

**An Integrated Optical-Waveguide Chip for
Measurement of Cold-Atom Clouds**

Manuel Succo

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

Imperial College London
Physics Department

February 2011

Referees

Referee: Professor Simon L. Cornish
Department of Physics
Durham University

Co-referee: Professor Leszek Frasiniski
Department of Physics
Imperial College London

Declaration

I declare that I have done this thesis on my own and without any aids other than the declared ones.

The work has not been submitted previously in same or similar form to another examination committee and has not yet been published.

Manuel Succo, London, 15th February 2011

Abstract

This thesis introduces the first demonstration of a monolithic, micro-fabricated, multi-channel, optical-waveguide chip to measure ultra-cold atomic clouds. The optics consist of an array of 12 independent junctions, which are separated by only $10\ \mu\text{m}$ and have large atom-photon coupling. The integrated and scalable design is presented, along with an atom chip for mounting the optical waveguide chip and magnetically trapping and handling ultra-cold atoms. The experimental apparatus which was built to accommodate this new chip set is described, along with a new experimental control programme which was developed to accommodate the scalability requirements of the new chip. The chip was optically, mechanically and magnetically characterised and cold atoms with densities up to $10^{-2}\ \mu\text{m}^{-3}$, corresponding to 1 atom at a time inside the waveguide mode, were detected with this new kind of chip using absorption and fluorescence techniques. Subsequently, the atoms were utilised to diagnose light polarisation and intensity within the optical-waveguide chip. For future use, various detection methods adapted to the optical-waveguide chip were considered to minimise photon scattering and thus heating of a trapped ultra-cold sample of atoms.

Contents

1	Introduction	7
2	Magnetic Trapping Basics	12
2.1	Current-Carrying Structures	13
2.1.1	The Z-Trap	14
2.1.2	Fragmentation of the Cloud	16
2.1.3	U-MOT	17
2.2	Magnetic Trap Parameters	17
2.2.1	Location in Space	18
2.2.2	Trap Bottom	20
2.2.3	Direction of Principal Axis	21
2.2.4	Oscillation Frequencies	22
2.2.5	Cloud Size	24
2.2.6	Trap Depth	25
2.3	Constraints	27
2.3.1	Current Limit	27
2.3.2	Size of the Waveguide Chip Trench	28
2.4	Tweaking the Trap	28
3	Experimental Setup	30
3.1	Lasers	30
3.1.1	Reference Laser	30
3.1.2	Cooling & Detection Laser	32
3.1.3	Repump Laser	33
3.1.4	Optical Pumping	33
3.2	Vacuum	33
3.2.1	Requirements & Design	34
3.2.2	Installation	38

3.3	Detection	40
3.3.1	Imaging	41
3.3.2	Waveguide Chip Detection	43
3.4	Computer Control	44
3.4.1	Architecture	44
3.4.2	TycheUI	45
3.4.3	TycheController	48
3.4.4	Data Acquisition Controllers	48
3.4.5	Advantages & Disadvantages	49
4	The Sub-Chip:	
	Manufacturing Process and Characterisation	51
4.1	Design	51
4.2	Manufacturing Process	53
4.2.1	How to put waveguides into a chip	53
4.2.2	How to make this atom-chip	54
4.3	Performance Tests	62
4.3.1	Sub-Chip	63
4.3.2	Waveguide Chip	68
4.4	Future Improvements	70
5	The Waveguide Chip:	
	Atoms in the Trench	72
5.1	Detection System	72
5.1.1	Corrections	73
5.1.2	Shot-Noise Limit	74
5.2	Moving Atoms into the Trench	76
5.3	Absorption Measurement	77
5.3.1	Technique	78
5.3.2	Probe Frequency	81
5.3.3	Probe Intensity	82
5.3.4	Probe Polarisation	85
5.3.5	More Channels	86
5.4	Fluorescence Measurement	87
5.4.1	Technique	87
5.4.2	Results	88

6 Photon-Shot-Noise-Limited Detection	90
6.1 Light-Matter Interaction	90
6.1.1 Condensate Density	90
6.1.2 Effective Thickness	93
6.1.3 Light-Condensate Overlap	93
6.1.4 Suseptibility	94
6.1.5 Effective Dipole Matrix Elements	97
6.2 Optical Model	98
6.2.1 Simulating the Interaction	98
6.2.2 Analytical Interaction Model	101
6.3 Prospective Experiments	102
6.3.1 Absorption Measurement	104
6.3.2 Mach-Zender Interferometer	107
6.3.3 Beat Note Detection	110
6.3.4 Summary of the Techniques	117
7 Conclusion and Outlook	119
List of Figures	121
List of Tables	129
Bibliography	130

Chapter 1

Introduction

The first experimental realisation and property studies of Bose-Einstein condensation (BEC) [1–3] in dilute atomic gases in 1995 by Eric Cornell, Carl Wieman and, independently, Wolfgang Ketterle earned them the 2001 Nobel price and created a novel field among atomic and condensed matter physics.

Following experiments in 1999 using current carrying, and in some cases charged, wires glued to glass slips to magnetically guide ultra-cold atoms [4,5], much study in recent years has focussed on progress in miniaturisation of atom control devices. Studying ultra-cold vapours employing so called *atom chips* [6] has grown in popularity ever since, and hit a milestone in 2001 with the first creation of a BEC on a chip [7,8]. These atom chips form a magnetic trapping field by sending electric currents through thin wires configured in a fashion leading to harmonic trapping potentials. Since then, more complex wire configurations lead to more adjustable trapping potentials which in turn lead to a wide range of applications for these kinds of traps [9,10], as well as traps using permanent magnetic fields [11–14]. Since the detection of cold atoms is almost always optical, bringing micro-fabrication of optics and atom chips together has become a major undertaking.

Several groups have explored single mode optical fibres with typical diameters of $125\ \mu\text{m}$ glued [15–18] or otherwise attached [19] to a chip. When aligned opposite each other, they can be used to detect a cloud of atoms with close to single-atom sensitivity [16]. To increase the interaction between the light and atoms and enter a region of strong atom-photon coupling, reflective coatings can be added to the two fibers [18,20,21] or one

can be coated and one replaced by a micro-fabricated mirror [17,20,22,23] to form a Fabry-Perot resonator (or optical cavity). These devices may prove powerful in applications in quantum processing and atom-detection has been demonstrated using optical fibres glued to a chip surface, by absorption [15,24] and fluorescence, the latter with single atom sensitivity [19,25]. However, the size of fibres places a limit on the level of miniaturisation and scalability of optical elements on atom chips. Using optical waveguides is the next step towards the goal of integrating atom traps and optical tools to probe and manipulate ultra-cold atoms.

The optical properties of waveguides, such as modes of propagation, attenuation and effects of the waveguide's geometry are well known [26]. The fabrication of miniature integrated optical waveguides and more complex optical structures on substrates has been studied even before the rise of Bose-Einstein condensation [27]. And as the telecommunication revolution advances, so do techniques for miniaturisation and mass fabrication of optical elements based on silicon [28], offering small sizes and the possibility of scalability of such devices.

The main advantage, though, of using optical waveguides in the context of atomic physics is that they are so small that they emit light with a small enough mode area to have a large coupling efficiency between light and atoms [18,29,30]. Also the large numerical apertures possible do not require the additional use of aberration-corrected multi-element lenses.

The technology of laser cooling and ultra-cold atoms has now matured to the point where applications have become feasible, for example as space-worthy time standards [31], gravitational-wave detectors [32] or surface magnetometers [33]. However, the reliability and robustness remain barriers to applications of ultra-cold atoms, although much progress in making laser sources [34] as well as detectors [35–37] has been made, taking experiments into micro-gravity environment through parabolic flights [38] and drop towers [39,40].

Micro-fabricated chips are widely used to manipulate ultra-cold atom clouds and Bose-Einstein condensates [18,41,42] and have recently also been extended to the control of ions [43] and molecules [44]. Vapour cells on the millimetre-scale have brought optical spectroscopy [45], clocks [46] and magnetometry [46] to atom chips. Furthermore, etched mirrors on a silicon wafer have integrated magneto-optical traps, a much-used tool for preparation of cold samples, on an atom chip [47]. Fresnel lenses [48]




Figure 1.1: Sketch of the gold-coated optical-waveguide chip, fabricated on a silicon substrate. Fibres are connected to both sides of the chip. The waveguides converge in the centre of the chip where a trench has been cut. There, atoms can be detected between opposing waveguides.

have been used to trap atoms less than 1 mm from a surface. Micro-disk resonators [49–51] and micro-torus resonators [49,50,52] have been used for exceptionally strong single atom-single photon interactions. Soon, micro-fabricated devices may take advantage of the elementary quantum nature of atoms [21, 53, 54], ions [55] and even molecules [44, 56] in conjunction with photons [57].

A big step towards this goal is building multi-element devices, such that detection can be built into the same architecture that generates the samples. The optical waveguide industry is now very advanced and amazing experiments have demonstrated on-chip switches for quantum metrology [58], fully-integrated controlled-NOT gates [57] and even Shor’s factoring algorithm [59] in silica-on-silicon waveguide circuits. The hope is that such integrated structures will lead to the desired reliability and robustness required for applications of cold atoms, eventually taking advantage of the single-mode, fibre-pigtailed telecommunications technology.

The optical-waveguide chip used in this thesis is using the same technology as [57–59], i.e. a doped SiO_2 . The waveguide chip is sketched in

figure 1.1. The chip is a silicon substrate with doped silica on top, forming 12 waveguides (only 6 shown) which converge in the centre region of the chip. There, a $16\ \mu\text{m}$ wide, $500\ \mu\text{m}$ long and $22\ \mu\text{m}$ deep trench has been cut to intersect the waveguides. The trench is just big enough to allow a cloud of magnetically trapped ultra-cold atoms to enter the trench, and also small enough to not lose too much light across the trench due to the divergence of the emitted light mode. Single-mode optical fibres are connected to both sides of chip, bringing laser light into the waveguides and collecting it after traversing the trench. This thesis is about this multi-channel optical-waveguide device, an accompanying paper of which has been published in Nature Photonics [60]. The other part is about the current-carrying atom-chip built to be able to trap atoms and bring them into the trench, i.e. the optical interaction region, of the optical-waveguide chip.

Description of Chapters

- 2. Trapping Atoms on a Chip** This chapter describes how cold atoms are trapped by means of the magnetic field of a current-carrying bulk structure which is to serve the waveguide chip. We start by looking at basic approximations and continue to numerical calculations to map out important parameters like trap depth, oscillation frequencies and principal axes of rotation.
- 3. Experimental Setup** This chapter contains an overview of the surrounding experimental apparatus, a big part of which has been built from scratch in this thesis. Information is given on the laser and vacuum system, detection hardware, supplementary magnetic fields, and finally the control software, which has also been re-written.
- 4. The Chip: Manufacturing Process and Characterisation** This chapter describes in much detail the underlying design and manufacturing process by which an atom chip including an optical waveguide chip was successfully produced. Furthermore, the atom chip was tested thermally, electrically, mechanically and magnetically and the results and conclusions of these tests are presented here. Results and conclusions on optical tests performed on the waveguide chip are also included. Possible improvements to the combined chip and the manufacturing process are discussed.

5. **The Chip: Experiments with Atoms** In this chapter first chip experiments with atoms are presented. The atoms were introduced into the waveguide chip and detected using their absorption or fluorescence response. Intensity and polarisation of the light inside the device were diagnosed and comparisons with a model are discussed.
6. **Photon-shot-noise-limited detection** This chapter takes a theoretical look at the interaction between atoms in a Bose-Einstein condensate brought inside the waveguide chip and the light as it is sent through the waveguide chip. It includes which signals and uncertainties are to be expected from intensity or phase measurement of the light. The latter is done first in the general context of a Mach-Zender interferometer, followed by the specific case of a beat-note interferometer. A discussion highlights advantages and disadvantages of either detection method.
7. **Conclusion and Outlook** This chapter concludes the thesis and covers some imminent goals achievable with the waveguide chip device in the present as well as in a modified setup.

Chapter 2

Magnetic Trapping Basics

This chapter gives a brief introduction on trapping atoms with magnetic fields, followed by calculations of the fields of the current structure built in the scope of this thesis. The process of building the current-carrying structure is described in detail in chapter 4. From the fields, the parameters of the trap, i.e. location in space, trap depth, direction of the principal axis, oscillation frequencies and size are inferred.

We describe the state of a rubidium 87 atom by its total angular momentum F composed of the total electron spin S , total orbital momentum L of the electrons and total nuclear spin I . The ground state of such a rubidium atom is denoted by $^2S_{1/2}$ and the magnitude of the nuclear spin I . In this state the magnitude F of F can take the values 1 or 2. Each of these hyperfine energy levels contains $2F + 1$ magnetic sublevels that determine the angular distribution of the electron wave function. In the absence of an applied magnetic field, the sublevels are degenerate. However, placing the atom in an external magnetic field \mathbf{B} lifts their degeneracy [61].

If the energy shift due to the magnetic field is small compared to the hyperfine splitting, then F is a good quantum number. For weak magnetic fields the energy levels split linearly. The potential energy of the atom in the field can be written as,

$$\mathcal{U} = g_F \cdot m_F \cdot \mu_B \cdot |\mathbf{B}|,$$

where μ_B is the Bohr magneton and g_F the hyperfine Landé factor. The sign of the product $g_F m_F$ determines the sign of the potential. States with $g_F m_F$ positive (negative) are weak (strong) field seeking states. Maxwell's equations do not allow for a static magnetic field maximum in free space

[62], hence trapping atoms requires a magnetic field configuration with a local field minimum and atoms in a weak-field-seeking state. For rubidium 87 the Landé factors are $-1/2$ and $+1/2$ for $F = 1$ and $F = 2$, respectively. This means, the states $|F, m_F\rangle$ that can be trapped are $|1, -1\rangle$, $|2, 1\rangle$ and $|2, 2\rangle$. Of these states $|2, 2\rangle$ is the state interacting strongest with the applied field: $g_F m_F = 1$. We primarily use this state for trapping atoms and exclusively treat its potential energy in this chapter.

There is one trade-off in using weak field seeking states. To preserve the Zeeman state in the magnetic trap, the magnitude of \mathbf{B} must not approach zero. In practice atoms in a small or zero magnetic field undergo Majorana spin-flip transitions to untrapped states and are lost from the trap, limiting the trap life time considerably. We require the atom's magnetic moment to follow the direction of \mathbf{B} . In other words, the change of the direction θ of the magnetic field experienced by a trapped atom is slower than the (classical) precession frequency of the magnetic moment, i.e. the Larmor frequency ω_{Larmor} [10, 61],

$$\dot{\theta} < \omega_{\text{Larmor}} = \frac{g_F m_F \mu_B}{\hbar} \cdot |\mathbf{B}|.$$

We have to take this into consideration when constructing the fields for magnetic trapping to avoid spin-flip loss and not limit the trap lifetime.

2.1 Current-Carrying Structures

Let us now calculate the field created by the current structures shown in figure 2.1. Here, the Z- and end structures are coplanar, the bias wires lie underneath the structure surface. For this treatment, however, we consider the trap to form above the chip. This structure cannot be approximated by infinitely thin wires, as opposed to calculating the field for wire traps where the approximation yields good results. This is because the distance between the trap location and the structure is on the order of the dimensions of the structure itself (especially the width of the Z-structure). We bypass this obstacle by using Radia, a software for 3d magnetostatics [63–65] in conjunction with Wolfram Mathematica to calculate the field and infer the trap parameters. Variables in this calculation are the five currents \mathcal{I}_Z , $\mathcal{I}_{\text{bias},1}$, $\mathcal{I}_{\text{bias},2}$, $\mathcal{I}_{\text{end},1}$, $\mathcal{I}_{\text{end},2}$ as they are displayed in figure 2.1, as well as a three-component magnetic field $\mathbf{B}_{\text{bkg}} = (B_x, B_y, B_z)$, which is homogeneous across the space we are considering for the calculations.

The dimensions of the structure used are as follows. The tracks are




Figure 2.1: The bulk current structure used to numerically calculate the total magnetic field. The rest of the H-structure, which the Z-structure (red) is part of, is not shown. The bias wires (blue) create the field in the x -direction to counter the field created by the Z-structure. Currents sent through the end structure (green) tweak the end points of the elongated trap, hence the name. The origin (white) of the coordinate system used lies in the centre of its symmetry but on top of the Z-structure surface.

0.5 mm thick and 1.2 mm wide. The centre bar is 6 mm in length and thus the centre-centre distance of the Z-tracks is 7.2 mm. The centre-centre distance of the end structure tracks is 10.6 mm. For the calculation, the coplanar structure spans over a distance of 20 mm from end to end. The bias wires have a centre-centre distance of 4 mm. The distance from their centre to the top surface of the structure ($y = 0$) is 1.2 mm. Although the bias-wires used in the experiment have a diameter of 0.9 mm the calculation approximates them as thin wires.

2.1.1 The Z-Trap

To get a feel for the general shape of the trap, we look at the example in figure 2.2. For simplicity, we only treat the field created by running a current $I_Z = 34$ A through the Z-structure and a current $I_{\text{bias},1} = I_{\text{bias},2} = I_{\text{bias}} = 80$ A through the bias wires. The background field and remaining currents




Figure 2.2: The unit for all spatial coordinates x , y and z is millimeter. The magnetic field for example currents of $I_Z = 34$ A and $I_{\text{bias}} = 80$ A in three cuts through the field minimum $y_{\text{min}} = 0.493$ mm. The field value at this point is $|\mathbf{B}_{\text{min}}| = 0.2$ mT. Its vector is lying all in the z -direction. (*left*) The magnitude of the field is plotted in units of mT and (*right*) the projection of the field vectors in arbitrary units. The elongated shape of the potential is visible, as well as the tilt of trap axis in the x - z cut. Overlaying the vector plots is the Z-structure in blue.

are zero: $\mathbf{B}_{\text{bias}} = \mathcal{I}_{\text{end},1} = \mathcal{I}_{\text{end},2} = 0$. The direction of the currents is chosen such that the magnetic field minimum forms above the origin for positive I_Z and I_{bias} .

In contrast to more common wire trap configurations where the bias

field is created by large coils far away from the trap, the bias field created by the bias wires quickly tends to zero when moving away from the field minimum. We have to keep this in mind for calculating the trap depth later in this chapter.

For the given values and a magnetic field minimum set to have a value of $|\mathbf{B}_{\min}| = 0.2 \text{ mT}$, the field minimum is located at $\mathbf{r}_{\min} = (0, 0.493, 0) \text{ mm}$. Figure 2.2 shows the total magnetic field strength $|\mathbf{B}|$ plotted on the left hand side and a 2d projection of the magnetic vector field \mathbf{B} on the right hand side. In both cases the cuts run through \mathbf{r}_{\min} along each of the three axes. Note also, that the axes are at different scales to fit the trap in the figure.

We can see a near-circular symmetry in the vicinity of the trap centre in the x - y -projection. Hence we can later recast the properties of the trap in the x - and y -direction into a radial and an angular part, with the angular part being approximately constant. Furthermore, the cloud is obviously elongated in the z -direction and the long axis of the trap is slightly rotated about the y -axis. This has the effect that the true confinement along the principal axial direction of the trap is slightly lower than what we calculate if we fit it simply along the z -direction. The opposite effect is true for the radial direction. The true confinement in the radial plane of the trap is slightly higher than what we calculate if we fit it in the x - y -plane. In addition, there is an S-shape quality to the trap (as viewed in the x - z -plane), which is, however, irrelevant for cold samples. In this case the size of the cloud is much smaller than the area displayed here and the S-shape becomes negligible. What is more, in the y - z -plane, the trap resembles a banana-shape with the ends pointing towards the chip. For currents of $I_Z = 50 \text{ A}$ and $I_{\text{bias}} = 50 \text{ A}$, when the cloud is relatively far away, the height difference at the ends of the cloud is only about 1% of its half length. Let us take a typical cloud with a half length of $150 \text{ }\mu\text{m}$ (c.f. table 2.1). The tips of the cloud are just above $1 \text{ }\mu\text{m}$ closer to the Z-structure than the centre of the cloud. This is well within the centre radius of such a cloud. For clouds closer to the surface, this effect is negligible. We ignore it for the remainder of the calculations presented here.

2.1.2 Fragmentation of the Cloud

Other groups, e.g [66–69], have reported problems of fragmentation of the trapped cloud close to the field creating structures, i.e. at typical distances of less than $100 \text{ }\mu\text{m}$. This effect stems from inhomogeneities in the magnetic

field, which in current carrying bulk structures are mainly attributed to edge corrugation (emerging when machining the structures). However, at the rather large distances ($> 500 \mu\text{m}$) between the trapped cloud and the surface of the current structure that we face, we do not expect any fragmentation effects to occur.

2.1.3 U-MOT

The current structure was designed to also form a field configuration that closely resembles a quadrupole by running the current in a U-fashion¹ in addition to the field the bias wires provide. Adding laser light close to atomic resonance will cool atoms within the light beam [71], in combination with the magnetic field, atoms will through the Zeeman-shift be pushed towards the magnetic field zero where they gather. This is called a magneto-optical trap² (MOT).

However, a U-MOT is not well suited to collect atoms from the background vapour because the field is only a true quadrupole near the field zero [72] and the capture volume of the trap is small for structures of our size. Nevertheless, such a MOT can be used to optimise the transfer of atoms from a MOT formed by coils into an elongated magnetic trap. This is due to the fact that compared to the gradient ratio of 1:1:2 for an ideal quadrupole, a typical U-MOT has 1:4:5 (1:3:4 in the optimised case). Although we do not present detailed calculations here for the U-MOT gradients, they were taken into account for the determination of the dimensions of the current structure.

2.2 Magnetic Trap Parameters

Now that we know in principle the shape of the potential formed by one configuration of our current structure, we can look closer at how to extract the aforementioned trap parameters for different currents applied. In this section we only look at how changes in I_Z and I_{bias} affect the trap param-

¹Actually the current configuration looks more like a rectangle with one of its sides missing. The two fields from the bent ends of the structure close the initial two-dimensional confinement along the central structure segment. The result is a three-dimensional quadrupole field, with a field minimum at $|\mathbf{B}| = 0$ [70].

²The inventors of the MOT, Steven Chu, Claude Cohen-Tannoudji and William D. Phillips, were jointly awarded the 1997 Nobel Prize in Physics.

eters. In section 2.4 we look at how to tweak parameters by changing the composition of accessible currents and fields.

As described in chapter 4 the Z-structure actually holds a waveguide chip with integrated trench for the atoms. The trench is at a distance of about 0.5 mm above the Z-structure. A typical experimental sequence comprises of three particular types of trap:

1. Loading Trap: atoms that have been pre-cooled in a magneto-optical trap are optically pumped into the $|2, 2\rangle$ state and transferred into this wide open magnetic trap.
2. Pre-Evaporation Trap: the trap is compressed to increase the collision rate in the sample to improve forced evaporative cooling and increase the phase-space density

$$PSD = \frac{\hbar\bar{\omega}}{k_B T} \cdot N$$

of the sample.

3. In-Trench Trap: the sample has been cooled down below the critical temperature

$$T_c = \frac{\hbar\bar{\omega}}{k_B} \cdot N^{1/3}$$

for condensate formation and moved to the location of the waveguide chip trench.

The parameters for these three points are shown in table 2.1. In the following section we discuss in detail the calculation of the quantities that are given in the table.

2.2.1 Location in Space

The location of the trap in space is calculated by using a simple minimum finding algorithm. Due to reasons of the symmetry of the geometry and the current flow, the minimum is bound to lie on the y -axis with $x = z = 0$ for $\mathbf{B}_{\text{bkg}} = 0$. Thus we need only take a look at the y -component of the location of the minimum, y_{min} .

Figure 2.3 shows a plot of y_{min} of the trap minimum as a function of the two currents I_Z and I_{bias} . Isodistance lines for distances 0.5, 1.0, and 1.5 mm distance from the Z-structure are drawn as well as the three points

Parameter	Unit	Post-Loading	Pre-Evaporation	In-Trench
Trap Parameters:				
I_Z	A	50	50	34
I_{bias}	A	50	70	80
Expected System Parameters:				
N_{atoms}		$5 \cdot 10^7$	$3 \cdot 10^7$	$1 \cdot 10^4$
T	μK	80	151	0.2
Results:				
y_{min}	mm	1.669	1.078	0.493
B_{min}	mT	0.200	0.200	0.100
T_{td}	μK	810	1712	12
B_z	mT	0.888	0.660	0.259
ω_r	$2\pi \times \text{Hz}$	178	417	1043
ω_a	$2\pi \times \text{Hz}$	16	20	16
$\bar{\omega}$	$2\pi \times \text{Hz}$	80	151	256
s_r	μm	10.9	6.1	0.7
s_a	μm	120.7	128.2	44.6
$\tan \vartheta$		0.174	0.117	0.056
AR		11	21	65
PSD		$5 \cdot 10^{-6}$	$3 \cdot 10^{-6}$	2.32

Table 2.1: Trap parameters for three important types of trap on the path to Bose-Einstein condensation. Symbols used: I_Z : Z-structure current, I_{bias} : bias wire current, y_{min} : vertical position of magnetic field minimum, B_{min} : value of magnetic field minimum, T_{td} : trap depth, B_z : shim field z-component, N_{atoms} : number of atoms, T : cloud temperature, $\omega_{r/a}$: radial/axial oscillation frequency, $\bar{\omega}$: geometric mean oscillation frequency, $s_{r/a}$: radial/axial cloud size, $\tan \vartheta$: relative cloud displacement, AR: aspect ratio, PSD: phase space density.

for which parameters were collected in table 2.1. We see that, similar to the case of a wire trap with homogeneous bias field, that the location of the trap is proportional to the ratio of the currents,

$$y_{\text{min}} \propto \frac{I_Z}{I_{\text{bias}}}.$$

This simply means that we have a wide range of current combinations




Figure 2.3: The distance r_{\min} is taken from the origin of the Z-structure (c.f. figure 2.1). For the considered geometry and $\mathbf{B}_{\text{bkg}} = 0$, all components but the y -component are zero. Therefore, we can express the location by plotting just y_{\min} , to specify the location of the trap minimum. Of special interest is the location of the waveguide chip trench and the three parameter sets from table 2.1. These points and isodistance lines for $y_{\min} = 0.5$ mm, 1.0 mm and 1.5 mm are drawn in this and all following parameter maps for comparison.

available to choose from to reach a certain distance for any other parameter values we require.

2.2.2 Trap Bottom

The trap bottom is the value of the local total magnetic field minimum,

$$B_{\min} = |\mathbf{B}(\mathbf{r} = \mathbf{r}_{\min})|.$$

The energy splitting between the hyperfine energy levels is $g_F \mu_B |\mathbf{B}|$. If we want to use evaporative cooling the evaporation ramp end frequency is

$$\omega_{\text{rf, end}} = \frac{k_B T}{\hbar} + \frac{g_F \mu_B}{\hbar} \cdot B_{\min}$$

to remove all atoms from the trap with an energy above $k_B T$. However, we primarily adjust B_z such that the resulting trap bottom field is 0.2 mT.




Figure 2.4: Plot of the relative horizontal displacement $\tan \vartheta$ of the axial principal axis of the magnetic trap due to its rotation around the y -axis. The amount of displacement is decreasing when closing in to the chip. This is mainly due to adjusting B_z to keep the trap bottom at 0.2 mT. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.

We do this to increase the radial trap frequencies or to hold the trap high enough to avoid trap loss due to spin flips.

2.2.3 Direction of Principal Axis

As we noted before in the plotting the total magnetic field in the z - x -plane in figure 2.2, the axial principal axis of the trap is rotated about the y -axis by an angle ϑ . This rotation should have a profound impact on the experiment, since the narrow interaction region of the waveguide chip is aligned along the z -axis.

To find the rotation angle, first, the minimum field location is found. Subsequently the minimum field location in the x - y -plane at a distance of $z = 100 \mu\text{m}$ is found. From these two locations, the rotation angle is derived³.

³Alternatively a matrix of the gradients of the field at the centre point can be found. Through diagonalising the matrix the principal axes (and oscillation frequencies along these) and thus the rotation angle of the field can be found.




Figure 2.5: Plot of the geometric mean radial oscillation frequency $2\pi \times \omega_r$. Due to setting B_z to yield a trap bottom of 0.2 mT throughout the map, the radial oscillation frequency is increasing with decreasing distance of the trap. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.

In figure 2.4 the tangent of the rotation angle is plotted. The value of the tangents translates into the relative horizontal displacement of the trap from the z-axis.

Comparing how the angle changes with distance to the chip (c.f. figure 2.3) we note that the amount of displacement decreases when closing in to the chip surface. For the an example value of $I_z = 34$ A and $I_{\text{bias}} = 80$ A, the rotation of the trap has a value of about $\tan \vartheta = 0.056$ ($\vartheta = 3.2^\circ$). With this value, a cloud that is 100 microns long experiences a displacement along the x -direction between its two ends of about $6 \mu\text{m}$. Although this cloud fits into the trench, its rotation is large enough to make it more difficult to manoeuvre it into the trench.

2.2.4 Oscillation Frequencies

The oscillation frequencies of the trap describe the curvature of the potential the atoms interact with. If we assume a harmonic potential in the vicinity of the centre of the trap the we can approximate the total magnetic field by

$$|\mathbf{B}| \approx B_{\min} + B_r'' \cdot \rho^2 + B_a'' \cdot \zeta^2.$$




Figure 2.6: Plot of the axial oscillation frequency $2\pi \times \omega_a$. Due to the control of B_z to a trap bottom of 0.2 mT the trap axially maintains the same shape, resulting in the axial oscillation frequency to plateau at about 11.5 Hz when coming closer to the chip. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.

The variable ρ and ζ are the coordinates across the principal plane and along the principal axis, i.e. they are rotated about the y -axis by ϑ . We combine the oscillation frequencies in the radial principal plane of the trap, by taking the geometric mean value of the oscillation frequencies of the two principal axes closed to the x - and y -axis.

Fitting the calculated field values to aforementioned form for $|\mathbf{B}|$ yields B_r'' and⁴ B_a'' as fitting parameters which can be translated into oscillation frequencies by identifying the potential energy of the simple harmonic oscillator with the approximation found above

$$\frac{1}{2}m_{\text{Rb}}\omega_k^2 x_k^2 = \mu_{\text{B}}B_k'' x_k^2$$

and thus

$$\omega_k = \sqrt{\frac{2\mu_{\text{B}}B_k''}{m_{\text{Rb}}}}$$

⁴And indeed B_{\min} . This value is preferred to the value found by the minimum-finding algorithm.

with m_{Rb} being the atomic mass of rubidium 87 and $k = \{r, a\}$. In figures 2.5 and 2.6 the radial and axial oscillation frequencies are plotted for different values of currents I_Z and I_{bias} . The radial oscillation frequency is increasing with the ratio of I_Z/I_{bias} . In contrast to this behaviour is the axial oscillation frequency which increases approximately with total current. What is also noted is that in controlling B_z we set the axial oscillation frequency to a value of about 11.5 Hz throughout a large region of the parameter map.

A good figure of merit for evaporative cooling is the collision rate after compression. In a 3d harmonic oscillator, the collision rate increases due to adiabatic compression proportional to the geometric mean of the three trap curvatures, which are proportional to the geometric mean,

$$\bar{\omega} = (\omega_r^2 \omega_a)^{\frac{1}{3}},$$

of the oscillation frequencies. In our labs we found a good value to start forced rf-evaporation is $\bar{\omega} \approx 2\pi \times 300$ Hz.

For the an example value of $I_Z = 34$ A and $I_{\text{bias}} = 80$ A (this time the B_z is set to 1 mT), the oscillation frequencies of the potential have a value of $\omega_r = 2\pi \times 1043$ Hz and $\omega_a = 2\pi \times 16$ Hz, resulting in $\bar{\omega} = 2\pi \times 260$ Hz which is close enough. Again the geometric constraints on the chip dimensions as well as the current limits limit the achievable values of $\bar{\omega}$.

2.2.5 Cloud Size

To calculate the size of the cloud we first need to understand what we mean by giving a diffuse vapour of atoms a hard property like size. In the limit of a classical thermal cloud of atoms, i.e. the phase density is smaller than one, $\rho_{\text{psd}} = n\lambda_{\text{dB}}^3 < 1$, the size is given by the rms radius S of the Gaussian density profile of the cloud,

$$S_k = \sqrt{\frac{2k_B T}{m_{\text{Rb}} \omega_k^2}}.$$

In this case the size of the cloud depends on the temperature of the cloud or in other words the energy per particle. Adding more atoms at the same temperature does not increase the size of the cloud.

Once the phase-space density rises above a certain point specified by $\rho_{\text{psd}} \approx 2.612$ [61] the cloud Bose-Einstein condenses. The density of the

cloud in the Thomas-Fermi limit⁵ takes on a parabolic profile. The size of a Thomas-Fermi condensate is given by the half size s of this profile,

$$s_k = \sqrt{\frac{2\mu}{m_{\text{Rb}}\omega_k^2}}.$$

In this case the size depends on the mean field energy or chemical potential μ of the cloud. Here the size of the cloud increase if we add one more atom, since μ depends on the number of atoms [73],

$$\mu = \frac{\hbar\bar{\omega}}{2} \left(\frac{15Na}{a_{\text{ho}}} \right)^{\frac{2}{5}},$$

where

$$a_{\text{ho}} = \sqrt{\frac{\hbar}{m_{\text{Rb}}\bar{\omega}}} \quad (2.1)$$

is the harmonic oscillator length and a the boson-boson scattering length⁶. With these relations the Thomas-Fermi size of a condensate can be expressed as

$$s_k = a_{\text{ho}} \left(\frac{15Na}{a_{\text{ho}}} \right)^{\frac{1}{5}} \cdot \frac{\bar{\omega}}{\omega_k}.$$

For high confinements in the radial direction, the kinetic energy of the atoms in that dimension may exceed the mean field energy of the condensate. In this case the condensate can be approximately treated in the limit of an ideal gas. Its rms size becomes the harmonic oscillator length from equation 2.1.

Since the sizes are dependent on the temperature and number of atoms in the cloud, and because they are so closely related to the oscillation frequencies, no map of cloud sizes is plotted.

2.2.6 Trap Depth

The trap depth is a parameter that describes how hot a sample can be trapped in the magnetic field. For a trap composed of a wire and a homogeneous bias field the trap depth is usually simply the difference between total background bias field and trap bottom field. In our case, the bias field is formed by wires close to the structure, hence the trap is surrounded

⁵This limit can be independently reached for the three different dimensions of the cloud. For each dimension the condition $\mu \gg \hbar\omega_k$ must be fulfilled.

⁶See for example [74] for scattering length values for different states. Let us use as a close approximation to the real value $a = 100 \cdot a_0 \approx 5.3$ nm.

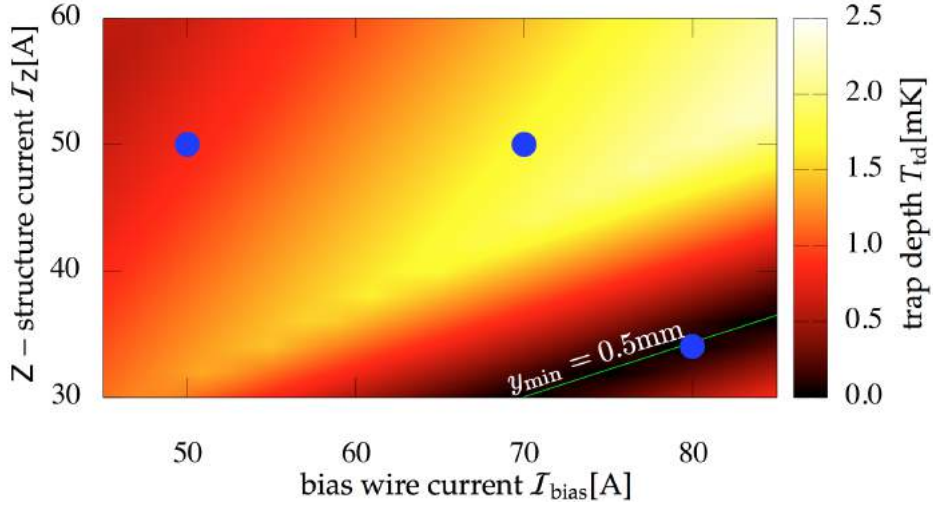


Figure 2.7: Plot of the depth of the trap given in units of temperature. The line denotes the waveguide chip surface. Without a chip the trap depth is limited by the bias field, however, closing in to the chip opens up a second escape route for the atoms by hitting the chip's surface. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.

by a local maximum field barrier in all directions instead of a height field plateau. There are two possible escape routes to drain the trap if the kinetic energy is increased or the potential energy barrier is lowered. First, the saddle point of weakest field in the surrounding barrier, and second, the point where the trap meets the surface⁷ of the waveguide chip at a distance of 0.5 mm.

As mentioned before, the trap centre lies on a line going straight away from the surface in the y -direction and it turns out that for this drain path we take the maximum of the total field from the centre of the trap along y ,

$$B_{\text{td},1} = \max(|\mathbf{B}(x_{\min}, y, z_{\min})|, y_{\min} < y < \infty).$$

For the surface escape route we take the value of the field at the distance of the chip just underneath the trap,

$$B_{\text{td},2} = |\mathbf{B}(x_{\min}, y = 0.5 \text{ mm}, z_{\min})|.$$

⁷To be precise, close to the surface ($< 1 \mu\text{m}$) the attractive Casimir-Polder force lowers the potential and therefore the trap depth. The effect, however, is the same. The atoms come in contact with the surface at room temperature and are expelled from the trap.

The trap depth is now just the lower of these two values minus the value of the trap bottom,

$$B_{\text{td}} = \min(B_{\text{td},1}, B_{\text{td},2}) - B_{\text{min}}.$$

The trap depth should be big compared to the mean atomic energy. A typical value is [10]

$$\mu_B B_{\text{td}} > \eta k_B T$$

with $\eta \approx 6$ to keep the trap loss negligible.

In figure 2.7, T_{td} is plotted with the line denoting the location of the waveguide chip surface. It is clearly visible that the trap depth, though it should be increasing closer to the Z-structure, is actually decreasing and vanishing at $y_{\text{min}} = 0$. This is due to the second escape route opening up for the atoms.

2.3 Constraints

The geometry of the current structure has emerged around several constraints. The most prominent one is the distance of the trench of the waveguide chip from the Z-structure, i.e. about 500 μm . Due to this distance using high⁸ currents in the Z-structure cancelled by a large bias field was the only available choice to achieve oscillation frequencies high enough for evaporative cooling and fitting the cloud into the trench.

2.3.1 Current Limit

For the calculation, we can of course specify any value of current that we want to gain the required results. However, in reality the working conditions of our atom chip depend on its temperature not exceeding a certain limit. The thing to note here is that the currents we can apply without endangering the performance of the system are about the same as the maximum current output of the power supplies presently in use:

$$I_Z^{\text{max}} = 50 \text{ A}, \quad I_{\text{bias}}^{\text{max}} = 80 \text{ A}, \quad \text{and} \quad I_{\text{end}}^{\text{max}} = 30 \text{ A}$$

For this reason, we concern ourselves in this chapter primarily with this regime.

⁸High, meaning higher than currents typically used in experiments where the cloud is trapped close to the current carrying structure.




Figure 2.8: The trap parameters: trap depth and relative displacement when adding a current in the end wires of the current structure. Current values for I_Z and I_{bias} , as well as B_z values set to control the trap bottom, are taken from the comparison table 2.1.

2.3.2 Size of the Waveguide Chip Trench

Taking the values from table 2.1 that a trap prepared to go into the trench is about $2s_r = 1.4 \mu\text{m}$ in diameter and $2s_a = 90 \mu\text{m}$ long. Since the trench is $16 \mu\text{m} \times 500 \mu\text{m}$ in size, a straight cloud should fit in fine. However, there is the issue of rotation. The ends of the clouds have a relative displacement or $\tan \vartheta = 0.056$. In other words, the tips of the cloud are about $2.5 \mu\text{m}$ off-centre, if the cloud is lying in the middle of the trench. Hence, for these typical parameters, the cloud fits in the trench.

2.4 Tweaking the Trap

In section 2.1 we have explored the magnetic field potential formed by a simple Z-structure and two bias wires. In section 2.3 we investigated the problems that emerge from the constraints imposed upon the structure by the presence, or more precisely the thickness, of the waveguide chip. Now we take a look at how one additional control knob, namely the additional end wire structure, can tweak the problematic parameters of the trap to correct for some of the effects encountered in the last sections.

By a running current I_{end} through the end wire structures (for simplicity we assume an equal current through both sides: $I_{\text{end},1} = I_{\text{end},2} = I_{\text{end}}$) we effectively only adjust the two trap parameters: depth and rotational displacement. By adjusting B_z , we hold the oscillation frequencies very close

to constant. Figure 2.8 shows plots of the changes in trap depth and relative displacement for the three parameter sets from table 2.1. The rotation of the cloud that would prohibit the cloud from fitting into the trench can be effectively reduced. Furthermore by adding to or cancelling the z-component of the magnetic field, the trap depth can be slightly tweaked.

Considerations

Although this chapter gives a rather detailed description of the magnetic field that would form the trap to load cold atoms into and cool them down to BEC, this weird state of nature was never achieved during the time of the thesis or until today. This is mostly due to vacuum issues (which are described in more detail in the next chapter). Because of this and because the waveguide chip is the main novel part in the setup, it was decided to rather characterise the chip with cold atoms, which were available in form of a magneto-optical trap, and aim for evaporative cooling and BEC at a later stage when the pressure was under better control.

The next chapter describes the experimental setup that was mostly⁹ build from scratch to accommodate the atom chip.

⁹And apart from the laser system.

Chapter 3

Experimental Setup

The major part of the machinery existed in the form of an experiment to investigate ultra-cold atom physics long before pursuing the waveguide chip setup, and most of it has been described in [75] or more recently in [76]. A brief overview of the components used is given but emphasis is put on the profound changes.

Due to the magnetic requirements discussed in the last chapter, this thesis uses a new atom chip for trapping and manipulating atoms. Although the atom chip is mentioned in the context of the new vacuum system, the chip is discussed in greater detail not here but in the next chapter.

3.1 Lasers

The laser system is used to trap, manipulate and detect rubidium 87 atoms. Trapping is done in a magneto-optical trap (MOT), whereas manipulation includes cooling and optical pumping. Detection is done via absorption imaging, scattering of resonant and off-resonant light. For the most part the system has been described in much detail in [75,76]. A brief overview mentioning a few changes made is given here, nonetheless.

There are three laser sources mainly used for reference, cooling and repumping.

3.1.1 Reference Laser

This laser produces light at a wavelength of 780 nm, corresponding to the rubidium 87 D_2 transition between the $5^2S_{1/2}$, $F = 2$ ground state and the

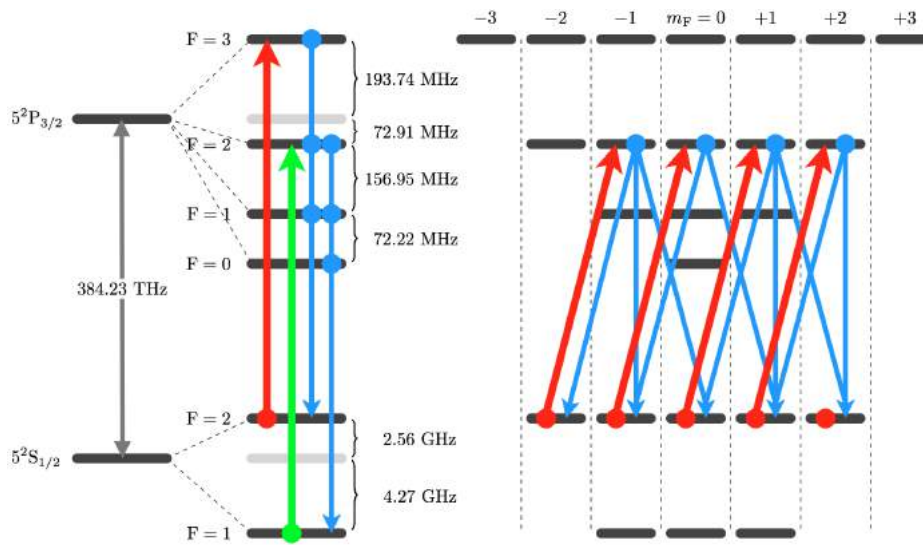


Figure 3.1: Hyperfine structure and transitions of the ground state ($5^2S_{1/2}$) and D_2 excited state ($5^2P_{3/2}$) of rubidium 87. (*left*) the red and green arrows denote the reference and re-pumping transitions, respectively. Blue arrows show allowed radiative decay routes. (*right*) the red arrows denote the σ^+ optical pumping light increasing m_F by 1 and the allowed radiative decay routes of each excited state as blue arrows. Since there is one state without an excitation arrow or a decay path, atoms are pumped into the $F = 2, m_F = 2$ state. Polarisation-resolved² repumper transitions, though not shown on the right, bring atoms which have through spontaneous emission decayed into the ground state back into the $5^2P_{3/2}$ $F = 2$ excited state.

$5^2P_{3/2}$, $F = 3$ excited state. The transition is drawn in the context of the energy structure of rubidium 87 in figure 3.1.

The reference laser is a home built extended-cavity grating-diode laser (ECDL) [78] in Littrow configuration to bring the linewidth below 1 MHz, compared to 6 MHz [79] natural linewidth of the addressed transition. The laser is locked to the aforementioned transition via polarisation spectroscopy on a rubidium vapour cell [77]. In this setup, the laser has a typical light output power of 16 mW using 54 mA drive current. The light output by this laser is divided into beams for the spectroscopy, to lock the cooling

²The repump laser is locked to a polarisation spectrometer which uses the light-induced birefringence and dichroism of rubidium vapour to increase the signal-to-background ratio [77].

laser, and for optical pumping.

3.1.2 Cooling & Detection Laser

A commercially available³ ECDL tapered-amplifier laser system is used to gain high light output power for an LVIS MOT [80] and science MOT cooling beams. An LVIS or low-velocity intense source is in practice a MOT with a small hole in one of the six mirrors forming the MOT. Thought this hole, atoms can escape from the MOT. The depth of the hole assures that only transversely cold atoms pass through the hole and into the science chamber where they are captured by a full MOT. In addition, it provides light for absorption imaging and in-waveguide-chip detection of the atomic samples.

The laser light can be locked 3 to 110 MHz red detuned from the light emitted by the reference laser. The type of the lock is an offset lock to the reference laser, which employs filter characteristics [81] in contrast to the interference lock described in [76]. In this setup, the laser has a typical output power of 490 mW at 94 mA diode and 2000 mA TA drive current with a linewidth of less than 1 MHz.

The power of the light is split by a ratio of 3:1 into two acousto-optical modulators (AOM) which shift the frequency of the light back 93 MHz. This results in light in the -1 st orders which have for the MOT phase a total red detuning from resonance of about 12 MHz (twice the natural linewidth), i.e. beat at 79 MHz. The range of accessible detuning values is +17 to -90 MHz.

The shifted higher power beam is sent through an optical fibre for stability, then superimposed with repump light and subsequently split into 3 beams which are used for cooling atoms in the LVIS MOT. The shifted lower power beam is sent through an optical fibre for stability, then superimposed with repump light and subsequently split into 6 beams which are used for cooling atoms in the science MOT. The unshifted 0th order of the higher power beam is sent through another AOM that shifts the light back to resonance. This light is split into three beams, which are in turn coupled into optical fibres for distribution around the chamber and pointing stability⁴.

³TOPTICA Photonics DL100 with LD785-0080 and TA100 system TA-0780, both from 2002.

⁴Here, pointing stability is traded for intensity stability. Coupling the beam into a fibre assures that the exiting beam is always pointing in one direction. However, slight directional variations on the entry side of the fibre will translate into a varying intensity through the

One is coupled back into fibres connected to the waveguide chip to detect atoms within the trench of the chip. The other two beams are used for absorption imaging along two axes.

3.1.3 Repump Laser

This laser serves the transition $5^2S_{1/2}, F = 1$ ground state to $5^2P_{3/2}, F = 2$ excited state (see figure 3.1) to repopulate $5^2S_{1/2}, F = 2$ ground state via radiative decay from the $5^2P_{3/2}, F = 2$ excited state⁵. The repump laser is a home built ECDL, same as the reference laser system. To increase the light output power this master module seeds a slave laser diode. The laser is locked via polarisation spectroscopy on a rubidium vapour cell to a dispersive feature about 80 MHz blue detuned from the aforementioned transition. In this setup, the laser has a typical output power of 34 mW at 69 mA master and 79 mA slave diode drive current with a linewidth of less than 1 MHz. The light is sent through a AOM to be shifted back 80 MHz to the repump transition, then divided to be superimposed with cooling light.

3.1.4 Optical Pumping

A part of the light coming from the reference laser is used for optical pumping by driving the transition from the $5^2S_{1/2}, F = 2$ ground state to $5^2P_{3/2}, F = 2$ excited state (see figure 3.1). The frequency is reduced by 246 MHz by sending the light through a double pass AOM. This leaves the light about 20 MHz blue detuned from the optical pumping transition. The optical pumping light is introduced into the system along the path of the absorption imaging. In this setup, the light power used is about 0.4 mW.

3.2 Vacuum

After careful consideration, we decided to not re-use the chamber that was used with the laser system [76]. Due to the geometry of the access flange and the minimum size requirements of the new atom chip, there was no way of bringing it into the chamber without developing a complicated mount to

fibre. In this case the stability in direction is more important than the stability in intensity, hence the trade-off.

⁵Note that the selection rules are $\Delta F = 0, \pm 1$ and $m_F = 0 \rightarrow m_F = 0$ forbidden for $\Delta F = 0$ [79].




Figure 3.2: Photograph of the chamber sitting on the table. Into the table a hole is cut to allow optical access to the chamber from below. In the background, on the left are the ion and getter pumps, as well as an ion pressure gauge. On the right sits the LVIS, separable from the main chamber via a gate valve. So far, the atom chip flange is not attached and replaced by a blind flange for initial hot baking.

thread the chip and the attached optical fibres through the flange without damaging them. Hence, although it is good practise to re-use successfully employed equipment (especially vacuum components), a new chamber had to be built around the requirements of the new atom chip and experiments utilising it. A picture of the nearly assembled vacuum chamber is shown in the photograph in figure 3.2.

3.2.1 Requirements & Design

The new vacuum system uses as its central part a spherical octagon chamber⁶ with eight 2.75" ports around the side, and two 8" ports on top and bottom. The atom chip, including waveguide chip, and feedthroughs were mounted on an 8" custom-made stainless steel blanking flange that can fit onto the top port of the chamber.

⁶Kimball Physics Inc. MCF800-SO2000800, material is non-magnetic, stainless steel type 316L.

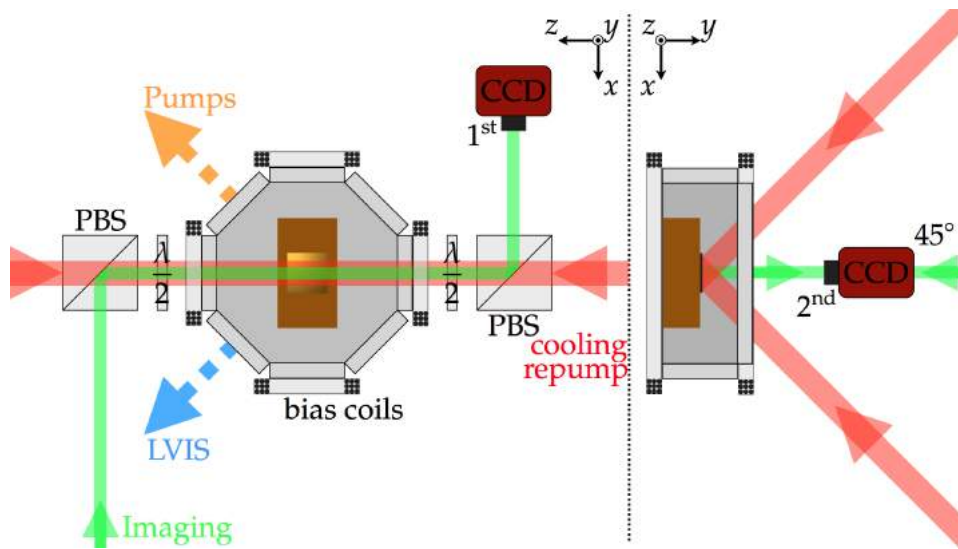


Figure 3.3: Sketch of the chamber seen from below (*left*) and the side (*right*). Except for the ports that lead to the LVIS and pumps, the sides of the chamber are closed with viewports. Cooling light is drawn in red, imaging light in green. Magnetic bias field coils are placed along the x , y and z axes. The first imaging beam traverses the chamber and is shone onto a CCD camera. The second imaging beam is goes through the bottom window and is reflected at 45° (not visible in this projection) from the waveguide chip before detection on a second CCD camera.

As discussed in chapter 2, the waveguide chip is 0.5 mm thick, so electric currents up to 80 A are needed for magnetic trapping. The high currents creating the fields, in turn, require not only thicker electrical feedthroughs than usual for atom chips but also a way to quickly dissipate the heat stemming from the intrinsic and contact resistances of the atom chip under current flow. To this end, the atom chip is mounted close to the top flange onto a block of copper, which transports the heat from the atom chip into the stainless steel flange. In addition, the contacts to the feedthroughs are distant from the central area of the chip.

To create a mirror-MOT with the waveguide chip as the mirror the quadrupole coils must be aligned on an axis at 45° to the chip plane. As a compromise, the coils were mounted inside the vacuum chamber. This way, they require a smaller current and don't obstruct any external features of the chamber or block optical access.

Figure 3.3 shows a sketch of the chamber. Cooling light (red in the figure) for a MOT comes in through two side view ports and at an angle of 45° from below through the bottom view port. To image atoms in a MOT, in or released from a magnetic trap, imaging light (green in the figure) is shone along the axis of two cooling light beams. The imaging light and the cooling light are superimposed on a polarising beam splitter (PBS) cube. After passing through the chamber, the beams are separated in the same way and the imaging light is collected with a CCD camera. A second imaging system shines imaging light at an angle of 45° from below through the bottom view port. After passing through the atoms the beam is reflected from the waveguide chip surface. A part of the light is again absorbed by the atoms and shone on a second CCD camera.

The entire chamber is sat on a large slab of aluminium with custom holes to optically access the chamber from below and mount optics on top of it in a way more stable than mounting them on long posts. The slab is supported at a height of 350 mm by four 120 mm thick legs. The dimensions of this extra optical table was chosen to minimise vibrations. A photograph of the chamber sitting on top of this slab is shown in figure 3.2.

Atom Chip Flange

The top flange, which holds the atom chip, also holds the electrical and optical feedthroughs to supply the atom chip, as well as dispensers⁷. A sketch of the layout of the atom chip flange and a side cross sectional view are shown in figure 3.4.

In the centre of the flange the atom chip is screwed down. The high-current feedthroughs are placed as close as possible to the atom chip and the coils. Two 1.33" half nipple ports are located at opposite sides of the flange to hold low current feedthroughs. Two additional 2.75" half nipple ports

⁷Dispensers are small folded sheet metal pockets containing a, in this case, rubidium salt (usually RbCl or RbF) and metallic Calcium powder. Sending a current through this pocket heats up the mixture. A chemical reaction now releases mono-atomic rubidium (72% ^{85}Rb , 28% ^{87}Rb , traces of other isotopes) through the slits in the pockets into the vacuum chamber. CaF_2 or CaCl_2 remains.

During the last opening of the chamber in March 2009 the dispensers were removed from the science chamber, to reduce a source of background pressure. The strict geometry of the internal quadrupole coils, together with the possibility of aligning the LVIS push beam via the opposite view port on the chamber, allows for finding a MOT without the need of an atom source inside the science chamber. The port is now closed with a blanking flange.




Figure 3.4: Layout of the atom chip flange that sits on top of the chamber. It features electrical feedthroughs for the quadrupole coils (a) that form the MOT, the end wire structures (b), the H-structure (partly covered by the waveguide chip) (c), the atom chip bias field wires (d), dispensers (only the feedthroughs shown here) (e). In addition, there are optical feedthroughs (f) to bring light to the waveguide chip. The side cross section view of the flange also shows that one of the quadrupole coils is embedded into the copper block (g), another is mounted on an aluminium support (h).

are placed at opposite sides of the flange to hold optical feedthroughs. Across the in-vacuum side of the surface of the flange, size M3 holes are tapped to hold support clips for fibres.

One of the quadrupole coil formers is embedded inside the copper block

of the atom chip. The former of the other coil is mounted to one side of the atom chip and is sticking out just enough to not touch the glass of the view port. The coils are separated to place their on-axis centre point 1 mm distant from the centre point on the waveguide chip surface⁸.

3.2.2 Installation

Prior to the installation of the new chamber the old was removed from the optical table. Subsequently, the LVIS section, pumps, gauges and view ports were removed to be re-used.

Chamber Preparation

In preparation of the installation of the chamber on the optical table the chamber system (including pumps and gauges) was baked at about 250° for 7 days. This was done under vacuum with the chamber ports closed but with blanking flanges instead of view ports, and connected to a turbo pump. After cool-down the view ports were attached and the pump system was removed. The atom chip flange had not yet been attached at this point.

Atom Chip Flange Preparation

The holes and pockets of the atom chip flange were all machined in house, as well as the welding of the half nipple ports and feedthroughs. The atom chip⁹ was attached and connected to the current feedthroughs. Three dispensers were attached to a 6-pin feed-through so that they would fit through one of the 1.33" ports. Subsequently, the waveguide chip and fibres were attached. Finally, thermocouples were attached to prominent points on the atom chip. A photograph of the assembled flange with atom chip and connections is shown in figure 3.5.

Optical Table Preparation

For the raised optical table, a 100 mm thick slab of aluminium was water-cut to open up holes for optical access to the LVIS and science chamber, as well

⁸Since the two coils have not the same number of windings (47 and 48), the magnetic field zero at equal current running through the coils is slightly pushed towards the chip surface.

⁹For details on the fabrication of the atom chip, attaching of the waveguide chip, and connecting fibres see chapter 4.




Figure 3.5: Photograph of the atom chip flange with atom chip and connections.

as bring beams up to the raised level. 10 mm thick plates of aluminium were patterned with a tapped holes in the fashion of optical tables, then anodized. Four 120 mm thick 350 mm long aluminium legs to hold the table were machined.

Bringing It All Together

After removing all optics from the old setup, the aluminium slab was placed on its legs on the main optical table. The patterned plates were screwed onto the top surface. The bottom bias coil was placed into the centre hole of the table, subsequently, placing the science chamber onto the coil. Following this, the LVIS was placed on the table and using positional adjustment aids, connected to the science chamber. The same was then done with the pump system. In this state, with a blanking flange instead of the atom chip flange in place and connected to a turbo pump, the chamber was baked at about 160° for 7 days.

To attach the atom chip flange, the blanking flange was removed and 4 long threaded rods were screwed into the top port of the science chamber. On nuts the atom chip flange was then carefully lowered down into the chamber, making sure the fibres were not destroyed. A photograph of this




Figure 3.6: Photograph of the installation of the atom chip flange. The flange is carefully lowered onto and the atom chip assembly into the octagonal vacuum chamber.

process is shown in figure 3.6.

After closing the flange, the remaining bias coils were put in their place, and the chamber was baked at about 130° for 18 days to bring the pressure further down. After a thorough leak-check¹⁰ the turbo pump was removed. A photograph of the final experimental setup with all optics and electrical connections installed is shown in figure 3.7.

3.3 Detection

We can detect atoms with two different detectors. First a camera, thus creating an image of the atoms, and second an avalanche photo diode connected to the waveguide chip to count photons scattered by a atoms within the trench.

¹⁰A small leak was found on one of the fibre feedthroughs.




Figure 3.7: Photograph of the experiment now. Underneath the table are some of the MOT optics already present. On the table, the camera can be seen. The atom chip flange is already in place.

3.3.1 Imaging

The two implemented absorption imaging setups are shown in figure 3.3. Two CCD sensor cameras¹¹, called Pike and Marlin, are available and they are placed depending on the requirements of the current measurement. The Pike uses a resolution of 1388×1038 pixels and can be used if required in 4x (hv) binning mode at a resolution of 692×518 with a 14 bit ADC (10 bit real SNR). Each of the sensor pixels has a size of $6.45 \mu\text{m} \times 6.45 \mu\text{m}$. The Marlin uses a resolution of 656×494 pixels without binning with an 8 bit ADC (7 bit real SNR). Each pixel has a size of $9.9 \mu\text{m} \times 9.9 \mu\text{m}$.

In the setup, which images horizontally along the z-direction, resonant light is collimated out of a fibre and superimposed with the path of the cooling light along this axis. Due to the aperture of the waveplate mounts and the PBS cubes in the path of the beam, its diameter is limited to about 11 mm. After the chamber, the imaging light is branched off from the cooling light and focussed onto a camera with a commercial macro lens. The magnification in this setup is between 0.35 and 0.39. In other words we look at an area 2.6 to 2.9 larger than the size of the camera sensor. The horizontal

¹¹Allied Vision Technologies (AVT) Pike F-145B [82] and Marlin F-033B [83].




Figure 3.8: Detection setup. The light is split into parts and injected into the fibres (*left*) connected to the waveguide chip (*top*). After traversing the trench, the light coming out of one fibre is collimated, then focussed onto an SPAPD (*right*). The second imaging axis consists of imaging light that is shone onto the atoms at an angle of 45° . The reflected light is focussed onto a CCD camera (*middle*).

imaging was tried along the x -axis instead of the z -axis. However, since the fibre reinforcements (c.f. chapter 4) obstruct optical access close to the chip surface, this axis was abandoned. The power of the horizontal imaging beam is typically $100 \mu\text{W}$ and a typical exposure is $100 \mu\text{s}$ long. A typical experimental sequence to image along this axis takes two to four exposures. One of which is a fluorescence image of the atoms in the MOT. Two are exposures taken of the imaging light beam, one with and one without atoms. A fourth exposure may image the background light. The latter can be used to calculate a background corrected absorption image of the atomic sample. However, for simplicity we now usually omit to take the background and fluorescence image.

The second setup takes an exposure through reflecting light off the waveguide chip. Due to the aperture of the telescope the diameter of the beam is limited to about 11 mm. The imaging light is reflected off the gold coating of the waveguide chip and passes through the atom cloud. The light that first passes through the atoms before it is being reflected is also imaged,

resulting in two shadows of the cloud appearing in the image. Using two achromatic lenses each with a focal length of $f' = 200$ mm mounted directly onto the camera, the magnification in this setup is 1 to get in close to the cloud (c.f. figure 3.8). The power of the horizontal imaging beam is $100 \mu\text{W}$ to $150 \mu\text{W}$ and a typical exposure is $100 \mu\text{s}$ long. A typical experimental sequence to image along this axis takes two exposures of the imaging light, one with atoms and one without atoms.

3.3.2 Waveguide Chip Detection

For the detection of light coming from the waveguide chip fibres, a fibre-coupled single-photon counting avalanche photo diode¹² (SPAPD) is used in conjunction with a 780 nm notch pass filter¹³. The SPAPD is not a commercially fibre-coupled device. The coupling has been done by the group. The SPAPD has a photon detection efficiency (quantum efficiency) at 780 nm of 50% and a detector dead time¹⁴ of $\tau_{\text{dead}} = 32$ ns. The dark count is 62 c / s (counts / second). It is fitted into a black box and covered entirely with black cloth which reduces the measured ambient count rate due to stray light to typically 130 c / s, even with the room lights on.

Upper power limits of interest are the following. The detector is a linear detector up to a point. Although we can calculate the number of photons on the detector from the number of recorded counts, the saturation effects reduce the sensitivity of the measurement. The sensitivity is reduced by a factor of 2 at a count rate of $0.293/\tau_{\text{dead}} = 9.2$ Mc / s or 4.7 pW on the detector. The next limit is detector paralysis due to the count rate reaching $1/\tau_{\text{dead}} = 31.3$ Mc / s or 15.9 pW on the detector. This is the limit above which a measurement can no longer be corrected. And the last limit is the destruction limit of the detector. This is reached for a pulse of 10^4 photons in a time of < 1 ns. This corresponds to a constant power of $2.5 \mu\text{W}$ on the detector. Typical count rates used for measurements are between 1 and 10 Mc / s.

A signal rate of 15 c / s in a measurement of 1 s has an uncertainty of $\sqrt{130 + 15}$ c / s = 12 c / s and can thus be just distinguished against the background with an equal uncertainty of about 12 c / s. At the upper limit of the detector, 15 Mc / s in our case, the dynamic range for a measurement of

¹²Perkin Elmer SPCM-AQRH-14 [84]

¹³Semrock LL01-780-12.5

¹⁴See chapter 5 for an in-depth analysis of dead time effects and methods of correction.

1 s duration becomes 10^6 (background counts become negligible for such a high count rate). For a measurement of 1 ms duration, however, the dynamic range reduces to just 10^3 . To extend the dynamic range of light intensities accessible for short time durations without damaging the SPAPD, additional neutral density filters can be fitted into this system to use high light intensities for measurement. A sketch of the SPAPD detection setup is shown in figure 3.8. The probe light coming from an AOM, with which the probe light frequency can be varied, is split to access all fibres. Blocking paths selects which fibre to use. The light, after interacting with the atoms, is focussed onto the detector.

3.4 Computer Control

Hermes, the software to control the hardware used with the experimental setup before the waveguide chip, was a JAVA programme that allows the user to specify the state of the hardware for a certain time and concatenate such time slots into a sequence. Each state is represented by an array of values that the hardware can put out. Though it is possible to see the progression of the state of the entire system on screen in this way, a bug in *Hermes* was the probable cause of the breaking of an atom chip in one of the prior experiments [76] and has several disadvantages that severely limit its usability. To name only a few, it was very slow in creating the values for the eight analogue channels. And although written object oriented, it is difficult to include new or different hardware to the set used at the time of writing. Because of these and other limitations, it was decided to write a new software suite for hardware control based on Microsoft's .NET framework, reusing a fair amount of the code base of the groups electron-dipole moment experiment [85]. The name of the suit is *Tyche*¹⁵.

3.4.1 Architecture

Tyche features a modular architecture to split up the tasks of hardware control and user interaction. In other words, for each task there is a single module, a single programme. The structure is schematically shown in figure 3.9. The user interaction is handled by the *Tyche* user interface (*TycheUI*). It

¹⁵Ancient Greek goddess of fortune and prosperity, daughter of *Hermes*. Lately also an acronym for 'Timing your control hardware events'.




Figure 3.9: Data flow chart of the experimental control. (*blue arrows*) the user interacts with the user interface as part of the software suite. The user interface serializes the input and hands it down to the main hardware controller, which in turn delegates tasks to several slave data acquisition controllers. All controllers then ready their hardware for the next experimental cycle. (*red arrows*) During the cycle, data is collected by the hardware, then processed by the slave controllers and passed on up through the chain back to the user interface to be accessed by the user. (*yellow arrows*) On request, all software modules save data and meta data to the computer hard drive.

processes the input and relays it to the main hardware controller (TycheController) which in turn programmes the analogue and digital output cards, as well as relay tasks from TycheUI to data acquisition controllers (DAC). All modules communicate with each other via .NET object publishing and individually save data to a hard drive.

A picture of the computer desktop whilst measuring is shown in figure 3.10 where the front-ends all of the modules are shown.

3.4.2 TycheUI

The Tyche user interface is a command line interface (CLI), as opposed to the graphical user interfaces typically used in cold atom experiments [86–88]. It is written in IronPython. Apart from direct and immediate control of each

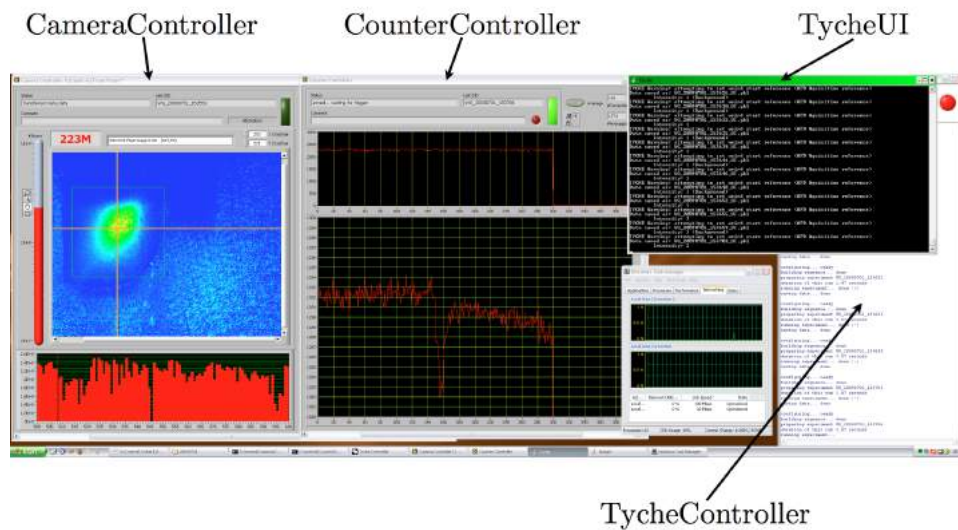


Figure 3.10: A screen grab of a typical window layout for performing experiments with Tyche. Two data acquisition controllers are shown on the left. One is for the camera, showing an absorption image of an atomic cloud. An approximate atom number is calculated on-the-fly (223 million in this case). The other DAC is for the counter input card connected to the SPAPD. Each shot is displayed on top and on-the-fly averaging of a number of shots is performed below. TycheUI is just a command prompt executing a loop in this grab. TycheController is really a backend to connect all other modules. It only has a window for status information output.

part of the experiment, control commands can be collated into so-called Events (which internally are objects). An Event includes commands which mark edges, i.e. changes in values. Thus it defines the behaviour of a part of the hardware from a starting time to an ending time (the next edge). This structure is much more generic than the time slot structure, i.e. any time slot can be written as one event, but not every event can be written as one time slot. On top of that, events can not only contain hardware behaviours, but also other events. Events on the same level are then related to each other by specifying the positions of their starting and ending times. An entire experimental sequence is eventually only one event containing many subevents. On execution, the relative timings of the event and branched sub-events are flattened out to be correctly interpreted by the sequentially working TycheController.

Apart from running an event with its default parameters, as they were specified on creation, parameters can be varied on execution. Let us look at an event that loads atoms into a MOT and then takes an image. We would write in the TycheUI command line

```
>>> go(mot_image)
```

which would run the experiment and take an absorption image of the atoms in the MOT. The Event `mot_image` contains the sub-Events:

- `initialise_all` initialises all hardware channels to known start values
- `init_cam` initialises the camera, i.e. resets buffers and makes sure the camera is connected
- `load_mot` switches magnetic fields and light shutters to form a MOT
- `molasses` switches magnetic fields and light detuning
- `abs_image` switches AOM to take an absorption image
- `fluo_with_atoms` takes an image
- `fluo_no_atoms` switches off all light and takes image

Each of these sub-Events specifies parameterised commands sent to the hardware. If we want to change one parameter of this sequence, e.g. the light detuning, we would simply change the line to execute to something like

```
>>> go(mot_image, parameters={"detuning":90})
```

One can guess at the raw potential of this approach. To take a number of images to scan over the detuning we only expand this to read

```
>>> for det in range(70,110,2):  
...     go(mot_image, pars={"detuning":det}, saving=True)
```

and would get 20 images of the MOT without hitting another key. This is only the beginning. With the scripting power of Python we assembled meta sequences that in addition to the execution also analyse the data on the fly to greatly simplify often needed measurements, like calculating the temperature in a magnetic trap by time-of-flight analysis, with one line of text on the command line.

TycheUI saves each executed event object to the hard drive and sends a set of meta data about the execution to be collated by TycheController.

3.4.3 TycheController

The Tyche main hardware controller is written in C# using National Instruments Measurement Studio libraries and is at the heart of the suite. It generates a shared experimental cycle identifier (EID) and distributes it throughout all active modules in the suite. The EID takes the form

```
ID_YYYYMMDD_HHMMSS,
```

with ID being an identifier for the experiment ('WG' in our case) and the rest notifying the year, month, day, hour, minute, and second of the start of the cycle.

TycheController receives commands from TycheUI. These can be immediate commands to change the state of the experiment, or timed commands in the form: 'change *this* output to *this* value at *this* time'. TycheController passes immediate commands straight away on to the hardware. Timed commands for the internal hardware, it saves in a time ordered layout. On execution, it interpolates this layout for each hardware channel into a byte stream. Commands for DACs are passed on to the corresponding controllers. All hardware control modules are synchronised by dedicated digital output channel.

TycheController saves the layout and meta data it receives from TycheUI and DACs to the hard drive.

3.4.4 Data Acquisition Controllers

Data acquisition controllers can be written in any programming language that can publish an object or be accessed via a programming language that can publish an object. In our case, each module is created in National Instruments LabVIEW. A small C# server connects to this module via the COM interface standard.

DACs control a dedicated bit of hardware, especially but not exclusively for acquiring data. This can be images from a camera, analogue voltages or a number from a counter card. Furthermore, they can act in a more abstract way, e.g. offering a text field to enter a note to each experimental sequence or manually enter data for which there is no direct link to the computer (like a current or temperature setting, or a light power).

Since each DAC is programmed as a separate module, it is left to the programmers discretion whether to include on the fly data presentation

and/or analysis. However, we included a mandatory interface protocol of method calls that includes receiving the EID, configuring and arming the device before the sequence, saving data, and relaying meta data to TycheController.

Let us take a look at two examples of DACs used in the experiment .

Camera Controller

The camera controller is a LabVIEW programme that communicates with a IEEE-1394 (FireWire) camera (Pike or Marlin). It receives configuration commands from TycheUI via TycheController and is armed for image acquisition by TycheController when running an experimental sequence. During the sequence, the camera is triggered by a digital output channel. After the sequence, the data is read-out of the camera buffer, together with the acquisition properties that were set. The images are processed and saved as a flexible image transport system (FITS) file to the hard drive. Depending on the subject of the measurement, different profiles can be fitted to the resulting image to calculate parameters like number of atoms or cloud size. The results are displayed within the graphical user interface of the programme.

Counter Controller

The APD that is connected to the waveguide chip fibres sends its signals to the counter input channel on a National Instruments PCI-6024E Multifunction Card. The counter controller is a LabVIEW programme. It configures and arms the counter input channel for acquisition. During the sequence, the counter input channel is triggered by a digital output channel. After the sequence, the data is read-out of the card buffer. The data is processed and saved as a comma separated values (CSV) file to the hard drive. The module can average the current acquisition with prior ones to approximate a running average of multiple experimental cycles.

3.4.5 Advantages & Disadvantages

The modular structure of the code base has several advantages. For one, nobody has to know the code of all of the modules to add a module to the suite. Furthermore, it is possible to simply adjust an already existing stand-alone hardware controller to work as a module within the suite. This greatly decreases development time for new modules.

The downsides to this approach are that indeed nobody knows the entire code. Hence, understanding unpredicted behaviour of the suite requires someone who understands the interconnection points of all of the modules. What is more, due to possible high number of modules that have to be managed by the operating system, running other programmes in parallel may influence the performance of execution. Finally, although the state of each experimental sequence is saved, changing modules may compromise the future-compatibility of scripts and events.

Nevertheless, the aforementioned advantages outweigh the disadvantages by far.

Considerations

Apart from the laser system, most parts of the experiment had to be designed, calculated, built, and tested from scratch, devouring a huge amount of time which as we will see in the following chapters left little time to set up experiments in a way that provides really new data.

However, a new experimental apparatus was realised to hold the new type of atom chip, which is described in more detail in the next chapter. Furthermore, a new and versatile control programme was developed, written and tested to automate the taking and analysis of data and generally make more time available to think of experiments and result, rather than take and analyse it. Unfortunately, the least bit of time was spent during this thesis with the application of it.

The next chapter describes the mechanics of the atom chip including the waveguide-chip, how it was built and characterised.

Chapter 4

The Sub-Chip: Manufacturing Process and Characterisation

Now that we have looked at the experimental infrastructure in the last chapter we shall have a closer, more detailed look at the centre piece of the experiment: *the chip*.

I explain briefly how the waveguide chip is manufactured, and describe in detail the process of building and assembling the sub-chip and wedding it to the waveguide chip. I present the results of several tests with the aim to judge and understand the chip's performance. And finally, I briefly comment on the whole endeavour and give some ideas for improvements in the future.

4.1 Design

The chip comprises an optical-waveguide chip attached to a current-carrying atom-chip device. An impression of the chip is shown in figure 4.1. A more detailed view of the waveguide chip and the centre region of the chip was already shown in figure 1.1 (page 9). A different view of the waveguide chip is shown in figure 4.2.

The waveguide chip has 12 pairs of parallel waveguides buried in silica, 3 of which are connected to fibres that leave the vacuum chamber via 6 FC-connector optical feedthroughs. The waveguides have a centre-to-centre pitch of $10\ \mu\text{m}$ and are intersected by a $16\ \mu\text{m}$ wide trench that stretches




Figure 4.1: The fully assembled atom chip (*artist's impression*).

over $500\ \mu\text{m}$ allowing for a cold cloud of atoms to enter the region where waveguide pairs face each other.

The 3 fibres connected are numbers #4, #9 and #10 as shown in figure 4.2. The distances between these fibre pairs are 10, 50, and $60\ \mu\text{m}$. This opens up two different length scales when probing long, thin atom clouds. The mode field diameter for the waveguides is $2.1\ \mu\text{m}$.

The design of the sub-chip boils down to the following points. An H-structure is used because it is a very versatile structure that allows for wire-like magnetic trapping in Z-configuration (plus axial confinement and zero-field offset come for free), as well as magneto-optical trapping by supplying a quadrupole field when used in U-configuration. Intermediate configurations allow for more efficient loading of the magnetic trap. The gradient coils and bias wires are placed inside the vacuum chamber instead of externally to preserve optical access around the chamber and limit the necessary current to reach the required magnetic fields and gradients. Two, instead of one, bias wires allow control of the horizontal position of the trapped cloud and produce less heat. End wires are employed for axial position correction, aspect ratio control, and as an rf-antenna for evaporative cooling. The solid copper mass of the sub-chip base maximises heat dissipation and helps buffering the heat transfer to the flange.

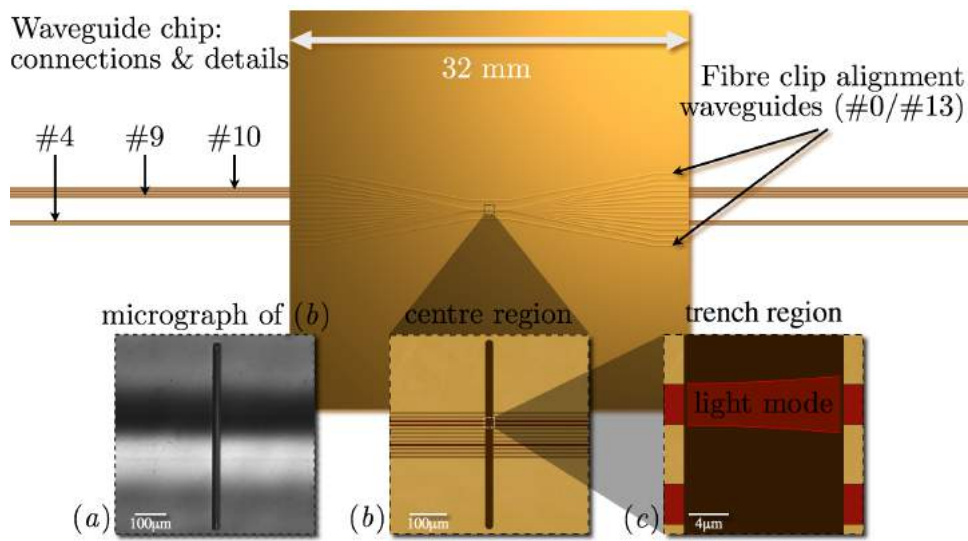


Figure 4.2: Bird's eye view of the waveguide chip. The lines hint at the buried waveguides converging in the centre. (a) DIC micrograph of the centre region. The light and dark lines show slopes on the surface. (b) Sketch of the same region with the waveguides clearly present. The red waveguides are the one that are connected to fibres. (c) A sketch of the even further enlarged trench region and inside a representation of the light mode field diameter diverging across the trench (for 780 nm).

4.2 Manufacturing Process

4.2.1 How to put waveguides into a chip

The chip was manufactured by the Centre for Integrated Photonics [89] on a 1 mm-thick silicon wafer. A 10 μm thick layer of silica, grown by thermal oxidation, formed the lower cladding for the waveguides. A second layer of silica, doped with Germanium and Boron to achieve a 0.75% refractive index contrast, was created by flame hydrolysis deposition. This was etched through a UV-lithography mask to create the waveguide cores. A further 10 μm silica layer of upper cladding, also deposited by flame hydrolysis, was doped with boron and phosphorus in order to match the refractive index of the lower cladding. A layer of 50 nm of chrome and then 100 nm of gold provides the reflecting surface needed by the MOT beams. Finally the central trench was cut by deep reactive ion etching to a depth of 22 μm

through a $16\ \mu\text{m} \times 500\ \mu\text{m}$ rectangular mask formed by UV-lithography. The finished chip was polished on the back to reduce its thickness to $500\ \mu\text{m}$. This reduced the distance between the atoms in the trench and the current-carrying wires underneath.

The chip was cut into a square shape with sides of length $32\ \text{mm}$. The square cross section size of the waveguide structure is $4\ \mu\text{m} \times 4\ \mu\text{m}$. This square size as well as the difference between the refractive indices is optimised for carrying a single light mode at a wavelength of $780\ \text{nm}$. The waveguide chip holds 12 primary and 2 secondary waveguides spanning the length of the chip. The cut trench is oriented such as to intersect the primary waveguides. At this intersection the centre to centre pitch of the primary waveguides is $10\ \mu\text{m}$. The secondary waveguides pass the centre uninterrupted just outside the gap. In both directions $1500\ \mu\text{m}$ from the gap, all waveguides fan out nearing the sides of the chip to a centre to centre pitch of $250\ \mu\text{m}$ to allow for fibre connection.

The waveguides are arranged in a plane, i.e. they form a line at the gap. This line consists of 12 facing $4\ \mu\text{m} \times 4\ \mu\text{m}$ waveguide pairs stretching over a length of $120\ \mu\text{m}$. The arrangement of the waveguides provides optical access to the gap. Although there are a total of 12 usable waveguides, in the first experimental setup that implements this chip, only 3 pairs are connected to 6 fibres. A magnified view of the chip's centre area is shown in figure 4.2.

4.2.2 How to make this atom-chip

This section contains detailed instructions to build the atom chip used in the context of this thesis. The instructions are grouped step-by-step and detailed in the figures 4.3 to 4.9. On a few general notes: all the screws used to build the chip have been vacuum relieved by either an on-axis through-hole or a wire eroder cut to the axis. Pockets liable to trap air have been relieved by small holes as well. Note that neither are shown in these instructions nor in the CAD drawings. The two component epoxy used to bond material layers is Epotek® H77. The index matching epoxy used for optical connection is Epotek® 353ND. The process of making the gradient coil formers and winding the coils as used in step (17) is not described. The entire process is to be understood more as a guideline than a perfect recipe to make this kind of atom chip.

Key to Materials and Textures

The colours and textures used in the figures try to be as close as possible to the colours and textures of the materials used in the process. This may lead to ambiguity but eventually helps better understanding of the process. This is however only practical since the different materials used offer a high enough contrast to be distinguishable. The list is in order of appearance.

- *Brown* represents the material OFHC (oxygen-free high thermal conductivity) copper.
- *Gray* (smooth) represents the material PEEK (polyetheretherketone).
- *White* may represent either the material PTFE (polytetrafluoroethylene) or a machinable aluminium nitride ceramic. Only the ceramic is used in the end product. PTFE is only used a spacer material.
- *Yellow* (translucent) represents uncured epoxy Epotek® H77.
- *Gray* (textured) represents aluminium.
- *Yellow* (dark, opaque) represents cured epoxy Epotek® H77.
- *Gold* represents the waveguide chip.
- *Green* (dark) represents Titanium.

17 Steps to Make the Chip

Step (1) We start with a solid block of OFHC copper of dimensions 90 mm × 70 mm × 19 mm and cut away material all around the block to leave a 70 mm × 50 mm platform in the centre. This is the heat sink. Changing to a small cutter, we machine a Union-Jack-type shape 0.1 mm deep out of the platform top surface. With the same tool, we cut 8 clearance holes 2 mm deep and 4 mm in diameter in the corners of the top surface. Subsequently, we drill 4 ∅ 3.2 mm through-holes around the platform and counter-sink them with 6.0 mm deep ∅ 6.0 mm holes. In addition, we drill 4 size M2 tapped holes 6 mm deep into the back-side of the block, as shown.

Step (2) We prepare a second OFHC copper (offset) block of dimensions 70 mm × 50 mm × 12 mm that features 4 protrusions that snugly fit into the clearance holes we drilled last in step (1). We join the two blocks. Now, we use a wire eroding machine with a 0.25 mm wire to cut (a) a straight fine




Figure 4.3: Manufacturing process: steps (1) and (2).

slit at $+45^\circ$, and (b) a rectangular inner pocket at -45° , as shown. These are to later on house a magnetic field coil and to reduce Eddy currents inside the bulk material stemming from the magnetic field formed by the coil. In preparation for gluing, we thoroughly degrease the block and placed it in an oven at 200°C for a few hours until the surface is mildly oxidised. *Degreasing* in this contexts means cleaning the surface for gluing. For this, consecutive ultrasonic baths of 4% Decon90 solution and deionised water are used. After drying, the surface is wiped with acetone. Note that acetone dissolves Epotek H77D!

Step (3) In this step we prepare 2 U-shaped pocket-structure brackets of polyetheretherketone (PEEK) to fit around the platform of the heat sink block, as well as several polytetrafluoroethylene (PTFE) pieces. We screw the PEEK brackets into place through the 4 holes in the heat sink block. The PTFE pieces we fit by firmly wedging them into place. At this point we remove the offset piece.

Step (4). We prepare an ample amount of epoxy and fill it into a syringe. With this we pour epoxy onto the copper and wet the entire surface with a spatula. After that, we thoroughly degreased a 2 mm thick rectangular tile of machinable aluminium nitride ceramic, wet one surface with epoxy, and place it into the alignment pocket formed by the PEEK brackets. We firmly press it down, spilling more epoxy into the side pockets. We then fill up the side pockets to the level of the slit between platform and tile.




Figure 4.4: Manufacturing process: steps (3), (4) and (5).

Step (5) We are now ready to bake the chip to cure the two-component epoxy. We place the chip onto a non-sticking surface, e.g. a sheet of PTFE, since some of the glue will run through the thin slit we cut in step (2), and we want to avoid gluing the chip to the oven. This precaution is only necessary for the initial bake.

We place a PTFE sheet and an aluminium stamp on top of the ceramic tile and firmly screw the stamp down with 8 screws. The raised Union-Jack-type shape on the platform allows exactly 0.1 mm epoxy to stay between the surfaces of the platform and the ceramic. The rest of the epoxy is expelled, thus leveling the surfaces parallel. The PTFE sheet prevents epoxy from sticking to the stamp and serves as a cushion in order to not break the ceramic. We are very careful to tighten the screws not to much, nevertheless. The other PTFE pieces prevent epoxy from running out from inside the PEEK brackets, and through the thin slit into the inner pocket.

We place the assembly an oven at room temperature and increase its temperature to 100 °C at a rate of 1 °C per minute. Here, we hold the temperature for 1 hour. Subsequently, we increase the temperature to 120 °C




Figure 4.5: Manufacturing process: steps (6), (7) and (8).

at the same rate and hold it there for another 2 hours [90]. Then, we switch the oven off and leave it closed overnight for the chip to cool down. After cool-down, we carefully remove the aluminium stamp and PTFE sheet.

This procedure is repeated for steps (8) in figure 4.5 and (11) in figure 4.6.

Step (6) To remove epoxy residue, we carefully face off the surface of the tile. Then, we cut two parallel channels 1.2 mm deep into the tile with a \varnothing 1.2 mm ball head cutter.

Step (7) We prepare a batch of epoxy and two OFHC copper wires of \varnothing 1 mm, about 200 mm long. We degrease the tile surface and channels, as well as the copper wire. Then, we fill the channels with epoxy, making sure no air pockets are trapped inside the glue. Now, we insert the wires into the channels with equal lengths of wire sticking out both sides of the channels. Also, we rotate the wires to fully wet their surfaces.

Step (8) We repeat the procedure as described in step (5) apart from the PTFE cushion this time having two elevations in place of the channels in the tile to press the wires firmly into the channels.

Step (9) To remove epoxy residue, we again carefully face off the surface




Figure 4.6: Manufacturing process: steps (9), (10) and (11).

of the tile and cut another 0.1 mm deep structure, similar in purpose to the one in step (1).

Step (10) For this step we prepare a conductor wire-eroded pre-structure cut into a 0.5 mm thick sheet of OFHC copper and a batch of epoxy. We degrease the tile surface and the conductor pre-structure. In addition, we oxidise the pre-structure as in step (2). We then, carefully wet one side of the pre-structure, as well as the tile surface, with epoxy. Subsequently, we place the pre-structure onto the tile and gently pressed down so that epoxy penetrates into the cuts in the pre-structure. We remove excess epoxy with a spatula, careful not to scratch the copper surface.

Step (11) We repeat the procedure as described in step (5). We remove all PTFE pieces after cool-down.

Step (12) With a $\varnothing 0.5$ mm cutting tool, we carefully crop the pre-structure to the size of the underlying tile, thus isolating the individual tracks, but omitting for now the area right above the copper wires. The cut is 4 mm deep from the conductor structure surface to liberate all layers from the surrounding hardened epoxy. After that, we crop the rest of the structure




Figure 4.7: Manufacturing process: steps (12) and (13).

by cutting the remaining bit about 0.45 mm deep with the cutting tool and the rest of the way with a scalpel. In this step we also cut alignment markers for later 0.1 mm deep into the copper surface. Subsequently, we remove the rest of the material (copper, epoxy, PEEK) surrounding the platform 4 mm downwards from the structure surface. This is done extra carefully.

Step (13) We now remove the screws that are holding the PEEK in place, pull it off the copper and remove excess glue around the copper wires with small hand tools. Here, we are extremely careful not to exert too much force against any of the bonded structures. Next, we cut away material around the platform to solely leave two protrusions that will help attach the chip to the CF vacuum flange. We hand-polish the copper structure to remove any hardened epoxy residue. To finish this step off, we put the offset block back in place.

Step (14) We drill 8 size M2 tapped holes 4 mm deep through the ceramic tile. We thoroughly rinsed out the holes with 4% Decon90 solution, deionised water and methanol, then dry with pressurised dry air. After degreasing 8 size M2 stainless steel grub screws, we prepare a small batch of epoxy and fill it into a syringe, fitted with a small diameter filling needle. We then apply a small drop of epoxy to the thread of each hole, careful not to scratch the thread in the soft ceramic, and a small drop to each thread of the grub screws. After inserting the screws we use a small amount of acetone on a clean wipe to remove excess epoxy from the exposed threads. We then




Figure 4.8: Manufacturing process: steps (14) and (15).

check for electrical insulation of each screw to the block and unscrew a bit if necessary. If all is good, we place the chip in the oven and repeat the baking instructions as described in step (5).

Step (15) After cool-down, we can start attaching the optics. Due to the temperature sensitivity of the optical components we do not use the oven for curing from now on, but a 150 W halogen lamp sitting beneath the block, heating it up to about 80 °C. We prepare a batch of the index matching two-component epoxy Epotek®353ND (which we use from now on) and apply several small drops to the bottom centre area of the waveguide chip. We then place it on the copper structure, minding the alignment markers from step (12), and hold it there, carefully pressed down, for 30 minutes [91].

Step (16) The fibres we now use are placed in a groove along titanium alloy splints (size 0.77 mm × 3.0 mm, 25/30/35 mm long) for structural support and to hold them steady in a micro-positioning stage. To wet the fibre facet we place a small drop of epoxy onto the side of the waveguide chip just outside the waveguides and immerse the tip of a fibre into it. Then, we move the fibre into position of the first waveguide and decrease the distance to the waveguide side until the epoxy touches the side. Whilst continually realigning the fibre to maximise transmission, we increase the temperature to 80 °C. Subsequently, we place a drop of glue between the splint and the copper structure, which distributes along the splint due to surface tension, and cures within a few minutes. During this time, we need




Figure 4.9: Manufacturing process: steps (16) and (17).

to monitor the transmission and quickly realign if necessary. After this first drop of glue has set, we fill the rest of the gap between the splint and the copper structure with epoxy and leave the chip baking for over 30 minutes. This procedure we repeat with 5 more splints and fibres.

Step (17) For the final step, we need two gradient coils. We slide the inner coil in from the underside of the sub-chip, being careful not to push too far and touch the fibres. Then, we slide the support in from the side, just below the fibres. We place the chip onto the prepared 8" CF flange and screw it tight. After it has been mounted on the flange, we can attach the outer gradient coil with its support and screw it to the sub-chip.

4.3 Performance Tests

Now that we have an atom chip in our hands, what can we expect from it? We really want to know how it behaves under experimental conditions. The following section describes what we want to know about the performance and properties of the chip and the results obtained in isolated tests. Most of these tests were done either in air or in a high vacuum test chamber. Therefore, most results are not directly transferable to the behaviour of the chip inside its experimental environment, but give a very good idea of its expected limits and performance. Furthermore, several specimens of sub-chip and waveguide chip were used in the tests, due to the destructive or

measured part	R [m Ω]	P_{\max} [W]	I_{design} [A]
inner gradient coil (47 turns)	145	5	6
outer gradient coil (48 turns)	151	5	6
each bias wire	≈ 1	< 7	80
both bias wires*	9	58	80
H-structure (any path)	≈ 1	< 3	50
H-structure* (any path)	4	10	50
each end wire**	≈ 1	< 1	30
each end wire*,**	4	< 4	30

*this is including connections.

**The resistance of the end wire structures has actually not been measured but due to the similarity to the H-structure is not expected to be greater than the value given.

Table 4.1: Results for the measured resistance R of the conducting sub-chip structures (uncertainties are due to the limit of the multimeter used: $\sigma_R = 1$ m Ω). The corresponding maximum power production P_{\max} is calculated for each resistance and design current I_{design} in our design.

damaging nature of some of the tests.

4.3.1 Sub-Chip

Resistance

The electrical resistances were measured with 6-digit multimeter (Agilent 34401A) via four-wire resistance measurement. Results for the different parts of the atom chip are given in table 4.1. Apart from the two coils, all intrinsic resistances are less than 1 m Ω . However, the connector resistance¹ (mostly contact resistance) plays a non-negligible role in the production of heat. If we take the resistivity of the connected parts and the maximal currents expected by our design, the resulting powers are mostly within normal bounds for in-vacuum currents (1 to 10 W). However, the bias wire current is expected to produce a large amount of heat, most of which originates at the connections. A part of this heat is conducted through the

¹The connections are made using either BeCu barrel connectors between wires and the feed throughs or thin copper stripes screwed directly to the chip surface and the feed through.

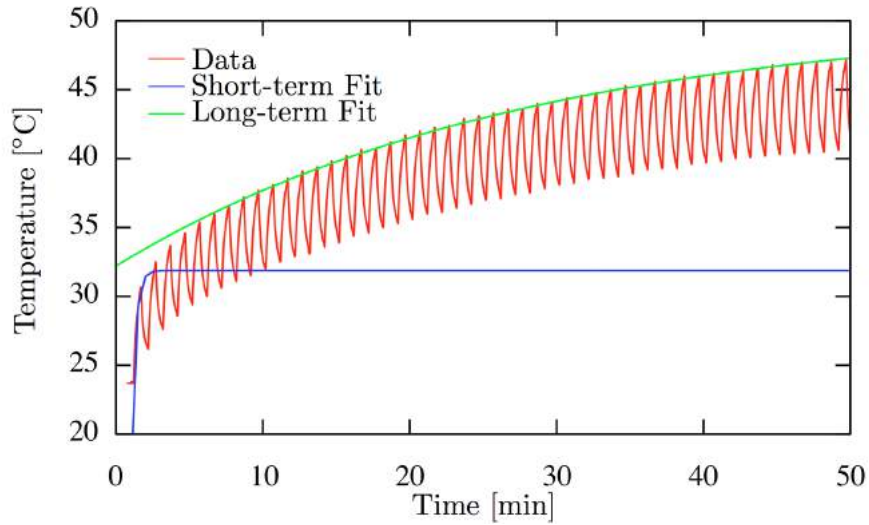


Figure 4.10: Temperature data with fitted curves to the short- (*blue*) and long-term (*green*) rise. The temperature was taken in the centre of the current-carrying structure without a waveguide chip attached.

feedthroughs, and part along the copper wires to the chip. If all currents are running at their maximum value, a total of 86 W is being produced.

Temperature

The temperature was measured with a K-type thermocouple connected to an NI 9211 thermocouple input module. Several tests in air were performed and one under high vacuum conditions. In the tests a sub-chip was used without a waveguide chip or fibres attached to it.

In general we see in the tests a very fast rise in the temperature local to the current. Furthermore when alternating the current on a 1 minute period we see a slow overall rise in temperature throughout the subchip and flange. We attribute the fast rise to short-term local heating before the heat can start to diffuse into the bulk material of the subchip which slowly heats up over a longer period of time. Fitting an exponential of the form²

$$T_0 + T_e \left(1 - \exp\left(-\frac{t}{\tau}\right) \right)$$

²We only take a look at the rise time and assume it is independent of the temperature at which the temperature rise starts. For the temperature range we were taking the measurements for, this was confirmed to hold true to within 10%.

to the short- and long-term rise yield time constants τ . Figure 4.10 shows an example plot of temperature rise data for 80 A running through both bias wires with short- and long-term fits from which the following time constants are derived.

The short-term time constant is an indicator for how long the subchip can continuously be used, the long-term time constant is an indicator for how long the subchip takes to equilibrate with the flange. The following tests have been performed in air with an aluminium heat sink attached to the bottom of the subchip. The uncertainties in time are the asymptotic standard errors of the fit. The uncertainty in temperature (measured with a K-type thermocouple) is on the order of 1°C.

- Running 80 A through each of the copper wires yield a rise time (temperature measured between the wires) of

$$\tau = 168 \text{ s} \pm 7 \text{ s},$$

though this has been measured with the current on for well over 250 s. The temperature did not rise over 60°C.

- Running 80 A through each of the copper wires but alternating the current between on and off every 30 s yield a short-term rise time of

$$\tau = 30 \text{ s} \pm 2 \text{ s},$$

and a long-term rise time of

$$\tau = 1700 \text{ s} \pm 90 \text{ s},$$

indicating that the former measurement is a superposition of the short and long-term rise. The temperature did not rise over 50°C.

- Running 50 A through the H-structure in a Z-configuration yield a rise time (temperature measured in the centre of the subchip) of

$$\tau = 225 \text{ s} \pm 9 \text{ s}.$$

The temperature did not rise over 34°C.

- Running 50 A through the H-structure in a Z-configuration but alternating the current between on and off every 30 s yield a short-term rise time of

$$\tau = 13 \text{ s} \pm 1 \text{ s},$$

and a long-term rise time of

$$\tau = 400 \text{ s} \pm 20 \text{ s}.$$

The temperature did not rise over 35°C.

- Running 80 A through the bias wires, 50 A through the H-structure in a Z-configuration and 6 A through each gradient coil yield a short time rise time (temperature measured in the centre of the subchip) of

$$\tau = 17 \text{ s} \pm 1 \text{ s}.$$

The long-term rise was nearly linear with a linear rise time coefficient of 11 s / K \pm 2 s / K from a temperature of 53°C onwards. Due to the linear rise time, the current was switched off just below 60°C not to damage the chip.

One test has been performed in high vacuum with the subchip attached to a stainless steel flange heat sink.

- Running 50 A through the H-structure in a Z-configuration yield a short-term rise time (temperature measured in the centre of the subchip) of

$$\tau = 750 \text{ s} \pm 30 \text{ s}.$$

The long-term rise was nearly linear with a linear rise time coefficient of 23 s / K \pm 3 s / K from a temperature of 66°C onwards. The current was switched off just below 120°C

The tests show that the chip is able to take a lot of heat produced by the currents run through it in various fashions without heating up to a destructive temperature. Even under vacuum condition, rise times are comparable to those in air, indicating the huge heat-sinking capabilities of the subchip.

Vacuum Issues

The sub-chip is made mostly out of low-out-gassing materials and in normal operation, however MOT-lifetime measurements show a lifetime of only 6 s. Also a pressure rise to a maximum value of 4×10^{-8} Torr can be observed when the gradient coils are switched on. We attribute this to the Kapton® insulation of the gradient coil wire. With regular baking (current left on

over night) we were able to bring the pressure down to 1.5×10^{-9} Torr and falling through 11 weeks of baking at 60°C when running 5 A through each of the coils and through the H-structure.

Mechanical Integrity

Apart from thermal degradation of the individual materials of the sub-chip, a change in temperature has another effect on this compound system. The different materials have different thermal expansion coefficients, hence expand relative to each other to different extents. The shear stress between the bonded layers of the sub-chip can lead to deformation or, in the worst case, de-lamination of layers.

Although the chip reached temperatures of up to 60°C in the temperature tests, we wanted to push it further. Since the acrylate coating of the fibres is the material with the lowest degradation temperature of 140°C , this seems to be a good level for a test. The sub-chip was placed in an oven and in 10 minutes heated from room temperature up to 143°C . The oven was then switched off and the door opened to bring the chip back to room temperature brutally. No active cooling was applied. An inspection showed no damage or other impairment to the mechanical integrity of the chip.

In this test, the temperature was increased near-homogeneously throughout the sub-chip. And that proved to be fine. On the other hand, temperature gradients may still have an impact because in this case not only different materials, but also different parts of the chip expand differently. However, the quick distribution of heat through the copper block and the huge heat capacity of the stainless steel flange prevent these gradients from increasing with increasing number of experimental runs. In other words: if the chip survives the first run, it should the next and so forth.

Magnetic Tests

For the magnetic test Hall probes type LakeShore MNA-1904-VH and MNT-4E04-VH were used. Moving the centre region of the H-structure of the sub-chip along the probe, stray field gradients of less than $300 \mu\text{T} / \text{m}$ were measured.

To check whether the subchip or any components thereof show magnetisation, a 5 mT permanent magnet was moved in the vicinity of the sub-chip

more than 10 times. Subsequently, a current of 80 A was cycled through the bias wires at a period of 1 s. In between and after the two field applications the relative magnetic field change across the centre region of the sub-chip was measured. Both cases showed neither hysteresis effects nor remnant fields (less than 1 μ T).

With a current of 6 A running through each of the gradient coils, the magnetic field along their common axis was measured. The gradient was measured to be 78 ± 26 mT / m at the point in between the coils, which is in agreement with the expected value (from the geometry and current) of 94 ± 9 mT / m. The huge error here comes from the error of positioning the Hall probe to measure the field strength.

Due to the dimensions of commercial Hall probes, it is challenging to measure the field created by the bias wires and the H-structure 0.5 mm from the sub-chip surface. Thus, only one point was taken at an effective distance of 4.3 mm from the sub-chip surface with a current of 80 A running through each bias wire in the same direction. A field of 5.61 ± 0.02 mT was measured in good agreement with the expected value³ of 5.5 ± 0.5 mT.

4.3.2 Waveguide Chip

Light Transmission

To measure the light transmitted by the waveguide chip we wetted the tip of a cleaved single-mode optical fibre with the epoxy Epotek®353ND [91]. Subsequently, the fibre is moved such that the coupling into the waveguide mode is maximised. Then an incandescent lamp underneath the waveguide chip is switched on to bring the temperature of the chip up slowly. The connection is optimised at the curing temperature of the epoxy used. Again due the different materials involved, we expect the coupling efficiency from the fibre into the waveguide (and vice versa) to decrease when the chip cools down and parts start moving relative to others. Hence we check if the temperature change the chip goes through has a reversible or permanent influence on the transmission efficiency.

In this test the sub-chip with connected waveguide chip is placed on a 150 W halogen lamp. With this lamp, the temperature is cycled between 30 °C and 85 °C. The transmission for each waveguide pair is recorded. We

³The uncertainty of the expected value stems from the positional uncertainty of the measurement which was put into a Radia calculation of the chip geometry.

can observe that in this temperature window the transmission recovers fully within a temperature cycle. However, there seems to be a small hysteresis effect.

Polarisation

The fibres used throughout the system are non-polarisation maintaining, so are the optical feedthroughs used. Due to their short length, though, depolarising effects should be small and are mainly due to stress in the air-side fibres. Measuring the polarisation fidelity with circular polarised light is 75%, i.e. if right-handed light is sent in a mixture of 75% right- and 25% left-handed light is measured at the output. If we could servo the input polarisation, i.e. lock the polarisation to a fixed value, more than 90% pure circular polarisation could be detected at the output. Linear polarised light remains 75% linearly polarised. However the axis is arbitrarily rotated. Polarisation drifts on time scales shorter than 10 minutes are less than 5%.

Although these tests tell us little about the state of the light inside the trench of the waveguide chip, we gain insight into overall polarisation performance of the system. Polarisation dependent measurements should be possible.

Waveguide Cross Talk

Cross talk between waveguides is mainly due to scattered light at the interface between fibre and waveguide. The light travels through the silica layer and some is coupled into the waveguides. By sending light into one waveguide pair and measuring the output at a second the cross talk was measured to be on the order of 10^{-3} . Cross talk between adjacent waveguides via the trench is negligible.

Etalon Effects

Although the waveguide chip was not designed for cavity-QED experiments, reflections at the interfaces of the optical system cause etalon effects. This leads to a modulation of the transmitted light as a function of frequency. Measuring this yields modulation of 4% amplitude at a frequency of about 1.1 GHz identified with the fibre feedthroughs of length 90 mm, and another modulation of 1% amplitude at a frequency of about 150 MHz identified with the the in-vacuum fibres of length 600 mm.

4.4 Future Improvements

In the end, the final design of the atom chip emerges from a process of test, result analysis and conclusion. For many features of the sub-chip (and the entire setup, at that) initial reason can get lost, or new ones emerge. This section tries to summarise a few changes to possibly improve certain aspects of this experiment. In this context some things are mentioned that went wrong for others not to make the same mistakes again.

Gradient Coils

The only source of out-gassing and therefore the present limit for cloud life time are the gradient coils. The H-structure on the chip allows for a U-MOT. Hence, if we can show that we can manage to load the magnetic trap efficiently right from the U-MOT, the next version of the chip will not require in-vacuum coils.

Apart from the out-gassing issue, removing the outer coil would open up even more optical access which in the direction of the waveguides is currently limited by the small window in the outer coil support. Removing the inner coil and not cutting the corresponding pocket would allow for even better heat sinking.

Feedthroughs and Fibre Connections

The fibre feedthroughs were specified to have only -0.3 dB attenuation, but we measured a loss of about -3 dB = 50%. They are also fairly expensive (about £300 for one connected piece) and we already detected a small leak in one of the welding seams. Using connector feedthroughs has certainly its advantages, but compared to the Swagelok®-ferrule technique [92, 93] routinely used in many cold-atoms experiments, the disadvantages prevail.

Not only do the fibre splints used to align and connect the fibres to the waveguide chip limit the camera window we have along this import detection direction, but also the the outer gradient coil support. The gradient coils could be replaced by a large Ustructure underneath the current carrying structure to form a MOT. Future splints can be manufactured to be only a few $100\ \mu\text{m}$ tall stretching across the sub-chip to the waveguides, but with a flag at the far end to hold them in the micro-positioning stage. This way, no part of the chip surface is blocked from view. Also, of course, better handling the standard V-groove blocks would be advantageous.

Temperature

Instead of gluing thermocouples to the chip surface, a hole can be drilled at interesting points and a couple inserted. Going in another direction, one could calibrate a 4-wire resistance measurement of one of the electrical structures on the chip. In other words have two electrical feedthroughs touch the connection of the H-structure and measure the voltage drop across them.

No Ceramic

The ceramic is a remnant of the idea of using thermally (and electrically) conductive epoxy (Epotek® H27D) to bond the structured layer to the heat sink block. In this case, an insulation layer between the two is necessary. Although the type of epoxy used changed to a non-conductive sort, the ceramic tile stayed and became the embedding for the bias wires.

It may be considered to omit a ceramic in the next version of the sub-chip. The bias wires can be replaced with an insulated type and placed directly into the copper heat sink. The structured layer can be insulated by the epoxy. However, note that the thermal conductivity of aluminium nitrate ceramic is up to 1000 times higher than that of Kapton®. To have the same performance as the 2 mm thick ceramic a Kapton® sheet no thicker than 2 μm should be used.

Replace screw on connections

To tap holes into aluminium nitride is easy because it is very soft. But because it is soft, the grub screws come out equally easy when applying too much pressure to the nuts. In turn, pressure is needed to make good and low contact resistance connection. So what to do apart from revisiting the entire connection design. The idea here is to use flat head screws instead of grub screws and bury the heads in the sub-chip material. To prevent electrical connection to the copper the head can be either surrounded with insulating material, e.g. Kapton®), or pulled up against the structured layer (or the ceramic if still in use) whilst gluing. The head sticks out into the clearance hole (see step (1) in section 4.2.2) and is covered with insulating epoxy. This second idea is tricky but one might give it a try.

Chapter 5

The Waveguide Chip: Atoms in the Trench

This thesis is about the waveguide chip as a new device to probe atoms. As a first demonstration, we have released cold atoms into the trench to measure the sensitivity of the waveguide chip system and to demonstrate the basics of its operation.

This chapter first investigates the properties of the detection system and subsequently describes the experimental sequence used to launch the atoms into the trench. The interaction of the atoms with light passing through a waveguide pair leads to the first ever signal of cold atoms with a monolithic waveguide device. We measure the spectral response of the atoms in the trench and infer the intensity and polarisation of the light traversing the trench.

Furthermore, we compare the 4 possible detection routes through the chip with the two different waveguide pairs present and we show that the waveguide chip can also pick up light scattered by atoms in the trench when excited by a resonant beam of light shone into the trench from outside.

5.1 Detection System

As previously described in chapter 3 the detection system comprises a single photon counting avalanche-photo-diode (APD) in front of which we can place neutral density filters. One fibre coming from the waveguide chip is connected to this APD and its light is focussed onto it. We assume that our counting system falls into the category of non-paralyseable (type

I) detectors. This means that each photon event is followed by a dead time of length τ_{dead} in which the detector will not record any other events. In other words, if a photon is recorded at time t , then any arrival in the interval $[t, t + \tau_{\text{dead}}]$ is not recorded¹.

In the dark, the APD generates random counts that follow a Poisson distribution. The number of dark counts is specified to be below 100 counts per second [84], however, together with the ambient light reaching the detector we measure 130 counts per second. In any case, the dark counts and ambient light can be neglected.

The typical dead time of the APD is specified with $\tau_{\text{dead}} = 32$ ns (40 ns max) [84]. Now, since no photons are recorded during the dead time, we expect the measured photon rate n_{recorded} to be smaller than the actual rate of photons n_{incident} incident on the detector. The effect is small for $n_{\text{incident}} \ll 1/\tau_{\text{dead}}$ (linear regime of the detector) but becomes more prominent for increasing n_{incident} . In the following we have a look at how to correct n_{recorded} to calculate n_{incident} .

5.1.1 Corrections

Let us assume that in a time interval T_{recorded} the detector records N_{recorded} photons. Now the real number of incident photons, which we refer to as N_{incident} is N_{recorded} plus the number of photons arriving while the detector is refreshing,

$$N_{\text{incident}} = N_{\text{recorded}} + N_{\text{deadtime}}.$$

Now the total time the detector is unresponsive is just the number of recorded counts times the deadtime $N_{\text{recorded}} \cdot \tau_{\text{dead}}$, and the total number of photons arriving in that time that are not counted is this time multiplied by the incident photon rate $N_{\text{incident}}/T_{\text{recorded}}$ where T_{recorded} is the measurement duration. Combining this yields

$$N_{\text{incident}} = N_{\text{recorded}} + \frac{N_{\text{incident}}}{T_{\text{recorded}}} \cdot N_{\text{recorded}} \cdot \tau_{\text{dead}}.$$

Solving this for N_{incident} and recasting in rates $n_{\text{incident}} = N_{\text{incident}}/T_{\text{recorded}}$ and $n_{\text{recorded}} = N_{\text{recorded}}/T_{\text{recorded}}$ we find [95]

$$n_{\text{incident}} = \frac{n_{\text{recorded}}}{1 - n_{\text{recorded}} \cdot \tau_{\text{dead}}}. \quad (5.1)$$

¹In detectors of type II a dead time interval is triggered every time a photon hits the detector, even if it is still in a dead time phase. Due to this behaviour, this type is also called paralysable [94].

Expressing the above equation in terms of power incident on the detector yields

$$P_{\text{APD}} = \frac{hc}{\lambda \eta_{\text{qe}}} \cdot \frac{n_{\text{recorded}}}{1 - n_{\text{recorded}} \cdot \tau_{\text{dead}}}. \quad (5.2)$$

Here λ is the wavelength of the probe light, η_{qe} is the quantum efficiency, and τ_{dead} is the dead time of the detector – in our case $\eta_{\text{qe}} \approx 0.5$ and $\tau_{\text{dead}} = 32 \text{ ns}$. Thus, a baseline count rate of $n_{\text{recorded}} = 7.8$ counts per μs corresponds to an actual power incident on the APD of $P_{\text{APD}} = 5.3 \text{ pW}$.

5.1.2 Shot-Noise Limit

Because we are dealing with coherent light of constant intensities over the time intervals we are interested in, we expect the temporal distribution of photon arrivals to obey Poisson statistics. In a Poissonian process, the variance $\text{Var}(n)$ is equal to the average count rate n . Including the effect of the dead time, the expected variance of the recorded number of photons expressed in terms of the incident rate is [94]

$$\text{Var}(n_{\text{recorded}}) = \frac{n_{\text{incident}}}{(1 + n_{\text{incident}} \cdot \tau_{\text{dead}})^3}. \quad (5.3)$$

This formula states that the measured variance is always smaller than the expected variance n_{recorded} . This can be understood in viewing the short paralysis of the detector on every incident photon acting as a low pass filter for times between two photons shorter than the dead time. With increasing n_{incident} the measured variance passes through a maximum and decreases further. With ever increasing number of incident photons, the time between the end of the paralysis phase and the next recorded event tends to zero. The recorded events appear in a more ordered fashion, therefore the variance tends to zero as well.

Inserting equation 5.1 into equation 5.3 and solving for the variance of the recorded values as a function of n_{recorded} yields

$$\text{Var}(n_{\text{recorded}}) = n_{\text{recorded}} \cdot (1 - n_{\text{recorded}} \cdot \tau_{\text{dead}})^2. \quad (5.4)$$

This result is very helpful, because it lets us validate whether technical noise sources are negligible and the signals we record are indeed shot noise limited.

To show this we measure 20 sample data sets of 10000 points each at 11 different light intensities. We calculate the mean and variance of each




Figure 5.1: Plot of the variance of the mean count rate of an intensity measurement $\text{Var}(n_{\text{recorded}})$ (blue) versus the measured count rate n_{recorded} itself. The lines are theory curves for shot noise (green) and dead-time-corrected shot noise (red) for a value of $\tau_{\text{dead}} = 32$ ns.

set and plot $\text{Var}(n_{\text{recorded}})$ versus n_{recorded} as shown in figure 5.1. The variance of the data points is well below the (green) line showing sub-Poissonian fluctuations. However, plotting equation 5.4 (red) and thus including correction for dead time effects fits the data (blue) very well. The value of maximum variance is reached at a value of $n_{\text{max}} = 1/2\tau_{\text{dead}} \approx 1.6 \times 10^7$ counts per second. The maximum measurable variance for this value of τ_{dead} is $\text{Var}(n_{\text{max}}) = 4/27\tau_{\text{dead}} \approx 4.6 \times 10^6$ counts per second. Beyond this point, the detector is oversaturated.

To express the variance of the incident photon rate, i.e. the variance of the incident power P_{APD} of equation 5.2, in terms of the variance of the measured photon rate, we need to find a solution to the equation 5.3 with n_{incident} substituted by $\text{Var}(n_{\text{incident}})$,

$$\text{Var}(n_{\text{recorded}}) = \frac{\text{Var}(n_{\text{incident}})}{(1 + \text{Var}(n_{\text{incident}}) \cdot \tau_{\text{dead}})^3}.$$

The solution to this equation is not a function but yields two solutions (for $n_{\text{incident}} > 0$). Which solution is applicable depends on the value of n_{incident} (whether n_{incident} is greater or less than n_{max}). Instead, because of this ambiguity, and because we have shown that our detection system is

indeed shot noise limited, i.e. on average $\text{Var}(n_{\text{incident}}) = n_{\text{incident}}$, we use equation 5.1 as estimator for n_{incident} . To specify the relative error on this estimate we use the square root of equation 5.4 and divide by n_{recorded} ,

$$\frac{\sigma_{n,\text{recorded}}}{n_{\text{recorded}}} = \frac{\sqrt{\text{Var}(n_{\text{recorded}})}}{n_{\text{recorded}}} = \frac{1 - n_{\text{recorded}} \cdot \tau_{\text{dead}}}{\sqrt{n_{\text{recorded}}}}. \quad (5.5)$$

5.2 Moving Atoms into the Trench

Every experimental sequence starts by cooling and collecting rubidium 87 atoms from a room-temperature vapour with the LVIS. The pre-cooled atoms are then transferred to the MOT which is located about 3 mm from the chip surface, where the peak atom density is 10^{16} m^{-3} and the temperature is $100 \mu\text{K}$. Loading atoms into the MOT typically takes 3 seconds. Once the atoms are loaded we can execute different sequences to measure the number of atoms in the MOT, or throw the atoms towards the trench and measure them via absorption of probe light or fluorescence. Let us first concentrate on how to move the atoms into the trench.

With the MOT light switched-on we ramp the shim field component pointing in the x -direction B_x from the value where the centre of the MOT is at a distance of about 3 mm from the chip, to a place where the centre of the MOT lies within or very close to within the trench. The positioning of the centre in the y - and z -direction is not very critical because the dimension of the cloud is much larger than the dimension of the trench. For a typical measurement, we do not change the remaining components of the shim field B_y and B_z .

The setup allows for a sequence cycle time of less than 4 s and for any measurement it is possible to specify a set of values for a parameter that we want to scan over. A number of sequences are run with randomized order of the value set for the given parameter. In addition, an associated control measurement without atoms (the quadrupole coils of the MOT are switched-off) is recorded, as well as an absorption measurement of the MOT at the start of each set to monitor the atom number throughout the sequences. This ensures that the data for all parameters is subject to equal long-term drifts of the apparatus, e.g. atom number or probe laser power, on time scales of more than one minute.

5.3 Absorption Measurement

Let us try to understand what is happening in the trench when atoms are introduced into the light field. The scattering rate of one atom in the trench (in cylindrical coordinates) for resonant light is given by [61]

$$R(r, z) = \frac{s(r, z)}{1 + s(r, z)} \cdot \frac{\Gamma}{2}, \quad (5.6)$$

with the transition linewidth (FWHM) Γ . The saturation parameter,

$$s(r, z) = s_0 \cdot \frac{w_0}{w(z)} e^{-2r^2/w^2(z)},$$

is included here in its explicit position dependent form. It comprises the polar coordinates r and z as well as the waist w_0 of the beam and the $1/e$ -field radius

$$w(z) = w_0 \left(1 + \left(\frac{z}{z_R} \right)^2 \right)^{1/2},$$

with the Rayleigh range

$$z_R = \frac{\pi w_0^2}{\lambda}.$$

This dependence is inherited from the definition as the ratio of local intensity and the saturation intensity parameter

$$s(r, z) = \frac{I(r, z)}{I_{\text{sat}}}$$

with $s_0 = s(r = 0, z = 0)$ the saturation parameter on axis immediately at the exit from the first waveguide. If we assume an atom number distribution of $\rho(r, z)$ inside the trench we can calculate the total light power scattered by the atoms,

$$P_{\text{scattered}} = \frac{hc}{\lambda_0} \cdot 2\pi \int_0^L dz \int_0^\infty dr r \cdot R(r, z) \cdot \rho(r, z).$$

Let us here make a simplification and assume $\rho(r, z) = \rho = \text{const.}$ ² and express the power

$$P_{\text{scattered}} = \frac{hc}{\lambda_0} \cdot \frac{\Gamma}{2} \cdot \rho \cdot \mathcal{S}$$

²This assumption is valid for a cloud of atoms pushed at the trench. It breaks down for a trapped sample of atoms.

with the scattering integral

$$\mathcal{S} = 2\pi \int_0^L dz \int_0^\infty dr \cdot \frac{r \cdot s(r, z)}{1 + s(r, z)}.$$

Now we want to relate this to the total power traversing the trench. The power in the trench is approximately, i.e. neglecting the divergence of the beam³,

$$P_{\text{trench}} = \frac{1}{2} \pi w_0^2 \cdot I_0 = \frac{1}{2} \pi w_0^2 \cdot s_0 I_{\text{sat}},$$

where $I_0 = I(r = 0, z = 0) = s_0 I_{\text{sat}}$ is the intensity on axis [96]. The relative absorption A of the light by the atoms is just the ratio of the scattered and total power,

$$A = \frac{P_{\text{scattered}}}{P_{\text{trench}}} = \frac{hc}{\lambda_0} \cdot \frac{\Gamma \cdot \rho_0}{\pi w_0^2 \cdot I_{\text{sat}}} \cdot \frac{\mathcal{S}}{s_0}.$$

At low saturation we can take the limit,

$$\lim_{s_0 \rightarrow 0} \frac{\mathcal{S}}{s_0} \approx \frac{1}{2} \pi w_0^2 L,$$

to approximate the integral, which is just the mode volume across the trench. With this result we can write the relative absorption as

$$A = \frac{hc}{\lambda} \cdot \frac{\Gamma L}{2I_{\text{sat}}} \cdot \rho.$$

Hence for an atom number density of $\rho = 10^{16} \text{ m}^{-3}$ and a saturation intensity⁴ of $I_{\text{sat}} = 24.3 \text{ pW} / \mu\text{m}^2$ we expect a peak absorption signal of $A = 3.2\%$.

5.3.1 Technique

To measure the absorption of probe light we add the following to the sequence that moves the atoms as discussed above. First, probe light is sent through the fibres, through the waveguide chip trench, and collected by the opposite waveguide. This light is then sent to the APD. The acquisition is set to take samples of APD-counts over a time of 80 ms at variable sample rates. The shim field ramp is started after 18 ms and ends after another 12 ms. Right after the ramp, the MOT light is switched-off. This is when

³This approximation is good since the Rayleigh range of the light in the trench is with about $19 \mu\text{m}$ well outside the opposite side of the trench.

⁴The given value I_{sat} is for a randomly chosen light polarisation for each experimental cycle. See section 5.3.3 below for a detailed derivation of this value.




Figure 5.2: Plot of the laser power transmitted through a waveguide pair (#9) to the APD detector. The dip in power shows the atoms blocking about 2.6% of the probe light power.

we expect the atoms to start arriving in the trench and blocking part of the probe light, forming a dip in the recorded transmitted light. 60 ms into the acquired signal, the probe light is switched-off to gather information about the ambient light and noise of the detection system.

An example of a signal is shown in figure 5.2 where we launched atoms at about $v = 0.4 \text{ m/s}$ into the trench. The light beams from the waveguides diverge only slightly as they cross the trench. With w , the $1/e$ -radius of the electric field, growing from $w = w_0 = 2.2 \text{ μm}$ to $w = 2.8 \text{ μm}$ over a width of the trench of $L = 16 \text{ μm}$. Hence, each atom crosses the light mode in about 7 μs , whilst scattering up to 130 photons.

Each experimental measurement records two sets of data points, $N_i(t)$, for the signals with atoms and $N_i^{(c)}(t)$ for the control measurements without atoms, where the index i spans over the number of experiments, and t over the number of data points in each set. In a first step, the sets are summed up to increase the signal-to-noise ratio to above 1,

$$S(t) = \sum_i N_i(t) \quad \text{and} \quad S^{(c)}(t) = \sum_i N_i^{(c)}(t).$$

In the first part of the data, the time interval $[t_0, t_1]$, the MOT light is still on, as well as the probe light. In the second part, the time interval $[t_1, t_2]$, is

the part where the probe light is still on but the MOT light is switched-off. In the third part, the time interval $[t_2, t_3]$, the background light and dark counts are recorded. This last part is not directly used in the signal analysis but as a diagnostic and discrimination tool. The contribution of the MOT light is calculated by taking the difference of the mean values of first and second part of the control measurement,

$$N^{(\text{MOT})} = \int_{t_0}^{t_1} dt S^{(\text{c})}(t) - \int_{t_1}^{t_2} dt S^{(\text{c})}(t).$$

Subsequently, this contribution of the MOT light to the total signal is subtracted from the first part of the data,

$$S(t) \rightarrow S(t) - N^{(\text{MOT})} \cdot \chi(t)$$

and

$$S^{(\text{c})}(t) \rightarrow S^{(\text{c})}(t) - N^{(\text{MOT})} \cdot \chi(t),$$

where $\chi(t)$ signifies the characteristic function,

$$\chi(t) = \begin{cases} 1 & t_0 \leq t \leq t_1 \\ 0 & \text{else.} \end{cases}$$

We can further normalize the data and thus express it in the form of a relative transmission,

$$T(t) = S(t) \cdot \left(\int_{t_0}^{t_2} dt S^{(\text{c})}(t) \right)^{-1}.$$

The integral is taken instead of point-wise normalising the data to avoid increasing the noise. In a last step we transform the data into relative absorption for simplicity,

$$A(t) = 1 - T(t).$$

Finally, to assign a value p to the relative absorption of the measurement we consider the data and select a time interval $[t', t'']$ and integrate the data over this interval,

$$p = \int_{t'}^{t''} dt A(t).$$

Typically, this interval is 2 ms long and $t' \gtrsim t_1$. In the case of measuring the relative absorption whilst scanning a parameter, the above process is applied to each measurement set.

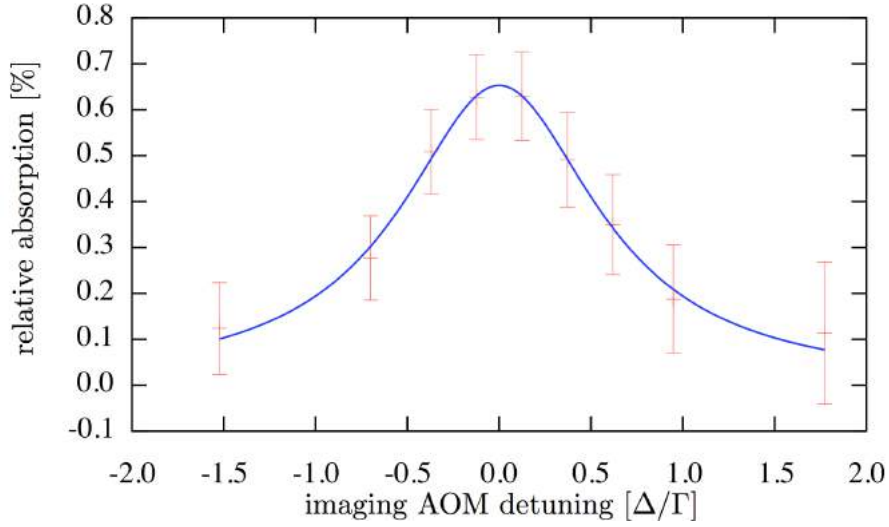


Figure 5.3: Plot of the absorption spectrum of the atom-waveguide-detector-system. The data in form of relative absorption signal for different values of the laser frequency detuning from resonance is superimposed with a Lorentzian fit that yields a linewidth of 7.9 ± 0.4 MHz. Error bars give 1σ statistical uncertainties for Poissonian light with detector dead time.

5.3.2 Probe Frequency

If we return to equation 5.6 and write it down for the case of non-resonant light,

$$R(r, z) = \frac{s(r, z)}{1 + s(r, z) + (2\Delta/\Gamma)^2} \cdot \frac{\Gamma}{2},$$

we note the direct dependence of the scattering rate R on the probe laser frequency detuning $\Delta = \omega_{probe} - \omega_0$, with the atomic transition frequency ν_0 . For this measurement the probe frequency ν_{probe} is scanned by varying the frequency of a voltage-controlled oscillator (VCO) which drives an acousto-optical modulator (AOM). 210 cycles per detuning value were taken and the data processed as outlined in section 5.3.1 with an integration time of $t'' - t' = 2$ ms. The resulting transition line including a Lorentzian fit is shown in figure 5.3 with a linewidth (FWHM) of

$$\Gamma_{\text{measured}} = 2\pi \times (7.9 \pm 0.4) \text{ MHz}$$

The error bars give 1σ statistical uncertainties (c.f. equation 5.5). The line broadening, as compared to the value of the natural linewidth of $\Gamma =$




Figure 5.4: Saturation Curve. For varying intensities the relative absorption is measured and plotted against the power incident on the detector. The solid line is a theoretical curve taking into account the spatial variation of light within the trench and optical pumping of the atoms. Error bars give 1σ statistical uncertainties for Poissonian light with detector dead time.

$2\pi \times 6.1$ MHz [79] most likely stems from a combination of laser frequency jitter (estimated to be around 1 MHz), and Zeeman shifts of different atomic magnetic sublevels in a magnetic background field (inferred to be less than 0.05 mT).

5.3.3 Probe Intensity

The atomic scattering rate is directly intensity dependent. By using calibrated neutral density filters and varying the probe laser power by changing the rf power sent into the AOM we are able to span almost 3 orders of magnitude of intensity in this measurement. A number between 160 and 1400 cycles per detuning value were taken⁵ and the data processed, again according to section 5.3.1 with an integration time of $t'' - t' = 2$ ms.

Figure 5.4 shows the relative absorption as a function of inferred power on the detector P_{APD} . Superimposed on the data points we fit a saturation

⁵Here we decided to take a higher number of sample cycles for lower intensities to account for the relatively higher noise that presents itself at lower incident photon numbers.

function of the form

$$A(P_{\text{APD}}) = \frac{A(P_{\text{APD}} \rightarrow 0)}{1 + P_{\text{APD}}/P_{\text{sat}}},$$

with a given value for the saturation intensity of $I_{\text{sat}} = 22.4 \text{ pW} / \mu\text{m}^2$, which is derived below and $A(P_{\text{APD}} \rightarrow 0)$ as a fit parameter. The error bars on the data points give 1σ statistical uncertainties (c.f. equation 5.5). We see good agreement between the measurement and theory, which indicates that we understand in this respect the interaction between the atoms and the light in the trench.

Let us have a closer look at the saturation intensity for the experiments we performed so far. The polarisation exiting the fibres may drift due to thermal and mechanical fluctuations in the fibres. So unless we make an extra effort to maintain polarisation, we can assume that all mixtures of right- and left-handed light are present during the experiment. The corresponding intensities are I_+ for right- and I_- for left-handed light. The light intensity I experienced by one atom in the trench can be expressed in the form

$$I = I_+ + I_-,$$

with $I_+ = X \cdot I$ and $I_- = (1 - X) \cdot I$ controlled by a splitting parameter X , $0 \leq X \leq 1$. In other words, a fraction X of the intensity is right-handed light and a fraction $(1-X)$ is left-handed light. We expect optical pumping into the extremal m_F states to occur, increasing the effective scattering as compared to an equal state distribution. The transitions we are looking at are between the hyperfine levels $|F = 2, m_F\rangle$ (5 levels) and $|F' = 3, m'_F = m_F + q\rangle$ (7 levels). The Einstein coefficients of these transitions are

$$A^{(m_F, q)} = \Gamma \cdot (2F' + 1) \cdot \left(\begin{array}{ccc} F' & 1 & F \\ -(m_F + q) & q & m_F \end{array} \right)^2$$

and

$$B^{(m_F, q)} = \frac{2\sigma_0}{\hbar\omega_0} \cdot \left(\begin{array}{ccc} F' & 1 & F \\ -(m_F + q) & q & m_F \end{array} \right)^2,$$

where $q = 0, \pm 1$ denotes the change in m_F for the transition. $q = -1$ for left-handed $q = 0$ for linear, and $q = +1$ for right-handed light polarisation. The inclusion of the Wigner $3j$ -symbols governs the relative transition strengths between the hyperfine levels. $A^{(m_F, q)}$ is the relative strength to relax from an excited state m_F to a ground state $m_F + q$. $B^{(m_F, q)}$ is the relative strength to

excite a ground state m_F to an excited state $m_F + q$, and for an excited state m_F to relax to a ground state $m_F + q$. $A^{(m_F, q)}$ is normalised so that the total decay rate into all sublevels is Γ ,

$$\sum_{q=-1}^1 A^{(m_F, q)} = \Gamma.$$

We express the population of the ground and excited state Zeeman levels for m_F by $N_g^{(m_F)}$ and $N_e^{(m_F)}$. The 12 rate equations for the system are then

$$\dot{N}_g^{(m_F)} = \left(\sum_{q=-1}^1 N_e^{(m_F-q)} \cdot A^{(m_F-q, q)} \right) + \left(\sum_{q=\pm 1} (N_e^{(m_F+q)} - N_g^{(m_F)}) \cdot B^{(m_F, q)} \cdot I_q \right)$$

and

$$\dot{N}_e^{(m_F)} = \left(\sum_{q=\pm 1} (N_g^{(m_F-q)} - N_e^{(m_F)}) \cdot B^{(m_F-q, q)} \cdot I_q \right) - \left(N_e^{(m_F)} \cdot \sum_{q=-1}^1 A^{(m_F, q)} \right),$$

with $m_F \in \{-2 \dots 2\}$ for $\dot{N}_g^{(m_F)}$ and $m_F \in \{-3 \dots 3\}$ for $\dot{N}_e^{(m_F)}$. Note, that we are not including excitations and stimulated emission by π -polarised light.

To find the steady state solutions to the state populations we set $\dot{N}_g^{(m_F)} = \dot{N}_e^{(m_F)} = 0$ and solve the coupled system of equations for $N_g^{(m_F)}$ and $N_e^{(m_F)}$. With the solutions we can now write an expression for the total scattering rate in steady state for a given right-handed light fraction X ,

$$\Gamma_{\text{sc}}^{(X)} = I \cdot \sum_{m_F=-2}^2 N_g^{(m_F)} \left(X \cdot B^{(m_F, +1)} + (1 - X) \cdot B^{(m_F, -1)} \right),$$

where we approximate the intensity with the saturation intensity $I = I_{\text{sat}}$. For the absorption spectrum and intensity measurements the value of X changed randomly from cycle to cycle and we assume all values are equally probable. In this case we can find an effective scattering rate by taking the mean over all X values,

$$\Gamma_{\text{sc,eff}} = \overline{\Gamma_{\text{sc}}^{(X)}}.$$

The value of saturation intensity for can be found by comparison with the minimum value for pure right- or left-handed light with a value of

$$I_{\text{sat}}^{(1)} = \frac{\pi h c}{3 \lambda_0^3} \Gamma = 16.7 \text{ W / m}^2$$

we find an effective saturation intensity of

$$I_{\text{sat,eff}} = \frac{\Gamma_{\text{sc}}^{(1)}}{\Gamma_{\text{sc,eff}}} \cdot I_{\text{sat}}^{(1)} = I_{\text{sat,eff}} = 24.3 \text{ W / m}^2.$$




Figure 5.5: Zeeman-split absorption spectrum. When pure left-handed light (red) is coupled into the input fibre the Zeeman shift is almost $-\mu_B B/h$. The green line is a theory obtained by summing over all the Zeeman components of the spectrum. This fit tells us that $(85 \pm 3)\%$ of the power in the trench is still in the left-handed mode. For right-handed light (blue) the shifts are opposite. The fit to these points gives $(81 \pm 3)\%$ of power in the right handed mode. Error bars give 1σ statistical uncertainties for Poissonian light with detector dead time.

5.3.4 Probe Polarisation

The prior measurements have been done with random polarisation from shot to shot. To measure the polarisation of the light with the atoms we use the fact that the Zeeman shift depends on the atomic magnetic sublevel, i.e. the frequency of the atomic transition depends on the polarisation of the atom, the light and the magnetic field. With a magnetic field along the propagation direction the quantisation axis of the atomic magnetic sublevels is along that direction. Hence, any light polarisation state can be described by a superposition of left- and right-handed light states.

In respect to the polarisation maintaining qualities of the waveguide chip system, we find that light prepared with positive helicity retains typically 87% of its power in that mode after traversing the system⁶. This value

⁶This is provided the fibres are securely held.

was measured using standard polarisation optics, such as quarter wave plates ($\lambda/4$), half wave plates ($\lambda/2$), and polarising beam splitters (PBS).

We apply a magnetic field of 0.78 mT along the direction of propagation by running a current of 8 A through the bias wires buried under the subchip. The current is switched-on when the MOT light is switched-off. A number of 30 cycles per probe frequency value were taken and the data processed as outlined in section 5.3.1 with an integration time of $t'' - t' = 2$ ms starting 2 ms after light off to leave time for eddy currents to ease off.

Figure 5.5 shows the relative absorption of the light against the normalised frequency shift in the presence of the applied magnetic field for input light of positive (open circles) and negative (filled circles) helicity. The expected Zeeman shift on this $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition is $\pm\mu_B B/h$ as expected by theory.

The data is fitted with 10 Lorentzian functions, one for each possible transition. The relative amplitude of each transition is weighted with its ground state population. The free parameters in the fit are X , the fraction of right-handed light, and the total intensity I . If the light enters the system with positive helicity, it goes through the trench with $X = (85 \pm 3)\%$ positive helicity. Left-handed light goes through the trench retaining $(1 - X) = (81 \pm 3)\%$ negative helicity.

5.3.5 More Channels

So far, we only took measurements using one waveguide pair sending the light through the chip in one direction. However, there are two working waveguide pairs, #9 and #10 (compare figure 4.2 on page 53), producing 4 possible directions of detection, and we would like to now how their performance compares to the one used so far. The saturation measurement of section 5.3.3 was repeated to determine the trench-to-detector efficiency for each of the 4 possibilities and compare it to the directly measured values.

Measuring waveguide #9 and sending the light from, say, the *left*⁷ waveguide, we measure an efficiency of about $18\% \pm 5\%$. Measuring the same waveguide from the other, the *right*, waveguide, we measure an efficiency of again about $18\% \pm 5\%$. Multiplying the two to get an estimate for the total transmission efficiency yields a value of $3\% \pm 2\%$, which is fully compatible with the measurement of directly calibrated transmission efficiency of 4.5% found in chapter 4.

⁷Only to give the direction the light is sent through a label.

The corresponding values for waveguide #10 are roughly $5\% \pm 2\%$ and less than $1\% \pm 2\%$ (nearly no signal was seen at all), yielding $0.1\% \pm 0.1\%$ which again agrees with the directly measured total transmission efficiency of 0.02%. We indeed have a multichannel device at hand, although with largely varying performance.

5.4 Fluorescence Measurement

In addition to utilizing the absorption of probe light by the atoms we can count their number within the junction by measuring their fluorescence into the waveguides. We tested this idea by moving atoms as outlined above almost into the trench, then turning on an intense excitation beam and detecting the fluorescence collected by the waveguide. We estimate the expected rate hitting the detector as follows.

The number of atoms in the light mode volume in front of each waveguide is this volume times the density ρ of the atoms, $\pi w_0^2 L \rho$, where $w_0 = 2.2 \mu\text{m}$ is the mode size and $L = 16 \mu\text{m}$ the width of the trench. Each atom fluoresces at a rate up to $\Gamma/2 = 2\pi \cdot 3.03 \text{ MHz}$. The waveguide accepts only the light emitted into the solid angle $\pi w_0^2 / 4\pi(L/2)^2 \approx 2\%$. In total, the peak fluorescence arriving at the detector is approximated by

$$F = \Gamma \frac{\rho \pi w_0^4 \eta}{2L},$$

where $\eta \approx 8.5\%$ accounts for transmission loss and detector quantum inefficiency. For an assumed atom number density of $\rho = 10^{16} \text{ m}^{-3}$ we expect thus a peak fluorescence of $F = 73$ counts per ms.

5.4.1 Technique

Similar to absorption measurement technique, we load atoms in a MOT. 18 ms after starting the data acquisition the atoms are moved towards the trench by ramping the shim fields. However, no probe light is sent through the waveguides to the detector. Instead the APD is measuring the ambient background light. Subsequently, 29 ms after start of detection an intense, resonant beam is switched-on. The direction of this beam is aligned perpendicular to the chip surface to impinge upon the centre of the waveguide chip and right into the trench. This beam pushes the atoms fast into the trench and is scattered into all directions by the atoms. While the atoms




Figure 5.6: Fluorescence signal collected by the waveguide and detected on an APD. This plot shows the averaged data of 34 shots. The measurement of the peak signal agrees very well with the estimated fluorescence of 73 counts per millisecond, indicating that the fluorescence detection methods works well in our device.

are in the mode volume of the waveguides, a part of this light enters the waveguide and travels down to the APD. Hence, we expect a short peak against the ambient light to appear in the signal. The peak width of $200 \mu\text{s}$ indicates a speed of the atoms of approximately 4 m/s . A further 60 ms into the signal acquisition, the fluorescence beam is switched-off to gather information about the ambient light and noise of the detection system.

5.4.2 Results

If we take a fluorescence measurement without atoms, the ambient light comprises mainly of the MOT light that is coupled into the fibres at different locations in the chamber. This light produces a background of about 10 counts per millisecond (compared to the 0.065 counts per second background from other sources). At 29 ms the MOT light is switched-off and the fluorescence beam switched-on, resulting in a change in ambient light to the fluorescence beam background level.

What is measured when adding atoms is shown in figure 5.6. The plot

is produced by averaging the data of 34 experimental cycles. We can clearly see the fluorescence signal on the APD rising out of the ambient background to the expected peak value when the excitation beam is turned on. Such good agreement shows that the waveguide indeed collects more than 1% of the fluorescence from atoms within the volume of the junction, making this another practical way to count the atoms, especially since measuring with this method, as compared to an absorption measurement, takes far fewer experimental cycles to reach the same signal-to-noise ratio.

Summary

To summarise, we detected cold atoms at a density of $10^{-2} \mu\text{m}^{-3}$. We roughly calculate the light mode volume between two waveguides as

$$2.5 \mu\text{m}^2 \cdot 16 \mu\text{m} = 100 \mu\text{m}^3.$$

Here we have taken the mean of beam waist and 1/e-radius across the trench to form a box across the length of the trench. Multiplying this volume with the measured density, we see that on average one atom per channel is residing in the detection light mode. This has been measured with a new kind of chip using absorption and fluorescence techniques. Subsequently, we utilised atoms to diagnose the polarisation and intensity of light within the waveguide chip trench.

Chapter 6

Photon-Shot-Noise-Limited Detection

This chapter takes a look at the interaction between atoms in a Bose-Einstein condensate brought inside the waveguide chip and the light as it is sent through the waveguide chip. In particular, non-destructive methods are discussed which minimise spontaneous emission that leads to heating and destruction of a sample condensate. Non-destructive methods are vital if one is to measure fast dynamics in condensates by reducing the sensitivity of the measurement, if possible down to the shot-noise limit [97,98]. Non-destructive measurements have been successfully demonstrated on atom clouds [99–101] and even single atoms [102].

6.1 Light-Matter Interaction

Figure 6.1 shows a model of the geometry of the initial problem. The light travels along the x -axis. It intersects the condensate, which is smaller than the light profile along the y -axis and larger along the z -axis.

6.1.1 Condensate Density

A condensate can only form inside a trapping potential. Thus, we must handle the shape of this trap as a parameter. We assume that the trap is harmonic with cylindrical symmetry specified by radial and axial oscillation frequencies ω_r and ω_a , respectively.




Figure 6.1: Model of cloud and light geometry. The shape of the condensate (*blue*) is parametrised by r_0 and z_0 . Its density ρ is a parabolic profile. The light beam (*red*) along the x -axis has a Gaussian electric field profile with waist w_0 .

In the Thomas-Fermi limit¹ a condensate can be described by a cloud with parabolic density profile along all three axes²,

$$\rho(r, z) = \rho_{\max} \left(1 - \frac{r^2}{r_0^2} - \frac{z^2}{z_0^2} \right),$$

with a peak atom density ρ_{\max} at the centre. The values for the equatorial radius r_0 and polar radius z_0 are given by the size of the ground state wave function and can be calculated from the oscillation frequencies, mass of the atomic species m_{Rb} and chemical potential μ of the cloud by [3]

$$r_0 = \sqrt{\frac{2\mu}{m_{\text{Rb}}\omega_r}} \quad \text{and} \quad z_0 = \sqrt{\frac{2\mu}{m_{\text{Rb}}\omega_a}}.$$

The surface of the cloud has the geometry of a prolate spheroid for $r_0 < z_0$.

¹We remember that in this limit of strong inter-atomic interaction the kinetic energy term in the Gross-Pitaevskii equation is neglected, leading to parabolic density profile in case of a harmonic trap [3].

²Note that the domain of ρ , as of all the functional properties of the condensate, is not all space but only the volume V_{bec} of the cloud.

The number of atoms in the condensate N_{atoms} can be calculated by integrating ρ over the volume of the condensate,

$$N_{\text{atoms}} = \iiint_{V_{\text{bec}}} dV \rho(r, z) = \int_0^{2\pi} d\phi \left\{ \int_{-z_0}^{z_0} dz \left\{ \int_0^{r'} dr [r \cdot \rho(r, z)] \right\} \right\},$$

with $r' = r_0(1 - z^2/z_0^2)^{1/2}$. We can simplify³ and perform the integration to arrive at

$$N_{\text{atoms}} = 4\pi \cdot \int_0^{z_0} dz \left\{ \int_0^{r'} dr [r \cdot \rho(r, z)] \right\} = \frac{8\pi}{15} \cdot r_0^2 \cdot z_0 \cdot \rho_{\text{max}} \quad (6.1)$$

or solving for ρ_{max} we find

$$\rho_{\text{max}} = \frac{15}{8\pi} \cdot \frac{N_{\text{atoms}}}{r_0^2 \cdot z_0}.$$

To simplify later on we recast the local density ρ into the form of an average density $\bar{\rho}$ of the cloud. The average density over the entire condensate is obviously the mean density, i.e. the number of atoms within the condensate divided by its volume. Using equation 6.1 we find

$$\bar{\rho} = \frac{N_{\text{atoms}}}{V_{\text{cloud}}} = \frac{3N_{\text{atoms}}}{4\pi r_0^2 z_0} = \frac{2\rho_{\text{max}}}{5}$$

However, we can do better. In a second approach we recognise that the cloud is stretched-out. If we assume that the cloud is stretched-out far, i.e. $z_0 \gg r_0$, and that r_0 is on the order or smaller than the $1/e^2$ -intensity diameter of the waveguide mode, i.e. $r_0 \leq w$, the light will only interact with a small portion of the entire condensate of near constant density along z . In other words we approximate the cloud as cylindrical shaped along the width of the light mode exiting the waveguide. We define an atom number density along the z -axis as

$$N'_{\text{atoms}} = 2\pi \cdot \int_0^{r'} dr [r \cdot \rho(r, z)] = \frac{\pi}{2} \cdot r_0^2 \cdot \rho_{\text{max}} \cdot \left(1 - \frac{z^2}{z_0^2}\right)^2.$$

³The integration over ϕ yields a factor of 2π . Due to the symmetry of ρ the lower bound of integration over z can be moved to zero, multiplying by a factor of 2.

Note that $N_{\text{atoms}} = \int_{-z_0}^{z_0} dz N'_{\text{atoms}}$. Furthermore, the volume density along the z -axis is just the area of the circle perpendicular to this axis,

$$V'_{\text{cloud}} = \pi r'^2 = \pi r_0^2 \cdot \left(1 - \frac{z^2}{z_0^2}\right),$$

thus the mean density in this case turns out to be

$$\bar{\rho} = \frac{N'_{\text{atoms}}}{V'_{\text{cloud}}} = \frac{\rho_{\text{max}}}{2} \cdot \left(1 - \frac{z^2}{z_0^2}\right) = \frac{15}{16\pi} \cdot \frac{N_{\text{atoms}}}{r_0^2 \cdot z_0} \cdot \left(1 - \frac{z^2}{z_0^2}\right)$$

for a cloud intersecting the waveguide pair at position z . The maximum value of

$$\bar{\rho} = \rho_{\text{max}}/2 \tag{6.2}$$

is reached if the light penetrates the centre of the cloud at $z = 0$, and we assume this to be the case for our model.

6.1.2 Effective Thickness

If the cloud has a local diameter d in the direction of light propagation, i.e. x , we can simplify by flattening out the density and providing a constant thickness across the interaction cross section, c.f. figure 6.3c. We transform the circular shape of area πr_0^2 into a rectangle with sides $h = 2r_0$ perpendicular, and \bar{d} parallel to the x -axis. Setting the two areas equal, we find

$$\bar{d} = \frac{\pi r_0}{2}, \tag{6.3}$$

which gives us an average or effective thickness across the cloud.

6.1.3 Light-Condensate Overlap

Let us calculate the overlap between the light mode and the condensate density profile by assuming a Gaussian beam profile [96],

$$E'(x, y, z) = E_{\text{max}} \cdot \frac{w_0}{w(x)} \cdot \exp\left(-\frac{y^2 + z^2}{w^2(x)} - i\theta\right),$$

with electric field amplitude E_{max} , beam waist w_0 , $1/e$ -field radius along the direction of propagation $w(x)$, and some phase θ . We need to account for the geometrical problem that the light interacts with some parts of the cloud stronger than others, due to the imperfect overlap of their spatial modes. We calculate the 1d-overlap integral between the light and the cloud. In

other words, we compare the profiles along a radial axis where $x = z = 0$ -axis. This is the density profile in the centre of the condensate with the beam profile as it leaves the waveguide mode, which is slightly smaller than the one at the position of the cloud,

$$O_{1d} = \frac{\int dy \rho(0, y, 0) \cdot E'(0, y, 0)}{\int dy E'(0, y, 0)} = \frac{1}{\sqrt{\pi}\xi} \cdot \exp(-\xi^2) + \left(1 - \frac{1}{2\xi^2}\right) \cdot \text{erf}(\xi),$$

for $\xi = r_0/w_0$, the size of the condensate relative to the initial beam radius.

To slightly improve upon this result we look at the problem in 2d. We first calculate the normalised column density along the x -axis for a circular cloud shape in the x - y -plane,

$$Q(y) = \frac{\int_{-x'}^{x'} dx \rho(x, 0, 0)}{\int_{-r_0}^{r_0} dx \rho(x, 0, 0)} = \left(1 - \frac{y^2}{r_0^2}\right)^{\frac{3}{2}},$$

with $x' = r_0(1 - y^2/r_0^2)^{1/2}$, and again convolve the result with the beam profile to obtain

$$\begin{aligned} O_{\text{quasi-2d}} &= \frac{\int_{-y_0}^{y_0} dy E(0, y, 0) \cdot Q(y)}{\int_{-\infty}^{\infty} dy E(0, y, 0)} \\ &= \frac{\sqrt{\pi}\xi}{2} \cdot \exp\left(-\frac{\xi^2}{2}\right) \cdot \left[\mathcal{I}_0\left(\frac{\xi^2}{2}\right) + \left(1 - \frac{1}{\xi^2}\right) \cdot \mathcal{I}_1\left(\frac{\xi^2}{2}\right) \right], \end{aligned}$$

where $\mathcal{I}_n(x)$ are the modified Bessel functions of the first kind [103]. Compared to the 1d case, we have now accounted for the changing density of the cloud along the direction of propagation.

Both O_{1d} and $O_{\text{quasi-2d}}$ are plotted in figure 6.2 for $w = w_0$. Note that including the change of the density along the x -axis decreases the overlap. Including the third dimension into the problem, i.e. calculating the overlap of the full 3d density profile and light mode, alters the result only negligibly, assuming as before that the cloud is highly elongated, i.e. $w_0 \lll z_0$.

6.1.4 Susceptibility

Assuming, the light passing through a medium is well below saturation, the response of the medium is going to be linear and characterised by the




Figure 6.2: Overlap integrals plotted against relative condensate size in the case $w = w_0$. In this case the curves reach their maximally possible values.

complex refractive index n . Let us derive an expression for n that depends on the parameters of the condensate and the condensate-light interaction, starting with the linear susceptibility [104],

$$\chi(\omega) = \rho \cdot \frac{e^2 \mu_{eg}^2}{3\epsilon_0 \hbar} \cdot \left(\frac{1}{\omega_0 - \omega - i\gamma} + \frac{1}{\omega_0 + \omega + i\gamma} \right),$$

of a medium with local density ρ . Furthermore, $\omega_0 = (E'_e - E_g)/\hbar$ is the frequency between the two energy levels that we are considering, γ the half-width-at-half-maximum (HWHM) of the transition line, and μ_{eg} the dipole matrix elements of the 5S ground state $|F, m_F\rangle$ to 5P excited state $|F', m'_F\rangle$ transition [61]. e is the electron charge, ϵ_0 the vacuum permittivity, and \hbar Planck's constant. If the light we shine in is close to resonance, i.e. $\omega \approx \omega_0$, and $\gamma \ll \omega_0$, the first term in the brackets exceeds the second. We neglect the latter. This is the rotating-wave approximation.

Since χ now only depends on the difference between the frequencies of light and atomic transition, we can express it as such. We define the unit-less detuning,

$$\Delta' = \frac{\Delta}{\Gamma} = \frac{\omega_0 - \omega}{\Gamma},$$

with $\Gamma = 2\gamma \approx 2\pi \cdot 6.1 \text{ MHz}$ for the D_2 -line [79] and use the relation $z^{-1} =$

$z^* \cdot (z \cdot z^*)^{-1}$, $z \in \mathbf{C}$ to write

$$\chi(\Delta') = \rho \cdot \frac{e^2}{3\epsilon_0 \hbar} \cdot \mu_{eg}^2 \cdot \left(\frac{2}{\Gamma} \cdot \frac{i + 2\Delta'}{1 + 4\Delta'^2} \right).$$

The dipole matrix elements are given by [104]

$$\frac{2J+1}{2J'+1} \cdot \mu_{eg}^2 = \frac{3\lambda_0^3 \epsilon_0 \hbar}{8\pi^2 e^2} \cdot \Gamma.$$

We consider the cycling transition $J = 1/2 \rightarrow J' = 3/2$. In this case, the degeneracy-ratio yields a factor of 1/2. We can thus write the susceptibility,

$$\chi(\Delta') = \frac{\rho \lambda_0^3}{2\pi^2} \cdot \mu_{\text{eff}}^2 \cdot \left(\frac{i + 2\Delta'}{1 + 4\Delta'^2} \right), \quad (6.4)$$

with some values of μ_{eff}^2 given in section 6.1.5 and values of μ_{eg} for other transition in [79] tables 7 to 20.

Approximating the local density by the average density $\bar{\rho}$ (equation 6.2) and using $\mu_{\text{eff}}^2 = 1/2$ as appropriate for pure σ^+ light⁴ we can recast equation 6.4 using the expression for the resonant scattering cross section⁵ $\sigma_0 = 3\lambda_0^2/2\pi$,

$$\chi(\Delta') = \frac{\sigma_0 \lambda_0 \bar{\rho}}{6\pi} \cdot \frac{i + 2\Delta'}{1 + 4\Delta'^2},$$

which we can also directly express in terms of the light, trap and condensate properties: wavelength, half-lengths and atom number,

$$\chi(\Delta') = \frac{15}{64\pi^3} \cdot N_{\text{atoms}} \cdot \frac{\lambda_0^3}{r_0^2 \cdot z_0} \cdot \frac{i + 2\Delta'}{1 + 4\Delta'^2}.$$

This is the well-known Lorentzian line shape function. From the derived expression for the susceptibility we can directly obtain the electric permittivity

$$\epsilon_r = 1 + \chi \quad (6.5)$$

and the complex index of refraction⁶

$$n = \sqrt{\epsilon_r} = \sqrt{1 + \chi}. \quad (6.6)$$

⁴Also assuming a quantisation axis along the direction of light propagation and confining to the $|F = 2, m_F = 2\rangle \leftrightarrow |F' = 3, m'_F = 3\rangle$ cycling transition.

⁵Used in the form for σ^+ -polarised light.

⁶We assume here, that the magnetic permeability μ of the medium (and later on the waveguide chip material) is equal to that of free space: $\mu \approx 1$. This claim is good as long as the frequency of light is close to an electric dipole transition.

6.1.5 Effective Dipole Matrix Elements

For the measurements performed in the last chapter we used the magnetic field of the bias wires to define the magnetic quantisation axis in the x -direction, i.e. the direction of light propagation. In this case, light with positive (negative) helicity is simply *seen* as σ^+ (σ^-) light by the atoms. In the case of atoms in a magnetic trap, the quantisation axis lies in the direction of the z -axis⁷. This results in different transition strengths, hence different transition matrix elements. We can distinguish four cases of light coming towards the atoms and write down the light's polarisation vector $\hat{\mathbf{p}}$ [61].

The cases are:

1. linear polarised light, oscillating along the y -axis: $\hat{\mathbf{p}}_1 = \hat{\mathbf{y}}$,
2. linear polarised light, oscillating along the z -axis: $\hat{\mathbf{p}}_2 = \hat{\mathbf{z}}$,
3. circular polarised light with positive helicity: $\hat{\mathbf{p}}_3 = (\hat{\mathbf{z}} - i\hat{\mathbf{y}}) / \sqrt{2}$,
4. circular polarised light with negative helicity: $\hat{\mathbf{p}}_4 = -(\hat{\mathbf{z}} + i\hat{\mathbf{y}}) / \sqrt{2}$.

The value of the dipole element depends on the wave functions of the ground and excited states and is generally complicated to calculate. Therefore, we want to find an approximate solution by expanding the light in terms of the simple spherical harmonics [105] with $l = 1$,

$$\begin{aligned}\hat{\mathbf{x}} &= \sqrt{\frac{2\pi}{3}} \cdot r \cdot (Y_{1,-1} - Y_{1,+1}), \\ \hat{\mathbf{y}} &= \sqrt{\frac{2\pi}{3}} \cdot ir \cdot (Y_{1,-1} + Y_{1,+1}), \\ \hat{\mathbf{z}} &= \sqrt{\frac{4\pi}{3}} \cdot r \cdot Y_{1,0}.\end{aligned}$$

This gives us the relative contribution of each transition $|F, m_F\rangle \rightarrow |F', m'_F\rangle$, which we can then multiply by the corresponding relative transition strength [79]. The relative contributions are given in table 6.1, together with the values of the resulting effective transition dipole elements. We limit the transitions to the ones originating in the ground state $|2, 2\rangle$ because that is the state prominently populated in our magnetic trap after optical pumping.

⁷In addition to the magnetic field component in the z -direction B_z , there is a small (typically 5%) contribution of the x -component B_x , resulting in a slightly tilted quantisation axis. Compare with calculations in chapter 2.

Case	Squared Hyperfine Dipole Matrix Elements			μ_{eff}^2
	$ 2, 2\rangle \rightarrow 3, 1\rangle$ $\equiv 1/30$	$ 2, 2\rangle \rightarrow 3, 2\rangle$ $\equiv 1/6$	$ 2, 2\rangle \rightarrow 3, 3\rangle$ $\equiv 1/2$	
1.	1/2	0	1/2	4/15
2.	0	1	0	1/6
3.	1/4	1/2	1/4	13/60
4.	1/4	1/2	1/4	13/60

Table 6.1: Effective hyperfine dipole matrix elements [79] from the $|2, 2\rangle$ state into $|3, m'_F\rangle$ exited states. For each case of incoming light polarisation (see text) the relative fraction of light in the $(\sigma^-, \pi, \sigma^+)$ basis is given, as well as the resulting effective transition strengths.

6.2 Optical Model

Having accounted for the geometry of the problem, we may treat the incoming light as a plane wave⁸,

$$E_i(x, t) = E' \cdot e^{i(\omega t - kx)},$$

with angular frequency ω and wave number k . Only the part $\mathcal{O}_{\text{quasi-2d}}$ will interact with the atoms, thus we can write the outgoing field as

$$E_o = (1 - \mathcal{O}_{\text{quasi-2d}}) \cdot E_i + \mathcal{O}_{\text{quasi-2d}} \cdot E_i \cdot e^{i\frac{\omega d}{c}(n-1)} = \tau e^{i\phi} \cdot E_i. \quad (6.7)$$

Here we characterise the atom cloud as a medium of thickness d with complex refractive index n . The interaction between atoms and light, i.e. the change in amplitude and phase of the incoming beam is parametrised by the transmission amplitude τ and shift in phase ϕ .

6.2.1 Simulating the Interaction

To solve the full wave equation [96] in the presence of a dielectric medium with relative electric permittivity ϵ_r ,

$$\nabla^2 \mathbf{E} - \frac{\epsilon_r \mu}{c^2} \ddot{\mathbf{E}} + \nabla (\mathbf{E} \cdot \nabla \ln \epsilon) = 0, \quad (6.8)$$

⁸We neglect the divergence of the beam. This assumption is valid because the radius of the cloud is much smaller than the Rayleigh length of the light, $r_0 \lll z_R$. This also means that we can use $\xi = r_0/w_0$.

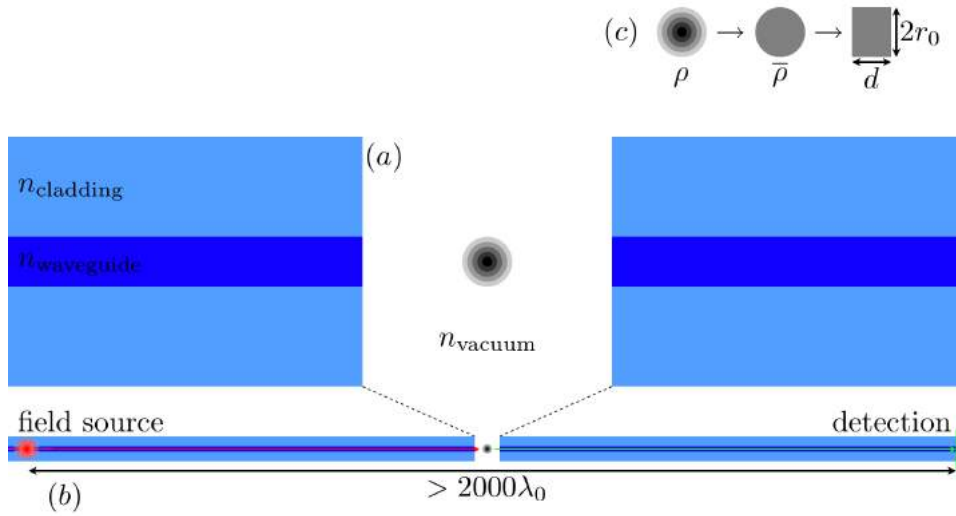


Figure 6.3: Geometry for the MEEP simulation. (a) the system is effectively treated in 2d. The waveguide materials are specified by constant refractive indices. The condensate is built out of shells of inwardly increasing complex refractive index. (b) To let the electric field settle down and stabilize after the field source and after the interaction region, the waveguides extend about 1000 wavelengths in both directions from the trench. (c) to simplify the calculation, the local density of the cloud is flattened into a mean density. Subsequently, the thickness is flattened to an effective thickness.

we use MEEP [106], a finite-difference time-domain simulation software package developed at MIT to model electromagnetic systems. Although it is possible to solve the full set of equations 6.8, we will make a couple of assumptions to reduce computational time.

1. Similar to equation 6.6, we assume that the magnetic permeability μ of the material is close to one: $\mu \approx 1$.
2. The geometry of the system is translational invariant along the z -axis. This effectively reduces the 3d problem to a 2d problem. However this is only true for the width of the waveguide of $4 \mu\text{m}$.
3. The electric field vector lies in the z -direction. We treat the field effectively as scalar and simulations placing the initial electric field vector in the y -direction yield similar results.




Figure 6.4: MEEP simulation example plots of the intensity for different detunings Δ' . The grey scale corresponds to energy density, the coloured background to the value of the relative permittivity ϵ_r . The light comes from the left, enters the trench and interacts with the condensate. These simulations were performed for 3200 atoms, and a condensate of size $r_0 = 1 \mu\text{m}$, $z_0 = 40 \mu\text{m}$.

So far we were only concerned with the light and the atoms. Now we bring in the geometry of the waveguide chip as shown in figure 6.3a. The chip material is specified with an index of refraction of $n_{\text{cladding}} = 1.4530$, the waveguide material with $n_{\text{waveguide}} = 1.0075 \cdot n_{\text{cladding}} = 1.4639$. The parabolic density profile of the cloud is approximated by a large number of circles of varying size and constant density.

To specify the response of a given material MEEP in fact requires the index of refraction to be cast into the real part of the relative electric permittivity (phase part). For the chip materials we use $\Re(\epsilon_r) = n^2$, for the atoms we use the real part of equation 6.5. To account for the imaginary part (extinction part), MEEP requires specification of the D-conductivity,

$$\sigma_D = \omega \cdot \frac{\Im(\epsilon_r)}{\Re(\epsilon_r)}.$$

Example Simulation

Figure 6.4 shows as an example plots of the light energy density (grey) overlain with the refractive index geometry (colours) to see the waveguides and the trench region. Comparing with figure 6.3b, the light source is located a long way to the left to stabilize the field in the waveguide mode. The light enters the trench, interacts with the atoms. A fraction of the light enters the accepting waveguide and is transmitted about 1 mm to the right until it forms a stable mode again. Here the intensity as well as the phase of the light is measured. If compared to the intensity and phase before the trench, we can measure τ^2 and ϕ .

For large detunings (effectively no atoms), we see that most of the light enters the accepting waveguide. However, for near or full on resonant detuning, the path of the light is heavily perturbed by the atoms in addition to random scattering. As a result, less light enters the accepting waveguide. Furthermore, a slight asymmetry can be seen between the two sides around resonance. This corresponds to the refractive index being greater or less than the surrounding medium. For $n > 1$, the medium acts as a converging lens, increasing τ . For $n < 1$, the medium acts as a diverging lens, decreasing τ . This behaviour is not treated in the semi-analytical model discussed next.

6.2.2 Analytical Interaction Model

To better understand the above simulations and to have access to analytical approximations, we create a simple model and fit it to the results of the simulation. The value for the refractive index n is, as in the simulation, given by equation 6.6 but now approximated assuming an optically sparse medium, i.e. $n \approx 1$,

$$n \approx 1 + \frac{\chi}{2}. \quad (6.9)$$

We already reduced the density profile of the cloud to that of a cylinder of mean density $\bar{\rho}$. Furthermore, we accounted for the fraction of light that interacts with the cloud by introducing $O_{\text{quasi-2d}}$ (which we abbreviate as O from now on). We begin with equation 6.7,

$$\tau e^{i\phi} = (1 - O) + O \cdot e^{i\frac{\omega d}{c}(n-1)}, \quad (6.10)$$

we now again approximate to first order taking the limit of an optically sparse medium,

$$\tau e^{i\phi} \approx \tau (1 + i\phi), \quad (6.11)$$

therefore

$$e^{i\frac{\omega d}{c}(n-1)} \approx 1 + i\frac{\omega d}{c}(n-1)$$

and from equation 6.9 follows

$$n-1 \approx \frac{\chi}{2}.$$

Taking this result and substituting d from equation 6.3 yields

$$e^{i\frac{\omega d}{c}(n-1)} \approx 1 - \frac{15}{128\pi} \cdot N_{\text{atoms}} \cdot \frac{\lambda_0^2}{r_0 \cdot z_0} \cdot \frac{1 - 2i\Delta'}{1 + 4\Delta'^2}, \quad (6.12)$$

again using $\omega \approx \omega_0$. Subsequently, we can compare real and imaginary parts of equation 6.10, approximated through equations 6.11 and 6.12, hence write for the transmission amplitude

$$\tau \approx 1 - \frac{N_{\text{atoms}}}{1 + 4\Delta'^2} \cdot K_1 \cdot \mathcal{O} \quad (6.13)$$

and for the phase shift

$$\phi \approx \frac{2\Delta' N_{\text{atoms}}}{1 + 4\Delta'^2} \cdot K_1 \cdot \mathcal{O}, \quad (6.14)$$

both with

$$K_1 = \frac{15}{128\pi} \cdot \frac{\lambda_0^2}{r_0 \cdot z_0}.$$

In figure 6.5 we compare the data gathered from MEEP simulations with the semi-analytical model we derived. The asymmetry around resonance that is apparent in the examples plots in figure 6.4 can be seen here, too. More light is gathered by the accepting waveguide for negative detuning.

The data points are taken for 3200 atoms, a cloud shape of $r_0 = 1 \mu\text{m}$, $z_0 = 40 \mu\text{m}$.

6.3 Prospective Experiments

Now that we have modelled the atom-light-system, we can include the details of the experimental measurement procedure and calculate how well we can estimate the number of atoms in a cloud, or its local density.

We do this for three types of measurements. First, measuring the amount of light absorbed by the atoms. Second, a two-path Mach-Zender interference measurement and finally a two-frequency interference beat-note detection. We then summarise the detection limits of the three types within some parameters limits.




Figure 6.5: Comparison between the simulation data and the semi-analytical model. The fit yields values for the overlap integral of $O = 0.58 \pm 0.02$ and $N_{\text{atoms}} = 6500 \pm 510$.

For all measurements we assume that the detection laser light is well approximated by a coherent state $|\alpha\rangle$ with average photon number expectation value $|\alpha|^2$. Moreover, the detection of a coherent state follows Poisson statistics. The root-mean-square deviation, or variance, of the expected photon number is therefore $|\alpha|$. The electric field expectation of a coherent state [104],

$$\langle \alpha | \hat{E} | \alpha \rangle = E_0 \cdot \cos(\omega t - kx),$$

is identical to the classical description of an oscillating electric field. Hence, we can treat the state of light at different locations in our experiment as an electric field amplitude.

To look at the noise in the signal that we expect on the detector, we recast the field in number of photons to use the fact that the expected variance of our measurement is going to be equal to the expected photon number itself. For a measurement interval of T_γ we write the number of photons penetrating the atomic cloud within the trench as

$$N_\gamma = E_0^2 \cdot T_\gamma \cdot K_2 \cdot K_{\text{qe}},$$

with K_{qe} the quantum efficiency of the detector and⁹

$$K_2 = \frac{\pi w_0^2 c \epsilon_0}{2 \hbar \omega_0} \approx 1.875 \cdot 10^4 \frac{\text{C}^2}{\text{N}^2 \cdot \text{s}}$$

for $w_0 = 1.07 \mu\text{m}$ and $\omega_0 = 2\pi \cdot 384 \text{ THz}$. The number of photons scattered is

$$N_{\text{sc}} = (1 - \tau^2) \cdot N_\gamma, \quad (6.15)$$

which means $\tau^2 N_\gamma$ photons pass through the medium. In the presence of N_{atoms} atoms the number of photons scattered per atom is

$$n_{\text{sc}} = \frac{N_{\text{sc}}}{N_{\text{atoms}}} = (1 - \tau^2) \cdot \frac{N_\gamma}{N_{\text{atoms}}}.$$

6.3.1 Absorption Measurement

We consider a measurement as shown in figure 6.6a. The incoming light field,

$$\mathcal{E}_0 = E_0 \cdot e^{i(\omega t - kx)},$$

is penetrating the atomic cloud. The light field just thereafter is

$$\mathcal{E}_1 = \mathcal{E}_0 \cdot \tau e^{i\phi}$$

and finally the light field just before the detector \mathcal{D}_1 is

$$\mathcal{E}_2 = \mathcal{E}_1 \cdot \eta = \mathcal{E}_0 \cdot \eta \tau e^{i\phi}.$$

Here η is the fractional transmission of the electric field.

Now we are looking at the problem in terms of photon numbers to determine the ultimate detection limit. To measure the photon flux of the probe light we measure for a duration T_0 and find

$$N_0 = \eta^2 \cdot E_0^2 \cdot T_0 \cdot K_2 \cdot K_{\text{qe}}.$$

With atoms in the trench absorbing a light pulse of duration T_1 , the photon signal is

$$N_1 = \eta^2 \tau^2 \cdot E_0^2 \cdot T_1 \cdot K_2 \cdot K_{\text{qe}}.$$

To estimate the absorption and assuming stable light intensities we can write

$$\frac{N_1}{N_0} = \tau^2 \cdot \frac{T_1}{T_0} \approx (2\tau - 1) \cdot \frac{T_1}{T_0} \quad (6.16)$$

⁹Derived from the average intensity of a gaussian beam $\bar{I} = I_{\text{peak}}/2 = P/(\pi w_0^2)$, the relation between local intensity and electric field $I_{\text{local}} = 0.5c\epsilon_0 E_0^2$, and $P = N_\gamma \hbar \omega_0 / T_\gamma$.




Figure 6.6: Principal geometry of measuring a cloud of atoms. (a) a direct absorption measurement. The light is sent through the cloud. (b) Mach-Zender interferometer measurement. The light is split, one part is sent through the cloud and one is used as a reference. After recombination, in phase and quadrature photon numbers are measured. The second beam splitter S_2 we assume splits 50:50.

for $\tau \approx 1$ to arrive at

$$1 - \tau = \frac{1}{2} \left(1 - \frac{N_1 T_0}{T_1 N_0} \right),$$

with a relative uncertainty assuming independent measurements and Poissonian statistics,

$$\frac{\sigma(1 - \tau)}{(1 - \tau)} = \sqrt{\frac{1}{N_0} + \frac{1}{N_1}}.$$

Assuming now that $N_0 \approx N_1$ and $T_0 \approx T_1$ we find¹⁰

$$\sigma(1 - \tau) \approx \sqrt{\frac{2}{N_0}}$$

for the uncertainty in the absorption signal and

$$\sigma(N_{\text{atoms}}) \approx \frac{1 + 4\Delta^2}{K_1 \cdot K_{\text{qe}} \cdot \mathcal{O}} \cdot \sqrt{\frac{2}{N_0}} \quad (6.17)$$

¹⁰If we wanted we could lower the uncertainty by a factor of $1/\sqrt{2}$ by increasing T_0 for the reference measurement, thereby increasing N_0 .

for the uncertainty in inferred number of atoms. This result is useful for a real measurement. If instead we want to express this result in terms of the number of photons scattered, we write

$$\sigma(N_{\text{atoms}}) \approx \frac{1 + 4\Delta'^2}{\eta \cdot K_1 \cdot K_{\text{qe}} \cdot \mathcal{O}} \cdot \sqrt{\frac{2}{N_\gamma}}$$

to include the number of photons hitting the atoms. Substituting N_γ with 6.15 in the limit $\tau \approx 1$ and using 6.13 we can rewrite the uncertainty as

$$\sigma(N_{\text{atoms}}) = \frac{2}{\eta K_{\text{qe}}} \sqrt{\frac{1 + 4\Delta'^2}{n_{\text{sc}} \cdot K_1 \cdot K_{\text{qe}} \cdot \mathcal{O}}}$$

One thing to notice is that using a longer interval for the reference measurement, the uncertainty can be lowered by a factor of about $1/\sqrt{2}$ compared with using equal measuring durations.

Non-Linear Regime

So far we were concerned with the linear regime of low optical densities. If Δ' approaches zero, though, the assumption $\tau \approx 1$ is not valid for realistic parameters, where the optical density can be high. We return to equation 6.10, set $\Delta' = 0$ and re-write it as

$$\tau = 1 - \mathcal{O} + \mathcal{O} \cdot e^{-\frac{1}{16} K_1 K_{\text{qe}} \mathcal{O} N_{\text{atoms}}} = 1 + \mathcal{O} \cdot (e^{-\frac{1}{16} K_1 K_{\text{qe}} \mathcal{O} N_{\text{atoms}}} - 1).$$

Solving this for N_{atoms} yields

$$N_{\text{atoms}} = -\frac{16}{K_1 K_{\text{qe}} \mathcal{O}} \cdot \ln\left(1 - \frac{1 - \tau}{\mathcal{O}}\right)$$

for the true estimator for τ^2 from equation 6.16,

$$\tau = \sqrt{\frac{N_1 T_0}{T_1 N_0}}$$

or $\tau = \sqrt{n_1/n_0}$ if expressed in photon count rates.

The uncertainty in atom number dependent on the measured photon numbers N_0 and N_1 is given by

$$\sigma(N_{\text{atoms}})^2 = \left(\frac{\partial N_{\text{atoms}}}{\partial N_0}\right)^2 \cdot \sigma(N_0)^2 + \left(\frac{\partial N_{\text{atoms}}}{\partial N_1}\right)^2 \cdot \sigma(N_1)^2.$$

The partial derivatives are

$$\begin{aligned}\frac{\partial N_{\text{atoms}}}{\partial N_0} &= -\frac{16}{K_1 K_{\text{qe}} \mathcal{O} N_0} \cdot \left(1 + (\mathcal{O}^{-1} - 1) e^{\frac{1}{16} K_1 K_{\text{qe}} \mathcal{O} N_{\text{atoms}}}\right) \\ &\approx -\frac{e^{\frac{1}{16} K_1 K_{\text{qe}} \mathcal{O} N_{\text{atoms}}}}{N_0} \cdot \frac{16}{K_1 K_{\text{qe}} \mathcal{O}} \cdot (1 - \mathcal{O})\end{aligned}$$

and

$$\begin{aligned}\frac{\partial N_{\text{atoms}}}{\partial N_1} &= \frac{16}{K_1 K_{\text{qe}} \mathcal{O} N_1} \cdot \left(1 + (\mathcal{O}^{-1} - 1) e^{\frac{1}{16} K_1 K_{\text{qe}} \mathcal{O} N_{\text{atoms}}}\right) \\ &\approx \frac{e^{\frac{1}{16} K_1 K_{\text{qe}} \mathcal{O} N_{\text{atoms}}}}{N_1} \cdot \frac{16}{K_1 K_{\text{qe}} \mathcal{O}} \cdot (1 - \mathcal{O})\end{aligned}$$

approximated for high optical densities, i.e. $K_1 N_{\text{atoms}} \gg 8$. Using Poissonian statistics, we find $\sigma N_0^2 = N_0$ and $\sigma N_1^2 = N_1$, resulting in a total uncertainty in atom number of

$$\sigma(N_{\text{atoms}}) \approx e^{\frac{1}{16} K_1 K_{\text{qe}} \mathcal{O} N_{\text{atoms}}} \cdot \sqrt{\frac{1}{N_0} + \frac{1}{N_1}} \cdot \frac{16}{K_1 K_{\text{qe}} \mathcal{O}} \cdot (1 - \mathcal{O}).$$

The exponential dependence in this result is what makes an absorption measurement at high optical densities useless. We could, however, still measure at low optical densities while keeping both sensitivity and number of scattered photons high. The result 6.17 we found in assuming low absorption by using off-resonance light is valid. However, maybe we can do better by looking at the phase shift 6.14 which the atoms impart on the light in addition to scattering photons.

6.3.2 Mach-Zender Interferometer

Let us look at an interferometer as shown in figure 6.6b. We do not need to perform a fully second-quantised analysis of the interferometer with photons. However, as long as the light is in a coherent state, i.e. a state whose dynamics most closely resemble the oscillating behaviour of a classical harmonic oscillator, the results match up with those of a classical treatment [95]. The incoming light field is again

$$\mathcal{E}_0 = E_0 \cdot e^{i(\omega t - kx)},$$

however this time split by a beam-splitter cube \mathcal{S}_1 . The fraction of light intensity reflected r can be adjusted using a waveplate in front of \mathcal{S}_1 . An

intensity fraction $(1 - r)$ is transmitted. The fields after the cube are

$$\begin{aligned}\mathcal{E}_1 &= \sqrt{1-r} \cdot \mathcal{E}_0, \\ \mathcal{E}_2 &= \sqrt{r} \cdot e^{i\pi} \cdot \mathcal{E}_0, \\ \mathcal{E}_3 &= e^{i\pi} \cdot \mathcal{E}_1 = \sqrt{1-r} \cdot e^{i\pi} \cdot \mathcal{E}_0, \\ \mathcal{E}_4 &= e^{i\pi} \cdot \mathcal{E}_2 = \sqrt{r} \cdot \mathcal{E}_0.\end{aligned}$$

Now the field \mathcal{E}_4 interacts with the atoms,

$$\mathcal{E}_5 = \tau e^{i\phi} \cdot \mathcal{E}_4 = \sqrt{r}\tau e^{i\phi} \cdot \mathcal{E}_0,$$

and leaves the chip in the form

$$\mathcal{E}_6 = \eta \cdot \mathcal{E}_5 = \eta \sqrt{r}\tau e^{i\phi} \cdot \mathcal{E}_0.$$

The fields \mathcal{E}_3 and \mathcal{E}_6 arrive on the second beam-splitter \mathcal{S}_2 which is assumed to be balanced ($r = 0.5$). The fields combine and after \mathcal{S}_2 are

$$\mathcal{E}_7 = \frac{1}{\sqrt{2}} \left(e^{i\pi} \cdot \mathcal{E}_3 + \mathcal{E}_6 \right) = \frac{\mathcal{E}_0}{\sqrt{2}} \left(\sqrt{1-r} + \eta \sqrt{r}\tau e^{i\phi} \right)$$

and

$$\mathcal{E}_8 = \frac{1}{\sqrt{2}} \left(\mathcal{E}_3 + e^{i\pi} \cdot \mathcal{E}_6 \right) = \frac{\mathcal{E}_0}{\sqrt{2}} \left(\sqrt{1-r} - \eta \sqrt{r}\tau e^{i\phi} \right).$$

The intensities are

$$I_7 = \frac{c\epsilon_0}{2} \cdot \mathcal{E}_7^2 = \frac{c\epsilon_0 \mathcal{E}_0^2}{4} \cdot \left(1 - r \left(1 - \eta^2 \tau^2 \right) + 2\eta\tau \sqrt{r(1-r)} \cos \phi \right)$$

and

$$I_8 = \frac{c\epsilon_0}{2} \cdot \mathcal{E}_8^2 = \frac{c\epsilon_0 \mathcal{E}_0^2}{4} \cdot \left(1 - r \left(1 - \eta^2 \tau^2 \right) - 2\eta\tau \sqrt{r(1-r)} \cos \phi \right).$$

From here we can calculate the total transmitted light,

$$I_{\text{total}} = I_7 + I_8 = \frac{c\epsilon_0 \mathcal{E}_0^2}{2} \cdot \left(1 - r \left(1 - \eta^2 \tau^2 \right) \right),$$

and fringe visibility,

$$\mathcal{V} = \frac{\max(I_8) - \min(I_8)}{\max(I_8) + \min(I_8)} = \frac{2\eta\tau \sqrt{r(1-r)}}{1 - r \cdot (1 - \eta^2 \tau^2)}.$$

Here we see the advantage of including the possibility to tune the first beam-splitter. We can find a value of r for a given loss coefficient η to optimise the fringe visibility. Choosing

$$r_{\max} = \frac{1}{1 + \eta^2}$$

leads to a maximum value of the fringe visibility of

$$\mathcal{V}_{\max} = \frac{2\tau}{1 + \tau^2}.$$

To improve the small-signal sensitivity, let us shift the phase of the two light beams, such that the intensities are on the slope of a fringe: $\phi \rightarrow \phi + \phi_0$. Using trigonometric addition theorems and again employing the assumption of an optically thin medium ($\phi \ll 1, \tau \approx 1$), we can set $\phi_0 = \pi/2$ and write

$$I_7 = \frac{c\epsilon_0\mathcal{E}_0^2}{4} \left(1 - r(1 - \eta^2) + 2\eta\phi\sqrt{r(1-r)} \right)$$

and

$$I_8 = \frac{c\epsilon_0\mathcal{E}_0^2}{4} \left(1 - r(1 - \eta^2) - 2\eta\phi\sqrt{r(1-r)} \right).$$

We write an estimator in terms of measured photon numbers, since the number of photons is directly proportional to the intensity,

$$\begin{aligned} \frac{N_8 - N_7}{N_8 + N_7} &= \frac{I_8 - I_7}{I_8 + I_7} = \frac{2\eta\sqrt{r(1-r)}}{1 - r(1 - \eta^2)}\phi \\ &= \frac{2\eta\sqrt{r(1-r)}}{1 - r(1 - \eta^2)} \cdot \frac{2\Delta'}{1 + 4\Delta'^2} \cdot N_{\text{atoms}} \cdot K_1 K_{\text{qe}} \mathcal{O}. \end{aligned}$$

Solving for the number of atoms yields

$$N_{\text{atoms}} = \frac{1 + 4\Delta'^2}{2\Delta'} \cdot \frac{N_8 - N_7}{N_8 + N_7} \cdot \frac{1}{\kappa K_1 K_{\text{qe}} \mathcal{O}'}$$

where we substitute

$$\kappa = \frac{2\eta\sqrt{r(1-r)}}{1 - r(1 - \eta^2)}.$$

Remember that we can always maximise κ by choosing $r_{\max} = 1/(1 + \eta^2)$, $\kappa(r = r_{\max}) = 1$. The uncertainty in the inferred phase is given by

$$\sigma(\phi)^2 = \left(\frac{\partial\phi}{\partial N_7} \cdot \sigma(N_7) \right)^2 + \left(\frac{\partial\phi}{\partial N_8} \cdot \sigma(N_8) \right)^2.$$

The partial derivatives are

$$\frac{\partial\phi}{\partial N_7} = -\frac{2N_8}{(N_8 + N_7)^2} \cdot \frac{1}{\kappa K_{qe}}$$

and

$$\frac{\partial\phi}{\partial N_8} = \frac{2N_7}{(N_8 + N_7)^2} \cdot \frac{1}{\kappa K_{qe}},$$

with $\sigma(N_7)^2 = N_7$ and $\sigma(N_8)^2 = N_8$. Inserting the derivatives into the expression for $\sigma(\phi)^2$ yields

$$\sigma(\phi)^2 = \frac{4N_8N_7}{(N_8 + N_7)^3} \cdot \frac{1}{\kappa^2 K_{qe}^2}.$$

For similar photon signals $N_7 \approx N_8 \approx N_\gamma/2$ we can simplify this expression to

$$\sigma(\phi) \approx \frac{1}{\kappa K_{qe} \sqrt{N_\gamma}}.$$

Similar to the absorption case, we can rewrite the uncertainty in atom number in terms of n_{sc} ,

$$\sigma(N_{\text{atoms}}) = \sqrt{\frac{1 + 4\Delta'^2}{4\Delta'^2}} \cdot \frac{1}{\sqrt{n_{sc}}} \cdot \sqrt{\frac{4}{\kappa^2 K_{qe}^2 K_1 \mathcal{O}}}.$$

Remember with this result, that $\kappa(\eta)$. Also note that for large values of the detuning Δ' , $\sigma(N_{\text{atoms}})$ is independent of the detuning.

6.3.3 Beat Note Detection

As shown, interferometry, as compared to the sole measurement of absorbed light, is a way to reduce the uncertainty in the estimation of e.g. the number of atoms measured in a sample. To improve even further on the generic model of the Mach-Zender interferometer, we have a closer look at a more specialised type of detection, i.e. beat note detection [107, 108], in this section.

A schematic for beat note detection is shown in figure 6.7 where monochromatic light at frequency ω_c is modulated with a voltage controlled signal at frequency ω_m to create two sidebands around the carrier frequency. The carrier itself in this scheme must be suppressed in this scheme to prevent resonant spontaneous emission. The sidebands interact with the atoms as shown in the previous section and are finally measured with a photo detector. The phase shift of the measured beat note signal is then compared to

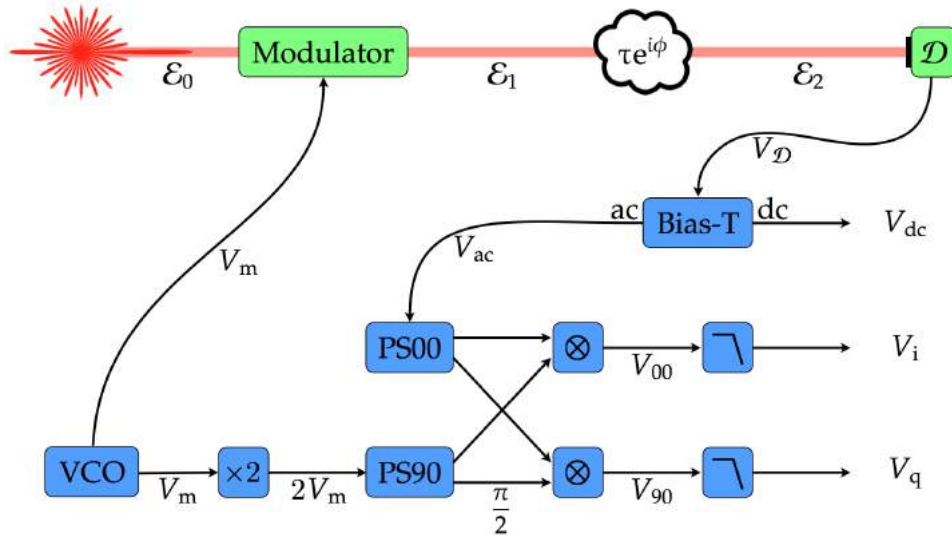


Figure 6.7: Possible setup of a beat-note phase detection scheme. (*top*) Monochromatic light (\mathcal{E}_0) passes through a modulation stage suppressing the carrier and creating two sidebands (\mathcal{E}_1), which in turn interact with the atomic sample. After the sample the light (\mathcal{E}_2) is detected (\mathcal{D}) and transformed into a voltage. (*bottom*) The measured voltage ($V_{\mathcal{D}}$) is split into its ac and dc parts (V_{ac} and V_{dc}). The ac part is split and mixed with a frequency-doubled original modulation signal ($2V_m$) to retrieve in phase and quadrature signals (V_i and V_q) after low pass filtering multiples of the modulation frequency ω_m .

the original modulation signal, effectively transforming the phase changes at optical frequencies into changes at radio or microwave frequencies. In the end, we find signals that correspond to the dc, in phase and quadrature parts of the beat note.

Beat note detection as compared to an optical interferometer has several advantages to offer. Due to the interferometry being performed with much longer wavelength, the stability of the system is much improved. The use of two balanced sidebands in this scheme doubles the signal without doubling the noise.

Disadvantages, or better problems that need addressing, are that for the beat note detection a fast photo detector and low-noise IQ-demodulation electronics are required. In addition, it is not trivial to generate the two side bands and at the same time suppress the carrier. The best step in that

direction may be a using double pass AOM as the modulator carrying two frequencies.

Optical Path

To derive the expected values of the beat note signal, we write the electrical field of the light as it enters the system as

$$\mathcal{E}_0(x, t) = E(x) \cdot \cos(\omega_c t) = \frac{E(x)}{2} \cdot e^{-i\omega_c t} + \frac{E(x)}{2} \cdot e^{i\omega_c t} = \mathcal{E}_0^{(+)} + \mathcal{E}_0^{(-)}.$$

For this derivation we can only consider the first part. However, when calculating the intensities we have to remember to include the complex conjugate part.

The light enters the modulator and is transformed into two sidebands. For simplicity we assume that the relative phase shift between the two sidebands is either 0 or π . If the carrier frequency is ω_c the new frequencies are $(\omega_c + \omega_m)$ and $(\omega_c - \omega_m)$, where ω_m is half the difference between the frequencies of the two sidebands. The field can be written as

$$\mathcal{E}_1^{(+)} = \frac{E(x)}{\sqrt{8}} \left(e^{-i(\omega_c + \omega_m)t} + P e^{-i(\omega_c - \omega_m)t} \right).$$

Here, P accounts for the two possible phase shift values between the sidebands. $P = +1$ if in phase, $P = -1$ if in antiphase. We pick here $P = +1$ without loss of generality.

Let us set the carrier frequency slightly off from the atomic resonance frequency ω_0 ,

$$\omega_c = \omega_0 + \Delta' \Gamma,$$

where $\Delta' = 0$ means, that the sidebands lie symmetrically around the atomic resonance frequency. We can rewrite this as

$$\mathcal{E}_1^{(+)} = \frac{E(x)}{\sqrt{8}} e^{-i(\omega_0 + \Delta' \Gamma)t} \left(e^{-i\omega_m t} + e^{i\omega_m t} \right).$$

Defining

$$\tau_{\pm} = \tau(\Delta' \pm \omega_m / \Gamma)$$

and

$$\phi_{\pm} = \phi(\Delta' \pm \omega_m / \Gamma)$$

as the transmission amplitude and phase shift at the frequencies of the sidebands, we can write down the field after the atomic sample as

$$\mathcal{E}_2^{(+)} = \frac{E(x)}{\sqrt{8}} e^{-i(\omega_0 + \Delta'\Gamma)t} (\tau_+ e^{-i\omega_m t + i\phi_+} + \tau_- e^{i\omega_m t + i\phi_-}).$$

This is the field that arrives at the photo detector, where it is measured in form of a voltage that is proportional to the intensity of the light. The proportionality factor is the gain G of the detection system, which includes a detector quantum efficiency K_{qe} ,

$$V_{\mathcal{D}} = G \cdot \langle \mathcal{E}_2 \cdot \mathcal{E}_2^* \rangle = G \cdot \langle (\mathcal{E}_2^{(+)} + \mathcal{E}_2^{(-)}) \cdot (\mathcal{E}_2^{(+)} + \mathcal{E}_2^{(-)})^* \rangle,$$

which expands to

$$V_{\mathcal{D}} = \frac{GE^2(x)}{4} (\tau_+^2 + \tau_-^2 + 2\tau_+\tau_- \cos(2\omega_m t + \Delta\phi)) \\ + \frac{GE^2(x)}{4} (\tau_+^2 \cos(2\omega_m t) + \tau_-^2 \cos(2\omega_m t) + 2\tau_+\tau_- \cos(\Delta\phi))$$

when we average over a time longer than $2\pi/\omega_0$ and assume $\Delta'\Gamma \ll \omega_0$ as well as $\omega_m \ll \omega_0$. We also use $\Delta\phi = \phi_+ - \phi_-$. We can rewrite above equation as

$$V_{\mathcal{D}} = \frac{GE^2(x)}{2} (\tau_+^2 + \tau_-^2 + 2\tau_+\tau_- \cos[2\omega_m t + (\phi_+ - \phi_-)]).$$

Electronic Path

Once we have the signal in electronic form we can split-off the dc part with a bias-T,

$$V_{dc} = V_0 (\tau_+^2 + \tau_-^2),$$

with $V_0 = GE^2(x)/2$. The remainder is the ac part of the signal,

$$V_{ac} = 2V_0 \tau_+\tau_- \cos(2\omega_m t + \Delta\phi).$$

We split the ac part of the signal into two parts with a balanced power splitter and mix one part with twice the frequency of the original modulation voltage $V_m = V_{m,0} \cdot \cos(\omega_m t)$. For the other part, we use this same voltage but shifted by $\pi/2$,

$$V_{00} = \frac{V_{ac}}{2} \cdot V_{m,0} \cos(\omega_m t) = \frac{V_0 V_{m,0}}{2} \tau_+\tau_- (\cos(\Delta\phi) + \cos(4\omega_m t + \Delta\phi))$$

and

$$V_{90} = \frac{V_{ac}}{2} \cdot V_{m,0} \sin(\omega_m t) = \frac{V_0 V_{m,0}}{2} \tau_+ \tau_- (\sin(\Delta\phi) - \sin(4\omega_m t + \Delta\phi)).$$

Sending these voltages through a low-pass filter to filter-out the components around $4\omega_m$ yields the in phase and quadrature voltages,

$$V_i = V' \cdot \tau_+ \tau_- \cos(\phi_+ - \phi_-)$$

and

$$V_q = V' \cdot \tau_+ \tau_- \sin(\phi_+ - \phi_-),$$

with $V' = V_0 V_{m,0}/2$. Figure 6.8 shows plots of the lineshapes created when scanning the laser frequency across resonance. The results shown are for relatively small atoms numbers, i.e. $N_{atoms} = 100$.

Linear Regime

If the absorption and phase shift induced by the atoms is small, we are back in the linear limit and can approximate $\cos \phi \approx 1 - \phi^2/2$ and $\sin \phi \approx \phi$. We use equations 6.13 and 6.14 and write

$$V_{dc} = V_0 \left[\left(1 - \frac{N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' + \Omega_m)^2} \right)^2 + \left(1 - \frac{N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' - \Omega_m)^2} \right)^2 \right],$$

$$V_i = V' \left(1 - \frac{N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' + \Omega_m)^2} \right) \left(1 - \frac{N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' - \Omega_m)^2} \right) \\ \times \left(1 - \frac{1}{2} \left(\frac{2(\Delta' + \Omega_m) N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' + \Omega_m)^2} - \frac{2(\Delta' - \Omega_m) N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' - \Omega_m)^2} \right)^2 \right),$$

and

$$V_q = V' \left(1 - \frac{N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' + \Omega_m)^2} \right) \left(1 - \frac{N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' - \Omega_m)^2} \right) \\ \times \left(\frac{2(\Delta' + \Omega_m) N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' + \Omega_m)^2} - \frac{2(\Delta' - \Omega_m) N_{atoms} K_1 \mathcal{O}}{1 + 4(\Delta' - \Omega_m)^2} \right),$$

with $\Omega_m = \omega_m/\Gamma$. To first order in N_{atoms} we can approximate to find

$$V_{dc} = 2V_0 [1 - N_{atoms} K_1 \mathcal{O} (\Omega_+ + \Omega_-)],$$

$$V_i = 2V' [1 - N_{atoms} K_1 \mathcal{O} (\Omega_+ + \Omega_-)],$$

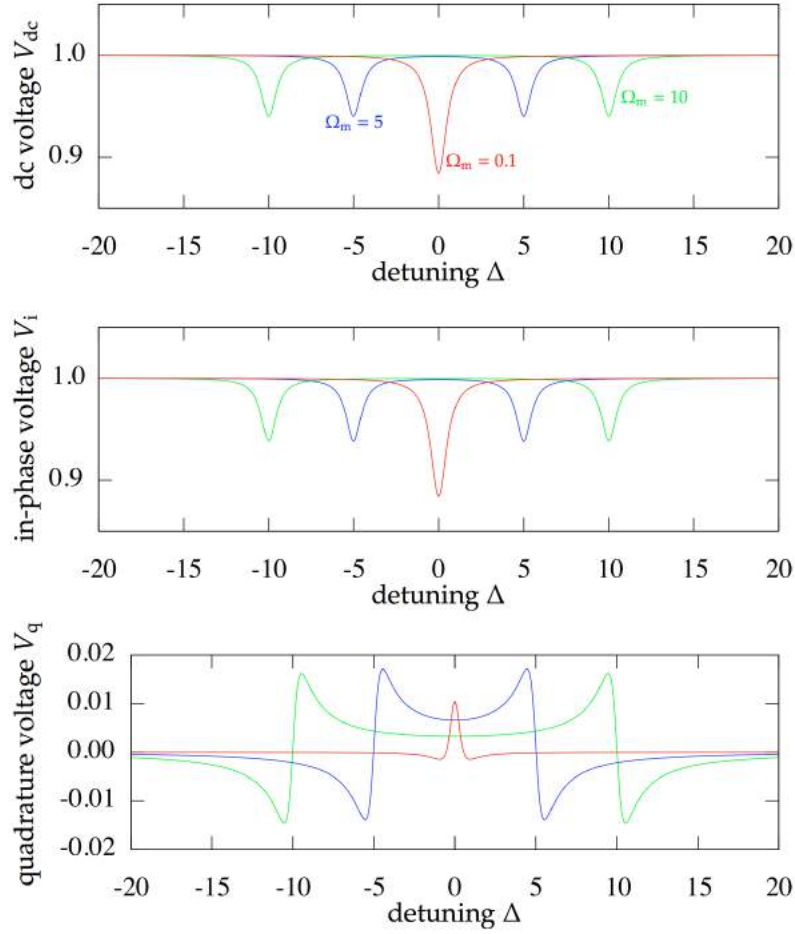


Figure 6.8: Normalised signals for 100 atoms in a BEC of dimensions $r_0 = 1 \mu\text{m}$ and $z_0 = 40 \mu\text{m}$, scanning the carrier frequency over resonance for three different modulation frequencies.

and

$$V_q = 4V'N_{\text{atoms}}K_1\mathcal{O}(\Delta'_+\Omega_+ - \Delta'_-\Omega_-),$$

using $\Delta'_\pm = \Delta' \pm \Omega_m$ and

$$\Omega_\pm = \frac{1}{1 + 4\Delta'_\pm{}^2}.$$

Note that for a resonant carrier, i.e. $\Delta' = 0$, the sidebands lie symmetrically around resonance and $\Omega_+ = \Omega_-$. Above results simplify to

$$V_{\text{dc}} = 2V_0 \left[1 - \frac{2N_{\text{atoms}}K_1\mathcal{O}}{1 + 4\Omega_m^2} \right]$$




Figure 6.9: Beat note signals for carrier frequency on resonance, i.e. $\Delta' = 0$, varying the number of atoms for three different modulation frequencies.

$$V_i = 2V' \left[1 - \frac{2N_{\text{atoms}}K_1\mathcal{O}}{1 + 4\Omega_m^2} \right]$$

$$V_q = 2V' \left[\frac{4N_{\text{atoms}}K_1\mathcal{O}}{1 + 4\Omega_m^2} \right] \cdot \Omega_m$$

In figure 6.9 the expected signals are plotted for $\Delta' = 0$ against atom number, together with the linear approximations. For small modulation frequencies, both the absorption and in phase measurements are sensitive to the atom number. However, since we expect real samples to be optically thick, we want to measure away from resonance, e.g. $\Omega_m > 2$. In this case, the quadrature signal is the most sensitive. The atom number inferred from the quadrature measurement is

$$N_{\text{atoms}}(V_q) = \left[\frac{1 + 4\Omega_m^2}{8V'K_1\mathcal{O}\Omega_m} \right] \cdot V_q,$$

where we have to remember that $V' \propto K_{qe}$. Assuming now that the uncertainty in the full-scale reading of a signal is the same as the fractional uncertainty in the total estimated number of photons times $\sqrt{2}$ due to using

two channels for the measurement, we can write the uncertainty as

$$\sigma(N_{\text{atoms}}) = \left[\frac{1 + 4\Omega_m^2}{8V'K_1O\Omega_m} \right] \cdot \sqrt{\frac{2}{N_\gamma}}$$

or in terms of photons scattered

$$\sigma(N_{\text{atoms}}) = \sqrt{\frac{1 + 4\Omega_m^2}{4\Omega_m^2}} \cdot \sqrt{\frac{1}{4K_1OV'^2}} \cdot \frac{1}{\sqrt{n_{\text{sc}}}}$$

6.3.4 Summary of the Techniques

Measuring the density of the atomic sample via its absorption of photons is a simple method to set up, but requires very low noise detection electronics to capture and register the low number of photons sent through the sample if one is to detect non-destructively. In contrast, the Mach-Zender interferometer has to be set up more carefully to limit technical noise introduced into the system but allows for a higher photon number to be detected and thus requires a less sensitive detector. A good compromise of the two methods, the beat note method uses standard low noise electronics to let the two signals interfere, offering a compromise between technical simplicity and detection requirements. In all three cases, if one chooses the detuning correctly, the uncertainty in atom number is proportional to the inverse square root of the number of photons scattered

$$\sigma(N_{\text{atoms}}) \propto \frac{1}{\sqrt{n_{\text{sc}}}}$$

with only a slight change in the pre-factor. In the case of absorption measurement the light is on resonance ($\Delta' = 0$), for the Mach-Zender interferometer and beat note detection it is far off resonance ($\Delta' \gg 1$).

In conclusion non-destructively detecting at the photon shot-noise limit (or at least close) is more promising with either Mach-Zender or beat note detection methods, because the single momentum transfers from each photon to an individual atom is smaller than by resonant detection where absorption of a photon always transfers the full photon momentum to the atom.

The stability intrinsic to most modulation techniques in comparison to light-interferometers makes them easier to scale to multiple devices which we like to use with the waveguide chip. Therefore, we are going to implement a modulation technique. The atomic density sensitivity resulting from

this technique may be higher than atomic shot noise for much less than 1 photon scattered per atom. This sensitivity may enable us to perform very interesting experiments indeed.

Chapter 7

Conclusion and Outlook

In conclusion, a machine to detect cold atoms with an integrated multi-channel optical-waveguide device was presented. The current structure on the atom chip was described which is used to trap cold atoms with magnetic fields, to cool them further down and move them into the waveguide chip trench. In this context was presented how the optical-waveguide atom chip, as well as the vacuum and detection systems were built. Also presented were experiments showing the waveguide chip as a device capable of detecting a vapour of cold atoms and used these measurements to calibrate the state of the light within the waveguide chip trench. Finally, the possibilities of photon-shot-noise-limited detection for our setup was theoretically examined and it was concluded that is possible with an adapted frequency modulation technique.

Looking at possible experiments in the imminent future, a better way to connect fibres to the waveguide chip, including replacing the fibre-vacuum feedthroughs, has the highest priority to increase detection efficiency and to have access to a higher number of waveguides simultaneously. In turn, replacing the MOT-coils with a U-plate should lead to better vacuum condition, due to decreased outgassing of insulated components inside the vacuum chamber. In addition, elimination of the outer MOT-coil former should maximise optical access to the trench and the space beneath it. Finally, characterising and then using frequency modulation detection, as concluded in chapter 6, should provide a powerful tool to near non-destructively detect atoms inside the waveguide chip trench.

These integrated multi-channel optical-waveguide device suggest new possibilities for engineering the quantum states of matter and light on a mi-

microscopic scale, opening the door to new information-processing methods. Using a scalable device, multi-cell traps can form single atom registers, a building block for quantum information processing. The sensitivity achievable allows to measure clock shifts, increasing the accuracy of the atomic gas cell standard, leading to a better clock and e.g. better location accuracy for space flight and global positioning systems. Non-destructive measurements on atoms together with fast computer response allow for feedback control of parameters of a BEC, like atom number, local density, etc. The tightly-spaced linear array of the multi-channel device together with non-destructive measurements facilitates the spatially resolved measurement of an elongated BEC aligned along the array. This way, dynamical local excitations within the cloud are possible to be induced and observed along the length of the condensate. Interesting times are ahead.

List of Figures

1.1	Sketch of the gold-coated optical-waveguide chip, fabricated on a silicon substrate. Fibres are connected to both sides of the chip. The waveguides converge in the centre of the chip where a trench has been cut. There, atoms can be detected between opposing waveguides.	9
2.1	The bulk current structure used to numerically calculate the total magnetic field. The rest of the H-structure, which the Z-structure (<i>red</i>) is part of, is not shown. The bias wires (<i>blue</i>) create the field in the x -direction to counter the field created by the Z-structure. Currents sent through the end structure (<i>green</i>) tweak the end points of the elongated trap, hence the name. The origin (<i>white</i>) of the coordinate system used lies in the centre of its symmetry but on top of the Z-structure surface.	14
2.2	The unit for all spatial coordinates x , y and z is millimeter. The magnetic field for example currents of $I_Z = 34$ A and $I_{\text{bias}} = 80$ A in three cuts through the field minimum $y_{\text{min}} = 0.493$ mm. The field value at this point is $ \mathbf{B}_{\text{min}} = 0.2$ mT. Its vector is lying all in the z -direction. (<i>left</i>) The magnitude of the field is plotted in units of mT and (<i>right</i>) the projection of the field vectors in arbitrary units. The elongated shape of the potential is visible, as well as the tilt of trap axis in the x - z cut. Overlaying the vector plots is the Z-structure in blue.	15

2.3	The distance r_{\min} is taken from the origin of the Z-structure (c.f. figure 2.1). For the considered geometry and $\mathbf{B}_{\text{bkg}} = 0$, all components but the y -component are zero. Therefore, we can express the location by plotting just y_{\min} , to specify the location of the trap minimum. Of special interest is the location of the waveguide chip trench and the three parameter sets from table 2.1. These points and isodistance lines for $y_{\min} = 0.5$ mm, 1.0 mm and 1.5 mm are drawn in this and all following parameter maps for comparison.	20
2.4	Plot of the relative horizontal displacement $\tan \vartheta$ of the axial principal axis of the magnetic trap due to its rotation around the y -axis. The amount of displacement is decreasing when closing in to the chip. This is mainly due to adjusting B_z to keep the trap bottom at 0.2 mT. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.	21
2.5	Plot of the geometric mean radial oscillation frequency $2\pi \times \omega_r$. Due to setting B_z to yield a trap bottom of 0.2 mT throughout the map, the radial oscillation frequency is increasing with decreasing distance of the trap. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.	22
2.6	Plot of the axial oscillation frequency $2\pi \times \omega_a$. Due to the control of B_z to a trap bottom of 0.2 mT the trap axially maintains the same shape, resulting in the axial oscillation frequency to plateau at about 11.5 Hz when coming closer to the chip. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.	23
2.7	Plot of the depth of the trap given in units of temperature. The line denotes the waveguide chip surface. Without a chip the trap depth is limited by the bias field, however, closing in to the chip opens up a second escape route for the atoms by hitting the chip's surface. The green line shows the isodistance line for $y_{\min} = 0.5$ mm.	26
2.8	The trap parameters: trap depth and relative displacement when adding a current in the end wires of the current structure. Current values for I_z and I_{bias} , as well as B_z values set to control the trap bottom, are taken from the comparison table 2.1.	28

3.1	Hyperfine structure and transitions of the ground state ($5^2S_{1/2}$) and D_2 excited state ($5^2P_{3/2}$) of rubidium 87. (<i>left</i>) the red and green arrows denote the reference and re-pumping transitions, respectively. Blue arrows show allowed radiative decay routes. (<i>right</i>) the red arrows denote the σ^+ optical pumping light increasing m_F by 1 and the allowed radiative decay routes of each excited state as blue arrows. Since there is one state without an excitation arrow or a decay path, atoms are pumped into the $F = 2, m_F = 2$ state. Polarisation-resolved ¹ repumper transitions, though not shown on the right, bring atoms which have through spontaneous emission decayed into the ground state back into the $5^2P_{3/2} F = 2$ excited state.	31
3.2	Photograph of the chamber sitting on the table. Into the table a hole is cut to allow optical access to the chamber from below. In the background, on the left are the ion and getter pumps, as well as an ion pressure gauge. On the right sits the LVIS, separable from the main chamber via a gate valve. So far, the atom chip flange is not attached and replaced by a blind flange for initial hot baking.	34
3.3	Sketch of the chamber seen from below (<i>left</i>) and the side (<i>right</i>). Except for the ports that lead to the LVIS and pumps, the sides of the chamber are closed with viewports. Cooling light is drawn in red, imaging light in green. Magnetic bias field coils are placed along the x, y and z axes. The first imaging beam traverses the chamber and is shone onto a CCD camera. The second imaging beam is goes through the bottom window and is reflected at 45° (not visible in this projection) from the waveguide chip before detection on a second CCD camera.	35

3.4	Layout of the atom chip flange that sits on top of the chamber. It features electrical feedthroughs for the quadrupole coils (<i>a</i>) that form the MOT, the end wire structures (<i>b</i>), the H-structure (partly covered by the waveguide chip) (<i>c</i>), the atom chip bias field wires (<i>d</i>), dispensers (only the feedthroughs shown here) (<i>e</i>). In addition, there are optical feedthroughs (<i>f</i>) to bring light to the waveguide chip. The side cross section view of the flange also shows that one of the quadrupole coils is embedded into the copper block (<i>g</i>), another is mounted on an aluminium support (<i>h</i>).	37
3.5	Photograph of the atom chip flange with atom chip and connections.	39
3.6	Photograph of the installation of the atom chip flange. The flange is carefully lowered onto and the atom chip assembly into the octagonal vacuum chamber.	40
3.7	Photograph of the experiment now. Underneath the table are some of the MOT optics already present. On the table, the camera can be seen. The atom chip flange is already in place.	41
3.8	Detection setup. The light is split into parts and injected into the fibres (<i>left</i>) connected to the waveguide chip (<i>top</i>). After traversing the trench, the light coming out of one fibre is collimated, then focussed onto an SPAPD (<i>right</i>). The second imaging axis consists of imaging light that is shone onto the atoms at an angle of 45°. The reflected light is focussed onto a CCD camera (<i>middle</i>).	42
3.9	Data flow chart of the experimental control. (<i>blue arrows</i>) the user interacts with the user interface as part of the software suite. The user interface serializes the input and hands it down to the main hardware controller, which in turn delegates tasks to several slave data acquisition controllers. All controllers then ready their hardware for the next experimental cycle. (<i>red arrows</i>) During the cycle, data is collected by the hardware, then processed by the slave controllers and passed on up through the chain back to the user interface to be accessed by the user. (<i>yellow arrows</i>) On request, all software modules save data and meta data to the computer hard drive.	45

3.10	A screen grab of a typical window layout for performing experiments with Tyche. Two data acquisition controllers are shown on the left. One is for the camera, showing an absorption image of an atomic cloud. An approximate atom number is calculated on-the-fly (223 million in this case). The other DAC is for the counter input card connected to the SPAPD. Each shot is displayed on top and on-the-fly averaging of a number of shots is performed below. TycheUI is just a command prompt executing a loop in this grab. TycheController is really a backend to connect all other modules. It only has a window for status information output.	46
4.1	The fully assembled atom chip (<i>artist's impression</i>).	52
4.2	Bird's eye view of the waveguide chip. The lines hint at the buried waveguides converging in the centre. (a) DIC micrograph of the centre region. The light and dark lines show slopes on the surface. (b) Sketch of the same region with the waveguides clearly present. The red waveguides are the one that are connected to fibres. (c) A sketch of the even further enlarged trench region and inside a representation of the light mode field diameter diverging across the trench (for 780 nm).	53
4.3	Manufacturing process: steps (1) and (2).	56
4.4	Manufacturing process: steps (3), (4) and (5).	57
4.5	Manufacturing process: steps (6), (7) and (8).	58
4.6	Manufacturing process: steps (9), (10) and (11).	59
4.7	Manufacturing process: steps (12) and (13).	60
4.8	Manufacturing process: steps (14) and (15).	61
4.9	Manufacturing process: steps (16) and (17).	62
4.10	Temperature data with fitted curves to the short- (<i>blue</i>) and long-term (<i>green</i>) rise. The temperature was taken in the centre of the current-carrying structure without a waveguide chip attached.	64

5.1	Plot of the variance of the mean count rate of an intensity measurement $\text{Var}(n_{\text{recorded}})$ (blue) versus the measured count rate n_{recorded} itself. The lines are theory curves for shot noise (green) and dead-time-corrected shot noise (red) for a value of $\tau_{\text{dead}} = 32$ ns.	75
5.2	Plot of the laser power transmitted through a waveguide pair (#9) to the APD detector. The dip in power shows the atoms blocking about 2.6% of the probe light power.	79
5.3	Plot of the absorption spectrum of the atom-waveguide-detector-system. The data in form of relative absorption signal for different values of the laser frequency detuning from resonance is superimposed with a Lorentzian fit that yields a linewidth of 7.9 ± 0.4 MHz. Error bars give 1σ statistical uncertainties for Poissonian light with detector dead time.	81
5.4	Saturation Curve. For varying intensities the relative absorption is measured and plotted against the power incident on the detector. The solid line is a theoretical curve taking into account the spatial variation of light within the trench and optical pumping of the atoms. Error bars give 1σ statistical uncertainties for Poissonian light with detector dead time.	82
5.5	Zeeman-split absorption spectrum. When pure left-handed light (red) is coupled into the input fibre the Zeeman shift is almost $-\mu_B B/h$. The green line is a theory obtained by summing over all the Zeeman components of the spectrum. This fit tells us that $(85 \pm 3)\%$ of the power in the trench is still in the left-handed mode. For right-handed light (blue) the shifts are opposite. The fit to these points gives $(81 \pm 3)\%$ of power in the right handed mode. Error bars give 1σ statistical uncertainties for Poissonian light with detector dead time.	85
5.6	Fluorescence signal collected by the waveguide and detected on an APD. This plot shows the averaged data of 34 shots. The measurement of the peak signal agrees very well with the estimated fluorescence of 73 counts per millisecond, indicating that the fluorescence detection methods works well in our device.	88

6.1	Model of cloud and light geometry. The shape of the condensate (<i>blue</i>) is parametrised by r_0 and z_0 . Its density ρ is a parabolic profile. The light beam (<i>red</i>) along the x -axis has a Gaussian electric field profile with waist w_0	91
6.2	Overlap integrals plotted against relative condensate size in the case $w = w_0$. In this case the curves reach their maximally possible values.	95
6.3	Geometry for the MEEP simulation. (<i>a</i>) the system is effectively treated in 2d. The waveguide materials are specified by constant refractive indices. The condensate is built out of shells of inwardly increasing complex refractive index. (<i>b</i>) To let the electric field settle down and stabilize after the field source and after the interaction region, the waveguides extend about 1000 wavelengths in both directions from the trench. (<i>c</i>) to simplify the calculation, the local density of the cloud is flattened into a mean density. Subsequently, the thickness is flattened to an effective thickness.	99
6.4	MEEP simulation example plots of the intensity for different detunings Δ' . The grey scale corresponds to energy density, the coloured background to the value of the relative permittivity ϵ_r . The light comes from the left, enters the trench and interacts with the condensate. These simulations were performed for 3200 atoms, and a condensate of size $r_0 = 1 \mu\text{m}$, $z_0 = 40 \mu\text{m}$	100
6.5	Comparison between the simulation data and the semi-analytical model. The fit yields values for the overlap integral of $\mathcal{O} = 0.58 \pm 0.02$ and $N_{\text{atoms}} = 6500 \pm 510$	103
6.6	Principal geometry of measuring a cloud of atoms. (<i>a</i>) a direct absorption measurement. The light is sent through the cloud. (<i>b</i>) Mach-Zender interferometer measurement. The light is split, one part is sent through the cloud and one is used as a reference. After recombination, in phase and quadrature photon numbers are measured. The second beam splitter S_2 we assume splits 50:50.	105

6.7	Possible setup of a beat-note phase detection scheme. <i>(top)</i> Monochromatic light (\mathcal{E}_0) passes through a modulation stage suppressing the carrier and creating two sidebands (\mathcal{E}_1), which in turn interact with the atomic sample. After the sample the light (\mathcal{E}_2) is detected (\mathcal{D}) and transformed into a voltage. <i>(bottom)</i> The measured voltage ($V_{\mathcal{D}}$) is split into its ac and dc parts (V_{ac} and V_{dc}). The ac part is split and mixed with a frequency-doubled original modulation signal ($2V_m$) to retrieve in phase and quadrature signals (V_i and V_q) after low pass filtering multiples of the modulation frequency ω_m .	111
6.8	Normalised signals for 100 atoms in a BEC of dimensions $r_0 = 1 \mu\text{m}$ and $z_0 = 40 \mu\text{m}$, scanning the carrier frequency over resonance for three different modulation frequencies.	115
6.9	Beat note signals for carrier frequency on resonance, i.e. $\Delta' = 0$, varying the number of atoms for three different modulation frequencies.	116

List of Tables

2.1	Trap parameters for three important types of trap on the path to Bose-Einstein condensation. Symbols used: I_Z : Z-structure current, I_{bias} : bias wire current, y_{min} : vertical position of magnetic field minimum, B_{min} : value of magnetic field minimum, T_{td} : trap depth, B_z : shim field z-component, N_{atoms} : number of atoms, T : cloud temperature, $\omega_{r/a}$: radial/axial oscillation frequency, $\bar{\omega}$: geometric mean oscillation frequency, $s_{r/a}$: radial/axial cloud size, $\tan \vartheta$: relative cloud displacement, AR: aspect ratio, PSD: phase space density.	19
4.1	Results for the measured resistance R of the conducting sub-chip structures (uncertainties are due to the limit of the multimeter used: $\sigma_R = 1 \text{ m}\Omega$). The corresponding maximum power production P_{max} is calculated for each resistance and design current I_{design} in our design.	63
6.1	Effective hyperfine dipole matrix elements [79] from the $ 2, 2\rangle$ state into $ 3, m'_F\rangle$ excited states. For each case of incoming light polarisation (see text) the relative fraction of light in the $(\sigma^-, \pi, \sigma^+)$ basis is given, as well as the resulting effective transition strengths.	98

Bibliography

- [1] M. Anderson, J. Ensher, M. Matthews, C. Wieman, and E. Cornell. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, 269:198, 1995.
- [2] K. Davis, M.-O. Mewes, M. Andrews, N. van Druten, D. Durfee, D. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Physical Review Letters*, 75:3969, 1995.
- [3] W. Ketterle, D. Durfee, and D. Stamper-Kurn. Making, probing and understanding Bose-Einstein condensates. In Bose-Einstein condensation in atomic gases, Proceedings of the International School of Physics "Enrico Fermi", Course CXL, edited by M. Inguscio, S. Stringari and C. E. Wieman (IOS Press, Amsterdam, 1999) pp. 67-176.
- [4] D. Müller, D. Z. Anderson, R. J. Grow, P. D. D. Schwindt, and E. A. Cornell. Guiding neutral atoms around curves with lithographically patterned current-carrying wires. *Physical Review Letters*, 83(25):5194–5197, 1999.
- [5] J. Denschlag, D. Cassettari, A. Chenet, S. Schneider, and J. Schmiedmayer. A neutral atom and a wire: towards mesoscopic atom optics. *Applied Physics B: Lasers and Optics*, 69(4):291–301, 1999.
- [6] R. Folman, P. Krüger, D. Cassettari, B. Hessmo, T. Maier, and J. Schmiedmayer. Controlling cold atoms using nanofabricated surfaces: Atom chips. *Physical Review Letters*, 84(20):4749–4752, 2000.
- [7] W. Hänsel, P. Hommelhoff, T. W. Hänsch, and J. Reichel. Bose-einstein condensation on a microelectronic chip. *Nature*, 413(6855):498–501, 2001.

- [8] H. Ott, J. Fortagh, G. Schlotterbeck, A. Grossmann, and C. Zimmermann. Bose-einstein condensation in a surface microtrap. *Phys. Rev. Lett.*, 87(23):230401, 2001.
- [9] R. Long, T. Steinmetz, P. Hommelhoff, W. Hänsel, T. Hänsch, and J. Reichel. Magnetic microchip traps and single-atom detection. *Royal Society of London Philosophical Transactions Series A*, 361:1375–1389, 2003.
- [10] J. Reichel. Microchip traps and bose-einstein condensation. *Applied Physics B: Lasers and Optics*, 74(6):469–487, 2002.
- [11] B. V. Hall, S. Whitlock, F. Scharnberg, P. Hannaford, and A. Sidorov. A permanent magnetic film atom chip for bose einstein condensation. *Journal of Physics B Atomic Molecular Physics*, 39:27–36, January 2006.
- [12] C. D. J. Sinclair, J. A. Retter, E. A. Curtis, B. V. Hall, I. Llorente Garcia, S. Eriksson, B. E. Sauer, and E. A. Hinds. Cold atoms in videotape micro-traps. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 35(1):105–110, 2005.
- [13] C. D. J. Sinclair, E. A. Curtis, I. Llorente Garcia, J. A. Retter, B. V. Hall, S. Eriksson, B. E. Sauer, and E. A. Hinds. Bose-einstein condensation on a permanent-magnet atom chip. *Physical Review A*, 72:031603, 2005.
- [14] S. Whitlock, B. V. Hall, T. Roach, R. Anderson, M. Volk, P. Hannaford, and A. I. Sidorov. Effect of magnetization inhomogeneity on magnetic microtraps for atoms. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 75(4):043602, 2007.
- [15] P. Quinto-Su, M. Tschernack, M. Holmes, and N. Bigelow. On-chip optical detection of laser cooled atoms. *Opt. Express*, 12(21):5098–5103, 2004.
- [16] S. Eriksson, M. Trupke, H. F. Powell, D. Sahagun, C. D. J. Sinclair, E. A. Curtis, B. E. Sauer, E. A. Hinds, Z. Moktadir, C. O. Gollasch, and M. Kraft. Integrated optical components on atom chips. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 35(1):135–139, 2005.
- [17] A. Takamizawa, T. Steinmetz, R. Delhuille, T. W. Hänsch, and J. Reichel. Miniature fluorescence detector for single atom observation on a microchip. *Optics Express*, 14(23):10976–10983, 2006.

- [18] P. Horak, B. G. Klappauf, A. Haase, R. Folman, J. Schmiedmayer, P. Domokos, and E. A. Hinds. Possibility of single-atom detection on a chip. *Physical Review A (Atomic, Molecular, and Optical Physics)*. *Rev. A*, 67(4):043806, 2003.
- [19] Marco Wilzbach, Dennis Heine, Sönke Groth, Xiyuan Liu, Thomas Raub, Björn Hessmo, and Jörg Schmiedmayer. Simple integrated single-atom detector. *Opt. Lett.*, 34(3):259–261, 2009.
- [20] Yves Colombe, Tilo Steinmetz, Guilhem Dubois, Felix Linke, David Hunger, and Jakob Reichel. Strong atom-field coupling for bose-einstein condensates in an optical cavity on a chip. *Nature*, 450:272–276, 2007.
- [21] S. Gleyzes, A. El Amili, R. A. Cornelussen, P. Lalanne, C. I. Westbrook, A. Aspect, J. Estève, G. Moreau, A. Martinez, X. Lafosse, L. Ferlazzo, J. C. Harmand, D. Maily, and A. Ramdane. Towards a monolithic optical cavity for atom detection and manipulation. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 53:107–111, 2009.
- [22] M. Trupke, J. Goldwin, B. Darquié, G. Dutier, S. Eriksson, J. Ashmore, and E. A. Hinds. Atom detection and photon production in a scalable, open, optical microcavity. *Physical Review Letters*, 99(6):063601, 2007.
- [23] Igor Teper, Yu-Ju Lin, and Vladan Vuletić. Resonator-aided single-atom detection on a microfabricated chip. *Physical Review Letters*, 97(2):023002, 2006.
- [24] Ying-Ju Wang, Dana Z. Anderson, Victor M. Bright, Eric A. Cornell, Quentin Diot, Tetsuo Kishimoto, Mara Prentiss, R. A. Saravanan, Stephen R. Segal, and Saijun Wu. Atom michelson interferometer on a chip using a bose-einstein condensate. *Phys. Rev. Lett.*, 94(9):090405, Mar 2005.
- [25] Dennis Heine, Marco Wilzbach, Thomas Raub, Björn Hessmo, and Jörg Schmiedmayer. Integrated atom detector: Single atoms and photon statistics. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 79(2):021804, 2009.
- [26] R. Olshansky. Propagation in glass optical waveguides. *Rev. Mod. Phys.*, 51(2):341–367, Apr 1979.

- [27] P. K. Tien. Integrated optics and new wave phenomena in optical waveguides. *Rev. Mod. Phys.*, 49(2):361–420, Apr 1977.
- [28] Graeme D. Maxwell. Photosensitivity & rare-earth doping in flame hydrolysis deposited planar silica waveguides. Invited paper, SPIE Functional Photonic & Fiber Devices, 30th Jan 1996, San Jose, California.
- [29] M. Bajcsy, S. Hofferberth, V. Balic, T. Peyronel, M. Hafezi, A. S. Zibrov, V. Vuletic, and M. D. Lukin. Efficient all-optical switching using slow light within a hollow fiber. *Phys. Rev. Lett.*, 102(20):203902, May 2009.
- [30] Meng Khoon Tey, Zilong Chen, Syed Abdullah Aljunid, Brenda Chng, Florian Huber, Gleb Maslennikov, and Christian Kurtsiefer. Strong interaction between light and a single trapped atom without the need for a cavity. *Nature Physics*, 4:924 – 927, 2008.
- [31] Ph. Laurent, M. Abgrall, Ch. Jentsch, P. Lemonde, G. Santarelli, A. Clairon, I. Maksimovic, S. Bize, Ch. Salomon, D. Blonde, J.F. Vega, O. Grosjean, F. Picard, M. Saccoccio, M. Chaubet, N. Ladiette, L. Guillet, I. Zenone, Ch. Delaroche, and Ch. Sirmain. Design of the cold atom pharao space clock and initial test results. *Applied Physics B: Lasers and Optics*, 84:683, 2006.
- [32] Keng-Yeow Chung, Sheng wey Chiow, Sven Herrmann, Steven Chu, and Holger Müller. Atom interferometry tests of local lorentz invariance in gravity and electrodynamics. *Physical Review D (Particles and Fields)*, 80(1):016002, 2009.
- [33] S. Aigner, L. Della 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*, 319:1226, 2008.
- [34] Haisheng Rong, Richard Jones, Ansheng Liu, Oded Cohen, Dani Hak, Alexander Fang, and Mario Paniccia. A continuous-wave raman silicon laser. *Nature*, 433:725–728, 2005.
- [35] A. Stibor, S. Kraft, T. Campey, D. Komma, A. Günther, J. Fortágh, C. J. Vale, H. Rubinsztein-Dunlop, and C. Zimmermann. Calibration of a single-atom detector for atomic microchips. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 76(3):033614, 2007.

- [36] H. Park, A. W. Fang, R. Jones, O. Cohen, O. Raday, M. N. Sysak, M. J. Paniccia, and J. E. Bowers. A hybrid algal-silicon evanescent waveguide photodetector. *Optics Express*, 15(10):6044–6052, 2007.
- [37] C. Hagleitner, A. Hierlemann, D. Lange, A. Kummer, N. Kerness, O. Brand, and H. Baltes. Smart single-chip gas sensor microsystem. *Nature*, 414:293–296, 2001.
- [38] G. Stern, B. Battelier, R. Geiger, G. Varoquaux, A. Villing, F. Moron, O. Carraz, N. Zahzam, Y. Bidet, W. Chaibi, F. Pereira Dos Santos, A. Bresson, A. Landragin, and P. Bouyer. Light-pulse atom interferometry in microgravity. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 53:353, 2009.
- [39] T. Könemann, W. Brinkmann, E. Göklü, C. Lämmerzahl, H. Dittus, T. van Zoest, E.M. Rasel, W. Ertmer, W. Lewoczko-Adamczyk, M. Schiemangk, A. Peters, A. Vogel, G. Johannsen, S. Wildfang, K. Bongs, K. Sengstock, E. Kajari, G. Nandi, R. Walser, and W.P. Schleich. A freely falling magneto-optical trap drop tower experiment. *Applied Physics B: Lasers and Optics*, 89:431, 2007.
- [40] F. Sorrentino, K. Bongs, P. Bouyer, L. Cacciapuoti, M. de Angelis, H. Dittus, W. Ertmer, A. Giorgini, J. Hartwig, M. Hauth, S. Herrmann, M. Inguscio, E. Kajari, T. Kænemann, C. Lämmerzahl, A. Landragin, G. Modugno, F. Pereira dos Santos, A. Peters, M. Prevedelli, E. M. Rasel, W. P. Schleich, M. Schmidt, A. Senger, K. Sengstok, G. Stern, G. M. Tino, and R. Walser. SAI: a compact atom interferometer for future space missions. *ArXiv e-prints*, March 2010.
- [41] József Fortágh and Claus Zimmermann. Magnetic microtraps for ultracold atoms. *Reviews of Modern Physics*, 79(1):235, 2007.
- [42] J. Reichel and V. Vuletic. *Atom Chips*. Wiley-VCH, first edition, 2011.
- [43] S. Seidelin, J. Chiaverini, R. Reichle, J. J. Bollinger, D. Leibfried, J. Britton, J. H. Wesenberg, R. B. Blakestad, R. J. Epstein, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, N. Shiga, and D. J. Wineland. Microfabricated surface-electrode ion trap for scalable quantum information processing. *Physical Review Letters*, 96(25):253003, 2006.

- [44] A. André, D. DeMille, J. M. Doyle, M. D. Lukin, S. E. Maxwell, P. Rabl, R. J. Schoelkopf, and P. Zoller. A coherent all-electrical interface between polar molecules and mesoscopic superconducting resonators. *Nature Physics*, 2:636 – 642, 2006.
- [45] Wenge Yang, Donald B. Conkey, Bin Wu, Dongliang Yin, Aaron R. Hawkins, and Holger Schmid. Atomic spectroscopy on a chip. *Nature Photonics*, 1:331 – 335, 2007.
- [46] Svenja Knappe, Vishal Shah, Peter D. D. Schwindt, Leo Hollberg, John Kitching, Li-Anne Liew, and John Moreland. A microfabricated atomic clock. *Applied Physics Letters*, 85(9):1460–1462, 2004.
- [47] S. Pollock, J. P. Cotter, A. Laliotis, and E. A. Hinds. Integrated magneto-optical traps on a chip using silicon pyramid structures. *Opt. Express*, 17(16):14109–14114, 2009.
- [48] O. Alloschery, R. Mathevet, J. Weiner, and H. J. Lezec. All-optical atom surface traps implemented with one-dimensional planar diffractive microstructures. *Opt. Express*, 14(26):12568–12575, 2006.
- [49] Barak Dayan, A. S. Parkins, Takao Aoki, E. P. Ostby, K. J. Vahala, and H. J. Kimble. A photon turnstile dynamically regulated by one atom. *Science*, 319:1062, 2008.
- [50] Takao Aoki, A. S. Parkins, D. J. Alton, C. A. Regal, Barak Dayan, E. Ostby, K. J. Vahala, and H. J. Kimble. Efficient routing of single photons by one atom and a microtoroidal cavity. *Physical Review Letters*, 102(8):083601, 2009.
- [51] Paul E. Barclay, Kartik Srinivasan, Oskar Painter, Benjamin Lev, and Hideo Mabuchi. Integration of fiber-coupled high-q \sin_x microdisks with atom chips. *Applied Physics Letters*, 89(13):131108, 2006.
- [52] D. K. Armani, T. J. Kippenberg, S. M. Spillane, and K. J. Vahala. Ultra-high-q toroid microcavity on a chip. *Nature*, 421:925–928, 2003.
- [53] D. Jaksch, H.-J. Briegel, J. I. Cirac, C. W. Gardiner, and P. Zoller. Entanglement of atoms via cold controlled collisions. *Phys. Rev. Lett.*, 82(9):1975–1978, Mar 1999.

- [54] Shuai Chen, Yu-Ao Chen, Thorsten Strassel, Zhen-Sheng Yuan, Bo Zhao, Jörg Schmiedmayer, and Jian-Wei Pan. Deterministic and storable single-photon source based on a quantum memory. *Physical Review Letters*, 97(17):173004, 2006.
- [55] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Phys. Rev. Lett.*, 74(20):4091–4094, May 1995.
- [56] D. DeMille. Quantum computation with trapped polar molecules. *Phys. Rev. Lett.*, 88(6):067901, Jan 2002.
- [57] Alberto Politi, Martin J. Cryan, John G. Rarity, Siyuan Yu, and Jeremy L. O’Brien. Silica-on-Silicon Waveguide Quantum Circuits. *Science*, 320(5876):646–649, 2008.
- [58] Jonathan C. F. Matthews, Alberto Politi, Andre Stefanov, and Jeremy L. O’Brien. Manipulation of multiphoton entanglement in waveguide quantum circuits. *Nature Photon*, 3:346–350, 2009.
- [59] Alberto Politi, Jonathan C. F. Matthews, and Jeremy L. O’Brien. Shor’s Quantum Factoring Algorithm on a Photonic Chip. *Science*, 325(5945):1221–, 2009.
- [60] M. Kohnen, M. Succo, P. G. Petrov, R. A. Nyman, M. Trupke, and E. A. Hinds. An array of integrated atom-photon junctions. *Nature Photonics*, 5:35–38, 2011.
- [61] H. J. Metcalf and P van der Straten. *Laser Cooling and Trapping*. Springer, 1999.
- [62] J. D. Jackson. *Classical Electrodynamics*. John Wiley & Sons, third edition, 1998.
- [63] O. Chubar, P. Elleaume, and J. Chavanne. Radia. [www.esrf.eu / Accelerators / Groups / InsertionDevices / Software / Radia](http://www.esrf.eu/Accelerators/Groups/InsertionDevices/Software/Radia) (Version 4.29, 7 October 2009).
- [64] O. Chubar, P. Elleaume, and J. Chavanne. Computing 3D Magnetic Field from Insertion Devices. Proceedings of the PAC97 Conference May 1997, p.3509-3511.

- [65] O. Chubar, P. Elleaume, and J. Chavanne. A 3D Magnetostatics Computer Code for Insertion devices. SRI97 Conference August 1997, J. Synchrotron Rad. (1998). 5, 481-484.
- [66] C J Vale, B Upcroft, M J Davis, N R Heckenberg, and H Rubinsztein-Dunlop. Foil-based atom chip for bose-einstein condensates. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 37(14):2959–2967, 2004.
- [67] M. Volk, S. Whitlock, B. V. Hall, and A. I. Sidorov. Scanning magnetoresistance microscopy of atomchips. *ArXiv Physics e-prints*, 2007.
- [68] T. Schumm, J. Estève, C. Aussibal, C. Figl, J.-B. Trebbia, H. Nguyen, D. Mailly, C. Westbrook, and A. Aspect. Atom chips in the real world: the effects of wire corrugation. *European Physical Journal D*, 32:171, 2005.
- [69] A. E. Leanhardt, Y. Shin, A. P. Chikkatur, D. Kielpinski, W. Ketterle, and D. E. Pritchard. Bose-einstein condensates near a microfabricated surface. *Physical Review Letters*, 90(10):100404, 2003.
- [70] J. Reichel, W. Hänsel, and T. W. Hänsch. Atomic micromanipulation with magnetic surface traps. *Phys. Rev. Lett.*, 83(17):3398–3401, Oct 1999.
- [71] W. D. Phillips and H. J. Metcalf. Cooling and trapping atoms. *Scientific American*, 256:50–56, March 1987.
- [72] S. Wildermuth, P. Krüger, C. Becker, M. Brajdic, S. Haupt, A. Kasper, R. Folman, and J. Schmiedmayer. Optimized magneto-optical trap for experiments with ultracold atoms near surfaces. *Physical Review A (Atomic, Molecular, and Optical Physics)*, 69(3):030901, 2004.
- [73] L. Pitaevskii and S. Stringari. *Bose-Einstein Condensation*. Oxford University Press, 2003.
- [74] E. G. M. van Kempen, S. J. J. M. F. Kokkelmans, D. J. Heinzen, and B. J. Verhaar. Interisotope determination of ultracold rubidium interactions from three high-precision experiments. *Phys. Rev. Lett.*, 88(9):093201, Feb 2002.
- [75] C. D. J. Sinclair. *Bose-Einstein Condensation in Microtraps on Videotape*. PhD thesis, Imperial College London, 2005.

- [76] I. Llorente-Garcia. *Advances in the design and operation of atom chips*. PhD thesis, Imperial College London, 2008.
- [77] C. Wieman and T. W. Hänsch. Doppler-free laser polarization spectroscopy. *Phys. Rev. Lett.*, 36(20):1170–1173, May 1976.
- [78] K. B. MacAdam, A. Steinbach, and C. Wieman. A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for cs and rb. *American Journal of Physics*, 60(12):1098–1111, 1992.
- [79] Daniel A. Steck. Rubidium 87 D Line Data. Available online at <http://steck.us/alkalidata> (revision 2.1.2, 12 August 2009).
- [80] 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):3331–3334, Oct 1996.
- [81] G. Ritt, G. Cennini, C. Geckeler, and M. Weitz. Laser frequency offset locking using a side of filter technique. *Applied Physics B: Lasers and Optics*, 79(3):363–365, 2004.
- [82] Allied Vision Technologies (AVT) Pike F-145B, Technical Data Sheet. [www.alliedvisiontec.com / fileadmin / content / PDF / Products / Data_sheet / Cameras / Pike / Pike_DataSheet_F145BC_fiber_V3.0.0_en.pdf](http://www.alliedvisiontec.com/fileadmin/content/PDF/Products/Data_sheet/Cameras/Pike/Pike_DataSheet_F145BC_fiber_V3.0.0_en.pdf).
- [83] Allied Vision Technologies (AVT) Marlin F-033B, Technical Data Sheet. [www.alliedvisiontec.com / fileadmin / content / PDF / Products / Data_sheet / Cameras / Marlin / Marlin_DataSheet_F033BC_V3.0.0_en.pdf](http://www.alliedvisiontec.com/fileadmin/content/PDF/Products/Data_sheet/Cameras/Marlin/Marlin_DataSheet_F033BC_V3.0.0_en.pdf).
- [84] Perkin Elmer SPCM-AQRH-14 – Si APD Single Photon Counting Module with less than 100 dark counts/second, Technical Data Sheet. [optoelectronics.perkinelmer.com / content / Datasheets / DTS_SPCMAQRH.pdf](http://optoelectronics.perkinelmer.com/content/Datasheets/DTS_SPCMAQRH.pdf).
- [85] J. J. Hudson, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds. Measurement of the electron electric dipole moment using ybf molecules. *Physical Review Letters*, 89(2):023003, Jun 2002.
- [86] A Laboratory Control System for Cold Atom Experiments. [iqoqi006.uibk.ac.at / users / c704250 / Control / index.html](http://iqoqi006.uibk.ac.at/users/c704250/Control/index.html).

- [87] Peter E. Gaskell, Jeremy J. Thorn, Sequoia Alba, and Daniel A. Steck. An open-source, extensible system for laboratory timing and control. *Review of Scientific Instruments*, 80(11):115103, 2009.
- [88] G. Varoquaux. Agile computer control of a complex experiment. *Computing in Science and Engineering*, 10:55–59, 2008.
- [89] R. McDougall. Mask Set A1241. QA Review, CIP Technologies, 2006.
- [90] Epotek®H77 Technical Data Sheet. [www.epotek.com / sscdocs / datasheets / H77.PDF](http://www.epotek.com/sscdocs/datasheets/H77.PDF).
- [91] Epotek®353ND Technical Data Sheet. [www.epotek.com / sscdocs / datasheets / 353ND.PDF](http://www.epotek.com/sscdocs/datasheets/353ND.PDF).
- [92] E. R. I. Abraham and E. A. Cornell. Teflon feedthrough for coupling optical fibers into ultrahigh vacuum systems. *Applied Optics*, 37(10):1762–1763, 1998.
- [93] D. L. Miller and N. T. Moshegov. All-metal ultrahigh vacuum optical fiber feedthrough. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 19(1):386–387, 2001.
- [94] D. F. Yu and J. A. Fessler. Mean and variance of coincidence counting with deadtime. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 488(1-2):362 – 374, 2002.
- [95] H.-A. Bachor and T. C. Ralph. *A Guide to Experiments in Quantum Optics*. John Wiley & Sons, second edition, 2004.
- [96] M. Born and E. Wolf. *Principles of Optics*. Pergamon Press, third edition, 1964.
- [97] J. E. Lye, J. J. Hope, and J. D. Close. Nondestructive dynamic detectors for bose-einstein condensates. *Phys. Rev. A*, 67(4):043609, Apr 2003.
- [98] J. E. Lye, J. J. Hope, and J. D. Close. Rapid real-time detection of cold atoms with minimal destruction. *Phys. Rev. A*, 69(2):023601, Feb 2004.
- [99] Steven Kasapi, Seema Lathi, and Yoshihisa Yamamoto. Sub-shot-noise fm noise spectroscopy of trapped rubidium atoms. *J. Opt. Soc. Am. B*, 15(10):2626–2630, 1998.

- [100] M. J. Snadden, R. B. M. Clarke, and E. Riis. Fm spectroscopy in fluorescence in laser-cooled rubidium. *Optics Communications*, 152(4-6):283 – 288, 1998.
- [101] V. Savalli, G. Zs. K. Horvath, P. D. Featonby, L. Cagnet, N. Westbrook, C. I. Westbrook, and A. Aspect. Optical detection of cold atoms without spontaneous emission. *Opt. Lett.*, 24(22):1552–1554, 1999.
- [102] H. Mabuchi, J. Ye, and H.J. Kimble. Full observation of single-atom dynamics in cavity qed. *Applied Physics B: Lasers and Optics*, 68:1095–1108, 1999. 10.1007/s003400050751.
- [103] Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, ninth dover printing, tenth gpo printing edition, 1964.
- [104] Rodney Loudon. *The Quantum Theory of Light*. Oxford University Press, second edition, 1983.
- [105] A. R. Edmonds. *Angular Momentum in Quantum Mechanics*. Princeton University Press, Princeton, NJ, 1957.
- [106] A. Farjadpour, D. Roundy, A. Rodriguez, M. Ibanescu, P. Bermel, J. D. Joannopoulos, S. G. Johnson, and G. Burr. Improving accuracy by subpixel smoothing in FDTD. *Optics Letters*, 31(20):2972–2974, 2006.
- [107] M. Saffman, D. Oblak, J. Appel, and E. S. Polzik. Spin squeezing of atomic ensembles by multicolor quantum nondemolition measurements. *Phys. Rev. A*, 79(2):023831, Feb 2009.
- [108] G. C. Bjorklund, M. D. Levenson, W. Lenth, and C. Ortiz. Frequency Modulation (FM) Spectroscopy. *Applied Physics B: Lasers and Optics*, 32:145–152, 1983.