


Matter Wave Interference on an Atom Chip

Rob Sewell



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

Imperial College London

University of London

June 29, 2009



Abstract

In this thesis I report on building an atom chip designed to operate as a matter wave interferometer. I show that we are able to coherently split a single Bose Einstein condensate (BEC) in two using a radio-frequency field to deform a magnetic trap smoothly from a single to a double well. We can then read this out from the interference pattern that is produced when they are released from the trap and allowed to overlap.

Acknowledgments

There are far too many people who have assisted me along the way to acknowledge properly here. However there are a few people who deserve particular thanks.

First and foremost my supervisor Prof. Ed Hinds has continually offered guidance throughout my time at Imperial College. I have appreciated his keen physical insight many times during our regular meetings. I have also appreciated the effort he has put in to help produce this document, particularly during the final stages of writing.

Dr. Jos Dingjan has similarly helped me enormously in producing this thesis, and in his guidance as the Post-Doc on the experiment over the past few years. I offer him heartfelt thanks.

My first Post-Doc was Dr. Stefan Eriksson, who has since moved to become a lecturer in Wales. He taught me all I needed to be able to run the experiment, and has regularly provided guidance in the years since, for which I want to thank him.

Florian Baumgärtner is the current student running the experiment. He has been a revelation in the lab, systematically gathering data about the interferometer. He was involved in gathering much of the data presented in this thesis. I would like to acknowledge his efforts here.

To the entire CCM group, I have enjoyed working with all of you and the friendships made along the way.

Finally my beautiful partner Alai has put up with much of the stress of writing a thesis without any of the benefits, and has been a constant source of encouragement and inspiration throughout it all. To her goes all my love and thanks.



Contents

1	Introduction	9
1.1	Our Experiment	10
1.2	Thesis Organisation	13
2	Some Background Theory	14
2.1	Introduction	14
2.2	Magnetic Trapping of Neutral Atoms	14
2.2.1	Trapping Configurations	16
2.2.2	Wire Trap	18
2.3	Bose Einstein Condensation	23
2.3.1	Noninteracting Bose Gas	23
2.3.2	Interacting Gas at Zero Temperature	26
2.3.3	Gross-Pitaevskii Equation	27
2.3.4	Elongated Condensates	31
2.3.5	Phase Fluctuations	33
2.4	Absorption Imaging	35
3	Experimental Apparatus	37
3.1	Laser System	37
3.2	Vacuum System	41
3.3	Top Flange Assembly	45
3.4	External Fields	48
3.5	Atom Chip	49
3.6	Current Control	54
3.7	Radio Frequency Generation	56
3.8	Experiment Control	57
3.9	Imaging System	57

4	Making BECs	60
4.1	Experimental Sequence	60
4.1.1	Low Velocity Intense Source	60
4.1.2	Mirror MOT	61
4.1.3	UMOT	63
4.1.4	Optical Molasses	65
4.1.5	Optical Pumping	65
4.1.6	Initial Magnetic Trap	66
4.1.7	Compressed Trap	66
4.1.8	Evaporative Cooling	67
4.2	Cold Thermal Clouds	67
4.2.1	Fragmentation	68
4.3	Bose Einstein Condensation	71
4.3.1	BEC Transition	72
4.3.2	BEC Parameters	72
4.3.3	Phase Fluctuations	78
5	RF Adiabatic Potentials	82
5.1	Introduction	82
5.2	Theoretical Approaches	86
5.2.1	Dressed Atoms	86
5.2.2	Full Numerical Calculation	90
5.3	Calculating Realistic Trapping Potentials	92
5.3.1	Accuracy of Analytical Approximations	94
5.3.2	Role of Different Polarisation Components	97
5.3.3	Effect of Current Imbalance	99
5.3.4	Effect of A Relative Phase in the RF Current in the Two Wires	101
6	Theory of Matter-Wave Interference	103
6.1	Matter Wave Interference	103
6.2	Two mode approximation	104
6.2.1	Number and Phase States	105
6.2.2	Correlation Functions	107
6.3	Phase Spreading	109
6.4	Longitudinal Phase Fluctuations	110
7	Matter Wave Interference	112
7.1	Introduction	112
7.2	Implementing the Double Well Potential	112
7.3	Data Analysis	117
7.3.1	Extracting Images	117
7.3.2	Density Modulation	118
7.3.3	Circular Statistics	122
7.3.4	Finite Imaging Resolution	124
7.4	Coherent Splitting	125
7.5	Some observations about BEC interference	129

7.5.1	The Effect of Interactions	130
7.5.2	Variation of Condensate Size	134
7.5.3	Longitudinal Phase Fluctuations	136
7.5.4	Interference of Zeeman Sublevels	137
8	Outlook	140
	Bibliography	141



List of Figures

1.1	Atom Chip Layout	11
2.1	Z-Wire Trap	18
2.2	Finite Wire	20
2.3	Rough Wire	21
3.1	^{87}Rb D2 line	38
3.2	Laser System Schematic	39
3.3	Vacuum System	42
3.4	Extraction Mirror	43
3.5	Vacuum Flange Assembly	45
3.6	Chip mount	46
3.7	Atom chip.	50
3.8	Atom chip fabrication	52
3.9	Scanning Electron Micrographs	53
3.10	Full Circuit	55
3.11	Absorption Imaging Set Up	58
4.1	LVIS	61
4.2	Mirror MOT	62
4.3	UMOT	63
4.4	Double Z Wires Trap	66
4.5	Fragmentation of Cold Thermal Clouds	68
4.6	Temperature and Magnification	69
4.7	Potential Roughness	70
4.8	Evaporation to BEC	73
4.9	BEC in Freefall	74
4.10	Static Trap Configuration	75
4.11	Trap Frequencies	76
4.12	Trap Minimum & Condensate Lifetime	77
4.13	Condensate Density Profile	77

4.14	Emergence of Phase Fluctuations	79
4.15	Elongated BEC	80
5.1	RF Adiabatic Potential	82
5.2	RF Adiabatic Double Well Potential	84
5.3	RF Adiabatic Potentials for the $F = 2$ System	86
5.4	Levels Coupled by RF Field Components	88
5.5	RF Adiabatic Potential Geometry	92
5.6	RF Coupling	93
5.7	Double Well Potential	94
5.8	Analytic Approximations	95
5.9	Double Well Potential Parameters	97
5.10	Role of Polarisation Components	98
5.11	Gradient of RF Field Amplitude	99
5.12	Balanced Potential	100
5.13	Effect of Imbalancing the RF Currents	101
5.14	RF Polarisation	101
7.1	Double Well Potential	113
7.2	In-trap Images of Separated BECs	115
7.3	Balancing the Double Well Potential	116
7.4	Interference	117
7.5	Fourier Filtering	118
7.6	Fringe Fitting	119
7.7	Filtered Density Modulation	120
7.8	Finite Resolution and Fringe Visibility	124
7.9	Phase Distribution	126
7.10	Phase Drift	127
7.11	Drift of Cloud Centre	127
7.12	Fringe Fit Data	128
7.13	Ensemble Average	129
7.14	GPE	131
7.15	Fringe Wavelength	132
7.16	Fringe Wavelength	133
7.17	Phase vs Well Separation	133
7.18	BECs	135
7.19	Fringe Visibility	135
7.20	Interference with Elongated BEC	136
7.21	Phase Distribution	137
7.22	Interference of Zeeman Sublevels	138
7.23	Correlation of Phases	139

Introduction

In this thesis I report on building an atom chip designed to operate as a matter wave interferometer. I show that we are able to coherently split a single Bose Einstein condensate (BEC) in two using a radio-frequency field to deform a magnetic trap smoothly from a single to a double well. We can then read this out from the interference pattern that is produced when they are released from the trap and allowed to overlap.

The double well BEC system has been the subject of enormous interest both from a theoretical and from an experimental standpoint. One of the earliest experiments attempted with a Bose Einstein condensate, soon after the phase transition itself was first experimentally realised, was to split the BEC in two in order to study the coherence of the system [1]. Interference between the two modes of the split condensate was observed in this experiment, but the splitting process itself was not coherent in the sense that it introduced a random relative phase between the two modes of the condensate. To date only a few groups worldwide have reported being able to split a single BEC *coherently* using a double well potential. The groups that have done so have either employed optical potentials [2] [3] [4] or RF dressed state potentials on an atom chip [5] [6]. The first such experiment was successfully performed only a few years ago.

We follow the approach pioneered by Schumm et al. [5] in which atoms trapped near the surface of an atom chip are dressed by RF fields produced by current conducting wires patterned on the surface of the atom chip. The close proximity of the atoms to the chip allows strong RF fields to be applied easily and with accurate control of the polarisation. By carefully controlling the amplitude, frequency and polarisation of the RF field, we gain precise control over the parameters of the trap and can deform it with great accuracy from a single to a double well.

Coherent splitting of a BEC using a double well potential is in many respects the equivalent of an optical beam splitter for a laser, and forms the basic building block of our matter wave interferometer. Using the double well potential, it is possible to implement

various different interferometric schemes using a BEC, such as Michelson and Sagnac type interferometers. The advantage of using a BEC over a laser is that the BEC is many orders of magnitude more sensitive to the small forces (such as gravity or a rotational acceleration) that we might be interested in measuring. Of course this sensitivity makes it much more difficult to implement a BEC interferometer. It is possible to use other methods to make a beam splitter for a BEC (e.g. via Bragg scattering from a resonant standing wave laser), but these are generally implemented in free fall, after the BEC has been released from its trapping potential. We would ultimately like to carry out the entire interferometry process with a trapped BEC, which would allow us to take advantage of a much longer interrogation time to increase the sensitivity of the interferometer.

There has been significant progress towards making a working trapped matter wave interferometer based on a double well scheme over the past few years. It has become possible to read out a small relative phase shift introduced to the system via measurements of the interference pattern seen in free fall [2] [6] and via a phase sensitive in-trap recombination scheme [7]. Studies of the statistics of interference patterns have also led to a form of noise thermometry for the double well system [8]. The underlying Josephson dynamics of a coupled double well system have also been observed [9], including non-linear self-trapping [3], and both the AC and DC Josephson effects [4]. Number squeezing due to atom-atom interactions in the splitting process has also been observed [6] [10]. This raises the possibility of greatly increasing the length of time during which the double well system remains phase coherent and perhaps engineering a system with sub-shot-noise limited sensitivity by exploiting entanglement between the two modes of the system. Atom chip based interferometers have also been able to study the coherence dynamics of very elongated BEC systems. At low temperature, these systems can cross over into a regime where the behaviour of the system is one-dimensional. This fundamentally changes the physics of the condensate, which can be seen in the the noise statistics (quantum and thermal) of the system [11], and in the relaxation dynamics of phase fluctuations that emerge along the length of the condensate [12]. The outlook for future progress thus looks very promising.

1.1 Our Experiment

A brief word of introduction to our experiment is worthwhile. Ours is an atom chip BEC experiment. An atom chip is a micro-fabricated device used to trap and manipulate ultracold atoms. The atoms are held trapped in free space close to the surface of the atom chip. In our case, the fields required to trap the atoms are produced by running DC and AC currents through gold wires that are patterned onto the surface of the chip. Other

atom chips have used permanent magnetic structures or static electric fields to trap the atoms [13]. Optical fibres, waveguides and cavities can be incorporated onto the atom chip to provide light for detecting and interacting with the atoms [14]. Ultimately laser diodes, detection devices such as charge-coupled device (CCD) arrays, optical elements, and everything required to run a BEC may one day be integrated into the same device.




Figure 1.1: Atom Chip Layout: Diagram of the layout of wires patterned on the surface of our atom chip. There are four parallel wires in the centre of the chip that are used to produce the necessary DC and RF magnetic fields for trapping and manipulating Bose Einstein condensates. The outer pair are $100\mu m$ wide and have a centre to centre separation of $300\mu m$. The inner pair are $50\mu m$ wide, and are separated by $85\mu m$. The ends of these wires are bent into a Z-shape to provide the field curvature along the length of the wires needed to close the trap along this axis. The diagram is to scale except that the length is considerably shortened. In the real atom chip the centre section of the wires is $7mm$ long. In addition to these trapping wires, two end wires are patterned onto the chip. These run perpendicular to the central trapping wires, and are there both to provide additional axial fields and to be able to run current in a net U-shape, which is used to make the quadrupole field for a magneto-optical trap that is aligned with the trapping fields of the chip wires. The principal axes of the trapping geometry are indicated, and the important external uniform bias fields necessary to produce the magnetic trap are illustrated. The chip is designed to carry DC and RF currents as indicated in the diagram, and the double well produced by the RF adiabatic potential is oriented so that the BEC splits along the x -axis of the trap, perpendicular to the surface of the chip as indicated by the two parallel clouds represented in the centre of the diagram.

The atom chip we use is illustrated in figure 1.1. There are four parallel wires in the centre of the chip that are used to produce the necessary DC and AC magnetic fields for trapping and manipulating Bose Einstein condensates. The outer pair are $100\text{ }\mu\text{m}$ wide and have a centre to centre separation of $300\text{ }\mu\text{m}$. The inner pair are $50\text{ }\mu\text{m}$ wide, and are separated by $85\text{ }\mu\text{m}$. The ends of these wires are bent into a Z-shape to provide the field curvature along the length of the wires needed to close the trap along this axis. The diagram is to scale except that the length is considerably shortened. In the real atom chip the centre section of the wires is 7 mm long. In addition to these trapping wires, two end wires are patterned onto the chip. These run perpendicular to the central trapping wires, and are there both to provide additional axial fields and to be able to run current in a net U-shape, which is used to make the quadrupole field for a magneto-optical trap that is aligned with the trapping fields of the chip wires. The principal axes of the trapping geometry are indicated, and the important external uniform bias fields necessary to produce the magnetic trap are illustrated. The chip is designed to carry DC and RF currents as indicated in the diagram, and the double well produced by the RF adiabatic potential is oriented so that the BEC splits along the x -axis of the trap, perpendicular to the surface of the chip as indicated by the two parallel clouds represented in the centre of the diagram.

In order to run the experiment, we also need an auxiliary source of ultracold atoms. In our experiment this comes from a dual magneto-optical trap (MOT) system consisting of a source MOT in a standard 6-beam configuration, which sends a beam of cold atoms towards a 4-beam surface or mirror MOT that exploits the reflective gold surface of our atom chip to trap atoms close to the chip wires, where they can be readily transferred to a magnetic trap. We refer to the source MOT as the LVIS, since it forms a Low-Velocity-Intense-Source of cold atoms for the mirror-MOT. The mirror-MOT uses a quadrupole field generated by a set of anti-Helmholtz coils that are placed at a 45° angle to the atom chip with the centre aligned a short distance from the chip surface. We also use a third MOT in the process of transferring atoms from the mirror-MOT to a chip-wire static magnetic trap. This is referred to as a UMOT, since the quadrupole field in this stage is provided by DC currents running through the chip wires. These currents are arranged so that the net DC current runs in a U-shape (effectively entering along one arm and running through the central section of a Z-wire, then exiting via one of the end wires illustrated in figure 1.1) which, with the addition of an external uniform bias field, creates a quadrupole field that is automatically aligned with the chip wires. These stages of the experiment are described in detail in chapter 4.

1.2 Thesis Organisation

This thesis is organised as follows: in chapter 2 I present a brief overview of the theory of trapping and manipulating ultracold atoms using static magnetic fields, and in particular some of the real-world effects of trapping atoms near the current carrying wires of an atom chip. I also introduce some of the theory of Bose-Einstein condensation, paying particular attention to topics relevant to describing the elongated condensates that we make in our experiments. This introduction provides the necessary background for understanding how we produce and characterise Bose Einstein condensates in our set-up. The experimental apparatus itself is described in chapter 3. Then in chapter 4 I describe how we make BECs using this apparatus, giving a brief description of each important step in a typical experimental sequence. I also pay careful attention to characterising the condensates that we make. These three chapters together make up the first half of the thesis, and provide the starting point for our interference experiments.

In chapter 5 I describe the RF dressed state potentials (also called RF adiabatic potentials) that we use to change smoothly from a single to a double well potential to split the BEC in two coherently. The theory of neutral atoms interacting with intense RF fields has a long history, but the idea of using RF adiabatic potentials to trap and manipulate atoms is relatively new. In this chapter I draw on some earlier literature to describe analytic approximations of the dressed state potentials beyond the standard rotating wave approximation, and discuss the full numerical calculation of realistic potentials, paying attention to the role of the different polarisation components of the RF field and to the effects introduced by using non-uniform RF fields to dress the atoms. Chapter 6 is an overview of the theory of matter wave interference in a double well system. It provides the necessary background for understanding our interference experiments. In chapter 7 I then present the data from our first interference experiments, and show that we are able to coherently split a single BEC in two using the RF adiabatic potential. This is the key result of this thesis. I also briefly discuss some more qualitative observations made in our early experiments.

Some Background Theory

2.1 Introduction

This chapter provides some basic background theory necessary for the following experimental chapters that describe how we make and characterise Bose Einstein condensates with our apparatus. It is organised as follows: In section 2.2 I give a brief overview of the principles behind trapping neutral atoms in static magnetic fields, and discuss in particular how we can make a Ioffe-Pritchard trap using an atom chip Z-wire. I also discuss how we calculate the magnetic field from a realistic wire of finite width and length, and the effect of imperfections in the wire on the fields it can generate. Then in section 2.3 I give a brief overview of the theory of Bose-Einstein condensation required to understand the characteristics of the condensates that we use in our interference experiments. Finally in section 2.4 I summarise how we extract data from absorption images of cold atomic clouds.

2.2 Magnetic Trapping of Neutral Atoms

A Note on Units

At the typical energy scale of cold atom experiments SI units are not a natural choice. I will typically express magnetic fields in terms of gauss and energy in terms of either the angular frequency or μK . We can easily convert between energy units using the Plank constant, $h = 6.6262 \times 10^{-34} \text{Js}$, and the Boltzmann constant, $k_B = 1.3807 \times 10^{-23} \text{J/K}$, where $E = k_B T = \hbar \omega = \mu_B B$ in temperature, frequency and magnetic field units respectively. The natural length scale for both BECs and the wires patterned onto the atom chip is μm . DC and AC electric currents are best expressed in mA.

Rubidium-87

We trap ^{87}Rb in the $5^2\text{S}_{1/2}$ electronic ground state of the atoms. Laser cooling operates on the D_2 line which couples $5^2\text{S}_{1/2}$ to $5^2\text{P}_{3/2}$. This line has a cycling transition, needed to scatter enough photons to effectively cool an ensemble of atoms to μK temperatures. The wavelength of the transition is $\sim 780\text{nm}$ with a natural linewidth $\Gamma \simeq 2\pi \times 6\text{MHz}$ and a Doppler temperature $T_D = 146\mu\text{K}$.

The $5^2\text{S}_{1/2}$ ground state is further split into two levels due to the hyperfine coupling of the total angular momentum of the electron $\mathbf{J}$ with the total nuclear angular momentum $\mathbf{I}$. The total angular momentum of the atom is $\mathbf{F} = \mathbf{J} + \mathbf{I}$, so the $5^2\text{S}_{1/2}$ state is split into $|F = 1\rangle$ and $|F = 2\rangle$ levels. Each state has a set of degenerate m_F sublevels. In ^{87}Rb the hyperfine splitting of the ground state is 6.8GHz [15]. We work with atoms optically pumped into the stretched state $|F = 2, m_F = 2\rangle$.

Some Orders of Magnitude

For ^{87}Rb atoms in the $|F = 2, m_F = 2\rangle$, a trap depth of 1G corresponds to a temperature of $67\mu\text{K}$, which is why some form of pre-cooling is required before we can magnetically trap an ensemble. Since the Doppler temperature of the laser cooling D_2 transition is $146\mu\text{K}$, we can readily trap a laser cooled ensemble with a magnetic trap with a depth of $10\text{-}15$ gauss.

The typical transition temperature for a BEC is of the order of $\sim 0.5\mu\text{K}$. The appropriate energy scale for both the BEC chemical potential and the magnetic trap parameters, such as trap frequencies and barrier height, is in terms of kHz .

The Zeeman shift of the $|F = 2, m_F = 2\rangle$ state is 1.4MHz G^{-1} . The energy shift due to gravity (a typical small force that we want to measure with the interferometer) is $\sim 2.1\text{kHz } \mu\text{m}^{-1}$. The magnetic field gradient at a distance of $100\mu\text{m}$ from a wire carrying a 1A current is $0.2\text{G } \mu\text{m}^{-1}$, which corresponds to $0.28\text{kHz } \mu\text{m}^{-1}$ for a ^{87}Rb atom in the $|F = 2, m_F = 2\rangle$ state.

Zeeman Interaction

In a static magnetic field the degeneracy of the hyperfine states is lifted. If the energy shift is small compared to the hyperfine splitting then F is a good quantum number and the interaction Hamiltonian is

$$H_{RF} = \mu_B g_F \mathbf{F} \cdot \mathbf{B}(\mathbf{r}), \quad (2.1)$$

where μ_B is the Bohr magneton and g_F the hyperfine Landé g-factor. The magnetic field

lifts the degeneracy of the Zeeman sublevels $-F \leq m_F \leq F$, which are classically interpreted as the projection of the angular momentum onto the local magnetic field. If the energy shift is small then m_F remains a good quantum number. To lowest order the Zeeman sublevels shift linearly according to

$$\Delta E_{F,m_F} = g_F \mu_B m_F B \quad (2.2)$$

where $\mu_B m_F g_F$ is the magnetic moment of the atoms (anti-)aligned with the field, $B = |\mathbf{B}(\mathbf{r})|$. For simplicity we will always follow the convention that the static magnetic field is aligned along $\hat{\mathbf{z}}$, which is the (local) quantisation axis for the atomic magnetic moment. It is this linear Zeeman shift that is exploited to magnetically trap neutral atoms. Equation (2.2) is simply a potential energy term in the Hamiltonian of the atom moving in a static magnetic field

$$H = \frac{\mathbf{p}^2}{2m} + V. \quad (2.3)$$

where $V = \mu_B m_F g_F B$. We assume the atomic quantisation axis adiabatically follows the orientation of the local magnetic field. We work with atoms optically pumped into the $|F = 2, m_F = 2\rangle$ stretched state.

Adiabaticity

The adiabatic approximation of equation (2.3) holds as long as the atomic spin can follow the direction of the magnetic field vector. Classically this means that

$$\frac{d\theta}{dt} < \omega_L \quad (2.4)$$

where θ is the angle between the spin and the local field vector, and $\omega_L = g_F \mu_B B / \hbar$ the Larmor precession frequency of the atom in the magnetic field. In regions where B vanishes, this condition can no longer be satisfied and atoms can undergo Majorana spin flip transitions, leading to the loss of atoms from a trap. In practice even a small residual field is enough to avoid Majorana losses on the time scale of the experiments that we perform.¹

2.2.1 Trapping Configurations

The potential energy term of equation (2.3) is proportional to the absolute value of the static magnetic field. The sign of the $m_F g_F$ prefactor determines the alignment of the

¹The loss rate scales as $\exp(-\frac{\omega_L}{\bar{\omega}})$ where $\bar{\omega}$ is the geometric mean trapping frequency. Typical numbers for our traps are $\omega_L = 0.7$ MHz and $\bar{\omega} = 480$ Hz, so the exponential term is negligibly small.

atomic spin with the local magnetic field vector. Aligned states with $m_F g_F < 0$ are attracted to regions of high field (*high field seekers*), while anti-aligned states with $m_F g_F > 0$ are attracted to regions of low field (*low field seekers*). Since Maxwell's equations do not allow for local maxima in free space, high field seekers cannot be trapped using static magnetic fields. In the $5^2S_{1/2}$ electronic ground state the trappable low field seeking states are $|F = 1, m_F = -1\rangle$ and $|F = 2, m_F = 1, 2\rangle$. These states are not the absolute ground state of the atom, so trapped atoms are in meta-stable states that can decay via inelastic 2- or 3- body spin-flip collisions. However for s-wave collisions, inelastic 2-body processes are strongly suppressed in the stretched states because of angular momentum conservation.

The lowest-order multipole trapping configuration, and the one offering the strongest confinement, is a quadrupole field, which creates a linear trapping potential. However a quadrupole necessarily has a zero magnetic field at which point the magnetic moment of the atom cannot adiabatically follow the change in field direction. The lowest order stable trap is thus a 3D harmonic potential

$$V(\mathbf{r}) = \frac{1}{2}m \sum_{i=x,y,z}^3 \omega_i x_i^2. \quad (2.5)$$

Ioffe-Pritchard Trap

A simple magnetic field configuration that is 3D harmonic near the origin is the Ioffe-Pritchard trap, first used in plasma confinement and then adapted to magnetic trapping of cold atoms. The Ioffe-Pritchard configuration combines a linear (2D) magnetic quadrupole with a longitudinal homogeneous field to give a non-zero minimum field. Longitudinal confinement comes from an additional curvature superimposed on the longitudinal field. The full Ioffe-Pritchard field is

$$\mathbf{B}_{IP}(\mathbf{r}) = B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + B' \begin{pmatrix} x \\ -y \\ 0 \end{pmatrix} + \frac{1}{2}B'' \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{1}{2}(x^2 + y^2) \end{pmatrix} \quad (2.6)$$

where $B' = \frac{\partial|B|}{\partial x} = \frac{\partial|B|}{\partial y}$ and $B'' = \frac{\partial^2|B|}{\partial z^2}$. The Ioffe-Pritchard trap is locally harmonic with trapping frequencies

$$\omega_{\perp} = \sqrt{\frac{g_F \mu_B m_F}{m} \frac{B'^2}{B_0}} \quad \text{and} \quad \omega_z = \sqrt{\frac{g_F \mu_B m_F}{m} B''}. \quad (2.7)$$

The harmonic approximation is very good in the longitudinal direction. In the transverse direction the potential is linear once $\sqrt{x^2 + y^2} > B_0/B'$. Since we have an axially symmet-

ric trapping geometry with $\omega_{\perp} \gg \omega_z$ we can usefully rewrite equation (2.5) in cylindrical coordinates centred on the trap minimum as

$$V(\mathbf{r}) = \frac{1}{2}m(\omega_{\perp}^2 \rho^2 + \omega_z^2 z^2). \quad (2.8)$$

where $\rho = \sqrt{x^2 + y^2}$.

2.2.2 Wire Trap




Figure 2.1: Z-Wire Trap: This figure illustrates the basic idea of how a Ioffe-Pritchard trap can be made with a single wire bent in a Z-configuration. The wire carries a DC current I . The field induced by this current is cancelled along the length of the central section of the wire at a height d by a uniform external bias field B_{bias} , forming a quadrupole guide along the wire. This guide is closed off along the axis by the fields from the ends of the wire, sketched in red in the figure. The curvature from these fields provides axial confinement, and they add to give a non-zero field at the height of the trap minimum. This field is augmented by another uniform external bias field B_0 which adjusts the magnitude of the field at the trap bottom.

The Ioffe-Pritchard configuration can be readily implemented on an atom chip using a single wire running a DC current I and an external uniform bias field perpendicular to the wire. Figure 2.1 is a schematic diagram illustrating such a trap. In our geometry the current I runs along $\hat{x}$ and $\hat{z}$, and the bias field along $\hat{x}$ (B_{bias}). The bias field will cancel the wire field along a line at a distance d above the wire, forming a 2D quadrupole with gradient B' giving tight confinement in the radial direction. These parameters are given by

$$d = \frac{\mu_0}{2\pi} \frac{I}{B_{bias}} \quad \text{and} \quad B' = \frac{\mu_0}{2\pi} \frac{I}{d^2} \quad (2.9)$$

Axial confinement is provided by provided the current along $\hat{x}$. Since both ends of the wire run current in the same direction, they produce a non-zero minimum of B_z at the centre of the trap. We typically add an additional uniform field B_0 along the $\hat{z}$ to increase this minimum to field to about 1 G.

If the central section of the wire is sufficiently long, the field from the ends of the Z-wire can be approximated as a quadrupole in the $\hat{y}$ - $\hat{z}$ plane with a gradient $B'_{ends} = \frac{\mu_0 I}{2\pi L^2}$ where L is the half-length of the central section of the wire. The effect of the quadrupole is to shift the trap minimum towards the wire and to rotate the longitudinal axis of the trap in the $\hat{x}$ - $\hat{z}$ plane. The trap minimum $\tilde{d}$ and rotation angle ϕ are given by

$$\tilde{d} = d - \frac{B'_{ends}(B_0 + B'_{ends}d)}{B'^2 + B'^2_{ends}} \quad \text{and} \quad \tan(\phi) = \frac{d^2}{L^2}. \quad (2.10)$$

Typically $d \ll L$ and $B'_{ends} \ll B'$ so these corrections are very small. With respect to the rotated axes, the curvature of the trap in the radial and longitudinal directions is given (to first order) by

$$\frac{\partial^2 B}{\partial \rho^2} = \left(\frac{2\pi}{\mu_0 I}\right)^2 \frac{B_{bias}^4}{B_0} \quad \text{and} \quad \frac{\partial^2 B}{\partial z^2} = 48 \left(\frac{\mu_0 I}{\pi}\right) \frac{d}{L^4}. \quad (2.11)$$

Note that B' and B_{bias} are fixed by the choice of d - see equation (2.9). The trap aspect ratio is

$$\frac{\omega_{\perp}}{\omega_z} = \sqrt{\frac{B_{bias}}{6B_0}} \frac{L^2}{d^2}. \quad (2.12)$$

Typical numbers corresponding to our experimental parameters ($I = 2$ A, $d = 150$ μm , $L = 7$ mm and $B_0 = 1$ G) are $\omega_{\perp} = 2\pi \times 2.2$ kHz and $\omega_z = 2\pi \times 3$ Hz. The radial trap frequency we measure in the lab is consistent with these numbers, corrected for the finite width of the trapping wires (see section 2.2.2 below). However the axial trapping frequency is considerably higher, which is a result of anomalous field components introduced by small deviations in the current flowing through inhomogeneous wires (see section 2.2.2 below).

Finite Size Effects

The above calculations were carried out to first order for infinitesimally thin ideal wires and do not take into account the effect of the finite size of realistic wires on the magnetic fields. We expect finite size effects to become important when the distance d separating the trap minimum and the wire becomes comparable to the wire dimensions. The relevant wire dimensions are illustrated in figure 2.2. Our chip wires have a thickness $h = 3$ μm and widths $w = 50$ and 100 μm . We typically work at distances $d = 150 - 200$ μm to the chip surface, so the effect of the finite wire thickness is negligible but we do need to take into account the finite width of the wire to accurately calculate trapping potentials.

The magnetic field of an infinitely long wire of finite width w and zero thickness carrying current I can be calculated analytically [16] [17]. It is




Figure 2.2: *Finite Wire:* This diagram illustrates the effect of the finite wire dimensions on the trapping field of a realistic wire. Atoms are trapped at a distance d above the wire, which has a height h and width w . In our experiment $h \ll d$ and we can ignore the effect of the finite thickness on the field at the position of the atoms. However $d \simeq 1 - 3w$ (depending upon the trap used) and we must take the finite width into consideration. The field produced by an infinitely long finite width wire can be calculated as illustrated in figure 2.2(b). The length of the central section of our wires is much longer than the length of our condensate, so we can safely approximate the fields from the central section of the trapping wires in this way.

$$\mathbf{B}(x, y) = \frac{\mu_0 I}{2\pi w} \left(\ln\left(\frac{r_1}{r_2}\right) \hat{\mathbf{y}} + (\theta_1 - \theta_2) \hat{\mathbf{x}} \right) \quad (2.13)$$

At our distance of closest approach $d \simeq 1.3w$, the difference in the field amplitude is less than 5%, and the difference in field gradient about 10%. The radial trap frequency is proportional to the field gradient (see equation (2.7)), so the finite wire width contributes roughly a 10% correction to the important trap parameters dependent on trap frequency, such as the barrier height and well separation in the double well potential.

We also calculate full numerical trap potentials for our wire geometry taking into account the finite wire thickness and lengths using the Radia plug-in for Mathematica.¹ In general their additional corrections to the trap parameters are negligible.

Edge Roughness

One of the early observations from atom chip experiments was that cold atomic ensembles broke up into fragments when brought close to the surface of the trapping wires [18] [19] [20]. On axis the trapping potential is determined by the field parallel to the wire. The fragmentation is due to anomalous field components $\delta B_z(z)$ that raise or lower this potential minimum, so that the ensemble is trapped in a rough potential landscape. These field components are caused by transverse components $\pm \delta I_z(z)$ in the current flowing

¹A package developed by the European Synchrotron Radiation Facility and available online at <http://www.esrf.eu/Accelerators/Groups/InsertionDevices/Software/Radia>




Figure 2.3: Illustration of a thin wire with fluctuating edges. The width, w , and height, h , of the wire are indicated. Boundary conditions on the current density require that the current follows the fluctuations along the edge of the wire, introducing current components $\pm \delta I_x$ which produce magnetic field components $\pm B_z(z)$ along z at the position of the atoms at a distance, d , above the wire. On axis, the $\hat{z}$ component of the field defines the energy of the trap, so these components lead to a rough potential. When the magnitude of the roughness is comparable to the temperature of an ensemble of thermal atoms (or the chemical potential of a BEC), the cloud will fragment into lumps along the length of the wire. The edge fluctuations in this figure are much larger than in a typical wire.

through the trapping wire introduced by imperfections in the wire which cause the current to deviate [21] [22]. The resulting potential minimum at a point along the wire is proportional to $|B_0 + \delta B_z(z)|$. When the magnitude of the anomalous field is comparable to the temperature of an ensemble of thermal atoms $\mu_B \delta B_z(z) > k_B T$ (or $\mu_B \delta B_z(z) > \mu$, the chemical potential of a BEC) the ensemble will break up into fragments.

The origin of these transverse current components has now been well studied. There are three possible sources of transverse current components: fluctuations in the edge of the wire, surface roughness, and bulk defects. Edge fluctuations typically dominate at distances large compared to the width of the wire, particularly for wires, such as ours, that are thin and wide [23]. Surface roughness and bulk defects are more important close to the wire, where the anomalous component due to edge fluctuations tends to saturate. As wire fabrication techniques have improved, allowing atoms to be trapped much closer to the trapping wire, the effects of surface roughness and bulk inhomogeneity have also been

studied [24] [25]. In our experiments atoms are trapped at a distance large compared to the height of the wires, $d \gg h$. At this distance we expect edge fluctuations to dominate, so we can ignore the contribution of current components along $\hat{y}$ and the contribution from surface and bulk defects.

The effect of edge fluctuation on the current flowing through the wire is illustrated in figure 2.3. Assuming that the fluctuations in the edges of the wire are small compared to the width of the wire, fluctuations in the current $\delta I_{(x,z)}$ around the average current I_0 can be readily calculated. We can express the fluctuations in the position and width of the wire in terms of Fourier components $f_k^\pm = \frac{1}{2} \int_{-\infty}^{\infty} dz e^{ikz} f_R(z) \pm f_L(z)$, where $f_{R(L)}(z)$ are the fluctuations in the right (left) edge of the wire. The fluctuations in the current follow from the boundary condition that no current flow out of the wire, so that the current locally follows the edges of the wire. The anomalous magnetic field then follows by direct integration from the Biot-Savart law. Assuming that the width of the condensate is much smaller than d , the rough potential comes from the z-component of the anomalous magnetic field $\delta V = g_F \mu_B m_F \delta B_z(d, 0, z)$ [26] [23] [27].

Under these conditions the power spectrum of the anomalous magnetic field is $B_k = (g_F \mu_B m_F)^2 B_{bias}^2 R_{k,d}^2 F_k$, where B_0 is the field at the trap bottom, $B_{bias} = \mu_0 I / 2\pi d$ is the field produced by an ideal wire running current, I , at a distance, d , from the surface, and F_k is the power spectrum of the fluctuations in the edge of the wire. The function, $R_{k,d}$ translates this power spectrum into noise in the magnetic field seen by the atoms. An expression for $R_{k,d}$ was first calculated in [26]:

$$R_{k,d}^2 = (kd)^2 \frac{2 \sinh[kw/2]}{(kw) \sinh[kw]} \sum_{n=0}^{\infty} \frac{(-1)^n K_{n+1}[kd]}{n! (2kd)^n} [\gamma_{n+1}[\frac{kw}{2}] - \gamma_{n+1}[\frac{-kw}{2}]] \quad (2.14)$$

where $K_{n+1}[kd]$ is the modified Bessel function and $\gamma_{n+1}[x]$ the incomplete Gamma function. The series only converges for distances $d > w/2$, and converges rapidly for $kw \ll 1$.

There are two important length scales that dictate how the fluctuations in the current $\pm \delta I_z(z)$ translate into anomalous field components $\pm \delta B_z(z)$ at the position of the atoms: the separation d between the trap minimum and the wire surface, and the wavelength of a given current component. The magnitude of the potential roughness scales with $d^{2/5}$ for a white noise power spectrum. At distances $d > w$ the term $\sim (kd)^2 K_1(kd)$ dominates. This term rises proportional to $(kd)^2$ at low frequency, decays exponentially at high frequency and peaks at $kd \simeq 1.3$. The physical reason for the cut-off at high frequency is that fluctuations with a wavelength shorter than the distance to the wire will average to zero in the far field. At long wavelengths the angle between the meandering current and the z-axis becomes negligible (the transverse component $I_x \propto k$), which gives

the cut-off at low frequency. If we assume that the spectrum of noise in the current is white, we would thus expect to see a dominant contribution to the anomalous field at a wavelength a little less than the distance of the cold atom ensemble to the wire.¹

2.3 Bose Einstein Condensation

We also need some basic background theory in order to describe the condensates that we use in our experiments. In the following I give a brief overview of what we need to describe our experiment. My major source is Pitaevski and Stringari [28].

2.3.1 Noninteracting Bose Gas

We start with a simple description of a non-interacting Bose gas trapped in a harmonic potential given by equation (2.5), which has a mean trapping frequency $\bar{\omega} = \prod_k (\omega_k)^{1/3}$. In this potential the single particle energy eigenvalues are given by

$$\varepsilon(n_k) = \sum_{k=1}^3 (n_k + \frac{1}{2}) \hbar \omega_k \quad (2.15)$$

where $n_k > 0$. The occupation number of a quantum state is

$$\bar{n}_i = \frac{1}{e^{\beta(\varepsilon_i - \mu)} - 1} \quad (2.16)$$

where $\beta = 1/k_B T$, ε_i is the energy of the i^{th} state, and μ is the chemical potential, fixed by the normalisation condition $N = \sum_i \bar{n}_i$. Equation (2.16) requires that $\mu < \varepsilon_0$, the lowest eigenvalue of the single particle Hamiltonian. When $\mu \rightarrow \varepsilon_0$ the occupation number $N_0 = 1/\exp[\beta(\varepsilon_0 - \mu) - 1]$ of the ground state becomes large. For a large number of particles N , Bose Einstein condensation is characterised by the macroscopic occupancy of the ground state φ_0 in the limit $\mu \rightarrow \varepsilon_0$.

The ground state of the system is obtained by putting all N particles in the single particle groundstate with energy $\varepsilon_0 = \hbar(\omega_x + \omega_y + \omega_z)/2$. The single particle groundstate wave function is

$$\varphi_0(\mathbf{r}) = \left(\frac{m\bar{\omega}}{\pi\hbar}\right)^{3/4} \exp\left[-\frac{m}{2\hbar}(\omega_x x^2 + \omega_y y^2 + \omega_z z^2)\right] \quad (2.17)$$

¹The expected spectrum of noise for realistic wires is in fact not white but self-affine fractal roughness. For wire fabricated like ours, the typical correlation length of this roughness is ~ 20 nm with a mean square roughness of ~ 3 nm. When the distance to the wires is large compared to the correlation length, as in our case, the dominant wavelength is again given by the distance to the wires [27].

which has widths $a_k = \sqrt{\hbar/m\omega_k}$. The overall size of the cloud is given by the harmonic oscillator length

$$a_{ho} = \sqrt{\frac{\hbar}{m\omega_{ho}}}. \quad (2.18)$$

In general not all atoms will be in the groundstate, but there will also be a thermal component to the ensemble (defined as the number of particles that are not in the ground state) $N_{th} = \sum_{i \neq 0} \bar{n}_i$. The normalisation condition is then set by the total number of atoms, $N = N_0 + N_{th}$.

Critical Temperature

In the thermodynamic limit the number of particles in the thermal component can be calculated from equation (2.16) analytically by setting $\mu = \epsilon_0$ and replacing the sum over the states by an integral

$$N = N_0 + N_{th} = N_0 + \int_0^\infty \rho(\epsilon) \frac{1}{e^{\beta(\epsilon-\mu)} - 1} d\epsilon \quad (2.19)$$

where $\rho(\epsilon)$ is the single particle density of states. This can be solve analytically, and gives a critical temperature T_c above which $N_0 = 0$

$$T_c = 0.94 \frac{\hbar\omega_{ho}}{k_B} N^{\frac{1}{3}} \quad (2.20)$$

which occurs at a phase space density $PSD = 2.612$, where the phase space density of an ensemble of particles with a peak number density n_0 is defined as

$$PSD = n_0 \lambda_{dB}^3. \quad (2.21)$$

and λ_{dB} is the thermal deBroglie wavelength of a particle with mass m at temperature T

$$\lambda_{dB} = \sqrt{\frac{2\pi\hbar^2}{mk_B T}}. \quad (2.22)$$

The fraction of particles in the condensate has the simple temperature dependence

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c}\right)^3. \quad (2.23)$$

Both of these results, equation (2.20) and equation (2.23), are derived in the thermodynamic limit $N \rightarrow \infty$ under the condition that the temperature of the ensemble is much greater than the spacing between the single particle energy eigenvalues, $k_B T \gg \hbar\omega_k$. For

a highly anisotropic trap with $\omega_{\perp} \gg \omega_z$ this condition may no longer be satisfied and the description of the BEC becomes more complicated [29].

At finite atom number equation (2.20) is no longer valid. The correction is given by

$$\frac{\delta T_c}{T_c} = -0.73 \frac{\bar{\omega}}{\omega_{ho}} N^{-1/3} \quad (2.24)$$

where $\bar{\omega} = 1/3(\omega_x + \omega_y + \omega_z)$ is the arithmetic mean trapping frequency [30]. The critical temperature is further modified by interactions. To lowest order the shift is given by [31]

$$\frac{\delta T_c}{T_c} = -1.3 \frac{a}{a_{ho}} N^{1/6} \quad (2.25)$$

where $a_{ho} = \sqrt{\hbar/m\bar{\omega}}$ is the harmonic oscillator wavelength. This correction is independent of the shape of the trap as long as the gas remains three dimensional.¹

Density Distribution

Most of our experimental data is extracted from measurements of the number density distribution of the ensemble $n(\mathbf{r}) = n_0(\mathbf{r}) + n_{th}(\mathbf{r})$, normalised to the total number of particles $\int n(\mathbf{r}) = N$. The number density distribution of a thermal Bose gas is $n_{th}(\mathbf{r}) = \int d^3\mathbf{p} n_{\mathbf{p}}(\mathbf{r}) / (2\pi\hbar)^3$, which can be calculated below T_c by setting $\mu = 0$ in the semi-classical particle distribution function $n_{\mathbf{p}}(\mathbf{r}) = 1/(\exp[\beta\epsilon(\mathbf{r}, \mathbf{p}) - \mu] - 1)$. The integration in momentum space gives [28] [35] [36]

$$n_{th}(\mathbf{r}) = \frac{1}{\lambda_{dB}^3} g_{3/2} \left[\tilde{z} e^{-\beta V(\mathbf{r})} \right] \quad (2.26)$$

where $\tilde{z} = e^{\beta\mu}$ is the fugacity of the ensemble and $g_{3/2}[z]$ is a Bose function defined by $g_p[z] = \sum_{l=1}^{\infty} z^l / l^p$. When $T < T_c$ $\tilde{z} = 1$, corresponding to a saturated thermal component. Above T_c we can treat the fugacity $\tilde{z}$ as a fit parameter.

For a thermal cloud in a harmonic trap, the leading term of the Bose distribution is a gaussian, and is close to the Boltzmann distribution of thermal atoms. However the Bose distribution has a higher peak and is slightly narrower in comparison to the Boltzmann distribution due to the bunching effect of indistinguishable bosons. The difference is negligible in the wings of the distribution where the thermal component is distinguishable from the BEC. We can therefore use the Boltzmann distribution to describe the thermal component

¹Equation (2.25) describes a mean-field effect. For a discussion of higher order corrections see the work of Baym et al. [32] [33] [28] [34].

$$n_{th}(\mathbf{r}) = \frac{N_{th}}{(2\pi)^{3/2}} \prod_{k=1}^3 \frac{1}{\sigma_k} \exp \left[-\frac{x_k^2}{2\sigma_k^2} \right]. \quad (2.27)$$

where the gaussian widths are defined as $\sigma_x = \sqrt{k_B T / m \omega_x^2}$. The raw images of our data are column densities integrated over the imaging axis (either the x -axis or the z -axis). If we integrate equation (2.27) over one spatial dimension, the widths σ_k are unchanged but the peak density becomes $n_0 = N_{th} / 2\pi \sigma_y \sigma_z$. We often also integrate over the y -axis to get a line density for analysis, in which case the peak density of equation (2.27) is $n_0 = N_{th} / \sqrt{2\pi} \sigma_z$.

Free Fall

We almost always image the cloud after releasing it from the trapping potential and allowing it to expand in free fall. The initial RMS velocity of the thermal cloud is

$$v_{RMS} = \sqrt{k_B T / m} = \sigma_k \omega_k \quad (2.28)$$

In the absence of interactions, the velocity distribution remains unchanged in free fall. If the trap is turned off suddenly at time $t = 0$ it will expand isotropically and the initial gaussian widths will be scaled according to

$$\sigma_k^2(t) = \sigma_k^2 + (k_B T / m) t^2 \quad (2.29)$$

which can be used to extract the temperature from the gradient of $\sigma_k^2(t)$ versus t^2 . In practice, since we cannot be sure that the trap turns off suddenly at $t = t_0$, we always calculate the temperature of a cloud using the second expression. This can then be used to calibrated since shot measurements of clouds at different temperatures released under the same experimental conditions. Note that the scaling of the σ_k will also scale the peak density n_0 given according to the expressions in the text above.

2.3.2 Interacting Gas at Zero Temperature

The Hamiltonian for interacting bosons in an external potential is

$$\hat{H} = \int d\mathbf{r} \hat{\Psi}^\dagger(\mathbf{r}) \left(-\frac{\hbar^2}{2m} \nabla^2 + V_{ext}(\mathbf{r}) \right) \hat{\Psi}(\mathbf{r}) + \frac{1}{2} \int d\mathbf{r} d\mathbf{r}' \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}^\dagger(\mathbf{r}') V(\mathbf{r} - \mathbf{r}') \hat{\Psi}(\mathbf{r}') \hat{\Psi}(\mathbf{r}) \quad (2.30)$$

where $V(\mathbf{r} - \mathbf{r}')$ is the interaction potential and $\hat{\Psi}^\dagger(\mathbf{r})$ and $\hat{\Psi}(\mathbf{r})$ are the usual field creation and annihilation operators respectively. The field operator $\hat{\Psi}(\mathbf{r})$ can be written in terms

of the single particle wave functions as

$$\hat{\Psi}(\mathbf{r}) = \varphi_0(\mathbf{r})\hat{a}_0 + \sum_{i \neq 0} \varphi_i \hat{a}_i \quad (2.31)$$

where the $\hat{a}_i$ are the single particle operators and the ground state term $i = 0$ representing the condensed fraction $N_0 = \langle \hat{a}_0^\dagger \hat{a}_0 \rangle$ has been separated from the remaining terms. Under the approximation that the eigenvalue N_0 is large, we can make the Bogoliubov approximation and treat $\varphi_0(\mathbf{r})\hat{a}_0$ as a classical field, rewriting equation (2.31) as

$$\hat{\Psi}(\mathbf{r}) = \Phi_0(\mathbf{r}) + \delta\hat{\Psi}(\mathbf{r}). \quad (2.32)$$

where we have defined the classical field $\Phi_0(\mathbf{r}) \equiv \sqrt{N_0}\varphi_0$ and $\delta\hat{\Psi}(\mathbf{r}) \equiv \sum_{i \neq 0} \varphi_i \hat{a}_i$. In dilute Bose gases at low temperature the condensed fraction can be well over 90% and we can typically ignore the non-condensed component $\delta\hat{\Psi}(\mathbf{r})$. In this case the field operator is equivalent to the classical field $\Phi_0(\mathbf{r})$ and the system behaves classically.

We can identify the wave function of the condensate $\Phi_0(\mathbf{r})$ as the order parameter of the system described in equation (6.2). The time evolution of the wave function is given by

$$\Phi_0(\mathbf{r}, t) = \sqrt{N_0}\varphi_0(\mathbf{r})e^{i\mu t/\hbar} \quad (2.33)$$

where $\mu = \partial E / \partial N$ is the chemical potential of the condensate.

2.3.3 Gross-Pitaevskii Equation

In the Heisenberg representation the equation for the field operator $\hat{\Psi}(\mathbf{r}, t)$ is

$$i\hbar \frac{\partial}{\partial t} \hat{\Psi}(\mathbf{r}, t) = [\hat{\Psi}(\mathbf{r}, t), \hat{H}] = \left[-\frac{\hbar^2 \nabla^2}{2m} + V_{ext}(\mathbf{r}, t) + \int d\mathbf{r}' \hat{\Psi}^\dagger(\mathbf{r}', t) V(\mathbf{r}' - \mathbf{r}) \hat{\Psi}(\mathbf{r}', t) \right] \hat{\Psi}(\mathbf{r}, t). \quad (2.34)$$

where $\hat{H}$ is defined in equation (2.30). To derive the Gross-Pitaevskii equation (GPE) we identify $\hat{\Psi}(\mathbf{r}) = \Phi_0(\mathbf{r})$ and replace the interatomic potential $V(\mathbf{r}' - \mathbf{r})$ with an effective delta-potential $g\delta(\mathbf{r} - \mathbf{r}')$ where g is a coupling constant determined by the s-wave scattering length a

$$g = \frac{4\pi\hbar^2 a}{m}. \quad (2.35)$$

Replacing the interatomic potential with an effective delta-function potential is equivalent to assuming that the condensate wave function $\Phi_0(\mathbf{r}, t)$ varies slowly over distances of the

order of the interatomic force. For ultra-cold atoms the s-wave scattering length a is typically of the order of a few nm. Under these conditions equation (2.34) reduces to the GPE

$$i\hbar \frac{\partial}{\partial t} \Phi_0(\mathbf{r}, t) = \left[-\frac{\hbar^2 \nabla^2}{2m} + V_{ext}(\mathbf{r}, t) + g|\Phi_0(\mathbf{r}, t)|^2 \right] \Phi_0(\mathbf{r}, t). \quad (2.36)$$

The wave function is normalised to the total number of atoms $\int |\Phi_0|^2 d\mathbf{r} = N_0$ and the density distribution is $n(\mathbf{r}) = |\Phi_0|^2$. The time evolution of the condensate wave function is given by equation (2.33), where $\phi_0(\mathbf{r})$ are stationary solutions of the GPE. If we further assume that the external potential V_{ext} is time-independent then we can derive the time-independent GPE

$$\left[-\frac{\hbar^2 \nabla^2}{2m} + V_{ext}(\mathbf{r}) + g|\phi_0(\mathbf{r})|^2 \right] \phi_0(\mathbf{r}) = \mu \phi_0(\mathbf{r}) \quad (2.37)$$

where μ is fixed by the normalisation condition $\int \phi_0^2(\mathbf{r}) d\mathbf{r} = 1$.

We can write an equivalent and useful description of the condensate wavefunction in terms of the density and phase of the system: $\Psi(\mathbf{r}, t) = \sqrt{n(\mathbf{r}, t)} e^{(i/\hbar)S(\mathbf{r}, t)}$. The time evolution of the wave function is determined by the coupled equations

$$\begin{aligned} \frac{\partial}{\partial t} n + \nabla \cdot (\mathbf{v}n) &= 0 \\ \text{and} \quad m \frac{\partial}{\partial t} \mathbf{v} + \nabla \left(\frac{m\mathbf{v}^2}{2} + V_{ext} + gn - \frac{\hbar^2}{2m\sqrt{n}} \nabla^2 \sqrt{n} \right) &= 0 \end{aligned} \quad (2.38)$$

where in general $\mathbf{v} = \mathbf{v}(\mathbf{r}, t)$ and $n = n(\mathbf{r}, t)$.

Thomas Fermi Approximation

Equation (2.37) is the starting point for calculating the ground state of the interacting condensate in an external trapping potential. In the following we assume $\phi_0(\mathbf{r})$ is real so that we can identify $\sqrt{N_0} \phi_0(\mathbf{r}) = \sqrt{n(\mathbf{r})}$, and that the external potential is harmonic, given by equation (2.5). In general equation (2.37) must be solved numerically, but analytical solutions are possible in certain limits.

In the limit that the interaction term $g|\phi_0(\mathbf{r})|^2$ in equation (2.37) is much larger than the kinetic energy term $-\hbar^2 \nabla^2 / 2m$, we have

$$(V_{ext}(\mathbf{r}) - \mu + g|\phi_0(\mathbf{r})|^2) \phi_0(\mathbf{r}) = 0. \quad (2.39)$$

This is called the Thomas-Fermi approximation. Equation (2.39) has the solution

$$n(\mathbf{r}) = \begin{cases} \frac{\mu_{TF} - V(\mathbf{r})}{g} & \text{if } V(\mathbf{r}) \leq \mu, \\ 0 & \text{if } V(\mathbf{r}) > \mu. \end{cases} \quad (2.40)$$

In a harmonic trap equation (2.40) is an inverted parabola

$$n(\mathbf{r}) = \frac{15}{8\pi} \frac{N_0}{\prod_k R_k} \text{Max} \left[\left(1 - \sum_{k=1}^3 \left(\frac{x_k}{R_k} \right)^2 \right), 0 \right] \quad (2.41)$$

where the wave function extends out to a radius $R_k = \sqrt{2\mu/m\omega_k^2}$ so that

$$\mu_{TF} = \frac{\hbar\omega}{2} \left(\frac{15N_0 a}{a_{ho}} \right)^{2/5}. \quad (2.42)$$

In a highly anisotropic cigar-shaped trap the Thomas-Fermi approximation may then be valid in the longitudinal direction but fail in the transverse direction.

The images that we extract of our condensates are column densities integrated over the imaging axis. For analysis we typically also integrate over the y -axis to return a line density $n(z)$. The shape of the inverted parabola changes for the corresponding column and line number densities to

$$n(y, z) = \begin{cases} \frac{5}{2\pi} \frac{N_0}{R_x R_y} \left(1 - \left(\frac{y^2}{R_y^2} + \frac{z^2}{R_z^2} \right) \right)^{3/2} & \text{if } y \leq R_y \text{ and } z \leq R_z \\ 0 & \text{otherwise} \end{cases} \quad (2.43)$$

$$\text{and } n(z) = \begin{cases} \frac{15}{16} \frac{N_0}{R_z} \left(1 - \left(\frac{z}{R_z} \right)^2 \right)^2 & \text{if } z \leq R_z \\ 0 & \text{otherwise} \end{cases} \quad (2.44)$$

Energy

The ground state solution of equation (2.36) minimises the energy functional of the system for a fixed number of particles

$$E = \int d\mathbf{r} \left(\frac{\hbar^2}{2m} |\nabla \Psi_0|^2 + V_{ext}(\mathbf{r}) |\Psi_0|^2 + \frac{g}{2} |\Psi_0|^4 \right) \quad (2.45)$$

where E is the total energy of the system. The energy is then a function of density only

$$E(n) = \int d\mathbf{r} \left(\frac{\hbar^2}{2m} |\nabla \sqrt{n}|^2 + n V_{ext}(\mathbf{r}) + \frac{gn^2}{2} \right). \quad (2.46)$$

The three terms in this equation correspond to the kinetic, harmonic oscillator, and inter-

action energies respectively. When analytic solutions do not exist, the ground state of the system can be found by numerically minimising equation (2.45). Direct integration of the GPE yields the expression

$$\mu = \frac{1}{N} (E_{kin} + E_{ho} + 2E_{int}) \quad (2.47)$$

which is exact within Gross-Pitaevskii theory and allows us to calculate the chemical potential numerically by integrating over solutions to the GPE.

Expansion in Free Fall

In the Thomas-Fermi limit the expansion of the condensate in free fall may be modeled analytically using a scaling approach introduced by Castin and Dum [37]. The time-evolution of the gas is governed by scaling factors in each dimension λ_k obtained by calculating the total force on a particle in the gas and using this in the Newtonian equation of motion. The solution obtained by this method is also a solution of the GPE. This allows us to calculate the evolution of a condensate with initial size $R_k(0)$ via coupled differential equations for the scale factors λ_k rather than by full numerical solutions of equation (2.36).

We take the time-dependent potential

$$V(\mathbf{r}, t) = (1/2) \sum_{k=x,y,z} m \omega_k^2(t) r_k^2 \quad (2.48)$$

with $\omega_{\perp}(t) \gg \omega_z(t)$. The scaling parameters for the condensate radii are solutions to the equations

$$\ddot{\lambda}_k = \frac{\omega_k^2(0)}{\lambda_k(t) \lambda_x(t) \lambda_y(t) \lambda_z(t)} - \lambda_k^2(t) \omega_k^2(t) \quad (2.49)$$

with the initial conditions $\lambda_k(0) = 1$. For our cylindrically symmetrical trap we have $\omega_x(0) = \omega_y(0) = \omega_{\perp}$ and $\omega_z(0) = \omega_z$. In the case of free expansion all the trap frequencies $\omega_k(t)$ go to zero at $t = 0$. We work with a dimensionless time scaled by the radial trap frequency $\tau = \omega_{\perp} t$. Under these conditions equation (2.48) reduces to

$$\frac{d^2}{d\tau^2} \lambda_{\perp}(\tau) = \frac{1}{\lambda_{\perp}^3(\tau) \lambda_z(\tau)} \quad (2.50)$$

$$\frac{d^2}{d\tau^2} \lambda_z(\tau) = \frac{\varepsilon^2}{\lambda_{\perp}^2(\tau) \lambda_z^2(\tau)} \quad (2.51)$$

where $\varepsilon = \omega_z/\omega_{\perp}$ is the aspect ratio of the initial trap. For our traps $\varepsilon \ll 1$. In the limit

$\omega_z \ll \omega_\perp$ these equations have the analytic solutions

$$\lambda_\perp(\tau) = \sqrt{1 + \tau^2} \quad (2.52)$$

$$\lambda_z(\tau) = 1 + \varepsilon^2(\tau \arctan \tau - \ln \sqrt{1 + \tau^2}). \quad (2.53)$$

In this case the condensate expands rapidly in the radial direction because of interactions, but the expansion in the axial direction is suppressed by a factor ε^2 so that for typical drop times of the order of 15-20 ms in our experiment there is very little axial expansion of the condensate, $\lambda_z \simeq 1$.

2.3.4 Elongated Condensates

The condensates we make in our experiment are highly elongated and can begin to cross over from the fully three-dimensional regime we have been discussing so far into a one-dimensional regime in which the radial modes of the system are kinematically frozen out. The condition for this to occur in a condensate is that $\mu < \hbar\omega_\perp$ and $k_B T < \hbar\omega_\perp$. Our condensates typically have $\mu \simeq 2 - 3 \times \hbar\omega_\perp$ in the static trapping potential. However in the RF adiabatic potential each mode is more elongated than in the static potential, the axial trapping frequency is relaxed and the number of atoms is halved. These condensate modes do enter into the cross over regime between 3D and 1D behaviour.

The following discussion is taken largely from [38], but see also [39] [40] [41] [42] and [43]. If we ignore the axial trapping potential, we can write the linear density of a uniform system as $n_1 = n\pi a_\perp^2$ where $a_\perp = \sqrt{\hbar/m\omega_\perp}$ is the radial harmonic oscillator length. When $an_1 \gg 1$ the system is in the 3D Thomas-Fermi regime in which many mode of the radial harmonic oscillator are excited. The radial density distribution will be an inverted parabola with turning points $R_\perp = \sqrt{2\mu/m\omega_\perp^2} = 2a_\perp(an_1)^{1/4}$. This system is properly 3D even though it is highly elongated. It is commonly referred to as a 3D cigar. In the opposite limit $an_1 \ll 1$ there is a perturbative regime in which mean-field theory is still applicable, but the radial wavefunction approaches the gaussian single particle ground state.

In the presence of an axial trapping potential, the equilibrium density distribution along z can be calculated in the local density approximation. In this approximation the condensate is assumed to be locally uniform at each point along the z -axis. We can then fix the energy of the condensate via the equation of state

$$\mu_z(n_1(z)) + V(z) = \mu \quad (2.54)$$

where μ is the chemical potential of the BEC, determined by the normalisation condition $\int n_1(z)dz = N_0$.

When $N_0aa_\perp/a_z^2 \gg 1$ we are in the radial Thomas-Fermi regime, and the axial density distribution is

$$n_1(z) = \frac{1}{16a} \left(\frac{15N_0aa_\perp}{a_z^2} \right)^{4/5} \left(1 - \frac{z^2}{Z^2} \right)^2 \quad (2.55)$$

where the axial Thomas-Fermi radius is $Z = (a_z^2/a_\perp)(15N_0aa_\perp/a_z^2)^{1/5}$. The local equilibrium chemical potential is $\mu_z(n_1(z)) = 2\sqrt{an_1}$. In the opposite limit $N_0aa_\perp/a_z^2 \ll 1$ and the radial motion is frozen out of the system. The axial density distribution is

$$n_1(z) = \frac{1}{4a} \left(\frac{3N_0aa_\perp}{a_z^2} \right)^{2/3} \left(1 - \frac{z^2}{Z^2} \right) \quad (2.56)$$

with the axial Thomas-Fermi radius $Z = (a_z^2/a_\perp)(3N_0aa_\perp/a_z^2)^{1/3}$. The local equilibrium chemical potential is $\mu_z(n_1(z)) = 1 + 2an_1$. The local density approximation can be applied when the axial Thomas-Fermi radius $Z \gg a_z$.

Gerbier calculated an analytical approximation that interpolates between the results of equation (2.55) and equation (2.56) [44].¹ His approach is based on minimising the chemical potential in equation (2.54) using a trial gaussian wave function for $f_\perp[n_1]$ where the variational parameter is the width of the gaussian. The optimised width is $\sigma_\perp = a_\perp(1 + 4an_1)^{1/4}$, which is slightly larger than the ground state harmonic oscillator width $a_\perp$. The local chemical potential is then

$$\mu_z(n_1(z)) = \hbar\omega_\perp(\sqrt{1 + 4an_1} - 1). \quad (2.57)$$

We can combine equations 2.54 and 2.57 along with the axial trapping potential to obtain the axial density distribution

$$n_1(z) = \frac{\alpha}{16a} \left(1 - \left(\frac{z}{Z} \right)^2 \right) \left[\alpha \left(1 - \left(\frac{z}{Z} \right)^2 \right) + 4 \right] \quad (2.58)$$

where $\alpha = 2(\mu/\hbar\omega_\perp - 1)$. The condensate length L is evaluated using equation (2.54) by setting $n_1(L) = 0$ so that $\mu_z(n_1(z)) = \hbar\omega_\perp$, giving $L = (a_z^2/a_\perp)\sqrt{\alpha}$. The axial density is normalised to the number of atoms in the condensate. Using equation (2.58), the normalisation condition gives an equation for α

¹Note: van Amerongen [45] pointed out that there are typographical errors in [44], which I have corrected in the text. Note also that equation (2.57) is incorrectly written in [45] and [43] but used correctly in [46] [47]. Numerical data from equation 5 in [38] can be found online at http://bec.science.unitn.it/data/data_3D_1D.txt and agrees with equation (2.57) as written.

$$\alpha^3 (\alpha + 5)^2 = (15\chi)^2 \quad (2.59)$$

where the parameter $\chi = N_0 a a_\perp / a_z^2$ is roughly the ratio of the interaction energy to the radial zero-point energy. This parameter governs how far into the 1D regime the BEC is. In the limit $\alpha \gg 5$ we are in the 3D Thomas Fermi regime and $\alpha = \alpha_{3D} \simeq (15\chi)^{2/5}$. In the limit $\alpha \ll 5$ we are in the 1D mean field regime and $\alpha = \alpha_{1D} \simeq (3\chi)^{2/3}$. The crossover occurs at $\alpha_{3D} = \alpha_{1D}$ when $\chi = \frac{1}{3}5^{2/3} \simeq 3.73$

2.3.5 Phase Fluctuations

In an elongated condensate such as we use in our experiments, density fluctuations are suppressed the BEC transition temperature, but phase fluctuations driven by thermal excitation can remain significant until the temperature drop much lower. These phase fluctuations have a practical consequence for our experiment since we are interested in measuring the relative phase between the two modes of the condensate.

The following analysis is outline in Petrov et al. [40] [48]. The field operator for the atoms in this regime can be written

$$\hat{\psi}(\mathbf{r}) = \sqrt{n_0(\mathbf{r})} \exp(i\hat{\phi}(\mathbf{r})) \quad (2.60)$$

where $\hat{\phi}(\mathbf{r})$ is the phase noise operator, defined in terms of Bogolyubov excitations of the BEC,¹ and the global phase $S(\mathbf{r})$ has been ignored.

The single particle correlation function is given by [48]

$$\langle \hat{\psi}^\dagger(\mathbf{r}') \hat{\psi}(\mathbf{r}) \rangle = \sqrt{n_0(\mathbf{r}) n_0(\mathbf{r}')} \exp\left(\frac{-i \langle \delta \hat{\phi}(\mathbf{r}, \mathbf{r}') \rangle}{2}\right) \quad (2.61)$$

with $\delta \hat{\phi}(\mathbf{r}, \mathbf{r}') = \hat{\phi}(\mathbf{r}) - \hat{\phi}(\mathbf{r}')$.

Phase fluctuations are dominated by thermal excitations. The dominant contribution is from low energy axial excitations (with energies $\varepsilon_v < \hbar\omega_\perp$). These typically have wavelengths longer than the radius R but shorter than the length L for a highly elongated BEC. The relevant parameter that can be related to experimental measurements is the mean square amplitude of the phase fluctuations. Near the centre of the BEC ($|z|, |z'| \ll L$)

¹Explicitly

$$\hat{\phi}_j(\mathbf{r}) = \frac{1}{\sqrt{4n_0(\mathbf{r})}} \sum_j f_j^+(\mathbf{r}) \hat{a}_j + h.c.$$

where $\hat{a}_j$ is the annihilation operator of an excitation with quantum number j and energy ε_j . The modes $f_j^\pm = u_j \pm v_j$ are determined by the Bogoliubov-de Gennes equations for the excitations of the BEC. See Pitaevski and Stringari, *Bose-Einstein Condensation*, Clarendon Press, 2003.

the mean square amplitude is given by [48]

$$\langle [\delta\hat{\phi}(z, z')]^2 \rangle = \delta_L^2 \frac{|z - z'|}{L}. \quad (2.62)$$

where δ_L is a measure of the phase fluctuations on a length scale $|z - z'| \sim L$. It is given by [48]

$$\delta_L(T) = \frac{32\mu k_B T}{15N_c (\hbar\omega_z)^2}. \quad (2.63)$$

We can rewrite this expression as

$$\delta_L(T) = \left(\frac{T}{T_c} \right) \left(\frac{N}{N_c} \right)^{3/5} \delta_c^2. \quad (2.64)$$

where

$$\delta_c^2 = \frac{32\mu}{15N_c^{2/3} \hbar\bar{\omega}} \left(\frac{\omega_\perp}{\omega_z} \right)^{4/3}. \quad (2.65)$$

For $T \ll T_c$ we have $N_c \simeq N$ and $\delta_L(T) = (T/T_c)\delta_c^2$. Measuring the mean square amplitude of the phase fluctuations can thus be a form of thermometry for a BEC well below the transition temperature where no discernable thermal component remains. Phase fluctuations will be significant for $\delta_c^2 \gg 1$. δ_c^2 strongly depends on the aspect ratio of the trap $\omega_\perp/\omega_z$ and weakly on the number of atoms in the BEC.

The characteristic temperature below which phase fluctuations disappear ($\delta_L^2 \simeq 1$) is defined as

$$k_B T_\phi = \frac{15(\hbar\omega_z)^2 N_c}{32\mu}. \quad (2.66)$$

Above this temperature the phase fluctuates on a length scale smaller than L . In this regime we have a quasi-condensate, with suppressed density fluctuations but a fluctuating phase along the length of the cloud. The phase coherence length is

$$l_\phi = L \left(\frac{T_\phi}{T} \right) \quad (2.67)$$

which is typically much longer than the correlation length (or healing length) of the BEC

$$\xi = \sqrt{\frac{\hbar^2}{2m\mu}}. \quad (2.68)$$

Below T_ϕ phase fluctuations are suppressed and we recover a true BEC.

Density Fluctuations in Free Fall

Phase fluctuations along the length of the BEC show up as density fluctuations in free fall because the velocity field of a BEC is proportional to the gradient of the phase. The noise in the velocity field is then

$$\mathbf{v}_s(\mathbf{r}, t) = \frac{\hbar}{m} \nabla \hat{\phi}(\mathbf{r}, t). \quad (2.69)$$

To first approximation, phase fluctuations can be mapped on to an initial velocity distribution along the length of the cloud. Density fluctuations at long drop times are the result of the interference of components of the BEC with different initial velocities and positions. Since the initial phase distribution is random, the density fluctuations will vary randomly from shot to shot. However average quantities, such as the mean square amplitude of the fluctuations, should remain the same.

An analytical effective operator for the relative density fluctuations valid for drop times $\mu/\hbar\omega_z^2 \gg t \gg \mu/\hbar\omega_\perp^2$ is given in [49] [50]:

$$\frac{\delta n}{n_0} = 2 \sum_j \sin \left(\frac{\varepsilon_j^2 \tau}{\hbar\omega_\perp \mu \left(1 - \left(\frac{z}{L}\right)^2\right)} \right) \tau^{-\left(\frac{\varepsilon_j}{\hbar\omega_\perp}\right)^2} \hat{\phi}_j(z) \quad (2.70)$$

where $\tau = \omega_\perp t$, $\varepsilon_j = \hbar\omega_z \sqrt{j(j+3)/4}$ and $\hat{\phi}_j$ is the phase operator for the j th mode.

The mean square amplitude of the fluctuations follow by averaging over different initial phases. Near the centre of the BEC the average is

$$\begin{aligned} \left\langle \left[\frac{\delta n(t)}{n_0(t)} \right]^2 \right\rangle &= \left(\frac{\sigma_{BEC}}{n_0} \right)^2 \\ &= \frac{T}{\lambda T_\phi} \sqrt{\frac{\ln \tau}{\pi}} \left(\sqrt{1 + \sqrt{1 + \left(\frac{\hbar\omega_\perp \tau}{\mu \ln \tau} \right)^2}} - \sqrt{2} \right). \end{aligned} \quad (2.71)$$

In the lab we measure $(\sigma_{BEC}/n_0)^2$. We can then use this to estimate the temperature of the BEC, given that we know the trap frequencies, chemical potential and drop time. We can also take into account the finite resolution of our imaging system.

2.4 Absorption Imaging

All of our quantitative data comes from absorption images of atoms that have been released from the trap and allowed to expand sufficiently to be at low optical density for the

probe beam. The probe beam intensity is well below the saturation intensity for the D2 transition in ^{87}Rb . Under these conditions the analysis of an absorption image is straightforward. The absorption of the probe beam in this regime is an exponential function of the optical thickness of the ensemble. For a collimated beam propagating along the x-axis the transmitted intensity profile is:

$$I'(y, z) = I_0(y, z)e^{-\sigma_L n_{2D}(y, z)} \quad (2.72)$$

where $I_0(y, z)$ is the intensity profile of the probe beam, σ_L the absorption cross-section and $n_{2D}(y, z) = \int n(x, y, z) dx$ the column density of the ensemble. The absorption cross section is a measure of the probability that a photon is absorbed by an atom in the ensemble. For a two level atom the it is

$$\sigma_L = \hbar\omega \frac{\Gamma}{2} \frac{1/I_{sat}}{1 + I/I_{sat} + 4(\frac{\delta}{\Gamma})^2} \quad (2.73)$$

where I_{sat} is the saturation intensity, δ is the detuning from resonance and Γ is the natural linewidth of the imaging transition. We image on the f_{23} transition using σ^+ polarised light. With this polarisation the transition is closed, with a saturation intensity of 1.67 mW cm^{-2} . The quantisation axis is provided by a uniform magnetic field generated by a set of external coils. The atoms are Zeeman shifted into resonance with the imaging beam by the imaging field so that $\delta = 0$.

We can recover the column density of the atomic ensemble by dividing $I'(y, z)$ by $I_0(y, z)$, where the imaging beam intensity profile $I_0(y, z)$ is measured under identical conditions during the same experimental run, and taking the natural log

$$n_{2D}(y, z) = -\frac{1}{\sigma_L} \ln(T(y, z)). \quad (2.74)$$

where $T(y, z) = I'(y, z)/I_0(y, z)$ is the transmission. The total number of atoms is then obtained by summing over the pixels of the CCD array

$$N = -\frac{\Delta y \Delta z}{\sigma_L} \sum \ln(T(y, z)) \quad (2.75)$$

where Δy and Δz are the dimensions of a pixel in the object plane.

Experimental Apparatus

In this chapter I describe the experimental set-up we use to produce a BEC in an atom chip microtrap. Making a BEC is today a routine procedure that can be achieved using standard textbook methods. We use an approach combining a double MOT system with techniques developed to efficiently load the atom chip microtrap and evaporate to BEC. With our set-up we can routinely produce nearly pure BECs with a few tens of thousands of atoms in a trap conveniently located $\sim 150\,\mu\text{m}$ from the surface of the atom chip.

There are many good introductions to the principles of laser cooling, magnetic trapping, and evaporative cooling of atomic gases. I rely principally on Metcalf and Van der Straten [51]. Rubidium data is taken from [15]. Many aspects of the apparatus have been described in previous theses from the group [52] [53] [54]. Where these have changed little I rely heavily on the descriptions given in these theses.

3.1 Laser System

We require lasers operating at several different frequencies to run the experiment. The lasers are locked to spectral lines in the D2 transition in ^{87}Rb , as illustrated in figure 3.1. The D2 transition connects the $5^2S_{1/2}$ electronic ground state of the atom to the $5^2P_{3/2}$ excited state (one half of a doublet). The ground and excited states have hyperfine structure, and I label transitions from $F \rightarrow F'$ as $f_{FF'}$. With appropriately polarised (σ^+) light the f_{23} transition is a cycling transition, required for efficient laser cooling. The D2 transition has a natural linewidth of $2\pi \times 6\text{MHz}$, giving a Doppler temperature of $147\,\mu\text{K}$.

During the initial magneto-optical trapping (MOT) stage we need a trapping laser red-detuned by a few linewidths from f_{23} , and a re-pump laser tuned to f_{12} . Before turning on the magnetic trap we optically pump the atoms into the $F = 2, m_F = 2$ state using σ^+ light tuned to f_{22} . Finally, to image the atomic gas we need a short pulse of σ^+ light




Figure 3.1: Schematic diagram of the ^{87}Rb D2 line showing the hyperfine splittings in the ground and excited states and the transitions used for laser cooling and absorption imaging. The F-number of the hyperfine levels is indicated. Transitions from the ground state hyperfine levels F to excited state hyperfine levels F' used in the experiment are labelled, specifically the principle transition f_{23} used for cooling and imaging, the repump transition f_{12} and the optical pumping transition f_{22} .

tuned to resonance with the f_{23} transition.

We use a Ti:Sapphire laser to provide trapping light, and two diode lasers as a reference and repump laser. The lasers are locked to atomic spectral lines using a combination of polarisation spectroscopy and a frequency offset lock. The frequency offset lock is also used to rapidly switch the frequency of the trapping/imaging laser.

Laser Schematic

A schematic diagram of the laser system is given in figure 3.2. The components are described below. The entire system is enclosed in a black box with holes for the trapping, optical pumping and imaging beams, to minimise the effect of stray resonant light reaching the experiment when the atoms are magnetically trapped. Blackout curtains around the optical table serve a similar purpose.

Ti:Sapphire Laser

The trapping and imaging laser beams are produced by a Coherent MBR-110 Ti:Sapphire laser pumped by an 8W Coherent Verdi V-8 diode-pumped solid state laser operating at 532nm. The output is linearly polarised, with a TEM_{00} spatial mode and a linewidth of




Figure 3.2: Schematic diagram of the laser system used in the experiment showing how light from the Ti:Sapphire and diode lasers is picked off for spectroscopy and the frequency offset lock, and where the imaging light is picked off from the trapping laser, and the optical pumping light from the reference laser. I also show where the repump light is combined with the trapping light and distributed to the LVIS and main chamber MOT beams, and the telescopes used to shape the imaging beam and MOT beams.

100kHz [52].

Once locked to the reference cavity, the MBR-110 laser can be scanned up to 40GHz by changing the cavity length, which can be externally controlled. Coarse wavelength selection is made by observing the fluorescence of the D2 line from a pick-off beam passing through a rubidium vapour cell (see figure 3.2). We lock the laser to the required frequency by controlling its reference cavity. The error signal for the lock is derived from the beat frequency of the MBR-110 mixed with a reference laser locked to f_{23} using a frequency offset lock based on a design published in Schünemann et al. [55]. Details are again in Matt Jones’s thesis [52]. The MBR-110 provides the trapping beams for in the LVIS and the main mirror-MOT. For laser cooling we red-detune these lasers 18MHz from f_{23} (three times the natural linewidth of the transition). We use a shutter to ensure that the beam is completely blocked when switched off.

Diode Lasers

The reference and repump lasers are home built external cavity diode systems. The external cavity is provided by a grating mounted in the Littrow configuration. The design is a slightly modified version of that published in Arnold et al. [56] and full details are given

in Matt Jones's thesis [52]. The frequency can be scanned through 3 GHz and the output power at 780 nm is around 30 mW.

The reference laser provides the reference frequency for locking the MBR-110 and locked to resonance with f_{23} . The repump laser is locked to f_{12} . We use a polarisation spectrometer based on a design by Wieman and Hansch [57] to lock each of these lasers. A detailed description of our set-up can be found in Matt Jones' thesis [52].

Switching and Shuttering

We switch the trap laser, imaging laser and optical pumping lasers using acousto-optical modulators (AOMs). These provide fast switching control for precise timing. In addition, to ensure that all resonant light is blocked from reaching the main chamber during the sensitive stages of the experiment, we block all beams with mechanical shutters.

MOT Beams

Light from the repump laser is combined with the trapping light using a polarising beam cube. This light is then sent to two MOT systems: a mirror-MOT in the main chamber, and a low-velocity intense source (LVIS) that is used to load the mirror-MOT. The light is first split into horizontal and 45° beams for the mirror-MOT, then the LVIS beam is picked off from the horizontal MOT beam. The beams are telescoped immediately after the combining beam cube. This set-up is illustrated in figure 3.2.

Optical Pumping Beam

The optical pumping beam is picked off from the reference laser and passed twice through an AOM operating at 140 MHz. The resulting beam is 13 MHz red-detuned from f_{22} . We use red-detuned light for optical pumping because the ensemble is optically thick for resonant light when the MOT is large. The AOM also controls the switching of the beam, and a shutter is in place to block it when not in use.

Imaging Beam

We pick off light from the MBR-110 for the imaging beam. This passes through an AOM operating at 100 MHz which provides precise timing control of the pulse used for absorption imaging. A shutter is in place to block stray light from the beam when switched off. The frequency of the imaging beam is detuned slightly from resonance with f_{23} . This detuning is compensated by the imaging magnetic field, as discussed in more detail below. The imaging beam is telescoped after passing through the AOM. We switch between the two imaging axes using a flip-mirror adjusted by hand between experimental runs.

3.2 Vacuum System

The vacuum system consists of two chambers separated by a gate valve and a differential pumping aperture. This aperture allows us to maintain ultra-high vacuum (UHV) conditions in the main chamber while dispensing rubidium in the LVIS chamber and provides a path for the LVIS beam to enter the main chamber. The LVIS design is based on the near identical system described in Chris Sinclair's thesis [53].

LVIS Chamber

The LVIS vacuum chamber consists of two $2\frac{3}{4}$ inch conflat six-way crosses connected by a T-piece (see figure 3.3). The LVIS MOT is made in the cross closest to the main chamber, which has four view-ports for the MOT lasers orthogonal to the cold atom beam. A mirror with a small hole drilled through the centre for the LVIS beam to pass through is mounted on a $2\frac{3}{4}$ inch conflat flange attached to the port nearest the main chamber. We refer to this as the extraction mirror. An additional conflat-mini viewport welded to the cross at 45° is used as a view port for observing fluorescence from atoms in the LVIS MOT. Two rubidium dispensers are mounted on an electrical feedthrough fitted to the connecting T-piece. A 20ls^{-1} ion pump (Varian VacIon Plus 20) is connected to the chamber via a short nipple. We shield the experiment from stray magnetic fields from the ion pump by encasing it in a soft iron shell. We have also attached two angle valves, one used to connect a turbo pump for the initial pump down of the vacuum chamber, the other connecting to a by-pass tube which joins the cross to the region of the LVIS chamber behind the extraction mirror. We used this by-pass tube when we set up the LVIS system to ensure that this region of the chamber was properly evacuated during the bake-out of the LVIS prior to opening the connecting valve between the LVIS and the main chamber. The angle valve was then closed leaving the by-pass tube on the UHV side of the vacuum system.

On the far side of the extraction mirror is a nipple with two conflat-mini viewports welded orthogonal to the LVIS beam. These were included in the design to allow for additional diagnostics on the LVIS beam after it passes through the aperture. In practice this has not been necessary to get the LVIS working. This nipple is then connected to a gate valve on the main vacuum chamber.




Figure 3.3: Illustration of the experiment vacuum system as seen from above. The LVIS is at the bottom left, separated from the main chamber by a gate valve and differential pumping aperture. Pumping is provided by two ion-pumps and a non-evaporable getter (NEG). Also shown is the ion-gauge used to monitor the vacuum during pump-down, angle valves used to attach a roughing pump for the bake-out and initial pump-down, and the feedthroughs for the LVIS dispensers. The atom chip in the main chamber is mounted upside down on a single flange which also houses the vacuum feedthroughs for all the intra-vacuum components. The windows used for the MOT light in both chambers, and the imaging and optical pumping beams, are indicated. The LVIS MOT and beam can also be monitored via the additional viewports as shown. The 45° MOT beams in the main chamber enter via a large viewport mounted on the underside of the vacuum chamber.




Figure 3.4: Diagram of the extraction mirror mounted in a $2\frac{3}{4}$ inch conflat flange. The position of the mirror is indicated in figure 3.3. The mirror is made from a $\lambda/4$ waveplate coated in aluminium to provide a reflective surface. A hole 1 mm in diameter in the centre of the mirror provides an aperture for the LVIS beam. The mirror is mounted into a recess machined into the flange and then attached with Epoxy (Bylapox 7285).

Extraction Mirror

The extraction mirror is made from a 6 mm thick $\lambda/4$ waveplate drilled with a 1 mm diameter hole at the centre. The back face of the waveplate is coated in aluminium to provide a reflective surface. The reflectivity is 70% at 780 nm. It is mounted in a modified $2\frac{3}{4}$ inch conflat flange (CF DN40).

Figure 3.4 is a schematic cross-section of the waveplate mounted in the modified flange. A 3 mm deep recess for the optic was machined in the centre of the flange, and a 3 mm hole drilled through the centre. The hole is wider than that in the optic to allow for easy alignment. We placed the waveplate into the recess and then applied epoxy (Bylapox 7285) around the edge to secure it to the flange. Once the epoxy had cured, we attached the modified flange directly to the six-way cross.

The aperture in the extraction mirror is the only connection between the main chamber and the LVIS chamber. The pressure differential that can be maintained between the two chambers is determined by the conductance of the aperture and the effective pumping speed on the UHV side of the aperture due to the main chamber vacuum pumps. Under UHV conditions, the molecular conductance C through the aperture is determined by its diameter d and length l [53]. For an aperture 1 mm in diameter and 6 mm in length, the conductance is approximately 0.02 ls^{-1} . The pressure differential that can be maintained across the aperture is approximately the conductance divided by the effective pumping speed of the main chamber pumps on the UHV side of the aperture. We assume that all atoms or molecules that pass through the aperture enter the main chamber, neglecting the effect of the larger hole in the mounting flange and the nipple and gate valve between the extraction mirror and the main chamber. The main chamber is pumped by a 20 ls^{-1} ion pump (VacIon 20) and a non-evaporable getter (NEG). The effective combined pumping speed gives a pressure ratio of about $\kappa = 5 \times 10^4$. In practice we observed no change in the pressure in the main chamber when dispensing atoms in the LVIS chamber, measured with the ion gauge attached to the main chamber. The pressure in the main chamber was 7×10^{-11} Torr at this point.

Because of the low conductance of the aperture, a by-pass tube connects to the nipple on the UHV side of the aperture to an angle valve near the LVIS ion pump. This by-pass was open during the initial pump down and baking of the LVIS vacuum system to ensure that this part of the vacuum system was properly evacuated. The angle valve was then closed before the gate valve connecting the LVIS to the main chamber was opened.

Dispensers

Two rubidium dispensers are attached to a flange on to the T-piece connecting the two six-way crosses in the LVIS vacuum system. We use one dispenser to release rubidium vapour into the chamber. The other was added as a back up in case of the first failing or becoming exhausted. The dispensers, SAES getters commonly used in cold atom experiments [58], release a rubidium vapour when resistively heated above a threshold temperature. During operation we maintain a current of 2 A through the dispenser to keep it warm, which allows the rubidium vapour to be switched on faster and more reproducibly. In order to make an LVIS MOT we typically run 7.5 A through the dispenser for 15 s.

Main Chamber

The main vacuum chamber is a spherical octagon made from 304 grade stainless steel by Kimball Physics inc. It has eight $2\frac{3}{4}$ inch (DN40 CF) conflat ports around the perimeter, and 8 inch (DN160 CF) conflat ports on the top and bottom. The bottom flange is sealed with a viewport which allows optical access for the two 45° mirror-MOT beams. Four of the perimeter ports are also sealed with viewports to give optical access to the horizontal MOT beams and for optical pumping and imaging the ensemble.

One of the other four ports is connected via the gate valve to the LVIS chamber. The cold atom beam enters the chamber through this port. One port is connected to an angle valve (Lesker VZCR40R), used to attach a turbo pump for the initial pump down of the chamber and the bake out. An ion gauge (Varian UHV 24p) is also connected to this port via a T-piece between the chamber and the angle valve. The remaining two ports are connected to vacuum pumps. One is a 20ls^{-1} ion pump (Varian VacIon Plus 20) mounted on the end of a 150 mm nipple, and surrounded by a soft-iron shell to shield the main chamber from its magnetic field. The other is a non-evaporable getter (NEG) (Saes Sorb-AC GP502F with an St-101 Alloy).




Figure 3.5: Entire vacuum flange assembly including atom chip, chip mount and auxiliary wires and coils used in loading atoms into the chip trap.

3.3 Top Flange Assembly

All of the components of the experiment inside the main vacuum chamber are mounted on a single 8 inch conflat flange (DN160 CF). A photograph of the flange assembly is shown in figure 3.5. The principle components are: the MOT coils, which provide the quadrupole magnetic field needed for the initial mirror-MOT; a SAES getter, used to load the mirror-MOT from a thermal vapour before the LVIS was brought on line; the atom chip mounted on a ceramic base plate and a copper block connected via three legs to the flange; and a copper H structure mounted in the ceramic base plate immediately below the atom chip, designed as an alternative source of the magnetic field required to make a Ioffe-Pritchard (IP) trap. The MOT coils are wound onto a stainless steel coil former with a 12.5mm inner radius and a 45mm centre to centre separation. We typically run 2A through the coils when we make a mirror-MOT, which gives a gradient of 11 Gcm^{-1} along the axis of the coils. The chip mount is described in more detail below.




Figure 3.6: Model of chip mount. The mount is designed to be modular, allowing the chip and substructures to be changed relatively easily. The chip is mounted onto a ceramic base plate made from Shapal-M and held in place with mechanical clamps. A macroscopic H-wire machined from oxygen-free copper is mounted into grooves in the Shapal base plate and sits underneath the centre of the chip. It is used to provide auxiliary magnetic fields for trapping and manipulating atoms. It also acts as a heat sink conducting heat away from the chip wires to the chip mount. The Shapal base plate is mounted onto a copper block that is attached to the vacuum flange via three thick copper legs. Further auxiliary wires are mounted below this copper block (not used in the experiments described in this thesis). Also shown are the clips that connect to the chip wires (described in the main text). The clips connected to the tightly spaced Z-wire flares are held in place with macor clamps. The clips are soldered to wires that connect to the vacuum feedthrough pins and are not shown in this illustration.

Chip Mount

The chip mount is illustrated in figure 3.6. It is designed to provide electrical connection to the chip wires and conduct heat away from the chip wires as efficiently as possible. The chip is mounted directly onto a 5 mm thick base-plate made from Shapal-M, a machinable ceramic made from aluminium-oxide chosen for its good thermal conductivity ($90 \text{ W m}^{-1} \text{ K}^{-1}$) and mechanical strength. It is held in place with copper clamps which screw into holes tapped into the Shapal. Embedded into the Shapal base-plate is an H-wire made from oxygen-free copper, which has a 2 mm square cross section, and a central cross-piece 2 mm long. This copper structure sits directly below the centre of the atom chip, where the majority of the heat generated by the chip wires is dissipated and plays an important role in conducting heat away from the chip when the experiment is running.

The H-wire sits in a groove machined out of the Shapal base-plate and is held in place with UHV compatible epoxy (Bylapox). In order to get a smooth surface for the chip to sit on, we initially made the H-wire 2.2 mm thick so that it extended $200 \mu\text{m}$ above the Shapal surface. Once it had been glued in place and the epoxy had cured, we hand polished the surface until the H-wire was level with the Shapal.

Below the Shapal piece is a 7 mm-thick copper block, which is connected to the top vacuum flange via three copper legs, each with a 5 mm square cross section. The legs are 21.5 mm long. The various pieces of the mount, Shapal base-plate, copper block and legs, are held in place with vacuum relieved stainless steel screws. Together they conduct heat away from the chip to the top flange, which acts as a heat sink for the experiment.

Chip Wire Connections

The chip wires flare out into pads at the corners of the chip to provide a surface that can be used to make electrical connections. These pads are closely spaced around all four corners of the chip. We make the electrical connection to these pads using 1.5 mm wide clips that mechanically clamp onto the edge of the chip - see figure 3.6.

The clips are commercially made (by Batten & Allen) from a copper-tin alloy (CuSn_6). The four clips that connect to the Z-wires at either end of the chip are held in a clamp made from two blocks of Macor screwed against one another. Grooves in one Macor block hold the clips at the correct pitch to connect to the closely spaced Z-wire pads. All the other clips are connected individually to the various chip wire pads. The clips are soldered to macroscopic wires that connect to the vacuum feedthrough pins via beryllium-copper in-line barrel connectors. The solder is UHV compatible fluxless solder from Allectra. We have not observed any adverse effects of the solder on the vacuum that we can achieve in the main chamber.

When we originally mounted the chip the clips made good electrical connection to the chip wire pads, adding no significant resistivity to the chip wires. We tested the response of the clips to several cycles of resistive heating and cooling and observed no change in either the electrical connection or the mechanical strength of the clip. The clips were protected from the conductive doped silicon substrate by the layer of silicon-oxide grown before the gold from the chip wires was evaporated onto the chip. Where the clips were placed very close to one another, or to the copper H-wire running under the centre of the chip, we used Kapton foil as an additional insulating layer. Mounted in this way, there were no observable cross-connections amongst the chip wires or between the wires and the chip mount.

Subsequently we have run into problems with these clips. We believe that they have not maintained their mechanical stability under the stress of repeated cycles of heating and cooling, which may have weakened the grip of the clips on the chip over time. We also believe that the clips may have scraped through the thin silicon-oxide layer that insulates the doped-silicon on the back and edges of the chip, creating cross-connections between the chip wires, and connection between the wires and the chip mount.

At the time of writing, there are cross-connections between the various chip wires with a resistance of a few hundred Ohms, and connections to the experiment common point (ground) via the chip mount of 1-2 k Ω . We believe that these connections are made via the silicon substrate where the silicon-oxide layer has been scraped off by the clips as they move under repeated thermal cycles. We have had to change the way that we control DC current through the chip wires because of these problems. However despite these problems, we are still able to successfully make BECs and observe matter wave interference using this chip, as described in subsequent chapters.

Electrical connection to the outside world is made via three smaller flanges machined directly into the top flange: two $1\frac{3}{4}$ inch conflat and one $2\frac{3}{4}$ inch conflat. We connected vacuum feedthroughs directly to these flanges. The two smaller feedthroughs are made by Ceramaseal. Each has 8 copper feedthrough pins 0.81 mm in diameter. The larger feedthrough is from MDC Vacuum Products. It also has 8 pins, each 2.39 mm in diameter. They can be seen in figure 3.5.

3.4 External Fields

LVIS Coils

The quadrupole field for the LVIS MOT is generated by anti-Helmholtz coils wound directly around the vertical ports of the six-way cross using insulated copper wire. There

are 100 turns in each coil, with an average diameter of about 4.5 cm and a mean separation of about 6.5 cm. With a typical operating current of 2 A, coils generate a field gradient of 15 G cm^{-1} along the axis of the coils.

There are also two sets of shim-coils wound in Helmholtz configuration around the vertical and horizontal ports of the six-way cross. Each pair can generate a field of approximately 1.5 G per Ampere. These fields from the shim-coils are used to align the LVIS with the hole in the extraction mirror. We can also use the shim fields to switch off the LVIS beam while the LVIS MOT is running by shifting the centre of the quadrupole away from the hole.

Main Chamber Coils

The most important fields for running the experiment are the X-bias and Z-bias. The X-bias field cancels the field from the trapping wires along a line to form the magnetic trap, and the Z-bias is aligned along the trapping wires and sets the magnitude of the field at the bottom of the trap, as shown in figure 2.1. The X-bias field is also used for absorption imaging through the diameter of the cloud, as in figure 3.11. These fields are provided by coils in near-Helmholtz configuration mounted around the main chamber. The X-bias coils provide a field of up to 30 G along $\hat{x}$. The Z-bias coils provide a field of a few gauss along $\hat{z}$. Additional shim coils in provide fields of up to a few gauss along $\hat{y}$.

3.5 Atom Chip

Fabrication

One of the basic motivations behind atom chip experiments is the possibility of controlling the motion of the atoms on small length scales and with large trapping frequencies. The minimum length scale over which a static magnetic field can vary significantly is set by the distance to the source of the field, a consequence of Laplace's equation. This suggests working with atoms trapped close to small wires is preferable. Similarly, the achievable field gradient (and thus trapping frequency) increases as the atoms move closer to a wire (saturating at a finite value close to a wire of finite width and height). The maximum field gradient for a wire of finite width is proportional to I/w^2 . The maximum current density achievable is limited by resistive heating of the wire, and is proportional to $\sqrt{hw^2}$ (assuming a rectangular wire of height h and with one surface of width w in contact with a heat reservoir of constant temperature). The gradient thus scales as $B'_{max} \propto \sqrt{h/w^2}$ [59] [23].




Figure 3.7: Figure 3.7(a) is a photograph of our atom chip before it was mounted in the vacuum chamber. The four central wires are not visible at this scale, but can be distinguished at the edge of the chip where they fan out into larger pads for connection to macroscopic wires. The two end wires, and surrounding gold pads can also be seen. Figure 3.7(b) is an optical microscope image of the four parallel Z-wires in the center of the chip. The outer wires are $100\mu\text{m}$ wide and have a centre to centre separation of $300\mu\text{m}$. The smaller inner wires are $50\mu\text{m}$ wide and have a centre to centre separation of $85\mu\text{m}$.

These criteria set the basic parameters of atom chip fabrication. Achieving high current density dictates using a wire with good bulk conductivity and homogeneity, and using a substrate with good thermal conductivity so the wire does not overheat. Gold, silver and copper wires have all been used in atom chip experiments, typically on silicon or sapphire substrates. We use an atom chip with gold wires patterned onto a p-doped (100) silicon wafer [60] with a resistivity of $17\text{--}33\Omega\text{cm}$. Silicon has advantages over other substrates, such as sapphire, which may have slightly better thermal conductivity. It can be prepared with an atomically smooth surface, which is important for minimising deviations in thin wires. It is also the basis of a mature fabrication industry, with many techniques available that are readily adapted to developing atom chip technology. The silicon is wet-oxidised in a furnace to create a 100nm layer of SiO_2 to electrically isolate the gold wires from the substrate. The oxide is kept thin to maintain good thermal conductivity between the wires and the silicon. Setting conservative limits on how much we can heat the chip wires (a 50% increase in the chip wire resistivity, or a temperature rise of about 100°C) we can run current densities of up to $9 \times 10^9 \text{ A m}^{-2}$ for over ten seconds, which is sufficient time to run our experimental cycle.

The limiting factor in trapping atoms close to the surface of a wire is the homogeneity of the wires themselves. Surface roughness, edge roughness and bulk defects lead to

small transverse current components that introduce an anomalous magnetic field along the length of the trap. When the magnitude of this field approaches the temperature of the atomic cloud, the cloud breaks up into lumps. This anomalous potential roughness also increases as the atoms come closer to the trapping wire.

There are three processes standardly used to fabricate an atom chip: wires can be electrochemically deposited into a mould defined by a patterned resist, etched out of a thin film of metal covering the substrate, or evaporated onto the substrate through a patterned mask [27] [59]. Detailed studies have suggested that the third method, also known as lift-off, results in considerably smoother wires than electrochemical deposition [20] [23] [24]. Our atom chip is prepared with a variation of the second technique in which the wires are etched with an ion-beam out of an evaporated gold film. Similar wires have been studied in other groups [25].¹ Other etching techniques studied produced considerably rougher wire edges and/or surfaces [60] [62].

Ion Beam Milling

The wires were patterned in a multi-step process, illustrated in figure 3.8. In the first step, a 3 μm layer of gold was evaporated onto the substrate. Since gold does not readily adhere to the oxide layer covering the silicon wafer, a 50 nm layer of chromium was first deposited onto the surface. The gold was then evaporated in six steps of 600 nm each. The result was a gold film with good bulk conductivity and an optical quality surface. There was no indication of any layering effect due to the multiple step evaporation process, and what cracks and grain can be seen in the sides of the wires cross multiple layers – see figure 3.9(a) and figure 3.9(b).

In the second step the wires were patterned by ion-beam milling. A 2.2 μm thick layer of HPR504 photoresist was spun onto the surface and patterned by ultra-violet (UV) lithography. The resist was then heated to allow it to reflow, which prevents slivers of resist forming during the ion-beam milling, which are then difficult to remove. The surface was then ion-beam milled, removing the gold layer under the gaps in the resist. Because of the reflow, the resulting wires have a trapezoidal cross-section, with the sides angled at 30° to the surface normal, as seen in figure 3.9(b).

The resulting wires have an optically smooth surface suitable for making a mirror-MOT. An unintended consequence of the ion-beam milling was that the silicon substrate between the wires was over etched by approximately 1 μm . Some of the silicon removed during the milling was redeposited on the sides of the wires, forming conductive bridges that crossed the thin insulating oxide layer. This redeposited silicon was removed with

¹However, note that Simon Aigner refers to focussed ion-beam milling in Figure 2.11 of his thesis [61].




Figure 3.8: Schematic diagram illustrating the process of fabricating our atom chip. In the first step, figure 3.8(a), the silicon substrate is wet-oxidised to create a 100 nm layer of Silicon-Dioxide. A 50 nm layer of chromium is then deposited, followed by a 3 μm evaporated gold layer. Finally, figure 3.8(b), the photoresist is spun onto the gold surface. In the second step, figure 3.8(c), the photoresist is patterned by ultra-violet lithography, removing photoresist from where the gaps between the wires will form. The photoresist is heated and allowed to reflow, figure 3.8(d), which prevents hard to remove fragments of resist sticking to the gold following the ion-beam mill. After the reflow, the cross section of the resist above the wires is trapezoidal. In the final step, figure 3.8(e) the entire structure is ion-beam milled. The wires have a trapezoidal cross-section, inherited from the photoresist. The ion-beam over mills the gaps between the wires, removing the chromium, silicon-dioxide, and some of the silicon substrate. Because of this, some silicon can be redeposited across the sides of the wires during the ion-beam milling, forming conductive bridges between the wires. This material is subsequently removed by a potassium hydroxide etch.

potassium hydroxide (KOH), which will etch silicon anisotropically but leaves the gold intact. This etch removes the unwanted bridges from the sides of the wires and a further $2\text{ }\mu\text{m}$ layer of silicon from the wafer, leaving behind a rough surface, which can be seen in the gaps between the wires in figure 3.7(b). The resulting cross-resistances between neighbouring wires were measured at $2\text{--}3\text{ M}\Omega$ prior to mounting the chip.

Wire Imperfections




Figure 3.9: Scanning Electron Micrographs (SEM) of the chip wires after ion-beam milling. Figure 3.9(a) shows the side of the wire. Structures such as grains and cracks in the gold cross the five 600 nm layers in which the gold was evaporated (which are not apparent in the micrograph). Similar structure can be seen in the bulk of the wire in figure 3.9(b), which shows a cut away cross-section of a wire. The 30° angle of the sides of the chip wires to the surface normal is indicated. The image in figure 3.9(c) is an optical microscope image of one of the $50\text{ }\mu\text{m}$ wide smaller wires near the centre of the chip.

Typical fabrication techniques, such as those used for our chip, minimise fluctuations at short (sub-micrometre) correlation length, which are negligible for our geometry. Longer wavelength fluctuations, which are difficult to observe with standard techniques for characterising wires, are less well understood.

At short wavelengths the power spectrum of a wire can be estimated from SEM images such as figure 3.9(a). Wires such as ours display what is referred to as self-affine fractal roughness,¹ described by an empirical correlation function $\langle f(z)f(z+r) \rangle = \sigma^2 \text{Exp}[-(r/\xi)^{2\alpha}]$, where α is referred to as the Hurst exponent, $\sigma^2 = \langle f(z)^2 \rangle$ is the mean square roughness, and ξ is the correlation length. For wires similar to ours, typical parameters are $\sigma = 3 \text{ nm}$, $\xi = 20 \text{ nm}$ and $\alpha = 0.5$ [27]. At long wavelengths it is more difficult to measure these parameters accurately - typically many SEM images must be carefully stitched together [23]. For our experiment parameters, with atoms trapped $150 \mu\text{m}$ above a $50 \mu\text{m}$ wide wire running 2 A the corresponding magnitude of the fluctuating field component $|\delta B_z|$ is $< 1 \text{ mG}$ (or $\sim 1 \text{ kHz}$), significantly smaller than the roughness we measure with our atoms. This suggests that the fragmentation we observe in cold atomic ensembles (see chapter 4) is due to fluctuations at longer wavelengths, which may be due to thermal processes, such as the re-flow of the resist, or acoustic vibrations during the ion beam milling.

3.6 Current Control

To create the magnetic fields required to run the experiment, we need to control the current running in various coils and in the chip wires. The current control circuit is discussed below. The circuits are controlled by analogue input voltages. For simple switching of coils we use 2-way and 4-way analogue level setter circuits that provide a fixed voltage depending on the state of TTL logic inputs (one or two digital input signals). The chip wire currents, X-bias field and Z-bias field are all ramped linearly during the transfer from one experimental stage to another. The control voltage for these ramps is provided by analogue voltages from a National Instruments analogue output board (see below).

Op Amp Circuit

We control DC currents through the chip wires using a simple home built operational amplifier (OpAmp) current source. The design is based on circuits previously used to run the experiment that used a high power field effect transistor to control the current through the load [52] [63]. These circuits connected the load directly to a voltage source, with the switch (a field effect transistor) between the load and ground. Our current chip wires connect to the experiment ground with a resistance of less than $1 \text{ k}\Omega$ (see the discussion in section 3.3 of chapter 3), so with these circuits we were unable to switch off the current in

¹A self-affine fractal is invariant under anisotropic transformations, so the scaling in each direction can be different.




Figure 3.10: Schematic diagram of the full circuit chip wire current control circuit.

the chip wires completely. The field gradient from this latent current not only adversely affected the shot-to-shot stability of the experiment, but made interference experiments impossible.

The control circuit uses a high-powered Operational Amplifier (OPA548 or OPA549 from Burr-Brown) as a current source. The circuit is illustrated in figure 3.10. It is stabilised by an integrator with variable gain. Feedback comes from a 0.5Ω sense resistor which is fed into a differential amplified for comparison with the control voltage. The circuit can be run in uni-polar or bi-polar mode, and we have tested it as an AC current driver running 1 A peak-to-peak current at a few 10s of kHz through the chip wires.

We typically operate the circuits in uni-polar mode. The OPA549 is powered by a single high-power linear laboratory power supply connected to V^+ . The control circuitry (and the V^- pin on the OPA549) is powered by a floating linear power supply (typically $\pm 15\text{ V}$ and 200mA). The entire circuit shares a single common point. We trim the input voltage with a separate circuit in order to accurately zero the current running through the chip wires in the off state. In this configuration the circuit will switch off the current in a resistive load like the chip wires in less than $100\mu\text{s}$. The circuits are built on aluminium heat-sinks to help dissipate the power that is dropped over the OPA549 during operation.

For applications requiring large currents, multiple OpAmps can be connected in a master-slave configuration, with the additional OpAmps connected to the output as voltage followers. An inductive load such as the X-bias coil can be run stably without the additional integral gain circuitry. A capacitor and resistor across the OPA549 provide sufficient integral gain for stable operation and proportional gain for fast switching.

3.7 Radio Frequency Generation

We need to apply radio-frequency fields both to evaporatively cool the atoms and to split the condensate, which we do by applying an intense RF field to dress the atoms. The RF adiabatic potentials that result are described in detail in the following chapters. We use two different types of frequency generator for the two applications. One drives RF-currents through the chip wires, and the other through the copper-H mounted beneath the chip.

Versatile Frequency Generator

The radio-frequency fields for forced evaporative cooling are provided by a Versatile Frequency Generator (VFG150). The VFG150 allows almost arbitrary control of pulses and pulse sequences at frequencies up to 150MHz with an output power range of -69dBm to 0dBm. The frequency, amplitude and phase of the waveform can be specified at each point in the sequence, which can be generated with simple scripting tools. The sequence is sent to the VFG150 via a USB interface, where it is stored on board. We then trigger the sequence at the appropriate moment using an external TTL signal.

DS345

The radio-frequency fields used to split the condensate are generated by two phase locked Stanford Research Systems DS345 Function Generators, which can output up to 24dBm at up to 30MHz. We set the amplitude of each DS345 separately using the rear panel amplitude modulation input, controlled by an analogue output channel (see below). Frequency sweeps are triggered with an external TTL signal.

Switching, Amplification and Connection

We control the switching of the RF fields using ZASWA-2-50DR switches from Mini-Circuits. These have a typical rise/fall time of 5ns and an in-out isolation of 100dB at frequencies of up to 100MHz. This gives us precise timing control over the RF field, and allows us to fully attenuate the field (the amplitude modulation input on the DS345 provides insufficient attenuation). We connect one switch to the output of each of the DS345s. In order to generate the large fields required to make the RF adiabatic potentials, we use one amplifier (Mini-Circuits ZHL-32A) for each function generator, connected to the output of the RF switches. These provide 25dB gain.

We couple the RF for the adiabatic potentials directly onto the chip wires used for DC trapping. We couple the RF to each wire using an RF Transformer (Mini-Circuits T1-

1T). We isolate the transformer from the DC current using a capacitor. This is illustrated schematically in figure 3.10. We connect the VFG output to the copper-H via a ZHL-32A amplifier. It is used to generate the chirped RF field needed for evaporative cooling in the chip wire magnetic trap.

3.8 Experiment Control

In order to run the experiment we need to control the switching of many currents, fields and lasers with sub-millisecond timing. We use hardware from National Instruments to provide this control: one digital pattern generator (National Instruments PCI-6543), and one analogue output board (National Instruments PCI-6713). Custom written software, originally written by Chris Sinclair to run the video-tape BEC experiment in our laboratory, provides the user interface. The software allows us to specify the state of digital and analogue output channels in a linear sequence of events. An entire waveform is generated and then sent to the relevant output boards. The timing of the sequence of events is then controlled by the boards. This allows us to run the experiment on a Windows computer system while still maintaining precise timing control. Details are given in Chris's thesis [53].

3.9 Imaging System

Camera

All of our experimental data comes from images acquired using a CCD camera. We have used two cameras over the course of the experiments described in this thesis, both Penta-MAX cameras from Princeton Instruments made with CCD arrays from Kodak (Kodak-768 and Kodak-1035). One camera has 768×512 pixels each with an area of $9 \times 9 \mu\text{m}$. The other has 1317×1035 pixels of $6.8 \times 6.8 \mu\text{m}$. The cameras are controlled by image capture software from the manufacturer (WinView). During an experimental run, the camera is initialised by a programme executed from within our control software, then waits for a TTL hardware trigger from the digital pattern generator. When the trigger is received, the camera shutter opens to expose the CCD array. We typically expose and read out the entire array in a single shot, which takes about 330ms for the larger array. The exposure time is set by adjusting a parameter within WinView prior to executing an experiment.

The cameras are electronically cooled to -30°C . The sensitivity of each camera at 780nm has been accurately calibrated against a power meter [64] [52]. The Kodak-768

chip has 0.019 counts per photon, and the Kodak-1035 chip 0.026 counts per photon. Each camera has a 12-bit DAC, so that each pixel records between 0 and 4095 counts.

There are two basic methods of imaging an ensemble of cold atoms: fluorescence imaging and absorption imaging. We capture fluorescence images of MOTs for qualitative diagnostic purposes when optimising the experimental parameters for making BECs. All quantitative data from the experiment comes from absorption images of cold atomic ensembles and BECs.

Imaging Set Up




Figure 3.11: Absorption Imaging Set Up.

The absorption imaging system is set up to allow us to image either along $\hat{x}$ or $\hat{z}$, as shown in figure 3.11. We refer to the former as imaging axis one, and the latter as imaging axis two. We switch between the two directions using a flip mirror that is adjusted by hand between experimental runs. The camera is aligned at an angle (25°) to imaging axis 1, chosen to avoid etaloning of the imaging beam by the glass casing around the CCD array of the camera. As a result, the horizontal and vertical magnifications are different. Imaging axis 2 is aligned parallel to the horizontal MOT beam, with which it shares its optics. It is combined with the MOT beam using a polarising beam cube, so that it has an orthogonal polarisation. After exiting the chamber the beam passes through a second $\lambda/4$ waveplate and is then picked off with a second polarising beam cube. The beam is directed onto the same camera as before, also at an angle of 25° , with two steering mirrors.

The imaging beam is picked off from the trap laser and passes through a 100MHz AOM (see figure 3.2) as described above. After the AOM the beam is linearly polarised with a polaroid linear polariser, spatially filtered with a $40\text{ }\mu\text{m}$ pinhole, and recollimated with a 12 mm full-width half-maximum gaussian profile. The large spot size ensures even illumination of the atom ensemble. The frequency of the trap laser during the imaging sequence is set using the VCO input on the frequency offset lock so that the imaging beam is tuned to resonance with the f_{23} transition of the imaged atoms in the appropriate uniform imaging magnetic field (aligned along the X- or Z-axis). The σ^+ polarisation with respect to the relevant imaging field is set by a $\lambda/4$ waveplate placed in the beam path immediately prior to entering the chamber. The AOM sets the duration of the imaging pulse at $30\text{ }\mu\text{s}$. The intensity of each imaging beam is controlled by a polarising beam cube and $\lambda/2$ waveplate in the beam path just in front of the flip mirror. We typically image with an intensity of $\sim 1\text{ mW cm}^{-2}$, which is a compromise between maximising the signal from the atoms (a greater percentage of light is absorbed at lower intensity) and using the full dynamic range of the camera.

Along axis one, the image is formed by a single spherical lens in a near- $2f$ configuration with magnification $M = 1.2$. The lens is a pair of achromatic doublets placed back to back with a nominal focal length of 80.5 mm (Comar 03 TT 25). It is mounted on a micrometer stage bolted to the optical table for fine focal adjustment, and typically placed a little under 150 mm from the centre of the chamber. The resolution of this imaging set-up is $6.5\text{ }\mu\text{m}$. It can also be used to capture fluorescence images of the MOT that are useful when optimising the experimental sequence.

Along axis two, the image is formed by two spherical lenses, both precision achromats (Melles-Griot 01 LAO 625 and 01 LAO 688). The first has a nominal focal length of 200 mm, the second a nominal focal length of 400 mm. The magnification along this imaging axis $M = 1.8$. The lenses are mounted on a micrometer stage bolted to the optical table for fine adjustment. The micrometer stage allows separate Z-axis adjustment of the two lenses, but common X-Y adjustment.

This imaging system is designed to give optimal resolution given the limitations of the optical access to the chamber along this axis. The numerical aperture of the imaging system is ultimately limited by the diameter of the viewport on the chamber and the distance to the centre of the chamber. The viewport is 110 mm from the centre of the chamber and has a diameter of 30 mm, which gives a diffraction limited resolution of $3.5\text{ }\mu\text{m}$ at 780 nm. The numerical aperture is currently limited by the 25 mm diameter of the beam cube shown in figure 3.11, which is 175 mm from the centre of the chamber. This introduces spherical aberrations that decrease the resolution of the imaging set-up to approximately $7\text{ }\mu\text{m}$.

Making BECs

In this chapter I describe the experimental sequence we use to produce a BEC in our atom chip microtrap, and characterise a typical BEC used in our interference experiments.

4.1 Experimental Sequence

4.1.1 Low Velocity Intense Source

Following the successful implementation of a double-MOT system using on a low velocity intense source of atoms (LVIS) and a mirror-MOT in the videotape experiment described in Chris Sinclair's thesis [53], which was based on the system described in [65], we decided to build a near identical system for the interferometer experiment.

The principle behind our double MOT system is to have one MOT, called the LVIS, operating at relatively high background pressure to load a second MOT maintained at low-pressure. The LVIS is used to trap and cool a large cloud of atoms, which are transferred to the low-pressure chamber via a small aperture. These atoms can be recaptured in the secondary MOT without a significant rise in the background pressure of the low-pressure chamber. This is critical for maintaining a long lifetime for the atomic cloud in the magnetic trap used for evaporative cooling to BEC.

The LVIS is a standard 6-beam MOT with one simple modification: one of the six laser beams has a shadow in the centre created by an aperture in one of the retro-reflecting optics, which we call the extraction mirror. Atoms in the centre of the LVIS feel a net force because of this shadow and are pushed out of the MOT, through the aperture, and towards the mirror-MOT in the main chamber, where the atoms are recaptured. The hole that defines the cold atom beam is made small to limit the molecular conductance between the LVIS vacuum chamber and the main chamber, allowing a large pressure differential to be




Figure 4.1: *Illustration of the operation of the LVIS. Shown are the MOT coils and quadrupole field, four of the six laser beams (the other two point into and out of the page), the trapped atoms at the centre of the LVIS, and the LVIS beam, which is pushed out through the hole in the extraction mirror towards the main chamber.*

maintained even during the operation of the double-MOT system.

We load the LVIS MOT from a Rubidium dispenser, typically run at 7.5 A for 15-17 s. In between experiments the dispenser is kept warm with a 2 A current to speed up the turn on of the dispenser and maintain a stable 'off' state. During this time the LVIS is centered on the shadow so that the LVIS beam is on and loading the mirror-MOT. After turning off the dispenser we allow the LVIS beam to run for 2 more seconds.

The LVIS beam is easily aligned once the two MOTs are both in operation, and can be optimised with minimal adjustment of the shim fields and push beam. It readily loads the mirror MOT with a physically large cloud with $\sim 10^9$ atoms. Operating in this way, the loading step from the mirror MOT to the magnetic trap is saturated, which aids the shot-to-shot stability of the experiment. Detailed diagnostics were carried out in by Chris Sinclair with the near-identical system used in the video-tape BEC experiment in our lab. He estimated the beam flux to be $\sim 1 \times 10^8$ atoms per second with a velocity of $20 \text{ ms}^{-1} < v < 10 \text{ ms}^{-1}$ and divergence of $\sim 40 \text{ mRad}$ [53].

4.1.2 Mirror MOT

The mirror-MOT approach illustrated in figure 4.2, used in the first atom-chip BEC experiments [66], takes advantage of the reflective surface of an atom chip such as ours. It employs 4 trapping beams, two parallel to the chip surface (which we refer to as the horizontal beams) and two that are reflected off the surface at 45° (the 45° beams). The




Figure 4.2: Illustration of the mirror MOT showing the two 45° MOT beams reflecting off the chip surface. The horizontal MOT beams are not shown (they point into and out of the page). I have labelled the MOT coils and quadrupole field, and the shim field used to position the MOT near the chip surface.

beams are circularly polarised as in a standard MOT, because the beams have their helicity reversed on reflection. The major advantage of this approach is that the mirror-MOT is semi-integrated into the atom-chip, which allows us to pass through a UMOT stage when loading the magnetic trap (discussed below).

The required quadrupole magnetic field comes from coils mounted in the chamber at 45° to the surface of the chip. The 45° MOT beams pass through one of these, which makes for easy alignment. These two beams enter the vacuum chamber through the large viewport below the chip. The diameter of the MOT beams is 25 mm, which is slightly larger than the width of the atom chip. We thus use the entire atom chip surface to make the mirror-MOT. The horizontal beams, perpendicular to the page in figure 4.2, are aligned with respect to the viewports on the vacuum chamber (using pinholes placed in front of the two viewports). Crude alignment of the beams with the MOT-coil former and the chamber viewports is enough for the mirror-MOT to work well. Fine adjustment of the alignment is made by tweaking the pointing of one of the horizontal beams to ensure that the UMOT is properly optimised (see below).

During the 15-17 s loading phase we run 2 A through the coils to give a gradient along the axis of the coils of 11 G cm^{-1} . The trap laser is 3Γ red-detuned from resonance with f_{23} during the loading phase, where Γ is the natural linewidth of the D2 cooling transition (6 MHz). The total power in the trapping lasers is approximately 15 mW per beam, adjusted to maximise the number of atoms in the mirror-MOT. We adjust the centre of the quadrupole field using the main chamber shim coils in order to maximise the number of atoms we load into the mirror-MOT. The initial MOT is loaded around 4 mm below the

surface of the atom chip. Previous studies in the group have indicated that the mirror-MOT loads inefficiently any closer to the surface than this [67] [52].

The mirror-MOT loading sequence, including the mirror-MOT and LVIS MOT parameters (beam pointing and shim fields), is adjusted empirically to give the largest possible mirror-MOT at the end of the loading sequence. In practice this is estimated from uncalibrated fluorescence images of the mirror-MOT. These images do not allow us to make absolute measurements of the number of atoms trapped in the mirror-MOT, but they do give reasonably accurate relative measurements of the mirror-MOT size. The absorption image, even at a short time of flight, extends beyond the field of view of the imaging system, which is framed by the chip and the MOT coils former. We estimate that we typically capture over 10^9 atoms at a temperature of a few hundred μK during the loading cycle.

4.1.3 UMOT

Once we have loaded the mirror-MOT with as many atoms as possible, we bring the trap closer to the surface in order to transfer the atoms to a magnetic trap. This step in the experimental sequence, efficiently transferring atoms from the dissipative MOT to the conservative magnetic trap, is perhaps the most difficult and important in making a BEC. The two traps must be well overlapped so that as many atoms as possible are loaded into the centre of the magnetic trap to prevent excessive heating in the transfer. The goal is to maintain as high a phase space density as possible.




Figure 4.3: Illustration of the UMOT, which uses the same trapping light as the mirror-MOT, but a quadrupole field generated by the chip wires which a net current that runs in a U-shape. The field from these wires produces the quadrupole shown near the surface of the chip with the addition of the bias field labelled in the diagram.

The vital step in our experiment in making this transfer is to use an intermediary MOT

in which the quadrupole magnetic field is provided by the chip wires rather than the MOT coils. We run 2 A in each of two chip Z-wires and 4 A in one of the end-wires. The net current from the Z-wires and the end-wire runs in a U-shape, illustrated in figure 4.2. If we add a uniform bias field along $\hat{x}$, the net magnetic field forms a quadrupole below the chip wire. The center of this quadrupole is not immediately below the central wire, but just on the outside the U. It can be rotated into closer alignment below the centre of the chip wire by adding a uniform field along $\hat{y}$. Since the same Z-wires are used to make the magnetic trap that catches the cloud when released from the UMOT, the two traps are automatically aligned provided that the MOT beams are properly balanced (which can be easily achieved).

The UMOT has the added benefit of compressing the MOT in the direction perpendicular to the chip wire because the magnetic field gradient from the chip wires is several hundred Gauss per centimeter in the radial direction, and less than 2 G cm^{-1} along the axis of the trap. The shape of the mirror-MOT thus changes from roughly spheroid to an elongated ellipsoid that more closely matches the contours of the Z-wire magnetic trap. This change in shape also allows the atom density in the centre of the UMOT to increase. The density in a MOT is ultimately limited by the radiation pressure of photons absorbed from the trapping lasers and multiply rescattered by the atoms. This pressure can be reduced by further detuning the trapping laser or reducing its intensity. It can also be reduced by changing the trap geometry: in an elongated UMOT photons do not need to travel very far to exit the trapping region radially, which reduces the probability of rescattering. Earlier studies in our group have shown that the density (and phase space density) of a MOT can be significantly increased by a combination of detuning the trapping laser and compressing the MOT into an elongated geometry [67] [52]. We follow both these strategies when transferring to the UMOT.

We transfer to the UMOT in two steps. The initial mirror-MOT is first brought closer to the surface by ramping on a uniform bias field along the X-axis (we refer to this as the X_{bias} field - it is also used to make the Z-wire magnetic trap). This ramp occurs over 50 ms, and is accompanied by adjustments of the shim fields to align the mirror-MOT with the UMOT. At this point the mirror-MOT is centered about 1 mm below the surface of the chip. We then simultaneously switch off the MOT coils and switch on the chip wires, replacing one quadrupole field with another. The trap lasers stay on during the switch. We hold the UMOT in this configuration for 20 ms to allow the atoms to be recaptured, then move the trap closer to the surface again by increasing the X_{bias} field with a linear ramp over 20 ms in preparation for transferring to the Z-wire magnetic trap. As we move the UMOT closer to the surface, we simultaneously red-detune the trapping laser to 6Γ away from resonance with the trapping transition, further increasing the density of the

UMOT near the centre of the trap.

At the end of this transfer sequence the UMOT is 0.5 mm below the surface of the chip. There are about 10^8 atoms in the cloud, at a temperature of $\sim 300 \mu\text{K}$.

4.1.4 Optical Molasses

At the end of the UMOT stage, we further cool the atoms with a short optical molasses stage lasting 3 ms. During this stage all the magnetic fields are switched off, the laser is further detuned to $\sim 20\Gamma$, and the laser power is reduced to $\sim 3 \text{ mW}$ per beam. An efficient molasses requires well balanced light forces and small residual magnetic fields (less than 100 mG). We make no attempt to cancel all stray bias fields and implement a full optical molasses stage. However this short molasses stage does increase our the efficiency of loading atoms into the magnetic trap. We can also observe the molasses dissipating in real time on the infra-red video camera set up to observe the main chamber. This provides a useful diagnostic for the balancing of light forces in the experiment, since the molasses will rapidly shoot off in one direction if the balance is poor.

4.1.5 Optical Pumping

Before catching the atoms from the UMOT in the magnetic trap, we optically pump them into the $F = 2$, $m_F = 2$ stretched state. We use σ^+ polarised light slightly red-detuned from the f_{22} transition, which rapidly transfers the atoms into the $F = 2$, $m_F = 2$ state. Since this state is dark for this polarisation, there is little residual heating of the atoms during this stage. The optical pumping beam is red-detuned 13 MHz from the f_{22} optical pumping transition using a double pass AOM, shown in figure 3.2. We use red-detuned light because the atomic cloud is optically thick for resonant light. After the AOM the beam is spatially filtered, expanded to 6 mm in diameter (FWHM), and combined with the imaging beam along imaging axis one using a polarising beam cube. It is circularly polarised by the same $\lambda/4$ waveplate immediately prior to entering the chamber, but with the opposite helicity. We use the Xbias field, which is aligned along the X-axis in the opposite direction to the imaging field, to set the quantisation axis for the atoms, with respect to which the optical pumping beam polarisation is σ^+ . The optical pumping cycle lasts for 1.2 ms. The optical pumping pulse itself is $400 \mu\text{s}$ long and comes in the middle of this sequence. The pulse transfers nearly all the atoms into the $F = 2$, $m_F = 2$ state.




Figure 4.4: Illustration of the double Z-Wire configuration that we use for magnetic trapping. Each wire carries 1 A DC current, and a uniform bias field is added to create the trap. A further bias field aligned along $\hat{z}$ provides the field at the minimum of the trap (not shown in the figure).

4.1.6 Initial Magnetic Trap

The spin-polarised atoms can now be caught in a Ioffe-Pritchard magnetic trap shown in figure 4.4. In order to match the shape and position of the magnetic trap as closely as possible to the shape of the cloud released from the UMOT, we make as deep a trap as we can at the position of the atoms, about 0.5 mm below the surface of the chip. The depth is limited by the amount of current that we can safely run through the chip wires without overheating them. We run 2 A through the same two Z-wires that were used to make the UMOT, which automatically aligns the main trapping axis of the two traps. This gives a trap depth of 15 G, limited by the presence of a second trapping minimum near the chip surface. This also sets the position below the surface that we catch the atoms.

With these parameters we catch approximately 5×10^7 atoms in the initial magnetic traps at a temperature of $\sim 100 \mu\text{K}$. The trap frequencies in this configuration are $\omega_z = 2\pi \times 6\text{Hz}$ and $\omega_{\perp} = 2\pi \times 1\text{kHz}$.

4.1.7 Compressed Trap

The initial 2-wire Ioffe-Pritchard trap is held for 50 ms to ensure the trapping fields have stabilised and to catch as many atoms as possible. We then ramp to a compressed trap with much larger radial trapping frequencies that is more favourable for evaporative cooling. We linearly ramp up the Xbias field from 15 G to 27.5 G in 100 ms. We cannot maintain 2 A in two chip wires without overheating them, so we ramp off the current in one of the wires as the Xbias field is increased. After the compression the trap frequencies are $\omega_z = 2\pi \times 28\text{Hz}$, $\omega_{\perp} = 2\pi \times 2\text{kHz}$, the field at the bottom of the trap is 0.9 G, and the

trap is $130\,\mu\text{m}$ below the surface of the chip.

During the compression the atoms are heated to $\sim 500\,\mu\text{K}$. The trap depth does not increase in proportion to the heating of the cloud, so the hotter atoms spill out of the compressed trap. This initial spilling occurs over a period of about 2 s, during which time we hold the trap without applying any RF. At the end of the 2 s we are left with 5×10^6 atoms in the trap at a temperature of $\sim 100\,\mu\text{K}$. This is the starting point for forced evaporative cooling to BEC. We try to operate the experiment so that we cannot increase the number of atoms left in the compressed trap at this point. In this sense we saturate the loading sequence, deliberately operating with a much larger MOT and initial magnetic trap than is needed, so that the starting point for evaporation to BEC is less sensitive to long term drifts in the experimental parameters such as the trap laser power or beam pointing. The background-collision limited lifetime of the trap is over 20 seconds, which is more than enough to evaporate to BEC.

4.1.8 Evaporative Cooling

Once we have allowed the hot atoms to spill out of the compressed trap, we begin forced evaporative cooling of the cloud. We typically apply an RF current to the copper-H underneath the atom chip (see figure 3.6 in chapter 3) to selectively remove the hot atoms via induced spin flips. We evaporate to BEC with a single exponential ramp from 15 MHz to 0.67 MHz over 4.6 seconds. During the final stage of the evaporation ramp we reduce the power in the RF to minimise the effect of power-broadening.

4.2 Cold Thermal Clouds

Towards the end of the evaporation ramp the cold thermal cloud begins to break up into lumps as the atoms begin to see anomalous field components $\delta B_z(z)$ along the length of the wire that have their origin in imperfections in the wires used to make the trapping fields (see the discussion in section 2.2 of chapter 2). These components give rise to a spatially varying trapping potential along the length of the wire, which the atoms see when they become sufficiently cold.

The practical consequences of this potential roughness for our experiment are twofold: we are limited in how close we can bring the cold atomic cloud to the wires because the anomalous components scale exponentially with distance, and they introduce both a random component to the field at the minimum of the trapping potential and a curvature $\partial B/\partial z$ in the trapping field that we cannot calculate a-priori. Neither of these is a major limitation. We do not need to bring our condensates closer to the trapping wires, and

we can measure empirically both the trap bottom and axial trapping frequency, and since the potential roughness is determined by material imperfections in the wires it does not change over time.

4.2.1 Fragmentation




Figure 4.5: *Fragmentation of Cold Thermal Clouds: Towards the end of the evaporation ramp, the cold thermal cloud begins to sample the rough potential at the bottom of the trap due to anomalous field components that have their origin in deviations in the flow of current through imperfect wires. Colder clouds break up into lumps along the length of the wire. We see here absorption images of successively colder thermal clouds in this rough potential, and the corresponding line density of each cloud. A condensate begins to form in the largest of these lumps in the final image. We can analyse the density distribution of the cold thermal clouds (the first three images) to extract information about the anomalous field components - see the main text.*

We can infer variations in the trapping potential along the z -axis from the density distribution of a thermal atomic cloud imaged in free fall. We take an absorption image of the cloud a short time after the trapping potential has been switched off, during which time there is no significant axial expansion of the cloud (typical numbers are an axial trapping frequency $\omega_z = 2\pi \times 6\text{Hz}$, a temperature of $1\mu\text{K}$ and a 2 ms time of flight, during which time the cloud has expanded axially less than 1%). The axial density distribution in free fall thus closely approximates the in-trap distribution.

The temperature of the cloud is determined from the transverse expansion of the cloud in free fall, where the width of the cloud is measured at various different times of flight (see figure 4.6 for an example measurement). By measuring the temperature in this way along the length of the cloud we can also check that the cloud is in thermal equilibrium, and that we are justified in treating the axial and radial trapping potentials as separable. We can also infer the distance of the cloud from the trapping wire by measuring the position of the centre of mass of the cloud as it falls.




Figure 4.6: Temperature and Magnification: We measure the temperature of the cold thermal cloud by fitting a gaussian to the transverse density distribution at various drop times and solving equation (2.29) to find the rate of expansion. The same measurement is used to cross-check the magnification of our imaging system by fitting the acceleration of the centre of mass of the cloud to the expected acceleration due to gravity. For the data presented here, the temperature of the cloud is $\sim 1.25\mu\text{K}$ and the magnification (on imaging axis one) is confirmed to be 1.2, which is what we calculate from the geometry of our imaging set up (see section 3.9 of chapter 3).

Since the cloud is in thermal equilibrium, the axial trapping potential can be inferred from the atomic density distribution using Boltzmann's law,

$$V(z) = -k_B T \ln(n_{th}(z)). \quad (4.1)$$

In the absence of fluctuations, the underlying potential is parabolic,¹ and we can fit to $V(z)$ with the cloud temperature as the fit parameter, and check the temperature of the cloud against the time of flight measurement. To extract information about the fluctuating component of the potential, we subtract the parabola from the measured density distribution, $\ln[n_{th}(z)]$. What remains is $\delta B_z/k_B T$, and since we know T we can directly infer δB_z . This procedure requires that the thermal cloud be continuous along the length of the trap so that we can accurately fit to the underlying potential.




Figure 4.7: This figure illustrates the steps in analysing the anomalous potential roughness from an absorption image of a cold atomic cloud. The three data sets correspond to the clouds shown in figure 4.5. The temperatures of the clouds, inferred from the gaussian fits to the axial density distribution shown in figure 4.7(a), were 7, 3.5 and 2 μK respectively. The number density is normalised to the total atom number. Figure 4.7(b) shows the axial trapping potential reconstructed from equation (4.1). The underlying harmonic potential is extracted from a fit to the data. After subtracting the harmonic potential, what remains is the anomalous potential roughness shown in figure 4.7(c). Only the centre section of this figure gives meaningful data because of the rapid drop off in signal to noise in wings of the distribution.

¹with a trapping frequency $\omega_z = 2\pi \times 6\text{Hz}$

This analysis is illustrated for three successively colder clouds in figure 4.7. The three data sets correspond to the clouds shown in figure 4.7. The temperatures of the clouds, inferred from the gaussian fits to the axial density distribution shown in figure 4.7(a), were $6.4\mu\text{K}$, $3.6\mu\text{K}$ and $1.9\mu\text{K}$ respectively. The number density is normalised to the total atom number. Figure 4.7(b) shows the axial trapping potential in arbitrary units, reconstructed from equation (4.1). The underlying harmonic potential from the ends of the Z-wires is extracted from a fit to the data. After subtracting the harmonic potential, what remains is the anomalous potential roughness shown in figure 4.7(c). To analyse the data we look at the centre section of the cloud where the signal to noise ratio is good. In the following I have taken the data from 0.9mm to 2.9mm in figure 4.7(c).

The scale of the potential roughness is set by the field of the wire at the position of the trap, $B_0 = \mu_0 I / 2\pi d = 24.8\text{Gauss}$. The root mean square roughness of the anomalous potential shown in figure 4.7(c) is $\delta B = 5.6\text{mG}$, which corresponds to $\delta B/B = 2.2 \times 10^{-4}$ at a distance of $156\mu\text{m}$ from the trapping wire. Therefore the angular variation of the current in the wire is $\sim 2 \times 10^{-4}\text{Rad}$, which, given the wavelength of the potential roughness is $\sim 0.5 - 1\text{mm}$, corresponds to a transverse wandering of the wires of $\sim 50 - 100\text{nm}$.

This shows that the roughness has nothing to do with microscopic wire imperfections of the type discussed in section 3.5 of chapter 3. It is due to uncontrolled factors in the fabrication process, probably during the reflow of the resist (such as acoustic noise during the heating phase). Whatever the origin, the practical consequence is that the axial trapping potential is dominated by contributions from the anomalous magnetic field components that cannot be calculated a priori, but must be measured empirically. Since the trapping wires are used to apply RF fields to create the adiabatic potentials we use in our interference experiments, these components will also exist in those fields. The fact that the dominant feature of the spectrum is at such a long wavelength is also good in the sense that it is much longer than the axial size of our condensate, which means that the condensate sits in a relatively smooth dip in the rough potential.

4.3 Bose Einstein Condensation

In the following section I describe and analyse the Bose-Einstein condensates that we use in our interference experiments. As with the majority of atom chip experiments, the BEC is highly elongated, which leads to important modifications to the properties of the BEC. In particular, we observe phase fluctuations along the long axis of the BEC that affect the coherence of the interferometer, as discussed in detail in chapter 7. The radial confinement is tight enough that our BECs are no longer fully 3D in character, but not so tight that they are rigorously 1D.

Our experimental data is extracted from measurements of the density distribution of the ensemble $n(\mathbf{r}) = n_0(\mathbf{r}) + n_{th}(\mathbf{r})$, normalised to the total number of particles $N = \int n(\mathbf{r}) = N_0 + N_{th}$. To analyse the density distribution of the falling cloud we integrate the 2D absorption image over the y -axis to obtain a line density $n(z)$. We then fit a gaussian to the wings of the thermal component and a Thomas-Fermi profile to the centre of the distribution

$$\begin{aligned} n_{th}(z) &= a_{th} \exp\left(-\frac{(z-z_0)^2}{2\sigma_z^2}\right) \\ n(z) &= \begin{cases} a \left(1 - \left(\frac{z}{L}\right)^2\right)^2 & \text{if } z \leq L \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (4.2)$$

4.3.1 BEC Transition

We typically reach BEC at a transition temperature $T_c \simeq 0.5 \mu\text{K}$. Figure 4.8 shows absorption images of a cloud after 10.4 ms time of flight as we cross the transition to BEC, along with corresponding slices through the centres of each cloud. The transition to BEC is clearly apparent in the emergence of a bimodal density distribution. The final slice shows the Thomas-Fermi density distribution of a nearly pure BEC with $\sim 1.5 \times 10^4$ atoms and a chemical potential $\mu = h \times 3 \text{ kHz}$. This is typical of the BECs that we use in our interference experiments.

In figure 4.9 we see absorption images of a BEC at different different times-of-flight. The anisotropic expansion of the BEC in the radial and axial directions is a clear marker of the transition to BEC, since a thermal cloud will expand isotropically, whereas the aspect ratio of an elongated BEC will be inverted in free fall (see the discussion in section 2.3 of chapter 2). The axial expansion of the condensate is barely apparent after 17.4 ms time of flight.

4.3.2 BEC Parameters

In order to control the parameters of the interference experiments that we run, it is important that we carefully characterise our static trapping potential and the parameters of our condensate. In the following section I summarise how we do this and present typical data from the day to day operation of our experiment.




Figure 4.8: *Evaporation to BEC:* We see here absorption images of the cold thermal cloud crossing the BEC phase transition. A pronounced bimodal density distribution typical of the BEC transition is evident in the second image. In the third image almost no thermal component is evident and we are left with a nearly pure condensate with $\sim 1.5 \times 10^4$ atoms and a chemical potential $\mu = h \times 3 \text{ kHz}$, which is typical of the condensates used in our interference experiments.




Figure 4.9: BEC in free fall after the trapping potential has been suddenly switched off. The time of flight is indicated at the top of each image. The aspect ratio of the BEC is clearly inverted in free fall.

Static Trapping Configuration

The static trapping configuration we use for our interference experiments is illustrated in figure 4.10. We use two broad wires to create the DC magnetic trap for our BEC: wire Z3 is $50\mu\text{m}$ wide and wire Z4 is $100\mu\text{m}$ wide. Both wires are $3\mu\text{m}$ thick, and the centre-to-centre separation of the trapping wires is $107.5\mu\text{m}$. To calculate the transverse components of our static trapping potential we use the analytical expression for the field of a wide wire given in equation (2.13). The quadrupole field around the trap minimum is slightly rotated because of the asymmetry between the two trapping wires. The finite thickness of the trapping wires makes no significant contribution to the trapping potential.

To make the DC trapping potential we run 1 A in each of the DC wires in series from a single current driver. In the illustration the current runs into the page, and the trap is formed by adding a bias field along the x -axis of 27.5 G. The trap forms $132\mu\text{m}$ below the surface of the chip. The transverse trap frequency is $\sim 2\pi \times 2 \pm 0.05\text{ kHz}$.

Trap Frequencies

We measure the axial trapping frequency by displacing the condensate with a field gradient, then releasing it and allowing it to oscillate in the trapping potential. The measured frequency, seen in figure 4.11(b) is $\omega_z = 2\pi \times 28.2(1)\text{ Hz}$.




Figure 4.10: *Static Trap Configuration:* We use two broad wires to create the DC magnetic trap for our BEC. Wire Z3 (on the left in the figure) is $50\,\mu\text{m}$ wide and wire Z4 (on the right in the figure) is $100\,\mu\text{m}$ wide. Both wires are $3\,\mu\text{m}$ thick. The centre-to-centre separation of the trapping wires is $107.5\,\mu\text{m}$. Note that because the trapping fields are generated by two wires, there are two local field minima. The BEC forms in the lower of the two trapping minima shown in the figure. To calculate the transverse components of our static trapping potential we use the analytical expression for the field of a wide wire given in equation (2.13). We run 1 A in each of the DC wires in series from a single current driver. In the figure the current runs into the page, and the trap is formed by adding a bias field along the x -axis of 27.5 G. The trap forms $132\,\mu\text{m}$ below the surface of the chip. The finite thickness of the trapping wires makes no significant contribution to the trapping potential. The quadrupole field around the trap minimum is slightly rotated because of the asymmetry between the two trapping wires.




Figure 4.11: Trap Frequencies: We see here the measured axial and radial trapping frequencies. The axial frequency is measured by displacing the condensate with a field gradient, then releasing it and allowing it to oscillate in the trapping potential. The radial frequencies are measured via parametric heating of the atoms using an AC current drive in the copper-H substructure below the atom chip (see section 3.3 of chapter 3), which produced an AC field oriented along the x -axis. Two frequency components are evident in figure 4.11(a), which is expected because the asymmetry of the trapping wires introduces a small asymmetry in the radial trapping potential.

The radial trapping frequencies are measured via parametric heating of the atoms using an AC current drive in the copper-H substructure below the atom chip (see section 3.3 of chapter 3), which produces an AC field oriented along the x -axis. This oscillating field shakes the trap. When the frequency of the driving AC field is resonant with the trapping frequency, the atoms will be heated out of the condensate. Two frequency components at about 2 kHz are evident in figure 4.11(a).

Trap Minimum & Condensate Lifetime

We measure the static magnetic field at the position of the atoms by outcoupling atoms with an RF knife. The RF field is generated by applying an RF current to the copper-H below the chip, just as in the parametric heating measurement of the radial trapping frequencies. We scan the RF frequency and monitor the outcoupling of atoms from the trap. Shown in figure 4.12(a) is a typical quick scan of the trap used to check the magnetic field the trap minimum before running an experiment. From day to day the trap minimum will change by up to ± 5 kHz, but it is stable over the course of a given experiment (run over a few hours) to within ± 1 kHz. In this data the trap bottom is at ~ 648 kHz. For our interference experiments we use a trap with a minimum at 630 kHz.

The condensate lifetime is measured with an RF knife held at the end frequency of the evaporation ramp so that any atoms that are heated out of the condensate are lost from the trap. The measured lifetime of ~ 1.25 s is much longer than the time needed to run an interference experiment. The lifetime of the condensate is consistent with measured




Figure 4.12: Trap Minimum & Condensate Lifetime: We measure the trap minimum by outcoupling atoms with an RF knife. Shown here is a typical quick scan of the trap used to check the position of the trap minimum before running an experiment. From day to day the trap minimum will change by up to ± 5 kHz, but it is stable over the course of a given experiment (run over a few hours) to within ± 1 kHz, which corresponds to a magnetic field change of ~ 1 mG. The minimum here is at ~ 648 kHz. For our interference experiments we use a trap with a minimum at 630 kHz. The condensate lifetime is measured with an RF knife held at the end frequency of the evaporation ramp so that any atoms that are heated out of the condensate are lost from the trap. The measured lifetime of ~ 1.25 s is much longer than the time needed to run an interference experiment.

technical noise in the apparatus. We have not been able to completely remove noise in the chip wire currents, perhaps because of the cross-connections amongst the wires and between the wires and the experiment common point via the chamber described in section 3.3 of chapter 3.

Density Profile




Figure 4.13: Condensate Density Profile: The figure on the left shows the axial line density of the BEC. A 3D Thomas-Fermi profile is a very good approximation of this density distribution. The figures on the right shows the radial line density. Two fit functions are shown: a 3D Thomas-Fermi profile (blue) and a gaussian (red), which is the expected radial profile of a 1D condensate. Both fit functions give a reasonable approximation to the data.

We typically fit the density profile of the BEC to the fit functions given in equation (4.2), which describe a three dimensional condensate. The axial line density closely fits the expected Thomas Fermi profile, as can be seen in figure 4.13(a). In this figure the peak linear number density is $\sim 300 \mu\text{m}^{-1}$. The radial line density is also closely approximated by a 3D Thomas Fermi profile. If we fit a gaussian distribution, the expected radial distribution in the 3D-1D crossover regime, the fit (shown in red) is slightly better, which suggests that we do start to sample some of the physics of the crossover regime.

For the BEC analysed in this chapter we calculate from equation (2.59) that $\chi \sim 4.53$ and $\alpha \sim 3.88$, which suggests that we are in the 3D Thomas Fermi regime. The crossover point between 3D and 1D behaviour in the mean field regime occurs when the chemical potential $\mu < \hbar\omega_{\perp}$, which leads to the condition $n_1 a \ll 1$, corresponding to a linear density of $\simeq 200 \mu\text{m}^{-1}$ for ^{87}Rb atoms in the $F = 2, m_F = 2$ state. The peak linear density in the BEC analysed in this chapter is $n_1 \sim 300 \text{atoms} \mu\text{m}^{-1}$, putting us in the 3D Thomas Fermi regime. However the average linear density is much lower (around $130 \text{atoms} \mu\text{m}^{-1}$).

4.3.3 Phase Fluctuations

Although the condensates we use do not cross over into the 1D regime, the aspect ratio is sufficiently large (typically $\omega_{\perp}/\omega_z \simeq 71$) to induce phase fluctuations along the length of the condensate. These are of practical consequence in our interference experiments since they can introduce a changing relative phase along the length of the condensate. In the following we characterising these longitudinal phase fluctuations, following the theoretical outline discussed in section 2.3 of chapter 2.

In order to measure the amplitude of the phase fluctuations in the lab, we follow the procedure outlined in the papers by Dettmer et al. [49] [50]. The elongated condensate is allowed to expand freely after the trapping fields are suddenly turned off. In free fall, phase fluctuations along the length of the cloud are turned into density fluctuations because the velocity field of the condensate is directly related to the phase via equation (2.69). We measure the amplitude of these density fluctuations and infer the amplitude of the phase fluctuations from this measurement.

Typical images of our condensates at 15 ms time of flight are shown in figure 4.14. Significant density modulations along the length of the cloud appear at long drop times. In the lab we measure the amplitude of these fluctuations $(\sigma_{\text{BEC}}/n_0)^2$, defined by equation (2.71). We can then use this to estimate the temperature of the BEC, given that we know the trap frequencies, chemical potential and drop time.

We measure $(\sigma_{\text{BEC}}/n_0)^2$ with the following procedure: we fit a Thomas Fermi profile $n(z) = b + a(1 - (z - c)^2/w^2)^2$ to the axial line density of the condensate, from which we




Figure 4.14: *Emergence of Phase Fluctuations:* We observe density modulations in free fall due to longitudinal phase fluctuations in our condensates. The top two figures show typical absorption images after 15 ms time of flight. In the figure on the left there are significant density modulations. The corresponding line density along the length of the cloud (integrated across the y -axis) shows pronounced deviations from the Thomas-Fermi profile of an ideal condensate. The amplitude of these fluctuations is larger in the centre of the cloud than in the wings (where imaging noise dominates). In the image on the right the evaporation ramp was taken to a lower end frequency so that the temperature and size of the resulting condensate is smaller. Some density modulations are still apparent, but the amplitude is much smaller. This is typical of the condensates we use in our interference experiments.

estimate the number of atoms in the BEC and the chemical potential; we then estimate $(\sigma_{\text{BEC}}/n_0)^2$ by calculating the mean square deviation of the density profile from the fit near the centre of the BEC ($|z| < L/2$); we estimate the background noise $(\sigma_T/n_0)^2$ by calculating the mean square deviation of the density profile from the fit in the wings of the profiles ($|z| > L$).

Significant density fluctuations $\sigma_{\text{BEC}} > 2\sigma_T$ appear at long drop times. The position of the modulations varies from shot to shot (since the phase fluctuations are random), but the amplitude remains fixed for a given condensate. From these measurements we can estimate the temperature of the BEC - in the condensate presented on the left of figure 4.14 the temperature estimate is $T \simeq 100 \text{ nK}$. Since we have no independent means of calibrating this measurement it is difficult to estimate the error in this figure. The

estimate of 100 nK takes into account the finite resolution of the imaging system ($6.5 \mu\text{m}$) - see [50] for details of how this is done.

When we evaporate further so that the condensate has fewer atoms, the amplitude of the density fluctuations diminishes, as in the image on the right of figure 4.14. This is typical of the condensates that we use for our interference experiments. For this BEC, we estimate a phase coherence length of $23 \mu\text{m}$, which is about half the condensate length ($41 \mu\text{m}$). So we expect there to be roughly two independent phase domains along the length of the condensate, each with half the atoms in the whole condensate. The effect of this on our interference experiments is discussed in chapter 7.

Elongated BEC




Figure 4.15: *Elongated BEC: Images of an elongated condensate made using a slightly different experimental set-up in our apparatus. The radial trapping frequency was $\omega_{\perp} = 2\pi \times 2.3 \text{ kHz}$ and the axial trapping frequency was $\omega_z = 2\pi \times 6.5 \text{ Hz}$. The resulting aspect ratio of ~ 350 is much greater than our typical BEC. The number of atoms in this condensate is typically 3×10^4 and the chemical potential $\mu = h \times 4 \text{ kHz}$. With these numbers the condensate is still not in the 1D regime, but enters deeper into the crossover region than our typical BEC. Significant density modulations at long drop times are due to the presence of longitudinal phase fluctuations. We have still however successfully carried out interference experiments with this BEC.*

As a point of comparison we have also carried out interference experiments with much more elongated condensates that exhibit far more pronounced phase fluctuations. A typ-

ical sequence of images taken at various different fall times are shown in figure 4.15. Significant density modulations are clear at long drop times. We estimate that the phase coherence length of this condensate is $3.5\text{ }\mu\text{m}$, which is much shorter than its Thomas Fermi length ($110\text{ }\mu\text{m}$).

This BEC is far more elongated than our typical condensate, with a radial trapping frequency of $\omega_{\perp} = 2\pi \times 2.3\text{ kHz}$ and an axial trapping frequency of $\omega_z = 2\pi \times 6.5\text{ Hz}$. The resulting aspect ratio is ~ 350 . The number of atoms in this condensate is 3×10^4 and the chemical potential $\mu = h \times 4\text{ kHz}$. The peak linear density is $\sim 95\text{ atoms }\mu\text{m}^{-1}$. We estimate the temperature at $\sim 300\text{ nK}$, well above the characteristic temperature for the disappearance of phase fluctuations (6 nK). I show data from interference experiments performed with this condensate at the end of chapter 7.

RF Adiabatic Potentials

5.1 Introduction




Figure 5.1: Illustration of how we produce a double well using RF adiabatic potentials. We use two parallel wires for both the DC and RF fields used to produce the trapping potentials. The static trapping field from a two wire trap is illustrated in black. Both wires carry the same DC current, and the field from these wires is cancelled by a uniform external bias field at a height h above the surface of the atom chip. A uniform bias field along the z -axis (not shown in this figure) sets the field at the minimum of the trapping potential. The RF current is added to the same two wires, but with a π phase shift between the two wires so that the vertical field component adds to give the RF field vector illustrated in red in the figure. This perturbs the atoms, forming a double well potential along the axis where the RF field is aligned with the DC quadrupole field so that the coupling of the atoms to the RF field is weakest.

The idea that cold atoms could be trapped and manipulated by RF adiabatic potentials was first put forward by Zobay and Garraway [68] [69]. The basic idea is simple: atoms sitting in a local static magnetic field are dressed by a strong near-resonant RF field. In this system the m_F states are thus no longer eigenstates. Instead, we can move to the basis of

dressed states, which are a combination of the internal state of the atom and the photons in the RF field. We can call these $\tilde{m}_F$ states, labelled by analogy with the Zeeman sublevels. If the energy of the $\tilde{m}_F$ states varies in space then the atoms feel a changing potential energy. By suitably engineering the static and RF magnetic fields we can then trap atoms in these dressed states. The $\tilde{m}_F$ states can be weak-field or strong-field seeking, but the restriction that only weak-field seekers can be trapped is relaxed because the potential that the atoms see is a result of the interaction between two spatially varying fields (the DC field and the RF field), which can have both local maxima and minima in free space.¹

Various configurations of RF adiabatic potentials have been used in cold atom experiments. The original idea of Zobay and Garraway was to produce a 2D trapping potential by dressing atoms in a quadrupole trap with a resonant RF field to produce a bubble trap. When gravity is added to the trapping potentials, the atoms will sit in a thin sheet at the bottom of the bubble. This trap has been successfully implemented by Perrin and co-workers [71] [72] and DeMarco and co-workers [73]. A similar configuration can be used to make a ring trap for atoms [74] (see also [75]). RF adiabatic potentials have also been used to alter the properties of static magnetic traps to reduce spin flip loss [76], to manipulate atoms in free fall, for example to make lenses [77].

To make a double well potential we follow the approach of Schumm et al. and dress atoms trapped in a Ioffe-Pritchard trap with a near-resonant linearly polarised RF field [5] [78]. In our experiment the RF fields are generated by AC currents added to the DC trapping wires, as illustrated in figure 5.1. The frequency is red-detuned with respect to the field at the minimum of the static trap. The polarisation dependence of the resulting RF-dressed state potential (or adiabatic potential) causes a double well to form along the axis of weakest coupling. This configuration has been studied extensively, both theoretically [79] [80] and experimentally [78] [81]. It is simple to implement and allows us to smoothly deform a single well into a double well potential such that we can coherently split a single condensate in two, as we will see in chapter 7.

Double Well Potential

The basic idea for creating the double well potential we use in our experiments is illustrated for a two level system in figure 5.2. We will assume that the RF field is linearly polarised along the y -axis, and the static trap is given by equation (2.6). The eigenenergies of the atoms in the static field (bare states) are given by the dashed lines, and the dressed state eigenenergies are given by the full lines.

There are two possible configurations that produce a double well, depending on the

¹Trapping configurations for strong field seeking $\tilde{m}_F$ states have been proposed - see for example [70].




Figure 5.2: RF Adiabatic Double Well Potential: There are two possible configurations for making an RF adiabatic double well potential depending on the detuning of the RF field with respect to the field at the minimum of the static magnetic trap. If $\nu_{RF} < \frac{|\mu_{BGF}|}{h} B_0$ the RF field is subresonant everywhere. This is illustrated in figure 5.2(a) where static magnetic potential is shown by the dashed line and the RF adiabatic potential by the solid line. The RF field vector coupling the bare levels is indicated by the arrow. In the dressed state basis, illustrated in figure 5.2(c), we see that the uncoupled states never cross. When coupled they repel each other, forming the double well potential. The 2D cross section of the resulting potential is plotted in figure 5.2(e). If $\nu_{RF} > \frac{|\mu_{BGF}|}{h} B_0$ the RF field is resonant at two points along the x -axis. The resonance points the uncoupled dressed states cross, and the coupling turns these into anti-crossings basis which form a double well for the atoms (figure 5.2(d)). The 2D cross section of the resulting potential is shown in figure 5.2(f).

detuning of the RF field with respect to the minimum of the local static field. If $\omega > \frac{|\mu_B g_F|}{\hbar} B_0$ (supraresonant) then the shape of the trapping potential is determined largely by the frequency of the RF field, which will be resonant with the local static field on a shell around the static trap minimum. Along the x -axis (where the double well forms) the RF is resonant at two points, indicated by the arrows in figure 5.2(b). In the dressed state basis (figure 5.2(d)) the bare state eigenenergies cross at these points. The RF field couples them so that this becomes an anti-crossing. If the atoms adiabatically follow the dressed state potential, they will be trapped in the minima formed at this anti-crossing. The RF field need only be strong enough to avoid spin flip loss at the trap minimum.

Alternatively $\omega < \frac{|\mu_B g_F|}{\hbar} B_0$ (subresonant), as in figure 5.2(a), in which case the shape of the trapping potential is determined more by the strength of the coupling term than the frequency. The bare states in this configuration do not cross, as can be seen in figure 5.2(c). However the RF field will still couple these states, pushing them apart. The coupling is strongest nearest to resonance, so the centre of the trapping potential will be pushed up further than the surrounding points. If the RF field is sufficiently intense, a double well potential will also form in this configuration because of this spatially varying coupling.

In both cases a double well potential will form along the axis of weakest coupling in the plane of the static quadrupole field. These two cases are illustrated in figure 5.2(e) and figure 5.2(f) respectively. This is the axis along which the RF field vector is parallel to the quadrupole field. For subresonant RF the centre of the trapping potential is a saddle point because the coupling term does not fall away sufficiently quickly along the y -axis to form a central barrier. The reason is that the RF field is always perpendicular to the local static field along this axis, so the change in coupling is only due to the change in the detuning between the local static field and the RF field. Along the x -axis the additional change in orientation between the RF field vector and local static field vector allows the double well to form. In the case of supraresonant RF a barrier forms along both axes, but the potential is lower along the x -axis because the coupling term is weaker due to the change in the orientation of the local static field vector and so a double well will still form. The potentials for an $F = 2$ system are formed similarly and are illustrated in figure 5.3.

Note that the atoms are coupled to the RF field across the entire trapping potential because of the non-zero z -component of the Ioffe-Pritchard field, which is everywhere perpendicular to the RF field vector. Along the y -axis the local static field rotates in a plane perpendicular to the RF field vector as we move away from the minimum of the Ioffe-Pritchard trap and the only change in coupling is due to the change in the magnitude of the local static field. Along the x -axis the static field rotates in the $x - z$ plane away




Figure 5.3: RF Adiabatic Potentials for the $F = 2$ System: On the left the potentials with a subresonant RF field are illustrated. On the right the potentials with a supresonant RF field are illustrated. We trap atoms in the $\tilde{m}_F = 2$ state using a subresonant RF field.

from the perpendicular, which strongly changes the coupling by introducing a significant parallel term to the RF field vector. For our interference experiments, the minimum of the double well is located about $2\,\mu\text{m}$ away from the trap centre along the x -axis, at which point the static field has increased by almost 10% (from $\sim 0.9\text{G}$ to $\sim 0.97\text{G}$) and the angle between the local static field and the RF field has changed from 90° to $\sim 69^\circ$.

In order to make a double well potential for the atoms with this configuration, the amplitude of the RF field must be large - in our case it is comparable to the magnitude of the static field at the minimum of the trap. Because of this, the commonly used rotating wave approximation to the interaction between the atoms and the dressing field breaks down. In the following I will describe the basic theory of the interaction between an atom and a strong RF field, and discuss analytical approximations beyond the simple rotating wave approximation, as well as full numerical calculations of the resulting dressed potentials. None of this theory is new, but some of it has not yet been addressed in this context.

5.2 Theoretical Approaches

5.2.1 Dressed Atoms

The most useful starting point for describing the interaction of an atom with spin $\mathbf{F}$ with an intense RF field is the dressed atom Hamiltonian, introduced by Haroche and Cohen-Tannoudji [82] [83] [84]

$$\begin{aligned}
H &= H_{static} + H_{field} + H_{RF} \\
&= \mu \mathbf{B}_s(\mathbf{r}) \cdot \mathbf{F} + \hbar \omega a^\dagger a + \mu \mathbf{B}_{RF}(\mathbf{r}, t) \cdot \mathbf{F}
\end{aligned} \tag{5.1}$$

where $\mathbf{B}_s(\mathbf{r})$ is the static magnetic field, $\mu = g_F \mu_B$ the scalar magnetic dipole moment, $\mathbf{F}$ is the atomic angular momentum, ω is the angular frequency of the RF field, and $\mathbf{B}_{RF}(\mathbf{r}, t)$ is the RF field. The operators $a^\dagger$ and a are the creation and annihilation operators of the RF field and satisfy the usual bosonic commutation relations.

We will assume that the static magnetic field and the parameters of the RF field vary slowly in space and time compared to the Larmor frequency of the dressed atom, so that the atom will adiabatically follow the dressed states of the Hamiltonian. For realistic experimental parameters this assumption will almost always be well satisfied. In this case, we can treat the interaction of the atom and the field described in equation (5.1) on a local point by point basis.

We can thus choose, for each point in space $\mathbf{r}$, a local coordinate system in which the static magnetic field is aligned along the z -axis so that $\mathbf{B}_s = B_0 \hat{\mathbf{z}}$, with $B_0 = |\mathbf{B}_s|$. The RF field can then be decomposed into a linearly polarised component B_z aligned along the z -axis, and two circularly polarised components B_a (rotating anticlockwise) and B_c (rotating clockwise) in the $\hat{x} - \hat{y}$ plane. The RF field is then written

$$\mathbf{B}_{RF} = \{B_a, B_c, B_z\}. \tag{5.2}$$

Dressed Atom Hamiltonian in the Local Frame

We can write the dressed atom Hamiltonian in terms of the local reference frame of the atom. The most useful basis states are the states $|n, m_F\rangle$, where n labels the number of photons in the RF field and m_F is the magnetic quantum number referred to the local z -axis $\mathbf{B}_s/|\mathbf{B}_s|$. We make the assumption that $\langle n \rangle = n$ is large and the variation around n is small. In this basis, the dressed atom Hamiltonian is

$$H = \mu B_0 m_F + \hbar \omega a^\dagger a + \frac{\mu}{2\sqrt{n}} \left(B_a a^\dagger F_- + B_c^* a^\dagger F_+ + B_z a^\dagger F_z + h.c. \right) \tag{5.3}$$

with $F_\pm = F_x \pm iF_y$. The first two terms of this Hamiltonian are diagonal. The others are all off diagonal terms representing couplings between the bare states of the system. The interaction part of the Hamiltonian has been split into three terms $H_{RF} = H_c + H_a + H_z$: the first term, containing $B_a a^\dagger F_- (+h.c.)$, will raise (lower) the number of RF photons in the field by one and lower (raise) the atomic spin state, the term containing $B_c a^\dagger F_+ (+h.c.)$ will raise (lower) both the spin state and the number of RF photons, while the term containing

$B_z a^\dagger (+h.c.)$ will raise (lower) the photon number but leave the spin state unchanged. We will refer to these terms as the co-rotating, counter-rotating and parallel respectively.

Coupling




Figure 5.4: Levels Coupled by RF Field Components: In this figure I illustrate examples of the levels coupled by each of the three different RF field components. The RF field is resonant with the local static field ($\omega = \omega_0$), so the $|+\rangle$ levels on the right are shifted by exactly one photon energy with respect to the $|-\rangle$ levels on the left. The co-rotating component (B_a) resonantly couples levels $|n, -\rangle$ and $|n-1, +\rangle$. The counter-rotating component (B_c) couples the levels $|n-1, -\rangle$ and $|n, +\rangle$ (which are $2\omega_0$ off resonance) and the parallel component (B_z) couples levels within one or the other spin manifold, which are always separated by ω_0 .

It is easy to describe the action of each of the interaction terms with a spin-1/2 system. We will use the notation $|n, \pm\rangle$ for the eigenstate corresponding to n photons and spin up (down). In the absence of interactions these terms have eigenenergies $n\hbar\omega \pm \hbar\omega_0/2$, where we have defined $\omega_0 = \mu B_0/\hbar$ and we have taken μ to be positive. The coupling terms are illustrated graphically in figure 5.4 for the case that the RF field is resonant with the (local) static magnetic field so that $\omega = \omega_0$. In this case the $|n, +\rangle$ levels are shifted in energy by exactly one photon with respect to the $|n, -\rangle$ levels.

The co-rotating component (B_a) is clearly the resonant (energy conserving) term of the interaction Hamiltonian. It couples levels $|n, -\rangle$ with $|n-1, +\rangle$. The counter-rotating component (B_c) couples levels $|n-1, -\rangle$ with $|n, +\rangle$, and the parallel component (B_z) couples levels within one manifold of states, $|\pm\rangle$. For systems of arbitrary spin these terms play the same role, though the corresponding diagram will become more complicated.

If we introduce a detuning $\Delta = \omega_0 - \omega$ between the RF field and the local static field, the difference in energy of the coupled levels for each term is Δ , $2\mu B_0/\hbar - \Delta$ and ω respectively. Thus for small detunings, the co-rotating term will remain closest to resonance. The counter-rotating term is always off-resonant, and the parallel term is resonant

at $\omega = 0$, i.e. when the RF field is just a constant addition to the static field.

Depending upon the terms that are kept in the interaction Hamiltonian, different manifolds of states will be coupled by the matrix elements of the dressed state Hamiltonian. The three interaction terms give rise to selection rules that describe which matrix elements are coupled by each term

$$\begin{aligned} \Delta m_F &= 0 \quad \text{for} \quad B_z \\ \Delta m_F + \Delta n &= 0 \quad \text{for} \quad B_a \\ \text{and} \quad \Delta m_F - \Delta n &= 0 \quad \text{for} \quad B_c. \end{aligned} \tag{5.4}$$

Rotating Wave Approximation

The rotating wave approximation consists of ignoring the far off-resonant terms in the Hamiltonian of equation (5.3) and keeping only the co-rotating terms. The resulting Hamiltonian can be diagonalised analytically.

The dressed Hamiltonian in the rotating wave approximation is

$$H_{RWA} = \mu B_0 F_z + \hbar \omega a^\dagger a + \frac{\mu}{\sqrt{2n}} (B_a a^\dagger F_- + h.c.) \tag{5.5}$$

which has eigenenergies

$$\begin{aligned} E_{RWA} &= \tilde{m}_F \sqrt{(\mu B_0 - \hbar \omega)^2 + \mu^2 B_a^2} + n \hbar \omega \\ &= \hbar \tilde{m}_F \sqrt{\Delta^2 + \Omega^2} + n \hbar \omega \end{aligned} \tag{5.6}$$

where $\tilde{m}_F = -F \dots F$ labels the pseudo- m_F states of the dressed system and I have introduced the detuning $\Delta = \mu B_0 / \hbar - \omega$ and the Rabi frequency $\Omega = \mu |B_a| / \hbar$. The field is approximately uniform across the trap, so we can drop the $n \hbar \omega$ term from the potential. The resulting RF adiabatic potentials are then

$$V_{RWA}(\mathbf{r}) = \hbar \tilde{m}_F \sqrt{\Delta(\mathbf{r})^2 + \Omega(\mathbf{r})^2}. \tag{5.7}$$

It is the spatial dependence of the detuning and the coupling term that gives rise to the double well potential described in section 5.1 above. We can think of the atom as moving in an effective static magnetic field $\Delta \hat{z} + \Omega \hat{e}_\perp$ with a magnitude $\sqrt{\Delta^2 + \Omega^2}$ and polar angle given by $\tan(\theta) = -\Omega/\Delta$. This identification of an effective static magnetic field for the dressed atoms leads to an alternative description of the effect of the dressing field, which can be said to modify the g -factor of the atom, changing how it responds to the static

field.

Recalling the adiabaticity criterion for trapping atoms in a static field (see equation (2.4)), we can also write down a semi-classical adiabatic criterion for the dressed states, requiring that the effective field vector changes slowly compared to the effective Larmor frequency $\sqrt{\Delta^2 + \Omega^2}$.

The rotating wave approximation captures all of the physics required to explain how a double well potential is formed. However it does not give an adequate quantitative description of the potentials that we make with our atom chip. This is because we use large amplitude RF fields so that the counter-rotating B_c and parallel B_z components of the RF field make a significant contribution to the real trapping potential. Higher order expansions of the eigenenergy of an atom dressed by a linearly polarised RF field (which is what we use in our experiment) can be found in the literature and include the contribution of these components [85] [86]. In our notation, the 4th order expansion for the potential is

$$V = \tilde{m}_F \sqrt{\Delta^2 + \Omega^2 \left(1 + \frac{\Delta}{\omega_0 + \omega}\right) \left(1 - \frac{\Omega^2}{4(\omega + \omega_0)^2} + \frac{3\Theta^2}{(\omega_0^2 - 4\omega^2)}\right)} \quad (5.8)$$

where I have defined $\Theta = \mu B_{\parallel} / 2\hbar$ in analogy with the definition of Ω . We have also used equation (5.8) to calculate our trap potentials. The agreement near the trap centre (where the detuning is small so that the atoms are near the resonance point) is excellent, even for RF fields with amplitudes larger than we typically use, and much better than within the rotating wave approximation. However, as we move away from the trap centre and the detuning is no longer small, equation (5.8) also breaks down. A comparison for realistic trapping potentials is made below.

5.2.2 Full Numerical Calculation

For accurate calculation of our trap potentials we need to numerically diagonalise the full Hamiltonian of equation (5.3) at each point in space. The basis states for the numerical calculation are the bare states of the undressed atom $|m_F\rangle$, where $m_F = -F \dots F$ and the RF field with n photons, i.e. $|n, m_F\rangle$. In this basis the first two terms of equation (5.3) are diagonal, with matrix elements $m_F \mu B_0$ and $n\hbar\omega$. All the other terms are off-diagonal. The matrix elements of the interaction terms of the Hamiltonian follow from the selection rules given in equation (5.4)

$$\langle n+1, m_F-1 | H_a | n, m_F \rangle = \frac{\mu}{2\sqrt{n}} B_a \sqrt{n+1} \sqrt{(F+m_F)(F-m_F+1)} \quad (5.9)$$

$$\langle n+1, m_F+1 | H_c | n, m_F \rangle = \frac{\mu}{2\sqrt{n}} B_c^* \sqrt{n+1} \sqrt{(F-m_F)(F+m_F+1)} \quad (5.10)$$

$$\langle n+1, m_F | H_z | n, m_F \rangle = \frac{\mu}{2\sqrt{n}} B_z \sqrt{n+1} m_F \quad (5.11)$$

plus complex conjugates. Under the assumption that $\langle n \rangle$ is large and the variation around $\langle n \rangle$ is small¹ $\sqrt{n+1} \simeq \sqrt{n}$. The matrix elements simplify to

$$\langle n+1, m_F-1 | H_a | n, m_F \rangle \simeq \frac{\mu B_a}{2} \sqrt{(F+m_F)(F-m_F+1)} \quad (5.12)$$

$$\langle n+1, m_F+1 | H_c | n, m_F \rangle \simeq \frac{\mu B_c^*}{2} \sqrt{(F-m_F)(F+m_F+1)} \quad (5.13)$$

$$\langle n+1, m_F | H_z | n, m_F \rangle \simeq \frac{\mu B_z}{2} m_F \quad (5.14)$$

plus complex conjugates. When all the components of the RF field are present, this matrix cannot be further broken down into smaller sub-system since all the states in the system are coupled. The Hamiltonian is infinite, so numerical diagonalisation requires truncation at some number of photons. To calculate the trap potentials we truncate the matrix at a suitable choice of $n \pm \Delta n$ around an arbitrary reference point n . Our states are then $|\Delta n, m_F\rangle$, with $|\Delta n| \leq \Delta n_{\max}$. The choice of $\Delta n_{\max}$ requires a compromise between accuracy and available computing power. Physically, the truncation involves ignoring higher order multiphoton transitions. We check the effects truncation numerically to make sure that the error introduced is small. With $\Delta n_{\max} = 4$ the errors introduced in our potentials are at the level of a few parts per million.

We can make the numerical calculations more tractable by using a method introduced by Majorana, who showed that a system with arbitrary angular momentum F can be represented by $2F$ spin-1/2 sub-systems [87]. In practice this means that in order to solve the Schrödinger equation for a spin- F system, we need only solve the corresponding spin-1/2 problem, and then write down the spin- F wavefunctions as a superposition of the spin-1/2 terms with appropriate amplitudes. The procedure for doing so is also outlined in a paper by Bloch and Rabi [88].

For the calculation of eigenenergies in the stretched state $\tilde{m}_F = 2$, which we need to calculate our trapping potentials, the procedure is simple. We calculate the eigenenergies

¹for a typical RF field amplitude of 1 G, the average number of photons in a Poisson distributed coherent field is $\langle n \rangle \simeq 10^{28}$, so this assumption is very good - see [84].

for the spin-1/2 system and then scale these by $\tilde{m}_F/\frac{1}{2}$. The calculations for the spin-1/2 system involve $(2F+1)^2/4$ less terms, which makes them far more tractable and allows us to include more photons in the matrix and avoid truncation effects in our calculations. The eigenvectors of the spin- F system similarly can be calculated following the method outlined in [88]. The procedure can also be usefully applied to analytical calculations.

5.3 Calculating Realistic Trapping Potentials

Chip Wire Geometry




Figure 5.5: *RF Adiabatic Potential Geometry:* We implement the RF adiabatic potentials in our experiment using a simple two-wire geometry that gives us control over the orientation of a linearly polarised RF field relative to the quadrupole of the static magnetic trap. The wires used in the experiment are highlighted in figure 5.5(a). One is $50\mu\text{m}$ wide, the other $100\mu\text{m}$ wide. Both carry 1 A of DC current to create the static trap (see section 4.3 in chapter 4 for more details). Relative to the sketch in figure 5.5(b) the DC current flows into the page. The RF current is added to both wires via the circuit illustrated in figure 3.10. The RF current is run with a π phase shift between the two wires so that the vertical components of the two RF fields add to give a linearly polarised field as illustrated in figure 5.5(b). Because of the different widths of the two wires, if we run the same RF current in each wire the resulting RF field is oriented at a small angle to the vertical. The double well potential forms along the axis of weakest coupling between the RF field and the atoms in the local static field, which is highlighted in blue in the figure 5.5(b).

We implement the RF adiabatic potentials in our experiment based on the two wire geometry we use for creating static magnetic potentials. The geometry is illustrated in figure 5.5. We use the same two wires to produce both the static and RF fields. One wire (labelled Z3) is $50\mu\text{m}$ wide and the other (labelled Z4) is $100\mu\text{m}$ wide. The centre to centre separation of the two wires is $107.5\mu\text{m}$. We did not design our experiment with

this asymmetric configuration, but were left with only these two trapping wires after the failure of the other two available trapping wires. As a result, we must be very careful in how we calculate our RF adiabatic potentials in order to understand our interference experiments.

We calculate our DC trapping fields using the analytic expression in equation (2.13) for a finite-width wire and cross check our calculations using the measured radial trapping frequencies (see section 4.3.2 of chapter 4) and known DC currents and bias field calibration. The RF field is similarly calculated using equation (2.13) with a given RF current in the wires and phase between the current in the two wires. We calculate the resulting RF adiabatic potential using the full numerical expression described in section 5.2.2. In the following sections I also compare the two analytic expressions for the RF adiabatic potentials given by equation (5.7) and equation (5.8).




Figure 5.6: *RF Coupling:* In our experiment the field at the minimum of the Ioffe-Pritchard trap is ~ 0.9 G (or ~ 630 kHz). The RF field is red-detuned by 90 kHz with respect to the trap minimum ($\nu_{RF} = 540$ kHz). The RF field amplitude is ~ 0.8 G.

In our experiment the field at the minimum of the Ioffe-Pritchard trap is ~ 0.9 G (or ~ 630 kHz). The RF field is red-detuned by 90 kHz with respect to the trap minimum ($\nu_{RF} = 540$ kHz). The RF field amplitude is ~ 0.8 G. The energy of the dressed states in this configuration is illustrated in figure 5.6, where zero of the energy scale is given by the energy of the $m_F = 0$ state of the undressed Hamiltonian.

A cross section through the centre of a potential with a horizontal splitting axis and a separation between the two minima of $\sim 4 \mu\text{m}$ is illustrated in figure 5.7. For reasons explained below, this potential must be adjusted to obtain the balanced double well that we use in our interference experiments.




Figure 5.7: *Double Well Potential: Full numerical calculation of a typical double well potential. Shown in the figure is a cross section through the centre of the trap (i.e. with $z = 0$). This is the potential that is formed with the same RF currents running through the two wires as illustrated in figure 5.5. The splitting axis is almost horizontal in this configuration, and the separation between the two minima is $\sim 4 \mu\text{m}$.*

5.3.1 Accuracy of Analytical Approximations

In order to compare the different analytical approximations discussed in section 5.2 and to highlight the role of the different polarisation components of the RF field (co-rotating, counter-rotating and parallel), I discuss two different double well potential configurations in the following paragraphs. The first is made with the RF red-detuned by 90kHz away from resonance with the trap bottom, which is how we make the double well potentials we use in our experiments. The current in each wire is adjusted so that the double well potential is balanced (see the discussion below). In the second trap the RF frequency is set to resonance with the trap bottom (all other parameters are kept the same) so that the effect of detuning is made smaller.¹

¹With an RF field blue-detuned with respect to the trap bottom a double well potential can be formed with much weaker RF fields and the rotating wave approximation can give an accurate description of the




Figure 5.8: Analytic Approximations: A comparison of the full numerical calculation of the potential shown in figure 5.7 with the analytic approximations of equation (5.7) and equation (5.8). In all figures the full numerical calculation is plotted with a solid line, the rotating wave approximation of equation (5.7) is plotted with a dotted line, and the approximation of equation (5.8) is plotted with a dashed line. On the left is a potential made with the RF field red-detuned by 90kHz with respect to the trap bottom (which is at 630kHz) and a field amplitude at the centre of the trap of 0.8G. On the right is a potential made with the RF field tuned to resonance with the trap bottom and a field amplitude at the centre of the trap of 0.45G. At the top I plot the approximations with the minimum of each potential set to zero. This is the double well that would be seen by the atoms. On the bottom I plot the potentials relative to the energy of the $m_F = 0$ bare state. We can see that the higher order approximation correctly predicts the absolute energy shift of the potential near resonance in both configurations. For the trap on the left, however, the approximation breaks down away from the centre of the trap because the detuning becomes too large (see the discussion in the main text). For the resonant trap, the detuning stays small across the entire potential and the higher order approximation is excellent. The rotating wave approximation underestimates the shift in the potential when the RF is subresonant. It also fails to correctly predict the resonant potential when the amplitude of the RF field is large, as in the figures on the right.

In figure 5.8 I compare the full numerical calculation of the potential shown in figure 5.7 with the analytic approximations of equation (5.7) and equation (5.8). In all plots in the figure the full numerical calculation is shown with a solid line, the rotating wave approximation of equation (5.7) with a dotted line, and equation (5.8) with a dashed line. On the left is a potential made with the RF field red-detuned by 90 kHz with respect to the trap bottom and field amplitude at the centre of the trap is 0.8 G. On the right is a potential made with the RF field tuned to resonance with the trap bottom and the field amplitude at the centre of the trap is 0.45 G.

At the top I plot the approximations with the minimum of each potential set to zero. This is the double well that would be seen by atoms trapped in the potential. On the bottom I plot the potentials relative to the energy of the $m_F = 0$ bare state so that we can compare the energy shift predicted by each approximation on an absolute scale. We can see that the higher order approximation correctly predicts the energy shift of the potential near resonance in both configurations. For the subresonant trap, however, the approximation breaks down away from the centre of the trap because the detuning becomes too large. The rotating wave approximation underestimates the energy shift across the entire potential. For the resonant configuration, the detuning remains small across the entire potential and the higher order approximation is excellent. The rotating wave approximation, however, fails when the RF field amplitude becomes too large. When the RF field is blue-detuned with respect to the trap bottom, the agreement with the higher order approximation is even better (not shown here). However for that configuration a large RF field amplitude is in general not required to make a double well potential, and so for realistic parameters the rotating wave approximation is still reasonable.

In figure 5.9 I show a plot of the important parameters of the double well (the trap separation, barrier height and trap frequency) calculated numerically (in black) and within the rotating wave approximation (in blue) and higher order approximation (in red). The calculations are made with the RF field red-detuned with respect to the trap bottom by 90 kHz just as in our experiment. The variable in the plots is the amplitude of the RF field at the centre of the trap, which is controlled by varying the RF current in the two wires. We can see that the higher order approximation correctly predicts the critical field amplitude at which the two minima begin to separate, and in fact correctly predicts the distance between the two minima across the full range of field amplitudes shown in the plot. The rotating wave approximation underestimates the critical field required to split the trap.

As a point of comparison, for the potentials calculated here the separation between

potential without needing to include the higher order terms of equation (5.8).




Figure 5.9: *Double Well Potential Parameters:* In this plot I show the important double well potential parameters calculated with the full numerics (in black) compared to the rotating wave approximation (in blue) and the higher order approximation (in red).

the trap minima in the full numerical calculation (with the red-detuned trap) is $3.9 \mu\text{m}$, the barrier height is 10 kHz and the trap frequency is $2\pi \times 1.4 \text{ kHz}$. In the rotating wave approximation, the same field produces a separation of $3.3 \mu\text{m}$, a barrier height of 5.1 kHz and a trap frequency of $2\pi \times 1.25 \text{ kHz}$. In the higher order approximation the numbers are $3.7 \mu\text{m}$, 7.5 kHz and $2\pi \times 1.37 \text{ kHz}$. Typical numbers for our interference experiments are similar, although the trap we actually use is slightly different (see below).

5.3.2 Role of Different Polarisation Components

If we compare the expression in equation (5.8) with that in equation (5.7), we can see that the semiclassical approximation introduces second order terms (in the amplitude of the RF field) due to the counter-rotating and parallel components of the RF field, as well as a first order shift in the resonance term which is due to the counter-rotating component. It is useful to get an estimate of the role of these various components of the RF field in forming the trapping potential. In figure 5.10 I plot the same full numerical potentials as in figure 5.8 (the solid lines), and the contribution of the various polarisation components to




Figure 5.10: Illustration of the role of the different polarisation components of $\mathbf{B}_{\text{RF}}$ in forming the double well potential.

the trapping potential. Again I show the potentials with the minima set to zero on the top, and on an energy scale with the zero set by the energy of the $m_F = 0$ bare state. The dotted lines plot the potentials with only the co-rotating component included in the Hamiltonian. This is the same as the rotating wave approximation. The dashed line includes the co-rotating and counter-rotating terms, and the dot-dashed line includes the co-rotating and parallel terms.

We can see that with the co-rotating and counter-rotating terms, the energy shift at the centre of the trap is correctly predicted. This is to be expected because the static field at the centre of the trap is aligned perpendicular to the RF field, so there is no parallel RF component. Away from the centre of the trap, this component becomes more important. At the position of the trap minima (2 μm away from the centre of the trap) the angle between the local static field and the RF field vector is about 69° , so the amplitude of the parallel component of the RF field is about 40% of the amplitude of the perpendicular component (which has been reduced by about 7% by the rotation). We can see that without including the parallel component of the RF field in the calculation, the approximation is increasingly worse away from the origin as the magnitude of the parallel component grows. The contribution of all the RF field components must be included in order to accurately calculate the full adiabatic potential.

Both approximations break down when the detuning becomes too large. For our trap-

ping configuration with an RF field amplitude of $\sim 0.8\text{G}$ and the frequency red-detuned by 90kHz with respect to ω_0 , the detuning is 335kHz at the trap minimum $2\mu\text{m}$ away from the centre of the trap, which is more than 50% of ω_0 , and the approximation begins to break down. In the trap made with an RF field resonant with the trap bottom, on the other hand, even though the detuning $2\mu\text{m}$ away from the trap centre is $\sim 245\text{kHz}$. The semiclassical approximation is still good in this region.

5.3.3 Effect of Current Imbalance




Figure 5.11: Illustration of the gradients of the RF field amplitude near the centre of the trapping potential. The solid lines are plotted against the y -axis and the dotted lines along the x -axis.

So far we have analysed potentials calculated with the same RF current running in each of the two trapping wires. However although the splitting axis for the double well produced by this configuration is nearly horizontal, it does not produce a balanced double well for the atoms. The reason is that the RF field amplitude is not uniform across the trapping potential because the RF field is generated by the two trapping wires. An infinitely thin wire produces a field with a gradient $\partial|B|/\partial r \propto 1/r$. The RF field we use is the vector sum of the fields from two broad wires, and the gradient of the amplitude along the y -axis is roughly proportional to $\partial|B_{RF}|/\partial y \propto 1/r^{0.72}$ in the vicinity of our trap. There is also a small gradient $\partial|B_{RF}|/\partial x$ around the centre of the trap.

In order to make a balanced double well potential we must rotate the trap away from the horizontal configuration presented here by adjusting the currents in the two wires

relative to one another. The double well potential becomes balanced when the offset introduced by this effect is balanced by the gravitational imbalance introduced by rotating the potential. It is thus important that we include gravity in our calculation of realistic trapping potentials (which I do for the calculation of the full trapping potential shown in chapter 7).

The amplitude of the RF field near the trap centre is plotted in figure 5.11(a). The solid line is a plot of $|B_{RF}|$ against y and the dashed a plot of $|B_{RF}|$ against x . The origin marks the centre of the trap in both axes. In figure 5.11(b) I plot the gradient of the RF field amplitude along each axis. In terms of energy, the shift due to the gradient along the y -axis is considerable. In figure 5.11(c) I re-plot the amplitude of the RF field in energy units near the centre of the trap. There is a shift of the order of $7 \text{ kHz } \mu\text{m}^{-1}$ along the y -axis and $1 \text{ kHz } \mu\text{m}^{-1}$ along the x -axis. This shifts the energy of the RF adiabatic potential by a similar amount. The gradient along the x -axis is due to the asymmetry of the two trapping wires. The result is that a horizontally split potential is not necessarily a balanced double well for the atoms.




Figure 5.12: *Balanced Potential:* I show here the balanced double well potential used in our interference experiments. There is $\sim 20\%$ more RF current in Z3 (the thinner of the two trapping wires) than Z4. The splitting axis of the double well potential is rotated 7.5° away from the horizontal. In this trap the well separation is $4 \mu\text{m}$ and the barrier height is 8.75 kHz . The frequency at the trap minima is $\omega_\perp \simeq 2\pi \times 1.2 \text{ kHz}$ (there is a 5% difference between the two wells).

In order to make a balanced double well, we must imbalance the RF current carried by the two wires, which will rotate the RF field vector at the trap position. A balanced double well in our configuration requires approximately 20% more RF current in the thinner of the two wires (Z3), which rotates the resulting double well potential by 7.5° away from the horizontal. The resulting potential has a separation between the two minima of $4 \mu\text{m}$

and a barrier height of 8.75 kHz. The frequency at the trap minima is $\omega_{\perp} \simeq 2\pi \times 1.2$ kHz (there is a 5% difference between the two wells). This is the potential that we use for our interference experiments. In figure 5.13 I show the effect of a 10% change in the relative RF current carried by each wire on the resulting trapping potential.




Figure 5.13: *Effect of Imbalancing the RF Currents:* I plot here the effect of changing the difference in the RF current carried by the two trapping wires by 10%. The solid line is a balanced potential with 70 mA in each wire (giving a field magnitude of ~ 0.81 G at the centre of the trap). The dotted and dashed lines are the potentials that come from adding 5 mA to one wire and subtracting the same from the other. An imbalance of ± 5 kHz between the two potential minima is the result.

5.3.4 Effect of A Relative Phase in the RF Current in the Two Wires




Figure 5.14: *Illustration of the effect of introducing a 10° phase offset between the RF current carried by the two trapping wires on the important trap parameters.*

We need to take into account one further effect in calculating our RF adiabatic potentials. If there is a small phase offset between the RF currents carried by the two wires, the resulting RF field seen by the atoms will no longer be linearly polarised but elliptically polarised. The major axis of the polarisation vector will be the same, so the splitting direction will be unchanged, but the relative contribution of the different polarisation components of the RF field will change. Since the co-rotating component couples much more

strongly to the atoms than the other components, this can result in a large change in the RF adiabatic potential.

We have discovered in our experiment that a phase offset of up to 10° can arise between the two RF currents. These can be measured and their effect calculated. The effect of a $\pm 10^\circ$ phase shift is illustrated in figure 5.14. Now that we are aware of this effect we can adjust the phase between the two wires so that the RF field is linearly polarised. However in the experiments that I describe in the following chapter, we had not yet made this correction. As a result an unknown phase of up to 10° may have been present between the RF currents in the two wires. Because of this, in the data I present in chapter 7 we cannot accurately calibrate the field produced by the wires with a given current. We can calculate a potential assuming that the polarisation is linear (i.e. that there is no phase offset in the RF current between the two wires), and know accurately how this potential scales with the RF current amplitude as a control parameter. Importantly, the phase offset discussed here does not change how the trap scales with the RF current. This means that, with our static trap parameters and an RF field red-detuned by 90kHz with respect to the trap bottom, a double well with a separation of $4\text{ }\mu\text{m}$ will have the same barrier height and trap frequency to within a few % regardless of the phase offset between the two wires. The only effect of this offset is to alter the critical field amplitude at which the double well potential begins to form (because of the increase or decrease in the relative amplitude of the co-rotating term). We treat this as an unknown offset in our calculations.

Theory of Matter-Wave Interference

6.1 Matter Wave Interference

The simplest description of the interference fringes that we see in this type of experiment is obtained by assuming that two initially well separated and independent BECs overlap in free fall after the trapping potential is (suddenly) turned off [28]. The initial separation of the two condensates is $\mathbf{d}$. In the mean field picture, the initial wave function of this system is a linear combination of the two BEC wave functions with an arbitrary but fixed relative phase $\phi = \phi_a - \phi_b$

$$\Psi(\mathbf{r}) = \Psi_a(\mathbf{r}) + e^{i\phi} \Psi_b(\mathbf{r}) \quad \text{where} \quad \int d\mathbf{r} \Psi_a^*(\mathbf{r}) \Psi_b(\mathbf{r}) \simeq 0. \quad (6.1)$$

Making the mean field approximation we write

$$\Psi(\mathbf{r}) = |\Psi(\mathbf{r})| e^{iS(\mathbf{r})}. \quad (6.2)$$

The mean phase of the condensate is related to the velocity field of the wave function via

$$\mathbf{v} = \frac{\hbar}{m} \nabla S. \quad (6.3)$$

In free fall the velocity field asymptotically approaches the classical velocity of the particle $\mathbf{v} = \mathbf{r}/t$. The phase of the wave function therefore has the asymptotic solution

$$S(\mathbf{r}, t) \rightarrow \frac{1}{2} \frac{m}{\hbar} \mathbf{k} \cdot \mathbf{r} \quad (6.4)$$

where $\mathbf{k} = \mathbf{r}/t$ is the wave vector of the expanding wave function, independent of the initial form of the wave function.

If we neglect interactions between the two condensates after we turn off the trapping potential (a good assumption if the condensates are initially well separated), the overlap-

ping wave functions in free fall have a total density which exhibits modulations of the form

$$n(\mathbf{r}, t) = n_a(\mathbf{r}, t) + n_b(\mathbf{r}, t) + 2\sqrt{n_a(\mathbf{r}, t)n_b(\mathbf{r}, t)} \cos(\phi + S_a(\mathbf{r}, t) - S_b(\mathbf{r}, t)). \quad (6.5)$$

where $n_{a,b}(\mathbf{r}, t) = |\Psi_{a,b}(\mathbf{r}, t)|^2$ are the densities of the two expanding BECs.

Using equation (6.4), the relative phase of the two condensates is

$$\begin{aligned} \phi + S_a(\mathbf{r}, t) - S_b(\mathbf{r}, t) &= \phi + \frac{1}{2} \frac{m}{\hbar} (\mathbf{k}_a - \mathbf{k}_b) \cdot \mathbf{r} \\ &= \phi + \frac{m}{\hbar t} \mathbf{d} \cdot \mathbf{r}. \end{aligned} \quad (6.6)$$

Equation (6.5) can thus be written

$$n(\mathbf{r}, t) = n_a(\mathbf{r}, t) + n_b(\mathbf{r}, t) + 2\sqrt{n_a(\mathbf{r}, t)n_b(\mathbf{r}, t)} \cos\left(\phi + \frac{m}{\hbar t} \mathbf{d} \cdot \mathbf{r}\right) \quad (6.7)$$

which is the interference pattern of two overlapping plane waves with straight line fringes orthogonal to the splitting axis with a fringe spacing

$$\Lambda = \frac{\hbar t}{m d} \quad (6.8)$$

where $d = |\mathbf{d}|$, which is equal to the de Broglie wavelength of an atom with velocity $\mathbf{v} = (\mathbf{r}/t)$. Since the position of the fringes depends upon the initial relative phase ϕ , we can read-out the relative phase of the two condensates at release by analysing the density distribution at long drop times. This is the principle of operation of the interferometer.

6.2 Two mode approximation

Consider two condensates that are completely independent, having no past connection. In the above analysis each condensate is assigned a global phase $S_{a,b}(\mathbf{r})$. The existence of such a phase can be explained as a form of spontaneous symmetry breaking in the formation of each condensate. In each run of the experiment the global phase of each condensate will be different, and the relative phase ϕ will vary randomly from $-\pi$ to $+\pi$ from shot to shot. In experiments with independently prepared BECs this is in fact what is observed. However when a single condensate is split in two and the halves allowed to interfere what we see is a (narrow) distribution of phases $\Delta\phi$ around a mean relative phase $\bar{\phi}$. In order to describe this distribution of phases, we have to employ a quantised

description of the the atom field, replacing the wave functions $\Psi_{a,b}(\mathbf{r}, t)$ with the relevant operators.

The simplest approach is a two mode theory in which each condensate represents a single (localised) mode of the system. We can define creation operators for the two modes

$$\hat{a}^\dagger, \hat{b}^\dagger = \int d\mathbf{r} \Phi_{a,b}(\mathbf{r}) \hat{\Psi}^\dagger(\mathbf{r}) \quad (6.9)$$

where $\Phi_{a,b}(\mathbf{r}) = \Psi_{a,b}(\mathbf{r})/\sqrt{N_{a,b}}$ are the single particle states of the two modes, which are orthogonal under the assumption $d \gg R$, the radius of the condensate. The field operator $\hat{\Psi}^\dagger(\mathbf{r})$ obeys the usual bosonic commutation relations and can be written in terms of density and phase operators

$$\hat{\Psi}(\mathbf{r}) = \sqrt{\hat{n}(\mathbf{r})} e^{i\hat{\phi}(\mathbf{r})} \quad (6.10)$$

in analogy to equation (6.2). Defined in this way, the particle number and phase of each mode are conjugate variables [28]

$$[\hat{n}(\mathbf{r}), \hat{\phi}(\mathbf{r}')] = i\delta(\mathbf{r} - \mathbf{r}'). \quad (6.11)$$

6.2.1 Number and Phase States

In this picture, two initially independent condensates with N_a and N_b atoms respectively are described by the number state (or Fock state)

$$|N_a, N_b\rangle = \frac{1}{\sqrt{N_a! N_b!}} (\hat{a}^\dagger)^{N_a} (\hat{b}^\dagger)^{N_b} |0\rangle \quad (6.12)$$

where $|0\rangle$ represents the vacuum. In this state each atom is localised in one of the two modes. In free fall, after a long enough fall time that the two single particle mode functions (almost) completely overlap, the expectation value of the particle density $n(\mathbf{r}, t)$ for this state is

$$\langle n(\mathbf{r}, t) \rangle = \langle N_a, N_b, t | \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}(\mathbf{r}) | N_a, N_b, t \rangle = n_a(\mathbf{r}, t) + n_b(\mathbf{r}, t) \quad (6.13)$$

which, unlike equation (6.5), has no interference term. However the ensemble average here refers to many repetitions of the same experiment and gives us no information about the outcome of a single phase measurement. To discuss a single shot experiment, we need to introduce another set of many-body states called binomial states

$$|\phi, N\rangle = \frac{1}{\sqrt{N! 2^N}} \left(a^\dagger + e^{-i\phi} b^\dagger \right)^N |0\rangle \quad (6.14)$$

in which all of the particles are in the same single particle state $(\Phi_a + e^{i\phi}\Phi_b)/\sqrt{2}$ which is a superposition of the two local mode functions. The relative phase ϕ between the two modes is fixed, as is the total number of atoms N . However the binomial state is not an eigenstate of the operators $\hat{N}_{a,b}$ for the number of atoms in each mode. The mean value of these operators acting on the binomial state is $\langle \hat{N}_{a,b} \rangle = N/2$, but the state itself corresponds to a linear combination of number states

$$|\phi, N\rangle = e^{-iN\phi/2} \sum_k C(k) e^{ik\phi} |k\rangle \quad (6.15)$$

where $k = (N_a - N_b)/2$ labels the number states with $N_{a,b} = N/2 \pm k$ and the coefficients are

$$C(k) = \frac{\sqrt{N!}}{\sqrt{2^N (N/2 + k)! (N/2 - k)!}}. \quad (6.16)$$

The probability distribution $|C(k)|^2$ is binomial, hence the name of the states.

The expectation value of the particle density $n(\mathbf{r})$ for the binomial state is

$$\begin{aligned} \langle n(\mathbf{r}, t) \rangle &= \langle \phi, N, t | \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}(\mathbf{r}) | \phi, N, t \rangle \\ &= n_a(\mathbf{r}, t) + n_b(\mathbf{r}, t) + 2\sqrt{n_a(\mathbf{r}, t)n_b(\mathbf{r}, t)} \cos\left(\phi + \frac{m}{\hbar t} \mathbf{d} \cdot \mathbf{r}\right) \end{aligned} \quad (6.17)$$

which is identical to the single shot density distribution described in equation (6.5). Again, the ensemble average here refers to multiple repetitions of the same experiment, so for this configuration we expect each run to pick out the same relative phase ϕ . In fact, for a pure binomial state we have $\langle \hat{\phi} \rangle = \phi$ and $\langle \phi, N | \hat{\phi}^2 | \phi, N \rangle - (\langle \phi, N | \hat{\phi} | \phi, N \rangle)^2 = 0$.

The number state given in equation (6.12) can be written in the basis of binomial states as

$$|N_a, N_b\rangle = \frac{1}{2\pi} \int_{-\pi}^{+\pi} d\phi |\phi, N\rangle. \quad (6.18)$$

This allows us to give a measurement-based interpretation of the observation of interference fringes in a single shot experiment with number states. The single shot can be interpreted as a projective measurement in which a single binomial state is observed with a random phase ϕ ranging from $-\pi$ to $+\pi$. Each measurement will provide a density profile given by equation (6.17), but the ensemble average over many experiments is given by equation (6.13). The details of how an interference pattern emerges in a single shot experiment is the subject of ongoing research. There are a number of papers discussing how the process of measuring each atom can induce interference patterns in a many-body

system [89] [90] [91], as well as descriptions of how interactions between the atoms can induce interference patterns [92] [93] [94] [95]. These two mechanisms may in principle be experimentally distinguishable [94].

For independent condensates the relative phase is random from experiment to experiment. For an (ideal) coherent source (such as two condensates produced from the same source) the relative phase will be the same in every experiment. In between the relative phases observed in a set of experiments will have some distribution about a mean relative phase, where the width of the distribution is a measure of the coherence of the source. The experimental challenge is to be able to produce a coherent source by splitting a single condensate in two. We do this by ramping up the barrier of the double well potential slowly enough that there is no excitation of bulk oscillation modes, but fast enough that tunneling between the two modes of the well plays no appreciable role.

A realistic splitting process can be approximately described in two steps [96]: an adiabatic splitting period in which the change of the potential is slow compared to the relevant energy scale of the system, and a non-adiabatic period when the adiabatic criterion breaks down. In the first period, we can think of the system as a single BEC spread across a double well potential with a barrier smaller than the chemical potential of the condensate. The condensate wavefunction adiabatically follows the deformation of the potential up to the point when the barrier increases above the BEC chemical potential. After this, tunneling between the two modes is very rapidly switched off, which is a realistic experimental approximation because the tunneling term is exponentially dependent on the overlap of the two mode functions.

6.2.2 Correlation Functions

The first order correlation function $g^{(1)}(\tilde{\mathbf{r}}, \tilde{\mathbf{r}}')$, where $\tilde{\mathbf{r}}$ and $\tilde{\mathbf{r}}'$ refer to the position of the atoms in the double well potential, describes the degree of spatial coherence between the two modes [97]

$$\begin{aligned}
 g^{(1)}(\tilde{\mathbf{r}}, \tilde{\mathbf{r}}') &= \frac{|\langle \hat{\Psi}^\dagger(\tilde{\mathbf{r}}') \hat{\Psi}(\tilde{\mathbf{r}}) \rangle|}{\sqrt{\langle \hat{\Psi}^\dagger(\tilde{\mathbf{r}}) \hat{\Psi}(\tilde{\mathbf{r}}) \rangle \langle \hat{\Psi}^\dagger(\tilde{\mathbf{r}}') \hat{\Psi}(\tilde{\mathbf{r}}') \rangle}} \\
 &= \frac{\langle \hat{a}^\dagger \hat{b} + \hat{b}^\dagger \hat{a} \rangle}{N} \\
 &= \alpha
 \end{aligned} \tag{6.19}$$

and has a direct experimental meaning: it corresponds to the mean fringe contrast of the ensemble-average interference pattern [9]. If we assume that the two modes have the

same spatial distribution after a long drop time, then

$$\langle n(\mathbf{r}, t) \rangle = 2n(\mathbf{r}, t) \left(1 + \alpha \cos\left(\phi + \frac{m}{\hbar t} \mathbf{d} \cdot \mathbf{r}\right) \right). \quad (6.20)$$

We can use this expression to extract information about the initial state of the system: $\alpha = 0$ for the number state and 1 for the binomial state, and in between it is a direct measure of the spatial coherence between the two modes.

It is also useful to analyse higher order correlation functions. The interference that we observe between the two modes of the condensate is a many-body effect. In a double slit interferometer with single atoms each atom interferes with itself, with the relative phase between the two possible paths that it travels set by the geometry of the problem. Two independent single atom sources do not interfere. In a many body system, interference terms can show up in higher order correlation functions, the classic example being the Hanbury-Brown and Twiss experiment [98]. We can similarly look for higher order interference terms to explain the appearance of interference fringes in a single experiment with the number state described by equation (6.12). Information about the exact many-body state of the initial system in the double well potential will be encoded in these higher order correlation functions, and in principle we can extract information about these states by studying the statistics of these correlations [99] [100].

The two-particle correlation function is the probability of detecting a particle at a position $\mathbf{r}'$ given that a particle has been detected at position $\mathbf{r}$. If we assume that at long drop times the wave functions describing the two single particle modes are identical up to an initial phase factor, then for the number state the two-particle correlation function is

$$\begin{aligned} \langle N_a, N_b, t | \hat{\Psi}^\dagger(\mathbf{r}) \hat{\Psi}^\dagger(\mathbf{r}') \hat{\Psi}(\mathbf{r}') \hat{\Psi}(\mathbf{r}) | N_a, N_b, t \rangle = \\ N(N-1) |\Phi_0(\mathbf{r}, t)|^4 + 2N_a N_b |\Phi_0(\mathbf{r}, t)|^4 \cos\left(\frac{m}{\hbar t} \mathbf{d} \cdot (\mathbf{r} - \mathbf{r}')\right). \end{aligned} \quad (6.21)$$

What this equation says is that if an atom is detected at a position $\mathbf{r}$, then the probability of finding another atom at a position $\mathbf{r}'$ has a spatial modulation just like the interference pattern of equation (6.7). Again the initial random relative phase ϕ does not appear in the ensemble average, because we cannot say where the peak of a single shot interference pattern will form, only that if there is a maximum at a point $\mathbf{r}$, then there will be another maximum at a distance $2\pi(\hbar t / md)$ away from this point.

6.3 Phase Spreading

We want to use the measurement of the relative phase between the two modes of the condensate as the read out scheme of our interferometer. The resolution of the interferometer will then be limited by the width of the phase distribution of the many body state used to make the measurement. This suggests that the coherent state described in equation (6.14) is the best state to use in the interferometer. However the usefulness of this state is limited by a process called phase spreading.

In the experiment we split a single condensate into two modes that are well separated, so that there is no tunneling between the modes. The initial state can be expanded in terms of number states as described in equation (6.15). Each state $|k\rangle$ in the expansion will evolve with a phase factor $e^{-iE(k)t/\hbar}$, where the energy of the state depends on the number of atoms in the each mode. We will assume that the the initial state has a narrow distribution of states $|k\rangle$ around $k = 0$. The time dependence of the many-body state is then

$$|\Psi(t)\rangle = e^{-iE_0 t/\hbar} \sum_k C(k) \exp\left(\frac{-iE_C k^2 t}{2\hbar}\right) |k\rangle. \quad (6.22)$$

where $E_C = 2\partial\mu_i/\partial N_i$ is the interaction energy (or on-site energy) of the modes. Since each component evolves in time with its own phase, the width of the initial phase distribution will increase. This is referred to as phase spreading, and is a consequence of the interaction between atoms in each mode (since for an ideal gas $E_C = 0$). The time evolution of the phase fluctuations is

$$\langle\Delta\phi^2\rangle_t = \langle\Delta\phi^2\rangle_0 + \langle\Delta k^2\rangle \left(\frac{E_C t}{2\hbar}\right)^2. \quad (6.23)$$

where $\langle\Delta k^2\rangle$ is the fluctuation in relative atom number, which remains fixed since there is no tunneling between the modes. Thus a state with a well defined initial phase ($\Delta\phi_0 \ll 1$) will become dephased ($\Delta\phi_t \simeq 1$) in a time [28] [90] [101] [102] [103]

$$t_D = \frac{\hbar}{E_C} \frac{2}{\Delta k}. \quad (6.24)$$

In the Thomas Fermi approximation $E_C = 2\mu/5N$ and if we assume Poissonian number fluctuations $\Delta k \simeq \sqrt{N}$ around $k = 0$, the dephasing time becomes

$$t_D = \frac{5\hbar\sqrt{N}}{2\mu} \quad (6.25)$$

Since $\mu \propto N^{2/5}$, the dephasing time is weakly dependent on the atom number ($t_D \propto N^{1/10}$),

and determined largely by the geometric mean trapping frequency $\bar{\omega}$. Typical experimental parameters in the double well potential are $\bar{\omega} = 2\pi \times 345 \text{ Hz}$ ($\omega_{\perp} = 2\pi \times 1.5 \text{ kHz}$ and $\omega_z = 2\pi \times 18 \text{ Hz}$), $\mu = 2\pi \times 2.2 \text{ kHz}$ and $N = 1.5 \times 10^4$ atoms, giving $t_D \simeq 16 \text{ ms}$.

6.4 Longitudinal Phase Fluctuations

As we have seen in section 4.3, we make quasi-condensates in the 3D-1D crossover range with phase fluctuations along the length of the BEC. A BEC with a phase coherence length $l_{\phi} < L$ (corresponding to a temperature $T > T_{\phi}$) will exhibit phase fluctuations along its length. As a first approximation, we can think of the BEC as breaking into L/l_{ϕ} phase domains, each with a random phase relative to each other. Consequently when we split the BEC transverse to the longitudinal axis, we split L/l_{ϕ} pairs of condensates, each with a relative phase ϕ_z . The integrated interference pattern will be the incoherent sum of the interference patterns of each of these pairs. Although this is a crude model, it leads to a reasonable approximation for the loss of contrast in the interference pattern from a phase fluctuating condensate [104]. The model has been developed by Demler and co-workers - see [105] [106] [107] [108] and the review paper [109].

The condensates are aligned along the z -axis and have a length L , a phase coherence length l_{ϕ} (defined in section 2.3 of chapter 2), and a separation d along the x -axis. We will ignore the radial modes of the elongated condensate. The interference pattern of the two overlapping condensates arises from the density-density correlation function in free fall

$$\langle n_a(\mathbf{r}, t) n_b(\mathbf{r}, t) \rangle = \langle |A_k|^2 \rangle \left(e^{i \frac{m}{\hbar} \mathbf{d} \cdot \mathbf{r}} + c.c. \right) \quad (6.26)$$

where $|A_k|^2$ is the mean square amplitude of the fluctuations with wavevector $k = md/\hbar t$. We can write the amplitude in terms of the single particle correlation function of the condensates (assuming that the two condensates are identical)

$$\langle |A_k|^2 \rangle = L \int dz \langle \hat{a}^{\dagger} \hat{a} \rangle. \quad (6.27)$$

From equation (2.61) and equation (2.62) we can see that $\langle \hat{a}^{\dagger} \hat{a} \rangle$ decays exponentially with distance along the z -axis with a correlation length l_{ϕ} . The mean square amplitude $|A_k|$ is then proportional to $\sqrt{l_{\phi} L}$. This system is equivalent to L/l_{ϕ} pairs of independent condensates, as in the simple description above. The total amplitude $|A_k|$ is the sum of L/l_{ϕ} independent vectors with length l_{ϕ} relative phases ϕ_z . The interference contrast is proportional to $|A_k|/L$, so the ratio of the fringe amplitude to the background signal is $\sqrt{l_{\phi}/L}$. If the two condensates are phase coherent along their entire length then

$\langle \hat{\psi}^\dagger(z) \hat{\psi}(z') \rangle$ is constant and $|A_k| \propto L$. In this case the vectors all have the same phase. We can study the noise statistics of the system by analysing shot-to-shot variations of the fringe amplitude.

The axial phase fluctuations also affect phase spreading time of the split condensates. In a phase fluctuating condensate, the long-range order in the initial relative phase decays exponentially along the length of the cloud. This is a random process driven by thermal excitation, so it will lead to fluctuation in the *relative* phase along the length of the condensate once the two modes are separated. The decay of longitudinal phase fluctuations in 1D Bose gases is the subject of ongoing theoretical [105] [106] [110] and experimental research [12] [11].

Matter Wave Interference

7.1 Introduction

In this chapter I present data from our first interference experiment. In section 7.2 I show how we create a balanced double well potential for our atoms. In section 7.3 I discuss how we analyse the data from the interference patterns that we observe. The most important data is presented in section 7.4 where I show that we are able to split a single condensate coherently in the double well potential, which is the pre-requisite for making a working atom chip based interferometer. This data is the core result presented in this thesis. To date, there are only a few groups around the world that have been able to successfully do this. In section 7.5 I present a few observations about the interference experiments we have run. This data is preliminary and mainly qualitative. Detailed quantitative analysis of how to improve and control the splitting process and make a working interferometer with this set-up is currently underway and will be the topic of future studies.

7.2 Implementing the Double Well Potential

The double well potential is formed using two trapping wires as described in section 5.3 of chapter 5. The potential is illustrated in figure 7.1. The DC currents in the circuit are controlled by home-built current drivers as described in section 3.6 of chapter 3. We run 1 A in each of the two trapping wires, which are connected in series to a single current driver. The RF current is connected to the wires as illustrated in figure 3.10 of the same chapter. The RF circuit is designed so that there is a π phase shift between the RF current in the wires, which is required to generate a linearly polarised RF field aligned closely to the y -axis of the trap. The exact orientation of the RF field is adjusted by altering the




Figure 7.1: Double Well Potential: The RF adiabatic potential that we use in our interference experiments is illustrated here. In figure 7.1(a) I highlight the chip wires (Z3 and Z4) that are used to generate the trapping fields. A cross-section of the geometry of the chip wires in relation to the trap is shown in figure 7.1(b). The two wires are $50 \mu\text{m}$ and $100 \mu\text{m}$ wide respectively. Both carry 1 A DC current (which runs into the page in figure 7.1(b)). The RF current in each wire is adjusted to that the resulting double well potential is balanced, which requires approximately 20% more RF current in Z3 than in Z4. A cross section through the resulting RF adiabatic potential is shown in figure 7.1(d), and a slice through the centre from trap minimum to trap minimum is shown in figure 7.1(c). The splitting axis of the double well potential is rotated 7.5° away from the horizontal. The well separation is $4 \mu\text{m}$ and the barrier height is 8.75 kHz. The frequency at the trap minima is $\omega_{\perp} \simeq 2\pi \times 1.2 \text{ kHz}$.

RF current in the two wires relative to one another until a balanced potential is formed. Recall from the discussion in chapter 5 that in order to make it balanced we must rotate the splitting axis of the double well potential away from the horizontal by about 7.5° . Assuming that the field is linearly polarised, this requires an RF current of 60 mA in Z3 and 73 mA in Z4, which produces an RF field amplitude of 0.78 G at the centre of the trapping potential.¹ To calculate the trapping potential, the finite width of the wires has been taken into account for both the static and RF fields, the RF adiabatic potential calculated with the full numerics described in chapter 5 and gravity included in the calculation.

For all the experimental data presented in section 7.3 and section 7.4, we use the following experimental procedure. We deform the single well into the double well by linearly ramping the amplitudes of the RF current in the two wires to their final values over 20 ms. The RF frequency is kept constant at 540 kHz, which is 90 kHz red-detuned away from resonance at the minimum of the static magnetic trap. During the first part of the amplitude ramp the single well potential is not significantly deformed. Calculations of the trapping potential indicate that it does not exhibit two minima until the currents have reached 60% of their final values. The two modes of the condensate do not begin to separate until the barrier height is larger than the chemical potential of the condensate. This point is difficult to specify since the condensate deforms along with the trap. In the single well potential the BEC has a chemical potential of $\mu = h \times 2\pi \times 3 \text{ kHz}$. In the final double well potential each mode has a chemical potential $\mu_i = h \times 2\pi \times 2.2 \text{ kHz}$. Using these numbers we estimate that the two modes begin to separate when the RF currents have reached between 75% and 80% of their final values. We cannot yet estimate the point at which the overlap between the two modes disappears and the deformation of the double well potential is no longer adiabatic. These numbers suggest that the condensate splits only at the last few ms of the amplitude ramp.

The final trap has a separation of $4 \mu\text{m}$ between the two trap minima, and a barrier height of $\sim 8.75 \text{ kHz}$. In this trap the two modes of the condensate are well separated (with no overlap and thus no tunneling) and have begun to evolve separately. At the end of the amplitude ramp we immediately release the atoms from the trapping potential and allow the two modes to interfere as they overlap in free fall. Images were captured after 14.2 ms of free fall.




Figure 7.2: *In-trap Images of Separated BECs:* Here I show in-trap absorption images of BECs separated by $\sim 50 \mu\text{m}$ in the RF adiabatic potentials using various configurations of RF currents in the two trapping wires to control the polarisation of the RF field. After linearly ramping the amplitude of the RF current to its final value, we then ramp the frequency of the field up to 3.5 MHz to make these traps. The image on the top left of the figure shows the trap split in the configuration that we employ to make a balanced double well potential for our interference experiments. We can see a small rotation of the splitting axis away from the horizontal in this image.

In-Trap Images of a Split Condensate

It is also possible to increase the separation between the two potential wells by sweeping the frequency of the RF field after the amplitude ramp is complete. This allows us to separate the wells far enough that we can see the two separate condensates by imaging them while still trapped in the double well potential. In figure 7.2 I show images in which the condensate modes have been separated by $\sim 50 \mu\text{m}$ using various different configurations of the RF field so that the axis along which the double well forms is changed. In all cases the polarisation of the RF field is linear, and the amplitude of the field at the centre of the trap is $\sim 0.8 \text{ G}$. The frequency of the RF field was swept linearly to 3.5 MHz at the end of the amplitude ramp, and the images taken with the atoms still trapped in the RF adiabatic potentials. We can see here that by controlling the relative amplitude of the RF current in the two wires, we can split the condensate along any axis that we wish. The image on the top left of the figure shows the trap split in the configuration that we employ to make a balanced double well potential for our interference experiments. The small rotation of the splitting axis away from the horizontal is apparent in this image.

¹Recall from the discussion at the end of chapter 5 that there may be a phase offset between the RF currents in the two wires, which would introduce an offset in the current needed to make the trap. However the other parameters of the trap are fixed by the known trap separation since they scale the same way regardless of this offset.

Balancing the Two Modes




Figure 7.3: *Balancing the Double Well Potential:* In order to check that the splitting process is balanced (so that an equal mean number of atoms ends up in each condensate mode) we first form the double well potential, then suddenly switch off the RF field, keeping the static magnetic trap on. This projects the atoms onto the bare m_F states. The $m_F = -1$ and $m_F = -2$ atoms are rapidly expelled from the trap and are not detected. The $m_F = 0$ atoms fall freely and can be seen at the bottom of the array of clouds in the figure. The $m_F = 1$ and $m_F = 2$ atoms are accelerated towards the centre of the trap because the projection occurs $2\mu\text{m}$ away from the centre of the static magnetic trap. $400\mu\text{s}$ later the static trap is also switched off. The clouds ballistically cross and separate far enough in free fall that they can be optically resolved. We then count the number of atoms in the various clouds and adjust the splitting potential until there are equal numbers of atoms in the left and right modes.

We would like to split our condensate equally into the two modes in order to maximise the visibility of the interference fringes we observe in our experiment. We empirically check that we split the condensate equally using a procedure described in Thorsten Schumm's thesis [23]. We cannot optically resolve the two modes of the BEC in the double well potential that we use for our interference experiments. In order to monitor that the splitting is balanced we therefore first split the BEC in the RF adiabatic potential, then keeping the static magnetic trap turned on, we rapidly switch off the RF field. This projects the atoms onto the m_F states of the bare trapping potential. Atoms projected onto $m_F = -1$ and $m_F = -2$ atoms are rapidly expelled from the trap and are not detected, while atoms projected onto $m_F = 0$ fall out of the trap. That leaves the atoms projected onto the low field seeking $m_F = 1$ and $m_F = 2$ states. Since the projection occurs $2\mu\text{m}$ away from the centre of the static magnetic trap, these atoms are accelerated towards the centre of the trap. The two different m_F states feel a differing acceleration during this period. The static magnetic trap is also rapidly turned off $\sim 400\mu\text{s}$ later. The clouds ballistically cross and move apart in free fall sufficiently far that we can optically resolve the difference components. A typical image from this balancing procedure is shown in figure 7.3. We can see atoms in both the $m_F = 1$ and $m_F = 2$ states originating from each of the two condensate modes. By counting the number of atoms that come from each

mode (ignoring the atoms in the $m_F = 0$ state that come from both modes) we can check whether the splitting was balanced. We then adjust the splitting parameters as necessary to get the same number of atoms in each mode.

Once the two modes are balanced using this procedure, we return to a small splitting of $3 - 4 \mu\text{m}$ and are able to observe interference fringes of the overlapping condensate modes with good visibility. To run the interference experiment we then simply release the condensate from the double well potential, ensuring that all the trapping fields are turned off simultaneously. We allow the condensate modes to expand and overlap in free fall and image the resulting cloud after 12-15 ms time of flight. The fall time is chosen so that the wavelength of the interference fringes is as long as possible, but there is also a high enough optical density to make high contrast images of the interference pattern. The rapid radial expansion of the condensate after the trap is released means that at longer drop times than ~ 15 ms the density of the cloud is too low to make good absorption images of the interference pattern.

7.3 Data Analysis

7.3.1 Extracting Images




Figure 7.4: *Interference.*

Data from a typical interference experiment is shown in figure 7.4(a). The BEC interference pattern is superimposed on a noisy background of fringes caused by diffraction of the imaging laser and etalon effects in the imaging set up. These fringes are clearly distinguishable from the interference pattern of the BEC. In the Fourier spectrum of a slice through the centre of the BEC (see figure 7.6(c) below), we can clearly separate the imaging noise from the peak at the wavelength of the BEC interference pattern. This

allows us to filter the image without losing information about the interference pattern. A typical filtered image is shown in figure 7.4(b).





Figure 7.5: *Fourier filtering of an interference pattern. The unfiltered line density is shown at the top left, and the filtered line density at the top right. The position and visibility of the interference pattern is not affected by the filtering. The gaussian envelope is shown at the bottom left with both the imaging noise and density modulation removed. The density modulation is shown at the bottom right.*


The filter cuts out frequency components in the region between the DC and fringe frequency peaks that can be clearly seen in the Fourier spectrum shown in figure 7.6(c). Figure 7.5 illustrates the effect of the filtering. On the top left of the figure the linear number density extracted from a slice through the centre of the interference pattern is plotted (see below). The filtered line density that we use for our fitting routine is plotted on the top right. The gaussian envelop can be extracted from the image by filtering all the high frequency components, which leaves the line density shown in the bottom left plot. Alternatively, if we filter out the gaussian envelope and the background noise, we can clearly see the modulation period of the interference pattern, shown in the plot on the bottom right.

7.3.2 Density Modulation


From an image we extract information about the visibility, wavelength and phase of the interference pattern. The analysis procedure is illustrated in figure 7.6. We first extract a line density from the filtered image by integrating over a slice through the centre of the interference pattern. The integration is over a region $y_0 \pm \sigma_y/2$, where the centre y_0 and



(a) Interference Pattern



(b) Density Modulation



(c) Fourier Spectrum

Figure 7.6: *Fringe Fitting:* We first extract a line density from the filtered image by integrating over a slice through the centre of the interference pattern, indicated by the red lines in figure 7.6(a). We fit a modulated gaussian (equation (7.1)) to the resulting line density (figure 7.6(b)) to extract the fringe amplitude d , wavelength Λ and phase ϕ . The same information can also be extracted from a Fourier spectrum of the interference pattern, illustrated in figure 7.6(c). The interference pattern spectrum in blue, and the spectrum from a background slice integrated over a region of the same size immediately above the interference pattern in red, which is indicative of the imaging noise.

width σ_y of the slices are determined by a gaussian fit to the density distribution in the direction perpendicular to the interference fringes. A typical region is outlined by the red lines in figure 7.6(a).

We then fit a modulated gaussian to the resulting line density

$$f(x) = A \exp \left[-\frac{(x - x_0)^2}{2\sigma^2} \right] \left(1 + d \cos \left[\phi + \frac{2\pi x}{\Lambda} \right] \right). \quad (7.1)$$

The fit returns the wavelength Λ and visibility d of the density modulation, where d is measured as a fraction of the amplitude A of the gaussian envelope. The visibility of the integrated line density is the same as the fringe amplitude defined in equation (6.20) up to a normalisation factor given by the total number of atoms in the region over which we integrate to get the line density.

We cannot fix the absolute phase of an interference pattern without a well defined

reference for zero relative phase between the two modes of the condensate. However, since any offset is the same for all interferograms in a given data set, we can still extract useful information about the distribution of relative phases without knowing the absolute phase between the components. A possible procedure for establishing a phase reference empirically (for a given data set) is outlined below.

A typical line density and fit are shown in figure 7.6(b). Between 5 and 10 modulations are visible in a typical line density. Our imaging resolution limits the fringe visibility to a maximum possible value of about 40%. We typically observe a visibility between 20-25%, which is sufficient to accurately fix the wavelength and phase of the density modulation. We can typically find the wavelength of each interference pattern to within 1%, the visibility to within 5% and the phase to within 20mRad.




Figure 7.7: *Effect of filtering the image. The unfiltered line density is shown in red, the filtered line density in blue. The filtering slightly improves the visibility without shift the position of the fringes. The fit to the gaussian envelope is significantly improved by removing long wavelength imaging noise.*

This fitting procedure is much more robust when we filter the image before extracting the line density. The reason is that long wavelength imaging noise interferes with fitting the gaussian envelope. In figure 7.7 I show the filtered and unfiltered data from a typical slice. In the filtered data the fringe visibility is slightly improved, but the position of the fringes (and thus the phase extracted) is unchanged.

We can extract the same information from a Fourier spectrum of the interference pattern. A typical spectrum is shown in figure 7.6(c). The unfiltered spectrum is shown in blue, and the filtered spectrum of the image in figure 7.4(b) is shown overlapping this in red. We extract the visibility, wavelength and phase from the unfiltered Fourier spectrum directly, which we can compare to the same quantities extracted from the fitting routine described above.

The wavelength of the interference pattern corresponds to the wavenumber of the peak indicated in figure 7.6(c). The absolute value of the Fourier spectrum at this wavenumber

is identical to the fringe amplitude defined in equation (6.20). The phase of the interference pattern is the phase of the (complex) Fourier coefficient at this wavenumber.

The mathematics behind extracting information from the Fourier spectrum is straightforward and readily extended to two dimensions [111] [112] [113] [92] [114]. We assume that the interference pattern has the form

$$I(\mathbf{r}) = a(\mathbf{r}) (1 + d \cos(\phi + \mathbf{k}_0 \cdot \mathbf{r})) \quad (7.2)$$

where $\mathbf{k}_0$ is referred to as the carrier frequency (in our case the relative velocity $\mathbf{k}_0 = \hbar \mathbf{v} / m d$ between the two overlapping wave functions), ϕ is the relative phase we want to know, and the envelope $a(\mathbf{r})$ is a gaussian. The Fourier transform of equation (7.2) is given by

$$F(\mathbf{k}) = A(\mathbf{k}) + C(\mathbf{k} - \mathbf{k}_0) + C^*(\mathbf{k} + \mathbf{k}_0) \quad (7.3)$$

where $C(\mathbf{k})$ is a gaussian envelope in $\mathbf{k}$ -space. The $C(\mathbf{k})$ components are displaced by $\pm \mathbf{k}_0$ around the DC component, and so can be clearly distinguished in the Fourier spectrum, as we see in figure 7.6(c).

The Fourier transform method is significantly faster to implement than fitting a modulated gaussian to the line density and involves fewer free parameters. However some care must be taken in identifying the carrier frequency $\mathbf{k}_0$, particularly in an experiment in which the fringe spacing is deliberately altered. In general the peak of the gaussian envelope $C(\mathbf{k})$ will not coincide precisely with a wavevector represented by a pixel in the discrete Fourier transform (DFT). This introduces an error in assigning a phase based on the peak of the DFT spectrum. If the phase in the DFT winds sufficiently slowly from pixel to pixel that we can unwrap the phase across the spectrum of interest we can in principle correct for this error. In general this is not possible with our data. However when we compare the relative phases extracted from a set of interferograms that all have the same wavelength (and thus carrier frequency), the error introduced by the discretisation of the DFT is the same for each data point. We can thus confidently use the DFT to analyse the phase distribution of a data set provided that the wavelength stays fixed. An example is provided in section 7.4 below comparing the two methods of data analysis. In experiments where the wavelength is changing we prefer to fit the modulated gaussian to the line density.

If the background noise has a significant contribution to the Fourier spectrum at the carrier frequency $\mathbf{k}_0$ this will also introduce an error in the estimation of the phase and visibility of the interference pattern. We have analysed the spectrum of background imaging noise in the region where we image the interference pattern. A typical spectrum is shown in figure 7.6(c). The amplitude of the background spectrum at the carrier frequency is typ-

ically $\sim 5\%$ of the peak corresponding to the interference fringes. The maximum phase error this component can introduce is $\leq 5\%$ (if the two complex Fourier components are perpendicular to one another in the complex plane, the phase error is maximal).

7.3.3 Circular Statistics

Since the relative phase ϕ that we measure has a periodicity of 2π , we cannot use standard definitions of the mean and standard deviation to analyse our data. Instead we use the equivalent definitions developed for analysing directional (or circular) statistics [115] [116] [117] [118].

We measure a set of N phases ϕ_k each represented as a point on a unit circle. From this set of observations we can calculate the phasor that represents the sample mean

$$r e^{i\bar{\phi}} = \frac{1}{N} \sum_{k=1}^N e^{i\phi_k}. \quad (7.4)$$

where

$$r = \sqrt{\bar{c}^2 + \bar{s}^2} \quad \text{and} \quad \bar{\phi} = \begin{cases} \arctan\left(\frac{\bar{s}}{\bar{c}}\right) & \text{if } \bar{c} \geq 0 \\ \arctan\left(\frac{\bar{s}}{\bar{c}}\right) + \pi & \text{if } \bar{c} < 0. \end{cases} \quad (7.5)$$

with

$$\bar{s} = \frac{1}{N} \sum_k \sin(\phi_k) \quad \text{and} \quad \bar{c} = \frac{1}{N} \sum_k \cos(\phi_k). \quad (7.6)$$

The mean length r is a measure of the concentration of points around $\bar{\phi}$. If the points are tightly clustered around $\bar{\phi}$ then $r \simeq 1$, and if they are widely dispersed $r \simeq 0$. We can calculate r from the mean of the sines and cosines. We can then define the sample circular standard deviation in terms of r

$$\sigma = \sqrt{-2\ln(r)}, \quad (7.7)$$

where $\sigma \in [0, \infty)$.

Circular Distributions

We would then like to draw inferences from the data. Hypothesis testing requires a model of the distribution of phases from which we took our measurements. The equivalent of

the Gaussian (normal) distribution in circular statistics is the von Mises distribution¹.

$$f(\phi|\mu, \kappa) = \frac{e^{\kappa \cos(\phi - \mu)}}{2\pi I_0(\kappa)}, \quad (7.8)$$

where $I_0(\kappa)$ is the zeroth order modified Bessel function of the first kind and $\phi \in [0, 2\pi]$. The parameter μ is the mean of the distribution, and κ concentration of points around μ . In the limit of large κ , the von Mises distribution becomes the usual normal distribution with $\sigma^2 \simeq 1/\kappa$, and in the limit of small κ it becomes a uniform distribution:

$$\lim_{\kappa \rightarrow \infty} f(\phi|\mu, \kappa) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\phi - \mu)^2}{2\sigma^2}\right) \quad (7.9)$$

$$\lim_{\kappa \rightarrow 0} f(\phi|\mu, \kappa) = U(\phi). \quad (7.10)$$

Unless the number of sample points is very large, the von Mises and wrapped normal distributions are in practice indistinguishable and either may be chosen for the purposes of statistical hypothesis testing. The von Mises distribution is typically used because it is easier to give simple expressions for confidence intervals and hypothesis tests.

Hypothesis Testing

For a given set of data we would like to know if the mean phase we measure is statistically significant. One way to do this is to test the null hypothesis that we have drawn our data from a uniformly distributed set of phases. The standard test of uniformity is called the Rayleigh test. We calculate the probability P that a set of N data points drawn from a uniform distribution produces a mean vector longer than r [118]

$$P = \exp\left[\sqrt{1 + 4N + 4N^2(1 - r^2)} - (1 + 2N)\right]. \quad (7.11)$$

For a data set of N sample points, we calculate the length of the mean phasor $\bar{r}$ from equation (7.5) and compare this to a radius r calculated from equation (7.11) for a chosen confidence interval. If $\bar{r} > r$, then we can reject the null hypothesis that the underlying distribution was uniformly distributed with a confidence $1 - P$.

It is also useful to be able to calculate confidence intervals $\bar{\phi} \pm \delta\bar{\phi}_\alpha$ around the mea-

¹A commonly used alternative is the wrapped normal distribution

$$f(\phi|\mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{(\phi - \mu - 2\pi k)^2}{2\sigma^2}\right),$$

where $\sigma = \sqrt{-2\ln(r)}$ is the circular standard deviation and μ the mean.

sured mean phase $\bar{\phi}$. In general, this requires making an assumption about the underlying distribution of phases (although model-independent techniques exist). Relatively simple expressions for confidence intervals exist assuming that the phases follow the von Mises distribution. When the data are tightly clustered the von Mises distribution can be approximated as a normal distribution with mean μ and variance $1/\kappa$. For less clustered data we can use instead the following expressions [116]

$$\delta\bar{\phi}_\alpha \simeq \begin{cases} \arccos \left[\frac{1}{r} \sqrt{\left(1 - (1 - r^2)e^{Z_\alpha^2/N}\right)} \right] & \text{if } r \geq 0.9 \\ \arccos \left(\frac{1}{r} \sqrt{\frac{2(2Nr^2 - Z_\alpha^2)}{4N - Z_\alpha^2}} \right) & \text{if } r < 0.9 \text{ and } r > \sqrt{\frac{Z_\alpha^2}{2N}} \end{cases} \quad (7.12)$$

where Z_α is the quantile of the χ_1^2 distribution corresponding to α . These approximations are derived under the assumption that κ is unknown but small and are reasonable for $N > 8$. For our typical data sets we can estimate the mean phase to within 5% (the 95% confidence interval assuming an underlying von Mises distribution), which is significantly smaller than the width of the phase distribution itself.

Various alternative statistical tests based on the von Mises distribution also exist, including tests for the concentration parameter κ of the fit, and tests for distinguishing the phase difference of two data sets [115]. For this reason the von Mises distribution is typically used when making statistical inferences from circular data.

7.3.4 Finite Imaging Resolution




Figure 7.8: Reduced fringe visibility due to the finite resolution of the imaging system. An interference pattern with a wavelength of $18\mu\text{m}$ has its visibility reduced by about 40% when we take the resolution of the imaging system into account.

The visibility of the fringes we observe is considerably reduced by the finite resolution

of the imaging system. We can account for this by convolving the density distribution with a Gaussian kernel¹

$$K(x) = \frac{1}{\sqrt{\pi}s} \exp\left(\frac{-x^2}{2s^2}\right) \quad (7.13)$$

where s is the imaging resolution. An ideal sinusoid convolved with a period $\Lambda = 2\pi/k$ convolved with $K(x)$ has a reduced visibility $\mathcal{V} = \exp(-k^2s^2/2)$ [119]. With our experimental parameters, $\Lambda \simeq 20\text{ }\mu\text{m}$ and our imaging resolution is $s \simeq 7\text{ }\mu\text{m}$, the expected visibility of an interference pattern with 100% visibility is about 0.4 - see figure 7.8. Note also that a misalignment of the imaging axis with respect to the fringes will also degrade the maximum possible fringe visibility, but this is not a significant effect in the experiments described here.

7.4 Coherent Splitting

One of the early experiments performed with a Bose-Einstein condensate involved the observation of interference fringes when the two separate components of a condensate split in a double well potential were allowed to overlap in free fall [1]. From shot to shot, the relative phase of the two modes was observed to be random. It was not until much later that a reproducible relative phase from the two components of a single split condensate was observed, originally with optical potentials [2] [3] and then subsequently with RF adiabatic potentials on an atom chip [5]. Achieving this step, which is a prerequisite for building an atom chip based BEC interferometer, is the main result reported in this thesis.

Phase Distribution

To test for phase coherence we repeat the experiment described above in section 7.2 many times, and analyse the resulting interference patterns. The following data set was drawn from 103 repetitions of the same experimental sequence.

The polar plot in figure 7.9(a) shows the phases and visibilities extracted from each of the interference patterns. The distribution of phases is clearly not uniform. The phases are peaked around a mean phase ϕ_0 with a width of $\Delta\phi = 40^\circ$. The Rayleigh test yields a negligible probability that the underlying distribution of phases was uniform ($\sim 2 \times 10^{-33}$). The average fringe visibility of the data set was $d = 0.21 \pm 0.03$. The wavelength of the density modulation was 5.15 ± 0.1 pixels, which corresponds to $17.5 \pm 0.3\text{ }\mu\text{m}$. The

¹The Gaussian kernel is a reasonable approximation to the central peak of an Airy diffraction pattern and far more tractable for analytical calculations.




Figure 7.9: Distribution of relative phases from 103 repetitions of the same experiment. Raw data above, data corrected for a slow linear phase drift below. The fringe visibility is plotted on the radial axis. In figure 7.9(b) the phase spread of the raw data is $\Delta\phi = 45^\circ$. In figure 7.9(e) the phase spread of the corrected data it is $\Delta\phi = 12^\circ$. In both cases it is clearly distinguishable from a random phase. The histograms on the right hand side are displayed with a fit to the von Mises distribution with a 95% confidence interval. The measured phase spread is derived according to equation (7.7). The concentration parameter κ of the fit to the von Mises distribution gives widths $\sigma = 1/\sqrt{\kappa}$ of $\sigma = 42^\circ$ and $\sigma = 14.5^\circ$ for the two histograms.

errors here correspond to the standard deviation of the distribution over the data set, not the standard deviation of the mean.

The data was collected over a period of just under 2.5 hours. In figure 7.10(a) I replot the same data colour coded in five sets, each of which represents about half an hour of data. There is a clear shift in the mean phase over time. In figure 7.10(b) the phase data is plotted against time. The drift of the mean phase over time is approximately 0.26 mRad s^{-1} . Correcting for this drift in the mean phase, I re-plot the same data set in figure 7.9(d). The width of this distribution is 12° ($\sim 200 \text{ mRad}$), which is much larger than the uncertainty of a single measurement, and compares favourably to data from similar experiments [5] [120]. In contrast, the observed spread of the other parameters is consistent with the uncertainty of a single measurement. This shows that the phase of the interfering condensate is fluctuating from shot to shot. If we assume that the two modes have been separated for a few ms before the trap was released, then the width of the observed distribution can be explained by phase spreading due to number fluctua-




Figure 7.10: Slow phase drift over the course of the experiment.

tions according to equation (6.23) (the distribution spreads by $\sim 100 \text{ mRad ms}^{-1}$ for our experimental parameters $\mu = 2\pi \times 2.2 \text{ kHz}$ and $N = 1.5 \times 10^4$ atoms).




Figure 7.11: Slow drift of the centre of the gaussian envelope over the course of the experiment.

The linear phase drift can be almost entirely explained by a corresponding slow drift in the centre of the gaussian envelope over the course of the experiment, which is shown in figure 7.11(c). The centre shifts 1.75 pixels ($4.4 \mu\text{m}$) over the course of the experiment. At a wavelength of 5.15 pixels ($17.5 \mu\text{m}$) this corresponds to a total phase shift of ~ 2.1 radians, compared to the measured total phase shift of ~ 2.3 radians. The corresponding drift in the wavelength and visibility of the interference fringes is negligible over the same period.

The phase of the fringes is supposed to be fixed by the geometry of the double well potential, independent of the cloud position in free fall. The correlation between phase and cloud centre suggests either that the imaging beam is slowly drifting due to a pointing instability in the optical path over the course of the experiment, or that the drift of the cloud centre is correlated with a change in the position of the double well potential. A shift of the trap position of $4.4 \mu\text{m}$ relative to the chip wires seems very unlikely (requiring a change in the lab field of close to 1 G, which is much larger than anything we measure near the experiment).

If we use the cloud centre as a reference point for the relative phase between the two

modes, then the shot to shot variation of phases is much greater than presented here. The reason is that the centres of the gaussian fits to the modulated density distributions have a standard deviation of ~ 1.35 pixels, which at a fringe wavelength of 5.15 pixels corresponds to an variation of the relative phase of $\sim \pi/2$ Rad.

Comparison of Fitting Procedures




Figure 7.12: Comparison of Fringe Fit and Fourier Analysis.

The data I have presented so far has been taken from fitting the modulated gaussian of equation (7.1) to the data. It is worthwhile comparing to this the results of the Fourier analysis routine. In figure 7.12 I plot the relative phases extracted from the two analysis methods against one another. There is an arbitrary offset of ~ 1 Rad between the two sets of phase measurements, but they align neatly on a line of slope one, as seen in the figure. The phase distribution extracted from the fringe fit has a slightly smaller width than that extracted from the fourier analysis ($\Delta\phi = 37^\circ$ compared to $\Delta\phi = 41^\circ$), and the mean visibility is slightly reduced (0.20 ± 0.03 compared to 0.21 ± 0.03). The wavelengths extracted are nearly identical (5.15 ± 0.07 compared to 5.16 ± 0.09). This strongly suggest that the two methods are interchangeable. Fourier analysis has the advantage of being much faster, and does not require filtering the image before applying the analysis routine. Obviously the experiment would be improved by reducing the intensity of the background fringes in the camera image. It is therefore preferable to use the Fourier analysis technique unless the wavelength of the interference pattern corresponds to a peak in the imaging noise, or cannot be clearly distinguished from the low frequency components of the gaussian envelope, in which case the procedure breaks down.

Ensemble Average




Figure 7.13: Ensemble average of 103 repetitions of the same experiment. The fit to the interference pattern has a visibility of $0.21(0)$, which, given the finite imaging resolution, is consistent with a first order spatial coherence (defined in equation (6.19)) of ~ 0.5 .

The coherence of the splitting procedure is also revealed in the ensemble average of the data set. In figure 7.13 I show the average density of the 103 experiments, and the corresponding line density and Fourier spectrum. The incoherent imaging noise has averaged away, and the sideband in the Fourier spectrum at the fringe frequency is prominent. The fitted wavelength is 5.17 ± 0.04 pixels, which is slightly longer than the individual images. The visibility is 0.21 ± 0.01 , which is statistically indistinguishable from the mean visibility of the individual fringe patterns in the ensemble (0.20 ± 0.03).

The fringe visibility of the ensemble averaged density distribution is a measure of the first order spatial coherence of the system, defined in equation (6.19) (see equation (6.20)). An ensemble averaged fringe visibility of 1 corresponds to a fixed relative phase between the two modes of the double well. Given the finite resolution of our imaging system, which limits us to observing a maximum fringe visibility of ~ 0.45 at a wavelength of $17.5 \mu\text{m}$, the measured visibility is consistent with $\alpha \geq 0.45$. In our experiment the measured coherence is limited by the presence of longitudinal phase fluctuations which lead to diffusion of the mean relative phase along the length of the condensate after the two modes are separated and reduce the average fringe visibility of a single shot of the experiment (see section 7.5.3 below). We have also seen evidence that the splitting process can excite axial oscillation modes of the system, which again leads to a loss of fringe visibility. Thermal fluctuations of the finite temperature system also lead to a loss of coherence. These effects are currently under further investigation.

7.5 Some observations about BEC interference

We have established that the splitting procedure described in section 7.2 is phase coherent, and that we can read out the relative phase between the two modes of the condensate by

analysing the interference pattern of the overlapping clouds in free fall. We are thus at the point where we can try to make a measurement using our BEC interferometer. In addition, we are studying the decoherence of the relative phase in the interferometer once the two modes are well separated, both from longitudinal phase fluctuations and phase spreading due to on-site interactions. We are also studying how to control these effects by varying the parameters of the splitting process such as the characteristic time take to split the two modes. These topics will be the subject of future reports.

To finish this thesis, I would like to present a few qualitative observation drawn from our early interference experiments.

7.5.1 The Effect of Interactions

So far we have assumed that the two modes of the condensate act as point sources for the observed interference pattern, and have ignored the contribution of interactions in the free fall expansion of the condensate. For an elongated BEC the radial expansion scales as $\sqrt{1+\tau}$, where $\tau = \omega_{\perp} t$. The initial expansion is dominated by the conversion of the release energy (largely the interaction energy $2\mu/5$ of the BEC) into kinetic energy. For our ~ 2 kHz radial trap frequency, this initial acceleration is over within the first 1 ms. It will show up as a small shift in the wavevector of the interfering condensates.

We have also implicitly assumed that the size of the two condensate modes is small compared to their separation. In our experiment the radial wave function is $\sim 0.5 \mu\text{m}$ in the double well potential, and the separation is $\sim 3.75 \mu\text{m}$, so this assumption may not be fully justified. In this case we should take into account the finite size of the two modes, which will modify the fringe spacing of the interference pattern given by equation (6.8). For extended sources, we should average over the density distribution of the two modes so that

$$\Lambda = \frac{ht}{m} \left\langle \frac{1}{\mathbf{d}} \right\rangle \quad (7.14)$$

where

$$\left\langle \frac{1}{\mathbf{d}} \right\rangle = \iint d\mathbf{r}_a d\mathbf{r}_b \frac{n(\mathbf{r}_a)n(\mathbf{r}_b)}{|\mathbf{r}_a - \mathbf{r}_b|} \quad (7.15)$$

In the original BEC interference experiment of Andrews et al. [1] the interference fringes were found to be inversely proportional not to d but to $\sqrt{d \langle \mathbf{r} \rangle}$ (via numerical extrapolation - see also the numerical studies of this experiment in [119] [121] [122]). Schumm et al. [5] also reported a deviation in the fringe spacing of equation (6.8) due to interactions. We believe that the difference between $\langle 1/\mathbf{d} \rangle$ and $1/\langle \mathbf{d} \rangle$ explains the observations.

An effect that we observe in our experiment (see below) is the appearance of a thick central fringe when the initial wave functions have a significant overlap [119]. Since the initial wave functions have a significant overlap, this fringe appears at *zero relative phase*, which means that we can use an images like that in figure 7.16(a) to provide a reference point for the relative phase which we otherwise would not have in our experiment (we cannot resolve the two modes of the double well in trap with our imaging system). The width of the central fringe is also a measure of the overlap of the condensate wave functions [123]. For strong interactions, the nonlinear term in the GPE can also broaden each fringe (while keeping the same fringe spacing) leading to a loss of contrast in an experiment with finite imaging resolution [124].




Figure 7.14: 1D GPE calculations of ground states and interference fringes for realistic double well potentials.

We see some of these effects in our experiments when we vary the separation between the minima of the double well potential. The simplest control parameter we can vary in making a double well potential is the RF current, which changes the magnitude of the RF field and thus the distance separating the two wells. The expected behaviour for an ideal interference experiment is shown in figure 7.14. When the trap separation is small compared to the size of the condensate wavefunction and there is still a large overlap between the two modes of the condensate, the interference pattern has a single broad central peak with small side lobes. As the separation increases, the full interference pattern becomes visible, and the wavelength decreases in proportion to the mean inverse separation between the two modes $\langle \frac{1}{d} \rangle$.

When we run this experiment, we do indeed see the wavelength of the interference pattern decrease as we increase the RF current in the wires. The results are plotted in figure 7.15. Below a critical point, the two modes of the double well potential are not




Figure 7.15: Fringe wavelength vs. RF current.

well separated, and we see a single broad interference fringe with small side lobes as in the interference pattern shown in figure 7.16(a) which corresponds to the point labelled *a*) in figure 7.15. The wavelength of the fringe fit ($30\ \mu\text{m}$) diverges because the interference pattern is no longer well described by equation (7.1). Above this current the wavelength decreases monotonically from $16.9\ \mu\text{m}$ to $13.8\ \mu\text{m}$ as we increase the current. Figure 7.16 shows the interference patterns corresponding to points *a*), *b*) and *c*) in figure 7.15. At the time of writing we have not been able to calibrate the RF current with a measured trap separation to make a quantitative study of the scaling of the wavelength (due to an unknown phase offset between the current in the two RF wires, discussed in section 5.3 of chapter 5). If we assume that the spacing is given by $\Lambda = \frac{h\tau}{m} \langle \frac{1}{d} \rangle$, the measured interference patterns in this data set correspond to separations ranging from $2.15\ \mu\text{m}$ for point *a*), when the two modes are not well separated, up to $4.75\ \mu\text{m}$.

Note that the central peak in figure 7.16(a) provides a reference for zero relative phase between the two modes of the condensate. We can potentially use this point to calibrate future phase measurements, whereas in the experiments described above, the mean phase extracted from the data set had an arbitrary offset determined by the choice of the integration region for the interference pattern. In figure 7.16(a) and figure 7.16(a) the fringe pattern has shifted due to a relative phase winding between the two condensate modes in the short time after they were fully separated and before the trap was turned off.




Figure 7.16: Fringe wavelength and phase vs. RF current.




Figure 7.17: Phase drift and spreading plotted against the amplitude of the RF current used to make the double well potential. A larger current corresponds to a greater final separation between the two traps, and a long time during which the two modes of the condensate have been well separated. We thus see a shift in the mean relative phase due to an imbalance in the trap potential, and a spreading of the relative phase distribution between the two modes which is due to number fluctuations in the split condensate.

In the data shown in figure 7.15 the only variable was the amplitude of the final RF current. The RF frequency, and the BEC and static magnetic trap parameters were the same for each experiment, so the two modes of the condensate separated (in all experiments) when the current reaches some critical amplitude (note that this is not the point when the double well starts to form, but the point when the adiabatic criterion for splitting breaks down). However the period over which the current was ramped up also stayed fixed, so that the speed of each linear ramp changed. As a result, the time at which the condensate modes separated is different for each of the experiments. In an experiment with a faster ramp (larger final current) the condensate has spent more time in the double well after separation than in an experiment with a slower ramp. We therefore expect to see a growth of the phase spread generated by the spread of number difference.

In figure 7.17 I plot the mean relative phases extracted from this data set, relative to the centre of the broad peak in figure 7.16(a). The error bars show the spread of phases around the mean. We see that indeed the spread of relative phases increases once the two components of the condensate are well separated. Since the two modes separate towards the end of the 20 ms linear amplitude ramp, this indicates that the phase spreading time in this experiment is of the order of 10 ms, consistent with the estimate of ~ 15 ms we calculated using in equation (6.24). In this data the mean relative phase also evolves as the trap is further separated, indicating that the double well potential was slightly asymmetric.

Note also that the visibility of the fringes shown in figure 7.16 initially increases as the two components become well separated, and then drops as the modes become further separated. This is due to the finite imaging resolution of our system, which washes out visibility of the fringes at shorter wavelengths.

7.5.2 Variation of Condensate Size

The effect of interactions on the interference pattern should show up if we vary the size of the condensate, keeping the parameters of the double well potential fixed. In figure 7.19 I plot the wavelength and visibility of the interference pattern against the end frequency of the evaporation ramp. For comparison the interference patterns of each data point are shown in figure 7.18 along with images of the corresponding condensates released from the static magnetic trap. We can see that as the condensate size is reduced the visibility of the interference fringes increases. The number of atoms in the smallest condensate is approximately $N \simeq 8 \times 10^3$ and in the largest $N \simeq 23 \times 10^3$.

We speculate that the increased visibility is due to a decrease in the amplitude of




Figure 7.18: *Fringe wavelength vs. condensate size. In this data we kept the parameters of the double well constant and changed the size of the condensate by evaporating further in the static magnetic trap before forming the double well. We see an increase in fringe visibility as the condensate becomes smaller. The number of atoms in the smallest condensate is approximately $N \simeq 8 \times 10^3$ and in the largest $N \simeq 23 \times 10^3$. We speculate that the increased visibility is due to a decrease in the amplitude of longitudinal phase fluctuations in the condensate as we reduce the size and temperature.*




Figure 7.19: *Fringe visibility and wavelength plotted against the end frequency of the evaporation ramp for the data presented in figure 7.18.*

longitudinal phase fluctuations in the condensate as we reduce the size and temperature, similar to what has been observed by Jo et al. [104]. There is also a hint that the wavelength of the interference pattern decreases as the condensate becomes smaller (and the mean inverse separation between the two modes larger), a correction due to the finite size of the initial condensate modes in the double well potential. We have a similar shift in two different data sets. The shift in both cases is small, but we are attempting to study this in more detail.

7.5.3 Longitudinal Phase Fluctuations

In the interference data I have presented so far, we chose trap parameters so as to avoid significant phase variation along the length of our condensate. However, we have also performed interference experiments with the much more elongated BEC that was described at the end of chapter 4. This BEC had a larger radial trapping frequency ($\omega_{\perp} = 2\pi \times 2.3 \text{ kHz}$) and smaller axial trapping frequency ($\omega_z = 2\pi \times 6.5 \text{ Hz}$). The resulting aspect ratio is about 350. These condensates typically had $\sim 3 \times 10^4$ atoms, a chemical potential $\mu \simeq h \times 4 \text{ kHz}$ and a phase coherence length of $3.5 \mu\text{m}$, which is much less than their length of $110 \mu\text{m}$, breaking them up into many more phase domains (about 30) than our standard BEC (as a point of comparison, recall that the length of our usual condensate is $41 \mu\text{m}$ and the phase coherence length $\sim 23 \mu\text{m}$, giving only two longitudinal phase domains on average).




Figure 7.20: We see here data from two interference experiments from the highly elongated BEC described in the text. Data for two different trap separations are shown. A clear interference pattern is observed, despite significant longitudinal phase fluctuations in this condensate. However the fringe visibility is reduced compared to typical interference experiments using our usual condensate.




Figure 7.21: Distribution of relative phases from 24 repetitions of the same experiment. The fringe visibility is plotted on the radial axis. The phase spread of the data is $\Delta\phi = 109^\circ$. The length mean phasor is less than the Rayleigh radius of a uniform distribution at the 95% confidence level.

Typical experimental data from an interference experiment with the elongated BEC, run in the same way as described in section 7.2, is shown in figure 7.20 and figure 7.21. The interference pattern was observed after 16ms time of flight.

In figure 7.20 I show typical interference patterns for experiments with two different amplitudes of the final RF field. The observed wavelengths of $30\mu\text{m}$ and $20\mu\text{m}$ correspond to a mean separation between the modes of $2.5\mu\text{m}$ and $3.7\mu\text{m}$. A clear interference signal is seen, but the visibility of the fringes is less than in our usual experiments. This was always the case with the elongated condensate. In addition, we never observed phase coherence between the two modes of the elongated BEC. A typical data set from 24 repetitions of the same experiment is shown in figure 7.21. We cannot distinguish this data set from a random sampling from a uniform phase distribution. We speculate that this is due to rapid decoherence due to longitudinal phase fluctuation after the two modes of the elongated condensate separate.

7.5.4 Interference of Zeeman Sublevels

Finally, we investigated the effect of projecting the atoms trapped in the RF dressed state potentials onto different Zeeman sublevels in the turn-off of the double well potential. We were interested in whether a projection onto different m_F -states might affect the contrast of our interference experiments. In order to study this effect we keep a small DC current flowing through the trapping wires for a short period when we switch off the trapping potential. The atoms were thus exposed to a vertical field gradient during the first $\sim 100 - 200\mu\text{s}$ of free fall, which separated the m_F components enough for them to be resolved at long drop times.

What we then observed was an interference pattern for each of the separate m_F -states.




Figure 7.22: *Interference of Zeeman sublevels separated by a magnetic field gradient in free fall. Here we see high visibility fringes on the $m_F = 1, 2$ components of the condensate, which have been separated during the trap turn off by a vertical field gradient.*

A typical example is shown in figure 7.22. Interference patterns with good fringe visibility are seen on two of the three components. The corresponding line densities extracted from the centre of the cloud via our usual procedure are also shown in the figure. The anti-trapped states are rapidly expelled from the trapping region during the trap turn off, but the $m_F = 0, 1, 2$ states can be seen in figure 7.22. The density of atoms in the $m_F = 0$ state in the figure is too small to observe interference fringes, but we have seen fringes on this component in other experiments. The relative populations of the m_F components depends on the phase of the RF dressing field at the moment when it is switched off. We have found that this is reproducible, allowing us to gather some statistics on the interference patterns that we observe. A typical data set is presented in figure 7.23. Histograms of the relative phases measured by the $m_F = 0, 1, 2$ components indicate that phase coherence between the two modes is preserved during the projection onto the m_F -states. In addition, there is a clear correlation between the phases observed in each of the two clouds.

From a practical point of view this is interesting because it allows us to gather information about the phase shifts that are acquired between the two modes after the projection due to stray magnetic field gradients in the turn off of the trap and in free fall. Each of




Figure 7.23: Relative phases of the different Zeeman components are clearly non-random and correlated.

the m_F —states will respond differently to these field gradients, potentially allowing us to quantify systematic errors introduced during this stage of the read-out process.

Outlook

I have demonstrated in this thesis that we are able to split a single Bose Einstein condensate coherently in a double well. This means that our atom chip can be used as a matter wave interferometer.

The outlook for this experiment is now extremely promising. A systematic study of the the phase spreading dynamics is now under way, with the aim of increasing the coherence time of the separated condensate modes. Once this is properly characterised, we aim to show that we can use the interferometer as a sensitive device for measuring small forces acting on the atoms. We are also studying how to control the splitting process in order to reduce the relative number fluctuations between the two modes, and thus further increase the coherence time. These studies will be reported in the thesis of the current graduate student running the experiment, Florian Baumgärtner. I wish him the best of luck.

Bibliography

- [1] M. Andrews, C. Townsend, H. Miesner, D. Durfee, D. Kurn, and W. Ketterle, ‘Observation of Interference Between Two Bose Condensates’, *Science* **275**, p. 637 (1997).
- [2] Y. Shin, M. Saba, T. A. Pasquini, W. Ketterle, D. E. Pritchard, and A. E. Leanhardt, ‘Atom Interferometry with Bose-Einstein Condensates in a Double-Well Potential’, *Phys. Rev. Lett.* **92**(5), 050405 (2004).
- [3] J. Fölling, S. Hunsmann, M. Cristiani, and M. Oberthaler, ‘Direct Observation of Tunneling and Nonlinear Self-Trapping in a Single Bosonic Josephson Junction’, *Phys. Rev. Lett.* 010402 (2005).
- [4] S. Levy, E. Lahoud, I. Shomroni, and J. Steinhauer, ‘The A.C. and D.C. Josephson Effects in a Bose-Einstein Condensate’, *Nature* **449**(7162), p. 579 (2007).
- [5] T. Schumm, S. Hofferberth, L. M. Andersson, S. Wildermuth, S. Groth, I. Bar-Joseph, J. Schmiedmayer, and P. Kruger, ‘Matter-Wave Interferometry in a Double Well on an Atom Chip’, *Nat. Phys.* **1**(1), p. 57 (2005).
- [6] G.-B. Jo, Y. Shin, S. Will, T. A. Pasquini, M. Saba, W. Ketterle, and D. E. Pritchard, ‘Long Phase Coherence Time and Number Squeezing of Two Bose-Einstein Condensates on an Atom Chip’, *Phys. Rev. Lett.* **98**(3), 030407 (2007).
- [7] G.-B. Jo, J.-H. Choi, C. A. Christensen, T. A. Pasquini, Y.-R. Lee, W. Ketterle, and D. E. Pritchard, ‘Phase-Sensitive Recombination of Two Bose-Einstein Condensates on an Atom Chip’, *Phys. Rev. Lett.* **98**(18), 180401 (2007).
- [8] R. Gati, B. Hemmerling, J. Fölling, M. Albiez, and M. Oberthaler, ‘Noise Thermometry with Two Weakly Coupled Bose-Einstein Condensates’, *Phys. Rev. Lett.* **96**(13), 130404 (2006).
- [9] R. Gati and M. Oberthaler, ‘A Bosonic Josephson Junction’, *J. Phys. B* **40**, p. R61 (2007).

- [10] J. Esteve, C. Gross, A. Weller, S. Giovanazzi, and M. K. Oberthaler, ‘Squeezing and Entanglement in a Bose-Einstein Condensate’, *Nature* **455**(7217), p. 1216 (2008).
- [11] S. Hofferberth, I. Lesanovsky, T. Schumm, A. Imambekov, V. Gritsev, E. Demler, and J. Schmiedmayer, ‘Probing Quantum and Thermal Noise in an Interacting Many-Body System’, *Nat. Phys.* **4**(6), p. 489 (2008).
- [12] S. Hofferberth, I. Lesanovsky, B. Fischer, T. Schumm, and J. Schmiedmayer, ‘Non-Equilibrium Coherence Dynamics in One-Dimensional Bose Gases’, *Nature* **449**(7160), p. 324 (2007).
- [13] C. Sinclair, J. Retter, E. Curtis, and B. Hall, ‘Cold Atoms in Videotape Micro-Traps’, *Eur. Phys. J. D* **35**, p. 105 (2005).
- [14] S. Eriksson, M. Trupke, H. Powell, and D. Sahagun, ‘Integrated Optical Components on Atom Chips’, *Eur. Phys. J. D* **35**, p. 135 (2005).
- [15] D. A. Steck, ‘Rubidium 87 D Line Data’, available online at <http://steck.us/alkalidata>, (revision 2.1.1, 30 April 2009).
- [16] J. Reichel, ‘Microchip Traps and Bose-Einstein Condensation’, *Appl. Phys. B* **74**(6), p. 469 (2002).
- [17] W. Hansel, *Magnetische Mikrofällen für Rubidiumatome*, Ph.D. thesis, Ludwig-Maximilians-Universität München (2001).
- [18] J. Fortágh, H. Ott, S. Kraft, A. Günther, and C. Zimmermann, ‘Surface Effects in Magnetic Microtraps’, *Phys. Rev. A* **66**, 041604 (2002).
- [19] M. Jones, C. Vale, D. Sahagun, B. Hall, C. Eberlein, B. Sauer, K. Furusawa, D. Richardson, and E. Hinds, ‘Cold Atoms Probe the Magnetic Field Near a Wire’, *J. Phys. B* **37**, p. L15 (2004).
- [20] J. Estève, C. Aussibal, T. Schumm, C. Figl, and D. Mailly, ‘Role of Wire Imperfections in Micromagnetic Traps for Atoms’, *Phys. Rev. A* **70**, 043629 (2004).
- [21] S. Kraft, A. Gunther, H. Ott, D. Wharam, and C. Zimmermann, ‘Anomalous Longitudinal Magnetic Field Near the Surface of Copper Conductors’, *J. Phys. B* **35**, p. L649 (2002).
- [22] A. Leanhardt, Y. Shin, A. Chikkatur, D. Kielpinski, W. Ketterle, and D. Pritchard, ‘Bose-Einstein Condensates near a Microfabricated Surface’, *Phys. Rev. Lett.* **90**, 100404 (2003).
- [23] T. Schumm, *Bose-Einstein Condensates in Magnetic Double Well Potentials*, Ph.D. thesis, Ruperto-Carola University of Heidelberg (2005).
- [24] P. Krüger, L. M. Andersson, S. Wildermuth, S. Hofferberth, E. Haller, S. Aigner, S. Groth, I. Bar-Joseph, and J. Schmiedmayer, ‘Potential Roughness near Lithographically Fabricated Atom Chips’, *Phys. Rev. A* **76**(6), p. 8 (2007).

- [25] S. Aigner, L. Pietra, Y. Japha, O. Entin-Wohlman, T. David, R. Salem, R. Folman, and J. Schmiedmayer, ‘Long-Range Order in Electronic Transport Through Disordered Metal Films’, *Science* **319**, p. 1226 (2008).
- [26] D. Wang, M. Lukin, and E. Demler, ‘Disordered Bose-Einstein Condensates in Quasi-One-Dimensional Magnetic Microtraps’, *Phys. Rev. Lett.* **92**(7), p. 076802 (2004).
- [27] Z. Moktadir, B. Darquie, M. Kraft, and E. A. Hinds, ‘The Effect of Self-Affine Fractal Roughness of Wires on Atom Chips’, *J. Mod. Opt.* **54**, p. 2149 (2007).
- [28] L. Pitaevskii and S. Stringari, *Bose-Einstein Condensation* (Oxford University Press, 2003).
- [29] N. J. V. Druten, ‘Two-Step Condensation of the Ideal Bose Gas in Highly Anisotropic Traps’, *Phys. Rev. Lett.* **79**(4), p. 549 (1997).
- [30] W. Ketterle and N. van Druten, ‘Bose-Einstein Condensation of a Finite Number of Particles Trapped in One or Three Dimensions’, *Phys. Rev. A* **54**(1), p. 656 (1996).
- [31] S. Giorgini, L. Pitaevskii, and S. Stringari, ‘Condensate Fraction and Critical Temperature of a Trapped Interacting Bose Gas’, *Phys. Rev. A* **54**(6), p. 4633 (1996).
- [32] G. Baym, J.-P. Blaizot, M. Holzmann, F. Laloë, and D. Vautherin, ‘The Transition Temperature of the Dilute Interacting Bose Gas’, *Phys. Rev. Lett.* **83**(9), p. 1703 (1999).
- [33] G. Baym, J.-P. Blaizot, M. Holzmann, F. Laloë, and D. Vautherin, ‘Bose-Einstein Transition in a Dilute Interacting Gas’, **24**(1), p. 107 (2001).
- [34] M. Holzmann, J. Fuchs, G. Baym, J. Blaizot, and F. Laloë, ‘Bose Einstein Transition Temperature in a Dilute Repulsive Gas’, *C. R. Phys.* **5**, p. 21 (2004).
- [35] W. Ketterle, D. Durfee, and D. Stamper-Kern, ‘Making, Probing and Understanding Bose-Einstein Condensates, Bose-Einstein Condensation in Atomic Gases’, *Proceedings of the International School of Physics Enrico Fermi* **140**, p. 67 (1999).
- [36] M. Naraschewski, ‘Analytical Description of a Trapped Semi-Ideal Bose Gas at Finite Temperature’, *Phys. Rev. A* **58**(3), p. 2423 (1998).
- [37] Y. Castin and R. Dum, ‘Bose-Einstein Condensates in Time Dependent Traps’, *Phys. Rev. Lett.* **77**(27), p. 5315 (1996).
- [38] C. Menotti and S. Stringari, ‘Collective Oscillations of a One-Dimensional Trapped Bose-Einstein Gas’, *Phys. Rev. A* **66**(4), p. 6 (2002).
- [39] G. Baym and C. J. Pethick, ‘Ground-State Properties of Magnetically Trapped Bose-Condensed Rubidium Gas’, *Phys. Rev. Lett.* **76**(1), p. 6 (1996).
- [40] D. Petrov, G. Shlyapnikov, and J. Walraven, ‘Regimes of Quantum Degeneracy in Trapped 1D Gases’, *Phys. Rev. Lett.* **85**(18), p. 3745 (2000).

- [41] V. Dunjko, V. Lorent, and M. Olshanii, ‘Bosons in Cigar-Shaped Traps: Thomas-Fermi Regime, Tonks-Girardeau Regime, and In Between’, *Phys. Rev. Lett.* **86**(24), p. 5413 (2001).
- [42] K. K. Das, ‘Highly Anisotropic Bose-Einstein Condensates: Crossover to Lower Dimensionality’, *Phys. Rev. A* **66**(5), 053612 (2002).
- [43] J. Fuchs, X. Leyronas, and R. Combescot, ‘Hydrodynamic Modes of a One-Dimensional Trapped Bose Gas’, *Phys. Rev. A* **68**, 043610 (2003).
- [44] F. Gerbier, ‘Quasi-1D Bose-Einstein Condensates in the Dimensional Crossover Regime’, *Europhys. Lett.* **66**(6), p. 771 (2007).
- [45] A. Amerongen, *One-Dimensional Bose Gas on an Atom Chip*, Ph.D. thesis, Universiteit van Amsterdam (2008).
- [46] S. Wildermuth, S. Hofferberth, I. Lesanovsky, S. Groth, P. Krüger, J. Schmiedmayer, and I. Bar-Joseph, ‘Sensing Electric and Magnetic Fields with Bose-Einstein Condensates’, *Appl. Phys. Lett.* **88**(26), 264103 (2006).
- [47] S. Wildermuth, *One-Dimensional Bose-Einstein Condensates in Micro-Traps*, Ph.D. thesis, Ruperto-Carola University of Heidelberg (2005).
- [48] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven, ‘Phase-Fluctuating 3D Bose-Einstein Condensates in Elongated Traps’, *Phys. Rev. Lett.* **87**(5), 050404 (2001).
- [49] S. Dettmer, D. Hellweg, P. Ryytty, J. J. Arlt, W. Ertmer, K. Sengstock, H. Kreutzmann, L. Santos, and M. Lewenstein, ‘Observation of Phase Fluctuations in Elongated Bose-Einstein Condensates’, *Phys. Rev. Lett.* **87**(16), 160406 (2001).
- [50] D. Hellweg, S. Dettmer, P. Ryytty, J. Arlt, W. Ertmer, K. Sengstock, D. Petrov, G. Shlyapnikov, H. Kreutzmann, L. Santos, and M. Lewenstein, ‘Phase Fluctuations in Bose-Einstein Condensates’, *Appl. Phys. B* **73**(8), p. 781 (2001).
- [51] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping* (Springer-Verlag, 1999).
- [52] M. Jones, *Bose-Einstein Condensation on an Atom Chip*, Ph.D. thesis, University of Sussex (2003).
- [53] C. D. James, *Bose-Einstein Condensation in Microtraps on Videotape*, Ph.D. thesis, Imperial College London (2005).
- [54] D. S. Sanchez, *Cold Atoms Trapped near Surfaces*, Ph.D. thesis, Imperial College London (2006).
- [55] U. Schünemann, H. Engler, R. Grimm, M. Weidemüller, and M. Zielonkowski, ‘Simple Scheme for Tunable Frequency Offset Locking of Two Lasers’, *Rev. Sci. Inst.* **70**(1), p. 242 (1999).

- [56] A. Arnold, J. Wilson, and M. Boshier, ‘A Simple Extended-Cavity Diode Laser’, *Rev. Sci. Instr.* **69**(3), p. 1236 (1998).
- [57] C. Wieman and T. W. Hänsch, ‘Doppler-Free Laser Polarization Spectroscopy’, *Phys. Rev. Lett.* **36**(20), p. 1170 (1976).
- [58] J. Fortagh, A. Grossmann, T. Hänsch, and C. Zimmermann, ‘Fast Loading of a Magneto-Optical Trap from a Pulsed Thermal Source’, *J. Appl. Phys.* **84**(12), p. 6499 (1998).
- [59] S. Groth, P. Krüger, S. Wildermuth, R. Folman, T. Fernholz, J. Schmiedmayer, D. Mahalu, and I. Bar-Joseph, ‘Atom Chips: Fabrication and Thermal Properties’, *Appl. Phys. Lett.* **85**(14), p. 2980 (2004).
- [60] E. Koukharenko, Z. Moktadir, M. Kraft, M. Abdelsalam, D. Bagnall, C. Vale, M. Jones, and E. Hinds, ‘Microfabrication of Gold Wires for Atom Guides’, *Sensor. Actuat. A* **115**, p. 600 (2004).
- [61] S. Aigner, *Magnetic Field Microscopy using Ultracold Atoms*, Ph.D. thesis, Ruperto-Carola University of Heidelberg (2007).
- [62] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. Cornell, ‘Simplified System for Creating a Bose-Einstein Condensate’, *J. Low Temp. Phys.* **132**(5-6), p. 309 (2003).
- [63] A. Arnold, *Preparation and Manipulation of an ^{87}Rb Bose-Einstein Condensate*, Ph.D. thesis, University of Sussex (1999).
- [64] J. Retter, *Cold Atom Microtraps Above a Videotape Surface*, Ph.D. thesis, University of Sussex (2002).
- [65] Z. T. Lu, K. L. Corwin, M. J. Renn, M. H. Anderson, E. A. Cornell, and C. E. Wieman, ‘Low-Velocity Intense Source of Atoms from a Magneto-Optical Trap’, *Phys. Rev. Lett.* **77**(16), p. 3331 (1996).
- [66] J. Reichel, W. Hänsel, and T. Hänsch, ‘Atomic Micromanipulation with Magnetic Surface Traps’, *Phys. Rev. Lett.* **83**(17), p. 3398 (1999).
- [67] W. Salzmann, *Atoms in a Long Magneto-Optical Trap*, Master’s thesis, University of Sussex (2001).
- [68] O. Zobay and B. M. Garraway, ‘Two-Dimensional Atom Trapping in Field-Induced Adiabatic Potentials’, *Phys. Rev. Lett.* **86**(7), p. 1195 (2001).
- [69] O. Zobay and B. Garraway, ‘Atom Trapping and Two-Dimensional Bose-Einstein Condensates in Field-Induced Adiabatic Potentials’, *Phys. Rev. A* **69**(2), 023605 (2004).
- [70] Y. Bo, L. Xiao-Lin, K. Min, and W. Yu-Zhu, ‘A Radio Frequency Field Guide Using Field Induced Adiabatic Potential’, *Chin. Phys. Lett.* **24**, p. 1260 (2007).

- [71] Y. Colombe, E. Knyazchyan, O. Morizot, B. Mercier, V. Lorent, and H. Perrin, ‘Ultracold Atoms Confined in RF-Induced Two-Dimensional Trapping Potentials’, *Europhys. Lett.* **67**(4), p. 593 (2007).
- [72] O. Morizot, L. Longchambon, R. K. Easwaran, R. Dubessy, E. Knyazchyan, P.-E. Pottie, V. Lorent, and H. Perrin, ‘Influence of the Radio-Frequency Source Properties on RF-Based Atom Traps’, *Eur. Phys. J. D* **47**(2), p. 209 (2008).
- [73] M. White, H. Gao, M. Pasienski, and B. Demarco, ‘Bose-Einstein Condensates in RF-Dressed Adiabatic Potentials’, *Phys. Rev. A* **74**, 023616 (2006).
- [74] W. Heathcote, E. Nugent, B. Sheard, and C. Foot, ‘A Ring Trap for Ultracold Atoms in an RF-Dressed State’, *New J. Phys.* **10**, 043012 (2008).
- [75] T. Fernholz, R. Gerritsma, P. Krueger, and R. J. C. Spreeuw, ‘Dynamically Controlled Toroidal and Ring-Shaped Magnetic Traps’, <http://arXiv.physics:0512017> (2005).
- [76] A. Günther, S. Kraft, C. Zimmermann, and J. Fortágh, ‘Atom Interferometer Based on Phase Coherent Splitting of Bose-Einstein Condensates with an Integrated Magnetic Grating’, *Phys. Rev. Lett.* **98**(14) (2007).
- [77] E. Maréchal, B. Laburthe-Tolra, L. Vernac, and J. K. Idots, ‘A Magnetic Lens for Cold Atoms Controlled by a RF Field’, *Appl. Phys. B* **91**, p. 233 (2008).
- [78] S. Hofferberth, I. Lesanovsky, B. Fischer, J. Verdu, and J. Schmiedmayer, ‘Radiofrequency-Dressed-State Potentials for Neutral Atoms’, *Nat. Phys.* **2**(10), p. 710 (2006).
- [79] I. Lesanovsky, T. Schumm, and S. Hofferberth, ‘Adiabatic Radio-Frequency Potentials for the Coherent Manipulation of Matter Waves’, *Phys. Rev. A* 033619 (2006).
- [80] S. Hofferberth, B. Fischer, T. Schumm, and J. S. Idots, ‘Ultracold Atoms in Radio-Frequency Dressed Potentials Beyond the Rotating-Wave Approximation’, *Phys. Rev. A* **76**, 013401 (2007).
- [81] J. J. P. van Es, S. Whitlock, T. Fernholz, A. H. van Amerongen, and N. J. van Druten, ‘Three-Dimensional Character of Atom-Chip-Based RF-Dressed Potentials’, *Phys. Rev. A* **77**, 063623 (2008).
- [82] C. Cohen-Tannoudji, J. Dupont-Roc, and C. Fabre, ‘A Quantum Calculation of the Higher Order Terms in the Bloch-Siegert Shift’, *J. Phys. B* **6**, p. L214 (1973).
- [83] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions* (Wiley-Interscience, 1998).
- [84] S. Haroche, ‘Étude théorique et expérimentale des propriétés physiques d’atomes en interaction avec des photons de radiofréquence’, *Ann. Phys.-Paris* **6**, p. 189 and 327 (1971).

- [85] D. T. Pegg, ‘Magnetic Resonance in Arbitrarily Oriented Fields’, *J. Phys. B* **10**, p. 1027 (1977).
- [86] W. A. McClean and S. Swain, ‘Crossings and Anti-Crossings for a Spin 1/2 System with Arbitrary Field Orientations’, *J. Phys. B* **9**, p. 1673 (1976).
- [87] E. Majorana, ‘Atomi Orientati in Campo Magnetico Variabile’, *Nuovo Cimento* **9**, p. 43 (1932).
- [88] F. Bloch and I. Rabi, ‘Atoms in Variable Magnetic Fields’, *Rev. Mod. Phys.* **17**(2-3), p. 237 (1945).
- [89] J. Javanainen and S. M. Yoo, ‘Quantum Phase of a Bose-Einstein Condensate with an Arbitrary Number of Atoms’, *Phys. Rev. Lett.* **76**(2), p. 161 (1996).
- [90] Y. Castin and J. Dalibard, ‘Relative Phase of Two Bose-Einstein Condensates’, *Phys. Rev. A* **55**(6), p. 4430 (1997).
- [91] G. S. Paraoanu, ‘Phase Coherence and Fragmentation in Weakly Interacting Bosonic Gases’, *Phys. Rev. A* **77**(4), 041605 (2008).
- [92] J. B. Liu and P. D. Ronney, ‘Modified Fourier Transform Method for Interferogram Fringe Pattern Analysis’, *Appl. Opt.* **36**(25), p. 6231 (1997).
- [93] D. J. Masiello and W. P. Reinhardt, ‘Time-Dependent Quantum Many-Body Theory of Identical Bosons in a Double Well: Early-Time Ballistic Interferences of Fragmented and Number Entangled States’, *Phys. Rev. A* **76**(4), 043612 (2007).
- [94] L. S. Cederbaum, A. I. Streltsov, Y. B. Band, and O. E. Alon, ‘Interferences in the Density of Two Bose-Einstein Condensates Consisting of Identical or Different Atoms’, *Phys. Rev. Lett.* **98**(11), 110405 (2007).
- [95] S. Anderloni, F. Benatti, R. Floreanini, and G. G. Guerreschi, ‘Noise Induced Interference Fringes in Trapped Ultracold Bosonic Gases’, <http://arXiv:0811.2857v1> (2008).
- [96] A. A. Burkov, M. D. Lukin, and E. Demler, ‘Decoherence Dynamics in Low-Dimensional Cold Atom Interferometers’, *Phys. Rev. Lett.* **98**(20), 200404 (2007).
- [97] R. Spekkens and J. Sipe, ‘Spatial Fragmentation of a Bose-Einstein Condensate in a Double-Well Potential’, *Phys. Rev. A* p. 3868 (1999).
- [98] R. Hanbury-Brown and R. Q. Twiss, ‘Correlation Between Photons in two Coherent Beams of Light’, *Nature* **177**, p. 27 (1956).
- [99] E. Altman, E. Demler, and M. D. Lukin, ‘Probing Many-Body States of Ultracold Atoms Via Noise Correlations’, *Phys. Rev. A* **70**(1), 013603 (2004).
- [100] S. Fölling, F. Gerbier, A. Widera, O. Mandel, T. Gericke, and I. Bloch, ‘Spatial Quantum Noise Interferometry in Expanding Ultracold Atom Clouds’, *Nature* **434**(7032), p. 481 (2005).

- [101] M. Lewenstein and L. You, ‘Quantum Phase Diffusion of a Bose-Einstein Condensate’, *Phys. Rev. Lett.* **77**(17), p. 3489 (1996).
- [102] E. M. Wright and D. F. Walls, ‘Collapses and Revivals of Bose-Einstein Condensates Formed in Small Atomic Samples’, *Phys. Rev. Lett.* **77**(11), p. 2158 (1996).
- [103] J. Javanainen, ‘Phase and Phase Diffusion of a Split Bose-Einstein Condensate’, *Phys. Rev. Lett.* **78**(25), p. 4675 (1997).
- [104] G.-B. Jo, J.-H. Choi, C. A. Christensen, Y.-R. Lee, T. A. Pasquini, W. Ketterle, and D. E. Pritchard, ‘Matter-Wave Interferometry with Phase Fluctuating Bose-Einstein Condensates’, *Phys. Rev. Lett.* **99**(24), 240406 (2007).
- [105] A. Polkovnikov, E. Altman, and E. Demler, ‘Interference Between Independent Fluctuating Condensates’, *Proc. N. Acad. Sci.* p. 6125 (2006).
- [106] V. Gritsev, E. Altman, E. Demler, and A. Polkovnikov, ‘Full Quantum Distribution of Contrast in Interference Experiments Between Interacting One-Dimensional Bose Liquids’, *Nat. Phys.* **2**(10), p. 705 (2006).
- [107] R. Bistritzer and E. Altman, ‘Intrinsic Dephasing in One-Dimensional Ultracold Atom Interferometers’, *Proc. N. Acad. Sci.* **104**(24), p. 9955 (2007).
- [108] A. Polkovnikov, ‘Shot Noise of Interference Between Independent Atomic Systems’, *Europhys. Lett.* **78**, p. 10006 (2007).
- [109] A. Imambekov, V. Gritsev, and E. Demler, ‘Fundamental Noise in Matter Interferometers’, <http://arXiv:cond-mat/0703766v1> (2007).
- [110] N. K. Whitlock and I. Bouchoule, ‘Relative Phase Fluctuations of Two Coupled One-Dimensional Condensates’, *Phys. Rev. A* **68**, 053609 (2003).
- [111] M. Takeda, H. Ina, and S. Kobayashi, ‘Fourier-Transform Method of Fringe-Pattern Analysis for Computer-Based Topography and Interferometry’, *J. Opt. Soc. America* **72**(1), p. 156 (1981).
- [112] K. A. Nugent, ‘Interferogram Analysis Using an Accurate Fully Automatic Algorithm’, *Appl. Opt.* **24**(18), p. 3103 (1985).
- [113] D. J. Bone, H. A. Bachor, and R. J. Sandeman, ‘Fringe-Pattern Analysis Using a 2-D Fourier Transform’, *Appl. Opt.* **25**(10), p. 1653 (1985).
- [114] K. A. Goldberg and J. Bokor, ‘Fourier-Transform Method of Phase-Shift Determination’, *Appl. Opt.* **40**(10), p. 2886 (2001).
- [115] G. J. G. Upton, ‘Single-Sample Tests for the von Mises Distribution’, *Biometrika* **60**(1), p. 87 (1973).
- [116] G. J. G. Upton, ‘Approximate Confidence Intervals for the Mean Direction of a von Mises Distribution’, *Biometrika* **73**(2), p. 525 (1986).
- [117] N. Fisher, *Statistical Analysis of Circular Data* (Cambridge University Press, 1993).

- [118] K. V. Mardia and P. E. Jupp, *Directional Statistics* (John Wiley & Sons, 2000).
- [119] A. Röhrig, M. Naraschewski, A. Schenzle, and H. Wallis, ‘Transition from Phase Locking to the Interference of Independent Bose Condensates: Theory versus Experiment’, *Phys. Rev. Lett.* **78**(22), p. 4143 (1997).
- [120] S. Hofferberth, *Coherent Manipulation of Bose-Einstein Condensates with Radio-Frequency Adiabatic Potentials on Atom Chips*, Ph.D. thesis, Ruperto-Carola University of Heidelberg (2007).
- [121] H. Wallis, A. Röhrig, M. Naraschewski, and A. Schenzle, ‘Phase-Space Dynamics of Bose Condensates: Interference Versus Interaction’, *Phys. Rev. A* **55**(3), p. 2109 (1997).
- [122] M. Naraschewski, H. Wallis, A. Schenzle, and P. Zoller, ‘Interference of Bose Condensates’, *Phys. Rev. A* **54**(3), p. 2185 (1996).
- [123] H. Wallis and H. Steck, ‘Inseparable Time Evolution of Anisotropic Bose-Einstein Condensates’, *Europhys. Lett.* **41**, p. 477 (1998).
- [124] W. Liu, B. Wu, and Q. Niu, ‘Nonlinear Effects in Interference of Bose-Einstein Condensates’, *Phys. Rev. Lett.* **84**(11), p. 2294 (2000).