 ms later turn on a 60 G/cm magnetic field gradient for 3 ms using the quadrupole coils. As the condensates fall, they experience a state-selective force from the magnetic gradient (the Stern–Gerlach effect). The three spin components of each component separate spatially to form six distinct clouds. We then take an absorption image of these clouds.

To form one image which shows the optical depth of the clouds, we record three images: one with the imaging light turned off (the 'dark field' image), one with the imaging light turned on with no atoms present (the 'flat field' image), and one with the imaging light passing through the atoms (the 'atom' image). From this data, the optical depth (OD) is calculated by

$$\mathrm{OD} = \ln\left(\frac{I_{\mathrm{atom}} - I_{\mathrm{dark}}}{I_{\mathrm{flat}} - I_{\mathrm{dark}}}\right), \tag{6.5}$$

where I_{x} is the intensity from the 'x' image. To minimise spatial fringes in the OD image from vibrations of optical components, the flat field and atomic frames are taken in rapid succession, as close together in time as possible. With our equipment, this interframe time is 40 μs; details of how we achieve this rapid "back-to-back" imaging are discussed in Reference 61.

To measure the spin projection of one condensate, we first measure the total atom number N_i in each spin component i, from which we can calculate the spin projection:

$$F_z = \frac{N_{+1} - N_{-1}}{N_{+1} + N_0 + N_{-1}}. \tag{6.6}$$

To measure N_i, we fit a Thomas–Fermi distribution to each spin component; hence the uncertainty in F_z is a result of the uncertainty in this atom number fit. We find that our fitting algorithm becomes less accurate for very low atom numbers. To account for this, if the uncertainty in the measured atom number is more than

30% we decide to record the atom number in that spin component as zero, with an uncertainty of 500 atoms (a typical uncertainty for successful atom number fits).

6.3 Individually addressing two condensates with a square-profile pulse

Here we use experimental parameters satisfying Equation 6.4 to successfully tip the spin of one condensate, while leaving the second condensate unchanged.

To find these suitable parameters, I first applied a differential current through one pair of bias coils, aiming for the largest gradient achievable with these coils while limiting the magnetic field offset to around 1 G (to keep the quadratic Zeeman shift negligible). I then set the condensate separation to be $\Delta z' = 840\,\mu\text{m}$; this is the largest separation we can achieve by splitting our optical dipole trap without the power loss through the AOMs significantly reducing the atom number in each condensate.

From previous calibrations of our bias coils, I could predict approximately what Rabi frequency Ω_0 would be required to achieve a π-pulse in this magnetic field. With this initial Rabi frequency, I found the resonance frequency ω_0 of one of the condensates by measuring the Rabi spectrum of the condensate. With the resonance frequency found, I then varied the pulse duration until I found the correct duration t_π (the pulse which maximised the spin projection) for this condensate with this Rabi frequency in this magnetic field environment.

I then measured the Rabi spectra of both condensates and observed two well-separated Rabi spectra. Unsurprisingly, the separation of the Rabi spectra was not exactly what we desired; the first minimum of one curve did not align with the peak of the other curve, they were separated by more than was necessary.

To remedy this, I reduced the physical separation Δz between the condensates and measured the Rabi spectra again, until the first Rabi minimum of condensate A coincided with the Rabi peak of condensate B. These measurements are shown in Figure 6.4. This allows us to tip the spin of one condensate while leaving the other unchanged; the first proof-of-principle demonstration of magnetic resonance control. Here the condensate spacing was $\Delta z' = 288.7(9)\,\mu\text{m}$, as measured in Figure 6.5(a). The uncertainty in this measurement is from the uncertainty in the centroid position from the Thomas–Fermi fit to the optical depth, and from the uncertainty in the magnification of the imaging system (which was measured by tracking the centre-of-mass motion of the condensate during free-fall for several known drop times). At this separation, the difference in resonance frequencies was $2\pi \times 4.43(4)\,\text{kHz}$. This gives a measure of the applied magnetic gradient: from Equation 6.2, the gradient along z' is $219(2)\,\text{mG/cm}$. The average of the Rabi frequency measured with each condensate is $2\pi \times 2.56(5)\,\text{kHz}$.

These measured values for the magnetic field gradient, condensate separation and Rabi frequency satisfy Equation 6.4 within uncertainty. Particular examples of absorption images from the two Rabi curves, shown in Figure 6.5, demonstrate our ability to individually address the two condensates with just a square-profile pulse.

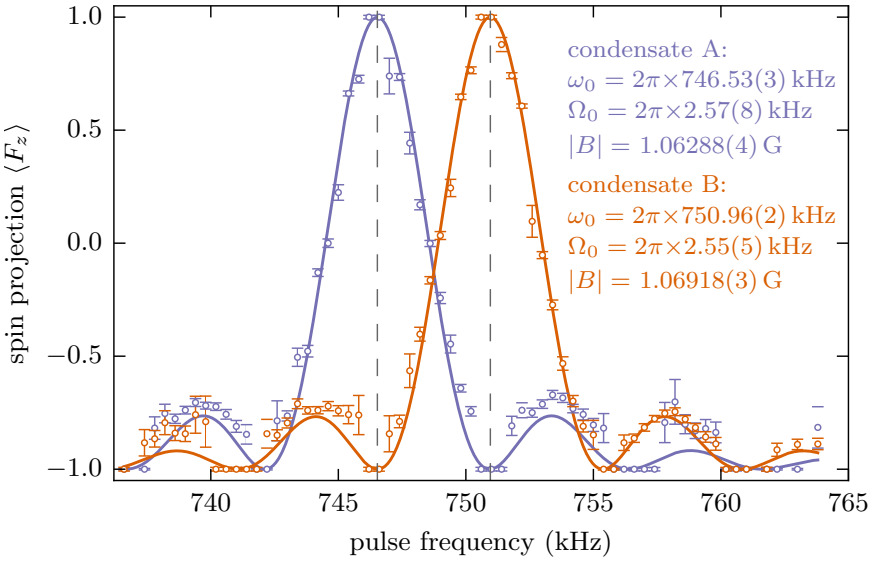

Figure 6.4: Using a square-profile pulse to individually address two condensates. These are the Rabi spectra of two condensates spatially separated by $\Delta z' = 288.7(9)$ µm in a magnetic gradient of $\mathrm{d}\,|B|\,/\mathrm{d}z' = 219(2)$ mG/cm. The duration of the pulse was $t_\pi = 196$ µs. Uncertainties in spin projection F_z are due to fits to atom number in each spin component, as discussed in Section 6.2.3. Solid lines are fits using Equation 6.1, with fit results inset. Dotted lines mark the measured resonance frequencies of each condensate, showing that the resonance frequency of condensate B coincides with the first Rabi minima of condensate A, and vice versa.

The spacing necessary to allow individual addressing of these condensates was larger than the separation of the crossed dipole traps when the condensates were formed, so after reaching condensation we needed to move the condensates further apart. To avoid introducing a sloshing mode while increasing the separation, the radiofrequency signals driving the dipole beam 1 AOM were smoothly ramped from 96.1 MHz to 95.7 MHz and from 105.1 MHz to 106.5 MHz, respectively, over two seconds.

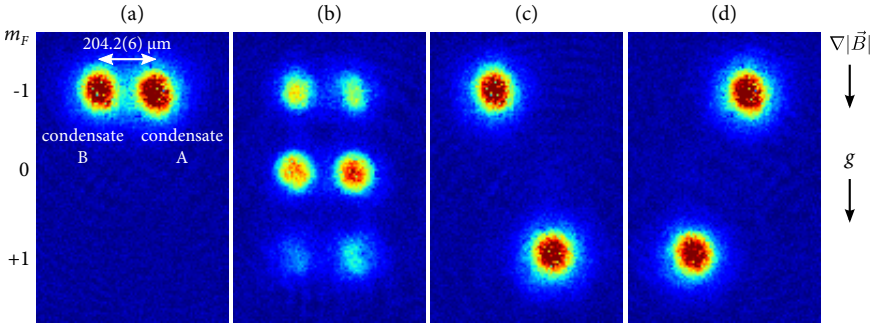

Figure 6.5: Examples of absorption images from the data shown in Figure 6.4. These images were taken after 20 ms of free-fall. (a) The two condensates after no pulses were applied, showing that the condensates are formed in $|-1\rangle$ (with measured spin projections of $F_{z(\mathrm{A})} = -1.000(3)$ and $F_{z(\mathrm{B})} = -1.000(4)$). Thomas–Fermi fits to the condensates determined a separation of $\Delta z = 204.2(6)$ µm, corresponding to $\Delta z' = 288.7(9)$ µm (since we have a 45° oblique view). (b) A π-pulse with frequency $\omega_0 = 2\pi \times 748.6$ kHz (mid-way between the resonance frequencies of the condensates) addresses both condensates, populating all three spin components ($F_{z(\mathrm{A})} = 0.00(1)$ and $F_{z(\mathrm{B})} = 0.16(2)$). (c) A π-pulse with frequency $\omega_0 = 2\pi \times 746.53$ kHz tips the spin of condensate A while leaving condensate B unchanged ($F_{z(\mathrm{A})} = 1.00(1)$ and $F_{z(\mathrm{B})} = -1.00(1)$). (d) A π-pulse with frequency $\omega_0 = 2\pi \times 750.98$ kHz tips the spin of condensate B while leaving condensate A unchanged ($F_{z(\mathrm{A})} = -1.00(1)$ and $F_{z(\mathrm{B})} = 0.99(1)$).

.

6.4 *Individually addressing two condensates with a hyperbolic secant pulse*

Here we demonstrate that under the same conditions a hyperbolic secant pulse can also individually address the two condensates. Equations 4.4a and 4.4b define the Rabi frequency $\Omega(t)$ and detuning $\Delta(t,\, z)$ of the hyperbolic secant pulse. We again use the dimensionless pulse parameters μ, Γ, and α to describe the hyperbolic secant pulse, as defined in Equations 4.14, 4.15, and 4.16, respectively.

We use the same magnetic field gradient, condensate separation and peak Rabi frequency used in Section 6.3 to enable a direct comparison between using a hyperbolic secant pulse and a square-profile pulse. These values are $\Delta z' = 288.7(9)$ µm, $\mathrm{d}\,|B|\,/\mathrm{d}z' = 219(2)$ mG/cm, and $\Omega_0 = 2\pi \times 2.56(5)$ kHz.

We select hyperbolic secant pulse parameters of $\Delta_0 = 2\pi \times 2.5$ kHz (corresponding to $\mu = 0.98(2)$), $\Gamma = 40$ and $\alpha = 0.003$. This pulse has a duration of $t_p = 33.1$ ms. These parameters were chosen so that the slice thickness was similar to the thickness of the π-pulse used in Section 6.3. Here the slice sharpness (as calculated for a line of stationary spins) is 4.9 µm, which is smaller than the width of each condensate.

Figure 6.6 shows the measured hyperbolic secant pulse spectra for the two condensates using these pulse parameters, along with theoretical curves. We see that the general form of the hyperbolic secant pulse spectrum is as expected; a rectangular profile without any side lobes. However, the slice sharpness is much broader than expected, as observed in Figure 6.7. We attribute this to the long pulse dura-

tion which makes us sensitive to magnetic field fluctuations during the pulse.

The dominant source of magnetic field fluctuations in our laboratory is the AC line cycle, which oscillates at 50 Hz (a period of 20 ms). We typically minimise the effects of this by synchronising the start of the coupling pulse with a zero-crossing of the AC line cycle. This is sufficient for short coupling pulses, but this hyperbolic secant pulse has a duration longer than the AC line oscillation period. This change in the background magnetic bias field during the pulse has the effect of blurring the position of the slice edge, causing the measured slice edge to be broader than expected.

The other possible explanation of the broad slice edge is the fact that the pulse duration is several orders of magnitude longer than the healing time of the condensate (typically on the order of 100 μs) which makes the condensate fluid flow significant during the pulse, as discussed in Section 4.4.2.

Even though the slice edge is broader than we aimed for, this hyperbolic secant pulse successfully tips the spin of one condensate while leaving the other unchanged. In fact, we can see from Figures 6.6 and 6.7 that there is a much wider frequency domain (3.5 kHz) in which this can be achieved, compared to using a square-profile pulse. This means that a hyperbolic secant pulse is dramatically more robust than a square-profile pulse against fluctuations in magnetic field and rf coupling strength. Key examples of absorption images from the data in Figures 6.6 and 6.7 are shown in Figure 6.8.

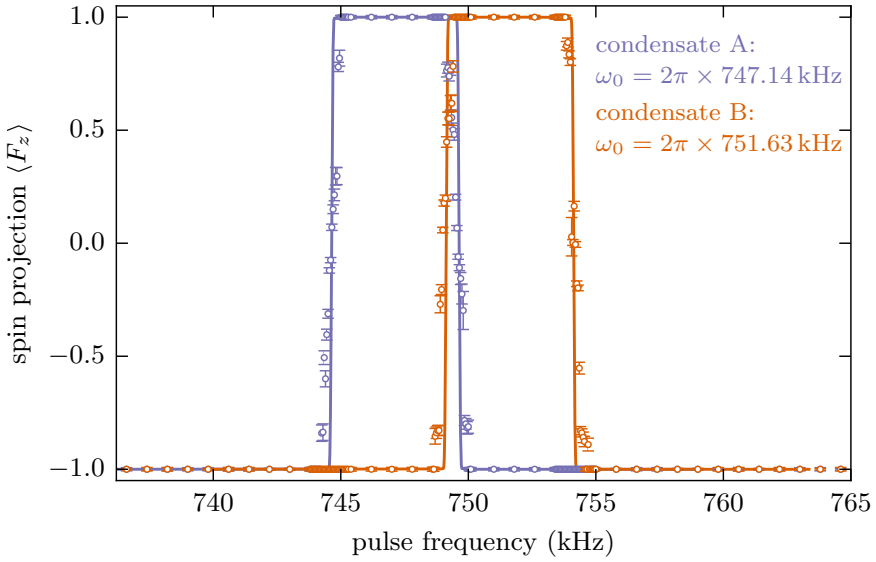

Figure 6.6: Using a hyperbolic secant pulse to individually address two condensates. Here the condensate separation and magnetic field gradient were the same as those used in Figure 6.4. The pulse parameters used were $\Delta_0 = 2\,\pi \times 2.5\,\mathrm{kHz}$, $\Gamma = 40$ and $\alpha = 0.003$. This data shows that there is a frequency region 3.5 kHz wide for each condensate in which the hyperbolic secant pulse can tip the spin of one condensate while leaving the other unchanged. Solid lines are the theoretical hyperbolic secant pulse spectra, with only the resonance frequency of each condensate (the horizontal offset) as a free parameter. This offset was found by eye for each condensate, and is stated in the figure. Later in Section 6.5 we use a least-squares fit to find the offset, but that algorithm fails here because the slice edge is broader than predicted by theory.

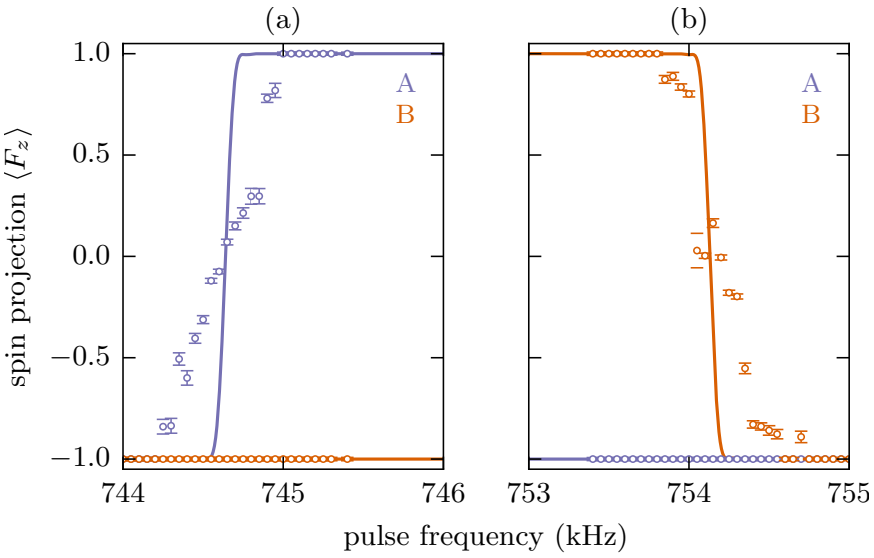

Figure 6.7: A subset of the data shown in Figure 6.6, revealing the shape of the edges of the hyperbolic secant pulse spectra. The edge sharpness of 0.7 kHz is larger than the theoretical sharpness of 0.1 kHz; we attribute this to the long pulse duration causing sensitivity to magnetic field fluctuations, and the fluid flow motion of the condensate, during the pulse. The uncertainty in Rabi frequency is insignificant, resulting in a change in the slice edge shape too small to be seen on this scale.

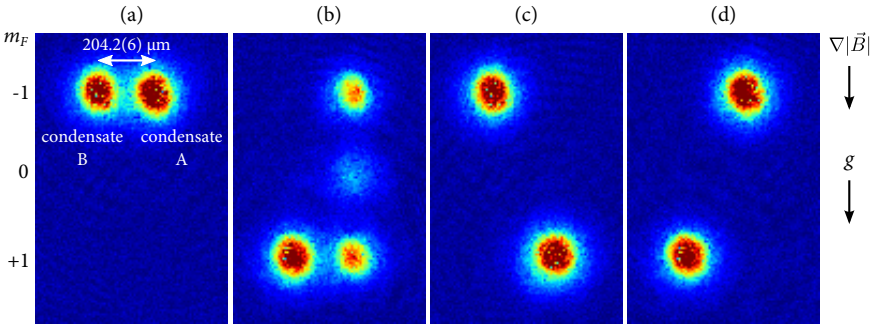

Figure 6.8: Examples of absorption images from the data shown in Figures 6.6 and 6.7. These images were taken after 20 ms of free-fall. (a) The condensate separation was the same as when we applied π-pulses in Section 6.3, as was the magnetic field gradient. (b) A hyperbolic secant pulse with centre frequency $\omega_0 = 2\,\pi \times 749.55$ kHz locates the slice edge at the centre of condensate A, partially tipping the spin of condensate A while completely tipping the spin of condensate B ($F_{z(\mathrm{A})} = 0.07(1)$ and $F_{z(\mathrm{B})} = 1.000(8)$). The state populations in condensate A do not follow the typical $[0.25, 0.5, 0.25]$ distribution that a π-pulse produces; it would be interesting to investigate this effect further. (c) A hyperbolic secant pulse with centre frequency $\omega_0 = 2\,]pi \times 747.00$ kHz tips the spin of condensate A while leaving condensate B unchanged ($F_{z(\mathrm{A})} = 1.000(7)$ and $F_{z(\mathrm{B})} = -1.000(8)$). (d) A hyperbolic secant pulse with centre frequency $\omega_0 = 2\,\pi \times 751.80$ kHz tips the spin of condensate B while leaving condensate A unchanged ($F_{z(\mathrm{A})} = -1.000(8)$ and $F_{z(\mathrm{B})} = 1.000(8)$).

6.5 A hyperbolic secant pulse succeeds where a square-profile pulse fails

Finally, we further demonstrate the advantages of using a hyperbolic secant pulse for magnetic resonance control instead of a square-profile pulse by moving to a regime where a square-profile pulse cannot individually address the two condensates but a hyperbolic secant pulse still can. We do this by increasing the Rabi frequency to $\Omega_0 = 2\pi \times 63(2)\,\mathrm{kHz}$, while keeping the same condensate separation and magnetic field gradient used in Section 6.3. As we expect from Figure 6.1, this broadens the Rabi spectra so that the resonance frequency of one condensate no longer coincides with the first Rabi minima of the other condensate (Figure 6.9).

Increasing the Rabi frequency has the additional benefit of reducing the duration of the pulse, reducing the effect of magnetic field fluctuations, and condensate fluid flow, during the pulse. The hyperbolic secant pulse parameters used here are $\Delta_0 = 2\pi \times 65\,\mathrm{kHz}$ (corresponding to $\mu = 1.03(3)$), $\Gamma = 40$ and $\alpha = 0.003$. The resulting pulse spectra for the two condensates are shown in Figure 6.10, with pulse edges shown in Figure 6.11. These demonstrate excellent agreement between measured data and theory, and show that unlike the square-profile pulse, the hyperbolic secant pulse can still individually address the two condensates. Additionally, using hyperbolic secant pulses to measure the condensate spectra produces a measurement of the resonance frequency, and hence the magnetic field, to an accuracy three orders of magnitude greater than can be achieved using square-profile pulses. Key examples of absorption images from the data in Figures 6.10 and 6.11 are shown in Figure 6.12.

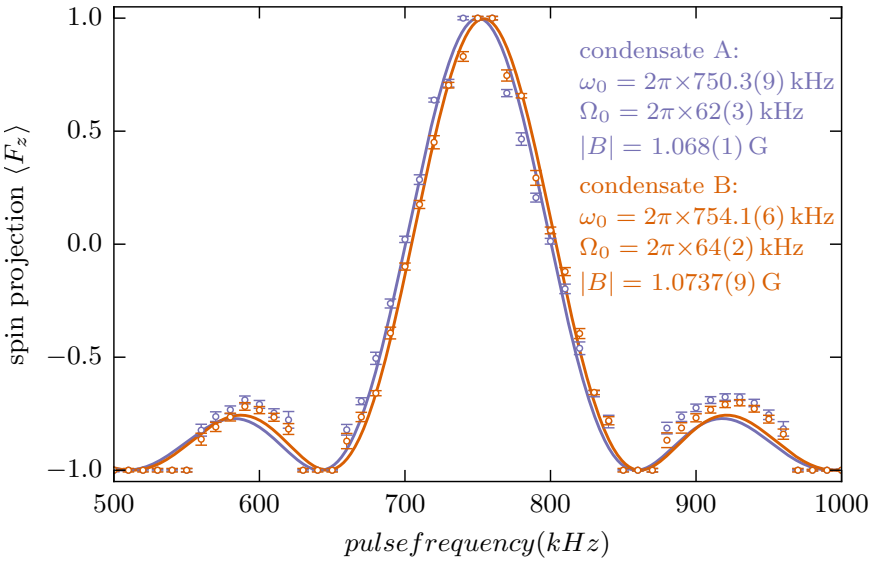

Figure 6.9: Increasing the Rabi frequency of the π-pulse increases the width of the Rabi spectra, so that for the same condensate separation and magnetic field gradient ($\Delta z' = 288.7(9)$ µm and $\mathrm{d}\,|B|\,/\mathrm{d}z' = 219(2)$ mG/cm) a π-pulse can no longer individually address the two condensates. The pulse duration was 8 µs. Solid lines are fits using Equation 6.1, with fit results inset.

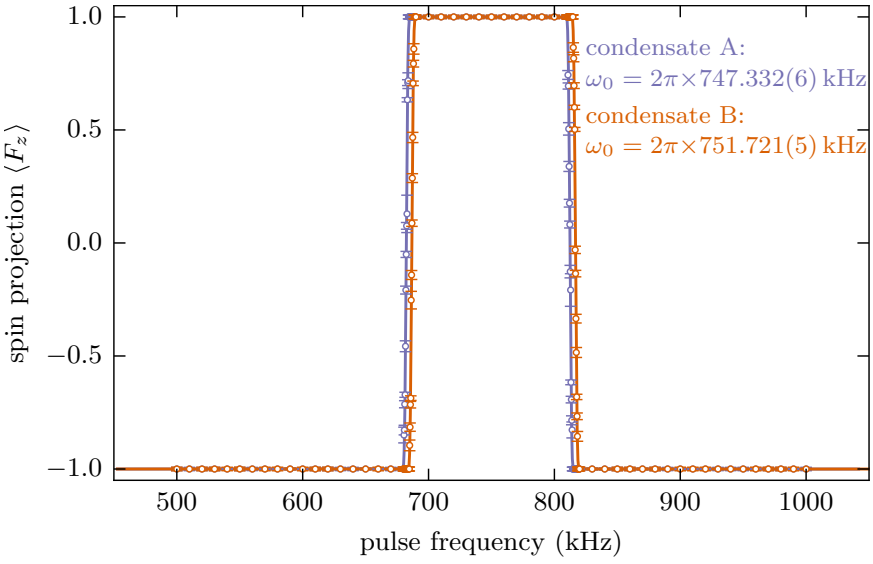

Figure 6.10: Hyperbolic secant pulse spectra using the pulse parameters $\Delta_0 = 2\,\pi\,\times\,65$ kHz, $\Gamma = 40$ and $\alpha = 0.003$. The pulse duration is 1.3 ms. Here the condensate separation and magnetic gradient is $\Delta z' = 288.7(9)$ µm and $\mathrm{d}\,|B|\,/\mathrm{d}z' = 219(2)$ mG/cm. Solid lines are the theoretical spectra, with only the centre frequency of each condensate left as a free parameter, found using a least-squares fitting algorithm, with fit results inset. There is a frequency region 0.9 kHz wide in which the spin of one condensate can be completely tipped without changing the spin of the other condensate.

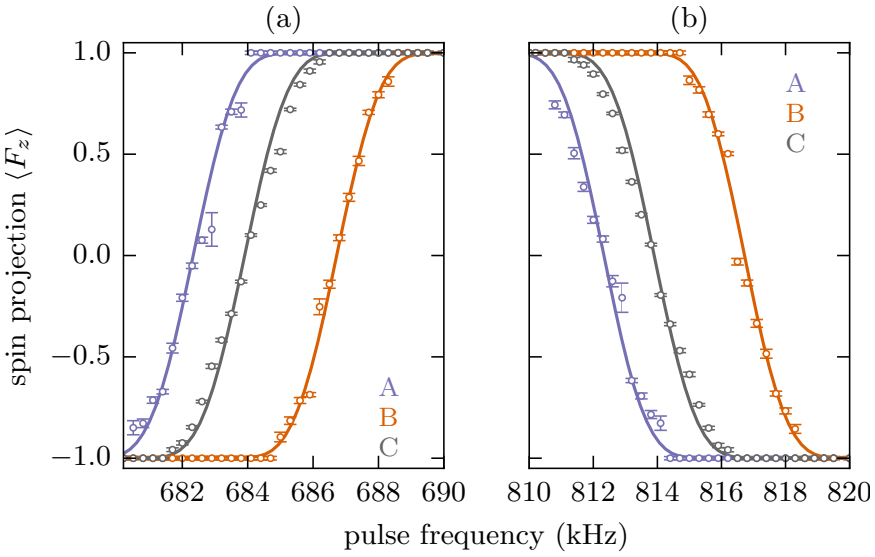

Figure 6.11: A subset of the data shown in Figure 6.10, revealing the shape of the edges of the hyperbolic secant pulse spectra. To better compare theory to experiment, included is data from condensate C, formed without splitting dipole beam 1. This results in a condensate with a higher atom number, and reduces the uncertainty in F_z. We can see that the data shows excellent agreement with theory; there is only a slight broadening of the slice edge, which we attribute to the pulse duration still being longer than the healing time of the condensate.

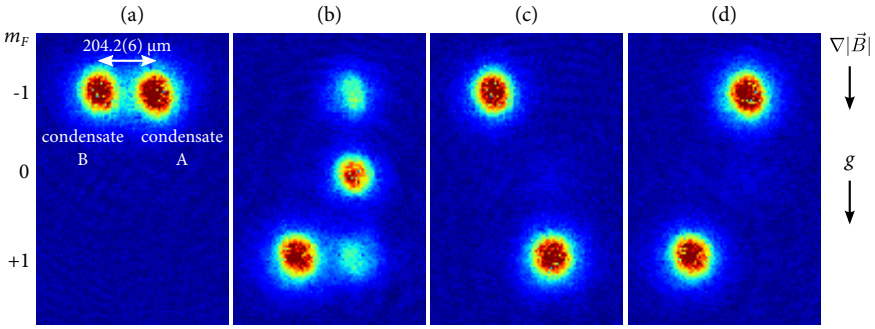

Figure 6.12: Examples of absorption images from the data shown in Figures 6.10 and 6.11. These images were taken after 20 ms of free-fall. (a) The condensate separation was the same as when we applied π-pulses in Section 6.3, as was the magnetic field gradient. (b) A hyperbolic secant pulse with centre frequency $\omega_0 = 2\,\pi \times 812.23$ kHz locates the slice edge at the centre of condensate A, partially tipping the spin of condensate A while completely tipping the spin of condensate B ($F_{z(\mathrm{A})} = 0.08(2)$ and $F_{z(\mathrm{B})} = 1.000(9)$). (c) A hyperbolic secant pulse with centre frequency $\omega_0 = 2\,\pi \times 684.40$ kHz tips the spin of condensate A while leaving condensate B unchanged ($F_{z(\mathrm{A})} = 1.000(8)$ and $F_{z(\mathrm{B})} = -1.000(9)$). (d) A hyperbolic secant pulse with centre frequency $\omega_0 = 2\,\pi \times 814.40$ kHz tips the spin of condensate B while leaving condensate A unchanged ($F_{z(\mathrm{A})} = -1.000(8)$ and $F_{z(\mathrm{B})} = 1.000(9)$).

6.6 Summary

These proof-of-principle experiments demonstrate that magnetic resonance control provides the ability to rotate the spin projection of one condensate without changing the spin projection of an adjacent condensate. The magnetic field gradients our apparatus is currently able to produce are weak compared to the gradients used in the simulations presented in Chapter 4. If we were to increase the magnetic field gradient, we would be able to selectively address regions within a single condensate with shorter duration hyperbolic secant pulses; one step closer to engineering a single black soliton in a Bose-Einstein condensate.

6.7 Record of experiment data files

We have a system in our research group of saving a `.h5` file for each experimental run (typically, each measurement of a condensate), containing a complete record of the hardware instructions and subsequent analysis results (such as the spin projection), with the filename in the format of a time and date stamp, the name of the experiment script file, and optionally a number indicating a sequence of shots or a repeated measurement. Since it is likely that there will be people in the research group continuing this work on magnetic resonance control, Table 6.1 contains a record of the `.h5` files used to create the plots shown in this chapter. To measure the spin projection of condensate A and condensate B, and the un-split condensate C, the regions of interest passed to fitting routines are recorded in Table 6.2.

Table 6.1: Record of the experiment `.h5` files which generated the results presented in this chapter. An asterisk in the filename indicates an entire sequence of shots.

Figure	File name
6.4	20150106T141775_mrc_*
	20150106T164514_mrc_*
6.5	20150106T170832_mrc_0
	20150106T141755_mrc_50
	20150106T170712_mrc_0
	20150106T170724_mrc_0
6.6	20150110T155231_mrc_*
	20150110T163324_mrc_*
	20150110T170254_mrc_*
	20150110T175321_mrc_*
6.7	20150110T155231_mrc_*
	20150110T163324_mrc_*
	20150110T175321_mrc_*
6.8	20150106T170832_mrc_0
	20150110T170254_mrc_03
	20150110T155231_mrc_19
	20150110T155231_mrc_46
6.9	20150110T145817_mrc_*
6.10	20150109T170249_mrc_*
	20150110T132244_mrc_*
	20150110T135926_mrc_*
6.11	20150109T170249_mrc_*
	20150110T132244_mrc_*
	20150110T135926_mrc_*
	20150109T143623_mrc_*
	20150109T150949_mrc_*
	20150109T153711_mrc_*
6.12	20150106T170832_mrc_0
	20150110T135926_mrc_01
	20150110T132244_mrc_18
	20150110T135926_mrc_20

Table 6.2: The regions of interest (subsets of the entire camera image) used when measuring the condensate spin projection, in units of pixels, defined by the coordinates $[(x_0, y_0), (x_1, y_1)]$ relative to the camera chip.

Condensate	Spin state	ROI
A	$m_F = -1$	$[(1074, 477), (1230, 613)]$
	$m_F = 0$	$[(1236, 477), (1380, 613)]$
	$m_F = +1$	$[(1386, 477), (1528, 613)]$
B	$m_F = -1$	$[(1074, 630), (1230, 759)]$
	$m_F = 0$	$[(1236, 630), (1380, 759)]$
	$m_F = +1$	$[(1386, 630), (1528, 759)]$
C	$m_F = -1$	$[(1082, 556), (1227, 688)]$
	$m_F = 0$	$[(1232, 555), (1365, 691)]$
	$m_F = +1$	$[(1366, 549), (1527, 694)]$

Chapter 7

Conclusion

For decades, the magnetic resonance imaging (MRI) community have combined magnetic field gradients with radiofrequency pulses to measure nuclear spin. I have shown in this thesis that we can apply similar techniques to operate on the atomic spin of Bose–Einstein condensates; not only to image condensates (as my colleague has shown [62]) but also to engineer the condensate density and phase. Our technique, which we have called magnetic resonance control, involves exposing the condensate to a magnetic field gradient, which causes the resonance frequency between internal states to have spatial dependence, so that we can use a time-dependent coupling between internal states to address selected points within the condensate, giving us the ability to engineer the density, or phase, or both, of the condensate.

Unlike other wavefunction engineering techniques which rely on an optical field to address different spatial regions of the condensate, the spatial resolution of magnetic resonance control is not limited to the wavelength of the light used; indeed, radiofrequency pulses have wavelengths many orders of magnitudes greater than the size of a condensate. Instead, the resolution of magnetic resonance control is limited by the strength of the magnetic field gradient you can generate, the speed at which this gradient can be switched, and the strength of the coupling; unlike the optical diffraction limit, these are limits which can be circumvented through successful grant applications.

I have shown through simulations that we should be able to use magnetic resonance control to engineer a single black soliton in a Bose–Einstein condensate (a feat requiring healing-length resolution) using simple hyperbolic secant radiofrequency pulses and magnetic gradients which, while strong, are not unusual for those who form Bose–Einstein condensates using chip traps. Achieving this in the laboratory would be remarkable because no one has engineered a single *black* soliton in a condensate before.

Developing this technique a step further, I have simulated using magnetic resonance control to engineer multiple solitons in rapid succession with control over their initial locations and trajectories. I achieved this through repeated application of a magnetic resonance control pulse sequence, engineering two solitons at a time,

but we believe magnetic resonance control could engineer more than two solitons simultaneously if we used multi-tonal coupling pulses; a waveform made of the superposition of hyperbolic secant pulses with different detuning offset, or different sweep range, or both, to address multiple regions of the condensate during the single pulse. This could allow the study of densely packed soliton lattices [120], dispersive shock waves in Bose–Einstein condensates [121] and soliton gases [122].

Another area of interest is the formation of spin domain walls [123], and the interactions between spin domain walls and spin waves [124,125]. We observe that during the magnetic resonance control sequence which forms a black soliton, after the first pulse we have formed a spin domain wall; each half of the condensate is in a different spin state, with the boundary extending over roughly one healing length. Magnetic resonance control is a technique which could be applied to create multiple spin domain walls, or even a mixture of spin domain walls, spin waves, and solitons, to observe how they interact.

In this thesis we have considered hyperbolic secant pulses, which are one of the simplest forms of pulses. Looking to the MRI field for inspiration, there are a wide array of intricate pulse shapes which could be investigated for use in the context of magnetic resonance control (e.g. [126–128]). It may be possible to find a pulse which could engineer the desired state in a much shorter time than a hyperbolic secant pulse, or, equivalently, a pulse which could engineer the desired state in the same length of time but requiring a much weaker coupling strength. There are likely to be different forms of pulses which could engineer target states difficult to produce using hyperbolic secant pulses alone. To find such a pulse, optimal control theory might be used to create an automated optimisation of pulse parameters.

The target states discussed so far have all been one-dimensional features, but we aim to extend magnetic resonance control to higher dimensions. For the study of topological defects, higher-dimensional magnetic resonance control would allow us to engineer vortices, skyrmions, and spin knots with precision. There is recent interest in the behaviour of caustics and diffraction catastrophes in Bose–Einstein condensates [129], where an exotic two-dimensional potential is applied to the condensate, causing it to contract (or focus) into intricate aberration patterns as though the condensate were an optical beam which had passed through a lens. In this paper, there are no suggestions for how these lensing potentials could be made; I suggest that magnetic resonance control could be used to modulate the phase across the whole condensate (rather than in only one sharp region, as we have considered in this thesis), in the same way as a lens modulates the phase of a light beam.

To achieve anything more complicated than a 'cylindrical lens' effect, we would need to extend magnetic resonance control to at least two dimensions. We might do this by using a rotating magnetic field gradient, or by using the vector light shift from an optical field. This optical field could be shaped using a spatial light modulator to synthesise a non-monotonic magnetic field gradient which varies in two dimensions (limited by the resolution of the spatial light modulator). A time-dependent coupling pulse could then be used to selectively address regions within

this gradient across scales smaller than the resolution of the spatial light modulator.

Before being able to perform magnetic resonance control experiments in the laboratory, we needed to build our spinor Bose–Einstein condensate apparatus. In this thesis I summarised my main contributions to this endeavour, including optical and electronic design and construction, assembling the ultra-high vacuum system, coordinating the vacuum bakeout, designing imaging systems, and creating the magneto-optical trap.

One of my tasks particularly pertinent to magnetic resonance control was designing and testing a high resolution diffraction-limited objective lens. Using this lens we should be able to resolve solitons in condensates *in situ*, a necessary ability in order to to confirm the success of engineering a black soliton in a condensate using magnetic resonance control. This objective lens consists of entirely standard lens elements, available from commercial catalogues, making this lens simpler and cheaper than using custom-ground objective lenses. Another feature of my design is that by changing the spacings between lens elements, the diffraction-limited performance is maintained across the imaging wavelengths for all the alkali gases. With these features, this objective lens design could become the 'work-horse' lens in the field of quantum gas experiments.

After constructing our Bose–Einstein condensate apparatus, and developing the technique of magnetic resonance control in simulations, I was able to perform the first proof-of-principle magnetic resonance control experiments on Bose–Einstein condensates. Limited by the weak magnetic field gradients we could produce with our bias coils I selectively transferred the spin state of only one of two adjacent condensates using a hyperbolic secant pulse, in a regime where we were unable to achieve this with a square-profile pulse.

The next step in these investigations will be to increase the magnetic field gradient, so that we can address the condensate at the healing length scale. We could increase the magnetic field gradient by either installing additional coils close to the science cell, or by using the vector light shift from a laser beam to synthesise the gradient. It may even be possible to use our bias coils to offset the strong magnetic trap produced by the quadrupole coils, to move a region of strong magnetic field gradient into the condensate. Of course, instead of pursuing these options ourselves, we might decide to form a collaboration with another research group who make Bose–Einstein condensates using chip traps easily capable of producing the requisite gradients for healing-length resolution.

Once we can use magnetic resonance control in the laboratory to engineer the condensate wavefunction with healing-length precision, we could investigate combining this technique with Faraday magnetic resonance imaging [62] which applies the techniques of MRI to construct images of condensates non-destructively. Using our laboratory control system software [87] it is possible to set up autonomous feedback control, where parameters for the next experimental procedure are automatically selected based on measurements of the system. With the combined powers of magnetic resonance control, Faraday MRI, and autonomous feedback control at our disposal, it could be possible to achieve real-time adaptive control

over the condensate wavefunction. Having the ability to modify the wavefunction immediately after non-destructively imaging some desired event would be ideal for using Bose–Einstein condensates as quantum simulators.

At the moment magnetic resonance control is still in its infancy, but with all these possible directions we can predict that it has a bright future ahead of it. My hope is that further development of magnetic resonance control will lead to it becoming a common technique in Bose–Einstein condensate experiments, whenever a fine degree of control over the full condensate wavefunction is desired.

References

[1] R. P. Feynman. *Simulating physics with computers*. International Journal of Theoretical Physics **21**, 467 (1982). DOI: 10.1007/BF02650179. [p 1]

[2] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press (2000). [p 1]

[3] J. Stolze and D. Suter. *Quantum computing: A short course from theory to experiment*. John Wiley & Sons (2004). [p 1]

[4] W. P. Schleich and H. Walther. *Elements of quantum information*. John Wiley & Sons (2007). [p 1]

[5] I. M. Georgescu, S. Ashhab, and F. Nori. *Quantum simulation*. Reviews of Modern Physics **86**, 153 (2014). DOI: 10.1103/RevModPhys.86.153. [p 2]

[6] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch. *Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms*. Nature **415**, 39 (2002). DOI: 10.1038/415039a. [p 2]

[7] D. Leibfried, B. DeMarco, V. Meyer, M. Rowe, A. Ben-Kish, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. *Trapped-ion quantum simulator: Experimental application to nonlinear interferometers*. Physical Review Letters **89**, 247901 (2002). DOI: 10.1103/PhysRevLett.89.247901. [p 2]

[8] A. Friedenauer, H. Schmitz, J. T. Glueckert, D. Porras, and T. Schaetz. *Simulating a quantum magnet with trapped ions*. Nature Physics **4**, 757 (2008). DOI: 10.1038/nphys1032. [p 2]

[9] M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, E. Lucero, A. D. O'Connell, D. Sank, H. Wang, J. Wenner, A. N. Cleland, M. R. Geller, and J. M. Martinis. *Emulation of a quantum spin with a superconducting phase qudit*. Science **325**, 722 (2009). DOI: 10.1126/science.1173440. [p 2]

[10] R. Gerritsma, G. Kirchmair, F. Zähringer, E. Solano, R. Blatt, and C. F. Roos. *Quantum simulation of the Dirac equation*. Nature **463**, 68 (2010). DOI: 10.1038/nature08688. [p 2]

[11] K. Kim, M.-S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G.-D. Lin, L.-M. Duan, and C. Monroe. *Quantum simulation of frustrated Ising spins with trapped ions*. Nature **465**, 590 (2010). DOI: 10.1038/nature09071. [p 2]

[12] I. Bloch, J. Dalibard, and S. Nascimbène. *Quantum simulations with ultracold quantum gases*. Nature Physics **8**, 267 (2012). DOI: 10.1038/nphys2259. [p 2]

[13] X. Yang, A. M. Wang, F. Xu, and J. Du. *Experimental simulation of a pairing Hamiltonian on an NMR quantum computer*. Chemical Physics Letters **422**, 20 (2006). DOI: 10.1016/j.cplett.2006.02.023. [p 2]

[14] C. A. Regal, M. Greiner, and D. S. Jin. *Observation of resonance condensation of fermionic atom pairs*. Physical Review Letters **92**, 040403 (2004). DOI: 10.1103/PhysRevLett.92.040403. [p 2]

[15] M. W. Zwierlein, J. R. Abo-Shaeer, A. Schirotzek, C. H. Schunck, and W. Ketterle. *Vortices and superfluidity in a strongly interacting Fermi gas*. Nature **435**, 1047 (2005). DOI: 10.1038/nature03858. [p 2]

[16] U. R. Fischer and R. Schützhold. *Quantum simulation of cosmic inflation in two-component Bose–Einstein condensates*. Physical Review A **70**, 063615 (2004). DOI: 10.1103/PhysRevA.70.063615. [p 2]

[17] C. Neuenhahn, A. Polkovnikov, and F. Marquardt. *Localized phase structures growing out of quantum fluctuations in a quench of tunnel-coupled atomic condensates*. Physical Review Letters **109**, 085304 (2012). DOI: 10.1103/PhysRevLett.109.085304. [p 2]

[18] B. Opanchuk, R. Polkinghorne, O. Fialko, J. Brand, and P. D. Drummond. *Quantum simulations of the early universe*. Annalen der Physik **525**, 866 (2013). DOI: 10.1002/andp.201300113. [p 2]

[19] J. Steinhauer. *Observation of self-amplifying Hawking radiation in an analogue black-hole laser*. Nature Physics **10**, 864 (2014). DOI: 10.1038/nphys3104. [p 2]

[20] Y. Kawaguchi and M. Ueda. *Spinor Bose–Einstein condensates*. Physics Reports **520**, 253 (2012). DOI: 10.1016/j.physrep.2012.07.005. [pp 2 and 63]

[21] M. Lewenstein, A. Sanpera, V. Ahufinger, B. Damski, A. Sen(De), and U. Sen. *Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond*. Advances in Physics **56**, 243 (2007). DOI: 10.1080/00018730701223200. [p 2]

[22] Arko Roy and D. Angom. *Fluctuation- and interaction-induced instability of dark solitons in single and binary condensates*. Physical Review A **90**, 023612 (2014). DOI: 10.1103/PhysRevA.90.023612. [p 3]

[23] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hänsch, and I. Bloch. *Coherent transport of neutral atoms in spin-dependent optical lattice potentials*. Physical Review Letters **91**, 010407 (2003). DOI: 10.1103/PhysRevLett.91.010407. [pp 3 and 10]

[24] M. O. Borgh and J. Ruostekoski. *Topological interface engineering and defect crossing in ultracold atomic gases*. Physical Review Letters **109**, 015302 (2012). DOI: 10.1103/PhysRevLett.109.015302. [p 3]

[25] J. S. Russell. *Report on waves*. Fourteenth meeting of the British Association for the Advancement of Science pages 311–390 (1844). [p 3]

[26] Y. S. Kivshar and G. Agrawal. *Optical solitons: From fibers to photonic crystals*. Academic Press (2003). [p 3]

[27] S. Burger, K. Bongs, S. Dettmer, W. Ertmer, K. Sengstock, A. Sanpera, G. V. Shlyapnikov, and M. Lewenstein. *Dark solitons in Bose–Einstein condensates*. Physical Review Letters **83**, 5198 (1999). DOI: 10.1103/PhysRevLett.83.5198. [pp 3, 4, and 5]

[28] J. Denschlag, J. E. Simsarian, D. L. Feder, Charles W. Clark, L. A. Collins, J. Cubizolles, L. Deng, E. W. Hagley, K. Helmerson, W. P. Reinhardt, S. L. Rolston, B. I. Schneider, and W. D. Phillips. *Generating solitons by phase engineering of a Bose–Einstein condensate*. Science **287**, 97 (2000). DOI: 10.1126/science.287.5450.97. [pp 3, 4, 5, and 7]

[29] K. E. Strecker, G. B. Partridge, A. G. Truscott, and R. G. Hulet. *Formation and propagation of matter-wave soliton trains*. Nature **417**, 150 (2002). DOI: 10.1038/nature747. [p 3]

[30] B. P. Anderson. *Dark solitons in BECs: the first experiments*. In Professor Panayotis G. Kevrekidis, Professor Dimitri J. Frantzeskakis, and Professor Ricardo Carretero-González (editors), *Emergent Nonlinear Phenomena in Bose–Einstein Condensates*, number 45 in Atomic, Optical, and Plasma Physics, pages 85–96. Springer Berlin Heidelberg (2008). [p 4]

[31] Z. Dutton, M. Budde, C. Slowe, and L. V. Hau. *Observation of quantum shock waves created with ultra- compressed slow light pulses in a Bose–Einstein condensate*. Science **293**, 663 (2001). DOI: 10.1126/science.1062527. [pp 4 and 92]

[32] H. Flayac, D. D. Solnyshkov, and G. Malpuech. *Spin effects in polariton condensates: From half-solitons to analogues of wormholes*. In A. Bramati and M. Modugno (editors), *Physics of Quantum Fluids: New Trends and Hot Topics in Atomic and Polariton Condensates*. Springer Science & Business Media (2013). [p 4]

[33] C. J. Pethick and H. Smith. *Bose–Einstein condensation in dilute gases*. Cambridge University Press (2002). [pp 5, 56, and 67]

[34] Th. Busch and J. R. Anglin. *Motion of dark solitons in trapped Bose–Einstein condensates*. Physics Review Letters **84**, 2298 (2000). DOI: 10.1103/PhysRevLett.84.2298. [p 5]

[35] Ł. Dobrek, M. Gajda, M. Lewenstein, K. Sengstock, G. Birkl, and W. Ertmer. *Optical generation of vortices in trapped Bose–Einstein condensates*. Physical Review A **60**, R3381 (1999). DOI: 10.1103/PhysRevA.60.R3381. [p 7]

[36] T. F. Scott, R. J. Ballagh, and K. Burnett. *Formation of fundamental structures in Bose–Einstein condensates*. Journal of Physics B: Atomic, Molecular and Optical Physics **31**, L329 (1998). DOI: 10.1088/0953-4075/31/8/001. [pp 7 and 8]

[37] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle. *Observation of interference between two Bose condensates*. Science **275**, 637 (1997). DOI: 10.1126/science.275.5300.637. [p 8]

[38] I. Shomroni, E. Lahoud, S. Levy, and J. Steinhauer. *Evidence for an oscillating soliton/vortex ring by density engineering of a Bose–Einstein condensate*. Nature Physics **5**, 193 (2009). DOI: 10.1038/nphys1177. [p 8]

[39] A. Weller, J. P. Ronzheimer, C. Gross, J. Esteve, M. K. Oberthaler, D. J. Frantzeskakis, G. Theocharis, and P. G. Kevrekidis. *Experimental observation of oscillating and interacting matter wave dark solitons*. Physical Review Letters **101**, 130401 (2008). DOI: 10.1103/PhysRevLett.101.130401. [p 8]

[40] L. D. Carr, J. Brand, S. Burger, and A. Sanpera. *Dark-soliton creation in Bose–Einstein condensates*. Physical Review A **63**, 051601 (2001). DOI: 10.1103/PhysRevA.63.051601. [p 9]

[41] G. Theocharis, A. Weller, J. P. Ronzheimer, C. Gross, M. K. Oberthaler, P. G. Kevrekidis, and D. J. Frantzeskakis. *Multiple atomic dark solitons in cigar-shaped Bose–Einstein condensates*. Physical Review A **81**, 063604 (2010). DOI: 10.1103/PhysRevA.81.063604. [p 9]

[42] M. F. Andersen, C. Ryu, P. Cladé, V. Natarajan, A. Vaziri, K. Helmerson, and W. D. Phillips. *Quantized rotation of atoms from photons with orbital angular momentum*. Physical Review Letters **97**, 170406 (2006). DOI: 10.1103/PhysRevLett.97.170406. [pp 9 and 62]

[43] K. C. Wright, L. S. Leslie, and N. P. Bigelow. *Optical control of the internal and external angular momentum of a Bose–Einstein condensate*. Physical Review A **77**, 041601 (2008). DOI: 10.1103/PhysRevA.77.041601. [p 9]

[44] K. C. Wright, L. S. Leslie, A. Hansen, and N. P. Bigelow. *Sculpting the vortex state of a spinor BEC*. Physical Review Letters **102**, 030405 (2009). DOI: 10.1103/PhysRevLett.102.030405. [pp 9 and 62]

[45] S. Franke-Arnold, J. Leach, M. J. Padgett, V. E. Lembessis, D. Ellinas, A. J. Wright, J. M. Girkin, P. Öhberg, and A. S. Arnold. *Optical ferris wheel for ultracold atoms*. Optics Express **15**, 8619 (2007). DOI: 10.1364/OE.15.008619. [p 9]

[46] T. Isoshima, M. Nakahara, T. Ohmi, and Kazushige Machida. *Creation of a persistent current and vortex in a Bose–Einstein condensate of alkali-metal atoms*. Physical Review A **61**, 063610 (2000). DOI: 10.1103/PhysRevA.61.063610. [p 9]

[47] A. E. Leanhardt, A. Görlitz, A. P. Chikkatur, D. Kielpinski, Y. Shin, D. E. Pritchard, and W. Ketterle. *Imprinting vortices in a Bose–Einstein condensate using topological phases*. Physical Review Letters **89**, 190403 (2002). DOI: 10.1103/PhysRevLett.89.190403. [p 9]

[48] H. Shibayama, Y. Yasaku, and T. Kuwamoto. *Vortex nucleation in Bose–Einstein condensates confined in a QUIC trap by topological phase imprinting*. Journal of Physics B: Atomic, Molecular and Optical Physics **44**, 075302 (2011). DOI: 10.1088/0953-4075/44/7/075302. [p 9]

[49] A. E. Leanhardt, Y. Shin, D. Kielpinski, D. E. Pritchard, and W. Ketterle. *Coreless vortex formation in a spinor Bose–Einstein condensate*. Physical Review Letters **90**, 140403 (2003). DOI: 10.1103/PhysRevLett.90.140403. [p 9]

[50] J. Choi, W. Kwon, and Y. Shin. *Observation of topologically stable 2D skyrmions in an antiferromagnetic spinor Bose–Einstein condensate*. Physical Review Letters **108**, 035301 (2012). DOI: 10.1103/PhysRevLett.108.035301. [p 9]

[51] Y. Kawaguchi, M. Nitta, and M. Ueda. *Knots in a spinor Bose–Einstein condensate*. Physical Review Letters **100**, 180403 (2008). DOI: 10.1103/PhysRevLett.100.180403. [p 9]

[52] M. Möttönen, V. Pietilä, and S. M. M. Virtanen. *Vortex pump for dilute Bose–Einstein condensates*. Physical Review Letters **99**, 250406 (2007). DOI: 10.1103/PhysRevLett.99.250406. [p 9]

[53] Z. F. Xu, P. Zhang, C. Raman, and L. You. *Continuous vortex pumping into a spinor condensate with magnetic fields*. Physical Review A **78**, 043606 (2008). DOI: 10.1103/PhysRevA.78.043606. [p 9]

[54] Z. F. Xu, P. Zhang, R. Lü, and L. You. *Quantum pump for countercirculation in a spinor Bose–Einstein condensate*. Physical Review A **81**, 053619 (2010). DOI: 10.1103/PhysRevA.81.053619. [p 9]

[55] S. Giorgini, L. Pitaevskii, and S. Stringari. *Theory of ultracold atomic Fermi gases*. Reviews of Modern Physics **80**, 1215 (2008). DOI: 10.1103/RevModPhys.80.1215. [p 10]

[56] D. Schrader, I. Dotsenko, M. Khudaverdyan, Y. Miroshnychenko, A. Rauschenbeutel, and D. Meschede. *Neutral atom quantum register*. Physical Review Letters **93**, 150501 (2004). DOI: 10.1103/PhysRevLett.93.150501. [p 10]

[57] C. Zhang, S. L. Rolston, and S. Das Sarma. *Manipulation of single neutral atoms in optical lattices*. Physical Review A **74**, 042316 (2006). DOI: 10.1103/PhysRevA.74.042316. [p 10]

[58] H. Godfrin and R. E. Rapp. *Two-dimensional nuclear magnets*. Advances in Physics **44**, 113 (1995). DOI: 10.1080/00018739500101516. [p 10]

[59] L. M. Bennie, P. T. Starkey, M. Jasperse, C. J. Billington, R. P. Anderson, and L. D. Turner. *A versatile high resolution objective for imaging quantum gases*. Optics Express **21**, 9011 (2013). DOI: 10.1364/OE.21.009011. [pp 11 and 39]

[60] L. M. Starkey, P. B. Wigley, S. S. Szigeti, M. Jasperse, J. J. Hope, L. D. Turner, and R. P. Anderson. *Precise wavefunction engineering with magnetic resonance*. arXiv:1412.6854 [quant-ph] (2016). [pp 11 and 51]

[61] A. A. Wood. *Spinor Bose–Einstein condensates in magnetic field gradients*. PhD thesis, Monash University (2015). [pp 13, 15, 16, 19, 25, 26, 36, 38, 73, 110, and 112]

[62] M. Jasperse. *Faraday magnetic resonance imaging of Bose–Einstein condensates*. PhD thesis, Monash University (2015). [pp 13, 15, 36, 38, 125, and 127]

[63] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto. *Rapid production of ^{87}Rb Bose–Einstein condensates in a combined magnetic and optical potential*. Physical Review A **79**, 063631 (2009). DOI: 10.1103/PhysRevA.79.063631. [p 13]

[64] S. C. Bell, M. Junker, M. Jasperse, L. D. Turner, Y.-J. Lin, I. B. Spielman, and R. E. Scholten. *A slow atom source using a collimated effusive oven and a single-layer variable pitch coil Zeeman slower*. Review of Scientific Instruments **81**, 013105 (2010). DOI: 10.1063/1.3276712. [p 13]

[65] L. M. Bennie. *Design and construction of a Zeeman slower for ^{85}Rb and ^{87}Rb*. 3rd year research project, Monash University (2010). [p 15]

[66] A. A. Wood. *A rubidium cold atom beam source*. Honours thesis, Monash University (2010). [p 15]

[67] L. M. Bennie. *A large atom number magneto-optical trap for BEC production*. Honours thesis, Monash University (2010). [pp 15 and 27]

[68] K. M. Birnbaum. *Ultra-high vacuum chambers*. Unpublished, California Institute of Technology, California Institute of Technology (2005). [p 18]

[69] V. Negnevitsky and L. D. Turner. *Wideband laser locking to an atomic reference with modulation transfer spectroscopy*. Optics Express **21**, 3103 (2013). DOI: 10.1364/OE.21.003103. [p 24]

[70] J. Appel, A. MacRae, and A. I. Lvovsky. *A versatile digital GHz phase lock for external cavity diode lasers*. Measurement Science and Technology **20**, 055302 (2009). DOI: 10.1088/0957-0233/20/5/055302. [p 24]

[71] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts. *Double-pass acousto-optic modulator system*. Review of Scientific Instruments **76**, 063112 (2005). DOI: 10.1063/1.1930095. [p 25]

[72] P. T. Starkey. *An optical, experiment control and data acquisition system for a BEC lab*. Honours thesis, Monash University (2010). [p 27]

[73] L. P. Maguire, S. Szilagyi, and R. E. Scholten. *High performance laser shutter using a hard disk drive voice-coil actuator*. Review of Scientific Instruments **75**, 3077 (2004). DOI: 10.1063/1.1786331. [p 30]

[74] P. C. D. Hobbs. *Building electro-optical systems: Making it all work*. John Wiley & Sons (2011). [p 32]

[75] P. Horowitz and W. Hill. *The art of electronics*. Cambridge University Press (2006). [p 34]

[76] W. Ketterle, D. S. Durfee, and D. M Stamper-Kurn. *Making, probing and understanding Bose–Einstein condensates*. In M. Inguscio, S. Stringari, and C. E. Wieman (editors), *Bose–Einstein condensation in atomic gases*, volume 140 of *Proceedings of the International School of Physics*. IOS Press (1999). [p 36]

[77] D. Malacara-Hernández and Z. Malacara-Hernández. *Handbook of optical design, second edition*. CRC Press (2014). [p 39]

[78] Y. R. P. Sortais, H. Marion, C. Tuchendler, A. M. Lance, M. Lamare, P. Fournet, C. Armellin, R. Mercier, G. Messin, A. Browaeys, and P. Grangier. *Diffraction-limited optics for single-atom manipulation*. Physical Review A **75**, 013406 (2007). DOI: 10.1103/PhysRevA.75.013406. [p 40]

[79] Waseem S. Bakr, Jonathon 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**, 74 (2009). DOI: 10.1038/nature08482. [p 40]

[80] J. W. Goodman. *Introduction to Fourier optics*. Roberts and Company Publishers (2005). [p 40]

[81] W. Alt. *An objective lens for efficient fluorescence detection of single atoms*. Optik - International Journal for Light and Electron Optics **113**, 142 (2002). DOI: 10.1078/0030-4026-00133. [p 40]

[82] T. B. Ottenstein. *A new objective for high-resolution imaging of Bose–Einstein condensates*. Diploma thesis, University of Heidelberg (2006). [p 40]

[83] K. D. Nelson, X. Li, and D. S. Weiss. *Imaging single atoms in a three-dimensional array*. Nature Physics **3**, 556 (2007). DOI: 10.1038/nphys645. [p 40]

[84] R. Bücker, A. Perrin, S. Manz, T. Betz, C. Koller, T. Plisson, J. Rottmann, T. Schumm, and J. Schmiedmayer. *Single-particle-sensitive imaging of freely propagating ultracold atoms*. New Journal of Physics **11**, 103039 (2009). DOI: 10.1088/1367-2630/11/10/103039. [p 40]

[85] 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**, 68 (2010). DOI: 10.1038/nature09378. [p 40]

[86] B Zimmermann, T Müller, J Meineke, T Esslinger, and H Moritz. *High-resolution imaging of ultracold fermions in microscopically tailored optical potentials*. New Journal of Physics **13**, 043007 (2011). DOI: 10.1088/1367-2630/13/4/043007. [p 40]

[87] P. T. Starkey, C. J. Billington, S. P. Johnstone, M. Jasperse, K. Helmerson, L. D. Turner, and R. P. Anderson. *A scripted control system for autonomous hardware-timed experiments*. Review of Scientific Instruments **84**, 085111 (2013). DOI: 10.1063/1.4817213. [pp 40 and 127]

[88] M. S. Silver, R. I. Joseph, and D. I. Hoult. *Selective spin inversion in nuclear magnetic resonance and coherent optics through an exact solution of the Bloch–Riccati equation*. Physical Review A **31**, 2753 (1985). DOI: 10.1103/PhysRevA.31.2753. [pp 52, 57, and 59]

[89] R. P. Anderson, C. Ticknor, A. I. Sidorov, and B. V. Hall. *Spatially inhomogeneous phase evolution of a two-component Bose–Einstein condensate*. Physical Review A **80**, 023603 (2009). DOI: 10.1103/PhysRevA.80.023603. [p 52]

[90] B. M. Dale, M. A. Brown, and R. C. Semelka. *MRI: Basic principles and applications*. John Wiley & Sons (2015). [pp 54 and 55]

[91] P. Deuar and P. D. Drummond. *First-principles quantum dynamics in interacting Bose gases: I. the positive P representation*. Journal of Physics A: Mathematical and General **39**, 1163 (2006). DOI: 10.1088/0305-4470/39/5/010. [p 54]

[92] M. A. Bernstein, K. F. King, and X. J. Zhou. *Handbook of MRI pulse sequences*. Elsevier (2004). [p 57]

[93] F. D. Giacomo and E. E. Nikitin. *The Majorana formula and the Landau–Zener–Stückelberg treatment of the avoided crossing problem*. Physics-Uspekhi **48**, 515 (2005). DOI: 10.1070/PU2005v048n05ABEH002804. [p 59]

[94] R. P. Anderson. *Non-equilibrium dynamics and relative phase evolution of two-component Bose–Einstein condensates*. PhD thesis, Swinburne University of Technology (2010). [pp 59, 63, and 72]

[95] J. E. Thomas and L. J. Wang. *Precision position measurement of moving atoms*. Physics Reports **262**, 311 (1995). DOI: 10.1016/0370-1573(95)00050-Q. [p 61]

[96] J. E. Williams and M. J. Holland. *Preparing topological states of a Bose–Einstein condensate*. Nature **401**, 568 (1999). DOI: 10.1038/44095. [p 62]

[97] 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*. Physical Review Letters **83**, 2498 (1999). DOI: 10.1103/PhysRevLett.83.2498. [p 62]

[98] B. P. Anderson, P. C. Haljan, C. A. Regal, D. L. Feder, L. A. Collins, C. W. Clark, and E. A. Cornell. *Watching dark solitons decay into vortex rings in a Bose–Einstein condensate*. Physical Review Letters **86**, 2926 (2001). DOI: 10.1103/PhysRevLett.86.2926. [pp 62 and 92]

[99] K. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. *Vortex formation in a stirred Bose–Einstein condensate*. Physical Review Letters **84**, 806 (2000). DOI: 10.1103/PhysRevLett.84.806. [p 62]

[100] E. P. Gross. *Structure of a quantized vortex in boson systems*. Il Nuovo Cimento Series 10 **20**, 454 (1961). DOI: 10.1007/BF02731494. [p 63]

[101] Lev P. Pitaevskii. *Vortex lines in an imperfect Bose gas*. Journal of Experimental and Theoretical Physics **13** (1961). [p 63]

[102] H. Zeng, W. Zhang, and F. Lin. *Nonclassical Bose–Einstein condensate*. Physical Review A **52**, 2155 (1995). DOI: 10.1103/PhysRevA.52.2155. [p 63]

[103] T.-L. Ho and V. B. Shenoy. *Binary mixtures of Bose condensates of alkali atoms*. Physical Review Letters **77**, 3276 (1996). DOI: 10.1103/PhysRevLett.77.3276. [p 63]

[104] L. P. Pitaevskii and S. Stringari. *Bose–Einstein Condensation*. Oxford University Press (2003). [p 63]

[105] W. Zhang and L. You. *An effective quasi-one-dimensional description of a spin-1 atomic condensate*. Physical Review A **71**, 025603 (2005). DOI: 10.1103/PhysRevA.71.025603. [pp 63, 64, and 65]

[106] M. Jasperse. *High resolution magnetic resonance imaging of Bose–Einstein condensates*. PhD confirmation report, Monash University (2011). [p 65]

[107] L. Lehtovaara, J. Toivanen, and J. Eloranta. *Solution of time-independent Schrödinger equation by the imaginary time propagation method*. Journal of Computational Physics **221**, 148 (2007). DOI: 10.1016/j.jcp.2006.06.006. [p 66]

[108] G. R. Dennis, J. J. Hope, and M. T. Johnsson. *XMDS2: Fast, scalable simulation of coupled stochastic partial differential equations*. Computer Physics Communications **184**, 201 (2013). DOI: 10.1016/j.cpc.2012.08.016. [p 66]

[109] J. C. Butcher. *An introduction to "Almost Runge–Kutta" methods*. Applied Numerical Mathematics **24**, 331 (1997). DOI: 10.1016/S0168-9274(97)00030-5. [p 66]

[110] W. Hänsel, P. Hommelhoff, T. W. Hänsch, and J. Reichel. *Bose–Einstein condensation on a microelectronic chip*. Nature **413**, 498 (2001). DOI: 10.1038/35097032. [p 72]

[111] C. Y. Park, H. Noh, C. M. Lee, and D. Cho. *Measurement of the Zeeman-like ac Stark shift*. Physical Review A **63**, 032512 (2001). DOI: 10.1103/PhysRevA.63.032512. [p 72]

[112] C. J. Whittle, R. P. Anderson, and L. D. Turner. *Limits to optically induced Zeeman shifts in the absence of scalar light shifts*. In preparation. [p 72]

[113] D. A. Steck. *Rubidium 87 D line data*. http://steck.us/alkalidata (2010). [p 72]

[114] A. Ramanathan, S. R. Muniz, K. C. Wright, R. P. Anderson, W. D. Phillips, K. Helmerson, and G. K. Campbell. *Partial-transfer absorption imaging: A versatile technique for optimal imaging of ultracold gases*. Review of Scientific Instruments **83**, 083119 (2012). DOI: 10.1063/1.4747163. [p 73]

[115] J. T. Mendonça and H. Terças. *Physics of ultra-cold matter*, volume 70 of *Springer Series on Atomic, Optical, and Plasma Physics*. Springer New York, New York, NY (2013). [p 87]

[116] A. E. Muryshev, H. B. van Linden van den Heuvell, and G. V. Shlyapnikov. *Stability of standing matter waves in a trap*. Physical Review A **60**, R2665 (1999). DOI: 10.1103/PhysRevA.60.R2665. [p 92]

[117] R. Grimm, M. Weidemuller, and Y. B. Ovchinnikov. *Optical dipole traps for neutral atoms*. Advances in Atomic Molecular and Optical Physics page 95 (1999). [pp 92 and 94]

[118] M. Sargent and P. Horwitz. *Three-level Rabi flopping*. Physical Review A **13**, 1962 (1976). DOI: 10.1103/PhysRevA.13.1962. [p 108]

[119] A. A. Wood, L. M. Bennie, A. Duong, M. Jasperse, L. D. Turner, and R. P. Anderson. *Magnetic tensor gradiometry using Ramsey interferometry of spinor condensates*. Physical Review A **92**, 053604 (2015). DOI: 10.1103/PhysRevA.92.053604. [pp 111 and 112]

[120] G. A. El, A. Gammal, E. G. Khamis, R. A. Kraenkel, and A. M. Kamchatnov. *Theory of optical dispersive shock waves in photorefractive media*. Physical Review A **76**, 053813 (2007). DOI: 10.1103/PhysRevA.76.053813. [p 126]

[121] M. A. Hoefer, M. J. Ablowitz, I. Coddington, E. A. Cornell, P. Engels, and V. Schweikhard. *Dispersive and classical shock waves in Bose–Einstein condensates and gas dynamics*. Physical Review A **74**, 023623 (2006). DOI: 10.1103/PhysRevA.74.023623. [p 126]

[122] G. A. El and A. M. Kamchatnov. *Kinetic equation for a dense soliton gas*. Physical Review Letters **95**, 204101 (2005). DOI: 10.1103/PhysRevLett.95.204101. [p 126]

[123] T. Isoshima, K. Machida, and T. Ohmi. *Spin-domain formation in spinor Bose–Einstein condensation*. Physical Review A **60**, 4857 (1999). DOI: 10.1103/PhysRevA.60.4857. [p 126]

[124] D.-S. Han, S.-K. Kim, J.-Y. Lee, S. J. Hermsdoerfer, H. Schultheiss, B. Leven, and B. Hillebrands. *Magnetic domain-wall motion by propagating spin waves*. Applied Physics Letters **94**, 112502 (2009). DOI: 10.1063/1.3098409. [p 126]

[125] X. S. Wang, P. Yan, Y. H. Shen, G. E. W. Bauer, and X. R. Wang. *Domain wall propagation through spin wave emission*. Physical Review Letters **109**, 167209 (2012). DOI: 10.1103/PhysRevLett.109.167209. [p 126]

[126] K. E. Cano, M. A. Smith, and A. J. Shaka. *Adjustable, broadband, selective excitation with uniform phase*. Journal of Magnetic Resonance **155**, 131 (2002). DOI: 10.1006/jmre.2002.2506. [p 126]

[127] J. M. Warnking and G. B. Pike. *Bandwidth-modulated adiabatic RF pulses for uniform selective saturation and inversion*. Magnetic Resonance in Medicine **52**, 1190 (2004). DOI: 10.1002/mrm.20262. [p 126]

[128] J.-Y. Park, L. DelaBarre, and M. Garwood. *Improved gradient-echo 3D magnetic resonance imaging using pseudo-echoes created by frequency-swept pulses*. Magnetic Resonance in Medicine **55**, 848 (2006). DOI: 10.1002/mrm.20821. [p 126]

[129] T. P. Simula, T. C. Petersen, and D. M. Paganin. *Diffraction catastrophes threaded by quantized vortex skeletons caused by atom-optical aberrations induced in trapped Bose–Einstein condensates*. Physical Review A **88**, 043626 (2013). DOI: 10.1103/PhysRevA.88.043626. [p 126]