Durham E-Theses

Splitting and recombination of bright-solitary-matter waves

WALES, OLIVER,JOHN

Splitting and recombination of bright-solitary-matter waves

Oliver John Wales

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

Department of Physics

Durham University

August 2019

Splitting and recombination of bright-solitary-matter waves

Oliver John Wales

Abstract

This thesis presents the very first experimental realisation of splitting and recombination of robust bright-solitary-matter waves on a narrow repulsive potential barrier. This system has intrinsic interest for fundamental studies of soliton phase and for the realisation of a soliton interferometer.

An upgraded imaging system is presented, which is capable of imaging the bright-solitary-matter waves *in-situ* and across a wide range of magnetic fields. The system uses a combination of an offset-locked imaging laser, a high-intensity probe beam and a microwave-transfer adiabatic rapid passage of the atoms to an auxiliary imaging state which has favourable transition properties.

BECs of ^{85}Rb are created by direct evaporative cooling, using an upgraded crossed optical dipole trap. Condensates of up to 7000 atoms are created, with greater than $80\,\%$ purity. We develop a wavefunction engineering protocol, which allows us to transfer the condensate into a quasi-1D potential and systematically demonstrate regions of interatomic interactions where bright-solitary-matter waves can be formed, as well as regions where condensate collapse or breathing-mode phenomena are observed. The bright-solitary-matter waves are very robust, propagating without measurable dispersion for over $20\,\mathrm{s}$.

The splitting of a soliton into two daughter solitons on a narrow blue-detuned optical potential is presented. We demonstrate full control over the transmission coefficient by varying the barrier height. Velocity selection between the outgoing daughter solitons is observed and quantified, whereby the transmitted daughter soliton always has a higher centre of mass kinetic energy than the reflected daughter soliton.

Velocity-selection-mediated recombination on a wide barrier is demonstrated, as well as interference-mediated recombination on a narrow barrier. We observe large shot-to-shot fluctuations for the narrow barrier which are fully consistent with independently-determined uncertainties in the barrier position. For the first time we explore this experimentally and theoretically, both with Gross-Pitaevskii simulations and analytical approximations, putting new limits on the required parameters for future phase-sensitive interferometric measurements.

Contents

List of Figures

List of Equations

Declaration

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

The data presented in this thesis are available at doi:10.15128/r13t945q78r.

Acknowledgements

Throughout the course of my PhD, I have had the privilege of working alongside a tremendous number of extremely talented, enthusiastic and supportive individuals, without whom it would have been impossible to have realised many of the successes documented within this thesis.

First and foremost, I would like to thank my supervisor, Simon Cornish. Simon is a fantastic scientist and mentor: he has an unwavering attention to detail and commitment to scientific rigour, towards which all of his students aspire. On a personal level, I am particularly grateful for the understanding and patience he showed me throughout the latter half of my PhD, when difficulties outside of academic life were impacting my work. Simon has taught me a great deal throughout my time as his student: from the complexities of experimental ultracold atomic physics, to the intricacies of nonlinear soliton dynamics, to the fact that the dispenser needs to be switched on in order for it to release any atoms (I am particularly indebted to him for this one).

As a member of the soliton lab, I have had the pleasure of working alongside Anna Marchant and Ana Rakonjac. Despite only overlapping with Anna for the first six months of my PhD, she was instrumental in ensuring that I was up-to-speed with the experiment and the fundamentals of running a laboratory. Although her taste in music was highly questionable, with far too many 90's "classics" making the lab playlist for my taste, she made me feel immediately welcome in the group and gave me the confidence to take over the day-to-day running of the laboratory. I am extremely grateful to Ana for all she has done throughout my PhD, both as a colleague and

a friend. Even though she was so valuable to the group that she ended up working on countless experiments simultaneously, she always took the time to ensure that the soliton experiment was progressing whenever I faced problems. Her tireless enthusiasm, drive and dedication to her work was a continuous inspiration for me. I wish her every success in her new career outside academia.

Beyond the academic side of PhD life, the friendships and camaraderie I encountered within the Quantum Light and Matter group (formally Atomic and Molecular Physics group) were vital for my successful completion of my PhD. I feel highly privileged to have worked and spent time with a great number of incredible people, with whom I have many fond memories of tea breaks, cake clubs, beach trips, pub trips, parkruns and even nugget-eating contests! I would particularly like to thank Will Hamlyn and Ryan Hanley for their constant good humour and support in the office.

I must give thanks to the technical staff in the physics department. Although I required relatively infrequent use of their facilities, compared with the requirements of other students, I encountered many occasions where the assistance of the mechanical, electrical, teaching or research technicians was invaluable. They are each a credit to the physics department in Durham and are essential to Durham's experimental physics research output.

Away from Durham, I wish to express how lucky I am to have such fantastic parents and siblings, I would never have been able to achieve half of what I have achieved without their constant love and support.

Finally I thank my fienée Joanna for her unwavering love, support and friendship, even after I convinced her to relocate with me to the north east. Thank you for always being there for me.

Chapter 1

Introduction

Solitons are of fundamental interest in a broad range of settings due to their ubiquity in nonlinear systems, which occur prolifically throughout nature [1, 2]. They were first observed to occur in shallow water in 1834 by John Scott Russell, when he termed them a *wave of translation* [3]. Since then, the field of soliton physics has been the subject of extensive theoretical and experimental investigation, as it has far-reaching applicability across a range of scientific disciplines [4–11].

1.1 Background

There are multiple, comprehensive accounts of the background and theory of Bose-Einstein condensates (BECs) and solitons. Most relevant for wider reading on the subjects around those presented in this work are the four theses produced by previous PhD students on our experiment [12–15]. Therefore, in the interests of brevity and focus, we will only cover a preliminary background in soliton theory here, to give context for theory and results to come in subsequent chapters.

Solitons emerge as solutions to 1D nonlinear equations. The exact form of the solution to the nonlinear equation gives rise to a variety of solitons, namely:

bright, dark and gap solitons. Regardless of which type, there are a strict set of mathematical properties that should be fulfilled in order to call an object a soliton. It is common to ascribe a set of three basic requirements of a soliton that encompass these properties [16], which are:

1. They are of permanent form.

2. They are localized within a region.

3. They can interact with other solitons, emerging from the collision unchanged, except for a phase shift.

A true 1D geometry is impossible to achieve in a real-world experiment, as it is impossible to avoid radial extent with a physical system and there is usually some weak axial harmonic trapping present. This removes the integrability of the system and prevents true solitons from being created. Fortunately, the above properties can be generalised to a *quasi-1D* geometry to give the 3D soliton analogue: the solitary wave. However, the terms *soliton* and *solitary wave* are often used interchangeably within the literature, including throughout the remainder of this thesis.

1.1.1   Gross-Pitaevskii equation

In BECs, the required nonlinearity is provided by interatomic interactions governed by the *s*-wave scattering length [17], which can be tuned using a magnetic Feshbach resonance [18]. The effects of this nonlinearity can be explored by considering the Gross-Pitaevskii mean-field model, which is an established description for dilute, weakly-interacting Bose gas systems. In this model, the collective wavefunction of the BEC, $\Psi(\mathbf{r}, t)$, obeys the Gross-Pitaevskii equation (GPE):

$$i\hbar \frac{\partial \Psi(\mathbf{r}, t)}{\partial t} = \left[\frac{-\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g_{\mathrm{3D}}|\Psi(\mathbf{r}, t)|^2 \right] \Psi(\mathbf{r}, t), \qquad (1.1)$$

with effective interaction strength:

$$g_{3D} = 4\pi\hbar^2 N a_s/m, \tag{1.2}$$

where m is the atomic mass, a_s is the s-wave scattering length and N is the atom number.

Throughout this work we will study the interactions between a soliton in a quasi-1D potential and a repulsive Gaussian barrier. This arrangement suggests a more natural formulation of the potential in Eq. 1.1 to be:

$$V(\mathbf{r}) = \frac{m[\omega_z^2 z^2 + \omega_r^2(x^2 + y^2)]}{2} + \mathcal{V}(x, z), \tag{1.3}$$

where $\mathcal{V}(x, z)$ denotes the barrier potential and ω_z and ω_r are the axial and radial harmonic trapping frequencies respectively.

As $\omega_r \gg \omega_z$ for a true quasi-1D potential, Eq. 1.3 may be simplified by assuming the atoms to be frozen in the ground state of the radial oscillator potential. This returns an expression for the quasi-1D GPE:

$$i\hbar\frac{\partial\psi}{\partial t} = \left[\frac{-\hbar^2}{2m}\frac{\partial^2}{\partial z^2} + \frac{m\omega_z^2 z^2}{2} + V(z) + g_{1D}|\psi|^2\right]\psi, \tag{1.4}$$

with effective interaction strength:

$$g_{1D} = 2\hbar\omega_r a_s N. \tag{1.5}$$

The success of this model relies upon the fact that the de Broglie wavelength of the atoms is larger than the short-range molecular potential associated with the two-atom bound state. Therefore, the atoms do not *see* the potential, only a phase shift which relates to the scattering length:

$$a_s = \lim_{k\to 0}\frac{\tan\delta_0(k)}{k}, \tag{1.6}$$

where δ_0 is the phase shift of the $l = 0$ partial wave and k is the wavevector.

Close to a single Feshbach resonance, the magnetic field dependence of the scattering length is given by the well-known expression:

$$a_{\text{s}} = a_{\text{bg}} \left(1 - \frac{\Delta}{B - B_{\text{peak}}} \right), \tag{1.7}$$

where a_{bg} is the background scattering length, Δ is the width of the resonance and B_{peak} is the magnetic field of the resonance centre.

Each atomic species has its own spectrum of Feshbach resonances. However, most studies with ^{85}Rb focus on one particular resonance. Using high precision bound-state spectroscopy [19], this resonance has been measured to have the properties: $a_{\text{bg}} = (-443 \pm 3)\,a_0$, $\Delta = (10.71 \pm 0.02)\,\text{G}$ and $B_{\text{peak}} = (155.051 \pm 0.018)\,\text{G}$.

1.1.2 Bright-solitary-matter waves

One particular solution to the GPE is of interest in this thesis: the solution that describes a bright-solitary-matter wave. Bright solitons occur in the presence of an attractive, self-focussing nonlinearity which is able to perfectly balance wavepacket dispersion. They manifest as a non-dispersive region of increased density: a *peak*. This contrasts with the dark soliton, which manifests as a non-dispersive region of decreased density: a *dip* [20]; and the gap soliton, which can only be realised in a lattice potential [21].

In the case of the BEC, interatomic interactions can be set to be slightly attractive by tuning the magnetic field in the vicinity of a Feshbach resonance to produce a negative scattering length. This can produce a bright-solitary-matter wave of the form:

$$\psi = \frac{a_{\text{r}}}{\sqrt{2|a_{\text{s}}|}\kappa}\text{sech}\left(\frac{z - vt}{\kappa}\right)\exp\left[i\frac{mv}{\hbar}z - \frac{i}{\hbar}\left(\frac{mv^2}{2} - \frac{\hbar^2\kappa^2}{2m}\right)t\right], \tag{1.8}$$

where $a_\mathrm{r} = \sqrt{\hbar/(m\omega_\mathrm{r})}$ is the characteristic radial confinement length scale, v is the soliton velocity and:

$$\kappa = \frac{a_\mathrm{r}^2}{|a_\mathrm{s}|N},\qquad(1.9)$$

is the soliton width. Typical values for the experiments performed with ^{85}Rb in this work are: $\omega_\mathrm{r} \sim 40\,\mathrm{Hz}$, $a_\mathrm{s} \sim -10\,\mathrm{a_0}$ and $N \sim 2500$, leading to a soliton width $\kappa \sim 2\,\mu\mathrm{m}$. This is approximately half the size of solitons created in experiments using Lithium [22], where the decreased atomic mass and slightly different soliton formation routine contributes to their larger width.

Although Eq. 1.9 is strictly a solution only for the true bright soliton, it indicates the form of the bright-solitary-matter wave in a quasi-1D geometry, despite the fact that the system's integrability has been removed by the relaxation of the dimensionality constraint and the addition of an axial potential.

It is important to note that a bright-solitary-matter wave can only occur in a quasi-1D setting, as the collapse instability of BECs requires that a stable BEC be repulsive in a 3D trap, as detailed below. Therefore, the usual experimental procedure for creating bright-solitary-matter waves is to first create a repulsive condensate (positive scattering length) in a 3D trap before changing the geometry to be quasi-1D and switching the interactions to be attractive (negative scattering length) [23]. This is the same procedure used throughout this thesis (see Ch. 4). Bright-solitary-matter waves in BECs of ^{7}Li, ^{85}Rb, ^{39}K and ^{133}Cs have so far been experimentally demonstrated [23–30].

1.1.3   Collapse instability

To minimise its interaction energy, a trapped condensate with attractive interatomic interactions will increase its density at the trap centre. This may be inconsequential if the interaction strength is sufficiently weak, or the atom

number is sufficiently low that the kinetic energy can balance it. However, beyond these limits, collapse occurs. The stability of the condensate in a 3D trap can therefore be parametrised by a term k [31], which is defined as:

$$k = \frac{N_{\mathrm{crit}}|a_s|}{a_{\mathrm{ho}}}, \tag{1.10}$$

where N_{crit} is the critical number of atoms, $a_{\mathrm{ho}} = \sqrt{\hbar/m\omega_{\mathrm{ho}}}$ is the average confining length scale and $\omega_{\mathrm{ho}} = (\omega_{\mathrm{x}}\omega_{\mathrm{y}}\omega_{\mathrm{z}})^{1/3}$ is the geometric average trapping frequency.

The value for k has been calculated numerically in a variety of trapping configurations. For spherically-symmetric traps, similar to those used for producing BECs in this work, a value of $k \sim 0.57$ has been calculated [31].

In a quasi-1D setting, Eq. 1.10 is modified to:

$$k = \frac{N_{\mathrm{crit}}|a_s|}{a_{\mathrm{r}}}. \tag{1.11}$$

For quasi-1D traps, similar to those used for producing solitons in this work, a value of $k \sim 0.675$ has been calculated [32]. Interestingly, a uniform condensate in free space is unstable at any negative scattering length, as $\omega_{\mathrm{ho}} \to 0$, whereas collapse is prohibited in a true 1D system.

1.2   Motivation

There are many targets for our experiment, some of which we have now achieved and others that may be addressed by future experimentalists. In brief, these targets are to:

- Create a machine and routine for reliably producing robust, long-lived solitons.

- Develop an imaging system for observing the solitons *in-situ*.

- Develop a non-destructive imaging system for observing soliton dynamics in real-time.

- Investigate the phase properties of solitons, in the context of soliton-soliton collisions and soliton interferometry.

- Investigate the post-quench dynamics of a soliton following an interaction quench, to identify possible beyond-mean-field effects.

- Investigate the collapse dynamics of a condensate.

- Investigate short-range atom-surface interactions.

To better explain why these targets are important, we provide motivation for each of those covered in this thesis below.

Robust, long-lived solitons

The foundation of any future soliton experiment is the robust generation of stable solitons. Therefore, we aim to create the most stable bright-solitary-matter waves possible. This requires the reliable production of a BEC with tunable interatomic interactions, along with a 3D density profile that can be mode-matched with that of the target bright-solitary-matter wave in the quasi-1D potential.

In-situ imaging of solitons

As bright-solitary-matter waves are small, dense objects, it has been a common practice to allow them to expand before imaging [23]. This has typically been achieved by switching the magnetic field off, causing the scattering length to jump to the background scattering length of $-443\,a_0$, which in turn triggers the soliton to collapse and rapidly expand [13, 23, 33]. After allowing the cloud to expand for some time, the result is an image of a much larger, less-dense object. However, this technique is extremely detrimental

to achieving an accurate or repeatable measurement of the soliton width, as the magnetic field sweeps across a Feshbach resonance during the switch-off. As the Feshbach resonance is nonlinear, it is highly non-trivial to relate the measured soliton width with the true in-trap value. Therefore, in order to study solitons in detail, as well as to quantitatively compare experiment with theory, it is important that we determine a technique for imaging the solitons *in-situ*.

Secondary to this, many aims of our experiment (and the interests of the field as a whole) relate to soliton dynamics. Historically, the workhorse imaging technique in the cold atoms community is a destructive method known as absorption imaging (see Ch. 3). This technique requires that each image is taken from a totally independent experimental run. Unfortunately, this means that any small-scale dynamics that are sensitive to the exact initial conditions (such as those associated with condensate collapse) are often lost amidst shot-to-shot noise. Therefore, it would be extremely beneficial to investigate a minimally-destructive imaging technique that could be applied to solitons, such as dark ground imaging [34], phase contrast imaging [35,36], Faraday imaging [37,38] or partial-transfer absorption imaging [39].

Phase properties of solitons

Understanding and probing the coherent phase carried by matter-wave solitons is an area of particular relevance for BEC physics. The reason for this is the impact that soliton phase has on two areas: phase-sensitive soliton-soliton collisions and soliton interferometry.

The relative phase between colliding solitons is attributed to be the stabilising force against collapse for soliton trains [29,40–43]. Though the mechanisms for creating the required alternating phase structure ($0 - \pi - 0$ etc.) between neighbouring solitons in a soliton train have been theoretically explored [44–50], there are currently no experiments that directly probe the soliton-soliton collision's dependence on the relative phase. The only experimental investigation of soliton-soliton collisions in BECs to date inferred the

relative phase from the observed outcome of the collision event [40]. There-fore, one aim for the experiment is to study phase-sensitive soliton-soliton collisions in a phase-controlled way.

Matter-wave interferometers have emerged as a means of achieving unprece-dented sensitivity in interferometric measurements [51–54]. However, they have typically been limited by either interatomic collisions or dispersion of the atomic wavepackets, which cause dephasing and a reduced signal to noise, respectively [55]. Previous works have successfully reduced the impact of in-teratomic collisions through the control of interatomic interactions [56, 57], or by generating squeezed states [58, 59]. However, dispersion remains a limita-tion. A soliton-based interferometer has the potential to overcome dispersion, allowing for much longer phase-accumulation times [22, 60–66], albeit for an increased quantum noise [67]. To date, only one experiment has demon-strated interferometry with a soliton [27], in which Bragg pulses were used for splitting and recombination. However, interferometer times were insuffi-cient to exploit the non-dispersive property of solitons. Therefore, another aim of the experiment is to create an interferometer that is able to make use of the beneficial properties of solitons.

1.3 Thesis context

This work is primarily focussed on the splitting and recombination of bright-solitary-matter waves of ^{85}Rb, as well as the optimised production and imag-ing of those same bright-solitary-matter waves. As such, we are mainly in-terested in the very final stages of the experimental sequence. Therefore, the work presented by the author (OJW) in this thesis builds directly upon the contributions made by previous phd students: Sylvi Händel (SH) [12], Anna Marchant (ALM) [13], Tim Wiles (TPW) [14] and Manfred Yu (MMHY). A short summary of their individual contributions to the relevant parts of the experiment for this thesis is listed below.

SH, AM and TW assembled, developed and characterised major aspects of the experimental apparatus, which ultimately led to the first creation of ^{85}Rb BEC through direct evaporative cooling [68]. The initial production of bright-solitary-matter waves was done by ALM, along with demonstrating reflection from a broad Gaussian barrier [23]. ALM, TPW and MMHY showed the reflection of a bright-solitary-matter wave from a narrow attractive barrier [33]. MMHY, along with Ana Rakonjac (AR) were involved in an upgrade of the trapping setup for the science cell stages of the experiment, as well as research into arbitrary potentials and ring traps.

Throughout this thesis we focus on the work undertaken by OJW, which we outline here. OJW constructed the laser and beat-locked laser system used for *in-situ* imaging, along with performing the low-field and high-field spectroscopy. AR constructed the microwave antenna which OJW used for microwave-transfer and partial-transfer imaging. AR and OJW contributed towards calibrating the imaging for high probe intensities. AR and OJW contributed to redesigning and aligning the new crossed-dipole trap lasers, as well as achieving the first ^{85}Rb BEC with the new setup. AR and OJW produced the first solitons in the N-S and E-W beams in this new setup. OJW added and explored the controllable harmonic potential along the quasi-1D potential, using magnetic coils installed by MMHY. OJW performed the field calibration, wavefunction engineering and width oscillation experiments that led to the production of robust, long-lived solitons. AR and OJW installed and aligned the wide barrier beam. AR profiled the width of the wide barrier. AR and OJW investigated controllable splitting and velocity-selection-mediated merging of daughter solitons on the wide barrier. AR and OJW installed the 30 mm lens to create the narrow barrier. OJW profiled the narrow barrier using a variety of different methods. OJW investigated controllable splitting and interference-mediated recombination on the narrow barrier. The study on velocity selection was performed by OJW. At various stages of this thesis there are computational simulations that have been provided by theorists: Tom Billam (TPB) for 1D-GPE and quasi-1D-GPE simulations and John Helm (JLH) for 3D-GPE simulations. TPB also

provided the approximate analytical model used to describe the observations of interference-mediated recombination with respect to barrier offset.

1.4 Thesis outline

The layout of this thesis is arranged such that it presents the results in an order that builds ever-increasing experimental sophistication. Therefore, each chapter depends entirely on the steps taken in the previous chapter. As an overview:

- Ch. 2 details the optimisation of our experimental sequence for producing condensates of ^{85}Rb.

- Ch. 3 gives a detailed account of our new imaging system, capable of imaging the solitons *in-situ* and across a wide range of magnetic fields.

- Ch. 4 provides information on how we ensure the production of robust, long-lived single bright-solitary-matter waves in a tunable quasi-1D potential.

- Ch. 5 reports the controllable splitting of a bright-solitary-matter wave on a narrow repulsive barrier, along with the observation of velocity selection.

- Ch. 6 reports the results of velocity-selection-mediated recombination on the wide barrier, as well as interference-mediated recombination on the narrow barrier. We explore the effect of tuning various experimental parameters and develop an intuitive interpretation for our observations, which encompasses the effects of velocity selection and interference. Finally, we look at the implications that our findings have for the field of interferometry.

- Ch. 7 summarises the work presented in this thesis and gives outlines for future experiments that may be performed with an upgraded apparatus, which could pave the way for fulfilling the unexplored targets set earlier in this chapter.

1.5   Publications

O. J. Wales, A. Rakonjac, T. P. Billam and S. L. Cornish, *Robust production and imaging of long-lived matter-wave solitons*, in preparation: provides the basis for Ch. 3 and Ch. 4.

[69] O. J. Wales, A. Rakonjac, T. P. Billam, J. L. Helm, S. A. Gardiner and S. L. Cornish, *Splitting and recombination of bright-solitary-matter waves*, arXiv e-prints, arXiv:1906.06083 (2019): provides the basis for Ch. 5 and Ch. 6.

Chapter 2

Bose-Einstein condensation of ^{85}Rb

In this chapter we describe the experimental setup and sequence used to produce ^{85}Rb Bose-Einstein condensates (BECs). As the majority of the setup and sequence has been extensively described elsewhere [12–15, 23, 33, 68, 70, 71] we will not provide a comprehensive description of each aspect, instead this chapter serves as a higher-level overview of the main points to put subsequent chapters of this thesis into context. It also highlights and explains important upgrades, changes and additions that have been made to the setup from previous works, such as the upgraded crossed optical dipole trap and dual photodiode servo.

We begin with an outline of the experimental apparatus before presenting sections on each of the two main components of our experiment, in which we cover the relevant stages of the experimental sequence and the background necessary to understand optimised evaporation.

2.1   Apparatus overview

The apparatus consists of a laser system for cooling, the vacuum system, the hardware for optical and magnetic trapping, as well as a computer control and data acquisition system.

The vacuum system (Fig. 2.1) is a two-chamber setup with a differential pumping connection that establishes a differential pressure of > 100 between the higher-pressure magneto-optical trap (MOT) chamber and the ultra-high vacuum (UHV) science cell. This is essential in order to achieve high MOT loads in the MOT chamber whilst still being able to efficiently cool in the science cell. Atoms are transported between the two chambers in a quadrupole magnetic trap generated by a pair of coils, the *transport coils*, which are mounted on a motorised translation stage and take around 2 s to transfer the atoms.

2.2   Cooling in the MOT chamber

Five stages of the experiment take place in the MOT chamber: MOT, compressed MOT (CMOT), optical molasses, optical pumping and transport trap loading. These stages are designed and optimised to load as many atoms as possible into the transport trap, at the highest phase space density.

The 6-beam MOT is made using three pairs of counter-propagating, off-resonant laser beams, along with a magnetic quadrupole that produces an axial magnetic field gradient of $10\,\mathrm{G\,cm^{-1}}$. The beams are taken from the same laser, the *cooling laser*, which is tuned to deliver light detuned by $-15\,\mathrm{MHz}$ from the $\mathrm{D_2}$ $5^2\mathrm{S}_{1/2}$ $F = 3 \to 5^2\mathrm{P}_{3/2}$ $F' = 4$ transition for the MOT stage. The cooling light passes through a tapered amplifier before reaching the experiment, resulting in 12 mW per beam at the vacuum chamber. The MOT beams have a $1/e^2$ radius of 15 mm, limited by the optical access to the vacuum chamber and by the required beam intensity to form a MOT. Despite the $F = 3 \to F' = 4$ being a closed transition, whereby atoms excited on this transition can only decay back to the $F = 3$ ground state, off-resonant scattering to the $F' = 2, 3$ states means that the $F = 2$ state can become populated. Therefore a weak *repump laser* beam, resonant with the $F = 2 \to F' = 3$ transition, returns atoms to the cycling transition for further cooling. The MOT stage typically lasts 20 s and loads up to 10^9 atoms.

Figure 2.1: Top-down view of the two-stage vacuum chamber setup. [a] The full vacuum system comprising the MOT chamber (red, also shown in [b]) and the science cell (purple, also shown in [c]). Atoms are transferred between these two chambers using the magnetic transport coils (blue). [b] The 3D MOT is created using three orthogonal pairs of counter-propagating beams. [c] The optical dipole trap comprises two orthogonal beams (solid red): the *N-S beam* and the *W-E beam*. Imaging beams also pass through the cell (faint maroon), one counter-propagating with the moving beam and one propagating vertically from below the science cell.

A 20 ms CMOT stage follows the MOT stage. Contrary to the findings of Ref. [72], where it is suggested that an increased magnetic gradient would result in a higher number of atoms being loaded into the transport coils, it has been found that a reduced gradient actually produces a larger atom number [12]. Therefore, the magnetic quadrupole gradient is relaxed to $5\,\mathrm{G\,cm^{-1}}$, the cooling beam detuning is increased to $-35\,\mathrm{MHz}$ and the repump intensity is reduced. The reduced repump intensity allows atoms to collect in the $F = 2$ dark state, decreasing the scattering rate and hence the heating rate. The increased cooling beam detuning also reduces heating through re-radiated photons. These changes result in a denser atomic sample to that formed in the MOT.

We cool the atoms further using optical molasses. The magnetic quadrupole field is removed and the cooling beam detuning is increased still further to $-90\,\mathrm{MHz}$ for 30 ms, cooling the atoms through a combination of standard Doppler cooling and polarisation gradient cooling.

To optically pump the atoms into the magnetically trappable, low-field seeking $|F, m_\mathrm{F}\rangle = |2, -2\rangle$ state, we use a 1.5 ms pulse of circularly polarised light such that it drives the required σ^-, $F = 2 \rightarrow F' = 2$ transitions.

The atoms are loaded directly into the transport potential following optical pumping. The transport coils are switched on simultaneously with the MOT beams turning off. Initially they provide a magnetic field gradient of $45\,\mathrm{G\,cm^{-1}}$ but, after a 25 ms hold, this gradient is ramped to $180\,\mathrm{G\,cm^{-1}}$ over 250 ms. It is important that this gradient is high to maintain trapping during the mechanical acceleration/ deceleration of the coils as they move towards the science cell. The magnetic dipole force felt by an atom placed in a magnetic field gradient $\nabla|B|$ is:

$$F_\mathrm{dip} = -m_\mathrm{F} g_\mathrm{F} \mu_\mathrm{B} \nabla|B|, \tag{2.1}$$

where g_F is the atom's hyperfine Landé g-factor and μ_B is the Bohr magneton. Therefore, for a potential with field gradient $180\,\mathrm{G\,cm^{-1}}$ vertically

$(90\,\mathrm{G}\,\mathrm{cm}^{-1}$ horizontally), an atom in the $|2, -2\rangle$ state allows a maximum trap acceleration of $\sim 40\,\mathrm{m}\,\mathrm{s}^{-2}$ to remain trapped. In practice, the accelerations used in the experiment fall well short of this, with the motorised translation stage only capable of accelerations up to $4\,\mathrm{m}\,\mathrm{s}^{-2}$.

A right-angled prism is located halfway along the path of the transported atoms. This blocks the line of sight between the two chambers, preventing stray hot atoms from entering the science cell. Additionally, it prevents damage of the MOT fluorescence camera from the high power dipole-trapping beams used in the science cell.

2.2.1 Cooling in the MOT chamber: optimisation

To maximise the number of atoms in the final condensate, it is essential to load as many atoms as possible into the transport trap. Thus, at each of the stages highlighted above, the cooling beam pair intensity balancing, cooling beam detuning, repump power and detuning, magnetic quadrupole gradient and magnetic quadrupole trap centre have to be optimised. The magnetic quadrupole trap centre is controlled using pairs of magnetic shim coils (labelled in Fig. 2.1[b]). Optimisation is achieved in two ways: direct observation of the MOT for coarse optimisation and a release-recapture method for precise optimisation.

The MOT, CMOT and molasses stages are able to cool atoms because of preferential scattering of photons by higher temperature atoms, which causes fluorescence. As the $F = 3 \to F' = 4$ transition is around $780\,\mathrm{nm}$, this fluorescence is faintly observable by the human eye and easily observable with commercial silicon-based photodiodes and cameras. We monitor the fluorescence with a camera to observe the position of the MOT within the vacuum chamber and with a photodiode to give a quantitative estimate of the total atom number. The output voltage of the photodiode is directly proportional to the atom number, for a constant MOT beam intensity and detuning [73].

An example of coarse optimisation using the photodiode is in the initial loading of the MOT. As we want to transfer the maximum number of atoms to the transport potential as possible, it is reasonable to assume that we want to load the largest MOT that we can. Therefore, early optimisations of the variables during the MOT stage were done to maximise the measured fluorescence. However, this measurement is insensitive to atomic temperature and so may not represent the optimal set of parameters for the eventual loading of the transport potential.

The camera is used for coarse optimisation of the cloud position during the various stages, as well as for balancing the cooling beams. As there are likely small differences in beam radii between the various cooling beams, a simple measurement of beam power is insufficient to properly balance the intensity between beam pairs. Instead, coarse optimisation is performed by observing the fluorescence from an extended molasses stage (after normal MOT and CMOT stages) with the camera. If the beam pairs are imbalanced the atomic cloud is rapidly pushed towards the weaker beam due to the radiation pressure differential. By carefully adjusting the power in each beam the atomic cloud instead disperses isotropically over $\sim 1\,\mathrm{s}$ timescales, remaining centred on its initial position. This allows for more efficient molasses cooling without compromising the spatial overlap with the transport trap.

More precise optimisation of all parameters is achieved through a release-recapture method [74]. Following the stage of interest, the atoms are loaded into the transport potential at a magnetic field gradient of $45\,\mathrm{G\,cm^{-1}}$, consistent with the gradient used during the normal experimental sequence. Using Eq. 2.1 we can see that this gradient is sufficient to support the $m_\mathrm{F} = -2$ and $m_\mathrm{F} = -1$ states against gravity, which require $\nabla B = 22.4\,\mathrm{G\,cm^{-1}}$ and $44.8\,\mathrm{G\,cm^{-1}}$ respectively. However, only the atoms in the $m_\mathrm{F} = -2$ state are trapped. After a hold of $400\,\mathrm{ms}$, the transport coils are switched off, releasing the atoms from the transport potential. The MOT stage is immediately switched on and some fraction of atoms are recaptured into it. The measured fluorescence at the start of the recaptured MOT stage, relative to that at the end of the initial MOT stage, gives a release-recapture efficiency. This is a

direct measurement of the efficiency of loading the transport potential and so is the important parameter to be optimised throughout all stages that take place in the MOT chamber.

Cooling in the MOT chamber: stability

Most of the above stages are fairly robust to small short-term drifts and rarely need reoptimising. However, for the benefit of future experimentalists, it's worth highlighting that the cooling laser beam balancing and the optical pumping stage require more frequent attention for the following reasons.

As described above, imbalanced beam intensities causes rapid loss during the molasses stage and, because the beam balancing is controlled through polarisation-sensitive optics, the purity and stability of the polarisation from the cooling laser must be high. However, although the light is passed through a polarisation-maintaining optical fiber before being split into the six beam paths, ambient temperature and pressure fluctuations can lead to imbalanced beam pairs and inefficient molasses.

Efficient optical pumping requires sufficient photons in order to transfer $\sim$ $100\,\%$ of the atoms to the $|2, -2\rangle$ state. Therefore, if the beam power is too low or the pulse is too short, there is a drastic reduction in recapture efficiency. The efficiency also depends on the purity of the polarisation to drive the necessary σ^- transitions. If the purity of polarisation is not high, pulses that are too long or too powerful can cause heating and a reduction in recapture efficiency. Assuming all atoms are initially equally distributed over the five hyperfine levels of the $F = 2$ state and with perfect optical pumping efficiency, the optical pumping stage changes the fraction of *theoretically* magnetically-trappable atoms from $20\,\%$ to $100\,\%$, at a field gradient of $\nabla B = 45\,\mathrm{G\,cm^{-1}}$. Consequently, the optical pumping stage has a drastic effect on the transport potential loading and has been found to require the most frequent optimisation.

2.3   Cooling in the science cell

We cool the atoms to degeneracy in the science cell, using the well-established techniques [75, 76] of radio frequency (RF) evaporation and optical evaporation. Fig. 2.2 demonstrates the evaporation trajectory taken, from loading the quadrupole trap to achieving condensation.

Figure 2.2: Evaporation trajectory to form a BEC of ^{85}Rb. The blue data show the values at the end of the two final RF ramps and the rest of the data show the values at the end of optical evaporation ramps. The diagonal dashed lines separate the regions with no applied bias field, $177\,\mathrm{G}$ and $162\,\mathrm{G}$ from right to left respectively. The green triangles are data measured during optimising the final evaporation ramp and so correspond to different final trap depths. The red dotted line highlights the BEC phase transition at $PSD = 2.61$.

The evaporation sequence is comprised of a series of linear ramps that approximate an exponential decay of the trap depth. For the RF evaporation these are frequency ramps, from high to low frequency; for the optical evaporation these are laser intensity ramps, from high to low intensity.

The specific scattering properties of ^{85}Rb, which will be discussed later in this section, make evaporating to condensation particularly challenging [68]. Therefore, it is essential that the evaporation procedure is well-optimised. We

ensure this by considering elastic and inelastic collisions, as well as by defining an evaporation efficiency parameter that we can optimise experimentally.

2.3.1 Elastic and inelastic collisions

The efficiency of evaporation depends on the ratio of elastic to inelastic collisions. Following the removal of the most energetic atoms, elastic collisions are responsible for rethermalising the remaining ensemble. This reduces the overall temperature and ideally increases the phase space density, PSD, given by [77]:

$$PSD = n_0 \lambda_{\mathrm{dB}}^3 = n_0 \left(\frac{2\pi\hbar^2}{mk_{\mathrm{B}}T} \right)^{\frac{3}{2}},$$ (2.2)

where λ_{dB} is the de Broglie wavelength and n_0 is the peak density of the atomic cloud. For a harmonic trapping potential with peak density given by $n_0 = N\omega^3 \left(\frac{m}{2\pi k_{\mathrm{B}}T} \right)^{3/2}$, Eq. 2.2 can also be recast to:

$$PSD = N \left(\frac{\hbar\omega_{\mathrm{ho}}}{k_{\mathrm{B}}T} \right)^3,$$ (2.3)

where $\omega_{\mathrm{ho}} = (\omega_{\mathrm{x}}\omega_{\mathrm{y}}\omega_{\mathrm{z}})^{1/3}$ is the geometric average of the three orthogonal trapping frequencies.

As described in Ch. 1, the s-wave scattering length varies nonlinearly with magnetic field close to Feshbach resonances. For atoms in the ultracold limit, the elastic scattering cross section is approximated by:

$$\sigma_{\mathrm{el}} = \frac{8\pi a_{\mathrm{s}}^2}{1 + k^2 a_{\mathrm{s}}^2},$$ (2.4)

where k is the wave vector defined by:

$$k = \left(\frac{16mk_{\mathrm{B}}T}{\pi\hbar^2} \right)^{\frac{1}{2}} = \frac{4\sqrt{2}}{\lambda_{\mathrm{dB}}}.$$ (2.5)

This approximation ignores an effective range correction, r_e, which modifies the scattering length to account for finite temperature effects. Therefore, Eq. 2.4 is strictly valid only when $r_e k^2 a_s \ll 1$, or equivalently when $\lambda_{dB}^2 \gg r_e a_s$, which occurs below $\sim 10\,\mu\text{K}$ for ^{85}Rb. However, this approximation is sufficient to understand that elastic collisions can be maximised by careful selection of the scattering length. The full elastic collision rate, Γ_{el}, also depends on the mean atomic density, $\langle n \rangle$, and the mean relative velocity, $\langle v_{rel} \rangle$, as:

$$\Gamma_{el} = \sigma_{el} \langle v_{rel} \rangle \langle n \rangle. \tag{2.6}$$

Unfortunately, inelastic collisions also depend on the scattering length (three-body loss scales as a_s^4) and cause loss without the increased PSD [78, 79]. Therefore, the optimal scattering length for maximising the elastic scattering rate is not necessarily the optimal scattering length for evaporation. This is particularly important because density dependent inelastic losses, such as two-body and three-body collisions, preferentially remove atoms from the colder, denser centre of the atomic cloud, leading to heating and *anti-evaporation*.

2.3.2 Evaporation efficiency

To achieve the best ratio of elastic to inelastic collisions and achieve the largest condensates, we experimentally optimise the evaporation efficiency at each stage of evaporation. The evaporation efficiency (γ) of a particular evaporation stage is characterised by:

$$\gamma = -\log\left(\frac{PSD_f/PSD_i}{N_f/N_i}\right), \tag{2.7}$$

where the subscripts 'i' and 'f' denote the initial and final values of the evaporation stage respectively.

Bose-Einstein condensation occurs when λ_{dB} exceeds the interatomic separation, such that there is significant wavefunction overlap, or equivalently when $PSD \geq 2.61$ [80]. Therefore, the target of the evaporation sequence is to reach $PSD = 2.61$ with as many atoms as possible, by optimising γ on a stage-by-stage basis.

2.3.3 Evaporation of ^{85}Rb to BEC

As mentioned in Sec. 2.3, evaporation to degeneracy with ^{85}Rb is particularly challenging [68]. Firstly, ^{85}Rb possesses a negative ($-443\,a_0$) background scattering length, therefore stable BECs can only be produced using Feshbach resonances. Secondly, the elastic cross section is strongly temperature dependent, with a Ramsauer-Townsend type minimum around $375\,\mu K$ that severely extends rethermalisation times and makes adiabatic compression ineffective [81]. Thirdly, in general the elastic cross section is low, as much as two orders of magnitude lower than that of ^{87}Rb for temperatures above $\sim 100\,\mu K$. This again slows rethermalisation and requires that evaporation ramps are much slower, giving more time for loss mechanisms to take place. Finally, far from Feshbach resonances, two-body and three-body inelastic losses are typically about two and five orders of magnitude larger than for ^{87}Rb respectively [78, 79]. This highlights how unfavourable direct evaporation to condensation could be with ^{85}Rb without careful consideration of elastic and inelastic collisions.

Fortunately, with sufficient optimisation of the evaporation sequence it is possible to reach degeneracy, albeit for $\sim 1\%$ of the number of atoms in ^{87}Rb BECs made using the same experimental apparatus [13]. By tuning the scattering length close to the Feshbach resonance at $155\,G$, it has been found that there are two peaks in evaporation efficiency (Fig. 2.3) that are sufficiently high for us to produce BECs.

Figure 2.3: Evaporation efficiency of a fixed evaporation sequence carried out at different magnetic fields close to the Feshbach resonance at $155\,\mathrm{G}$. [a] is the explored region of the Feshbach resonance and [b] is the evaporation efficiency. The dotted vertical lines in both figures show the magnetic field at $a_\mathrm{s} = 0$ but the solid line in [b] is a guide to the eye. Data for [b] is taken from Ref. [68] and the parameters for [a] are taken from Ref. [19].

2.3.4 Cooling in the science cell: RF evaporation

The first method of evaporation used is RF evaporation. After the atoms have been transported to the science cell, they are transferred to a different magnetic quadrupole trap (the *quadrupole trap*) by simultaneously ramping down the transport coils and ramping up the quadrupole trap coils to $179.36\,\mathrm{G\,cm^{-1}}$, matching the gradient provided by the transport coils before the ramp. We then apply an RF-sweep, from high to low frequency, which couples the hyperfine states for atoms within a particular shell of the quadrupole field gradient, forcing atoms to spin-flip through the $m_\mathrm{F} = -2 \rightarrow -1 \rightarrow 0 \rightarrow 1 \rightarrow 2$ transitions. Eq. 2.1 shows that atoms transferred to the $m_\mathrm{F} = 0, 1, 2$ states are no longer magnetically trapped (in fact $m_\mathrm{F} = 1, 2$ are anti-trapped) and are lost from the trap.

RF evaporation achieves temperature selectivity because the more energetic atoms explore more of the trap, reaching further up the slopes of the magnetic quadrupole potential. This results in temperature-dependent Zeeman

splitting which means that the more energetic atoms are the first to be removed by the RF sweeps. If the elastic-inelastic collision ratio is favourable, the remaining atoms rethermalise to a lower temperature and higher PSD.

We optimise a series of three RF ramps by aiming to halve the temperature of the cloud at each stage, whilst removing as few atoms as possible. Unfortunately, the limited field of view and large spatial extent of the clouds makes it difficult to accurately determine N and the PSD for the initial loading of the quadrupole trap and first stage of RF evaporation, for the current imaging magnification. Therefore, these points are missing from Figs. 2.2 and 2.4. Instead, these stages are optimised by maximising the optical density.

The magnetic quadrupole gradient remains at $179.36\,\mathrm{G\,cm^{-1}}$ throughout RF evaporation and we find an optimised set of RF ramps to be: $15\,\mathrm{MHz} \rightarrow 7\,\mathrm{MHz}$ over 15 s, $7\,\mathrm{MHz} \rightarrow 4\,\mathrm{MHz}$ over 5 s and finally $4\,\mathrm{MHz} \rightarrow 2.5\,\mathrm{MHz}$ over 1 s (see Fig. 2.4). This results in 3.7×10^7 atoms at $PSD = 7.4 \times 10^{-5}$ and $T = 40\,\mathrm{\mu K}$. At this point, Majorana spin flips become significant in limiting the trap lifetime because of the inverse square dependence of the Majorana loss rate on temperature [82]. Therefore, instead of performing further RF evaporation in the magnetic quadrupole trap, we transfer the atoms to a hybrid trap to perform optical evaporation, similar to the methods used in Ref. [82].

2.3.5   Loading the hybrid trap

After the final RF ramp, the magnetic quadrupole trap is ramped down from $179.36\,\mathrm{G\,cm^{-1}}$ to $20.85\,\mathrm{G\,cm^{-1}}$ over 3 s, transferring the population into a hybrid optical dipole + magnetic quadrupole trap (the *hybrid trap*). The crossed optical dipole trap is on throughout the RF evaporation but does not contribute significantly to the potential because of the large magnetic potential. The crossed optical dipole trap is positioned $\sim 100\,\mathrm{\mu m}$ below the magnetic quadrupole field zero to prevent Majorana spin flips.

Figure 2.4: The evaporation sequence used to form a BEC of ^{85}Rb in the science cell. The left hand side of the figure shows the three RF evaporation stages as well as the loading into the hybrid trap. The right hand side shows the eight optical evaporation stages, along with the changes in magnetic field necessary to optimise evaporation efficiency and to create a positive scattering length. The upper panel shows the RF frequency used throughout RF evaporation. The following four panels display the magnetic quadrupole gradient, the magnetic field at the atoms and the N-S and W-E dipole trap beam powers, from top to bottom respectively. The bottom two panels display the number and *PSD* throughout the evaporation trajectory.

The crossed optical dipole trap

The crossed optical dipole trap is comprised of two independent laser beams that propagate in the horizontal plane: the N-S and W-E beams (as defined in Fig. 2.1) with $1/e^2$ radii of 97 µm and 70 µm at the atoms respectively. The N-S beam is produced by a single frequency IPG 15 W fibre laser at 1064 nm (IPG: YLR-15-1064-LPSF) whilst the W-E beam is produced by a 50 W multimode IPG laser at 1070 nm (YLR-50-LP-AC-Y12). This is an upgraded setup from that used in previous works [12–15, 23, 33, 68, 70], where limitations in trapping power, flexibility and loss were encountered.

It is worth noting, for future experimentalists, that the losses occurred when using the above 1070 nm laser for both beams in the crossed dipole trap setup. As the laser is multimode, enhanced two-body loss rates (similar to those in Refs. [83, 84]) were observed. These losses appear to have originated from a two-photon Raman process that coupled the ground hyperfine states. As the laser is multimode, it is likely that the resonant frequency between the two ground states exists among the spacings between two laser modes. Therefore, in our realisation, one of the 1070 nm beam paths has been swapped out and replaced by the above 1064 nm laser, preventing the two-photon loss mechanism.

2.3.6 Cooling in the science cell: optical evaporation

Optical evaporation is performed by reducing the intensity of the optical dipole trap beams, lowering the trap depth and allowing the most energetic atoms to escape the trap. After the trap depths are lowered, the beams are kept at a constant value for a *hold* time, allowing the atoms to rethermalise and reach a lower equilibrium temperature.

During the first optical evaporation stage only the W-E beam is ramped down (see Fig. 2.4). This is because both beams begin the stage on full power but the W-E is narrower and has a higher power limit than the N-S beam. At

the end of the first ramp both beams have the same peak intensity, thus the same trap depth. Further evaporation ramps are performed by ramping down both beams simultaneously, matching peak intensities throughout.

Atoms lost during the first two evaporation stages spill out of the trap vertically (see Fig. 2.5[a]) as the curvature of the magnetic quadrupole trap prevents escape horizontally along the optical dipole beams but the vertical magnetic field gradient of $20.85\,\text{G}\,\text{cm}^{-1}$ is insufficient to cancel gravity ($22.4\,\text{G}\,\text{cm}^{-1}$ would be required to levitate).

Figure 2.5: Trap potentials for stages of the optical evaporation sequence. [a] With only the crossed optical dipole trap and magnetic quadrupole trap evaporation occurs vertically, aided by gravity, as the horizontal directions are closed off by the magnetic quadrupole. [b] With the addition of a uniform bias field the quadrupole centre is shifted far away vertically, causing the potential curvature in the horizontal directions to become greatly diminished. The blue lines indicate the vertical potential and the red dotted lines show the horizontal potential along one dipole trapping beam. The black arrows demonstrate the primary evaporation channel.

The strength of the horizontal confinement depends on the vertical offset (z_{vert}) between the optical and magnetic trap centres ($z_{\text{vert}} \sim 100\,\mu\text{m}$ at this point). The offset can be controlled by the bias field, which moves the trap centre $\Delta z_{\text{vert}} = B/\nabla B$ vertically. Separating the contributions of the optical and magnetic traps to the full potential [82] we find the magnetic component:

$$U_{\text{mag}} = m_{\text{F}} g_{\text{F}} \mu_{\text{B}} \nabla B \sqrt{\frac{x^2 + y^2}{4} + z_{\text{vert}}^2}.$$ (2.8)

To evaluate the contribution made along a single dipole beam, aligned parallel with either x or y, we assume the beam is perfectly centred on the trap in the transverse direction ($y = 0$ for a beam aligned parallel with x). This allows us to simplify and recast Eq. 2.8 into the following, for a beam aligned parallel with x:

$$U_{\mathrm{mag}} = m_{\mathrm{F}} g_{\mathrm{F}} \mu_{\mathrm{B}} \nabla B z_{\mathrm{vert}} \sqrt{\frac{x^2}{4z_{\mathrm{vert}}^2} + 1}, \tag{2.9}$$

which, for $x < z_{\mathrm{vert}}$, can be expanded to:

$$U_{\mathrm{mag}} \approx m_{\mathrm{F}} g_{\mathrm{F}} \mu_{\mathrm{B}} \nabla B \left(z_{\mathrm{vert}}^2 + \frac{x^2}{8z_{\mathrm{vert}}^2} \right). \tag{2.10}$$

Comparing this expression with the potential of a harmonic oscillator, $U = \frac{1}{2} m \omega^2 x^2$, we find that the second term denotes the contribution by the magnetic quadrupole trap to the axial trapping frequency along a single dipole beam ($\omega_{\mathrm{x,y}}$):

$$\omega_{\mathrm{x,y}} = \frac{1}{2} \sqrt{\frac{m_{\mathrm{F}} g_{\mathrm{F}} \mu_{\mathrm{B}} \nabla B}{m z_{\mathrm{vert}}}}. \tag{2.11}$$

It is important to note that this expression is only valid close to $x, y = 0$ as the quadrupole potential tends towards a linear relation for larger x and y.

Following the low-field evaporation stages we apply a $177\,\mathrm{G}$ uniform bias field to the atoms to exploit the window of efficient evaporation seen in Fig. 2.3. The additional field also has the effect of pushing the magnetic quadrupole trap centre far away ($\Delta z_{\mathrm{vert}} = B/\nabla B \approx 8.5\,\mathrm{cm}$). This reduces the curvature of the magnetic trap along the optical dipole beams from $\sim 24\,\mathrm{Hz}$ without the bias field, to $< 1\,\mathrm{Hz}$ with the bias, as calculated using Eq. 2.11. Therefore, atoms lost through further evaporation will leave the trap horizontally along the optical dipole trap beams (see Fig. 2.5[b]).

A further four stages of evaporation are performed at this bias field, resulting in 1.85×10^5 atoms at $PSD = 9.7 \times 10^{-2}$. As the scattering length at this

field is negative ($\sim -225\,\mathrm{a_0}$), it isn't possible to produce a stable BEC here. Therefore, we sacrifice some evaporation efficiency and change the scattering length to $219\,\mathrm{a_0}$, ramping the bias field to $\sim 162\,\mathrm{G}$ over $50\,\mathrm{ms}$. A final two stages of evaporation enable us to form BECs. These stages necessarily have longer ramp and hold times than previous stages because of the reduction in elastic collisions and low $\langle v_\mathrm{rel} \rangle$ (see Eq. 2.6).

Trap stabilisation

To vary the beam power over two orders of magnitude, whilst still achieving precision at low powers, we have upgraded the N-S beam servo to a dual photodiode servo setup. A small portion of the N-S beam is passed onto a non-polarising 50:50 beamsplitter, where each outgoing beam is detected by an independent photodiode (Fig. 2.6[a]). One photodiode has a gain that is $\sim 13\times$ greater than the other. This gives much greater control at low powers (to $\sim \mathrm{mW}$ precision) whilst still providing control around the laser's maximum output of $15\,\mathrm{W}$. A calibration of this scheme is shown in Fig. 2.6[b], demonstrating a characteristic response *elbow*.

As the optical evaporation ramps used in the experimental sequence are linear intensity ramps, we make special consideration for this response elbow. The fourth evaporation ramp is shorter and shallower than neighbouring stages (see Fig. 2.4) as the N-S power reaches the elbow at the end of this stage. This allows us to maintain a linear intensity ramp across each stage.

The W-E beam has a different arrangement for achieving a wide range of controllable powers. For high powers we directly control the power through the analog input of the laser controller and it is only for $P \lesssim 1\,\mathrm{W}$ that we control the power through a single photodiode servo. Future upgrades to the experiment could include installing a similar servo for the W-E beam as used for the N-S beam, giving even greater control at low powers.

The quadrupole trap and bias field coils are controlled using servos that receive signals from Hall sensors. A full description of the magnetic field servos are available in Ref. [12] and will not be repeated here.
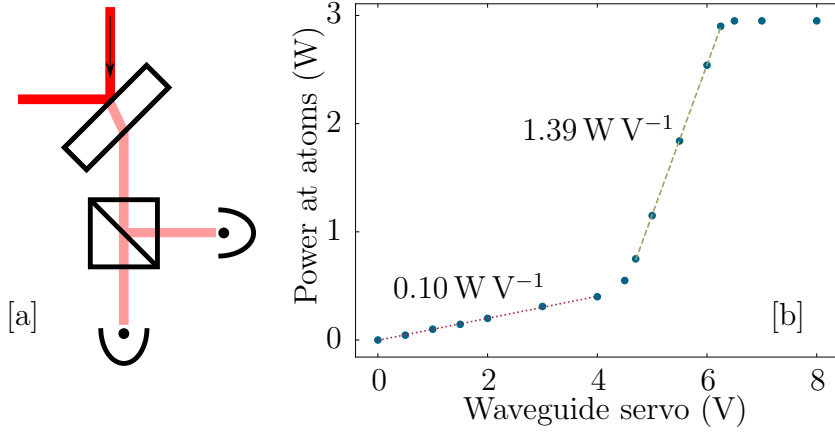
Figure 2.6: A figure to show the dual photodiode servo. [a] the optical layout: a small fraction of the N-S beam is taken from the leakage through a mirror before being passed onto the 50:50 beamsplitter and photodiodes. [b] is a calibration of the dual photodiode servo used for controlling the N-S power. Greater control of the laser power is achieved at low powers owing to the high-gain photodiode (red dotted line). This doesn't comprimise the ability to servo at high powers because the lower-gain photodiode (green dashed line) takes over once the high-gain photodiode reaches its $\sim 4.2\,\mathrm{V}$ voltage limit. The power saturates at servo voltages above $\sim 6.5\,\mathrm{V}$ as the maximum output power of the laser has been reached.

2.3.7   BEC optimisation

We explore the BEC transition by varying the final trap depth of the eighth evaporation ramp (see Fig. 2.4) with fixed ramp and hold durations of $3\,\mathrm{s}$ and $2\,\mathrm{s}$ respectively. When the final trap depth is lowered a BEC fraction emerges because part of the cloud crosses the $PSD = 2.61$ threshold (see Fig. 2.7). However, once the trap depth goes below $\sim 30\,\mathrm{nK}$, atoms from the BEC itself are lost because the crossed optical dipole trap becomes too weak to make up for the difference between magnetic field gradient and gravity described in Sec. 2.3.6.

Unlike the Gaussian density distribution of thermal atoms during time-of-flight expansion (TOF), BECs manifest as an inverted parabola when released from a harmonic potential in the Thomas-Fermi approximation [17,85], which is valid for $Na_{\mathrm{s}}/a_{\mathrm{ho}} \gg 1$. When both thermal and condensate atoms are present, a bimodal density distribution is observed (as in the upper panels of Fig. 2.7).

To distinguish the two distributions more easily we jump the scattering length to zero during TOF, slowing the expansion of the BEC component whilst leaving the thermal fraction essentially unaffected because the atomic kinetic energy is much larger than the interatomic interaction energy. Whilst this does break the requirements of the Thomas-Fermi approximation, which favours strong interatomic interactions, it produces a smaller, more dense BEC density profile and a broader, more sparse thermal density profile. Atom numbers are extracted for each component in order to calculate a BEC fraction, which is used to optimise the final evaporation ramp. We are able to achieve BECs of ~ 7000 atoms with $> 80\,\%$ purity, in a trap with $(\omega_{\mathrm{r}}, \omega_{\mathrm{z}}, \omega_{\mathrm{vert}}) = 2\pi \times (50, 30, 50)\ \mathrm{Hz}$ and a trap depth of around $50\,\mathrm{nK}$. Assuming the condensate is in the Thomas-Fermi limit here, this corresponds to a chemical potential of $1.6\,\mathrm{nK}$ [86]. We have found this to be more than sufficient for soliton production (see Ch. 4). It is worth noting that $^{85}\mathrm{Rb}$ requires a relatively weak trap and low density because of high loss rates, compared with condensates of other species.

Figure 2.7: Thermal to BEC transition as a function of the trap depth at the end of the final optical evaporation ramp. The raw images are shown above for trap depths of $5.8\,\mathrm{nK}$, $36\,\mathrm{nK}$, $130\,\mathrm{nK}$, $192\,\mathrm{nK}$ and $237\,\mathrm{nK}$ from left to right respectively, along with bimodal fits to the column/row sums in order to determine the atom number. The BEC fraction is shown in each image. The column/row sum of the BEC density component is fitted with a Thomas-Fermi distribution (green parabola) and the thermal density component is fitted with a Gaussian distribution (red). The black line shows the sum of the two fits. The inset shows the atom number in the thermal and BEC components (red circles and blue triangles respectively) highlighting that the thermal fraction diminishes as the trap depth is reduced, whilst the BEC component grows. All error bars correspond to the standard error from the fits to atomic densities.

We can also gain an understanding of the BEC transition by measuring the temperature of the atoms relative to the critical temperature, T_c. For temperatures below T_c, an atomic cloud in a harmonic trap has a condensate fraction given by [80]:

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c}\right)^3. \tag{2.12}$$

In Fig. 2.8 we combine the condensate fraction from Fig. 2.7 with temperature measurements of the thermal component.

Figure 2.8: The condensate fraction as a function of atomic temperature across the thermal to BEC transition. The temperature is determined from a fit to the thermal component of the cloud, assuming ballistic expansion. The red dotted line is a theoretical prediction of the BEC fraction for temperatures below T_c. The discrepancies between theory and experiment here are likely dominated by systematic uncertainties of the imaging system in determining atom number.

Temperature measurements are performed only on the thermal component because it undergoes ballistic expansion during a free time of flight, whereas the BEC does not [17]. The thermal atomic cloud expands during time of flight, of duration τ, as:

$$\sigma(\tau) = \sigma(0)\sqrt{1 + \frac{k_B T}{m}\tau^2}, \tag{2.13}$$

where $\sigma(0)$ and $\sigma(\tau)$ are the initial and final widths of the cloud respectively.

Chapter 3

Imaging

Our experimental setup includes two independent imaging systems: one orientated horizontally and one vertically. The horizontal imaging system uses weak-probe absorption imaging techniques. It is used predominantly for field calibrations, BEC characterisations, evaporation ramp optimisations etc., as its lower magnification ($4.98\,\mu\mathrm{m\,px}^{-1}$ on the CCD) provides a wide field of view for imaging larger clouds. The vertical imaging system has a number of additional features to enable better imaging of solitons and other small atomic samples: a higher magnification (either $2.69\,\mu\mathrm{m\,px}^{-1}$ or $0.83\,\mu\mathrm{m\,px}^{-1}$ depending on the choice of objective lens), a highly-tunable beat lock and microwave antenna setup for high-field imaging, as well as more power for high-intensity imaging. In this chapter we detail the various imaging methods used in our experiment and highlight why they are beneficial for precise measurements of solitons.

3.1   Absorption imaging

Absorption imaging is the workhorse imaging technique in BEC experiments, used extensively since the first experimentally-realised BEC [87]. In principle, absorption imaging works by exposing an atomic sample with light that is

resonant with an atomic transition and observing the shadow it casts onto a CCD. If the transition of interest is *closed*, such that any atom that absorbs a photon will decay back to the same state, the observed image can be used to calculate the number of atoms in the sample. Such a transition is also known as a *cycling transition*, as the atom will continuously cycle between a ground and excited state as it absorbs and emits photons.

The intensity of light that penetrates through a cloud of atoms with a line-of-sight density given by the optical depth (OD) can be determined from the Beer-Lambert law:

$$I = I_0 e^{-OD}, \tag{3.1}$$

where I_0 and I are the intensities before and after the light has passed through the atomic cloud respectively.

The optical depth is calculated by taking independent measurements of I_0 and I. We take one image with the atoms present, one without and one without any imaging light to detect dark counts and ambient background counts, measuring I_atoms, I_probe and I_bg respectively. Rearranging Eq. 3.1, the OD in our experiment is found by calculating:

$$OD = \ln\left(\frac{I}{I_0}\right) = \ln\left(\frac{I_\text{probe} - I_\text{bg}}{I_\text{atoms} - I_\text{bg}}\right). \tag{3.2}$$

In practice, off-resonance scattering leads to loss from the imaging state, meaning that the transition used for low-field imaging in our experiment is not fully closed. Therefore, a weak *repump* laser is used to transfer atoms from the imaging *dark* state back into the imaging state. This maintains the validity of the 2-level model.

3.1.1 Determining atom number

As the OD is a measure of the line-of-sight density of the atomic sample, we can use it to determine the total number of atoms. Described below are two methods for calculating the atom number, each having advantages in different situations.

Fitting method

For an atomic cloud with a 2D Gaussian density profile, which is generally a good approximation for thermal clouds and solitons, the number of atoms can be calculated from orthogonal Gaussian fits:

$$N_{\text{fit}} = \frac{2\pi OD_{\text{peak}} w_{\text{x}} w_{\text{y}}}{\sigma_0},\tag{3.3}$$

where OD_{peak}, w_{x} and w_{y} are the fitted maximum optical depth and the fitted widths for the two orthogonal directions respectively. $\sigma_0 = 3\lambda^2/2\pi$ for a closed, 2-level system.

For the avoidance of any confusion, the width is defined as the $1/e^2$ radius throughout this thesis unless specified otherwise, where the $1/e^2$ radius, w, is defined by a Gaussian function of the form:

$$f(x) = Ae^{-2\left(\frac{x-x_0}{w}\right)^2}.\tag{3.4}$$

This is not to be confused with the standard deviation of a Gaussian function, which is smaller by a factor of 2.

Pixel-sum method

As solitons are typically very small and only occupy a few pixels in an image, we often encounter significant uncertainty in w_{x} and w_{y} due to poor fitting. We also find poor fitting, hence a poor measure of the number of atoms, when the image contains imaging artefacts such as fringes. Therefore it is often more appropriate to calculate the number of atoms by the *pixel-sum* method:

$$N_{\text{px}} = \sum_{ROI} \frac{OD_{\text{px}} A_{\text{px}}}{\sigma_0},\tag{3.5}$$

where OD_{px} is the optical depth of a single pixel and A_{px} is the area of a single pixel in the plane of the atoms. The sum is performed over all of the pixels in a selected region of interest (ROI).

This technique is very useful for calculating the atom number of small features and with images that contain imaging artefacts, but it also comes with limitations. It is likely that the selected region of interest will contain pixels that are just imaging background, therefore it is essential that the I_{atoms} and I_{probe} images have the same counts level away from the atoms. In other words, the OD images found using Eq. 3.2 must have a zero offset (see Fig. 3.1). This is important to prevent artificial inflation/ reduction of N_{px}. We ensure that the probe intensity remains constant throughout the exposures and that the probe durations are consistent, using feedback from an acousto-optic modulator (AOM), allowing us to accurately determine the number of atoms using this method.

Comparing methods

Both of the atom number measurement techniques benefit from a large number of camera counts to reduce uncertainty. For the fitting method we attain better Gaussian fits when the features are more pronounced. For the pixel-sum method we find that the uncertainty in the zero offset, $\alpha_{OD=0} \propto I^{-1/2}$. Therefore, we always image with intensities and durations to give a readout close to the single-pixel bit-depth limit of the CCD (*Andor LUCA* with a 14-bit output $\equiv 2^{14} = 16,384$ possible different readout values).

Essentially, both methods of determining the atom number are equivalent for large atomic clouds with a zero background (Fig. 3.1[a]); the fitting method should be used for large clouds with a non-zero offset (Fig. 3.1[b]) and the pixel sum method for small objects, non-Gaussian objects or for images with large imaging artefacts (Fig. 3.1[c]).

3.1.2   Imaging corrections

As our camera has a 14-bit output, it should be able to measure ODs of up to $ln\left(2^{14}\right) = 9.7$. However, by imaging dense thermal clouds and identifying

Figure 3.1: An example set of OD images, along with the column/ row sums and their associated Gaussian fits. [a] shows a large atomic cloud with $\sim 10^5$ atoms and a Gaussian fit that well characterises the atomic distribution. Both atom number measurement techniques would work in this situation. [b] is the same image as [a] but with an artificial OD offset to demonstrate a situation where the pixel sum would be inappropriate for measuring the number of atoms. [c] is an image of a single soliton. This is a situation where the Gaussian fits demonstrate considerable deviation from the column sum, due to the small soliton size and additional imaging artefacts. Therefore the pixel sum method should be used to determine the number of atoms most accurately. The pixel size for [a] and [b] is $4.98\,\mu\mathrm{m}$ whereas for [c] it is $0.83\,\mu\mathrm{m}$.

the flat top in a cross-cut through the cloud, we find that the OD saturates to ~ 3.8. The reasons of this limit are two-fold: readout noise of the CCD and imperfections of the probe beam.

We typically measure a background $I_{\text{bg}} \sim 550\,\text{counts}\,\text{px}^{-1}$, which gives a Poissonian shot-noise uncertainty: $\alpha_N = \sqrt{N} \sim 23.5\,\text{counts}\,\text{px}^{-1}$. Therefore, assuming that we can access the full dynamic range of the CCD and that a SNR of at least two is required, the maximum OD that can be measured is ~ 4.1. This is clearly much closer to the experimentally-measured value than if we neglect camera noise.

Despite having a linewidth less than that of the transition, the probe beam does not have an infinitely-narrow linewidth. As a consequence, some non-resonant component of the imaging light will always exist and will pass through the atoms without interacting, causing a drop in the measured OD. Furthermore, any drifts in magnetic field will cause variations in the resonant component of the imaging light because of Zeeman splitting.

It is possible to modify Eq. 3.2 to account for these OD saturation effects, arriving at the following:

$$OD_{\text{mod}} = \ln \frac{1 - e^{-OD_{\text{sat}}}}{e^{-OD_{\text{meas}}} - e^{-OD_{\text{sat}}}}, \tag{3.6}$$

where OD_{mod}, OD_{meas} and OD_{sat} are the modified, measured and saturation optical depths respectively [88,89]. It is important to note that this correction is only valid for probe intensities in the *weak-probe regime* [90].

3.2 Alignment of the imaging system

For atomic clouds with large densities, the refractive index deviates significantly from unity because the atomic susceptibility is non-zero. This has the effect of lensing the probe beam, distorting the image [75, 91–93]. For blue-detuned light, the refractive index becomes less than one, effectively

creating a diverging (concave) lens. The converse is true for red-detuned light, creating a converging (convex) lens.

This effect adds an extra layer of complexity and sensitivity to imaging, if lensing is to be avoided. Either the entire probe has to be perfectly resonant with the atomic transition of interest, or the optics need to be perfectly aligned with the atoms, which themselves should be narrowly distributed to remain within a Rayleigh length of the focus. To minimise the risk of unnecessary noise, we optimise both the frequency and alignment of the probe beam.

We focus both imaging systems using two different techniques: on-resonance imaging alignment and off-resonance imaging alignment. Both methods involve adjusting the position of the imaging objective lens using a translation stage.

3.2.1   On-resonance imaging alignment

We take resonant absorption images of a soliton in order to identify the lens position where the smallest width is measured. This occurs when the atoms are in the focal plane of the objective lens. It is important to use as small an object as possible for this measurement in order to minimise the proportion of atoms that fall outside the Rayleigh length of the imaging system and maximise precision.

An example of a focussing dataset is shown in Fig. 3.2. It also demonstrates that the OD increases towards the focus, which is to be expected because the same number of atoms are being imaged onto a smaller region of the CCD.

3.2.2   Off-resonance imaging alignment

The off-resonance technique exploits the unwanted lensing property itself. Firstly, a weak-probe detuning curve is performed to tune the probe laser onto

Figure 3.2: The measured $1/e^2$ radius of a soliton (blue circles) and the corresponding peak OD (red triangles) as the imaging objective lens is moved. The blue solid line and red dashed line are fits of a Lorentzian profile to the width and OD data respectively, in order to extract the optimal lens position.

resonance. This involves imaging the atoms with various probe frequencies, to locate the frequency that results in the highest OD. Once found, the probe is detuned above and below resonance to observe lensing effects (Fig. 3.3) [91]. By setting the lens position to minimise lensing effects across all detunings, we are able to verify the optimal lens position found using the on-resonance focussing technique. As we have already found the resonance frequency of the transition with the detuning curve, the laser frequency can be set for future imaging also.

3.2.3 Nulling stray fields for low-field imaging

Previous work on this experiment has used low-field (close to zero magnetic field) imaging exclusively [23, 33, 68, 70]. To do this effectively, it is necessary to apply a small ($\sim 2\,\mathrm{G}$) quantisation magnetic field along the direction of the probe propagation. This field is provided by one of three pairs of orthogonal magnetic field shim coils, depending on the propagation direction of the

Figure 3.3: Off-resonance absorption images for different objective lens positions. [a-f] and [g-h] are taken with a probe beam red-detuned by 4 MHz and blue-detuned by 4 MHz respectively, at objective lens positions 0 mm, 5 mm, 7.5 mm, 9 mm, 10 mm and 13 mm from left to right. All positions are relative to the position in image [a].

probe. The quantisation field lifts the degeneracy of the hyperfine states through Zeeman splitting and, as long as light of the correct polarisation is used to probe the atoms, selection rules prohibit the atoms from leaving the cycling transition through resonant atom-light interactions. Therefore, using a probe beam of high polarisation purity is essential to driving the correct transition. It is also important to maximise the purity of the quantisation axis, as any stray transverse magnetic fields will permit undesired transitions.

To null any stray transverse magnetic fields we first perform a detuning curve with a weak probe at a smaller quantisation field (~ 0.15 G) to precisely locate the central frequency of the transition. The reduced quantisation field amplifies the effect of any stray fields, improving the precision to which they can be nulled. We then vary the two pairs of shim coils transverse to the probe, in order to maximise the imaged OD (as shown in Fig. 3.4). We attribute this point to be where the field is correctly nulled, as loss to unwanted transitions is minimised. The calibration is repeated for both imaging directions.

We check the nulled-field values of the shim coils in the following ways. For the shared transverse shim (labelled N-S in our experiment) we directly compare the nulled value. For the other two directions (labelled E-W and

Figure 3.4: Nulling stray fields using three pairs of orthogonal magnetic field shim coils, labelled N-S, U-D and E-W. [a] and [b] ([c] and [d]) are taken using the horizontal (vertical) imaging system, nulling stray fields using the two transverse coil pairs.

U-D) we first calculate the required shim current to flip the direction of the field whilst maintaining the magnitude. This is done by calculating the difference between the value used for applying the quantisation axis and the nulled value determined above. A probe with polarisation of the opposite handedness is used to take a detuning curve at this flipped field and, if the field is well-nulled, the same peak frequency will be recovered (as in Fig. 3.5).
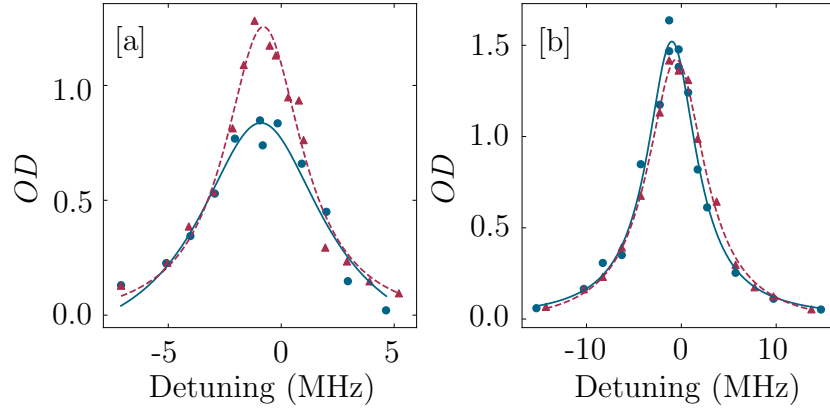
Figure 3.5: Detuning curves to determine the accuracy of the nulled field. [a] and [b] both show one curve (red triangles) where the quantisation axis is aligned with the probe propagation direction and one where it directly opposes it (blue circles), for the horizontal imaging system and vertical imaging system respectively. The red-dashed lines and blue-solid lines are Lorentzian fits to the co-propagating and counter-propagating data respectively, used to find the peak detuning frequency for comparison.

3.3   Beat-locked imaging laser

Our horizontal imaging system uses probe light taken from the cooling laser as we only require resonant, low-field, weak-probe imaging in this direction. However, the vertical imaging system uses a purpose built, extended-cavity diode laser (ECDL) along with an offset beat lock to image across a wide range of frequencies, allowing much greater flexibility.

The ECDL has a mode-hop-free tuning range of over 8 GHz. Standard temperature stabilisation via a negative feedback loop with a peltier and thermistor is used to minimise temperature-dependent frequency shifts [94]. Frequency tuning is provided by a piezo actuator mounted on the rear of the diffraction grating.

3.3.1   Beat note

There are an abundance of widely-used and long-established techniques for the frequency stabilisation of laser systems [95–99]. Each of these methods contain many similarities, almost all differing in one important aspect only: the form of spectroscopy used to interrogate an atomic sample. Essentially, all laser locks consist of three components: a reference source (almost always a sub-Doppler atomic transition feature) a laser that can be frequency modulated via an external input (usually by modulating the laser current or piezo voltage of an ECDL) and a way of comparing the laser and reference frequencies to produce an *error signal* used for feedback. Unfortunately, these methods all suffer similar detuning limitations as their error signals are based entirely on the frequencies of atomic transitions. Without more complex schemes that induce Stark or Zeeman shifts in the atomic sample, these frequencies are fixed. Therefore, they are only tunable to the extent of the optics that are used to modulate the laser output once the laser has been locked to an atomic transition. However, in the case of a beat-locked system, we are able to circumvent these limitations by employing two reference sources.

The first reference source is the master laser, to which we wish to lock our slave laser (the home-built ECDL). We can use either the cooling laser or the repump laser (see Ch. 2) as the master laser, depending on the required probe frequency. A flipper mirror is placed in the setup for ease of swapping between the two. We beat the master and slave lasers on a photodiode to extract the beat note and compare this frequency with another fixed frequency source, the second reference. The difference between the beat note and the second

reference generates the error signal, which is fed back to the slave laser. By controlling the ratio between the second reference and the beat note we can lock to any arbitrary frequency within the range of the slave laser.

Light emitted from the slave laser enters a pair of anamorphic prisms that reduce the ellipticity of the beam (nominally to zero eccentricity) and an optical isolator that prevents back-reflections into the laser, vital for laser stability. A portion of the light is then reflected off a polarising beam splitter (PBS) to establish a beat note. Before the light from the slave laser enters the fiber-coupled photodiode (*Vescent Photonics* D2-160) it is mixed with light from the master laser on a non-polarising 50:50 beam splitter. These two lasers produce a beat note on the photodiode at a frequency:

$$f_{\text{beat}} \quad = \quad |f_{\text{slave}} - f_{\text{master}}|. \qquad (3.7)$$

Beyond the optics described above, a double-passed acousto-optic-modulator (AOM) is used to allow sub-microsecond switching of the probe beam. This is important for performing multiple imaging pulses in quick succession, which will be relevant in Sec. 3.8. A shutter is included to prevent any light leakage through the AOM from reaching the experiment. However, as it has a much slower switching time than an AOM, it is left open for the duration of consecutive probes.

3.3.2 Locking electronics

The beat-note photodiode outputs a digitised signal at the beat frequency between the two lasers. In order to lock a slave laser, one must derive an error signal that is sensitive to drifts in frequency and feed it back to the laser, either via the diode current or piezo voltage. We use a frequency discriminator chip (*Analogue Devices* ADF4107) mounted on an evaluation board (*Analogue Devices* ADF411XSD1Z) along with a USB-interfaced controller board

(*Analogue Devices* SDP-S). The beat note is fed to the ADF4107 where the frequency discriminator divides the beat frequency by a factor N and the on-board 10 MHz reference signal by a factor R. The ADF4107 allows comparison of frequencies for any beat frequency up to 7 GHz, though this could be extended still further by adding a pre-divider to the beat signal, if required [99].

This setup is capable of providing both a frequency and phase lock. For small phase differences, the frequency discriminator outputs a signal proportional to the phase difference between the two divided frequencies. For frequency differences and larger phase differences, the chip outputs either the rail voltage or 0 V, depending on the sign of the frequency difference between the two. Therefore, the effective *capture range* of the lock is entirely limited by the range that the slave laser can scan without a mode-hop, assuming that the bandwidth of the beat note photodiode is sufficient and the gain of the feedback loop is sufficiently high.

To use the output signal of the frequency discriminator as feedback for the slave laser, we first run it through a proportional-integral-derivative (PID) controller with variable gain. The variable gain gives us control over the feedback magnitude, which is essential when optimising the stability of the lock (see below). The output of the PID controller is the error signal, which is passed directly into either the piezo driver or current driver as a modulating signal, completing the feedback loop.

The everyday running of the beat lock is relatively straightforward. First, the piezo of the slave laser is scanned and the Doppler-broadened spectroscopy of a Rb vapour cell is observed by a photodiode. We offset the piezo voltage such that the centre of the scan is approximately at the desired locking frequency, before turning off the scan. By monitoring the N-divided beat frequency we can confirm when the laser is locked, as the N-divided frequency will be stable and equal to the on-board reference frequency divided by R. It is important to be aware that the slave laser frequency will be forced away from the desired locking frequency if the polarity of the error signal is incorrect relative to

the sign of $f_{slave} - f_{master}$. This is a particularly important consideration with low-frequency beat notes as the small frequency difference is difficult to distinguish in Doppler-broadened spectroscopy.

3.3.3 Stability

There are two main aspects to consider when trying to optimise the stability of the beat lock. Firstly, we need to ensure that the beat note itself is stable against long-term drifts, relative to the master laser, so that the output frequency of the laser does not vary throughout an experimental run. Drifts could be caused by temperature changes altering the response of the beat note photodiode, for example.

Figure 3.6: These graphs show the spectra of the beat note when the gain of the PID controller is [a] well optimised and [b] too high. [c]-[h] show the drift in the beat note when the gain is set too low. There is a 5 s interval between each graph, from [c] to [h] respectively.

By monitoring the beat note on a spectrum analyser, the lock is found to be completely robust against long-term drift throughout full-day experimental

runs. In fact, it only unlocks if the master laser itself is unlocked or the N divider value is set to an extreme, such that it pushes the lock frequency beyond the mode-hop-free tuning range of the slave laser.

Secondly, we need to ensure that the gain of the PID controller is set correctly, which we show in Fig. 3.6. If the gain is set too low, any drift in the slave laser would not be sufficiently compensated for and so the laser frequency will continue to drift. Conversely an overly-large gain will cause any small drift in the slave laser output frequency to be over-compensated for and make the output frequency oscillate between being above and below the desired frequency: so-called *ringing*. Therefore, the gain is optimised to make certain that the laser remains emitting at the correct frequency, without introducing any high-frequency instability.

3.4   Low-field imaging

In this section we present the spectroscopy of ^{85}Rb at low magnetic fields, from different ground states. These spectra ensure that the imaging system is behaving correctly, before the possibility of imaging the atoms *in-situ* at high magnetic fields is investigated.

A cloud of 8×10^5 ^{85}Rb atoms at $2.3\,\mu$K in the $|F, m_\mathrm{F}\rangle = |2, -2\rangle$ state are released from the crossed-dipole trap and allowed to expand freely for $2\,$ms. They are then imaged by the vertical imaging system with an $80\,\mu$s probe pulse, using one of the following imaging techniques. A background image is taken after a further $800\,$ms, followed by a dark image $800\,$ms later. This delay is due to the finite read-out time of the camera and imaging program. The optical depth of the cloud is then calculated using Eq. 3.2.

Imaging from $F = 3$

The common approach to absorption imaging in ^{85}Rb is to use the $F = 3 \rightarrow F' = 4$ transition for the imaging transition, with a repump resonant with

the $F = 2 \rightarrow F' = 3$ transition. By repeating the above sequence multiple times and using different imaging laser detunings (beat frequency settings) we build up a detuning curve for the $F = 3 \rightarrow F' = 4$ transition (Fig. 3.7). This allows confirmation of the required beat frequency for setting the probe on the central frequency of the transition.

Figure 3.7: Low magnetic field detuning curve for the $F = 3 \rightarrow F' = 4$ transition to determine the required beat frequency for setting the probe on the central frequency of the transition.

In this case the vertical imaging laser is locked to a master laser, which is itself locked $-230\,\text{MHz}$ from the $F = 3 \rightarrow F' = 4$ transition. As the double-passed AOM is set to $-110\,\text{MHz}$, the slave laser will be $220\,\text{MHz}$ higher than the frequency reaching the experiment. Eq. 3.7 tells us that a beat frequency of $450\,\text{MHz}$ is required to produce light resonant with the $F = 3 \rightarrow F' = 4$ transition. In Fig. 3.7 we measure this value to be $(450.1 \pm 0.1)\,\text{MHz}$, in good agreement. We also determine the width of the detuning curve to be $\Gamma = (8.4 \pm 0.5)\,\text{MHz}$, slightly larger than the $\Gamma = 6.07\,\text{MHz}$ in the literature [93].

Imaging from $|2, -2\rangle$

It is also possible to remove the repump and instead image directly from the trapped $|2, -2\rangle$ state. However, the selection rules for electric dipole transitions in the coupled basis are $\Delta F = \pm 1, 0$. Therefore, if we image on any of the $F = 2 \rightarrow F' = 2, 3$ transitions, the atoms are quickly lost into the $F = 3$ state, whereupon they become *dark* to the imaging laser and are no longer measured (Fig. 3.8). Furthermore, imaging on any $F = 2 \rightarrow F' = 1$ transitions will result in population being pumped into a dark $F = 2$ state, with the particular dark hyperfine state/states depending on the polarisation of the probe beam.

Figure 3.8: Zero magnetic field detuning curve across the $|2, -2\rangle \rightarrow F' = 1, 2, 3$ transitions of ^{85}Rb (from left to right). The blue round points are taken using the imaging laser to pump atoms into the $F = 3$ state for imaging on the $F = 3 \rightarrow F' = 4$ transition with the horizontal imaging system whereas the red triangle data points are taken by direct imaging. The solid blue and dashed red lines are double Lorentzian fits to the blue and red data respectively.

Therefore, to create detuning curves for the $F = 2 \rightarrow F' = 2, 3$ transitions, the proportion of atoms that are lost into the $F = 3$ ground state is measured instead. The vertical imaging system is used to promote atoms into the

dark state before the horizontal imaging system takes the final image on the $F = 3 \rightarrow F' = 4$ transition. By varying the detuning of the vertical imaging system, the fraction of atoms transferred to $F = 3$ is changed. However, in this configuration the $F = 2 \rightarrow F' = 1$ transition becomes effectively dark, hence no peak is observed (Fig. 3.8). Therefore, it is necessary to measure both directly and indirectly to map out the $F = 2 \rightarrow F' = 1, 2, 3$ transitions completely (Fig. 3.8).

3.5 High-field imaging

In this thesis, ^{85}Rb solitons are produced using the Feshbach resonance at $\sim 155\,\mathrm{G}$, which has a zero crossing at around $165\,\mathrm{G}$ [19, 100–102]. As we are largely interested in imaging soliton dynamics, it is essential that solitons can be reliably imaged *in-situ* at magnetic fields close to the zero crossing.

Unlike for the zero-field case described in Sec. 3.4, where the selection rules are easily described in the coupled basis $|F, m_{\mathrm{F}}\rangle$, the introduction of a magnetic field forces the energy states to undergo Zeeman splitting. This removes the degeneracy of the states and complicates selection rules. Note that for large magnetic fields (higher than $\sim 1\,\mathrm{T}$) the splitting can be so great that the states enter the *hyperfine Paschen-Back regime*, whereupon the preferred basis is the uncoupled $|J, m_{\mathrm{J}}\rangle$ basis [93]. These states have very similar selection rules to the zero-field case, so understanding a spectroscopy signal in this regime is similarly straightforward.

Unfortunately, for intermediate fields, such as those relevant here for soliton production, we find that the preferred basis for the excited states is the uncoupled $|J, m_{\mathrm{J}}\rangle$ basis, whereas the preferred basis for the ground state remains the coupled $|F, m_F\rangle$ basis. Determining a set of selection rules for this regime becomes impossible, so instead we again use spectroscopy to characterise the allowed transitions at $(165.71 \pm 0.02)\,\mathrm{G}$ experimentally.

To minimise the fraction of atoms that can be lost to dark states, in Fig. 3.9 we set the probe pulse to $2.5\,\mu\mathrm{s}$ (considerably shorter than the $80\,\mu\mathrm{s}$ pulse

Figure 3.9: Spectroscopy of ^{85}Rb atoms in the $|F, m_{\mathrm{F}}\rangle = |2, -2\rangle$ state at $(165.71 \pm 0.02)\,\mathrm{G}$, interrogated by $780\,\mathrm{nm}$, circularly-polarised light that drives σ^+ and σ^- transitions (blue and red respectively). The solid blue and dashed red lines are double Lorentzian fits to the blue and red data respectively.

used in Sec. 3.4) and perform direct spectroscopy on 8×10^5 atoms. As the energy levels of the excited states are well separated at this field, σ^+ and σ^- transitions are easily distinguishable by using circularly-polarised light of opposing handedness.

Despite being able to image large clouds at high field, the peak optical depths in Fig. 3.9 are far lower than that of a comparable zero-field, closed-transition image (see Fig. 3.7). This, along with the reduced probe duration, leads to much greater noise levels in images, particularly for small objects. Therefore, it is worth considering the possible transitions from each ground state and identifying those with the most favourable transition strengths to excited states.

We determine the transition strengths for the relevant transitions, probe polarisations and magnetic field using the electric susceptibility calculator *ElecSus* [103]. The transition strengths are shown in Fig. 3.10, along with the detuning of each transition from the zero-field frequency and the branching ratios for the ground states from each accessible excited state.

3.5.1 Branching ratios

Using the transition strengths calculated above, it is possible to determine the number of photons that an atomic cloud can scatter before the atoms are lost to a dark state. Normalising the transitions strengths A from a particular excited state to each of the states in the ground state manifold ($\sum_g A_{e \to g} = 1$) returns the decay branching ratios. These are represented in Fig. 3.10 by the size of the data points in the ground states for each of the allowed excited states.

Consider an idealised three-level system (Fig. 3.11) with two states on which to image and one dark state ($|g\rangle$, $|e\rangle$ and $|d\rangle$ respectively). When an atom absorbs a photon and is excited from $|g\rangle$ to $|e\rangle$ it has a probability L of being lost to $|d\rangle$. The remaining $1 - L$ will decay back to $|g\rangle$. When a second
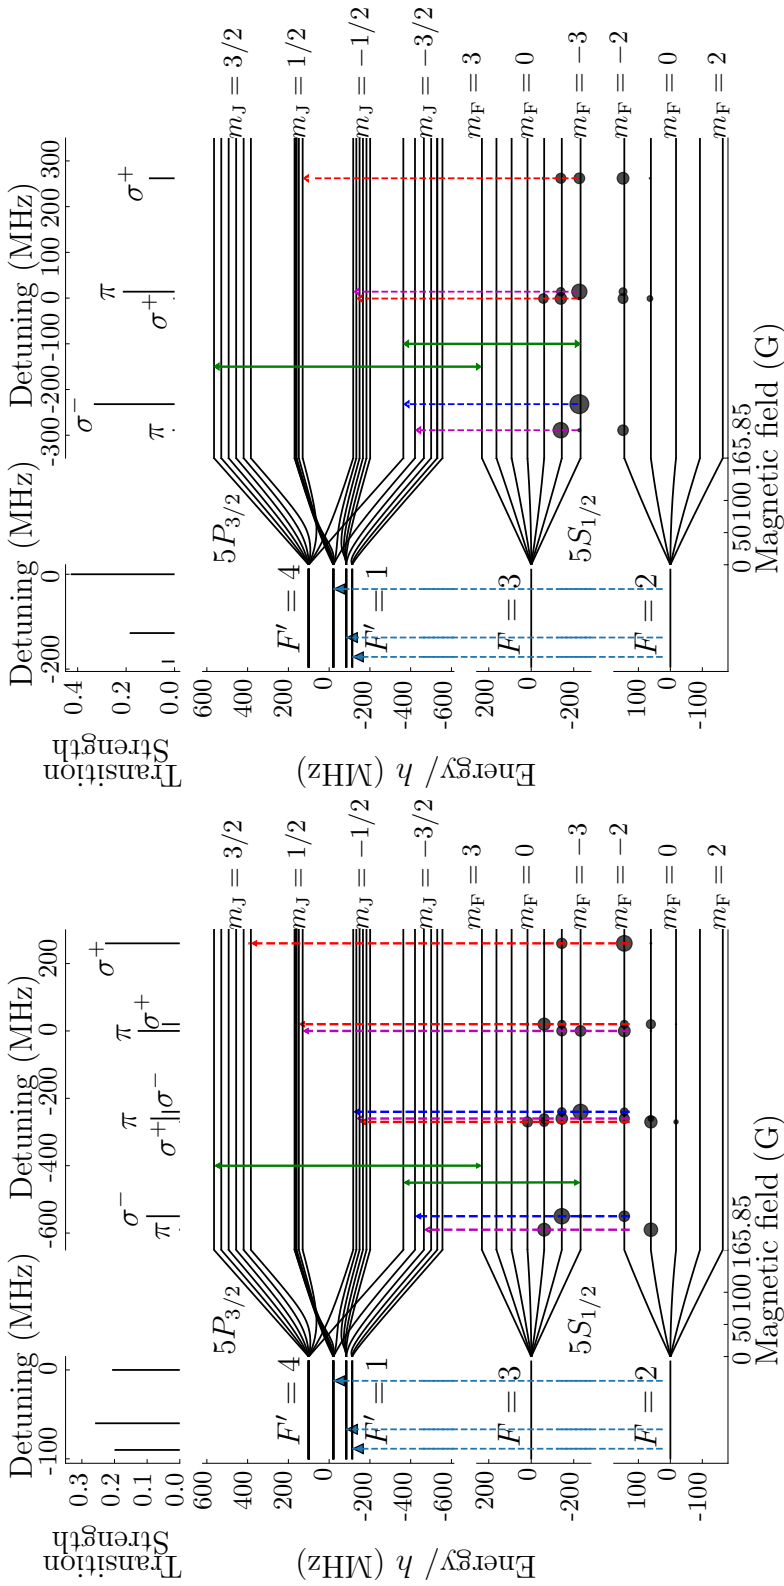
Figure 3.10: Left: transitions from the $|F, m_F\rangle = |2, -2\rangle$ state. Right: transitions from the $|F, m_F\rangle = |3, -3\rangle$ state. The allowed zero-field transitions are shown, along with those at 165.85 G. Red, blue and pink arrows in the high-field regime indicate σ^+, σ^- and π transitions respectively. The intermediate regime shows the energy level splitting as the field is increased. The solid green arrows indicate closed transitions and the size of the points in the ground state manifold indicate the branching ratio to that state from the particular excited state.

Figure 3.11: A simple 3-level system to describe an imaging transition $|g\rangle \rightarrow |e\rangle$ with a branching ratio L to a dark state $|d\rangle$.

photon is absorbed a further $(1 - L) \times L$ of the total population will decay to $|d\rangle$. On the n^{th} iteration of this cycle a further $(1 - L)^{n-1} \times L$ will be lost to $|d\rangle$. Therefore, the average number of photons each atom will scatter before being lost to $|d\rangle$ is given by:

$$\bar{N}_\gamma = \sum_{n=1}^{\infty} n(1 - L)^{n-1}L = \frac{1}{L}. \tag{3.8}$$

We trap our atoms in the state $|F, m_F\rangle = |2, -2\rangle$ and, as can be seen in Fig. 3.10, the *least open* transition is a σ^+ transition to the $|F', m_F\rangle = |4, -1\rangle$ low-field basis state (located in the $m_J = 3/2$ excited state manifold). The branching ratio from this excited state to dark states (states other than $|2, -2\rangle$) is 31.32%. Therefore, if used as an imaging transition at this field, an atom would only scatter an average of 3.19 photons before being lost. See Appendix A for tables of frequencies and transition strengths of all transitions with non-negligible transition strengths (> 0.0005).

However, is it possible to identify two completely closed transitions from the $F = 3$ ground state manifold: $|3, -3\rangle \rightarrow |4, -4\rangle$ and $|3, 3\rangle \rightarrow |4, 4\rangle$ (see the green arrows in Fig. 3.10). Our naive picture would imply that we should be able to scatter an infinite number of photons on this transition, giving us an infinite signal-to-noise ratio (SNR). However, our model clearly ignores

off-resonant scattering on other transitions, as well as atom heating and loss. These factors limit the number of photons that can be scattered, even for a theoretically closed transition.

Despite its limitations, this picture provides a good approximation of the relative absorption signals we should expect when imaging on transitions that are far-detuned from neighbouring transitions. As this is the case for the two ground states considered above, there should be a significant improvement in SNR if the atoms are efficiently transferred to the $|3, -3\rangle$ ground state before imaging.

3.6   Microwave transfer imaging

To transfer the atoms from the trapped $|2, -2\rangle$ state to the required $|3, -3\rangle$ imaging state we use microwaves, as the $|2, -2\rangle \rightarrow |3, -3\rangle$ transition is $\sim$ 2.66 GHz at 165 G. We use a home-built microwave antenna (Fig. 3.12), based on those used at JILA [104]. It is positioned directly below the science cell, $\sim$ 1 cm away from the trapped atoms. The hole is aligned directly underneath the atoms, maximising the circular polarisation output which is required for the σ^- microwave transition. This arrangement also allows the probe beam of the vertical imaging system to pass through the hole.

The antenna stub is tuned to provide maximum emission efficiency around the $|2, -2\rangle \rightarrow |3, -3\rangle$ transition frequency. This is achieved by varying the stub's position and length and observing the electrical power reflected by the antenna whilst the frequency is varied (Fig. 3.13).

3.6.1   Stern-Gerlach separation

To determine the state transfer efficiency we image the population in each of the two states ($|2, -2\rangle$ and $|3, -3\rangle$) simultaneously. This is achieved through

Figure 3.12: The microwave antenna design, showing views of both the top and bottom. The orange areas show the copper traces what are positioned on both sides of a thin phenolic substrate. The black circle is a circular aperture which allows the probe beam to pass through and is aligned to sit below the atoms. The antenna stub's length and position is tuned to adjust the peak emission wavelength.

Figure 3.13: The reflectance spectra of our microwave antenna across the frequency range of interest. The vertical dotted line highlights the required 2.66 GHz frequency of the microwave transition.

Stern-Gerlach separation, whereby two states are spatially-separated by a magnetic field gradient.

As described in Ch. 2, the magnetic quadrupole trap produces a field gradient of 20.85 G during the later parts of the experimental sequence, which almost cancels gravity for atoms in the $|2, -2\rangle$ state as they are low-field-seeking (Eq. 2.1). The crossed optical dipole trap accounts for the remainder of the levitating potential. However, atoms transferred to the $|3, -3\rangle$ state by the microwaves experience a downwards force of $\sim 2.5mg$, as they have become high-field-seeking. The additional $0.5mg$ arises from the change of m_F. The optical dipole trap is now insufficient to levitate these atoms and so they spill out of the trap. After a 4 ms hold to allow the two states to spatially separate by $\sim 80\,\mu$m, both clouds are imaged simultaneously using the horizontal imaging system. The atom number in each cloud is extracted simultaneously from the image, using the methods outlined in Sec. 3.1.

3.6.2   Rabi oscillations

In Fig. 3.14 we measure Rabi oscillations to identify the pulse duration that produces the most efficient state transfer. To achieve a high transfer efficiency, the Rabi frequency needs to be large relative to any sources of decoherence. Unfortunately, we are limited by microwave power and the finite size of the cloud, which leads to dephasing as the Rabi frequency varies across it. We find a maximum transfer efficiency of $\sim 60\,\%$ for a probe duration of $\sim 15\,\mu$s. This is far too low to be used in a robust imaging sequence, as it creates similar limitations to imaging directly from the $|2, -2\rangle$ state. Fortunately, the efficiency of the microwave state transfer can be vastly improved by performing adiabatic rapid passage between the two states.

3.6.3   Adiabatic rapid passage

Adiabatic rapid passage (ARP) can be most easily understood in the *dressed-state* picture of a two-level system. The energies of the pair of levels, coupled

Figure 3.14: A measurement of the state transfer efficiency as a function of microwave pulse duration. Some evidence of Rabi oscillations is visible, though it is heavily suppressed.

by an electric dipole interaction of Rabi frequency Ω, are given by:

$$E_\pm = \pm(\hbar/2)\sqrt{\Delta^2 + \Omega^2}, \tag{3.9}$$

where Δ is the detuning from resonance of the coupling field [105].

In the limit of zero coupling field, the two eigenenergy surfaces are degenerate at zero detuning and the energy levels are able to cross. In this limit a state transfer would not occur and the atom would simply remain in the ground state and follow the eigenergy (see Fig. 3.15). For any non-zero coupling field the energy levels are subject to the coupling perturbation Ω and the level crossing is avoided (an *avoided crossing*). This allows a coherent state transfer by adiabatically ramping the detuning across the resonance (adiabaticity is satisfied for $\Omega^2 \gg |\frac{d\Delta}{dt}|$). Fig. 3.15 shows how the avoided crossing energy difference increases with a stronger coupling field.

Using the Stern-Gerlach technique described earlier, we can optimise our ARP sweeps. The microwave frequency is tuned around the expected resonance to maximise efficiency, producing a microwave detuning curve. From

Figure 3.15: An example of an avoided crossing between a ground state (red) and an excited state (green) caused by an inter-state coupling field Ω as the frequency Δ is swept across resonance. Ω, Δ and E are all shown in arbitrary units. The dashed arrows highlight the limit of zero coupling, where state transfer does not occur, and the case of non-zero coupling, where an avoided crossing exists.

this we determine the transition line centre. To optimise the ARP, we also
take a series of symmetric sweeps centred on the resonance position for var-
ious sweep amplitudes, durations and powers (Fig. 3.16).

Figure 3.16: The left figure shows the ARP efficiency as a function of
sweep amplitude for various sweep durations: red, blue, green, pink and
black are $200\,\mu s$, $100\,\mu s$, $50\,\mu s$, $40\,\mu s$ and $15\,\mu s$ sweeps respectively. The right
figure shows the ARP efficiency as a function of sweep amplitude, for various
microwave powers: black, blue and red represent $13\,dBm$, $10\,dBm$ and $7\,dBm$
out of our microwave source respectively. The sweep duration is kept constant
at $100\,\mu s$ for the data. The lines in both figures are only intended as guides
to the eye.

Fig. 3.16 shows that slower scans at higher microwave powers produce a
higher efficiency. This is to be expected as slower sweeps improve the adi-
abaticity of the state transfer and higher coupling powers split the energy
levels more and so reduces the curvature of the avoided crossing, again im-
proving adiabaticity.

For larger frequency sweeps over a fixed sweep time, we observe a drop in
efficiency because the frequency gradient increases, reducing the adiabaticity.
The reason behind the rapid reduction in efficiency for small frequency am-
plitudes is slightly more subtle. When the sweep begins at a large detuning
the energy levels are well separated, ensuring atoms are initially addressed
in a pure eigenstate. However, for small frequency sweeps the atoms begin in

an admixture of the ground and excited eigenstates, due to the finite chance of microwave photon absorption at the initial detuning. This means that the atoms remain in an admixture at the end of an ARP sweep that is symmetric about the microwave transition line centre, preventing 100 % transfer efficiency. Therefore the optimum parameters are found to be a balance between these regimes, whilst keeping the ARP sweep as short as possible to minimise cloud expansion. This is especially important for imaging solitons, as the change of internal state will cause a rapid quench of the interatomic interactions. We routinely use a 100 µs, 300 kHz sweep at maximum source power to achieve $\sim 90\,\%$ transfer efficiency.

3.7   Partial-transfer imaging

The attraction of taking multiple weak measurements of the same atomic sample is clear: cold atom experiments are typically very slow because of the time it takes to cool the atoms, so getting more than one data point from a sequence is highly appealing. *Non-destructive* imaging, or more accurately *partially-destructive* imaging, has been achieved in a variety of ways, through both dispersive [34–37] and resonant imaging techniques [39].

Our microwave transfer technique lends itself to partial-transfer imaging, whereby a small fraction of the atoms are transferred to the intermediate $|3,-3\rangle$ state and imaged, leaving the remainder of the population virtually unperturbed. Fig. 3.17 shows a calibration of the number of atoms in each state, as a function of the ARP sweep duration. This can be used to determine the number of atoms that have to be removed from the trapped $|2,-2\rangle$ state in order to achieve the required SNR in the images of the $|3,-3\rangle$ state. Partial-transfer imaging is not used in this thesis because of the undetermined impact it may have on soliton phase. However, it may provide a solution for future studies of soliton dynamics.

Figure 3.17: Atoms in the $|2, -2\rangle$ state (red triangles) are partially transferred to the $|3, -3\rangle$ state (blue circles) by microwave adiabatic rapid passage of various sweep durations. The frequency sweep amplitude ($300\,\text{kHz}$) and microwave power are kept constant.

3.8   High-intensity imaging

For high OD clouds such as BECs, the probe beam can be attenuated almost completely as it passes through. If the OD is high enough, it can become impossible to accurately apply Eq. 3.6, as the Poissonian noise in I_{atoms} and I_{bg} dominates the modified OD. Fortunately, this can be overcome by increasing the probe beam intensity.

Consider a two-level system in the presence of a saturation effect, Beer-Lambert's law for resonant light can be written as:

$$\frac{dI}{dz} = -n\frac{\sigma_0}{\alpha}\frac{1}{1 + I/I_{sat}}I, \tag{3.10}$$

where $\sigma_0 = 3\lambda^2/2\pi$ is the resonant cross-section, n is the atomic density, α accounts for non-perfect polarisation or magnetic field orientation and $I_{\text{sat}} = 1.16\,\text{mW}\,\text{cm}^{-2}$ for the ^{85}Rb cycling transition [93].

Eq. 3.10 shows that the fraction of absorbed probe light can be modified by increasing the probe intensity, which will in turn reduce the fractional noise in the measured OD. However, because more light is transmitted through

the atomic sample, the OD will be artificially reduced. To compensate for this saturation effect, a more complete equation for the optical depth which includes the effects of the probe beam intensity must be used:

$$OD_{\text{act}} = OD_{\text{mod}} + \left(1 - e^{-OD_{\text{mod}}}\right) \frac{I}{I_{\text{sat}}^{eff}},\qquad(3.11)$$

where OD_{mod} is that given by Eq. 3.6 and I_{sat}^{eff} is the *effective* saturation intensity in the units of the measurement of I (camera counts in our case).

If we re-cast Eq. 3.11 in terms of camera counts, in the limit of infinite dynamic range and perfectly-resonant light ($OD_{\text{sat}} = \infty$) we reach:

$$OD_{\text{act}} = \ln\left(\frac{C_{\text{probe}} - C_{\text{bg}}}{C_{\text{atoms}} - C_{\text{bg}}}\right) + \frac{C_{\text{probe}} - C_{\text{atoms}}}{C_{\text{sat}}^{eff}},\qquad(3.12)$$

The linear term in Eq. 3.12 becomes increasingly relevant for the higher probe intensities discussed here as $C_{\text{atoms}} \sim C_{\text{probe}}$ when $C_{\text{probe}} \gg C_{\text{sat}}$, for all optical depths relevant to us.

The camera counts (C) can be calculated by:

$$C = \frac{I A_{\text{pix}} \lambda}{hcM^2} \times T \times QE \times \tau,\qquad(3.13)$$

where A_{pix} is the camera pixel area, M is the system magnification, T is the transmission through the imaging system, QE is the quantum efficiency and τ is the probe duration. An additional factor should be included if the camera is not calibrated to give one count per photo-electron.

3.8.1   High-intensity imaging calibration

Experimentally, the imaging system must be calibrated to take the effects of the high probe intensity into account. There are currently three main methods that can be used to do this.

Method I

The first method is to simply calculate the intensity of the probe beam at the atoms, taking into account losses at optical surfaces, using a beam power calibration taken outside of the vacuum chamber. However, this method relies on knowing many experimental parameters to high precision, namely: the transmission coefficients for every optic in the probe beam path after the point at which the power calibration was taken, the quantum efficiency of the camera, the magnification of the imaging system, the camera pixel area and the width of the probe beam at the atoms. It is also vital to have an accurate method of measuring the probe beam power, which itself is a challenge when using semiconductor-based or thermopile-based power meters [106]. Instead, using the images themselves to determine the intensity of the probe at the atoms on a shot-by-shot basis is a much more desirable technique.

Method II

The second method for calibrating the OD is to take multiple images of an identical thermal cloud for several different probe intensities, as shown in Fig. 3.18. As expected, the observation of a lower OD corresponds with an increased probe intensity. By fitting Eq. 3.11 to the data in Fig. 3.18, with I and OD_{actual} as fitting parameters, we can extract a measurement of I_{sat}^{eff}. This can be used to calibrate the optical depths of any cloud at any probe intensity.

However, there are two downsides to using this method. Firstly, this method requires that OD_{act} is stable from shot-to-shot, which effectively means that the atom number must be stable. Secondly, atoms can become Doppler-shifted out of resonance if they are exposed to the probe for too long, due to momentum transfer from the probe beam photons to the atoms. Consequently probe durations are kept short ($< 5\,\mu$s) when this method is used.

Figure 3.18: The camera counts per pixel from the region of interest of the *probe* image, as a function of the measured optical depth when the probe intensity is varied.

Method III

The final method relies directly on this momentum transfer [107]. The atoms are released from the trap and a very short ($2\,\mu s$) probe pulse is applied. After $2\,ms$ of free time of flight an image of the atoms is taken in the transverse direction. The momentum transferred to each atom is directly proportional to the number of scattered photons (N_γ),

$$N_\gamma = \frac{\Gamma}{2}\frac{s_0}{1+s_0}\tau, \qquad (3.14)$$

where Γ is the natural linewidth of the imaging transition, τ is the probe duration and:

$$s_0 = \frac{I}{I_{\text{sat}}} = \frac{C}{C_{\text{sat}}}. \qquad (3.15)$$

Compared with methods I and II, this method has an important advantage: it requires neither a stable atom number nor a power measurement of the imaging beam as we are only interested in the flight distance of the atoms.

To identify the saturation intensity count rate per pixel the following equation is fitted to the final position of the atom cloud as a function of camera counts (Fig. 3.19).

$$f\left(C\right) = \eta N_\gamma C = \eta\frac{\Gamma}{2}\frac{1}{1+C_{\text{sat}}/C}, \qquad (3.16)$$

where η is a proportionality constant that converts the number of scattered photons into a position change. C_{sat} and η are the fitting parameters.

Figure 3.19: Distance traveled by an atomic cloud after an impulse provided by an resonant imaging beam pulse. We find a saturation intensity count rate per pixel of $(240 \pm 30)\,\mu\text{s}^{-1}\text{px}^{-1}$ for the vertical imaging system, using the 30 mm objective lens.

As the probe is 2 μs long, the C_{sat} measured here gives the saturation counts per pixel for a 2 μs probe pulse. To retrieve the saturation intensity count rate per pixel we simply divide by two. This method is used in all imaging hereafter to compensate for any saturation effects. A typical high-intensity imaging sequence involves a probe duration of less than 10 μs.

3.8.2 Fast kinetics

Due to the increased count rate of high-intensity imaging, images are more sensitive to mechanical vibrations. This is particularly true when waiting 800 ms between each pulse, as was the standard in low-intensity imaging (Sec. 3.4). However, this wait time was dictated by the time it takes to read-out the full CCD, so reducing it is not possible in this configuration. To overcome this, we make use of the *fast kinetics* mode that is available on the *Andor LUCA* camera.

In *fast kinetics* mode a home-built mask is placed in front of the CCD array, so that only some rows will be exposed once the probe beam is pulsed on and the camera is triggered. These rows are aligned to be those that contain the image of the atoms. Instead of reading out the full CCD between each image, the CCD is clocked vertically by a set number of rows (with a clock speed of $0.9\,\mu s\,px^{-1}$) so that the image goes behind the mask and is no longer being exposed. The full CCD is read out once all the images have been taken.

We typically use a mask that exposes ~ 250 pixels to the probe beam and allow $300\,\mu s$ between images for the CCD array to be shifted. Whilst shorter hold times could be possible (a ~ 250 pixel shift should only require a $225\,\mu s$ hold) we find that $300\,\mu s$ is sufficient to prevent fringes in the images from mechanical vibrations, which are typically $< 1\,kHz$ in laboratory environments.

3.9   Imaging resolution

The resolution limit of the vertical imaging system is expected to be of the order of the soliton width. Therefore, the observed soliton in the image plane is the convolution of the soliton in the object plane with the point spread function of the imaging system. Thus, the observed width $k_{N,observed}$ is the true soliton width k_N summed in quadrature with the resolution r of the imaging system:

$$k_{N,observed} = \sqrt{k_N^2 + r^2}. \tag{3.17}$$

We can estimate the resolution limit of the imaging system by imaging the soliton before and after 50:50 splitting as, according to 1D-GPE theory, the soliton width should double when the atom number is halved for a constant interaction strength. Mathematically, this implies:

$$k_{N/2,observed} = \sqrt{k_{N/2}^2 + r^2} = \sqrt{4k_N^2 + r^2}, \tag{3.18}$$

where $k_{\mathrm{N}/2,observed}$ and $k_{\mathrm{N}/2}$ are the observed and actual widths of the split soliton respectively. Rearranged, Eqs. 3.17 and 3.18 give the resolution of the imaging system to be:

$$r = \sqrt{\frac{4k_{\mathrm{N}}^2 - k_{\mathrm{N}/2}^2}{3}}. \tag{3.19}$$

As mentioned at the beginning of this chapter, we have used two different objective lenses throughout this thesis, each providing a different magnification. For the $f = 30\,\mathrm{mm}$ lens, we measure a soliton width of $(3.6 \pm 0.2)\,\mathrm{\mu m}$ and a split soliton width of $(4.9 \pm 0.3)\,\mathrm{\mu m}$. This implies a resolution limit of $(3.0 \pm 0.3)\,\mathrm{\mu m}$ and a true un-split soliton width of $(2.0 \pm 0.5)\,\mathrm{\mu m}$. For the $f = 100\,\mathrm{mm}$ lens we do not measure a difference in width following splitting, which implies that the resolution limit is larger than the soliton width in this case. Therefore, we estimate the resolution limit to be $\sim 10\,\mathrm{\mu m}$, determined by the ratio in magnification between the two lenses.

Chapter 4

Producing optimised solitons

In this chapter we detail the steps we take to create stable solitons of ^{85}Rb BEC with unprecedented lifetimes. We describe two methods of field calibration that are used to precisely determine the s-wave scattering length, before demonstrating how we optimise the interatomic interactions of a soliton through a systematic study of BEC width oscillations. We then present observations of solitons oscillating in a weak harmonic potential for longer than 20 s without measurable dispersion.

4.1   Magnetic field calibration

As detailed in Sec. 2.3, precise control over the magnetic field is of central importance for efficient BEC production of ^{85}Rb. However, it is especially important for stable and reliable soliton production as inter-atomic interactions, governed by the s-wave scattering length, depend strongly and non-linearly on magnetic field via Feshbach resonances [19, 100–102].

We calibrate the magnetic field at the atoms using two methods: RF field calibration and microwave-transfer field calibration. Both of these methods rely on having a suitably strong magnetic field gradient so as to remove magnetically anti-trapped atoms from any optical dipole trap.

4.1.1 RF field calibration

RF field calibrations work by coupling the trapped $|2, -2\rangle$ state with anti-trapped states using RF radiation, in a similar fashion to the RF evaporation detailed in Sec. 2.3.4. A 50 ms RF pulse is applied to the atoms, followed by a 200 ms hold. If the RF frequency is resonant with a $\Delta m_{\mathrm{F}} = 1$ transition, atoms will be pumped into the $|2, 1\rangle$ and $|2, 2\rangle$ un-trapped states and lost. As the $\Delta m_{\mathrm{F}} = 1$ transition frequency depends on Zeeman splitting, we can use this loss feature as a measure of the local magnetic field at the atoms.

Initially, 100 kHz frequency sweep modulations are performed during the 50 ms pulse for coarse calibration, as this promotes loss for any transition with a $\Delta m_{\mathrm{F}} = 1$ transition frequency in this range. Fixed-frequency pulses are used for higher-precision spectroscopy, once the approximate RF transition frequency has been located. This method allows us to either calibrate an unknown magnetic field (Fig. 4.1[a]) or to set the coil current to achieve a required magnetic field (Fig. 4.1[b]).

Figure 4.1: RF magnetic field calibrations. [a] Atoms in an unknown magnetic field are probed with RF radiation which, when resonant with a $\Delta m_{\mathrm{F}} = 1$ transition at that magnetic field, causes atom loss from the trap. This is used to determine the magnetic field at the atoms (165.89 G in this case). Conversely, in [b] the RF frequency is set to be resonant with the $\Delta m_{\mathrm{F}} = 1$ transitions at 165.75 G ($a_{\mathrm{s}} = 0$) and the bias field coil current is varied instead.

4.1.2 Microwave-transfer field calibration

The second field calibration technique exploits microwave-transfer and Stern-Gerlach methods. In Sec. 3.6, we saw that atoms can be transferred from the trapped $|2, -2\rangle$ state to the anti-trapped $|3, -3\rangle$ state using microwave pulses or microwave ARP. If the atoms in the $|3, -3\rangle$ state are allowed to fall from the trap, both states can be imaged simultaneously. This gives a measurement of the atom number in each state without having to assume the total number of atoms remains constant from shot to shot, as was necessary with our realisation of RF field calibration. Note that it is possible to implement a similar Stern-Gerlach method for the RF field calibration.

A typical calibration involves a 15 µs microwave pulse followed by a 4 ms hold to allow the atoms in the $|3, -3\rangle$ state to spatially separate from the trapped atoms vertically by $\sim 80\,\mu$m. We extract the magnetic field value of a fixed magnetic coil current by varying the microwave frequency from shot-to-shot and optimising the transfer efficiency. The optimal microwave frequency gives the transition line centre, from which the field can be calculated using the well-characterised Zeeman splitting [103]. Conversely we can fix the magnetic field value by setting the calculated microwave frequency and varying the coil current until the transfer efficiency is maximised (Fig. 4.2).

Microwave-transfer field calibrations provide knowledge of the magnetic field at the atoms to 1 mG precision, which equates to a scattering length of $\sim 0.04\,a_0$ at $a_\mathrm{s} = -10\,a_0$. This uncertainty is more than precise enough to produce stable solitons, as we will see in the following sections. For comparison, the RF field calibration gives a precision of 3 mG ($\sim 0.12\,a_0$ at $a_\mathrm{s} = -10\,a_0$) which is still sufficient. However, the practicalities of the microwave-transfer calibration make it much faster and easier to perform, hence it is the method used throughout the remainder of this thesis.

4.2 Experimental sequence

As described in Sec. 1.1, *true* solitons require a strict 1D geometry in order to persist as a stable state. Clearly, a truly 1D geometry is impossible to realise
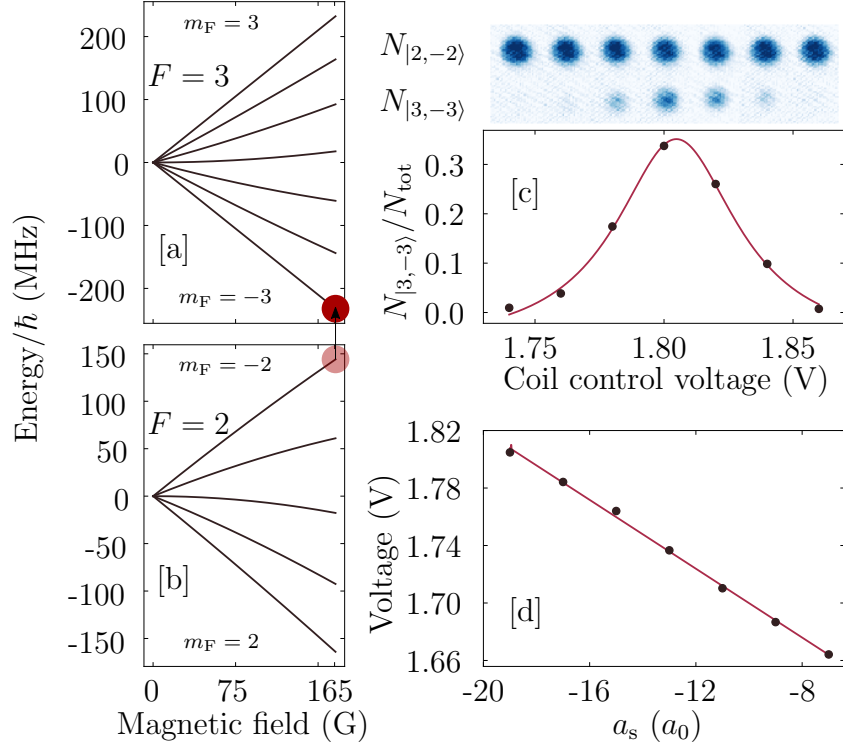
Figure 4.2: Magnetic field calibration by microwave spectroscopy and Stern-Gerlach separation. Figures [a] and [b] show Breit-Rabi diagrams of the $F = 3$ and $F = 2$ ground state manifolds respectively, up to 165 G, from which the $|2, -2\rangle \rightarrow |3, -3\rangle$ transition frequency is calculated (red circles). An example calibration is shown in [c], for a fixed microwave frequency, which is determined by the target magnetic field. The raw images are shown above: the upper (lower) cloud is made up of atoms in the $|2, -2\rangle$ ($|3, -3\rangle$) state. Using appropriate microwave frequencies we characterise the magnetic field coils across a range of scattering lengths close to a Feshbach resonance [d].

in the laboratory, as any trapping must have radial extent. Fortunately it is still possible to form solitons when a *sufficiently* 1D geometry has been created, such that there is a separation of scales between the radial and axial directions and the GPE integrability isn't badly broken [108]. This is deemed to be a *quasi*-1*D* geometry, which we achieve experimentally by loading the BEC into a single dipole trapping beam.

4.2.1   Loading the waveguide

Initially, we perform a linear ramp of the scattering length from $a_\mathrm{s} = 219\,\mathrm{a_0}$ (the value at the end of the BEC evaporation ramps) to $a_\mathrm{s} = (0.00 \pm 0.04)\,\mathrm{a_0}$, over 100 ms. This reduces the size of the condensate and creates a BEC that is approximately in the harmonic oscillator ground state of the crossed dipole trap. We then simultaneously jump the magnetic field to a slightly negative scattering length and remove the *N-S* beam. This releases a condensate into the *W-E* beam, hereafter labelled the *waveguide* beam, which is creating the required quasi-1D geometry for soliton production (Fig. 4.3). In Sec. 4.3 we explore varying the final scattering length to produce solitons, by engineering a soliton wavefunction that is spatially mode-matched to that of the initial condensate.

An additional magnetic coil pair, the *curvature coils*, produce a weakly confining harmonic potential along the waveguide:

$$V = \frac{1}{2}m\omega_\mathrm{z}^2 \left(z - z_0\right)^2,\tag{4.1}$$

where ω_z is the axial trapping frequency and z_0 is the offset of the curvature trap from the crossed-dipole trap centre. If $z_0 \neq 0$, the BEC undergoes centre of mass oscillations once released into the waveguide.

These steps change the trapping frequencies of the crossed dipole trap from $(\omega_\mathrm{r}, \omega_\mathrm{z}, \omega_\mathrm{vert}) = 2\pi \times (50, 30, 50)$ Hz, where ω_r is the radial trapping frequency, to $2\pi \times (50, 0.5 \to 1.5, 32)$ Hz for the waveguide setup. The radial and axial directions here are defined in terms of the waveguide trap. The range in ω_z is due to the range of achievable curvature coil currents.

Figure 4.3: An example magnetic field (scattering length) ramp and jump sequence used for loading the waveguide and soliton creation. Inset plots [a] and [b] are sketches of the trapping geometries for crossed dipole trap and waveguide trap respectively. The red shaded regions are the optical dipole trap and the blue contours show the off-centre magnetic quadrupole trap formed by the curvature coils.

Centre of mass oscillation control

To control the amplitude of the centre of mass oscillations of the condensate, the distance between the soliton release position and the harmonic curvature potential centre is controlled (z_0). This can be achieved by either varying the release position, or by moving the centre of the harmonic potential. In practice, it is more convenient to shift the position of the harmonic potential centre relative to the fixed crossed dipole trap release position using magnetic shim coils (the same coils used to provide quantisation axes in Sec. 3.4). This effectively initiates the release position closer/further away from the harmonic potential centre, resulting in smaller/larger centre of mass oscillations (see Figs. 4.3[a] and [b]). The centre of mass oscillation frequency is controlled by the harmonic potential produced by the curvature coils, with frequencies in the range $2\pi \times 0.5 \, \text{Hz} < \omega_z < 2\pi \times 1.5 \, \text{Hz}$ explored in this work.

Loading into N-S beam

Our initial exploration of loading a condensate into a single beam was centred around the N-S beam, not the W-E (waveguide) beam described above. This was because we want to study atom-surface interactions, using a prism mounted in the science cell to the north side of the atoms (see Ch. 7). However, whilst we were able to load a condensate (and form a soliton) in the N-S beam, the atoms rapidly accelerated outside of the field of view of the imaging system (Fig. 4.4). We believe this was due to an geometric asymmetry in the bias coils, caused by the point at which they start/ finish winding. As there are only four windings per coil and large currents of $\sim 200 \, \text{A}$ are required to produce the $\sim 160 \, \text{G}$ field necessary for soliton production, the small asymmetry pushed the centre of the harmonic potential far away. Fig. 4.4 shows that this offset was $(16 \pm 2) \, \text{mm}$ from the release position. Despite having a set of shim coils in the N-S direction, with which we could manipulate the trap centre, this effect was too large to be compensated for.

Fortunately, because the bias coils start/ finish winding on the north side of the science cell, as defined in Fig. 2.1, the W-E (waveguide) beam does not

Figure 4.4: A BEC loaded into the N-S beam is seen to rapidly leave the field of view of the imaging system. The solid line is a sinusoidal fit to the data, from which we extract the centre of the oscillation to be $(16 \pm 2)\,\text{mm}$ away from the release position.

have such a significant offset. Therefore, using the E-W shims and curvature coils, we can observe full sinusoidal centre of mass oscillations of the condensate and have complete control of the oscillation amplitude over a wide range, as we will see in the following sections.

4.3 Tunable interactions: BEC width oscillations

In 1D-GPE theory, any negative scattering length should produce a soliton, until the collapse threshold is reached [109]. However, as solitons are produced in our experiment following a rather violent change in trap geometry and the width of the BEC in the crossed dipole trap is quite narrow ($\sim 7\,\mu\text{m}$), the scattering length has to be sufficiently negative to engineer a soliton with a similar equilibrium width. If the width of the Thomas-Fermi BEC profile in the crossed trap and target soliton width in the waveguide are too dissimilar then width oscillations will be observed because of the rapid change of equilibrium. Here we demonstrate how these width oscillations may be used to ensure good spatial mode-matching and robust soliton production.

Once the BEC has been released into the waveguide, following the routine in Sec. 4.2.1, the condensate undergoes centre of mass oscillations (see Fig. 4.5). Superimposed on these oscillations the BEC also undergoes width oscillations, similar to the *breathers* observed when the BEC experiences a magnetic field quench [110]. However, for repulsive/ weakly-interacting BECs, these width oscillations are primarily due to the atoms bunching up at the edges of the trap and dispersing as they travel towards the centre.

Figure 4.5: Width oscillations of a BEC undergoing centre of mass oscillations. [a] A sketch to show the expansion/ contraction of a BEC as it oscillates in a quasi-1D weakly-harmonic potential. [b] Image sequence showing a non-interacting condensate ($a_s = 0$) oscillating in a quasi-1D harmonic potential with axial trapping frequency $\omega_z = 2\pi \times 1.4\,\text{Hz}$ and centre of mass oscillation amplitude $z_0 = 90\,\mu\text{m}$. The line is a least-squares fit to the extracted condensate centres.

By varying the final scattering length, we observe that the width oscillation amplitude can be controlled (Fig. 4.6). We determine three distinct situations, depending on the scattering length: dispersion and contraction of the condensate with frequency $2\omega_z$, no observable width oscillations and collapse with subsequent contraction, again with frequency $2\omega_z$. In Fig. 4.7, we show the dependence of the width oscillation amplitude on scattering length and highlight three regions that link directly with the above observations: the weakly-interacting region, the soliton region and the strongly-attractive region respectively.

Figure 4.6: [a]-[c] show BEC width oscillations for $a_\mathrm{s} = 0$, $-12\,a_0$ and $-15\,a_0$ respectively. The error bars indicate the standard error of 3 repeats for [a] and [b]. The data in [c] has been split into two depending on measured condensate atom number: red triangles for $N > 720$, blue circles for $N < 720$. The vertical scales are consistent across [a]-[c] for ease of comparison. The visible oscillations in [a] and [c] are at a frequency $2\omega_\mathrm{z}$.

Figure 4.7: Width oscillation amplitude as a function of scattering length. The vertical dotted lines highlight the crossover between the weakly-interacting (i), soliton (ii) and strongly-attractive regimes (iii). The black points in the strongly-attractive regime show the data with $N > 720$.

In the weakly-interacting region, the interatomic interactions are insufficient to prevent the expansion and contraction of the BEC, imposed on it by dispersion and the trapping potential respectively. This is caused by the spatial mode of the target soliton wavefunction being too large, relative to the condensate in the crossed dipole trap, such that the rapid change in equilibrium results in width oscillations.

In the soliton region, the inter-atomic interactions perfectly balance dispersion and the width oscillations are nulled. In this region the wavefunction of the soliton is sufficiently close to the initial BEC, such that the soliton is able to remain stable.

In the strongly attractive region, collapse is observed when the BEC collects at the edges of the harmonic trap and exceeds a critical density. This explosion in width following collapse leads to *apparent* width oscillations, as the BEC comes back together again after a half trapping period. This effect is similar to that observed during the first investigations of BECs with attractive interactions [111] where collapse was followed by a rapid ejection of atoms. However, our additional confining harmonic potential is sufficient to capture most of the ejected atoms and refocus them after half a trapping period.

The fixed width oscillation frequency of $2\omega_z$ across the whole of the weakly-interacting region may appear surprising, as a recent experiment has shown experimentally that the breathing frequency of a soliton depends on the combined interaction strength $N|a_s|/a_r$ [112]. It is also expected that a BEC in the 1D Thomas-Fermi regime has a width oscillation frequency $\sqrt{3}\omega_z$ [113]. In fact, the frequency $2\omega_z$ is consistent with non-interacting condensates and thermal atoms [114, 115]. However, the fixed frequency of $2\omega_z$ in our experiment is caused by the dispersion and contraction of the atomic wavepacket as it travels down and up the harmonic potential. This is the dominant effect in the weakly-interacting region because the interatomic interactions are insufficient to compensate for the mode-mismatch between the spatial density profiles of the initial BEC and the target soliton. Positioning the harmonic

potential to overlap the crossed dipole trap would prevent his from occurring, as the condensate will have no centre of mass velocity, allowing the recovery of the frequency dependencies observed in Ref. [112]. However, as we are only interested in minimising the magnitude of the width oscillations, we leave this unexplored.

Interestingly, in the absence of collapse, we observe a distinct reduction in atom number as the scattering length is made more negative across the soliton region (Fig. 4.8). The result of this is that the combined interaction strength $N|a_{\rm s}|/a_{\rm r}$ is approximately constant across the soliton region, at 0.41 ± 0.05. As expected, this is far short of the 0.675 collapse threshold from Eq. 1.11, which may indicate that the soliton undergoes some procedure of self-stabilisation during its formation. For the solitons we know to have collapsed (marked in black in Fig. 4.8), $Na_{\rm s}/a_{\rm r} = 0.80 \pm 0.01$, which is indeed above the numerically-determined collapse threshold. Note that we cannot accurately determine the frequency of any width oscillations that might exist across the soliton region, as the shot-to-shot noise in the width is generally larger than the oscillation amplitude. However, as $N|a_{\rm s}|/a_{\rm r}$ is approximately constant, it is expected that any width oscillations would also be of constant frequency [112].

In Fig. 4.7 we found that scattering lengths below $\sim 7\,a_0$ were sufficient to form a soliton in our realisation, however this is completely dependent on the spatial profile of the BEC in the crossed dipole trap, as explained above. If the crossed dipole trap was relaxed slightly before releasing into the waveguide, it may be possible to create a wider BEC, which would allow the stable production of solitons with less negative scattering lengths. This could be achieved by simply reducing the dipole trap beam powers. Alternatively, stable solitons with less negative scattering lengths may be created by starting with a BEC of higher atom number, due to the $1/N$ dependence of the soliton width. These options were unnecessary for the experiments performed later in this thesis and so were left unexplored. However, they may provide even deeper insight into the nature of soliton formation in the future.

Figure 4.8: The variation in atom number with scattering length during the width oscillation experiment. The vertical error bars indicate the standard error in atom number across the oscillation data at each scattering length. The vertical dotted lines highlight the crossover between the weakly-interacting (i), soliton (ii) and strongly-attractive regimes (iii). The black points in the strongly-attractive regime show the data with $N > 720$.

To better understand the structure and form of the soliton we study the effect that varying the scattering length has on the soliton width itself. Instead of observing the width at multiple hold times, in Fig. 4.9 we perform multiple runs for the same hold time. We find that the soliton width narrows for more negative scattering lengths, comparing well with the $1/|a_\mathrm{s}|$ relationship given by soliton solutions to the 1D-GPE (Sec. 1.1).

Figure 4.9: The measured soliton width over the soliton-forming regime. Blue points show the measured soliton width after a fixed hold time $(0.2\,\mathrm{s})$, with error bars indicating the standard error of 10 experimental runs. Red points show the mean width measured across the soliton-forming regime in Fig. 4.7, with error bars indicating the standard error. The line is a guide to the eye of $w \propto 1/|a_\mathrm{s}|$.

4.3.1 Single soliton or soliton train production?

It is interesting to note that we only ever create a single soliton with our method, over all scattering lengths investigated. This is contrary to other soliton experiments [24, 26, 29, 116, 117] where soliton trains are produced.

Soliton trains are formed as a result of modulation instability in an elongated condensate with attractive interatomic interactions, whereby small perturbations from noise or self-interference seed the exponential growth of the local

density [46–48,116,117]. As these density peaks grow, the attractive interactions and random phases induce a second stage that includes collapse and/or merging between pairs of solitons, followed by a subsequent resurgence of density peaks. This self-regulates into the formation of a stable series of solitons with an alternating (0-π-0) phase structure, preventing further collapse.

Modulation instability in a condensate promotes the most rapid density growth for the wavenumber:

$$k_{\mathrm{MI}} = \frac{\sqrt{4\,|a_{\mathrm{s}}|\,n_{\mathrm{1D}}}}{a_{\mathrm{r}}}, \tag{4.2}$$

where $a_{\mathrm{r}} = \sqrt{\hbar/(m\omega_{\mathrm{r}})}$ is the characteristic radial confinement length scale and n_{1D} is the density along the line of sight. From the wavenumber we can calculate the modulation instability wavelength, λ_{MI}:

$$\lambda_{\mathrm{MI}} = \frac{2\pi}{k_{\mathrm{MI}}} = \frac{2\pi a_{\mathrm{r}}}{\sqrt{4\,|a_{\mathrm{s}}|\,n_{\mathrm{1D}}}}, \tag{4.3}$$

which is equivalent to the distance between solitons in the train at the point of stable formation. This has been found to agree well with the spacings measured in experiments [116,117]. For example, in the experiment at Rice, an extremely elongated ^{7}Li condensate with $\sim 3 \times 10^5$ atoms is formed in a single beam dipole trap with $(\omega_{\mathrm{r}}, \omega_{\mathrm{z}}) = 2\pi \times (346, 7.4)$ Hz. A magnetic field quench is performed to rapidly change the scattering length from $a_{\mathrm{s}} = 3\,a_0$ to $a_{\mathrm{s}} = -0.18\,a_0$, creating a soliton train. Eq. 4.3 predicts a soliton spacing of $\sim 15\,\mu$m, in good agreement with that measured in Refs. [116] and [24].

Applying the same formula to our experiment yields essentially the same λ_{MI}, using typical experimental values. However, as we have far fewer atoms in our condensate and we simultaneously quench the interaction strength and the trap geometry (from the crossed-dipole trap to the waveguide trap), the condensate is nearly spherical and spatially smaller than those produced at Rice. In fact, Fig. 4.6 shows that we typically measure initial condensate widths of $\sim 7\,\mu$m, which is a factor of two smaller than the λ_{MI}. Therefore, there is insufficient space to seed more than a single soliton through

modulation instability, so we only ever produce one soliton at the point of formation.

A calculation of λ_MI also explains the recent observations at the Jožef Stefan Institute [29], where a near-spherical condensate still formed a soliton train. Their experiment is based on a ^{133}Cs condensate with $\sim 10^4$ atoms, in a trap with $(\omega_\mathrm{r}, \omega_\mathrm{z}) = 2\pi \times (107, 3.33i)$ Hz, at comparable scattering lengths to our experiment. The 'i' denotes an imaginary trapping frequency, as the curvature of their harmonic trap is negative. The increased mass, atom number and trapping frequencies cause their λ_MI to be a factor of ~ 4 smaller than in our experiment, at $\sim 3\,\mathrm{\mu m}$. As they measure initial condensate sizes of $\sim 15\,\mathrm{\mu m}$, we can expect modulation instability to result in ~ 5 stable solitons, which is similar to the numbers observed [29].

4.4 Soliton oscillations

As described in Sec. 4.2.1, the centre of mass oscillations of the BEC (or soliton) in the waveguide trap can be controlled: the curvature coils are used to set the trapping frequency and the magnetic shim coils are used to set the oscillation amplitude. In Fig. 4.10 we show examples of soliton trajectories as they oscillate in harmonic potentials with different parameters, illustrating the highly-tunable nature of our experimental configuration.

The importance of this control will be fully explained in Ch. 5, but essentially it allows us to set the velocity, or kinetic energy, of the soliton at specific points on its trajectory. This will prove to be vital for exploring soliton splitting and recombination.

4.4.1 Soliton lifetime

Our spatially mode-matched solitons are exceptionally long-lived, oscillating for longer than 20 s without measurable dispersion (see Fig. 4.11). This

Figure 4.10: Trajectories of solitons undergoing centre of mass oscillations in the harmonic trap produced by the curvature coils. [a]-[d] show trajectories of solitons oscillating in a trap with $\omega_z/2\pi \sim 0.55\,\text{Hz}$, $0.82\,\text{Hz}$, $1.13\,\text{Hz}$ and $1.4\,\text{Hz}$ respectively, for various oscillation amplitudes. The blue datapoint at $0.8\,\text{s}$ in [a] was outside the field of view of the camera and so is missing.

corresponds to ~ 30 centre of mass oscillations and a total distance covered of over $2\,\mathrm{cm}$. To achieve this, the quasi-1D harmonic potential must be very stable and, because each image is a discrete experimental run, the solitons must be reliably reproduced from shot-to-shot.
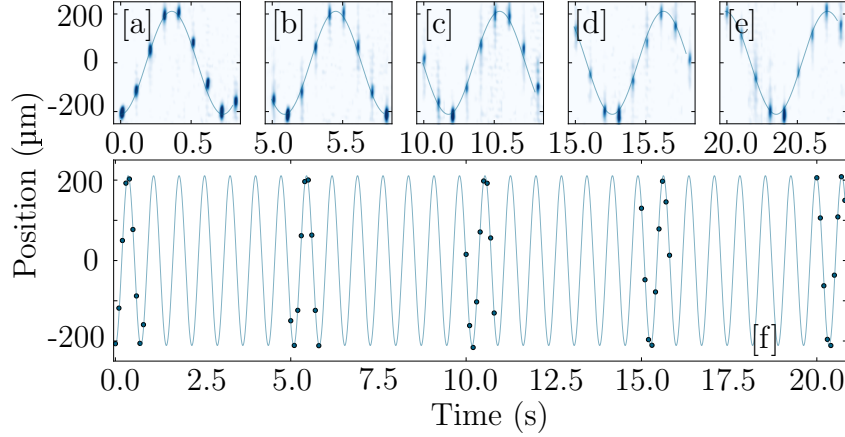
Figure 4.11: A sequence of images showing a soliton with $a_\mathrm{s} = -8.4\,\mathrm{a_0}$, undergoing centre of mass oscillations in a quasi-1D harmonic potential. [a]-[e] The raw images, overlaid with a fit to the soliton centre position. [f] The soliton is observed to oscillate for longer than $20\,\mathrm{s}$ with a centre of mass oscillation amplitude of $225\,\mathrm{\mu m}$ and a oscillation frequency of $\omega_\mathrm{z} = 2\pi \times 1.4\,\mathrm{Hz}$.

To achieve a harmonic potential with the curvature trap, the magnetic field must vary across the trap. However, if the variation is too large, the soliton may exceed the critical threshold of scattering length (seen in Fig. 4.7) and become unstable. Fig. 4.11 shows centre of mass oscillations of a soliton, for an amplitude of $225\,\mathrm{\mu m}$ and a trapping frequency of $\omega_\mathrm{z} = 2\pi \times 1.4\,\mathrm{Hz}$, which are typical maximum values used in our experiment. Despite these (relatively) large values, the variation of magnetic field is calculated to be less than $0.06\,\mathrm{mG}$ using Eq. 2.8. This is equivalent to a scattering length variation of less than $2 \times 10^{-3}\,a_0$ at $-10\,\mathrm{a_0}$. Such a small variation (below our field calibration resolution limit) does not impact the stability of the soliton. Shot-to-shot soliton atom number fluctuations of $\sim 10\,\%$ are a far bigger concern for soliton stability.

We measure the lifetime of the solitons in the quasi-1D harmonic potential by extracting the fitted peak optical depth of the images shown in Fig. 4.11[a-e]. As the peak optical depth is directly proportional to the number of atoms (Eq. 3.3), we can use these values to determine the atom loss decay curve. We model the OD as an exponential decay:

$$OD = OD_0 e^{-t/\tau}, \qquad (4.4)$$

where OD_0 and τ are the initial OD and lifetime respectively. A lifetime of $(28 \pm 3)\,\mathrm{s}$ is determined from the fit shown in Fig. 4.12.

Figure 4.12: The optical depth of solitons, measured at various stages of a centre of mass oscillation in the quasi-1D harmonic potential. The data correspond to the peak optical depths of fits to the images seen in Fig. 4.11. The solid line is a least-squares fit of Eq. 4.4 to the data, which returns a lifetime of $(28 \pm 3)\,\mathrm{s}$.

This incredibly long lifetime demonstrates one of the potential benefits of a soliton-based interferometry device, as longer phase accumulation times allow for higher sensitivity measurements.

Chapter 5

Soliton splitting

In this chapter we describe the realisation of a matter-wave beam splitter, where the splitting ratio can be fully controlled. We begin with a recap of existing systems that have shown soliton splitting, before providing details of the repulsive Gaussian barriers used in our experiment, as well as how they are characterised. We demonstrate the tunability of the splitting ratio before discussing the nature of the splitting mechanism in more detail. Finally, we explore a feature of soliton splitting that occurs in our experiment which has drastic consequences for the feasibility of a soliton interferometer: velocity selection.

5.1   Splitting techniques

To create an interferometer, or to investigate soliton-soliton collisions, two coherent solitons are required. There are several experiments that have demonstrated the formation of two solitons, each using different techniques. Here we provide a brief outline of each of technique, along with their respective advantages and disadvantages.

Begin with two condensates

In the experiment at Rice, two solitons are formed from two separate ^{7}Li condensates [40]. Initially a single condensate is created, before it is exposed to a cylindrically-focussed blue-detuned laser beam, which creates a double well potential and cuts the condensate in half. Following a magnetic field ramp from $a_{\mathrm{s}} = 140\,\mathrm{a_0}$ to $-0.57\,\mathrm{a_0}$, two solitons with ~ 28000 atoms in each are produced, spatially separated by $26\,\mu\mathrm{m}$. The barrier is removed and these solitons are used to study soliton-soliton collisions.

In the experiment at the Jožef Stefan Institute, two solitons are again formed from two separate condensates [29]. However, in this case, the solitons are of ^{133}Cs and the method of creating two condensates differs greatly. Instead of splitting a single BEC, simultaneous evaporation to degeneracy in two independent traps is performed. This method requires particular attention to the number of atoms in each trap, due to the soliton stability conditions discussed in Sec. 4.3 and the modulation instability thresholds discussed in Sec. 4.3.1.

However, these methods of producing two solitons are both ill-suited to phase-sensitive experiments, such as interferometry or soliton-soliton collisions. As shown by the soliton-soliton collision experiment performed by group at Rice [40], the relative phase between the solitons cannot be known *a priori* and must be inferred by the outcome of the collision. This is because the two solitons have no interaction before the collision and so are initiated with a random phase difference. Even if the two BECs can be created to be coherent, it is unlikely that coherence could be maintained during the soliton formation process.

Begin with two solitons

As both of the above experiments have demonstrated the creation of soliton trains seeded by modulation instabilities (along with the experiments

described in Refs. [26, 117]) the atom number and scattering length could, in principle, be chosen to create two solitons only. As long as each individual soliton is sufficiently large that it is close to the collapse threshold, such that the sum of the two would exceed the critical number of atoms, this method has the significant attraction that the relative phase between the solitons is fixed (owing to the phase-stabilising secondary collapse discussed in Sec. 4.3.1). However, to reliably produce only 2 solitons with each run would require exceptionally high levels of control over shot-to-shot fluctuations in the evaporation trajectory, which may be unattainable. Furthermore, it would be technically challenging to separate and control the two solitons without disrupting the phase coherence.

Splitting: Bragg Pulses

In the experiment at the ANU, whilst also being able to create soliton trains [117], they have demonstrated the splitting of a soliton into two daughter solitons using Bragg diffraction [27]. Splitting was achieved by performing a $\pi/2$ Bragg pulse, diffracting their initial soliton off a 1D optical lattice. The lattice was formed collinear with their waveguide, using a pair of counter-propagating beams detuned $100\,\mathrm{GHz}$ from the ^{85}Rb D2 transition. This method was also proven to retain some phase coherence, as interferometric fringes were recovered following a standard $\pi/2$-π-$\pi/2$ set of Bragg pulses, with the phase of the final $\pi/2$ pulse being shifted to recover the fringes.

However, this experiment suffered severe limitations in coherence times: the visibility of the interferometric fringes was negligible after only a few milliseconds of phase-acquisition between the Bragg pulses. Furthermore, because of the short phase-acquisition times, the split solitons remained spatially overlapped at all times. This precludes applications for interferometry or soliton-soliton collisions as it is impossible to affect a relative phase shift between the two solitons. It is not clear what was the cause of these short coherence times, though it is possible that the Bragg pulse splitting method causes unwanted phase noise.

Splitting: narrow attractive barrier

In a previous investigation from our experiment, splitting of a soliton was achieved using a narrow attractive barrier [33]. The red-detuned barrier was formed from an 852 nm laser beam, focussed down to intersect the waveguide with a width of 1.9 µm. When incident on this barrier, the soliton was observed to break into transmitted, reflected and confined components. Despite the classical particle-like appearance of solitons, reflection from an attractive barrier (or from above a repulsive barrier) is made possible because of quantum reflection, exposing the wave-like properties of solitons [118]. The splitting fractions were seen to be sensitive to the barrier potential depth: deepening the barrier resulted in larger reflected and confined fractions, with a reduced transmitted fraction. Although the experiment was limited by barrier potential depth and was thus unable to demonstrate 50 % reflection, with a more powerful barrier laser it would be possible to split the initial soliton in half.

Whilst this experiment probed interesting questions about the fundamental quantum mechanical nature of solitons, it faces several drawbacks as a means to achieving phase-coherent splitting. Firstly, as images of the atoms were typically resolution-limited, it was impossible to determine whether the split components were themselves solitons [15], which would be required for soliton-soliton measurements. Furthermore, the fraction of the soliton that is confined by the attractive potential would likely prove to be a nuisance for collision measurements and would need to be removed before the collision.

Splitting: narrow repulsive Gaussian barrier

There has been a great deal of theoretical interest into splitting solitons using narrow repulsive barriers [22, 60–66, 119–122]. Contrary to the attractive barrier case, splitting on a repulsive barrier can be either quantum or classical, through quantum tunneling or classical velocity-distribution splitting respectively. As we will see during the remainder of this chapter, the exact

nature of the splitting mechanism has drastic consequences for the utility of the atomic beamsplitter as a phase-coherent splitting mechanism.

This method of splitting a soliton has been realised by the group at Rice in one previous experiment [123]. Similar to the attractive barrier described above, the repulsive barrier was focussed down to intersect the waveguide and the soliton was allowed to collide with it. A fraction of the soliton was observed to transmit beyond the barrier and the remainder was reflected. Unlike in the attractive barrier case, note that there is no confined fraction as the potential is repulsive. The splitting fraction was shown to be coarsely controllable by changing the height of the potential barrier. Unfortunately, shot-to-shot noise in the soliton position and velocity proved overwhelming and made further investigations infeasible.

Nevertheless, in this work we investigate splitting with narrow repulsive barriers. As shown in Ch. 4, particularly in Fig. 4.11, our solitons have highly stable and reproducible trajectories, with shot-to-shot position fluctuations as low as $1.3\,\mu m$ (see Sec. 6.2). Therefore, we may be able to overcome the shot-to-shot uncertainties experienced by the Rice experiment and better exploit the narrow repulsive barrier as an atomic beamsplitter.

5.2   Soliton splitting: theory

There are a wealth of theoretical works that explore the mechanics and dynamics of splitting a soliton on a narrow repulsive barrier. Some works analyse an idealised, δ-function barriers [63, 66, 121], where analytical expressions are more-easily derived and simulations are less computationally expensive. However, most focus on simulations based on finite-width barriers, with various profiles [22, 60–66, 119, 122]. Here we present an overview of the theoretical aspects of soliton splitting, in order to provide context to the results that follow.

As mentioned in the previous section, upon reaching a narrow repulsive barrier a soliton will be either transmitted, reflected or some combination of the

two. The exact outcome will be heavily influenced by the size and shape of the barrier.

5.2.1 δ-function barriers

A δ-function barrier manifests as an additional potential term in the GPE (Eq. 1.3) of the form:

$$V_\delta(z) = q\delta_0(z), \tag{5.1}$$

where q is the strength of the barrier. The parameter q can also be interpreted as the *area* of the barrier, as $\int_{-\infty}^{\infty} V_\delta(z)\,dz = \int_{-\infty}^{\infty} q\delta_0(z)\,dz = q$. In Ref. [121], it is shown that this potential leads to a quantum transmission coefficient of:

$$T_\delta = \frac{v^2}{v^2 + q^2}, \tag{5.2}$$

where v is the soliton velocity at the point of collision with the barrier. It is important to note that Eq. 5.2 is strictly only a valid description of soliton splitting on a δ-function barrier for high soliton velocities ($v \to \infty$). In this limit the soliton-barrier interaction time is minimised, giving it a transmission coefficient that is approximately equal to the transmission coefficient for a plane wave on an identical δ-function barrier. This allows the splitting to be treated as a linear process, which is described by Eq. 5.2.

Eq. 5.2 is plotted in Fig. 5.1 and clearly shows that as the kinetic energy of the soliton grows with respect to the barrier area, a larger fraction of the soliton will be transmitted. However, as δ-function barriers are experimentally impossible to realise, Eq. 5.2 does not offer a realistic model of our experiment.

It is worth commenting on another feature of splitting a soliton on a δ-function barrier: the daughter solitons acquire a $\pi/2$ relative phase during the process. Therefore, despite the δ-function barrier being a practical impossibility, it would be beneficial for future phase-sensitive experiments to reach a regime as close as possible to it.

Figure 5.1: The quantum transmission coefficient for a δ-function barrier, as a function of the soliton velocity normalised by the barrier area. The solid blue line shows the transmission coefficient and the dotted red line shows the corresponding reflection coefficient ($R_\delta = 1 - T_\delta$).

5.2.2 Finite-width barriers

The barrier beam in our experiment is expected to be of Gaussian profile. Therefore, instead of the δ-function described by Eq. 5.1, the Gaussian function modifies the potential term to:

$$V_{\text{Gauss}} = V_0 e^{\frac{-2z^2}{w_z^2}}, \tag{5.3}$$

where V_0 is the potential in the centre of the Gaussian and will be defined later (Eqs. 5.7 and 5.8).

A full GPE simulation (in either 1D, quasi-1D or 3D) can be used to numerically evaluate the transmission coefficient for various barrier parameters using this Gaussian function. However, there exist a few functions for which an analytic expression of the transmission coefficient is known, at least for non-interacting atoms. One such profile is the sech^2 potential [122]:

$$V_{\text{sech}^2} = V_0 \, \text{sech}^2\left(\frac{\pi z}{\alpha w_z}\right), \tag{5.4}$$

where $\alpha = (\pi/2)^{3/2}$. The sech^2 potential also has the advantage that the profile is very similar to a Gaussian, particularly across the central region of the function (see Fig. 5.2).

Figure 5.2: A comparison between the narrow ($w_z = 3.6\,\mu\text{m}$) Gaussian beam profile used in the experiment and a sech^2 approximation, shown by the solid blue line and dotted red line respectively.

The transmission coefficient of the sech^2 potential is given by:

$$T_{\text{sech}^2} = \frac{\sinh^2\left(w_z \pi \sqrt{\kappa E_k}\right)}{\sinh^2\left(w_z \pi \sqrt{\kappa E_k}\right) + \cosh^2\left(\frac{\pi}{2}\sqrt{4 w_z^2 \kappa V_0 - 1}\right)}, \tag{5.5}$$

where $\kappa = 2m/\hbar^2$ and E_k is the kinetic energy of an atom within the soliton. It is shown in Ref. [63] that this expression recovers good agreement with numerical results for low nonlinearities.

Example transmission coefficient curves are shown in Fig. 5.3, clearly demonstrating an increased sharpness in the reflection-transmission crossover for wider barriers. This figure can be intuitively understood in terms of quantum tunneling in the following way. For a very narrow barrier, there is a high likelihood of tunneling for a wide range of soliton kinetic energies, hence the transmission function is broadened and the centre of mass kinetic energy of the soliton may be far lower than the barrier height to achieve 50 % splitting (Fig. 5.4). However, tunneling is far less likely for wide barriers of the

same barrier height, owing to the exponential decay of tunneling probability with distance through a potential barrier. Therefore, in order to successfully pass over/through the barrier, the kinetic energy of the soliton must be much closer to the amplitude of the barrier potential. This creates a much narrower crossover between reflection and transmission, hence the transmission function is narrowed. The transmission coefficient curve becomes a step function in the limit of an infinitely wide barrier.

Figure 5.3: The transmission coefficients for a sech^2 barrier, as a function of the normalised kinetic energy of the soliton, for various barrier widths. The barrier height is $V_0 = 5.14\,\mathrm{nK}$ for all of the curves, which is typical of the barrier heights used in the experiment.

It is important to note that, even though in the limit of infinitely-wide barriers we recover a step function for the transmission function, in the limit of infinitely-narrow barriers we do not find a uniform transmission probability across all kinetic energies, which might be the initial expectation. Instead, Eq. 5.5 recovers the δ-function barrier result predicted by Eq. 5.2 [63].

The splitting process splits the soliton into two atomic ensembles, each of which may itself be a soliton [121]. Vitally, these daughter solitons are expected to remain coherent, with a fixed relative phase. However, the relative phase is not necessarily the $\pi/2$ of the δ-function barrier. Instead, it has been found to be sensitive to the velocity of the initial soliton at the barrier, the

Figure 5.4: The dependence of the normalised kinetic energy required to achieve 50 % splitting on the barrier width. For barrier widths below $\sim 3\,\mu\text{m}$ there is a sharp drop, which indicates the appearance of a significant quantum tunneling contribution, whereas classical centre of mass splitting dominates for wider barriers. Inset are pictorial representations of a soliton incident on $0.5\,\mu\text{m}$ and $10\,\mu\text{m}$ barriers, with the red circle indicating a soliton at the kinetic energy required to split into two solitons of equal atom number, for each barrier respectively.

interaction strength and the barrier width [63]. Specifically, the phase difference is expected to increase with increasing velocity, interaction strength and barrier width. In the limit of a high-velocity soliton and a very narrow barrier, the phase difference tends towards the $\pi/2$ of the δ-function barrier. Although the relative phase between the daughter solitons in our experiment is not expected to be $\pi/2$, by ensuring that the soliton velocity, interaction strength and barrier width are sufficiently stable from shot-to-shot, we can expect the daughter solitons to remain coherent with a fixed relative phase. Creating two solitons that retain coherence is vital for the application of the narrow barrier as a beamsplitter in phase-sensitive experiments, as will be explored extensively later in this thesis.

5.3 Repulsive Gaussian barrier

We create a repulsive Gaussian barrier using a blue-detuned highly-elliptical laser beam focussed down to bisect the waveguide through the centre of the quasi-1D harmonic potential (see Fig. 5.5).

Figure 5.5: A blue-detuned (532 nm) highly-elliptical beam is focussed onto the waveguide, by an objective lens, such that the narrow axis aligns with the axial direction and the wide axis aligns transverse to the waveguide. Also shown are the circular curvature coils and rectangular shim coils that are used to create the harmonic trapping potentials. The coordinate axis displays the axial (z), transverse ($\perp$) and vertical (vert) directions.

The optical potential for the Gaussian barrier is produced by a 532 nm laser beam. An AOM is placed in the beam path for intensity control and switching. The light is delivered to the experiment by an optical fiber, whereupon it is collimated and passed through a cylindrical lens to form a highly elliptical beam. This is focussed onto the waveguide such that the narrow axis is oriented along the axial direction of the trap. The barrier position in the horizontal plane (along/ transverse to the waveguide) can be adjusted via a piezo-actuated mirror before the objective lens (see Fig. 5.6). Without the limitations of imaging resolution (see Sec. 3.9), the piezo-actuated mirror could allow barrier alignment with a precision of (16.0 ± 0.5) nm.

In this work we investigate two barrier widths, which are generated using two different objective lenses. The *wide* barrier is produced by an objective lens with focal length 100 mm, whereas the *narrow* barrier is produced by an objective lens with focal length 30 mm. These are the same lenses that are introduced in Ch. 3.

It could be possible to vary the distance between the objective lens and the atoms to gain essentially arbitrary control over the barrier width, as defocussing the barrier beam would effectively broaden it from the point of view of the atoms. However, such a method would make reliable barrier characterisation difficult and may cause additional stability noise on measurements, as the barrier width becomes more sensitive to vertical position fluctuations away from the focus. This is due to how the width of a Gaussian beam w varies with distance from the waist position x_{off}:

$$w = w_0 \sqrt{1 + \left(\frac{x_{\mathrm{off}}}{l_{\mathrm{R}}}\right)^2}, \tag{5.6}$$

where l_{R} is the Rayleigh length and w_0 is the waist. A fixed mechanical fluctuation of just $l_{\mathrm{R}}/4$, where $l_{\mathrm{R}} \sim 50\,\mu\mathrm{m}$ for the narrow barrier, leads to a possible shift in width of $\sim 3\,\%$ when at the focus. However, if the barrier were to be intentionally defocussed by l_{R}, this changes to $\sim 13\,\%$. In fact, for the full value of the uncertainty, an additional factor of approximately two

Figure 5.6: The optical beam path of the barrier beam. [a] shows the optics required for producing, switching and controlling the power of the barrier beam, which are all placed far from the main experiment to prevent unwanted vibrational noise from the laser's cooling fans. The light is coupled through an optical fiber to the main experiment, whereupon it passes through a cylindrical lens to create an elliptical beam profile [b]. The beam is then deflected upwards by the horizontal/ vertical (H/V) mirror, onto optics mounted on a vertical breadboard [c]. Here some light is collected for the servo and the rest is focussed onto the atoms through the top of the science cell, using a piezo mirror for precise pointing of the beam. An optional neutral density (ND) filter can be included after the servo for better intensity resolution at low powers. The beam path through the cell is shared with the counter-propagating vertical imaging beam, which is reflected by a dichroic mirror (THORLABS DMSP650L) onto the camera.

arises for the defocussed barrier. This is due to fact that the width can either increase or decrease, whereas it can only increase at the focus. Therefore, we always align the waist of the narrow direction of the barrier such that it intersects the centre of the waveguide vertically. This ensures that the atoms only interact with the barrier where it is narrowest.

Figure 5.7: The barrier width uncertainty increase associated with defocussing away from the centre of the waveguide, for a fixed position uncertainty of $\pm l_{\mathrm{R}}/4$. For a barrier aligned so that the waist is in the centre of the waveguide, the uncertainty in beam width due to mechanical instabilities is only $^{+3}_{-0}$ % of its value ($w_1 \equiv w_0$ for the centred barrier case). However, if the barrier is defocussed by l_{R}, the uncertainty in barrier width increases to $^{+13}_{-12}$ % of its new value. The solid blue line shows the width of a Gaussian beam as it propagates close to the position of its waist.

Control over the potential height of the barrier V_0 is extremely important for performing tunable splitting of a soliton. V_0 is defined as:

$$V_0 = -\frac{1}{2\epsilon_0 c} I_{\mathrm{peak}} \alpha, \qquad (5.7)$$

where I_{peak} is the peak intensity of the barrier in the plane of the atoms and α is the polarisability of ^{85}Rb at $532\,\mathrm{nm}$ ($-4.19 \times 10^{-39}\,\mathrm{C\,m^2\,V^{-1}}$, calculated

using Ref. [124]). For an elliptical Gaussian beam with axial and transverse widths w_z and $w_\perp$ (in the reference frame of the waveguide) I_{peak} is given by:

$$I_{\mathrm{peak}} = \frac{2P}{\pi w_z w_\perp},\tag{5.8}$$

where P is the total beam power. Therefore, accurate determinations of the barrier power and barrier widths are necessary to gain a good understanding of the barrier height.

5.3.1 Power calibration

As the barrier width is fixed for a particular objective lens, the barrier potential height is controlled predominantly by changing the total beam power. A fraction of the barrier beam is taken from the beam path after the fiber and is passed onto a photodiode (see Fig. 5.6[c]). This provides a monitor feedback for a servo that compares the photodiode value with a controllable setpoint and uses the AOM to increase/ lower the power allowed through it in response, to minimise the difference between the monitor and the setpoint.

In order to ascertain the absolute barrier height, it is important to precisely calibrate the barrier power at the atoms, as a function of the servo setpoint. However, as we are unable to directly measure the power inside the (uncoated) science cell, we instead perform a calibration by measuring the power before and after the beam passes through the science cell to compensate for losses from the glass surfaces. Unfortunately, the 30 mm lens (used to create the narrow barrier) is mounted very close to the science cell and so it is impractical to measure the barrier power directly after it. Instead, we measure the power before the lens and use an independent measurement of the transmission of the lens at 532 nm ($T = 0.48\pm0.01$) to make the compensation. We determine the transmission of the science cell to be 0.96 ± 0.04, which is consistent with the 0.92 transmission expected with an uncoated glass slab.

5.3.2 Determining barrier widths

Similar to the practical difficulties in determining the barrier power, profiling the barrier beam at the atoms using the standard methods of knife-edging or direct imaging are impossible. Therefore, we have developed several other methods for beam profiling, some of which are performed outside the vacuum chamber and others that use the physical response of the atoms to the barrier potential as a measuring device.

Gaussian beam propagation

One method of deducing the barrier widths is to profile the beam at various stages of the accessible parts of the optical setup before the science cell and use the known expressions for Gaussian beam propagation (Eq. 5.6) and beam propagation through lenses to calculate the expected beam width at the atoms. The beam path for the 30 mm lens is shown in Fig. 5.8, where a piece of 2 mm thick glass has been inserted to simulate the science cell. This figure clearly shows the effect of the cylindrical lens: the unfocussed direction goes on to produce the narrow, axial barrier width, whereas the focussed direction is smaller at the objective lens and so produces the wide, transverse barrier width at the atoms.

However, this method is largely unreliable as it neglects any aberrations or imperfections in the optics or input beam, giving only idealised diffraction-limited values. Therefore, the calculated prediction of 2.5 μm for the axial barrier waist is likely an underestimate of its true value. Furthermore, as the narrow beam width diverges strongly and there is no way of ascertaining the exact position of the waist relative to the atoms with this method, its predictions are too naive to be more than a rough guide.

Figure 5.8: Simulated beam paths and measurements for the axial and transverse dimensions of the *narrow* barrier beam. The black data points are beam profile measurements that are used to constrain the optical parameters and determine approximate beam dimensions at the atoms. The vertical dotted line shows the focus position of the axial beam dimension, highlighting the 10 mm offset in focus position for the transverse direction.

Direct beam profiling

The next step towards reliably measuring the beam waists is to create an identical setup of the barrier system outside of the vacuum chamber, using duplicate optics. From this we can perform direct measurements on a beam that should be very similar to the barrier beam inside the chamber.

For the wide barrier, images of the beam are taken close to the axial direction's waist, using a standard CMOS camera (Thorlabs DCC1545M). Fitting Gaussian distributions to cross-cuts taken from the centre of these images, we ascertain an axial waist of $(10.6 \pm 0.1)\,\mu\text{m}$ and a transverse width of $(434 \pm 5)\,\mu\text{m}$ at the axial waist position. To account for uncertainties in alignment and variations of optics we increase the upper limit of the axial waist uncertainty to give $(10.6\,^{+0.5}_{-0.1})\,\mu\text{m}$. Equivalent measurements for the narrow barrier are found to be below the $5.2\,\mu\text{m}$ pixel size of the camera. Therefore, a higher resolution camera is needed for an accurate measurement of the narrow barrier widths.

To profile the narrow barrier beam, we use an automated beam profiler that is based around a Pi NoIR CMOS camera. Importantly, this camera has a $1.12\,\mu\text{m}$ pixel size and so should be able to resolve narrow barrier waist, even at the $2.5\,\mu\text{m}$ prediction made by idealised system above. Fig. 5.9 shows an example image, along with the extracted beam propagation profiles. To determine a width on this profile, the raw image is first summed along the beam's axial dimension (Fig. 5.9[b]). This allows us to locate the centre of the beam which, so long as we have aligned the barrier correctly, is the section of the beam that the atoms interact with. We then take the strip of the central 10 pixels from the raw image and sum along the long axis of the beam (Fig. 5.9[c]). This gives an intensity profile of what the barrier beam looks like along the waveguide's axial direction. The width of the beam is determined from a Gaussian fit and this process is repeated for various positions along the barrier beam path. We fit Eq. 5.6 to the measured widths in order to calculate a waist, which gives $(3.91 \pm 0.04)\,\mu\text{m}$. Switching the order of vertical/ horizontal sums and performing the same routine also allows us

to determine the width of the transverse direction across this range. We measure a transverse width of $(118.3 \pm 0.3)\,\mu m$ at the axial waist position.

Figure 5.9: Beam profiling the barrier beam outside of the vacuum chamber using duplicate optics. A horizontal pixel sum [b] is performed on the raw image [a] to locate the centre of the beam. A vertical sum of the raw image is taken [c], over the central ten pixels (shown by the green shaded region in [b]). From this we determine the barrier width. By sampling the beam at various points along the propagation direction, we can locate and measure the axial direction waist [d] and associated transverse width [e].

Atom depletion

For the narrow barrier, we also determine the barrier width by directly imaging the imprint left by the barrier on an atomic cloud's density profile. A thermal cloud of 2×10^5 atoms is released into the quasi-1D waveguide trap with trapping frequencies $(\omega_r, \omega_z, \omega_{vert}) = 2\pi \times (80, 1.4, 80)$ Hz. The axial harmonic trap centre position is set so that it approximately aligns with the barrier position. Initially, the barrier beam is switched off and the atoms are allowed to freely expand and propagate along the waveguide. After allowing the cloud to expand for quarter trapping period (the point at which it

110

is widest) the barrier is pulsed on and the cloud is imaged *in-situ*. Example images are shown in Fig. 5.10 and clearly demonstrate the effects of the repulsive dipole force acting on the atoms in the depleted region.

Figure 5.10: Determining the width of the narrow barrier using an atom depletion method. For low powers the depth of the depleted region varies with power, but the width does not. The converse is true for higher powers, where a region of atoms has been completely depleted. The barrier beam width is determined by the average and standard deviation of the measured widths in the partial depletion region, shown by the solid green line and hatched area respectively. The upper panel shows example images of a thermal cloud after the barrier beam has been exposed onto it for 1 ms, for barrier powers ranging from 2.64 mW (far left) to 23.72 mW (far right).

By varying the power of the barrier beam, using the calibrated servo discussed in Sec. 5.3.1, we can identify two different regions of atom depletion: partial depletion and total depletion. For high powers ($P \gtrsim 8$ mW in this case) the atomic cloud is completely depleted in the central part of the barrier beam. Increasing the power further will not change the amplitude of the depletion *dip*, as this is determined by the atomic density in the centre of the atomic cloud. Instead, increasing the power beyond ~ 8 mW results in a broadening of the depleted region, as more of the Gaussian beam passes above a critical intensity that causes total depletion, hence the fitted Gaussian widths in

Fig. 5.10 are erroneously large.

For lower powers ($P \lesssim 8\,\text{mW}$) the atomic cloud is only partially depleted in the centre of the barrier beam. Reducing the power still further changes the dip amplitude but not the dip width, as fewer atoms from the central region of the barrier are displaced but the barrier width remains constant. Therefore, fitted widths of the density dip in this partial depletion region are a true reflection of the imprint left by the barrier, hence the true profile of the barrier beam.

The particular powers required for partial and total depletion in Fig. 5.10 are dependent on the pulse duration, as the barrier power determines the (anti) trapping frequency but the pulse duration sets how long the atoms have to be partially or fully expelled from the depletion region. For the range of barrier powers explored in Fig. 5.10, the trapping frequency ranges from 130 Hz to 1.25 kHz. Therefore, a 1 ms pulse duration was chosen to allow the full transition from no depletion to total depletion to be observed.

To profile the transverse width of the barrier beam, which is essential for calculating the barrier height, we use an acousto-optic deflector (AOD) placed in the beam path of the waveguide. The horizontal position of the waveguide in the science cell is controlled by changing the driving frequency provided to the AOD. This shifts the waveguide, hence the atomic cloud, along the wide dimension of the barrier. Starting with a barrier power in the partial depletion region, we measure the dip depth for multiple cloud positions. As we image the atoms using the vertical imaging system, we determine the exact position of the cloud on each image. We recover the expected Gaussian profile of the transverse dimension of the barrier beam (see Fig. 5.11).

This method is also very useful for precisely focussing the barrier, such that the axial waist overlaps vertically with the centre of the waveguide. As the intensity of the barrier beam is inversely proportional to w_z, focussing is achieved by maximising the amplitude of the density dip. Note that the beam's transverse width does not change appreciably over this range (see Fig. 5.9) and so does not significantly alter the intensity, thus it can be

Figure 5.11: Measuring the transverse (wider) direction of the barrier beam using the atom depletion method. The transverse position of the atoms is relative to the normal position of the waveguide, thus the barrier was slightly misaligned in this example.

ignored for the purposes of this optimisation. The transverse width does not vary much in this range because it has a much longer Rayleigh length than the axial width and the waist position is $\sim 10\,\text{mm}$ away from the waist position of the axial dimension.

Through the atom depletion method, we determine an axial barrier width of $(4.7 \pm 0.3)\,\text{µm}$ and a transverse barrier width of $(117 \pm 9)\,\text{µm}$. However, as discussed in Sec. 3.9, our imaging system is limited to a resolution of $(3.0 \pm 0.3)\,\text{µm}$. As this is comparable to the length scale of the axial barrier width, it has to be accounted for. For a measured width of w_{meas}, the true width is given by:

$$w_{\text{true}} = \sqrt{w_{\text{meas}}^2 - r^2}, \tag{5.9}$$

which gives an axial barrier width of $(3.6 \pm 0.4)\,\text{µm}$. The transverse barrier width remains at $(117 \pm 9)\,\text{µm}$. This agrees well with the direct offline measurement of $(3.91 \pm 0.04)\,\text{µm}$.

Throughout the remainder of this thesis we use the following width values for all calculations and simulations. For the wide barrier, the axial and transverse widths are $(10.6\,^{+0.5}_{-0.1})\,\text{µm}$ and $(434 \pm 5)\,\text{µm}$ respectively. For the narrow

113

barrier, the axial and transverse widths are $(3.6 \pm 0.4)\,\mu\mathrm{m}$ and $(117 \pm 9)\,\mu\mathrm{m}$ respectively.

5.4   Controllable splitting

To set the transmission of a soliton through the barrier, control over both the barrier potential and soliton kinetic energy is required, as shown in Eq. 5.4. In the previous sections we have highlighted the high level of control available with our experimental setup for both of these parameters independently. Here we show how the relationship between the barrier and the kinetic energy of the soliton allows controllable splitting of a soliton.

5.4.1   Soliton kinetic energy

The centre of mass kinetic energy of the soliton is determined by the centre of mass velocity, $E_{\mathrm{k}} = \frac{1}{2}mv^2$. As the soliton is in a harmonic potential, it undergoes harmonic oscillations:

$$z(t) = z_0 \cos\left(\omega_{\mathrm{z}} t\right). \tag{5.10}$$

Therefore, the velocity is also time varying:

$$v(t) = -z_0 \omega_{\mathrm{z}} \sin\left(\omega_{\mathrm{z}} t\right) = -v_{\mathrm{max}} \sin\left(\omega_{\mathrm{z}} t\right), \tag{5.11}$$

where $v_{\mathrm{max}} = z_0 \omega_{\mathrm{z}}$ is the maximum velocity of the soliton, which occurs at the centre of the harmonic potential. Consequently, the kinetic energy of the soliton at the point of collision with the barrier will differ slightly if the barrier is offset from the centre of the trap.

As it is experimentally challenging to create a perfectly-stable, perfectly-aligned barrier setup, we instead amend the kinetic energy to account for a barrier offset:

$$E_{\mathrm{k}} = \frac{1}{2}mv(t)^2|_{t=t_{\mathrm{barr}}}, \tag{5.12}$$

where t_barr is the time of the soliton-barrier collision. For a barrier that is offset z_off from the centre of a harmonic potential, a soliton with $|z_0| > |z_\text{off}|$ will reach the centre of the barrier at:

$$t_\text{barr} = \frac{1}{\omega_\text{z}} \cos^{-1}\left[\frac{z_\text{off}}{z_0}\right].$$

(5.13)

Fortunately, the effects of an offset barrier on the soliton velocity become vanishingly small for small offsets, as found by substituting Eq. 5.13 into Eq. 5.11 to find the soliton velocity at the barrier v_barr:

$$v_\text{barr} = -v_\text{max} \sin\left(\cos^{-1}\left[\frac{z_\text{off}}{z_0}\right]\right) = -v_\text{max}\sqrt{1 - \left[\frac{z_\text{off}}{z_0}\right]^2}.$$

(5.14)

Therefore, $v_\text{barr} \approx v_\text{max}$ for $z_\text{off} < z_0$. For example, a fractional offset of $\frac{z_\text{off}}{z_0} = 0.02$ results in a fractional velocity change of only 0.0002 (0.02 %). These values are typical of our experimental uncertainties, as will be described in Sec. 6.2.

Experimentally, the position of the barrier is found using the atom depletion method from Sec. 5.3.2. The quasi-1D harmonic potential, which defines the kinetic energy in Eq. 5.12, is found by tracking trajectories of solitons oscillating in the potential, as in Sec. 4.4.

5.4.2 Splitting a soliton

We split the soliton by allowing it to accelerate towards and interact with the repulsive Gaussian barrier (see Fig. 5.12). Crucially, the split components are found to retain the non-dispersive characteristics of solitons, over the observed $\sim 10\,\text{s}$ timescales. Therefore, we attribute each split component to be a *daughter soliton* in its own right. This is also expected from theoretical studies [22, 60–66, 119–122].

Figure 5.12: A sequence of images taken to show the splitting of a soliton by the narrow barrier. The barrier height is set to give $\sim 50\%$ transmission. The upper panel shows the column density sum across the final image. Each image is taken from a single independent experimental run.

Importantly, the splitting fraction (i.e. the number of atoms in each of the reflected and transmitted daughter solitons) can be controlled by varying the barrier transmission, as shown in Sec. 5.2. This is achieved experimentally by varying the barrier beam power, which directly alters the barrier height as the width is fixed by the optical setup (Fig. 5.13).

To determine the shot-to-shot stability in the barrier transmission, data is compiled from many splitting events that were initially calibrated to give 50% transmission (Fig. 5.14). We find both barriers to have similar uncertainties in transmission, with standard deviations of $(8.5 \pm 0.6)\%$ and $(6.4 \pm 0.4)\%$ for the wide and narrow barriers respectively.

In Sec. 6.2, we will determine the shot-to-shot fluctuation in barrier position relative to the centre of the harmonic potential to be $1.3\,\mu\text{m}$. For a soliton in a $1.3\,\text{Hz}$ harmonic potential, initially displaced $50\,\mu\text{m}$ away from the narrow barrier, this position uncertainty produces a $\sim 0.05\,\text{nK}$ uncertainty on a $\sim 0.85\,\text{nK}$ kinetic energy at the barrier. This set of parameters is chosen

Figure 5.13: Controllable splitting of a soliton into two daughter solitons by a repulsive Gaussian barrier. [a] The transmission of a slow (blue) and fast (red) soliton through the wide barrier as the barrier power is varied, with kinetic energies $E_k/k_B = (15.5 \pm 0.3)\,\text{nK}$ and $E_k/k_B = (41 \pm 1)\,\text{nK}$ respectively. [b] The transmission through the narrow barrier, for solitons with kinetic energies at the barrier $E_k/k_B = (0.84 \pm 0.08)\,\text{nK}$ (blue) and $E_k/k_B = (16.8 \pm 0.6)\,\text{nK}$ (red). The solid lines are quasi-1D GPE simulations that are fit to the data by varying the barrier width.

Figure 5.14: The distribution of splitting fractions across many experimental runs with a target transmission of 50 %. The height of each bar indicates the probability of occurring in that bin. [a] and [b] show the distributions for the wide and narrow barriers, with standard deviations of $(8.5 \pm 0.6)\,\%$ and $(6.4 \pm 0.4)\,\%$ respectively. The red lines are Gaussian fits to the data.

to match that of the slow soliton in Fig. 5.13[b], such that we can use the gradient of the transmission curve $((-0.60 \pm 0.07)\,\mathrm{nK}^{-1})$ to find an associated transmission uncertainty: $\sim 3\,\%$. Therefore, the uncertainty in barrier offset accounts for a large fraction of the uncertainty in transmission seen in Fig. 5.14, but cannot account for it all.

Fitting with 1D-GPE

Further to the methods described in Sec. 5.3, we also measure the barrier widths by fitting quasi-1D GPE theory to transmission curves, as seen in Fig. 5.13. The axial barrier width is the only free parameter for these fits; the transverse direction is constrained by the other beam profiling methods. Using this technique, we determine barrier widths of $w_{\mathrm{z}} = (4.8 \pm 0.2)\,\mu\mathrm{m}$ and $w_{\mathrm{z}} = (11.9 \pm 0.3)\,\mu\mathrm{m}$ for the narrow and wide barrier beams respectively.

The uncertainty in each value is the standard error in the fitted widths across several sets of splitting data, taken for various soliton velocities.

The discrepancy between these values and the smaller values measured experimentally in Sec. 5.3 is likely due to the barrier beam not being of perfect Gaussian profile. As shown in Ref. [33], the presence of relatively small diffraction fringes in the wings of the beam profile of the barrier can greatly enhance reflection. These fringes are not included in the GPE theory and so may cause an erroneous broadening of the Gaussian beam profile fit. Therefore, we continue to use the experimentally-measured widths of $(10.6 \, ^{+0.5}_{-0.1}) \, \mu\mathrm{m}$ and $(3.6 \pm 0.4) \, \mu\mathrm{m}$, for the wide and narrow barrier respectively.

5.5   Quantum versus classical splitting

The nature of the splitting mechanism depends critically on the barrier width. In the limit of a δ-function barrier, quantum tunneling dominates and the barrier *area* determines the transmission probability [65]. However, for barriers wider than the soliton width, the transmission probability instead depends primarily on the incident centre of mass kinetic energy of the soliton relative to the barrier *height*. Therefore, for an infinitely-wide barrier, where the transmission coefficient curve is a step function, 50 % transmission is achieved when the kinetic energy per atom in the incident soliton exactly matches the barrier height.

For our wide barrier, quasi-1D GPE simulations yield 50 % transmission when the barrier height is only 1 % higher than the kinetic energy of the soliton at the barrier, implying that the splitting mechanism is almost entirely a classical process. However, equivalent simulations for the narrow barrier yield a barrier height that is 11 % higher than the kinetic energy, indicating that quantum tunneling plays a small role.

Experimentally, we measure 50 % transmission when the barrier height is $(11 \pm 2) \, \%$ and $(26 \pm 12) \, \%$ higher than the soliton kinetic energy at the

barrier, for the wide and narrow barriers respectively (Fig. 5.15). These calculations depend on the barrier widths determined in Sec. 5.3 to convert beam power to energy, using Eq. 5.7.

Figure 5.15: Barrier height for 50% splitting, as a function of the centre of mass kinetic energy per atom at the barrier interaction. The red (blue) points and lines show the data and fit for the wide (narrow) barrier. These fits give gradients of 1.11 ± 0.02 and 1.26 ± 0.12, for the wide and narrow barriers respectively. Both axis are scaled by k_{B} to give the energy in nK. The dashed line shows $V_0 = E_{\mathrm{k}}$ which, as shown in Fig. 5.4, is the classical limit.

As suggested in Sec. 5.4.2, the quantitative discrepancy between theory and experiment is likely due to a non-Gaussian beam profile. It may also be caused by remnant thermal atoms inflating our measurement of the transmitted fraction, as thermal atoms are more likely to be transmitted than condensate atoms. Nevertheless, our measurements verify that quantum tunneling is more relevant for the narrow barrier than the wide barrier, though classical splitting remains the dominant effect in both cases.

5.6 Velocity asymmetry

Interestingly, the steep energy dependence of the measured transmission functions (Figs. 5.3 and 5.13) produces a velocity filtering effect, whereby the transmitted daughter soliton always has a higher centre of mass kinetic energy than the reflected daughter soliton [22,63,119]. We define this asymmetry to be the ratio between the velocities of the transmitted to reflected solitons, namely:

$$\zeta = v_{\mathrm{T}}/v_{\mathrm{R}}. \tag{5.15}$$

More precisely, v_{T} and v_{R} are the maximum velocities of the daughter solitons (their velocity in the centre of the harmonic potential). If the frequency is fixed, Eq. 5.15 is equivalent to:

$$\zeta = z_{0,\mathrm{T}}/z_{0,\mathrm{R}}, \tag{5.16}$$

where $z_{0,\mathrm{T}}$ and $z_{0,\mathrm{R}}$ are the centre of mass oscillation amplitudes for the transmitted and reflected daughter solitons respectively. This difference in amplitude of the centre of mass oscillation between the daughter solitons is directly observable (Fig. 5.16).

5.6.1 Kinetic energy distribution of a soliton

The velocity asymmetry can be understood in terms of an initial kinetic energy distribution of the parent soliton (Fig. 5.17). For 50 % transmission at the barrier, the part of the soliton with an energy spectrum above (below) 50 % level of the transmission function has a higher probability of being transmitted (reflected). Therefore, the transmitted daughter soliton will have a higher centre of mass kinetic energy than the reflected daughter soliton.

It is worth commenting on the origin of the energy distribution of the parent distribution, as it may appear non-trivial that a condensate at close to $0\,\mathrm{K}$ should possess one. A quantum mechanical interpretation may be realised

Figure 5.16: Velocity selection during the splitting of a soliton. The initial soliton has ~ 2500 atoms at a scattering length of $\sim -8.5\,a_0$. The transmitted daughter soliton (blue trajectory) leaves the wide barrier (green) with a higher kinetic energy than the reflected soliton (dotted red trajectory) and so oscillates with a clearly larger amplitude. The barrier height is set to give $\sim 50\,\%$ transmission and is removed after the splitting to allow the solitons to oscillate, as in sequence of the previous figure. The amplitudes of the transmitted and reflected solitons are $(290 \pm 10)\,\mu m$ and $(226 \pm 6)\,\mu m$ respectively, implying a velocity asymmetry of 1.28 ± 0.06. The barrier is offset from the centre of the $0.5\,Hz$ trap by $\sim 30\,\mu m$ in this case, although this is not expected to play a significant role in the observed velocity filtering. The black dashed line is the expected trajectory leading up to splitting, which was not recorded for this specific sequence. It is calculated by assuming kinetic energy is conserved at the barrier interaction, implying in an initial release position $\sim 250\,\mu m$ from the centre of the harmonic potential.

Figure 5.17: A sketch to describe the velocity selection that occurs when a soliton interacts with a barrier of finite width. [a] The parent soliton, with a kinetic energy distribution shown by the blue curve, interacts with a finite-width barrier, which has a transmission function shown by the red dotted curve. [b] The interaction results in a splitting of the parent soliton into two daughter solitons. As the transmission function varies appreciably across the width of the energy distribution of the soliton, the transmitted daughter soliton has a higher mean kinetic energy than the reflected daughter soliton (as seen by the blue curves).

by taking the Fourier transform of the soliton wavefunction (Eq. 1.8) and deriving an expression for the soliton width in momentum space:

$$k_{\mathrm{v}} = \frac{2\omega_{\mathrm{r}}|a_{\mathrm{s}}|N}{\pi}. \tag{5.17}$$

Therefore, a soliton with $N = 2000$, $a_{\mathrm{s}} = -10\,\mathrm{a_0}$ and $\omega_{\mathrm{r}} = 2\pi \times 40\,\mathrm{Hz}$ has a velocity distribution of width $\sim 0.2\,\mathrm{mm\,s^{-1}}$. For a soliton released $100\,\mu\mathrm{m}$ away from the centre of a $1.5\,\mathrm{Hz}$ potential, this produces a kinetic energy distribution at the centre of the harmonic potential of width $\sim 2\,\mathrm{nK}$.

One may be tempted to consider an alternative, classical, interpretation. Specifically, consider the soliton to be an ensemble of particles with some spatial extent. As the harmonic potential has a curvature, different particles will begin their centre of mass harmonic oscillations with different potential energies, producing an energy distribution. For a soliton of spatial width $3\,\mu\mathrm{m}$ released into the same trap as before, atoms released at $103\,\mu\mathrm{m}$ away will be $\sim 0.03\,\mathrm{mm\,s^{-1}}$ faster than those that begin in the centre of the soliton

once they reach the centre of the harmonic potential. This produces a kinetic energy distribution at the centre of the harmonic potential of width $\sim 0.3\,\mathrm{nK}$.

Therefore, as the quantum mechanical interpretation of the energy distribution is ~ 7 times wider than the classical interpretation in our experiment, we can largely neglect the curvature of the trap in our system. This is an important distinction to make because the quantum and classical interpretations have drastically different implications for properties of the velocity asymmetry: a wider soliton in the classical picture produces a wider energy distribution, whereas a wider soliton in the quantum picture results in a narrower energy distribution. Therefore, a wider soliton creates a larger kinetic energy difference between the daughter solitons in the classical picture, but a smaller difference in the quantum picture.

5.6.2 Transmission gradient

Another important contributor to the magnitude of the asymmetry is the gradient of the transmission function. As shown in Sec. 5.2, the gradient of the transmission function can be altered by changing the width of the barrier. For example, the transmission curves at $E_\mathrm{k}/k_\mathrm{B} = (15.5 \pm 0.3)\,\mathrm{nK}$ in Fig. 5.13[a] and $E_\mathrm{k}/k_\mathrm{B} = (16.8 \pm 0.6)\,\mathrm{nK}$ in Fig. 5.13[b] have peak gradients of $(-0.15 \pm 0.01)\,\mathrm{nK}^{-1}$ and $(-0.11 \pm 0.01)\,\mathrm{nK}^{-1}$ respectively, despite splitting solitons with nearly the same kinetic energy. We explore the consequences of this for velocity selection in Fig. 5.18.

For very narrow barriers, the transmission function does not discriminate strongly when it comes to kinetic energy, hence the effects of velocity selection are minimal. However, as we saw in Sec. 5.2: the wider the barrier, the steeper the transmission function. Consequently, increasing the barrier width leads to a larger asymmetry between the kinetic energies of the outgoing daughter solitons, as the effects of velocity selection are exaggerated. In other words, the higher-velocity part of the momentum-space wavefunction of the
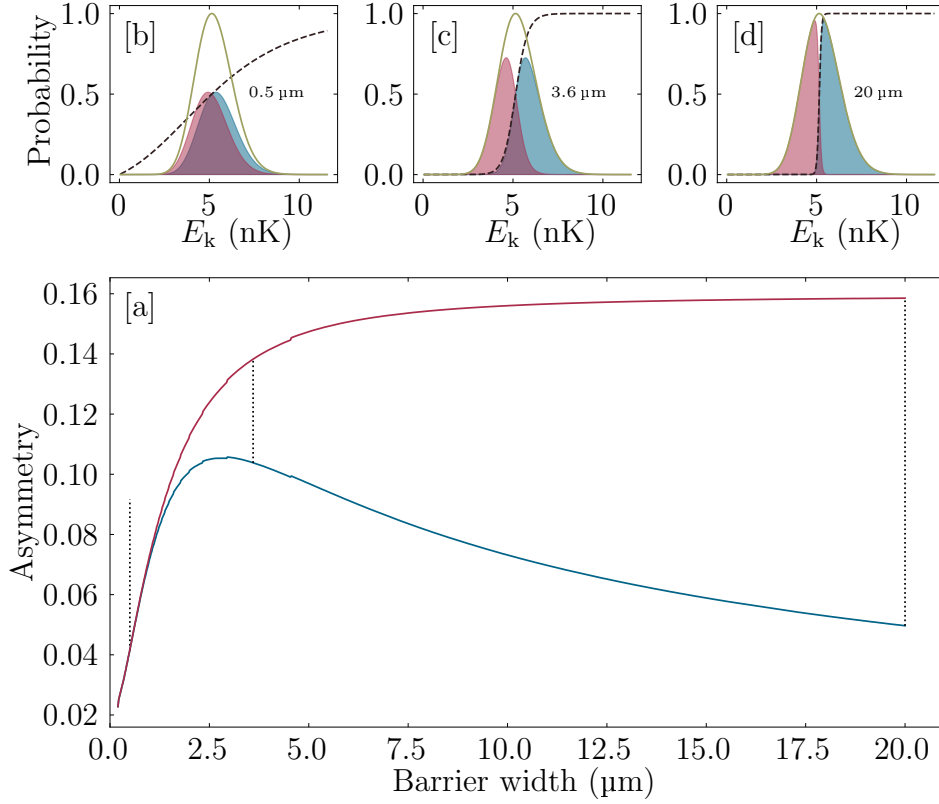
Figure 5.18: The asymmetry in velocity between the daughter solitons, as a function of the barrier width. The red curve in [a] shows the asymmetry in mean kinetic energy between the outgoing daughter solitons after the first barrier interaction. The red (blue) curve shows the asymmetry in kinetic energy between the outgoing daughter solitons after the first barrier interaction, assuming they re-form around the mean (peak) of the momentum space wavefunctions of the split solitons. The barrier height is adjusted at each barrier width to ensure 50 % splitting, as in Fig. 5.4. Figures [b], [c] and [d] show the transmission functions calculated using Eq. 5.5 (black dashed) and momentum-space wavefunctions of the initial, transmitted and reflected solitons (green, blue and red respectively). The barrier widths of [b], [c] and [d] are 0.5 μm, 3.6 μm and 20 μm, as indicated in [a] by the vertical dotted lines, from left to right respectively. The centre of mass kinetic energy and kinetic energy distribution of the initial soliton are constant throughout to represent a centre of mass velocity of $1\,\mathrm{mm\,s^{-1}}$ and width of $0.2\,\mathrm{mm\,s^{-1}}$.

soliton becomes far more likely to be transmitted than the lower-velocity part, causing it to have a higher mean kinetic energy.

However, for barrier widths larger than the soliton width, the expected nature of the asymmetry-width dependence is not as clear. Intuitively, one may expect a monotonic increase in asymmetry with barrier width, as the increasing sharpness of the transmission function must result in an increase of the difference between the mean kinetic energy of the daughter solitons (shown by the red curve in Fig. 5.18[a]). However, as we will see in Sec. 6.4.3, GPE simulations predict that the asymmetry actually *decreases* for widths beyond the soliton width (Fig. 6.16). The reasons for this are not entirely clear, however it may be that the daughter solitons actually re-form around the *peak* of the momentum-space wavefunction during nonlinear evolution, not the *mean* value one might first expect. This produces the initial increase and subsequent decrease in asymmetry (blue curve in Fig. 5.18[a]), as predicted by GPE simulations.

5.7 Velocity asymmetry experiments

In practice, we have three levers for altering the width of the soliton: the radial trapping frequency, the scattering length and the number of atoms (Eq. 5.17). In this section we explore the observable effects of varying the scattering length and barrier width, as well as the kinetic energy of the soliton.

5.7.1 Extracting the asymmetry parameter

To calculate the asymmetry parameter for the trajectories of some split solitons, such as in Fig. 5.16, the oscillation amplitudes of the daughter solitons must be extracted (see Eq. 5.16). Naively, one can fit the following equations for the transmitted and reflected daughter solitons respectively:

$$z_\mathrm{T}(t) = z_{0,\mathrm{T}} \cos\left(\omega_{\mathrm{z,T}} t\right) + z_\mathrm{cent,T}, \tag{5.18}$$

$$z_{\mathrm{R}}(t) = z_{0,\mathrm{R}} \cos\left(\omega_{z,\mathrm{R}}\left(t - 2t_{\mathrm{barr}}\right)\right) + z_{\mathrm{cent,R}}, \qquad (5.19)$$

where $z_{0,\mathrm{T}}$ and $z_{0,\mathrm{R}}$ are the amplitudes of the centre of mass oscillations, $\omega_{z,\mathrm{T}}$ and $\omega_{z,\mathrm{R}}$ are the frequencies of the centre of mass oscillations, and $z_{\mathrm{cent,T}}$ and $z_{\mathrm{cent,R}}$ are the fitted positions of the centre of the harmonic potential. The additional $2\omega_{z,\mathrm{R}}t_{\mathrm{barr}}$ phase shift in Eq. 5.19 accounts for a barrier offset (see Fig. 5.19).

Figure 5.19: A picture to show the splitting of a soliton in a harmonic potential. The reflected daughter soliton (red line) follows a trajectory which is out of phase with the transmitted daughter soliton by $2\omega_z t_{\mathrm{barr}}$.

However, there are clear redundancies in these equations, which will lead to unnecessary over-fitting and an artificial increase in the perceived uncertainty in ζ. Firstly, $\omega_{z,\mathrm{T}}$ and $\omega_{z,\mathrm{R}}$ should be identical as the daughter solitons oscillate in the same trap. Secondly, $z_{\mathrm{cent,T}}$ and $z_{\mathrm{cent,R}}$ should be identical, for the same reason. Finally, $z_{0,\mathrm{T}}$ and $z_{0,\mathrm{R}}$ are not independent, as they both depend linearly on z_0: the centre of mass amplitude of the initial soliton.

5.7.2   Lissajous curves

To avoid having redundant parameters, we instead look at the relationship between $z_T(t)$ and $z_R(t)$ directly. If the barrier is well-centred in the harmonic potential, the resultant lineshape of $z_R(t)$ against $z_T(t)$ is a perfectly straight line, as the daughter solitons are exactly π out of phase (Fig. 5.20[a]). However, if the barrier is off-centre (through imperfect or intentional misalignment) the resultant lineshape is an ellipse (Fig. 5.20[b]). This result can be explained by considering Eqs. 5.18 and 5.19 to be a set of parametric equations, which can be used to fit the observed ellipses.

Figure 5.20: Fitting Lissajous curves to the centre of mass positions of the transmitted and reflected daughter solitons as they undergo free oscillations, in order to extract the asymmetry parameter. [a] In the case that the barrier is centred in the harmonic potential, the daughter solitons are almost exactly π out of phase and the asymmetry parameter ζ is simply the gradient of the line. [b] In the case that the barrier is far from the centre of the harmonic potential, the daughter solitons are less than π out of phase and ζ can be understood as $\Delta z_T / \Delta z_R$, where Δz_T and Δz_R are the ranges in z_T and z_R swept out by the Lissajous curve respectively. In practice, ζ is found in both cases by fitting parametric equations (Eqs. 5.20 and 5.21).

As the fitting can now be performed with respect to the positions of both the transmitted and reflected daughter solitons simultaneously, we are able to

drop the redundant parameters that appear in Eqs. 5.18 and 5.19. This provides the following set of parametric equations in 2D Cartesian coordinates (x, y):

$$x(\theta) = z_{0,\mathrm{R}} \zeta \cos(\theta) + z_{\mathrm{cent}}, \qquad (5.20)$$

$$y(\theta) = z_{0,\mathrm{R}} \cos(\theta + \phi) + z_{\mathrm{cent}}, \qquad (5.21)$$

where $z_{0,\mathrm{T}}(t) = z_{0,\mathrm{R}} \zeta$ by definition, $\phi = 2\omega_{\mathrm{z}} t_{\mathrm{barr}}$, z_{cent} is the position of the centre of the harmonic potential and θ is the independent variable that defines the angle of the ellipse. Note: as we are only interested in the asymmetry, we do not independently fit for ω_{z} and t_{barr}, only ϕ. Formally, Eqs. 5.20 and 5.21 define a Lissajous curve with a fixed frequency and asymmetric amplitudes.

With these simplifications we have reduced the number of fitting parameters from at least 7 to a maximum of 4, ensuring we can get the best possible measure of the asymmetry from a given set of daughter soliton trajectories.

5.7.3 Experimental asymmetry

Applying this method to daughter solitons with a range of oscillation, barrier and internal parameters, we explore whether the properties of the kinetic energy asymmetry predicted in Sec. 5.6 can be observed experimentally. We study the effect of varying the scattering length on the asymmetry in Fig. 5.21, for the narrow barrier.

Whilst there is some apparent increase in the asymmetry for more negative scattering lengths, as we would expect for the reduced spatial width, the relationship is not very strong. This is due to varying centre of mass kinetic energies of the initial soliton between sets of experimental runs, which was not fixed. In Fig. 5.21 the kinetic energies range from $\sim 0.8\,\mathrm{nK}$ to $\sim 17\,\mathrm{nK}$. As described by Eq. 5.17, the width of the momentum distribution of the soliton does not depend on the centre of mass kinetic energy, in the GPE picture.

Figure 5.21: The kinetic energy asymmetry between the transmitted and daughter solitons (Eq. 5.15) as a function of scattering length, for the narrow barrier.

Therefore, the *absolute difference* in kinetic energy between the outgoing daughter solitons should be independent of the kinetic energy of the initial soliton, though the *fractional difference* (the asymmetry parameter) is not.

A further limitation of Fig. 5.21 is that the momentum-space width of the soliton is not only dependent on a_s, but also on the number of atoms. Recasting Eq. 5.17 in terms of N/N_crit, where N_crit is given by Eq. 1.10, we arrive at:

$$k_\mathrm{v} = \frac{2\omega_\mathrm{r} k a_\mathrm{ho}}{\pi} \frac{N}{N_\mathrm{c}}. \tag{5.22}$$

This transformation creates a single interaction parameter, N/N_crit, which includes both the scattering length and the number of atoms.

In Fig. 5.22, we observe a strong positive correlation between the absolute difference in velocity between the daughter solitons and N/N_crit, in good agreement with the momentum-space width of the soliton. Interestingly, we do not measure a difference between the wide and narrow barriers, which indicates that the splitting is classical in both cases. This supports the observations of Sec. 5.5 and is due to both barriers being relatively wide on

Figure 5.22: The absolute kinetic energy difference between the transmitted and daughter solitons, as a function of N/N_{crit}. The red and blue points represent the wide and narrow barriers respectively. The green line is the velocity width calculated directly from Eqs. 5.17 and Eq. 1.10, assuming constant axial and radial trapping frequencies of $2\pi \times 1.4\,\mathrm{Hz}$ and $2\pi \times 40\,\mathrm{Hz}$ respectively.

the scale of a $\sim 3\,\mu$m wide soliton, meaning that little difference is expected between them (Fig. 5.18).

Chapter 6

Soliton Recombination

In this chapter we describe the recombination of bright-solitary-matter waves on a narrow repulsive Gaussian barrier. We begin with an overview of similar systems that have made progress towards this goal, before providing details on the experimental sequence which allows us to realise soliton recombination on a narrow barrier. By exploring various barrier widths, offsets and heights, we uncover the impact of the velocity asymmetry on phase-sensitive recombination, before developing an intuitive understanding of the results using an approximate theoretical interpretation. Finally, we discuss the implications that our findings hold for any future practical soliton interferometry device.

6.1   Background

As described in Sec. 1.2, a soliton-based interferometer has the potential to bring advancements across a range of phase-sensitive measurements. Much like conventional optical interferometers, a soliton-based interferometer requires three main steps: phase-coherent splitting, phase accumulation time and phase-coherent interference-mediated recombination (Fig. 6.1). Following the recombination process, differences in atom number between the outgoing populations indicate the accumulated differential phase between the daughter solitons.

Figure 6.1: A soliton interferometry scheme based on a repulsive Gaussian barrier. [i] The initial soliton (blue) is split into two daughter solitons by the barrier (green). [ii] The daughter solitons then acquire a relative phase difference (ϕ). [iii] They return to the barrier and interfere, with resultant population fractions on the left and right (N_L and N_R respectively) determined by ϕ.

At the Gross-Pitaevskii level, a soliton incident on a barrier is split cleanly into transmitted and reflected daughter solitons, provided the incident velocity is sufficiently fast that the effects of interatomic interactions can be neglected during the splitting [125]. When these daughter solitons are subsequently made to spatio-temporally overlap at the barrier, total or partial interference-mediated recombination occurs, depending on their relative phase [66].

6.1.1 Current experiments

To date, the experiment at the ANU is alone in demonstrating phase-sensitive recombination with solitons [27]. As described in Sec. 5.1, a $\pi/2$-π-$\pi/2$ series of Bragg pulses allowed them to split and subsequently recombine their soliton, with the ability to resolve interference fringes for interferometer times below $\sim 3\,\mathrm{ms}$.

The repulsive barrier in the experiment at Rice has also been used to demonstrate recombination of solitons [123]. However, they were unable to measure any definitive phase-related effects. This was attributed to shot-to-shot noise being too large for the sensitivity of recombination to the interaction strength N/N_{crit}. As we will see later in this chapter, it is probable that their experiment was also in a region of parameter space where the velocity asymmetry heavily suppresses interference, such that interference effects would have been unlikely anyway.

As described in the previous chapter, we also use a repulsive Gaussian barrier setup. However, the reduced shot-to-shot uncertainty and increased tunability of our experiment may be sufficient to overcome the difficulties encountered by the Rice experiment.

6.2   Experimental overview

The experimental setup for recombination is identical to that of the splitting experiments. Briefly, we form a soliton of approximately 2500 ^{85}Rb atoms in a quasi-1D waveguide where an additional harmonic magnetic potential produces axial trapping. A narrow repulsive Gaussian barrier is focussed down to bisect the waveguide (Fig. 6.2). Upon reaching the barrier, the soliton is either reflected, transmitted, or split into two daughter solitons (Fig. 5.13). The barrier height and therefore the transmission through the barrier is tuned experimentally by varying the total barrier power.

Interference-mediated recombination is achieved by simply allowing the two daughter solitons to come back to the barrier and interfere with each other. Therefore, the only additional complexity to the recombination experiments is to wait longer between releasing the soliton into the waveguide and taking the image, relative to the splitting experiment. More precisely, the image is taken at time [iii] in Fig. 6.1, rather than time [ii].

Figure 6.2: A sketch of the experimental implementation of the barrier-based interferometer. The 3D plot and yz-plane highlight the optical waveguide (red) and the magnetic harmonic potential along the waveguide, respectively, with the barrier shown in green for both. The xz-plane shows an example image of a soliton in this potential.

6.2.1 Trap stability

As interference-mediated recombination relies on the interference between the two daughter solitons, it is imperative that they both arrive perfectly simultaneously at the barrier for the second barrier interaction. Therefore, the experiment has to be either insensitive to small misalignment, or sufficiently stable and calibrated to a high-enough precision, so that the daughter solitons will always perfectly spatio-temporally overlap at the barrier. There are two possibilities for ensuring one of these conditions: either the solitons should be allowed to perform a full oscillation between splitting and recombination (Fig. 6.3), or the barrier should be exactly centred in the harmonic potential.

In the first case, shown in Fig. 6.3, the barrier is removed after splitting to allow each daughter soliton to undergo a full centre of mass oscillation before interacting with the barrier a second time. This guarantees that the

Figure 6.3: A sequence of images showing the splitting of a soliton into two daughter solitons, which each undergo a full centre of mass harmonic oscillation before interacting with the barrier for a second time. The barrier power is the same at the first and second barrier interactions, giving 50 % transmission at the first barrier interaction.

solitons will each experience exactly the same potential and will reach the barrier at exactly the same time, assuming the trap is stable over a trapping period. However, this method has the severe limitation that the daughter solitons will not develop a relative phase shift, unless a strongly time-varying interaction is present, as they each explore the same full harmonic potential. Therefore, we focus instead on the second method.

To set the barrier to be exactly in the centre of the harmonic potential, it is essential that we understand the stability of the trap and the barrier position. We determine the shot-to-shot fluctuation in the position of the centre of the axial potential by taking ten repeat measurements of the soliton position both immediately after release and after half a trap period, finding shot-to-shot standard deviations of $(0.3 \pm 0.1)\,\mathrm{\mu m}$ and $(2.6 \pm 0.6)\,\mathrm{\mu m}$ respectively. These variations imply shot-to-shot uncertainty in the position of the centre of the harmonic potential of $(1.3 \pm 0.3)\,\mathrm{\mu m}$. As the uncertainty is dominated by position fluctuations after a half trapping period, we attribute it to be dominated by magnetic potential instability. By measuring the position of the harmonic potential as the magnetic field of the E-W shims is varied, we determine that a stray field of only $\sim 3.0\,\mathrm{mG}$ along the axial direction would account for this shot-to-shot fluctuation (Fig. 6.4). Using the atom depletion method of Sec. 5.3 and another set of ten images, we determine the barrier position to fluctuate from shot-to-shot with a standard deviation of $(0.3 \pm 0.1)\,\mathrm{\mu m}$.

6.3   Recombination

In the limit of a δ-function barrier, theoretical studies of soliton splitting indicate that there is an intrinsic $\pi/2$ phase difference between the daughter solitons [66]. In our harmonic potential, it is expected that this phase difference is maintained and that we should ideally achieve completely constructive (destructive) interference on the right (left) of the barrier, resulting in a fully recombined soliton appearing on the *right*. However, velocity filtering confounds this ideal outcome.

Figure 6.4: The position of a soliton after half a trapping period, as a function of the applied magnetic field from the E-W shims. This shows the sensitivity of the harmonic potential along the waveguide to axial magnetic fields. The gradient of the straight line fit is $(0.876 \pm 0.009)\,\mu\text{m}\,\text{mG}^{-1}$, which implies that the centre of the harmonic potential moves $(0.438 \pm 0.004)\,\mu\text{m}\,\text{mG}^{-1}$. Therefore, a stray field of only $\sim 3.0\,\text{mG}$ is sufficient to explain the $(1.3 \pm 0.3)\,\mu\text{m}$ trap instability.

6.3.1 Velocity filtering

If we remove the effects of interference and consider velocity filtering alone, the reflected (transmitted) daughter soliton is always primarily reflected from (transmitted through) the barrier at the second barrier interaction, resulting in a single soliton appearing on the *left*. This can be easily understood in terms of the kinetic energy difference seen in Fig. 5.17. As the reflected (transmitted) daughter soliton always leaves the first barrier interaction with a lower (higher) centre of mass kinetic energy than the initial soliton, it has a lower (higher) transmission coefficient at the second barrier interaction, for the same barrier height. The joining of two solitons into one in this case should be considered a *merging* of the two daughter solitons, rather than true *recombination*, as it is mediated by classical velocity filtering and not interference.

To isolate and expose the effects of velocity filtering experimentally, the barrier is intentionally offset from the centre of the harmonic potential, pre-

venting the daughter solitons from spatio-temporally overlapping during the second barrier interaction (Fig. 6.5 and Fig. 6.6 for the wide and narrow barrier respectively). This prevents any possibility of interference.

Figure 6.5: Trajectory of a soliton interacting twice with the wide barrier (green), which is offset by $10\,\mu\text{m}$ from the centre of the harmonic potential. The barrier power is held constant throughout to give $50\,\%$ transmission at the first barrier interaction. The upper panels are vertical sums over the final image in the sequence. [b] is a quasi-1D GPE simulation of [a].

Figs. 6.5 and 6.6 show that the population appears almost entirely on the left after the second barrier interaction, for both barrier widths, which is consistent with strong velocity filtering. We observe good agreement between experiment and theory, in both the form of the final density distributions and in the total populations on each side of the barrier after the second barrier interaction: $(32 \pm 1)\,\%$ $((24 \pm 4)\,\%)$ and $24\,\%$ $(24\,\%)$ on the right of the narrow (wide) barrier, for the experimental data and theory respectively. The width of the barrier makes little difference to the final populations, implying that the velocity filtering effects are similar for both barrier widths. This corroborates the findings of the direct velocity asymmetry measurements in

Figure 6.6: Trajectory of a soliton interacting twice with the narrow barrier (green), which is offset by $10\,\mu m$ from the centre of the harmonic potential. The barrier power is held constant throughout to give $50\,\%$ transmission at the first barrier interaction. The upper panels are vertical sums over the final image in the sequence. [b] is a quasi-1D GPE simulation of [a].

Sec. 5.7.

Another interesting feature of the offset barrier trajectories is the soliton-soliton collision that occurs after the second barrier interaction. In the 1D-GPE simulation (see Fig. 6.6[b]) there is a clear repulsive elastic interaction between the solitons which pushes the initially reflected (transmitted) daughter soliton back away from (towards) the barrier. However, within a 3D-GPE framework, the additional transverse degrees of freedom mean that the soliton-soliton collision may not be purely elastic and population transfer is possible. Therefore, this technique may prove to be an excellent method of producing low-velocity soliton-soliton collisions with a well-defined initial phase, which is of intrinsic interest [41]. The collisional velocity of Fig. 6.6 can be found by calculating the velocities of the transmitted and daughter solitons at the collision point using Eq. 5.14, with the barrier offset z_{off} replaced by the collision position and z_0 replaced by the centre of mass oscillation amplitudes for the transmitted and reflected solitons respectively. This results in a collisional velocity of $\sim 0.25 \, \text{mm s}^{-1}$, which is far slower than any we have achieved with the collision occurring in the centre of the harmonic potential, such as in Figs. 5.16 and 6.3. To achieve the same relative velocity between daughter solitons in the centre of the harmonic potential, the oscillation amplitude would have to be $z_0 \sim 15 \, \mu\text{m}$, which is impractical in the current experimental configuration.

6.3.2 Interference-mediated recombination

To study interference-mediated recombination, the barrier is aligned with the centre of the harmonic potential to ensure maximal spatio-temporal overlap of the daughter solitons with the barrier. For the wide barrier, we observe that the majority of the population still appears on the *left* after the second barrier interaction (Fig. 6.7), suggesting that velocity filtering dominates the recombination process.

The outcome for the narrow barrier can be markedly different, as shown in Fig. 6.8. In this case, we can clearly see that the majority of the population

Figure 6.7: Trajectory of a soliton interacting twice with the wide barrier (green), which is centred in the harmonic potential. The barrier power is held constant throughout to give 50 % transmission at the first barrier interaction. The upper panels are vertical sums over the final image in the sequence. [b] is a quasi-1D GPE simulation of [a].

is *able* to finish on the *right*, which can *only* be explained by interference-mediated recombination. This occurs despite the presence of measurable velocity filtering effects (Fig. 6.6).

Figure 6.8: Trajectory of a soliton interacting twice with the narrow barrier (green), which is centred in the harmonic potential. The barrier power is held constant throughout to give 50 % transmission at the first barrier interaction. The upper panels are vertical sums over the final image in the sequence. [b] is a quasi-1D GPE simulation of [a].

Large shot-to-shot fluctuations in the final populations are observed when recombining on the centred narrow barrier, as seen across the final five images of Fig. 6.8. These are due to the finite shot-to-shot stability of the barrier position with respect to the minimum of the harmonic potential, as we discuss in detail later in this section. To observe the full extent of the shot-to-shot fluctuations, 15 repeats at 720 ms were taken (Fig. 6.9).

As detailed in Sec. 5.6, narrower barriers produce a reduced velocity filtering effect. However, in reality neither complete interference nor total velocity filtering can be achieved, as it is impossible to realise a δ-function barrier and interactions preclude total velocity filtering.
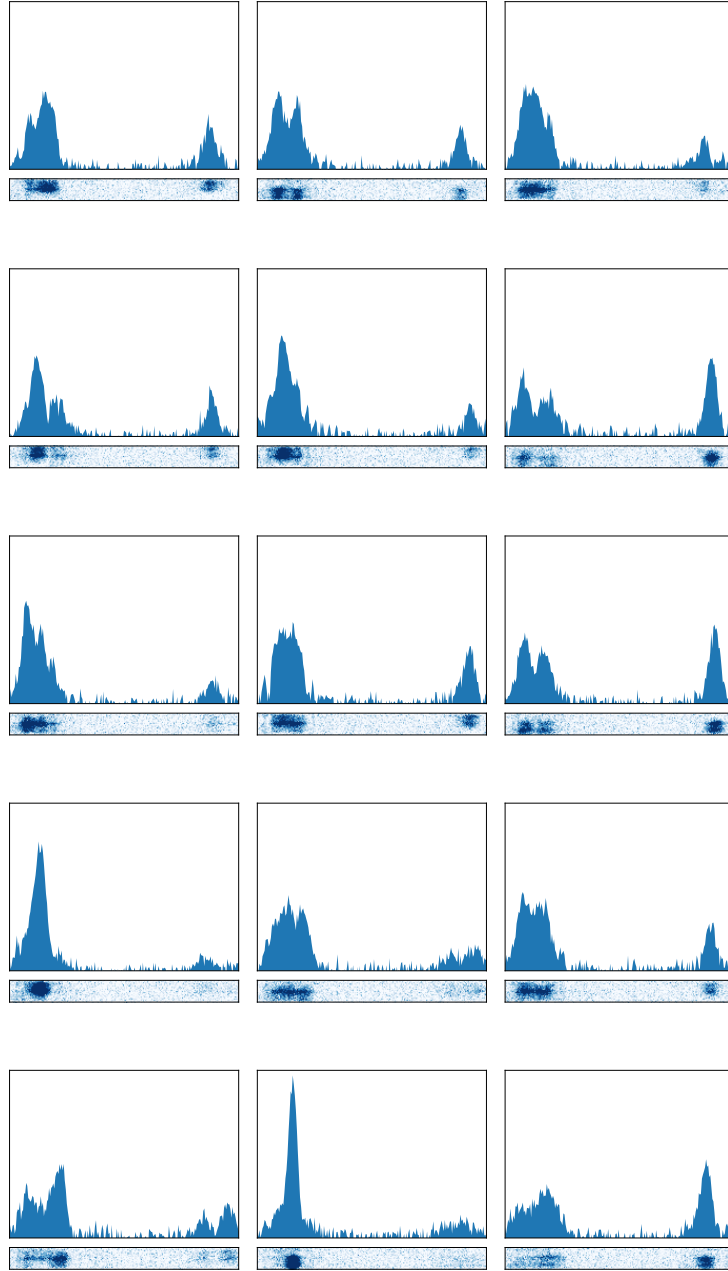
Figure 6.9: Images to show the variation in shot-to-shot recombination observed for the narrow centred barrier. Images were taken 720 ms after the initial release of the soliton into the waveguide (see Fig. 6.8), following identical experimental sequences. Each image is shown along with the profile of the vertical sum.

6.3.3 Importance of the barrier

It is important to note that the barrier is integral to interference-mediated recombination; in the absence of the barrier, the two daughter solitons simply pass through one another. This is seen in Fig. 6.10, where we have split the initial soliton into daughter solitons with unequal atom number, in order to *tag* particular daughter solitons.

Figure 6.10: Image sequences showing that the two daughter solitons simply pass through one another in the absence of the barrier. The initial soliton is split by the wide barrier into two daughter solitons with asymmetric atom number, allowing us to track their individual trajectories. In [a] the barrier was calibrated to give $\sim 20\%$ transmission and in [b] the barrier was calibrated to give $\sim 80\%$ transmission. The image for $0.55\,\mathrm{s}$ in [a] was not taken and so a space is left blank. The dashed black lines are sinusoidal fits to the soliton trajectories.

As in Ref. [40], we cannot disregard the possibility that atom exchange happens during the soliton-soliton collisions, causing the observed population

asymmetries after soliton-soliton collisions. However, these plots show the importance of the barrier for facilitating interference, as the populations after the soliton-soliton collision only depend on the populations prior to the collision, with no apparent phase-sensitivity.

6.4 Extended recombination experiments

It is important that we understand how sensitive interference-mediated recombination is to barrier parameters, namely the barrier transmission and barrier offset, so that we can calculate whether the experimental uncertainties measured in Sec. 6.2 are sufficiently small for performing useful phase-sensitive recombination. This is the first experimental investigation of either dependency, along with the first theoretical exploration of the effects of barrier offset. Only the effects of barrier transmission have been previously considered in GPE theory [61]. Therefore, in this section we explore the dependencies of interference-mediated recombination on barrier offset and barrier transmission, as well as build an intuitive understanding of our observations.

6.4.1 Barrier offset

We intentionally offset the barrier across a range of $\sim 20\,\mu\mathrm{m}$ around the harmonic trap centre, in order to understand the true sensitivity of recombination to offset noise (Fig. 6.11). The barrier transmission is set to give $50\,\%$ transmission at the first barrier interaction throughout, across all offsets.

Theoretically, we observe oscillations when the barrier offset is varied because the barrier offset introduces a position shift of the transmitted and reflected wavepackets, which in turn leads to velocity-induced phase gradients across the wavepackets when they recombine. This results in the observed interference fringes. These fringes are modulated by an envelope caused by the changing spatio-temporal overlap between the wavepackets.

Figure 6.11: The fraction of atoms on the right of the narrow barrier following the second barrier interaction, as a function of barrier offset. The black circles show the mean and the dashed grey envelope (blue shaded region) indicates the standard deviation (maximum and minimum values) across the 5–10 measurements taken at each offset position. The solid purple regions (red dotted lines) are quasi-1D (3D) GPE simulations. Theory lines have no free fitting parameters and experimental values for the barrier offset are determined from independent measurements of the trap and barrier positions.

Experimentally, we observe an increased shot-to-shot fluctuation when the barrier is closer to the centre of the harmonic potential, shown further in Fig. 6.12, within an envelope that is in good qualitative agreement with those predicted by quasi-1D and 3D GPE simulations. However, the oscillatory behaviour seen in the GPE theory is not resolved experimentally, implying that the shot-to-shot variation of the axial harmonic potential relative to the barrier position is large compared to the fringe spacing. In Sec. 6.2, the stability of the barrier offset was measured to be $\pm 1.3\,\mu$m, which is similar to the $1.8\,\mu$m peak-to-peak fringe period predicted by 1D GPE simulations in Fig. 6.11, fully explaining the shot-to-shot variation in $N_{\mathrm{R}}/N_{\mathrm{tot}}$.

Figure 6.12: Histograms demonstrating the increased shot-to-shot fluctuation in $N_{\mathrm{R}}/N_{\mathrm{tot}}$ when the barrier is well centred. [a] shows all data in Fig. 6.11 with $|z_{\mathrm{off}}| > 5\,\mu$m, whereas [b] shows all data with $|z_{\mathrm{off}}| < 1.3\,\mu$m. The height of each bar indicates the probability of occurring in that bin.

However, as the final populations also depend on the relative phase of the incoming daughter solitons, along with the barrier offset, any shot-to-shot variation in the acquired relative phase will manifest as a fluctuation of the final populations (Fig. 6.13). Such a fluctuating phase may be caused by incoherence during the splitting process, or even time-varying anharmonic imperfections of the harmonic potential.

Figure 6.13: Comparing phase noise and barrier offset noise as the source of shot-to-shot fluctuations observed in the experiments with the narrow barrier. 1D-GPE simulations predict oscillatory behaviour when the barrier is offset. A sampling of these fringes is performed by interpolation to simulate the $1.3\,\mu m$ uncertainty in barrier position relative to the centre of the harmonic potential. From this, we determine a mean (labelled 'z_{off}, mean') and an envelope defined by the standard deviation (labelled 'z_{off}, SD'). We achieve a similar mean and envelope if we instead ascribe a random phase to each daughter soliton following splitting (labelled 'ϕ, mean' and 'ϕ, SD' respectively).

In the experiment, phase noise has a similar effect to the offset noise: it averages out the oscillatory behaviour over many shots and creates a shot-to-shot spread of the data which is greatest when the barrier is centred. Computationally, this can be modelled by ascribing each daughter soliton a random phase before recombination, the results of which are also shown in Fig. 6.13.

By comparing the two sources of noise to the measured standard deviation across each barrier offset in Fig. 6.11, we find that the effects of offset noise are experimentally indistinguishable from phase noise (Fig. 6.14). Therefore, we are unable to definitively assign a source for the noise that prevents us from measuring the fringes. Nevertheless, the fact that we measure any sensitivity of the recombined populations to barrier offset is definitive evidence of interference-mediated recombination.

Figure 6.14: Comparing the standard deviation of the fraction of atoms on the right of the narrow barrier with that of the random offset and random phase models. The vertical error bars on the data are estimates of the fractional uncertainties in the error [126].

Note that the small offset of the standard deviation is fully accounted for by the $(6.4 \pm 0.4)\,\%$ shot-to-shot fluctuation in transmission seen in Fig. 5.14,

though a varying atom number in a small thermal population may also contribute to the discrepancy. Whilst these effects are clearly absent in the GPE simulations, the offset does not alter the interpretations of our findings as we still observe a definitive offset-dependent variation.

6.4.2   Barrier transmission

We measure the final populations as a function of barrier transmission for the centred narrow barrier (Fig. 6.15[b]). For a range of barrier heights, images were taken after the first and second barrier interactions in order to determine the transmission and recombination populations. Similar to the study of barrier offset, we again observe good qualitative agreement with the theory, with a higher population on the right after the second barrier interaction for 50 % transmission than for the transmission extremes.

Unlike in Fig. 6.11, the oscillatory behaviour observed in Fig. 6.15 is due to a phase-winding effect, whereby the phase acquired by each of the daughter solitons depends on the number of atoms [61].

6.4.3   Interpretation

As our data shows good agreement with GPE simulations, within experimental uncertainty, it would be useful to gain an intuition as to the relevant experimental parameters for improving our measurements further. Here we outline analytical theoretical interpretations for the observations reported in Fig. 6.11 and Fig. 6.15 respectively. Each model provides an approximation to the GPE solution, in the range of typical experimental parameters.

Figure 6.15: The fraction of atoms on the right of the narrow barrier following the second barrier interaction, as a function of the barrier transmission for a centred barrier. The error bars in $N_{\mathrm{R}}/N_{\mathrm{tot}}$ are calculated by assuming the single-shot uncertainties in N_{R} and N_{L} are given by $\sqrt{N_{\mathrm{R}}}$ and $\sqrt{N_{\mathrm{L}}}$ respectively, where N_{L} is the number of atoms on the left of the barrier following the second barrier interaction. The error bars in transmission are the difference between consecutive transmission measurements which sandwich the interference measurement, at the same barrier height. The solid purple regions (red dotted lines) are quasi-1D (3D) GPE simulations.

Barrier offset

As presented in Sec. 6.4.1, theoretical simulations of our system suggest the occurrence of oscillatory behaviour in the fraction of atoms on the right of the barrier after the second barrier interaction as the barrier offset is varied. This is due to the position shift of the transmitted and reflected wavepackets, which leads to velocity-induced phase gradients across the wavepackets when they recombine.

To better understand the nature of recombination as z_{off} is varied, particularly to attain useful insights for future experimental realisations, we outline a simple approximate analytic model for the fraction of the atoms found on the *right* of the barrier ($z > z_{\text{off}}$) following the second barrier interaction. For convenience, we denote this quantity $A \equiv N_{\text{R}}/N_{\text{tot}}$ and our analytic estimate of it A_{est}. It is convenient to develop the following in dimensionless *soliton units* defined by length $\ell_{\text{s}} = \hbar/(2m\omega_{\text{r}}|a_{\text{s}}|N)$, time $t_{\text{s}} = \hbar/(4m\omega_{\text{r}}^2|a_{\text{s}}|^2N^2)$ and energy $E_{\text{s}} = 4m\omega_{\text{r}}^2|a_{\text{s}}|^2N^2$. In these units, the quasi-1D GPE (Eq. 1.3) becomes:

$$i\frac{\partial\tilde{\psi}}{\partial\tilde{t}} = \left[-\frac{1}{2}\frac{\partial^2}{\partial\tilde{z}^2} + \frac{E_{\text{h}}^2\tilde{z}^2}{2E_{\text{s}}^2} + \frac{\tilde{q}}{\tilde{\ell}}\sqrt{\frac{2}{\pi}}\exp\left(\frac{-2\tilde{z}^2}{\tilde{\ell}^2}\right) - |\tilde{\psi}|^2\right]\tilde{\psi}, \qquad (6.1)$$

where $\tilde{\ell} = w_{\text{z}}/\ell_{\text{s}}$ and

$$\tilde{q} = \frac{\alpha P}{2\sqrt{2\pi}\hbar\epsilon_0 c w_{\text{x}}\omega_{\text{r}}|a_{\text{s}}|N}. \qquad (6.2)$$

Dimensionless velocities are given by:

$$\tilde{v} = \frac{v}{2\omega_{\text{r}}|a_{\text{s}}|N}. \qquad (6.3)$$

Ignoring the relative phase of the solitons, the timing of the various barrier interactions, the interatomic interactions, and assuming the waveguide

potential to be constant over the width of the barrier, A_{est} can be approximated by Eq. 5.5 [122]. In soliton units, the sech^2 approximation to the optical barrier potential is given by:

$$V(z) = \frac{\tilde{q}}{\tilde{\ell}}\sqrt{\frac{2}{\pi}}\exp\left(\frac{-2z^2}{\tilde{\ell}^2}\right) \approx \frac{\tilde{q}}{\tilde{\ell}}\sqrt{\frac{2}{\pi}}\text{sech}^2\left(\frac{\pi z}{\kappa\tilde{\ell}}\right),\qquad(6.4)$$

where $\kappa = (\pi/2)^{3/2}$. From this we obtain the approximate transmission probability *for a single barrier interaction* in soliton units:

$$T(\tilde{v}) = \frac{\sinh^2(\kappa\tilde{\ell}\tilde{v})}{\sinh^2(\kappa\tilde{\ell}\tilde{v}) + \cosh^2\left(\frac{\pi}{2}\sqrt{4\kappa\tilde{q}\tilde{\ell}/\pi - 1}\right)}.\qquad(6.5)$$

We assume that the atoms impact the barrier at the first interaction in the form of an ideal soliton with dimensionless velocity $\tilde{v}_{in} = z_0\omega_z/(2\omega_r|a_s|N)$, chosen such that in the quasi-1D numerics the transmitted fraction is $1/2$. The velocity $\tilde{v}_{half}$ at which $T(\tilde{v}_{half}) = 1/2$ is typically very close to this. After the first barrier interaction, the amplitudes of the outgoing wavepackets in momentum space are approximately:

$$|\tilde{\psi}_{z>z_{off}}(\tilde{v})| \approx t(\tilde{v} - \tilde{v}_{in} + \tilde{v}_{half})\text{sech}[\pi(\tilde{v} - \tilde{v}_{in})]\sqrt{\pi/2},\qquad(6.6)$$

$$|\tilde{\psi}_{z<z_{off}}(\tilde{v})| \approx r(\tilde{v} - \tilde{v}_{in} + \tilde{v}_{half})\text{sech}[\pi(\tilde{v} - \tilde{v}_{in})]\sqrt{\pi/2},\qquad(6.7)$$

where $t(\tilde{v}) = T(\tilde{v})^{1/2}$ and $r(\tilde{v}) = [1 - T(\tilde{v})]^{1/2}$. Between the first and second barrier interactions the wavepackets undergo nonlinear evolution. We empirically approximate this by assuming that the wavepackets re-form into solitons with half of the initial amplitude, preserving the location in momentum space of the *peak* of the transmitted or reflected amplitude, as obtained from Eqs. (6.6–6.7). This gives wavepackets incoming to the second barrier interaction:

$$\tilde{\psi}'_{z>z_{off}}(\tilde{v}) \approx \text{sech}[2\pi(\tilde{v} - \tilde{v}_{t,peak})]\sqrt{\pi/2},\qquad(6.8)$$

$$\tilde{\psi}'_{z<z_{off}}(\tilde{v}) \approx \text{sech}[2\pi(\tilde{v} - \tilde{v}_{r,peak})]\sqrt{\pi/2},\qquad(6.9)$$

where $\tilde{v}_{\mathrm{t,peak}}$ and $\tilde{v}_{\mathrm{r,peak}}$ denote the numerically-determined locations of the peaks described. Finally, the fraction of the total population to the right is approximated by:

$$A_{\mathrm{est}}^{(0)} = \int \left| t(\tilde{v} - \tilde{v}_{\mathrm{in}} + \tilde{v}_{\mathrm{half}}) \, \psi'_{z<z_{\mathrm{off}}}(\tilde{v}) \right.$$
$$\left. + r(\tilde{v} - \tilde{v}_{\mathrm{in}} + \tilde{v}_{\mathrm{half}}) \, \psi'_{z>z_{\mathrm{off}}}(\tilde{v}) \right|^2 dz. \quad (6.10)$$

The estimate above will be modulated by interference between the solitons. We model this using the approach for δ-function barriers described in Ref. [66]. This is expected to be a good approximation for narrow barriers when the dimensionless soliton velocity at the barrier $\tilde{v}_{\mathrm{in}} \gtrsim 1$. Assuming $z_0 \gg z_{\mathrm{off}}$ the dimensionless velocity of all solitons is approximately $\pm\tilde{v}_{\mathrm{in}}$ when they are in the neighbourhood of the barrier. In this approximation the two solitons to the right of the barrier are separated by $4z_{\mathrm{off}}$. Thus, the wavefunction to the right ($z > z_{\mathrm{off}}$) immediately after recombination can be written as:

$$\tilde{\psi}_{\mathrm{right}}(\tilde{z}) \approx \frac{\sqrt{A_{\mathrm{est}}^{(0)}}}{4} \left[e^{i\tilde{v}_{\mathrm{in}}\tilde{z}} \mathrm{sech}\left(\frac{\tilde{z} - \tilde{v}_{\mathrm{in}}\tilde{t}}{4} \right) \right.$$
$$\left. + e^{i\tilde{v}_{\mathrm{in}}(\tilde{z}-4\tilde{z}_{\mathrm{off}})} \mathrm{sech}\left(\frac{\tilde{z} - 4\tilde{z}_{\mathrm{off}} - \tilde{v}_{\mathrm{in}}\tilde{t}}{4} \right) \right], \quad (6.11)$$

where we have assumed equal-amplitude solitons. Note that we have omitted numerous irrelevant phase factors compared to the expressions in Ref. [66] for simplicity, and the phase shift of $\pi/2$ gained by the soliton transmitted at the second barrier interaction simply cancels the phase that the soliton reflected at the second barrier interaction previously acquired by being transmitted at the first barrier interaction. Integrating $|\tilde{\psi}_{\mathrm{right}}(\tilde{z})|^2$, we obtain:

$$A_{\mathrm{est}} = A_{\mathrm{est}}^{(0)} \left[1 + \cos(4\tilde{v}\tilde{z}_{\mathrm{off}}) \frac{\tilde{z}_{\mathrm{off}}}{\sinh(\tilde{z}_{\mathrm{off}})} \right]$$
$$= A_{\mathrm{est}}^{(0)} \left[1 + \cos\left(\frac{4z_0 z_{\mathrm{off}}}{\ell_0^2} \right) \frac{z_{\mathrm{off}}/\ell_{\mathrm{s}}}{\sinh(z_{\mathrm{off}}/\ell_{\mathrm{s}})} \right], \quad (6.12)$$

where the final expression is in real units, the velocity v_{in} has been replaced with $z_0 \omega_{\mathrm{z}}$, and $\ell_0 = (\hbar/m\omega_{\mathrm{z}})^{1/2}$ is the axial harmonic oscillator length.

As shown in Fig. 6.16, this analytic approximation gives a good qualitative picture of the behaviour across a wide regime of parameters and provides a quantitatively accurate value for the fringe spacing across a considerable fraction of this. More specifically, the analytic model provides good estimates when the velocity of the soliton at the barrier $v \gtrsim 2\omega_{\mathrm{r}}|a_{\mathrm{s}}|N$; slower velocities ($v \ll 2\omega_{\mathrm{r}}|a_{\mathrm{s}}|N$) generate much more complex interference fringes due to the nonlinear interactions, which are more significant during longer recombination times.

Eq. 6.12 and Fig. 6.16 demonstrate that the spacing of fringes as a function of barrier offset depends on the soliton velocity and the atomic mass, the amplitude of the fringes depends on the barrier width and the extent of the envelope depends on the soliton width. These are important insights for the feasibility of future phase-sensitive interferometric measurements, which we discuss further in Sec. 6.5.

Barrier transmission

The fringes in the GPE theory in Fig. 6.15 arise from the differing number-dependent chemical potentials [17] between the daughter solitons, inducing a phase-winding effect. As an approximation, we model the chemical potential μ in the Thomas-Fermi limit:

$$\mu = \frac{\hbar\omega_{\mathrm{ho}}}{2} \left(\frac{15N|a_{\mathrm{s}}|}{a_{\mathrm{ho}}} \right)^{\frac{2}{5}}, \tag{6.13}$$

where $a_{\mathrm{ho}} = \sqrt{\hbar/m\omega_{\mathrm{ho}}}$ and $\omega_{\mathrm{ho}} = (\omega_{\mathrm{x}}\omega_{\mathrm{y}}\omega_{\mathrm{z}})^{1/3}$, as defined in Sec. 2.3. Whilst this approximation is not strictly valid for our system, as the Thomas-Fermi limit requires $Na_{\mathrm{s}}/a_{\mathrm{ho}} \gg 1$, it can still be used to gain an insight into the relevant experimental parameters for optimising the experiment in the future.

Figure 6.16: Comparing the recombined fraction as a function of offset predicted by GPE theory with that predicted by our simple analytical model. The variation as a function of barrier width is shown on the left, for a fixed initial soliton displacement of 75 µm. The variation as a function of initial soliton displacement is shown on the right, for a fixed barrier width of 4.74 µm.

By considering the solitons to behave as plane waves, we derive the fractional population on the right of the barrier following the second barrier interaction to be:

$$N_{\mathrm{R}}/N_{\mathrm{tot}} = \frac{\cos\left(\Delta_\mu \tau\right) + 1}{2},$$ (6.14)

where Δ_μ is the difference in chemical potential between the transmitted and reflected components and τ is the time between the first and second barrier interactions.

To retrieve the full theoretical lineshape, Eq. 6.14 is modulated by an envelope, caused by the varying spatial overlap of the daughter soliton wavepackets as one becomes larger than the other. In Fig. 6.17[a] we show the effects of phase winding and the spatial overlap independently, before combining them in Fig. 6.17[b].

Figure 6.17: The fraction of atoms on the right of the narrow barrier following the second barrier interaction, showing the effects of phase winding with barrier transmission. [a] Asymmetric splitting of the initial soliton into daughter solitons creates different chemical potentials, which in turn leads to a phase winding effect in their relative phase (blue). $N_{\mathrm{R}}/N_{\mathrm{tot}}$ is also limited by the spatial overlap of the daughter soliton wavepackets (red). The product of these curves creates the theory curve in [b].

The differences between this basic interpretation and the results seen from GPE simulations (Fig. 6.15) are due to velocity filtering and finite-barrier-width effects, which limit the maximum possible $N_{\mathrm{R}}/N_{\mathrm{tot}}$ and cause the observed asymmetries respectively. Despite the limitations of this basic model,

it is clear that reducing the magnitude of the chemical potential results in a broadening of the observed fringes. Therefore, reducing the chemical potential could lead to less-stringent requirements of the shot-to-shot reproducibility of transmission.

6.5 Implications for phase-sensitive measurements

The experimental observations of the previous sections are definitive signatures of interference-mediated recombination, as the strong increase in population on the right of the barrier (Fig. 6.8) and the increase in population fluctuations (Figs. 6.11 and 6.15) can only be explained by interference effects. Without interference, there is no significance to the increased spatio-temporal overlap of the daughter solitons when the barrier is centred, so the same populations and fluctuations would be recovered for all offsets. To an extent, this is seen between the centred and offset cases for the wide barrier in Figs. 6.5 and 6.7 respectively.

An additional artefact of interference is seen in the double-peaked structure on the left of the barrier following the second barrier interaction in Fig. 6.8; further theoretical simulations demonstrate that introducing an additional phase difference between the daughter solitons can merge the two solitons into one, leaving only a small population on the right of the barrier. This phase dependence is explored theoretically in Fig. 6.18; changes in phase difference between the daughter solitons result in small variations in the final populations for the wide barrier, whereas they result in far more significant variations for the narrow barrier.

It is clear that improved stability in the relative position between the barrier and axial harmonic potential is required to create a viable interferometer; to resolve the oscillatory behaviour in Fig. 6.11 requires the barrier to be controllable at the level of $\sim 0.1\,\mu\mathrm{m}$ with respect to the harmonic potential.

Figure 6.18: 1D GPE simulations of the population on the right following the second barrier interaction, across a range of daughter soliton relative phases (ϕ). The red (green) curves show predictions for the wide (narrow) barrier, the solid blue line is for an idealised δ-function barrier at high velocity and the horizontal dotted line is for a δ-function barrier with non-interfering daughter solitons.

Currently, the axial potential in our experiment is generated magnetically, making it susceptible to the ambient magnetic field. Remarkably, a shot-to-shot variation of only 3 mG is sufficient to fully account for the observed fluctuations.

This sensitivity to magnetic field could be removed altogether using an all-optical potential generated by, for example, acousto-optic deflectors [127,128] or a digital micromirror device [129]. Optical methods are also experimentally attractive because the shot-to-shot stability of our current optical potential is $\sim 0.3\,\mu\text{m}$ and would probably be sufficient to observe the oscillatory behaviour in Figs. 6.11 and 6.15. Furthermore, these methods offer the flexibility for more complicated geometries, such as a ring-shaped trap for a soliton Sagnac interferometer [64].

Alternatively, other atomic systems may prove to be more resilient to barrier position. For instance, the lower mass of lithium would result in broader fringes, leading to less-stringent requirements of the relative position stability. The lower mass has the added benefit of making the solitons themselves larger, so the barrier is comparatively narrower and velocity filtering is suppressed.

In Sec. 6.4.1 we discussed the possibility that our observations may be explained by incoherent soliton splitting, which would prevent our setup from realising a phase-sensitive interferometer in the current configuration. A measurement of either the fringes seen in Figs. 6.11 or 6.15, or the observation of controllable, deterministic soliton-soliton collisions would be definitive proof of coherence between the daughter solitons. However, it is important to stress that all of our results are fully consistent with the independently-determined uncertainty in barrier offset.

Chapter 7

Conclusions and outlook

In this chapter we conclude by reviewing the major achievements of this work, before detailing their implications for the field. We then provide an indication of possible future work of interest for the experiment, along with our initial investigations into some of these areas.

7.1   Summary

In this thesis we have demonstrated several key steps towards the ability to make phase-sensitive measurements and the realisation of a soliton interferometer, namely:

- A tunable, robust imaging system that is capable of imaging the solitons *in-situ* and across a wide range of magnetic fields.

- A reliable method for producing robust bright-solitary-matter waves in a tunable quasi-1D potential.

- Bright-solitary-matter waves with extremely long lifetimes.

- Controllable splitting of a bright-solitary-matter wave on a narrow repulsive barrier.

- Interference-mediated recombination of bright-solitary-matter waves on a narrow repulsive barrier.

- An understanding of important sensitive experimental parameters for achieving interference-mediated recombination.

In Ch. 2, we gave an overview of the steps we take to create BECs of ^{85}Rb through direct evaporation. This was a multi-stage process involving: magneto-optical trapping, optical cooling, optical pumping, magnetic transport, radio-frequency evaporation and optical evaporation, all inside a dual-chamber vacuum system.

In Ch. 3, we demonstrated a range of imaging techniques and apparatus, which were used extensively throughout this thesis. We detailed our tunable beat-locked imaging laser system which, along with our microwave-transfer imaging technique, allows us to reliably image solitons *in-situ* across a wide range of magnetic fields.

The BECs of Ch. 2 were then used in Ch. 4 to create robust and reproducible bright-solitary-matter waves in a quasi-1D potential, by wavefunction engineering and mode-matching of the spatial profile of the condensate. We showed that the solitons performed centre of mass harmonic oscillations in an additional harmonic potential, applied along the quasi-1D axis, which were controllable over a wide range of frequencies and amplitudes. The solitons had lifetimes of over 20 s and were observed to propagate for over 2 cm, far exceeding that of similar experiments.

In Ch. 5, we demonstrated the controllable splitting of a soliton on a narrow repulsive Gaussian barrier, formed by a blue-detuned laser beam. The soliton was split into two daughter solitons, with populations in each controlled by the barrier height and the centre of mass kinetic energy of the initial soliton

at the barrier, both of which could be precisely tuned. We observed a velocity selection phenomenon, which caused the transmitted daughter soliton to always leave the barrier with a higher centre of mass kinetic energy than the reflected daughter soliton. This was explained by a simple model which compared the quantum mechanical momentum distribution of the soliton with a barrier-width-dependent transmission function.

In Ch. 6, we reported on our observations of soliton recombination. Our initial investigation used a relatively wide barrier which, as seen in Ch. 5, lead to a large degree of velocity filtering. This was found to cause a velocity-filtering-mediated *merging* of the daughter solitons at the second barrier interaction. However, by using our narrow barrier, aligned with the centre of the harmonic potential, we showed that interference-mediated *recombination* is possible. We further explored the effects of varying the barrier position and transmission, observing that the shot-to-shot variation in the recombined fraction increased for a centred barrier. We found excellent qualitative agreement with theory, within independently-measured levels of experimental noise. This confirmed that our measurements could only be explained by interference effects. Finally, we developed a set of simple analytic models to allow an intuitive interpretation of the observations, in terms of controllable parameters. These models gave insight into how a practical phase-sensitive interferometric measurement, using solitons and a narrow repulsive barriers, may be made in the future.

7.2   Future work and initial explorations

This work has shown a proof of principle that invites a wide array of exciting future experiments, using or building upon our current setup. Here we outline a few that have the potential to bring interesting and significant results.

7.2.1   BEC-soliton mode-matching

In Ch. 4 we reported on the production of highly-stable solitons, which was only possible because of the closely-matched spatial mode between the BEC in the crossed dipole trap and the soliton in the waveguide. Whilst we proposed that mode-matching using a different condensate may be exploited to create solitons with very different parameters from those used in this work, it was left unexplored. However, in light of the findings made later in this thesis, specifically those pertaining to the sensitivity of interference-mediated recombination to the size and interaction strength of the soliton (Eqs. 6.12 and 6.14), it may be beneficial to optimise the experimental sequence for a soliton with different parameters.

7.2.2   Velocity filtering studies

The work presented in Secs. 5.6 and 5.7 of this thesis identified the difference in kinetic energy between the daughter solitons following the first barrier interaction. Whilst we explored certain areas of parameter space, predominantly by varying the interatomic interactions, there is still more to be gained from a systematic study with a range of barrier widths. Along with additional experimental work, a thorough theoretical investigation may provide a deeper understanding into the nature of soliton splitting on finite-width barriers, as well as the classical to quantum crossover.

7.2.3   Deterministic phase control and measurement

Once the technical challenges highlighted in Sec. 6.5 have been overcome and interference-mediated recombination with a repeatable outcome has been shown, it should be possible to have full control over $N_{\mathrm{R}}/N_{\mathrm{tot}}$ by applying an arbitrary phase difference between the daughter solitons with an off-resonant light-shift. By imaging a knife edge onto half of the harmonic potential after

splitting, only one daughter soliton will be exposed to the light shift. The phase shift acquired by an off-resonant pulse of light is given by:

$$\phi(x,y) = \frac{U(x,y)\tau}{\hbar} = \frac{-I(x,y)\alpha\tau}{\hbar\epsilon_0 c}. \tag{7.1}$$

Therefore, it is trivial to calculate the required power to create a fixed phase shift, as long as the polarisability of the atoms for the light-shift wavelength is known. For example, Fig. 7.1 shows the required 850 nm laser power to produce a full 2π phase shift, for various different light-shift beam widths.

Figure 7.1: The power of 850 nm light required to produce a 2π phase shift between the daughter solitons, when imaged onto one of them. As the width of the light-shift beam increases, the power required also increases, in order to maintain the same intensity. This figure assumes a pulse duration of 360 ms, which was typical of a half trapping period of the experiment.

A slight complication with this method is caused by the curvature of the light-shift beam itself, as this produces a spatially-varying potential (Eq. 7.1) which alters the centre of mass harmonic oscillation frequency of the soliton. If this is shifted too far then the two daughter solitons will no longer spatio-temporally overlap at the barrier, preventing interference-mediated recombination. Therefore, a suitably large beam with sufficient power to create

the required phase shift will be necessary to effectively engineer the relative phase of the daughter solitons without impacting the functionality of the interferometer.

7.2.4   Solitons in a ring

The ability to create, split and recombine matter-wave solitons in a ring-shaped waveguide is of interest for interferometry [64], but also has potential as a platform for atomtronics. Previous work from our experiment has explored the creation of ring and arbitrary optical potentials using a spatial light modulator [15]. However, it proved difficult to create a sufficiently smooth potential. Therefore, the use of a digital micromirror device [129], acousto-optic deflector [127, 128] or intensity mask, all of which are proven means of achieving arbitrary potentials, could provide a more effective solution.

Sagnac interferometer

Coupled with the narrow barrier that has been the subject of this thesis, it is straightforward to envisage a setup such as Fig. 7.2 that utilises the ring-shaped potential as a Sagnac interferometer. Here the path length of each daughter soliton is influenced by rotational forces, which creates an effective rotation speed Ω. This in turn results in a phase difference between the daughter solitons of:

$$\phi = \frac{4\pi}{v\lambda}\Omega A, \tag{7.2}$$

where v is the centre of mass velocity of the soliton and A is the area enclosed by the ring.

However, Sagnac interferometry may prove challenging, as the path length of each soliton will be altered by rotational forces, by definition. Therefore, for large rotations the crossover point of the daughter solitons may become

Figure 7.2: A Sagnac interferometer using bright-solitary-matter waves. The initial soliton (dark blue) is incident on the narrow repulsive barrier (green), whereupon it is split into two daughter solitons (light blue) which each travel around the ring potential. Once they reach the barrier for a second time they undergo interference-mediated recombination. This arrangement is sensitive to rotational forces that impart an effective rotational speed Ω on the relative path length of the daughter solitons.

significantly misaligned from the barrier position, causing them to no longer spatio-temporally overlap at the barrier. This would prevent interference-mediated recombination from occurring. Fortunately, theory work suggests that the misalignment problem may not be encountered for the small rotations necessary to achieve significant phase shifts. Note that misalignment in this fashion is not a consideration for the linear case presented in this thesis, as the harmonic potential ensures that any path length variations do not alter the time at which the daughter solitons arrive at the barrier.

Atomtronics

With the ability to create arbitrary waveguides that can maintain and guide a soliton, along with a reliable technique for performing reliable splitting and recombination operations, one may envisage the creation of a wide range of matter-wave circuits. This could have significant impact for the field of

atomtronics, as the beneficial properties of solitons (no dispersion, long coherence times etc.) may be exploited for the efficient processing of packets of information [130].

7.2.5 Soliton-soliton collisions

Our apparatus also lends itself to soliton-soliton collision experiments. As the relative phase of the daughter solitons is expected to be well-defined and controllable, from GPE simulations, the outcome of soliton-soliton collisions could be controlled completely deterministically, unlike in other reported experiments [29, 40]. This could be achieved using the same light-shift technique described above. The ability to manipulate the relative velocity and population fractions of the daughter solitons also allows us to access a wide parameter range of interest [41, 42].

Another potential scheme for low-velocity soliton-soliton collisions was discussed in Sec. 6.3: by offsetting the barrier from the centre of the harmonic potential and waiting until after the second barrier interaction, velocity selection allows most of the population to leave on one side of the barrier. These outgoing solitons will collide close to the edge of the harmonic potential at low velocity. This scheme could also be adapted so that the barrier is very high for the second interaction. By switching this high barrier off before the transmitted soliton reaches it, the reflected daughter soliton will be fully re-reflected by it but the transmitted soliton will not see it. This arrangement has the advantage of only one splitting event, so no atoms are lost to the wrong side of the barrier, as well as minimising coherence loss between the daughter solitons.

All of the above schemes rely on coherence between the daughter solitons. We have already observed the daughter solitons to undergo many soliton-soliton collisions in the absence of the barrier without any instances of mergers or collapse, as shown in Fig. 7.3, which is a strong experimental marker for long coherence times. However, we cannot yet conclusively determine whether the

splitting process truly retains coherence between the daughter solitons, as the shot-to-shot fluctuations seen in Fig. 6.11 could be equally explained by phase noise (see Fig. 6.14). Interestingly, a measurement of either the fringes seen in Fig. 6.11 or controllable, deterministic soliton-soliton collisions would be definitive proof of coherence between the daughter solitons.
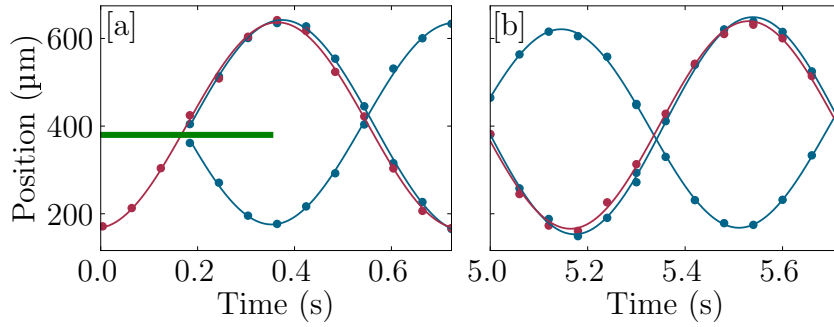
Figure 7.3: The trajectories of two daughter solitons (blue) oscillating in the harmonic potential after splitting on the wide barrier (green), along with the trajectory of an un-split soliton (red). [a] and [b] are 720 ms windows taken 5 s apart to precisely determine the frequency of the harmonic oscillation.

If two solitons collide with a relative internal phase of $\phi = \pi$, they will undergo a jump in the phase of the harmonic oscillation [40]. Experimentally, the consequence of this jump is an apparent increase in the oscillation frequency of the solitons, relative to the trapping frequency. Therefore, in Fig. 7.3, we attempt to measure this effect by alternately imaging a single soliton and split daughter solitons, at the same time intervals, in the same harmonic potential. From a sinusoidal fit of the trajectories, we determine the oscillation frequency of the single soliton to be $\omega_z = 2\pi \times 1.3548(7)$ Hz and the oscillation frequency of the split daughter solitons to be $\omega_z = 2\pi \times 1.3561(6)$ Hz. Though the central frequency of the split daughter solitons is higher than that for the single soliton, they are both consistent with each other within uncertainty.

This observation is also consistent with the analytical predictions in Ref. [40],

where they determine the frequency shift to be given by:

$$\Delta\omega_z = -\frac{0.67(N/N_{crit})a_z^4\omega_z}{\pi z_0^3 a_r},$$

(7.3)

where $a_z = \sqrt{\hbar/m\omega_z}$ is the axial confining length scale. For a soliton with 2000 atoms at $a_s = -10a_0$, initially displaced $225\,\mu m$ away from the centre of a $1.4\,Hz$ harmonic potential ($\omega_r = 2\pi \times 1.4\,Hz$), as in Fig. 7.3, Eq. 7.3 predicts a frequency shift of only $\omega_r = 2\pi \times 0.0001\,Hz$. This is below the uncertainty limit of the measurements in Fig. 7.3. Smaller values of z_0 may allow frequency shifts to become measurable. For example, reducing z_0 to $50\,\mu m$ is expected to produce a $\omega_r = 2\pi \times 0.009\,Hz$ frequency shift, which would be easily detectable in the current configuration.

7.2.6   Post-quench dynamics

Many theoretical works have striven to locate the limits of applicability of the GPE for the 1D Bose gas. For example, so called 'fragmentons': objects that have soliton-like dynamical properties (dispersion-less) but do not retain coherence, arise in many-body calculations [131] and do not exist within the GPE framework.

Another area where the validity of the GPE is questioned is in the dynamics of a soliton following a quench of the interatomic interaction strength. Mean-field calculations give specific predictions for interaction quenches [132]: a factor of four quench is predicted to create robust oscillatory behaviour for all times. These objects are deemed to be higher-order solitons [133]. However, it has been suggested that beyond-mean-field effects may lead to the dynamics being markedly different. For the same quench, beyond-mean-field effects may lead to a central soliton core with an emission of single particles [134].

As a first step towards investigating the dynamics that follow an interaction quench of a soliton, we have performed a series of preliminary experiments

(Fig. 7.4). We began with a soliton at $a_s = -8\,a_0$ in the same quasi-1D potential used throughout the latter parts of this thesis. After a fixed free-oscillation time, the magnetic field was rapidly jumped to produce $a_s = -16\,a_0$. We performed *in-situ* imaging of the solitons to observe the resultant dynamics.

Figure 7.4: A series of images of solitons to show the dynamics following a quench of scattering length from $-8\,a_0$ to $-16\,a_0$. In [a] the quench takes place at 80 ms; in [b] it happens at 180 ms. The absolute position values in each figure have an arbitrary offset and are only intended to convey scale. The jitter in position away from the expected harmonic trajectory is probably due to fluctuations in the harmonic potential, as encountered in Ch. 6.

Although we have not performed the full factor of four quench described above, Fig. 7.4 hints at the potential of a rich set of dynamics to explore. Many images show the breakup of the original soliton into three objects, particularly in series [b]. This may be suggestive of a higher-order soliton, a set of fragmentons, or simply the creation of a soliton train. If these objects are a soliton train, the breakup may be linked to the reduction of the modulation instability wavelength with the increased interaction strength, as discussed in Sec. 4.3.1.

Conversely, although there are some images in series [a] that display multiple

objects (particularly later in the series), most appear to show an increased background of atoms surrounding a central object, which may be more indicative of the beyond-mean-field predictions. Interestingly, very few images in either series show the formation of two objects, most appear to be of only one or three.

7.2.7 Collapse dynamics using microwave transfer

Around the time of the early works on bright-solitary-matter waves, it was identified that a quench of the scattering length from positive to negative resulted in a *Bosenova* effect. This manifested as an anisotropic burst of atoms exploding from a collapsing condensate, leaving behind a highly-excited remnant of multiple solitons [26,135]. In Ref. [135], the authors explain that they were able to perform these quenches in as little as $100\,\mu s$.

Whilst our experiment would not be able to improve on this timescale with magnetic fields, the microwave-transfer technique we use for *in-situ* imaging (Sec. 3.6) may be re-purposed to create a 10-fold improvement. This could allow for more precise studies of post-quench dynamics. As the Feshbach resonance used throughout this thesis is for the $|2, -2\rangle$ state and the background scattering length for ^{85}Rb is large and negative, transferring the population from $|2, -2\rangle$ to $|3, -3\rangle$ will produce a strong quench of the interaction strength.

Alongside this, it may be possible to use the current partial-transfer imaging system in reverse (transfer some atoms back to $|2, -2\rangle$ for imaging) in order to observe the full collapse dynamics of a single condensate. Even though the imaging transitions from $|2, -2\rangle$ are not fully-closed, we have demonstrated direct imaging from $|2, -2\rangle$ earlier in this thesis (see Sec. 3.5) and so do not expect it to cause too many problems. The rapid state transfer offers the possibility for extremely fast interatomic interaction quenches, changing the scattering length from highly-repulsive to highly-attractive in $\sim 10\,\mu s$ timescales.

7.2.8 Dynamics of chaos

Classical chaos is often associated with the ergodic filling of regions in phase space. This criterion cannot be easily applied to matter-wave systems, as dispersion usually dominates over the divergence of neighbouring trajectories. Therefore, it is clear that solitons offer a neat solution for studying chaotic dynamics. It has been shown that as few as three solitons, confined to a 1D potential (such as that used in this thesis), are sufficient to expect chaotic dynamics [136].

As a proof-of-principle, we have demonstrated the splitting of a soliton into three daughter solitons, shown in Fig. 7.5. The barrier was aligned on the soliton-release side of the harmonic potential, so that it could be turned off after the second barrier interaction. The barrier transmission was initially set to $\sim 33\,\%$, causing the majority of the population to be in the reflected soliton. The barrier height was then reduced to give $\sim 50\,\%$ transmission for the second barrier interaction, resulting in the initially-reflected daughter soliton being split into two granddaughter solitons. The barrier was removed and the final three solitons were allowed to oscillate freely in the trap. Unfortunately, we were unable to dedicate the time for a systematic investigation into chaotic dynamics. However, Fig. 7.5 demonstrates the technical ability of the current experiment to produce the required system.

7.2.9 Short-range atom-surface interactions

As mentioned in Sec. 4.2.1, a Dove prism is mounted inside the science chamber, such that the N-S dipole trapping beam passes through it to create half of the crossed-dipole trap. The main motivation for including the prism was for the study of atom-surface interactions using novel soliton-based experiments.

Atoms close to a surface experience the short-range, attractive Casimir-Polder and van der Waals potentials. However, by aligning a blue-detuned

Figure 7.5: Splitting a soliton into three daughter/ granddaughter solitons with a narrow repulsive Gaussian barrier. The barrier height is adjusted between the first and second splitting events to produce three solitons with approximately equal atom number.

Figure 7.6: Picture of the proposed short-range atom-surface experiment. A soliton is launched along the N-S beam towards the prism, whereupon it interacts with the short-range atom-surface potential. This potential can be tuned using an evanescent field produced by a blue-detuned laser beam.

laser beam through the side of the Dove prism, it is possible to create a repulsive evanescent field in the vicinity of the soliton-surface interaction (Fig. 7.6). In this way the net potential may be tuned between attractive and repulsive.

Experimentally, the soliton can be launched towards the surface using an off-set curvature field, as used in this thesis; or a transverse guiding beam controlled by an acousto-optic deflector. By observing the subsequent trajectory of the soliton, one can extract information about short-range atom-surface potentials, as well as classical and quantum reflection [137–139]. Solitons are expected to offer benefits over other methods for such an experiment, owing to their long lifetimes and stability against collisions [140, 141].

7.3   Concluding remarks

The controllable splitting and interference-mediated recombination of long-lived bright-solitary-matter waves are significant advances towards a greater understanding of the phase and coherence properties of bright-solitary-matter waves, as well as being fundamental steps in the development of a soliton interferometer. Furthermore, the knowledge that was acquired on the way to making recombination a success will be vital considerations for future experiments with solitons and narrow barriers. For example, we have presented the first experimental observations of the effects of barrier width, offset and transmission on interference-mediated recombination, which all carry important implications for a possible future soliton interferometer.

This chapter has highlighted some of the many exciting experiments that could be performed using apparatus similar, if not the same, to that used in this thesis. From a pure physics perspective, the possibility of studying soliton-soliton collisions to gain a deeper understanding of soliton phase is a particularly interesting prospect. However, through the applications to atomtronics, experimental studies of soliton control in arbitrary potentials could offer significant applied impact on a relatively young field.

Appendix A

Tables of transitions

In this Appendix, the frequencies and transition strengths of various transitions of interest are displayed. All values are determined using the electric susceptibility calculator *Elecsus*, which was developed by former members of the Quantum Light and Matter group at Durham University [103].

	Ground State			Excited State				
F	m_F	Δ (MHz)	F	m_F	Δ (MHz)	F	m_F	Δ (MHz)
2	-2	144.12	1	-1	-387.168	3	1	185.7
2	-1	60.97	1	0	-417.09	3	2	188.79
2	0	-17.69	1	1	-442.69	3	3	190.23
2	1	-92.53	2	-2	-380.855	4	-4	-464.35
2	2	-164.04	2	-1	-117.19	4	-3	-218.50
3	-3	-232.23	2	0	-99.12	4	-2	29.48
3	-2	-143.85	2	1	-81.84	4	-1	283.56
3	-1	-60.83	2	2	-65.3	4	0	319.12
3	0	17.69	3	-3	-400.59	4	1	355.85
3	1	92.39	3	-2	-112.84	4	2	392.61
3	2	163.76	3	-1	169.18	4	3	428.87
3	3	232.23	3	0	179.82	4	4	464.35

Table A.1: Tables showing the frequency shifts (Δ) of the ^{85}Rb $5S_{1/2}$ ground and $5P_{3/2}$ excited state energy levels (the $D2$ transition energy levels) at 165.85 G, relative to their zero-field value. The coupled basis (F, m_F) is not well defined in the excited states at this magnetic field, so the labels represent the energy level's low-field state, which can be tracked to the high-field regime using a Breit-Rabi diagram (Fig. 3.10)

<div align="center">— σ^+ transitions —</div>

$\lvert g\rangle$ F	m_F	$\lvert e\rangle$ F	m_F	δ (MHz)	A	$\lvert g\rangle$ F	m_F	$\lvert e\rangle$ F	m_F	δ (MHz)	A
2	2	3	3	352	0.1163	3	-2	2	2	-106	0.0049
2	2	4	3	712	0.0237	3	-2	3	-1	192	0.0562
2	1	2	-1	-90	0.0019	3	-2	4	-1	428	0.0981
2	1	3	2	279	0.1151	3	-1	2	1	-205	0.0033
2	1	4	2	604	0.0542	3	-1	3	0	120	0.0266
2	0	2	0	-147	0.0028	3	-1	4	0	380	0.1715
2	0	3	1	201	0.1077	3	0	2	0	-301	0.0017
2	0	4	1	493	0.0949	3	0	3	1	47	0.0111
2	-1	2	1	-208	0.0027	3	0	4	1	338	0.2263
2	-1	3	0	117	0.0894	3	1	2	-1	-394	0.0006
2	-1	4	0	377	0.1509	3	1	3	2	-24	0.0038
2	-2	2	2	-275	0.0016	3	1	4	2	300	0.2688
2	-2	3	-1	23	0.0548	3	2	3	3	-94	0.0009
2	-2	4	-1	258	0.2289	3	2	4	3	265	0.3036
3	-3	3	-2	-1	0.0053	3	3	4	4	232	0.3333
3	-3	4	-2	262	0.1058						

Table A.2: Table to show the relative transition frequency detunings (δ) and strengths (A) of the σ^+ transitions at $165.85\,\mathrm{G}$. The coupled basis labels of the ground state ($\lvert g\rangle$) and excited state ($\lvert e\rangle$) indicate the low-field state, which can be tracked to the high-field with a Breit-Rabi diagram. The detuning of the transitions from the $F = 2$ states are relative to the zero-field $F = 2 \rightarrow F = 3$ transition frequency and those from $F = 3$ are relative to the zero-field $F = 3 \rightarrow F = 4$ transition frequency.

$|g\rangle$		$|e\rangle$		δ	A	$|g\rangle$		$|e\rangle$		δ	A
F	m_{F}	F	m_{F}	(MHz)		F	m_{F}	F	m_{F}	(MHz)	
2	2	1	1	-373	0.3001	3	-1	2	-2	-504	0.1543
2	2	2	0	0	0.0042	3	-1	3	-2	-173	0.0849
2	1	1	0	-419	0.2632	3	-1	4	-2	91	0.0039
2	1	2	1	-55	0.0094	3	0	1	-1	-618	0.1052
2	1	3	0	270	0.0006	3	0	2	2	-267	0.0946
2	0	1	-1	-464	0.2209	3	0	3	-1	31	0.0055
2	0	2	2	-113	0.0169	3	1	1	0	-723	0.0654
2	0	3	-1	185	0.0012	3	1	2	1	-358	0.1003
2	-1	2	-2	-507	0.1699	3	1	3	0	-33	0.0054
2	-1	3	-2	-176	0.0297	3	1	4	0	227	0.0002
2	-1	4	-2	88	0.0017	3	2	1	1	-820	0.031
2	-2	3	-3	-547	0.1033	3	2	2	0	-447	0.1046
2	-2	4	-3	-244	0.0559	3	2	3	1	-99	0.0042
3	-3	4	-4	-232	0.3333	3	2	4	1	192	0.0002
3	-2	3	-3	-377	0.2214	3	3	2	-1	-534	0.1087
3	-2	4	-3	-74	0.0639	3	3	3	2	-164	0.0024

Table A.3: Table to show the relative transition frequency detunings (δ) and strengths (A) of the σ^- transitions at $165.85\,\mathrm{G}$. The coupled basis labels of the ground state ($|g\rangle$) and excited state ($|e\rangle$) indicate the low-field state, which can be tracked to the high-field with a Breit-Rabi diagram. The detuning of the transitions from the $F = 2$ states are relative to the zero-field $F = 2 \rightarrow F = 3$ transition frequency and those from $F = 3$ are relative to the zero-field $F = 3 \rightarrow F = 4$ transition frequency.

$|g\rangle$		$|e\rangle$		δ	A	$|g\rangle$		$|e\rangle$		δ	A
F	m_{F}	F	m_{F}	(MHz)		F	m_{F}	F	m_{F}	(MHz)	
2	2	2	-1	-19	0.2116	3	-2	2	-2	-421	0.0068
2	2	3	2	351	0.01	3	-2	3	-2	-90	0.1229
2	2	4	2	676	0.0006	3	-2	4	-2	174	0.0925
2	1	1	1	-445	0.0016	3	-1	1	-1	-540	0.0041
2	1	2	0	-72	0.195	3	-1	2	2	-189	0.075
2	1	3	1	276	0.024	3	-1	3	-1	109	0.1393
2	1	4	1	567	0.0016	3	-1	4	-1	345	0.0039
2	0	1	0	-494	0.0028	3	0	1	0	-648	0.0019
2	0	2	1	-130	0.1725	3	0	2	1	-284	0.0453
2	0	3	0	196	0.0445	3	0	3	0	41	0.1668
2	0	4	0	456	0.0025	3	0	4	0	302	0.0082
2	-1	1	-1	-543	0.0031	3	1	1	1	-748	0.0006
2	-1	2	2	-192	0.1403	3	1	2	0	-376	0.0251
2	-1	3	-1	106	0.0764	3	1	3	1	-27	0.1862
2	-1	4	-1	342	0.0024	3	1	4	1	264	0.0104
2	-2	2	-2	-590	0.0022	3	2	2	-1	-465	0.0106
2	-2	3	-2	-259	0.0905	3	2	3	2	-96	0.202
2	-2	4	-2	4	0.1295	3	2	4	2	229	0.0096
3	-3	3	-3	-289	0.0086	3	3	3	3	-163	0.2162
3	-3	4	-3	14	0.2136	3	3	4	3	197	0.0061

Table A.4: Table to show the relative transition frequency detunings (δ) and strengths (A) of the π transitions at $165.85\,\mathrm{G}$. The coupled basis labels of the ground state ($|g\rangle$) and excited state ($|e\rangle$) indicate the low-field state, which can be tracked to the high-field with a Breit-Rabi diagram. The detuning of the transitions from the $F = 2$ states are relative to the zero-field $F = 2 \to F = 3$ transition frequency and those from $F = 3$ are relative to the zero-field $F = 3 \to F = 4$ transition frequency.

Bibliography

[1] Y. S. Kivshar and B. A. Malomed, Rev. Mod. Phys. **61**, 763 (1989).

[2] Y. V. Kartashov, B. A. Malomed, and L. Torner, Rev. Mod. Phys. **83**, 247 (2011).

[3] J. S. Russell, Report on waves. report of the fourteenth meeting of the british association for the advancement of science held at york in september 1844, 1845.

[4] T. Dauxois and M. Peyrard, *Physics of solitons*, Cambridge University Press, 2006.

[5] L. F. Mollenauer, R. H. Stolen, and J. P. Gordon, Phys. Rev. Lett. **45**, 1095 (1980).

[6] A. Hasegawa and Y. Kodama, Opt. Lett. **15**, 1443 (1990).

[7] Y. S. Kivshar and G. Agrawal, *Optical solitons: from fibers to photonic crystals*, Academic press, 2003.

[8] D. Y. Tang, H. Zhang, L. M. Zhao, and X. Wu, Phys. Rev. Lett. **101**, 153904 (2008).

[9] K. Stasiewicz et al., Phys. Rev. Lett. **90**, 085002 (2003).

[10] Y. Togawa et al., Phys. Rev. Lett. **111**, 197204 (2013).

[11] T. Heimburg and A. D. Jackson, Proceedings of the National Academy of Sciences **102**, 9790 (2005).

[12] S. Händel, *Experiments on ultracold quantum gases of 85Rb and 87Rb*, PhD thesis, Durham University, 2011.

[13] A. Marchant, *Formation of bright solitary matter-waves*, PhD thesis, Durham University, 2012.

[14] T. Wiles, *Dynamics of bright solitary matter-waves*, PhD thesis, Durham University, 2013.

[15] M. Yu, *Towards interferometry with bright solitary waves in a ring*, PhD thesis, Durham University, 2016.

[16] P. G. Drazin, *Solitons*, London Mathematical Society Lecture Note Series, Cambridge University Press, 1983.

[17] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari, Rev. Mod. Phys. **71**, 463 (1999).

[18] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, Rev. Mod. Phys. **82**, 1225 (2010).

[19] N. R. Claussen et al., Phys. Rev. A **67**, 060701 (2003).

[20] D. J. Frantzeskakis, Journal of Physics A: Mathematical and Theoretical **43**, 213001 (2010).

[21] O. Zobay, S. Pötting, P. Meystre, and E. M. Wright, Phys. Rev. A **59**, 643 (1999).

[22] J. Cuevas, P. G. Kevrekidis, B. A. Malomed, P. Dyke, and R. G. Hulet, New Journal of Physics **15**, 063006 (2013).

[23] A. L. Marchant et al., Nature Communications **4**, 1865 (2013).

[24] K. E. Strecker, G. B. Partridge, A. G. Truscott, and R. G. Hulet, Nature **417**, 150 (2002).

[25] L. Khaykovich et al., Science **296**, 1290 (2002).

[26] S. L. Cornish, S. T. Thompson, and C. E. Wieman, Phys. Rev. Lett. **96**, 170401 (2006).

[27] G. D. McDonald et al., Phys. Rev. Lett. **113**, 013002 (2014).

[28] S. Lepoutre et al., Phys. Rev. A **94**, 053626 (2016).

[29] T. Mežnaršič et al., Phys. Rev. A **99**, 033625 (2019).

[30] A. Di Carli et al., arXiv e-prints , arXiv:1905.03690 (2019).

[31] P. A. Ruprecht, M. J. Holland, K. Burnett, and M. Edwards, Phys. Rev. A **51**, 4704 (1995).

[32] N. G. Parker, S. L. Cornish, C. S. Adams, and A. M. Martin, Journal of Physics B: Atomic, Molecular and Optical Physics **40**, 3127 (2007).

[33] A. L. Marchant et al., Phys. Rev. A **93**, 021604 (2016).

[34] M. R. Andrews et al., Science **273**, 84 (1996).

[35] M. R. Andrews et al., Phys. Rev. Lett. **79**, 553 (1997).

[36] R. Meppelink, R. A. Rozendaal, S. B. Koller, J. M. Vogels, and P. van der Straten, Phys. Rev. A **81**, 053632 (2010).

[37] M. Gajdacz et al., Review of Scientific Instruments **84** (2013).

[38] Kaminski, F. et al., Eur. Phys. J. D **66**, 227 (2012).

[39] A. Ramanathan et al., Review of Scientific Instruments **83** (2012).

[40] J. H. Nguyen, P. Dyke, D. Luo, B. A. Malomed, and R. G. Hulet, Nature Physics **10**, 918 (2014).

[41] N. G. Parker, A. M. Martin, S. L. Cornish, and C. S. Adams, Journal of Physics B: Atomic, Molecular and Optical Physics **41**, 045303 (2008).

[42] T. P. Billam, S. L. Cornish, and S. A. Gardiner, Phys. Rev. A **83**, 041602 (2011).

[43] S. Choudhury, A. Sreedharan, R. Mukherjee, A. Streltsov, and S. Wüster, arXiv e-prints , arXiv:1904.11878 (2019).

[44] U. Al Khawaja, H. T. C. Stoof, R. G. Hulet, K. E. Strecker, and G. B. Partridge, Phys. Rev. Lett. **89**, 200404 (2002).

[45] L. Salasnich, A. Parola, and L. Reatto, Phys. Rev. A **66**, 043603 (2002).

[46] L. Salasnich, A. Parola, and L. Reatto, Phys. Rev. Lett. **91**, 080405 (2003).

[47] L. D. Carr and J. Brand, Phys. Rev. Lett. **92**, 040401 (2004).

[48] L. D. Carr and J. Brand, Phys. Rev. A **70**, 033607 (2004).

[49] A. I. Streltsov, O. E. Alon, and L. S. Cederbaum, Phys. Rev. Lett. **106**, 240401 (2011).

[50] B. J. Dabrowska-Wüster, S. Wüster, and M. J. Davis, New Journal of Physics **11**, 053017 (2009).

[51] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard, Rev. Mod. Phys. **81**, 1051 (2009).

[52] C. S. Adams, M. Sigel, and J. Mlynek, Physics Reports **240**, 143 (1994).

[53] R. Godun, M. D'Arcy, G. Summy, and K. Burnett, Contemporary Physics **42**, 77 (2001).

[54] D. Pritchard, A. D. Cronin, S. Gupta, and D. Kokorowski, Annalen der Physik **10**, 35 (2001).

[55] T. Berrada et al., Nature communications **4**, 2077 (2013).

[56] M. Fattori et al., Phys. Rev. Lett. **100**, 080405 (2008).

[57] M. Gustavsson et al., Phys. Rev. Lett. **100**, 080404 (2008).

[58] M. Kitagawa and M. Ueda, Phys. Rev. A **47**, 5138 (1993).

[59] G.-B. Jo et al., Phys. Rev. Lett. **98**, 030407 (2007).

[60] A. D. Martin and J. Ruostekoski, New Journal of physics **14**, 043040 (2012).

[61] H. Sakaguchi and B. A. Malomed, New Journal of Physics **18**, 025020 (2016).

[62] F. K. Abdullaev and V. A. Brazhnyi, Journal of Physics B: Atomic, Molecular and Optical Physics **45**, 085301 (2012).

[63] J. Polo and V. Ahufinger, Physical Review A **88**, 053628 (2013).

[64] J. L. Helm, S. L. Cornish, and S. A. Gardiner, Physical review letters **114**, 134101 (2015).

[65] J. L. Helm, S. J. Rooney, C. Weiss, and S. A. Gardiner, Physical Review A **89**, 033610 (2014).

[66] J. L. Helm, T. P. Billam, and S. A. Gardiner, Physical Review A **85**, 053621 (2012).

[67] S. A. Haine, New Journal of Physics **20**, 033009 (2018).

[68] A. L. Marchant, S. Händel, S. A. Hopkins, T. P. Wiles, and S. L. Cornish, Phys. Rev. A **85**, 053647 (2012).

[69] O. J. Wales et al., arXiv e-prints , arXiv:1906.06083 (2019).

[70] S. Händel, A. L. Marchant, T. P. Wiles, S. A. Hopkins, and S. L. Cornish, Review of Scientific Instruments **83**, 013105 (2012).

[71] S. Händel et al., Phys. Rev. A **83**, 053633 (2011).

[72] W. Petrich, M. H. Anderson, J. R. Ensher, and E. A. Cornell, J. Opt. Soc. Am. B **11**, 1332 (1994).

[73] A. M. Thomas, *Ultra-cold collisions and evaporative cooling of caesium in a magnetic trap*, PhD thesis, University of Oxford, 2004.

[74] S. Chu, L. Hollberg, J. E. Bjorkholm, A. Cable, and A. Ashkin, Phys. Rev. Lett. **55**, 48 (1985).

[75] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn, arXiv e-prints , cond (1999).

[76] W. Ketterle and N. Van Druten, Evaporative cooling of trapped atoms, in *Advances in atomic, molecular, and optical physics*, volume 37, pages 181–236, Elsevier, 1996.

[77] C. G. Townsend et al., Phys. Rev. A **52**, 1423 (1995).

[78] J. L. Roberts, N. R. Claussen, S. L. Cornish, and C. E. Wieman, Phys. Rev. Lett. **85**, 728 (2000).

[79] E. A. Burt et al., Phys. Rev. Lett. **79**, 337 (1997).

[80] V. Bagnato, D. E. Pritchard, and D. Kleppner, Phys. Rev. A **35**, 4354 (1987).

[81] J. P. Burke, Jr., *Theoretical investigation of cold alkali atom collisions*, PhD thesis, University of Colorado at Boulder, 1999.

[82] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, Phys. Rev. A **79**, 063631 (2009).

[83] T. Lauber, J. Küber, O. Wille, and G. Birkl, Phys. Rev. A **84**, 043641 (2011).

[84] W. Hung et al., J. Opt. Soc. Am. B **32**, B32 (2015).

[85] C. Eberlein, S. Giovanazzi, and D. H. J. O'Dell, Phys. Rev. A **71**, 033618 (2005).

[86] C. J. Foot et al., *Atomic physics*, volume 7, Oxford University Press, 2005.

[87] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, Science **269**, 198 (1995).

[88] R. J. Wild, *Contact Measurements on a Strongly Interacting Bose Gas*, PhD thesis, University of Colorado, 2012.

[89] H. J. Lewandowski, *Coherences and correlations in an ultracold Bose gas*, PhD thesis, University of Colorado, 2002.

[90] B. E. Sherlock and I. G. Hughes, American Journal of Physics **77**, 111 (2009).

[91] D. Jenkin, *Feshbach spectroscopy of an ultracold Rb-Cs mixture*, PhD thesis, Durham University, 2012.

[92] M. R. Matthews, *Two-component Bose-Einstein condensation*, PhD thesis, University of Colorado at Boulder, 1999.

[93] D. A. Steck, Rubidium 85 d line data, Available online, 2013.

[94] K. B. MacAdam, A. Steinbach, and C. Wieman, American Journal of Physics **60**, 1098 (1992).

[95] J. E. Debs, N. P. Robins, A. Lance, M. B. Kruger, and J. D. Close, Appl. Opt. **47**, 5163 (2008).

[96] G. Lancaster, R. Conroy, M. Clifford, J. Arlt, and K. Dholakia, Optics Communications **170**, 79 (1999).

[97] V. Negnevitsky and L. D. Turner, Opt. Express **21**, 3103 (2013).

[98] A. Ratnapala et al., Opt. Lett. **29**, 2704 (2004).

[99] J. Appel, A. MacRae, and A. I. Lvovsky, Measurement Science and Technology **20**, 055302 (2009).

[100] S. L. Cornish, N. R. Claussen, J. L. Roberts, E. A. Cornell, and C. E. Wieman, Phys. Rev. Lett. **85**, 1795 (2000).

[101] C. L. Blackley et al., Phys. Rev. A **87**, 033611 (2013).

[102] J. L. Roberts et al., Phys. Rev. Lett. **81**, 5109 (1998).

[103] M. A. Zentile et al., Computer Physics Communications **189**, 162 (2015).

[104] R. S. Bloom, *Few-Body Collisions in a Quantum Gas Mixture of 40K and 87Rb Atoms*, PhD thesis, University of Colorado, 2014.

[105] T. Lu, X. Miao, and H. Metcalf, Phys. Rev. A **71**, 061405 (2005).

[106] K. Agatsuma et al., Opt. Express **22**, 2013 (2014).

[107] K. Hueck et al., Opt. Express **25**, 8670 (2017).

[108] I. E. Mazets, T. Schumm, and J. Schmiedmayer, Phys. Rev. Lett. **100**, 210403 (2008).

[109] T. Billam, A. Marchant, S. Cornish, S. Gardiner, and N. Parker, Bright solitary matter waves: formation, stability and interactions, in *Spontaneous Symmetry Breaking, Self-Trapping, and Josephson Oscillations*, pages 403–455, Springer, 2012.

[110] F. Chevy, V. Bretin, P. Rosenbusch, K. W. Madison, and J. Dalibard, Phys. Rev. Lett. **88**, 250402 (2002).

[111] J. M. Gerton, D. Strekalov, I. Prodan, and R. G. Hulet, Nature **408**, 692 (2000).

[112] A. Di Carli et al., arXiv e-prints (2019).

[113] C. Menotti and S. Stringari, Phys. Rev. A **66**, 043610 (2002).

[114] E. Haller et al., Science **325**, 1224 (2009).

[115] H. Moritz, T. Stöferle, M. Köhl, and T. Esslinger, Phys. Rev. Lett. **91**, 250402 (2003).

[116] J. H. Nguyen, D. Luo, and R. G. Hulet, Science **356**, 422 (2017).

[117] P. J. Everitt et al., Phys. Rev. A **96**, 041601 (2017).

[118] C. Lee and J. Brand, Europhysics Letters (EPL) **73**, 321 (2006).

[119] P. Manju et al., Phys. Rev. A **98**, 053629 (2018).

[120] Z.-Y. Sun, P. G. Kevrekidis, and P. Krüger, Physical Review A **90**, 063612 (2014).

[121] J. Holmer, J. Marzuola, and M. Zworski, Communications in Mathematical Physics **274**, 187 (2007).

[122] L. Landau and E. Lifshitz, *Quantum Mechanics: Non-Relativistic Theory*, Course of Theoretical Physics, Elsevier Science, 1981.

[123] E. J. Olson, S. E. Pollack, D. Dries, and R. G. Hulet, Experiments with a Single Bright Matter-Wave Soliton Incident on a Barrier, in *APS Division of Atomic, Molecular and Optical Physics Meeting Abstracts*, volume 55, page OPE.10, 2010.

[124] M. S. Safronova, B. Arora, and C. W. Clark, Phys. Rev. A **73**, 022505 (2006).

[125] J. Holmer, J. Marzuola, and M. Zworski, Communications in Mathematical Physics **274**, 187 (2007).

[126] I. Hughes and T. Hase, *Measurements and their uncertainties: a practical guide to modern error analysis*, Oxford University Press, 2010.

[127] K. Henderson, C. Ryu, C. MacCormick, and M. Boshier, New Journal of Physics **11**, 043030 (2009).

[128] S. Schnelle, E. Van Ooijen, M. Davis, N. Heckenberg, and H. Rubinsztein-Dunlop, Optics Express **16**, 1405 (2008).

[129] G. Gauthier et al., Optica **3**, 1136 (2016).

[130] L. Amico, G. Birkl, M. Boshier, and L.-C. Kwek, New Journal of Physics **19**, 020201 (2017).

[131] A. I. Streltsov, O. E. Alon, and L. S. Cederbaum, Phys. Rev. Lett. **100**, 130401 (2008).

[132] C. Weiss and L. D. Carr, arXiv e-prints , arXiv:1612.05545 (2016).

[133] L. D. Carr and Y. Castin, Phys. Rev. A **66**, 063602 (2002).

[134] L. Barbiero and L. Salasnich, Phys. Rev. A **89**, 063605 (2014).

[135] E. A. Donley et al., Nature **412**, 295 (2001).

[136] A. D. Martin, C. S. Adams, and S. A. Gardiner, Phys. Rev. Lett. **98**, 020402 (2007).

[137] M. A. Kasevich, D. S. Weiss, and S. Chu, Opt. Lett. **15**, 607 (1990).

[138] V. I. Balykin, V. S. Letokhov, Y. B. Ovchinnikov, and A. I. Sidorov, Phys. Rev. Lett. **60**, 2137 (1988).

[139] R. Côté, B. Segev, and M. G. Raizen, Phys. Rev. A **58**, 3999 (1998).

[140] S. Cornish et al., Physica D: Nonlinear Phenomena **238**, 1299 (2009), Nonlinear Phenomena in Degenerate Quantum Gases.

[141] H. Bender, P. W. Courteille, C. Marzok, C. Zimmermann, and S. Slama, Phys. Rev. Lett. **104**, 083201 (2010).