


Magnetometry and microscopy of cold atom clouds

Ottó Elíasson
PhD Thesis
November 2019

Department of Physics and Astronomy
Aarhus University
Denmark



Magnetometry and microscopy of cold atom clouds

A Dissertation
Presented to the Faculty of Science and Technology
of Århus University
in Partial Fulfilment of the Requirements
for the PhD Degree

by
Ottó Elíasson
Sunday 3rd November, 2019

Cover design by Sophie Christine Haack Andersen

*Þolinmæði þrautir vinnur allar.**

**Patience conquers all challenges.*

Abstract

This thesis covers experimental work conducted within two distinct fields of cold-atom physics. The first part describes an experiment capable of spatially-selective dispersive measurements of a cold-atomic system in part or in whole, by virtue of a digital micromirror device (DMD). Atom clouds are trapped in optical tweezers made and controlled by an acousto-optical deflector. Two types of high-precision magnetometers are realised. One is vectorial and reaches a single-shot precision of $\delta\mathbf{B} = (100, 200) \mu\text{G}$ for the field components parallel and transversal to the probe light, exceeding a previous realisation by two orders of magnitude. The other is a scalar magnetometer based on Larmor precession and reaches $\delta B = 30 \mu\text{G}$, putting it on equal footing with other state-of-the-art cold-atom magnetometers.

The second part is devoted to the characterisation and the first results of a new quantum gas microscope experiment. It features a 0.69 NA microscope objective for high-resolution fluorescence imaging of individual atoms trapped in deep optical lattices. The imaging system is shown to perform near the diffraction limit. By means of DMD-generated off-resonant tight optical tweezers, projected through the high-resolution optics, we can load only a few planes of a co-propagating 1D lattice. In a single realisation of the experiment we acquire multiple fluorescence images, where the objective is translated between images, bringing different planes of the optical lattices in focus. In this way we can tomographically reconstruct the atom distribution in 3D.

Resumé

Denne afhandling omhandler eksperimentelt arbejde lavet i to forskellige felter inden for fysik med kolde atomer. Den første del beskriver et eksperiment der kan lave en rumligselektive dispersiv måling af dele eller hele systemer af kolde atomer ved brug af en digital lysprojektor (DMD). Atomskyer er fanget i optiske pincetter, som er kontrolleret af en akustisk-optisk deflektor (AOD). To typer af højpræcisions magnetometre er realiseret. Den ene er vektorbaseret og opnår en præcision på $\delta\mathbf{B} = (100, 200) \mu\text{G}$ af feltets henholdsvis parallelle og transversale komponenter i forhold til sonde lyset udbredelsesretning. Dette er to størrelsesordner bedre end tidligere realiseringer. Det andet er et skalar-magnetometer baseret på Larmor-præcession og opnår $\delta B = 30 \mu\text{G}$, hvilket er tilsvarende med andre avancerede magnetometre der bruger kolde atomer.

Anden del er dedikeret til karakteriseringen og de første resultater af et nyt kvantegas mikroskop eksperiment. Dette har et mikroskop objektiv med en NA på 0,69 til at tage fluorescens billeder med højopløsning af individuelle atomer fanget i optiske gitre. Det bliver vist at afbildningssystemet præstere målinger tæt på diffraktionsgrænsen. Ved brug af DMD genererede små optiske pincetter lavet af laserlyser med en frekvens væk fra atomets resonans der er projekteret igennem højpræcisionsoptik, kan vi fylde få plan af et 1D gitter der propagerer parallelt med de optiske pincetter. I en enkelt realisering af eksperimentet tager vi adskillige billeder, hvor objektivitet er flyttet i mellem billederne, således at forskellige plan af det optiske gitter er i fokus. På den måde kan vi tomografisk rekonstruere atomfordelingen i 3D.

Samantekt

Þessi ritgerð gerir grein fyrir tilraunum innan tveggja sviða eðlisfræði kaldra atómskýja. Fyrri hlutinn lýsir tilraunum sem gera kleifar mælingar á köldum atómskýjum að hluta eða í heild. Til þessa er beitt tvísturhluta víxlverkunar atóms og ljóss, og er mæliljósíð mótað með stafrænum ljósvarpa (e. *digital micromirror device*). Atómskýin fanga smáar ljóstangir sem stýrt er með mótara hljóð-ljóshrifa (e. *acousto-optical deflector*). Í þessu kerfi voru útbúnar tvær gerðir nákvæmra segulsviðsmæla. Annar mælir vigursvið og nær nákvæmni $\delta B = (100, 200) \mu G$ í einni mælilotu, á þáttum segulsviðsins samsíða og hornrétt á stefnu mæliljóssins. Þar með bætir mælirinn nákvæmni þessarar tækni um tvö stærðarþrep. Hinn mælirinn byggir á Larmor pólveltu og mælir því aðeins stig sviðsins, en hann er í eðli sínu hittinn. Sá nær $\delta B = 30 \mu G$ nákvæmni, til jafns við það sem náðst hefur með öðrum segulmælum í köldum atómskýjum.

Seinni hluti ritgerðarinnar fjallar um uppbyggingu og fyrstu niðurstöður úr nýrri smásjártilraun fyrir skammtagös. Í smásjánni er hlutgler með ljósopstölu 0,69, nýtt til myndunar í hárri upplausn á flúrljósi frá einstökum atómum, sem fönguð eru í djúpu mætti þrívíðra ljósgrinda. Gæði flúrmyndanna eru því sem næst eingöngu takmörkuð af öldubeygju ljóssins. Örsmáum ljóstöngum með ljóstíðni fjarri hermu atómsins, er varpað gegnum hlutglerið með hjálp stafræns ljósvarpa. Í þeim sitja lítil atómský sem færa má yfir örfáar sléttur einvíðrar ljósgrindar. Í einni mælilotu eru teknar nokkrar flúrmyndir þar sem hlutglerið er flutt milli mynda, sem færir mismunandi sléttur ljósgrindarinnar í brennipunkt myndkerfisins. Þannig má endurskapa rúmlæga dreifingu atómanna í þrívídd.

Preface

This thesis accounts for some of the work that has been carried out during the past five years in the high-resolution laboratory (Hires lab), buried deep underneath the parking lot of the Institute for Physics and Astronomy at Århus University. The Hires lab has undergone many changes since I first got acquainted with the group around the summer of 2014. At the time, there was no science chamber present and the first Bose-Einstein condensate (BEC) had just been created in the cube chamber earlier that year. The group was working on two fronts: building up and planning for the first experiments of using the Faraday interaction for probing the BEC phase transition, and for the eventual high-resolution experiment. The main results of these first experiments were that Faraday measurements could be used to reduce experimental fluctuations, and narrow down the transition point from a thermal cloud to a BEC. For details on this work I refer to ch. 5 and 6 of ref. [131] and the journal article [24].

The science chamber was added to the experiment early in the year of 2015. Experimentally, the focus was still on Faraday imaging and trapping of atoms in microscopic potentials, now in a dual microscope configuration in the new chamber that allowed for simultaneous trapping and probing of small atom clouds. That work culminated with the realisation of two kinds of magnetometers, one vectorial and the other based on Larmor precession. The bulk of that data were acquired during the summer of 2016, and later that made it to a published article [84]. The first part of this thesis describes this setup and the experiments in detail.

Alongside those measurements we kept on using the cube chamber for different experiments. Inspired by more general work on optimisation of complex problems ongoing in the theoretical part of the group, we dove into the control landscape of BEC creation. There are standard ways of making a BEC, but there are many knobs available in the experiment that can be turned to affect its size. We wanted to find out if there were any non-trivial settings of the apparatus that could provide us with larger BECs. To start with we inspected that landscape by a multitude of 1D, 2D and 3D scans, as enabled by our new (home-grown) experimental control system. During this time we collaborated with theoretical physicists from the University of Ulm in Germany, that applied state-of-the-art algorithms to the challenge of optimising the size of the BEC. They were allowed to control the intensity of two laser beams of a crossed dipole trap, and the current in a magnetic quadrupole trap during the last cooling stage at the end of the experimental sequence. The cherry on the top (d. *rosinen i pølseenden*) was the gamification of the challenge of making the biggest BEC: *the Alice Challenge*. For a couple of weeks time in the fall of 2016, about 600 participants from all over the world were allowed to take control of the experiment remotely. This was achieved with an openly available game interface, constructed by the developer team of ScienceAtHome. The participants could shape the intensity of two crossed laser beams and the strength of a magnetic trap during the last cooling stage, by dragging around curves on the screen. Afterwards they would submit their solution, which was run immediately (depending on the queue) in the experiment and the score (the BEC atom number) would be sent back to the player. About 7500 such solutions were realised in the lab during this time. The biggest BEC we ever saw in the experiment was created during these days by someone from the general public. This study was covered in ch. 4 of ref. [131] and has also been presented in a scientific publication [132].

In October 2016 we disassembled all the existing experimental machinery, as the high-resolution viewport (the window through which we look at the atoms) had to be replaced before any experiments with the high-resolution microscope could take place. This required another bakeout of the system. In the seasons to come we would build up the experiment again, learn how to use the flexible

trapping potentials projected through the objective, construct optical lattices and use the high-resolution imaging system. We got first light through the microscope in May 2017 but it was a full year later, during the last days of May 2018 when we unambiguously succeeded to collect clear fluorescent signals from individual atoms trapped in deep optical lattice potentials, and thereby realising the 11th quantum gas microscope experiment in the world.

During the time when the microscope was finally operational our atom clouds were plagued by heating from one of the lattice axes and after about ten weeks of run-time a broken optical fibre halted the operation of the experiment for a while. Soon after we were running again our main lattice laser broke down, leaving us with almost four months of down-time. The microscope had only been working properly for a weeks time when an interlock for the high-current coils on the cube chamber failed with catastrophic consequences.

One Monday in early February 2019, I arrived in the morning to warm up the BEC machine. And boy did I warm it up. I failed to turn on the cooling water, which would usually not have caused any harm but this time the interlock system did not work as intended (as is accounted for in the outlook of the thesis). The heat stress and the fire that must have been present, broke the vacuum of the chambers and they were filled with soot and dirt. Within the months to come the experiment would rise from the ashes like the Phoenix. To our surprise most elements could be cleaned and re-used, and luckily none of the critical high-resolution components were damaged. Late summer we got the MOT properly working again and in the fall all chambers were in place and under vacuum. At the time of writing the experiment is still recovering from that incident, with a few months of work still ahead to reach where we were before the fire.

* * *

There is more to the research group than just the experiment. It is a part of bigger ecosystem, *the ScienceAtHome team*, which is an international group of people with backgrounds in physics, social science, psychology, game development, graphical design and more. The aim is to create video games and interfaces

to involve the general public in scientific research and to develop new ways of teaching science through such media.

The project came alive with the game *Quantum Moves*, where players solve the quantum optimal control problem of picking up an atom with an optical tweezer and moving it to a designated target area. The game was a success both scientifically [258] and publicly [262]. The players sought out different parts of the solution landscape than the algorithms that had been previously put to the task of numerically optimising the transport. In this way the players inspired new ways of searching. This journey opened the eyes of the team to how computer games can be used to study how humans solve complex problems. Highlights since *Quantum Moves* are the Alice Challenge mentioned above, *Skill Lab* which is a suite of mini games that explores the cognitive abilities of its players, and *Quantum Moves 2* where the players are given access to an optimiser during the gameplay. The new version also includes optimal control problems of BECs, that obey a non-linear Schrödinger equation.

All these games engage players in citizen science and in some ways strive to understand how people play and solve the challenges they offer. In the Alice Challenge the behaviour of people at the top of the leader board (those that had created the largest BECs) was notably different from that of those with lower scores. They focussed rather on *exploiting* the qualities of their obtained solutions by making small amendments to it, whereas those that had not reached high score were unafraid to *explore* and take larger jumps. The players' task of creating a BEC, is by no means something we encounter in the everyday. The method of creating an intuitive/gamified user interface and its success, emphasises the importance of how problems are represented [7]. As of June 2019 the games developed by ScienceAtHome had been played more than 8 million times by 300.000 players from all over the world. This concept will undoubtedly be a part of the experiment and shape its future.

Apart from my participation in the Alice Challenge my role within the team has been more related to didactic and outreach activities. *Quantum Composer* is another fruit of the ScienceAtHome tree. Its a flow based computational environment for quantum mechanics and is an ideal supplement to quantum me-

chanics teaching at the University level. I have taken part in creating exercises and other teaching material using Quantum Composer and carried out those exercises with students on a graduate level at AU. Another project focuses on building a frame around groups of high-school students from local schools that come occasionally to visit AU and our team to learn of our hopes and dreams. This is done with the aim of igniting the flame of curiosity in their hearts. In addition to my standard teaching duties which are a mandatory part of my studies, I have instructed many Bachelor and Masters students to a varying degree, who have chosen to make projects in our lab.

List of publications

During my time in the group I have taken part in the following three scientific publications from the Hires lab, with the fourth now available as a preprint. The topics of the first two articles has been accounted for in the theses of Romain [199] and Robert [131], and the topics of the latter two (in all their gory details) form the backbone of this thesis.

Bason, M., R. Heck, M. Napolitano, O. Eliasson, R. Müller, A. Thorsen, W. Zhang, J. Arlt, J. Sherson (2018). Measurement-enhanced determination of BEC phase diagrams. *Journal of Physics B*, 51(17), 175301. doi: 10.1088/1361-6455/aad447.

Heck, R., O. Vuculescu, J. J. Sørensen, J. Zoller, M. G. Andreassen, M. G. Bason, P. Ejlertsen, O. Eliasson, P. Haikka, J. S. Laustsen, L. L. Nielsen, A. Mao, R. Müller, M. Napolitano, M. K. Pedersen, A. R. Thorsen, C. Bergenholtz, T. Calarco, S. Montangero, and J. Sherson (2018). Remote optimization of an ultra-cold atoms experiment by experts and citizen scientists. *PNAS*, 115(48) E11231-E11237. doi: 10.1073/pnas.1716869115.

Eliasson, O., R. Heck, J. S. Laustsen, M. Napolitano, R. Müller, M. G. Bason, J. Arlt and J. Sherson (2019). Spatially-selective *in situ* magnetom-

etry of ultracold atomic clouds. *Journal of Physics B*, 52 (7), 075003. doi: 10.1088/1361-6455/ab0bd6.

Elíasson, O., R. Heck, J. S. Laustsen, R. Müller, C. Weidner, J. Arlt and J. Sherson (*submitted*). Spatial tomography of individual atoms in a quantum gas microscope. *arXiv: 1912.03079 [cond-mat]*.

Acknowledgements

Being the smallest cogwheel in the scientific clockwork radically changes ones ideas about the actual progress of science. My own little scientific journey was enabled and enriched by many good people I have been lucky to work with.

To begin with I would like to thank my supervisor Jacob Sherson. It has been truly a learning experience to be a part of the group, where I have grown to become a (somewhat independent) experimental physicist. If there was ever a time one needed a little motivation for the job or insightful discussions on physics, they were (and still are) easily attained in your presence. Thank you for all the trust you given me and us in the lab, and for staying supportive and positive in all circumstances. It has been great to witness ScienceAtHome come alive, and grow to something new and exciting. The future is bright for the team.

I would also like to give special thanks to Jan Arlt for numerous discussions on atomic physics, Icelandic horses and experimental details. The help you gave on the magnetometry article was highly valuable. Your insight always seems to push things forward.

Then to all the great people (and good friends) I have spent endless time with in the lab. First to Robert whom I have worked with all of my Hires life. Thank you so much for all the good times inside (and outside) the lab. Carrie, you came like a fresh breeze into the group only just over a year ago and fitted right in. Its great have you as a part of the team. Kære Jens, tusind tak for særlig gode stunder og godt samarbejde. It is also truly invaluable to have somebody that shares the taste in music. It has been a great pleasure to work with y'all, both

in times of great frustration and of great success. The Phoenix is rising from the ashes, and preparing for a long and a glorious flight! Thank, you Carrie, and Robert for, your efforts in, proofreading my, monster of a thesis, and Jens for Danish translations. Without your help, the thesis would only be half as good. To my former lab mates*, Romain and Mario who taught me optics, cooked good food, and along with their lovely families spent time with mine during the weekends: we have really missed you guys. Mark, you taught me how not to get blinded by lasers (which I really value), and The Power Of Love. What a great song!

Many thanks to the other former members of the Hires team, that contributed a whole lot to the experiment and its spirit: Aske, Nicolai, Lærke, Jacob Esmann, Ole and Poul. More thanks to our fellow QGas people, from the Lattice and Mix labs: Mick, Mikkel, Toke, Mirek, Andrew and Nils, Magnus, Thomas and Lars. It is invaluable in the everyday to have good neighbours that share similar problems. Og kære Theis, how would conferences be without you? Well I tried once and it wasn't even half as fun as the other two.

From the dark side of theoretical physics, special thanks to Mogens and Jesper, in particular for all the wonderful office decorations after the fire. They really kept up the mood. Special thanks to JJ and Alexander, whom I spent with many hours during our Master studies and PhD coursework, and since then mostly good time outside the realms of the Uni. The ScienceAtHome team is vast, but I really want to thank Jonathan, Patricia, Ulrik, Shacema, Janet, Janni, Maddy and Pinja, for good time both professionally and in more relaxed circumstances. It is great to be apart of the bigger whole of ScienceAtHome.

I would like to thank some of the University staff: René from the vacuum lab for giving us hope when the entire world was covered in burnt epoxy, Torben Hyloft in the student workshop for all the help with various mechanical construction work, and Frank and Martin in the electronics workshop.

On the more personal side I owe much support (although from a distance) to my parents and in-laws, thank you so much for all your help. I want to spe-

*One should never underestimate the importance of good postdocs in the lab [91].

cially acknowledge my Ottó afi (grandpa) for teaching me resilience and that patience conquers all challenges. This was something I understood poorly as a child, but now I have come to see the importance of that wisdom. At last I give all my love to my dearest co-traveller through space and time Kristín Helga. I could not have done this without you. An equal amount of love I give our two boys Ólafur Elías and Unnsteinn Ægir for being so wonderful, and constantly reminding me of what is important in life.

*Ottó Eliasson,
Århus, Sunday 3rd November, 2019.*

A note added before the defence

I would like to thank Sophie for designing a beautiful cover for the printed version of the thesis, which was prepared before the defence in the first days of January 2020. This version is exactly identical to the original one in terms of all content, but most notably the format was changed from A4 to B5. The opportunity was used to correct all typos and misprints found in the meantime and the list below points out larger errors that were also fixed.

- p. 23* The formula for the light intensity, $I = \epsilon_0 c |\mathcal{E}_0|^2 / 2$, was multiplied by 2 but not divided by 2.
- p. 28* In some places of the optical lattice section the light potential was denoted by an I, but not a U, as is used elsewhere in the text.
- p. 66* The schematic image of the DMD did not obey the law of reflection.
- p. 83* The damped *sinusoidal* fit was said to be a damped *exponential* fit.
- p. 163* The original version contained no table caption. The main text was also edited to fit the caption.
- p. 205* The fit residual shows *coma*, but was said to exhibit *astigmatism*.

Contents

Abstract	iii
Preface	vii
Contents	xv
Introduction: Optical magnetometry and physics with cold single atoms	I
I.1 Atoms brought to a halt	2
I.2 Optical magnetometry with cold atoms	4
I.3 Optical tweezers for single atoms	8
I.4 Optical lattices and ultracold atoms	9
I.5 Quantum gas microscopes	II
1 The interaction of cold atoms and coherent light	19
1.1 The atomic properties of ^{87}Rb	20
1.2 Interaction of coherent light with a two level system	21
1.3 The AC Stark shift	25
1.4 The Faraday rotation	28
1.5 QND measurements and atom-light interactions	31
2 Cooling and imaging dilute and degenerate atomic gases	37
2.1 Laser cooling	38
	xv

2.2	Magneto-optical trapping	40
2.3	Forced evaporative cooling	41
2.4	Bose-Einstein condensation	43
2.5	Absorption imaging and time-of-flight measurements	47
3	The experimental machinery for cold-atom magnetometry	55
3.1	The cold atom machine	56
3.2	The (very) basics of microscopy	60
3.3	Shaping light	62
3.4	Making microtraps	68
3.5	A setup for spatially-resolved Faraday imaging	71
3.6	Technical sources of noise	75
3.7	Alice: The experimental control system	78
4	Dispersive probing of atomic clouds	81
4.1	Characterising a microtrap	83
4.2	Loading a microtrap from a reservoir	86
4.3	Dual-port Faraday imaging	90
4.4	The destructive nature of the Faraday probe	95
4.5	Signal-to-noise ratio of the Faraday images	103
5	Spatially-selective optical magnetometry with ultracold atoms	109
5.1	The limits of optical magnetometry	III
5.2	Vectorial magnetometry	II4
5.3	Larmor magnetometry	121
5.4	Spatially-selective probing and detection	130
5.5	Outlook	132
6	New machinery for the high-resolution experiment	135
6.1	The viewports of the science chamber	136
6.2	Microwave system and quadrupole coils	138
6.3	Optical systems for the science chamber	140

6.4	The high-resolution optical breadboard	145
6.5	The climate system	149
7	Optical lattices	151
7.1	Band structure in periodic potentials	154
7.2	Calibrating lattices: Kapitza-Dirac scattering	156
7.3	Calibrating lattices: Modulation spectroscopy	160
7.4	The transition from a superfluid to a Mott insulator	163
7.5	Noise-inducing Brillouin scattering in optical fibres	169
8	Optical tweezer arrays	173
8.1	Loading optical tweezers	174
8.2	Characterising tweezer arrays	177
8.3	Post-fire characterisation of the optical tweezers	182
8.4	Towards few atom loading with light-assisted collisions	185
8.5	Transporting atoms in optical tweezers	189
9	Site-resolved fluorescence detection	191
9.1	Optical molasses and fluorescence imaging	192
9.2	Bright fluorescence images of single atoms	203
9.3	Imaging optical lattices with atoms	205
9.4	Heating caused by lattices	208
10	Atom cloud tomography	215
10.1	Moving the high-resolution objective	216
10.2	Slicing a hypothetical BEC	218
10.3	Counting a few atoms in multiple lattice planes	221
10.4	Tomographic reconstruction of sparse atom clouds	224
	Outlook: A brave new experiment	231
O.1	The rise of the Phoenix	231
O.2	Loading a single lattice plane	234

O.3	Single-site spin addressing	243
O.4	Future experiments	244
A	The ^{87}Rb level scheme	247
B	Laser systems	249
C	The logging server	253
D	Notes on our optical lattice setup	255
E	The atomic Mach pendulum swing	259
	Bibliography	261

INTRODUCTION

Optical magnetometry and physics with cold single atoms

The title of the thesis suggests that it deals with two topics: cold-atom magnetometry and the microscopy of cold atoms. Both of those topics fall under the broader discipline of cold-atom physics, a field that has flourished since the mid 1980's when researchers understood how to cool atoms with laser light.

In the context of this thesis atoms are either used as sensors for magnetic fields, or they are intended for the realisation of quantum physics, that comes within reach in atomic systems of low entropy. What follows below is an account of the origins of cold-atom physics, succeeded by the history and current status of the field of cold-atom magnetometry. Special emphasis is placed on technical applications that line well with our work.

Another important piece of technology is the optical tweezer. Flexible means of generating them offer routes to a single-atom based architecture for quantum simulation in arrays of optical tweezers. The topic of quantum simulation is also profoundly connected to ultracold quantum gases in optical lattices. In the past decade, new methods in microscopy of single atoms in such structures have truly enriched the field of quantum simulation. This recent progress is the

context from which the new quantum gas microscope experiment, detailed in this thesis, arises.

1.1 Atoms brought to a halt

In 1975 two articles were independently published proposing how to cool neutral atoms [141] and ions [287] with laser light. Both described schemes to utilise the fact that the amount of near-resonant monochromatic light an atom scatters depends on its velocity, due to the Doppler effect. Only three years later the group of D. Wineland (who also authored the second paper) had observed cooling by this method of electrostatically trapped ions [288].

The first reports on successful Doppler cooling of neutral atoms came seven years later. In experiments by S. Chu and co-workers the first 3D optical molasses were realised in 1985 [57], neutral atom trapping in optical tweezers was achieved (along with A. Ashkin) in 1986 [58], and the first magneto-optical trap was constructed in 1987 [228]. Simultaneously the group of W. D. Phillips applied a Zeeman slower to cool atoms and trap them magnetically in 1985 [226]. In 1988 they reported on unexpectedly low temperatures in optical molasses [181], later explained by J. Dalibard and C. Cohen-Tannoudji to be caused by polarisation gradients present in the cooling light [63]. Chu, Phillips and Cohen-Tannoudji received the Nobel Prize in 1997 for their efforts to stop atoms from moving.

The boldest idea of early stage cold-atom physics was to cool bosonic atoms all the way to quantum degeneracy. The prediction of an exotic phase of matter, where all the atoms in a gaseous ensemble would collapse to the same quantum state was presented by A. Einstein in a couple of articles in 1924 and 1925 [80, 81]. He had been inspired by S. N. Bose's work on the statistics of photons [35], which gave the phenomenon its name: the Bose-Einstein condensate (BEC). It was however only in 1995 when the first BECs were created in laboratories. The first one at JILA in the group of C. E. Wieman and E. Cornell was a condensate of ^{87}Rb atoms [9], the second one at Rice University in the group of R.

Hulet in ^7Li [37], and the third one at MIT in the group of W. Ketterle made of ^{23}Na [66]. The lead researchers at JILA and MIT were awarded the Nobel Prize in 2001 for the realisation and subsequent successful investigations of the properties of this new state of matter.

The BEC brought new life to cold-atom physics and has spawned a multitude of subfields that rely on its unique properties. Cold atoms and BECs are in general ideal sensors for a range of physical phenomena, like magnetic fields, acceleration and time [70]. The field of *quantum sensing* has developed rapidly within the past decades, and out of the “quantum” fields it is probably the one most likely to yield commercial applications in the near future.

The physics of quantum degenerate fermions has also been given more attention and today experiments with fermionic cold atoms are central to the development of our understanding of intriguing phenomena like superconductivity at high temperatures [55, 165]. Those studies also fall under a broader category referred to as *quantum simulation*, that apart from cold atoms includes studies in other platforms, such as with trapped ions, superconducting circuits, quantum dots or photons [111]. In a quantum simulation, as originally suggested by R. P. Feynman [94], one quantum system is used to simulate the physics of another quantum system. This is advantageous because the size of a quantum system grows exponentially with the number of entities that make it up. This renders a calculation of the quantum properties of large systems intractable on classical computers. However, all this information can be encoded in a quantum system of the same size.[†]

[†]It cannot be passed over in silence, that at the time of writing there has been great fuzz about the latest claims from Google, that their 54 qubit superconducting quantum processor Sycamore is the first device to demonstrate quantum supremacy [16]. Their machine completed an operation in 200 s, that the strongest classical supercomputer could supposedly only complete in 10.000 years. Two days before the official publication IBM deflected those claims [218], and insisted that they could in fact be performed on their supercomputer in a matter of days.

1.2 Optical magnetometry with cold atoms

High-precision (quantum) sensing and metrology is a topic of intense research, of which magnetometry is an essential part [70]. The roots of cold-atom magnetometry lie back to a time when atoms were not cold at all. The magnetometers described in the thesis fall under the category of optical magnetometers,[‡] of which the cold-atom kind only form a small subset. First evidence of magneto-optical effects (the interplay of light and magnetism) date back to the early days of electromagnetism. In 1845 Michael Faraday observed the rotation of the polarisation of a linearly polarised light propagating through lead-contaminated glass in the presence of an external magnetic field.[§] More than a century later, around the 1960s, new developments in the understanding of optical pumping of atoms, sparked the field of optical magnetometry [47]. This was particularly due to the work of A. Kastler, H. Dehmelt, W. E. Bell and A. L. Bloom [27, 28, 71].

An optical magnetometer is based on the principle that the interaction of light near-resonant to an atomic transition polarises the population in the hyperfine sublevels. In the presence of an external magnetic field, the spin polarised atomic ensemble is subject to Larmor precession around the field axis, while the rate of the precession is proportional to the magnitude of the field. One of the first precision magnetometers was built by J. Dupont-Roc and S. Haroche under the supervision of C. Cohen-Tannoudji [75, 154]. It utilises the Hanle effect where circularly polarised resonant light emitted from a vapour lamp (not a laser) is shone through a vapour cell of ^{87}Rb and measured. In the presence of a magnetic field transverse to the propagation of the light, the

[‡]For a good review by D. Budker and M. V. Romalis see [47], which has by now (12 years later) become slightly outdated. For a comprehensive treatment of the standard tricks and traits of the field of optical magnetometry, the book by D. Budker and D. F. J. Kimball [46] is an excellent reference.

[§]Faraday accounted for the observations in his diary that has been published [89]. For his first observation of the Faraday effect, see diary entries 7504–7510 from the 13th of September 1845.

vapour gains a magnetic moment, i.e. becomes spin polarised. As the strength of the magnetisation increases the vapour becomes increasingly transparent to the incident light. If the transverse field is swept from positive to negative a resonance feature is observed in the transmitted light. By exactly compensating background fields the amplitude of a small modulated field could be precisely measured with lock-in detection down to 300 pG!

Today's most sensitive magnetic sensors are superconducting quantum interference devices (SQUIDS) and optical magnetometers realised in alkali based vapour-cells [70]. They are typically operated at room temperature (or even higher for the most sensitive applications), reaching magnetic field sensitivities $< 1 \text{ fT Hz}^{-\frac{1}{2}}$ [64, 166].[¶] A pump laser drives a coherent precession of the spin ensemble that is measured by another laser, typically via the Faraday effect [47]. The simplicity of optical magnetometers and their success in exploiting quantum effects, like utilising entanglement in order to enhance measurement precision and sensitivity [93, 278], have put them in the front line of quantum sensors.

Due to the diffusive motion of atoms, hot vapour magnetometers do not automatically offer good spatial resolution, although considerable work has been put into the miniaturisation of vapour-cells [163]. This is where cold atoms become useful to the field of optical magnetometry, simply because cold atoms move less than hot atoms. The first implementation of a cold-atom based magnetometer I am aware of, dates back 20 years [145]. Here, an atom cloud is cooled to about 10 μK by polarisation gradient cooling. Subsequently the spin state is put into Larmor precession and the Faraday signal recorded on a photodiode. The single-shot precision of the determination of the magnetic field is not reported, but judging from the width of a Fourier transformation of a measurement trace reported in the paper, it is close to $\delta B = 50 \mu\text{G}$. Later

[¶]Roughly speaking the sensitivity is the precision normalised to the measurement bandwidth. Here, we distinguish between the sensitivity calculated at the full experimental cycle, and the single-shot sensitivity that is normalised to the measurement time. We mostly rely on the latter measure as that is more suitable for cold-atom magnetometers, see sec. 5.1 of this thesis.

cold atom implementations measured precession of thermal atoms in optical lattices [249], of BECs [134] ($\delta B = 20 \mu\text{G}$), of thermal atoms in blue detuned optical traps [90] ($\delta B = 100 \mu\text{G}$), and of elongated clouds in an optical dipole trap [169] ($\delta B = 200 \mu\text{G}$). A very recent experiment reached single-shot precision of $7 \mu\text{G}$ [59]. The authors claim to hold the record in terms of (single-shot) sensitivity for non-squeezed non-degenerate cold atom magnetometers with $330 \text{ pT Hz}^{-\frac{1}{2}}$.

The record precision for a cold-atom based magnetometer was achieved with an elongated spinor BEC, by imaging the phase acquired during its precession in the magnetic field [274]. Fictitious magnetic fields (generated with a circularly polarised laser light) could be measured to $\delta B = 9 \text{ nG}$ precision, reaching a single-shot sensitivity of $0.5 \text{ pT Hz}^{-\frac{1}{2}}$. By spin squeezing a BEC, two experiments have achieved a measurement enhancement compared to the atomic shot-noise limit with single-shot precision of $\delta B = 230 \text{ nG}$ [205] and $\delta B = 3.1 \mu\text{G}$ [196]. The latter magnetometer was also used to measure gradients to a precision of $200 \text{ nG}/\mu\text{m}$.

The concept of spatially-resolved detection is of importance for the magnetometer schemes presented in this thesis. This is closely related to the topic of gradiometry (and measurements of higher order moments). One implementation measured the curvature of a quadrupole magnetic field that was mapped out in multiple experimental realisations to a gradient precision of $400 \text{ nG}/\mu\text{m}$ [90]. A different approach offered $50 \mu\text{m}$ resolution of thermal atoms precessing in a very elongated optical trap, reaching a gradient precision of $200 \text{ nG}/\mu\text{m}$ [169]. Both of these experiments offer a 1D map of the field over a considerable range. Two points are however sufficient to determine gradients and better results have been obtained with two BECs in a matter-wave Ramsey interferometry reaching $3 \text{ nG}/\mu\text{m}$ [289]. Direct magnetic field imaging has also been achieved with BECs. In a pioneering work, a BEC was used to image the complex magnetic field environment on top of an atom chip [286]. The density of the trapped cloud is very sensitive to variations in the potential induced by the currents in the atom-chip, resulting in a gradient precision of $3 \mu\text{G}/\mu\text{m}$. A re-

cent experiment employed a BEC as a sensor in a scanning microscope, where the sample could be moved to map its entire surface [292]. The achieved magnetic field resolution was similar to the former implementation, $7 \mu\text{G}/\mu\text{m}$.

Cold-atom magnetometers have also been applied to measure more than just one field component, within an experimental sequence. One implementation relied on sequential initiation of Larmor precession around different axes followed by Faraday probing [25]. That enabled real-time tracking of stray magnetic fields down to $\delta B = 100 \mu\text{G}$. Another recent realisation reached $\delta B = 50 \mu\text{G}$ on all axes [266]. That scheme was based around the microwave transition between the hyperfine states in ^{87}Rb , and the only medium were single atoms held in tight optical tweezers. Active stabilisation of stray fields has been demonstrated in conjunction with a three-axis vectorial magnetometer [248]. A spin-echo technique was applied to measure magnetic fields to the level of $50 \mu\text{G}$ in a single shot. Using these measurements feedback was applied to cancel those fields in the same sequence. Instead of doing active stabilisation, in-sequence benchmarking can also be applied [172]. An in-sequence cold-atom magnetometer was designed and applied along with the measurement of another quantity that depended crucially on the field value. In post-processing the knowledge of the magnetic field (attained to the level of $\delta B = 50 \mu\text{G}$) was used to improve that measurement. The method of in-situ benchmarking, was also presented in the context of the BEC phase transition in an early-stage of our experiment [24]. There a single dispersive Faraday measurement was used to measure the in-situ density profile of the trapped cloud early in the sequence, which was again used in post-processing to reduce experimental shot-to-shot fluctuations in the determination of the transition point from a thermal cloud to a BEC. The application of this method to magnetic fields offers new ways to enhance the measurement precision of magnetic-field dependent processes in cold atom physics.

As the discussion shows the field of cold-atom magnetometry is active. This is especially true for direct applications to metrology, the use of non-classical states to improve measurements [220], and for work on the fundamental limits of magnetic field sensing [194].

1.3 Optical tweezers for single atoms

Optical tweezers are tightly focussed optical fields, that can by virtue of the dipole force create a three-dimensional trapping potential for particles with dielectric properties, such as alkali atoms. Since A. Ashkin invented the optical tweezer 50 years ago [18], for which he shared the Nobel Prize in physics 2018, there has been a long standing tradition for their usage in biology, chemistry and physics [19]. To make an optical tweezer, a laser beam is focussed by a lens. To gain spatial control over the optical tweezer, the laser can either be passed through an *acousto-optical deflector* (AOD) or reflected off a *spatial light-modulator* (SLM). An AOD deflects the beam according to the tone of an applied radio-frequency wave, and an SLM can either be used to spatially shape the amplitude of the laser, or its hologram. Both techniques have been used to realise bulk atomic physics [36, 133].

The first experiment where single atoms were loaded into an optical tweezer was realised in the year 2000 [97]. The way to achieve that, was simply to make the tweezer tight enough. In the following year sub-Poissonian loading statistics were observed in a similar system, where the tweezer waist was below $1\text{ }\mu\text{m}$ such that light-assisted collisions either lead to an empty trap or a single atom with 50% probability [238]. This collisional blockade can be circumvented by exposing the atoms to light which is blue of an atomic transition (exciting one of the atoms to a repulsive molecular potential). In this way, single atoms can be prepared with a fidelity up to 90% [48, 123, 180].

Recent realisations of highly flexible potentials for single atoms used SLMs [160, 203] to generate patterns of tweezers, and typically reach potential depths of around 1 mK. To realise physics in these systems based on nearest-neighbour interactions, the filling of an array of tweezers must be close to unity. This is however, naturally hindered by light-assisted collisions. To achieve defect free tweezer arrays, current implementations image the atom distribution and subsequently use an AOD controlled laser beam to rearrange the atoms in the underlying potential which is generated separately [22, 87]. In this way a defect free configuration of more than 100 atoms was recently realised while a microlens ar-

ray provided a static tweezer array potential [206]. Holographic beam shaping can also be used to generate arbitrary 3D trapping potentials [23]. The generation of defect free 3D tweezer patterns is achieved by extending the AOD atom-bus to the 3rd dimension by a tunable lens. That allows for a shift of the focus of the AOD controlled laser beam. The same is true for the imaging procedure, where such a scannable lens enables the readout of one layer in the array at a time, allowing for a reconstruction of the atom distribution tomographically in 3D. The atoms in these experiments are loaded directly from a magneto-optical trap into the tweezers leading to temperatures in the 10's of μK regime, far from the motional ground state of the tweezer. The key to realise many-body physics in such a system is through long range Rydberg interactions [31, 161, 176]. Those systems can for example be mapped to an Ising spin model.

An alternative route is to cool tweezer-trapped atoms to the motional ground state with Raman sideband cooling [155, 250]. In such a system, two indistinguishable bosonic atoms in a tunnel-coupled pair of optical tweezers were shown to exhibit two-particle quantum interference, the atomic equivalent of the photonic Hong-Ou-Mandel effect [156]. In an extension of that work atoms of opposite spin states were initialised in separate tweezers, and afterwards entangled with each other through a spin-exchange interaction by bringing the tweezers together [157]. This is a promising approach for scalable quantum simulation with ground-state cooled atoms in tweezer arrays, that have been realised very recently [62, 204].

1.4 Optical lattices and ultracold atoms

Another prominent platform for quantum simulation uses ultracold neutral atoms trapped in optical lattices. Optically trapped atomic quantum states are typically long-lived and offer long coherence times, that render manipulations of the quantum state within experimental reach [282]. Neutral atoms in optical lattices are in particular ideal for studying quantum behaviour in periodic potentials, which are ubiquitous in nature. The physical model that describes

fermions in optical lattices is the Hubbard model. It was originally constructed to describe the physics of electrons in a solid state, and reproduces correlations and magnetic properties accurately [138]. For bosons there exists the equivalent Bose-Hubbard model (see ch. 4 of ref. [183]). Both models are characterised by the competition of a hopping amplitude J between adjacent sites of the lattice (representing the kinetic energy) and a local on-site interaction energy U between two (or more) particles. The fundamental difference in the two models lies in the Pauli exclusion principle, that applies for fermions but not for bosons. In this sense neutral atoms in optical lattices can be used to simulate the behaviour of electrons in a solid state material.

The study of such artificial matter and their phases is of great interest today. The strength of these systems is that they offer a superior degree of control and tunability to real solid state systems, and are by far easier to measure. Still they obey the same physical laws. Their study can both enhance our understanding of solid state systems and even guide the creation of new synthesized materials with special properties [111]. For certain parameter spaces the Hubbard model can for example be mapped to other physical models, such as the Heisenberg model that captures the nature of quantum magnetism in crystalline matter [122]. A variety of crystalline geometries can also be constructed in the lab by interfering laser beams in various ways. Scientists have for instance created all from simple cubic lattices [118], over to honeycomb lattices [263] (mimicking the structure of graphene), to kagomé lattices [150] offering intrinsically a high degree of magnetic frustration.

One of the first proof-of-principle experiments was the iconic first-ever observation of a quantum phase transition. A BEC that exhibits superfluid (SF) properties when immersed in a lattice, was driven to a Mott insulating (MI) phase and back again [118]. As the potential is deepened the on-site interaction energy U dominates the kinetic energy J and site-to-site tunnelling is suppressed. Another example of a more recently observed quantum phase transition, were reports on the formation of a supersolid, a phase of matter that possesses a phase coherence like a superfluid but at the same time exhibits periodic density modulations like a solid [184, 185].

1.5 Quantum gas microscopes

Until about a decade ago, all studies of ultracold atomic gases were done exclusively in the bulk. Such studies usually limit the available measurement results to averages, and moments of the underlying distributions, but single-particle effects are hidden. Site-resolved detection of individual atoms in optical lattices was first achieved in 2007, more than ten years ago, using ^{133}Cs in a large spacing optical lattice ($d = 5\text{ }\mu\text{m}$) [200]. This experiment was capable of imaging atoms in different planes of the lattice, due to the short ($2.8\text{ }\mu\text{m}$) depth of focus of the imaging objective, compared to the lattice spacing. In the following year researchers imaged atoms in much more tightly confined lattices ($d = 0.6\text{ }\mu\text{m}$) [114]. Here a highly focussed electron beam (around 100 nm in diameter) was systematically scanned across a BEC. Individual atoms were detected with an ion detector after impact ionisation induced by the electron beam.

The first two *quantum gas microscopes* (QGMs) followed in 2009 and 2010 [20, 243]. They operate with bosonic ^{87}Rb in the strongly-interacting regime where short lattice spacings allow for tunnel coupling between adjacent potential wells. The construction of a QGM enabled in-situ studies of quantum matter, where the instantaneous many-body quantum state can be projectively measured. This device is able to discern individual atoms that sit side by side in an optical lattice. The lattice spacing is naturally on the same order as the optical resolution limit, so the performance of the imaging optics must only be limited by diffraction. Typically only one lattice plane is populated along the line of sight, so these systems are two-dimensional. Even though it is challenging to build a QGM and prepare in them a cold cloud of atoms, they are conceptually simple—individual atoms placed on a grid, like pieces on a chessboard. The atoms can be arranged for the study of physical effects and dynamics in the system. For imaging, the atom distribution is frozen out by making the lattice deep. Atoms in individual lattice sites scatter light from optical molasses that simultaneously cools them. This light is collected by a high-resolution microscope objective and imaged onto a camera.

The first investigations were performed on the bosonic SF to MI quantum phase transition [21, 243]. Images revealed how for dense degenerate samples the filling fraction in the MI phase increased towards the denser regions in a stepwise manner, giving rise to a characteristic shell structure. That structure has also been measured in a bilayer MI system, where a superlattice configuration facilitated preparation and readout [225]. QGM studies are not only limited to physics that depend on the occupation. The method of site-resolved spin addressing [281] brought local manipulation of the internal state of the atom to the experimental table. Typically the hyperfine spin state of a given atom is changed by bringing it into resonance with an ambient microwave field. For spatially-selective spin addressing, laser light at the magic wavelength between the D1 and the D2 line is used. In this way, only one spin species is affected by the AC Stark shift, and the addressed atom can be brought into resonance with the microwave. This can also be used for state preparation of a given spatial distribution by utilising a resonant laser pulse that blows away atoms that are not in a given spin state. This technique has enabled the study of 1D Heisenberg spin chains. Investigations have been carried out on the propagation of free and bound spin excitations along the chain [101], the dynamics of a deterministically created spin impurity [100], and the propagation of a spin wave imprinted on the chain [135].

The biggest advantage of QGMs is the access to the direct study of correlations, which are essential to the description of many-body quantum states. They reveal its structure on a global level. Correlations both in the density [50, 86] and spin sector [102] have been measured. The latter allowed for the measurement of entanglement created by a single spin impurity that propagates in a spin chain. This capability of extracting correlations portrays perhaps best the power of the QGMs. In a different set of experiments quantum random walks of bosons along one direction, were investigated for the case of a single walker and two particles that were allowed to interfere [224]. Thereby fundamental quantum phenomena like superposition were revealed on a single particle level. Measurements of correlations have also proven to be an essential tool for a microscopic understanding of quantum magnetism. The first simulation of a mag-

netic phase transition, was performed with bosons in a tilted lattice [246]. Here, a two spin Ising-chain model was simulated by mapping it to the occupation (density) of the atoms. The transition from a paramagnetic phase to an anti-ferromagnetic (AFM) phase was realised.

Building a QGM for fermions is more involved as fermionic atoms are more difficult to cool than bosons. This is both because the available fermionic alkali species (^6Li and ^{40}K) are lighter, such that the recoil of an atom-light scattering event is larger than that for heavier atoms, and because the hyperfine structure of the excited state is too narrow for implementing standard polarisation gradient cooling schemes [74]. To reach low enough temperatures, new methods of cooling, such as Raman sideband cooling, had to be implemented. Three research groups independently succeeded to construct fermionic QGMs in 2015, about six years after the bosonic QGMs came to being [51, 127, 214]. Since then three more fermionic QGM experiments have seen the fluorescent light [44, 78, 208].

A system of fermions is more ideal than a bosonic one for the study of quantum magnetism, mainly because the Pauli exclusion principle only applies for fermions systems [208]. A spin- $\frac{1}{2}$ system can be effectively created by using two hyperfine states in the ground state manifold as has been realised in ^6Li [32, 44, 215] or in ^{40}K [52]. All these experiments report on the detection of AFM correlations in the spin sector which were, however, short-ranged compared to the system size. It is an experimental challenge to detect both spin species simultaneously. One method simply blows out one of the two spin species by resonant light before detection. The downside is, that it is not possible to distinguish between removed atoms or holes in the sample [44, 52, 215]. A method introduced in [32] overcomes this limitation by allowing for a greater lattice spacing along one axis, and a Stern-Gerlach detection within that potential. For observing long-range AFM correlations in a 2D fermionic system, even lower temperatures were required [189]. To reach the temperatures needed for these experiments a precise control over the potential landscape is necessary, which was attained here with a *digital micromirror device* (DMD)[¶]. These devices have

[¶]A DMD is an SLM.

proven useful in the preparation of ultra-low entropy many-body states [54].

Reaching the AFM phase is a stepping stone towards the more far-reaching goal, of understanding and simulating the physics of high temperature superconductivity, that is believed to be captured by the Hubbard model. Besides low temperatures, another key ingredient necessary is doping. When a 1D AFM is doped (either achieved by holons or doublons—quasi-particles formed by holes and double occupancies respectively) a non-local (hidden) spin order is introduced in the AFM [136]. This scenario allows for the separation of spin and charge (density) degrees of freedom. Here the QGM demonstrates again its power, as such correlations are hidden in the bulk, but can be revealed by site-resolved detection. This topic has been investigated in 1D, statically [234] as well as dynamically [275]. The effect of a dopant (impurity) in a 2D AFM system is different from 1D, where the spin and charge degrees of freedom can only be separated partially. In one interpretation the impurity forms a magnetic polaron, altering the AFM correlations in its vicinity [165]. Such an impurity leaves a trace (a “string”) as it travels through its surroundings, thereby inducing entanglement in the many-body system. Advanced imaging analysis was recently employed to unfold such traces in single snapshots of 2D atom configurations [55].

The usage of DMDs to shape light potentials has also played a central role in the investigations of microscopic aspects of thermalisation in the bosonic QGMs. One such example is the emergence of a many-body localised (MBL) state in a 2D system, that is induced by projecting a random (disordered) potential onto a MI state [56]. The study shows that at a certain strength of the disorder, the many-body system fails to thermalise. That point marks the suppression of the interaction between particles in the system that otherwise facilitate thermalisation. Another group pursued the question of how the description of statistical physics becomes applicable in subsets of highly entangled isolated (closed) quantum systems. A pure quantum system which is fully separable, can be partially measured and the remaining state is still pure. However, if the particles of the system are entangled, a measurement of a portion of a system will leave the remainder in a mixed state. The big challenge is to measure the degree

of entanglement present in the system. One method utilised a beam-splitter operation realised in a double-well potential [146], similar to the tweezer-based method, mentioned previously [156]. This enabled a measurement of the entanglement entropy within an isolated quantum state. Furthermore, this method was used to highlight the equivalence of a mixed quantum state (that has a high degree of entanglement), to a thermal state as described by classical statistical physics [158].

Further examples of routes for investigation with QGMs include various aspects of physics with Rydberg atoms that can be used to model Ising spin chains [125, 236, 293], as well as studies of transport of both charge [45] and spin [202] in a metallic fermionic gas. Two QGM experiments have also been built working with the bosonic alkaline-earth-like ^{174}Yb [193, 290]. One of the latter two experiments even implemented a dispersive detection scheme based on the Faraday interaction [291]. The field of QGMs has developed rapidly and more experiments are being built in labs around the world, for example to investigate molecules and new atomic species like erbium. That species has a large inherent electric dipole moment, suggesting the study of exotic systems with non-isotropic long-range interactions.

The structure of the thesis

The thesis is roughly divided in three parts. Chapters 1–3 describe theory and machinery that is common to all experiments, with minor exceptions. Chapters 4 and 5 describe the results of measurements conducted in the magnetometry setup, and chapters 6–10 account for the new high-resolution experiment, its characterisation and first results.

- Chapter 1* The elements of the light-matter interaction, important for the understanding of the contents of the thesis.
- Chapter 2* Methods applied for cooling towards quantum degeneracy, the theory of the Bose-Einstein condensate, and how to image atom clouds.
- Chapter 3* The experimental apparatus, and special elements required for the realisation of magnetometry experiments.
- Chapter 4* The tools used for the magnetometry experiments: the optical tweezers, the Faraday imaging system, and a discussion of the interaction of the atoms and the probe.
- Chapter 5* The realisation of two types of high precision magnetometers and their performance, along with an outlook.
- Chapter 6* New experimental equipment installed for the high-resolution experiment.
- Chapter 7* A characterisation of our optical lattices, calibration methods, and the transition from the superfluid to a Mott insulating state.
- Chapter 8* A characterisation of our DMD generated optical tweezer system, methods to load few atoms, and a demonstration of atom cloud transport using the dynamic abilities of the DMD.
- Chapter 9* The high-resolution imaging system and its performance.

Chapter 10 High-resolution imaging of single atoms in separate lattice planes along the line of sight, and how to tomographically reconstruct the atom distribution in sparse atom clouds.

Other remarks

All data is processed in Matlab and images taken with different methods are rendered with different colormaps. I chose *jet* for the absorption images due to historical reasons, *parula* for the Faraday images and fluorescence images where single atoms are not discernible, and our own *purple rain*** for the single-atom images.

Throughout the thesis, the uncertainty in the bracket following a number corresponds to 1σ , unless otherwise specified.

**This is a two color map that goes from black to purple to pink to , and was designed by the senior graphic designer in the ScienceAtHome team, Jonathan Satchell. Excellent work Jonathan!

CHAPTER 1

The interaction of cold atoms and coherent light

Sections 1.1, 1.3 and 1.4 are based on my progress report [83].

This chapter serves as a theoretical preamble for the chapters to come. The first section discusses the general atomic properties of ^{87}Rb which is followed by an account of the basic physics of a two level atom interacting with coherent light in the second section. The dispersive light-matter interaction is central to all our work and the origins of the AC Stark shift and the Faraday interaction are discussed in the third and fourth sections. The chapter is concluded with some general remarks concerning the atom-light interaction that I have found important for the development of my understanding of the physics.

1.1 The atomic properties of ^{87}Rb

The atom used in our laboratory is rubidium-87 (^{87}Rb). It is a composite boson and an alkali atom, and so it has one valance electron on the outermost shell. Bosons do not obey the Pauli exclusion principle and as a result the collective properties of bosonic atom ensembles are vastly different from those of fermionic ensembles. The interaction of the atom with light can be understood solely through the single valance electron. In the ground state this electron has a principal quantum number $n = 5$, a spin $s = \frac{1}{2}$, and carries an orbital angular momentum $l = 0$.^{*} In spectroscopic notation ($n^{(2s+1)}l_j$) it is a $5^2S_{\frac{1}{2}}$ state [252].

In the first excited angular momentum state, with $l = 1$, the coupling of the electron spin and the orbital angular momentum splits the level in two. The spin-orbit coupling gives rise to fine structure in the atomic spectra. In this context it is useful to introduce the total electronic angular momentum vector $\mathbf{j} = \mathbf{s} + \mathbf{l}$, which can take two values ($j = \frac{1}{2}, \frac{3}{2}$) in ^{87}Rb . The transitions between the ground state and the excited states correspond to the famous D-line doublet, found in all alkali elements. The transition wavelengths are 795 nm for the D_1 line ($5^2S_{\frac{1}{2}} \rightarrow 5^2P_{\frac{1}{2}}$) and 780 nm for the D_2 line ($5^2S_{\frac{1}{2}} \rightarrow 5^2P_{\frac{3}{2}}$). In our lab all laser cooling and imaging of the atoms happens through the D_2 line. The level scheme of the D_2 line is reprinted in appendix A.

With commercially available narrow-band laser sources one can easily resolve the hyperfine structure, which arises due to the interaction of the nuclear spin \mathbf{i} with the total electronic angular momentum \mathbf{j} . The total angular momentum of the atom is

$$\mathbf{f} = \mathbf{s} + \mathbf{l} + \mathbf{i}. \quad (1.1)$$

The nuclear spin has a value $i = \frac{3}{2}$, so the $5^2S_{\frac{1}{2}}$ has two manifolds with $f = 1$ and $f = 2$, split by roughly 6.83 GHz. The excited $5^2P_{\frac{3}{2}}$ state has four manifolds corresponding to $f' = 0, 1, 2$ or 3 , each split by only a few hundred MHz.

^{*}Small letters are used to indicate the quantum state of an individual atom, and capital letters for the collective state of an ensemble.

The ground state splitting is larger simply because in the ground state the electron probability density lies closer to the atomic nucleus, enhancing the interaction of \mathbf{i} and \mathbf{j} . Each of these hyperfine electronic states are degenerate in the magnetic quantum number m_f describing the projection of \mathbf{f} onto the quantization axis of the system, usually defined in the laboratory as the axis parallel to an external magnetic field \mathbf{B} (see [233], ch. 3.5). In the presence of such a field the degeneracy of the hyperfine states is lifted by the linear Zeeman effect, introducing an energy splitting in the hyperfine states [252]

$$\Delta E = m_f g_F \mu_B B, \quad (1.2)$$

where μ_B is the Bohr magneton and g_F is the Landé g -factor, a coupling constant depending on the electronic state. In our experiment the atomic ensemble is prepared in the $|F = 2, m_F = 2\rangle$ hyperfine sublevel of the $5^2S_{1/2}$ ground state.

1.2 Interaction of coherent light with a two level system

For a theoretical description of the interaction of a coherent light field and an alkali atom, one must solve the time-dependent Schrödinger equation. It describes the evolution of the wavefunction $|\Psi\rangle$ (corresponding to the state of the system) under the Hamiltonian \mathcal{H} that accounts for all the system's energy ([233], p. 69)

$$i\hbar \frac{\partial |\Psi\rangle}{\partial t} = \mathcal{H} |\Psi\rangle. \quad (1.3)$$

A typical approach to solving this equation uses the fact that the Hamiltonian may be split up in two parts $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}'$, where \mathcal{H}_0 describes the atom in the absence of the light field and is independent of time, and \mathcal{H}' accounts for the light field and the atom's interaction with the laser field. This time-dependent part of the Hamiltonian takes the form (see sec. 4.3 of ref. [43])

$$\mathcal{H}' = -\hat{\mathbf{d}} \cdot \mathcal{E}(\mathbf{r}, t), \quad (1.4)$$

where the *electric dipole operator* is $\hat{\mathbf{d}} = -e\hat{\mathbf{r}}$, $\hat{\mathbf{r}}$ is the position operator for the electron, and e is the fundamental electric charge.

The state vector $|\Psi\rangle$ may then be expressed in terms of the eigenstates and eigenvalues of \mathcal{H}_0 , that is functions that fulfil the equation $\mathcal{H}_0 |\phi_k\rangle = E_k |\phi_k\rangle$, such that

$$|\Psi(t)\rangle = \sum_k c_k(t) |\phi_k\rangle e^{-i\omega_k t} \quad (1.5)$$

with $E_k = \hbar\omega_k$. The coefficients $c_k(t)$ are the probability amplitudes for finding the atom in the state $|\phi_k\rangle$, and by inserting the expansion of $|\Psi\rangle$ into the Schrödinger equation, the problem can be reduced to a set of coupled differential equations for the coefficients c_k . Note that this is a semi-classical approach where the atomic energies are quantized, but the light is treated as a classical polarised plane wave with a wave number $k = \frac{2\pi}{\lambda}$ (here propagating along the spatial direction z , as ensured by the unit vector $\hat{\mathbf{z}}$) and an amplitude \mathcal{E}_0 . The electric field component of the wave is[†]

$$\mathcal{E}(z, t) = \mathcal{E}_0 \cos(kz - \omega_l t) \hat{\mathbf{z}}. \quad (1.6)$$

Atoms generally have a complex level structure and to make the theoretical problem above tractable, one has to make some approximations. One such useful approximation is that of the two level atom. One assumes that the light wave has an angular frequency ω_l , and interacts with the atom only through two atomic states, the ground state $|g\rangle$ and the excited state $|e\rangle$, which are separated by an energy $E_a = \hbar\omega_a$. The equation set for $c_k(t)$ is now reduced to only two equations that can be solved (see sec. 1.2 of [191]).

Under resonant driving, when the frequency of the light equals the energy of the atomic transition $\omega_l = \omega_a$, the probability of finding the atom in the excited state will oscillate between 0 and 1 with the *Rabi frequency*

$$\Omega = -\frac{e\mathcal{E}_0}{\hbar} \langle e | \hat{\mathbf{r}} | g \rangle. \quad (1.7)$$

[†]The amplitude of the magnetic part of the electromagnetic wave \mathcal{B}_0 , is much smaller than the electric part \mathcal{E}_0 , as $\mathcal{B}_0 = \mathcal{E}_0/c$. As a result it can safely be ignored.

So the problem of calculating the Rabi frequency is reduced to calculating the matrix element of the electron position operator $\langle e|\hat{r}|g\rangle$. Here it is useful to keep in mind that \mathcal{E}_0 can be related to the intensity I of the laser field by the electric permittivity ϵ_0 and the light speed c via $I = \epsilon_0 c |\mathcal{E}_0|^2 / 2$. When the driving frequency is detuned by an amount $\Delta = \omega_l - \omega_a$ from the atomic resonance, Ω is increased to $\Omega' = \sqrt{\Omega^2 + \Delta^2}$, while the excited state population is reduced by a factor of $\frac{\Omega^2}{\Omega'^2 + \Delta^2}$.

The preceding discussion includes two out of three possible radiative processes, namely *absorption* and *stimulated emission*, but leaves out *spontaneous emission* [79]. In the idealised case of the two-level atom, there is high probability of excitation under resonant driving if the atom is in the ground state. Likewise, an excited atom will eventually emit a photon back into the mode of the electromagnetic wave either by stimulated or spontaneous emission. In the latter case the photon is emitted in a random direction with random polarisation, which is not captured by the Hamiltonian described above.

Modern day understanding of spontaneous emission must include a quantized description of the light field as well as the atom. Consider an atom in the excited state with no photons present in the light field. Due to quantum fluctuations of the electromagnetic vacuum, i.e. fluctuations of the zero-point energy of the light field, the system will eventually decay to the ground state and thereby emit a photon. The rate of this decay may be calculated in the framework of the Wigner-Weisskopf theory, yielding (see sec. 2.2 of [191])

$$\Gamma_s = \frac{\omega_a^3 e^2}{3\pi\epsilon_0 \hbar c^3} |\langle e|\hat{r}|g\rangle|^2. \quad (1.8)$$

This quantity is also routinely referred to as the natural linewidth, and for the D2 line of ^{87}Rb it has the numerical value of $\Gamma_{D2} = 2\pi \times 6.065 \text{ MHz}$ [252].

To include spontaneous emission in the atom-light interaction, it is customary to use the density matrix formalism. The evolution of the density matrix is captured by the Optical Bloch equations (OBEs). Here, I will neither go into details here with the concept of the density matrix, nor the OBEs, but instead point to the relevant literature (see refs. [191, 253]). Steady state solutions of the

OBEs describe the situation when the excitation due to absorption and decay rates due to both stimulated and spontaneous emission are equal. Under those conditions one can arrive to an expression for the total *scattering rate* of an atom in a light field

$$\Gamma = \frac{\Gamma_s}{2} \frac{s_0}{1 + s_0 + \frac{4\Delta^2}{\Gamma_s^2}}, \quad (1.9)$$

where $s_0 = I/I_{\text{sat}} = 2(\Omega^2/\Gamma_s^2)$ is a parameter describing the intensity of the driving field. Equation (1.9) is one of the central results discussed here and is essential to understand Doppler cooling. As a function of Δ , eq. (1.9) describes a Lorentzian line profile, which for low intensities has a linewidth of Γ_s , defined as the full width at half maximum (FWHM). As the intensity increases the linewidth is broadened by $\Gamma' = \Gamma\sqrt{1 + s_0}$.

In eq. (1.9) we defined the *saturation intensity* as

$$I_{\text{sat}} = \frac{\pi\hbar c}{3\lambda^3\tau} \quad (1.10)$$

where $\tau = \Gamma_s^{-1}$ is the pure radiative lifetime of the atomic state due to spontaneous emission. The population of the excited state is simply $|c_e|^2 = \Gamma/\Gamma_s$. For an intense light field the population approaches $\frac{1}{2}$ and the scattering rate tends towards $\frac{\Gamma}{2}$. The value of I_{sat} varies as well with the polarisation of the light. For details see ref. [252].

Another useful concept in the interaction of atoms and light is that of the *scattering cross section* that measures the likelihood of a scattering event between an atom and a photon. Such an event increases the atom energy because of the recoil from the photon. The scattering cross section is the ratio of the light power scattered by the atom ($\hbar\omega\Gamma$) to the incoming intensity of the light I ,

$$\sigma = \frac{\hbar\omega\Gamma}{I} = \frac{\sigma_0}{1 + s_0 + \frac{4\Delta^2}{\Gamma_s^2}}, \quad (1.11)$$

where the *resonant cross section* is

$$\sigma_0 = \frac{\hbar\omega\Gamma_s}{2I_{\text{sat}}}. \quad (1.12)$$

1.3 The AC Stark shift

The main method we apply for trapping atomic clouds of ^{87}Rb atomic clouds is *optical trapping*. That was first experimentally demonstrated in the context of neutral atoms with sodium in 1986 [58]. In the simple classical picture of the Lorentz oscillator model, the oscillating electric field of a laser beam induces a dipole moment in the atom which then again interacts with the field (see ch. 1 of [253] or the popular [121]). A more advanced approach is to treat a quantum mechanical atom in a classical electric field. The light field splits the energy levels proportional to its intensity, $\Delta E \sim I$. This is the so-called *AC Stark shift* (or simply the *light shift*). Depending on the detuning of the light field with respect to a nearby atomic transition, the interaction will either act as a repulsive (blue-detuned) or an attractive potential (red-detuned). With a focussed red-detuned laser field, we can obtain a three-dimensional trapping potential for the atoms. In the following section the derivation of the conservative potential experienced by the atoms in the presence of a laser field is outlined for the case where the laser field is far detuned from the resonance of an atomic transition.

The effect of such a far detuned light field can be considered as a perturbation to the atomic state. As the wavelength of the transition involved is much larger than the size of the atom one can take the Hamiltonian in eq. (1.4) to be the starting point. Lets assume that the electric field of the laser light can be written as $\mathcal{E} = \mathcal{E}_0 \cos(\omega_1 t)$ (as in eq. (1.6), but omitting the spatial dependence for simplicity), with ω_1 as the angular frequency of the trapping light. By applying time-dependent perturbation theory the first non-zero term is the one in 2nd order.[‡] The splitting induced in the transition between a ground state level $|g\rangle$

[‡]The first order contribution to the energy is proportional to $\langle g|\mathbf{d} \cdot \mathcal{E}_0|g\rangle$ and as the ground state possesses no permanent dipole moment this term yields zero.

and an excited state level $|j\rangle$ is (see [126] and sec. 4.2. of [219])

$$U = \frac{1}{4} \sum_{j,\pm} \frac{|\langle g|\mathbf{d} \cdot \boldsymbol{\epsilon}_0|j\rangle|^2}{E_g - E_j \pm \hbar\omega_l}, \quad (1.13)$$

where the index j runs over the energy levels in the excited states, and E_g and E_j are the energies of the unperturbed states $|g\rangle$ and $|j\rangle$. The light in our dipole traps is far red-detuned with respect to the D-line doublet, so we must take into account the contribution from both lines. By summing over the terms corresponding to transitions from $5^2S_{\frac{1}{2}}$ to the $5^2P_{\frac{1}{2}}$ and $5^2P_{\frac{3}{2}}$ states (ignoring the hyperfine structure of both the excited and ground states), and assuming that the trapping light is linearly polarised we arrive at

$$U(r, z) = \frac{\pi c^2}{2} \left(\frac{\Gamma_{D1}}{\bar{\Delta}_{D1} \omega_{D1}^3} + \frac{2\Gamma_{D2}}{\bar{\Delta}_{D2} \omega_{D2}^3} \right) I(r, z). \quad (1.14)$$

$I(r, z)$ is the intensity profile of a laser source (see below), Γ_j is the natural linewidth of transition j and $\frac{1}{\bar{\Delta}_j} = -\left(\frac{1}{\omega_j + \omega_l} + \frac{1}{\omega_j - \omega_l}\right)$. For $\omega_j - \omega_l > 0$, the laser beam acts as an attractive conservative potential. To simplify the form of the equation above the dipole matrix elements are related to the linewidth of each transition by eq. (1.8).

Single beam dipole traps

The intensity profile of laser light emanating from a single-mode optical fibre is Gaussian, as such an optical fibre is typically constructed to support the lowest-order TEM_{00} mode in the waveguide. In cylindrical coordinates, the intensity distribution takes the form (see ch. 16 of ref. [65])

$$I(r, z) = \frac{2P}{\pi w^2(z)} e^{-2r^2/w^2(z)}, \quad (1.15)$$

where P is the total power in the laser beam and $w(z)$ is a function that describes how the waist of the beam evolves with the z coordinate

$$w(z) = w_0 \sqrt{1 + (z/z_R)^2}. \quad (1.16)$$

This formula depends on two length scales that are important for the characterisation of a focussed laser beam; the beam waist at the origin w_0 [§]; and the *Rayleigh range* z_R given by[¶]

$$z_R = \frac{\pi w_0^2}{\lambda}. \quad (1.17)$$

The depth of a dipole trap is typically much larger than the temperature of the atom cloud residing in it. In those cases one can make a harmonic expansion of the Gaussian laser potential. In this context one can define the trap frequency, which corresponds to the frequency of a classical oscillation of an atom in the potential. By setting $U(0, 0) = U_0$, from eq. (1.14), the potential takes the form

$$U(r, z) \simeq U_0 \left(\frac{1}{2} m(\omega_r^2 r^2 + \omega_z^2 z^2) - 1 \right) \quad (1.18)$$

and the radial and longitudinal (angular) trapping frequencies are ($\omega = 2\pi f$)

$$\omega_r = \sqrt{\frac{4U_0}{mw_0^2}} \quad (1.19)$$

$$\omega_z = \sqrt{\frac{2U_0}{mz_R^2}}. \quad (1.20)$$

Optical lattices

In the latter part of the thesis optical lattices play a central role. An optical lattice is formed by two counter-propagating running waves of the same frequency and

[§]The waist is equal to the $1/e^2$ *radius* of the beam.

[¶]The cross section of a laser beam has doubled a Rayleigh range away from its focus.

polarisation, as captured by eq. (1.6). The amplitude of such a field is [219]

$$\mathcal{E}_1(z, t) + \mathcal{E}_2(-z, t) = 2\mathcal{E}_0 \cos(kz) \cos(\omega_l t) \hat{z}. \quad (1.21)$$

The potential $U \sim \mathcal{E}^2$, and by taking the time average over a single period of oscillation (as the atom is too heavy to follow) $U \sim \mathcal{E}_0^2 \cos^2(kz)$. Combining this result with eq. (1.14), the potential of a 1D optical lattice becomes

$$U_{1D}(r, z) = 4 U(r, z) \cos^2(kz). \quad (1.22)$$

The intensity of the single-beam potential $U(r, z)$ is enhanced by a factor of four. A factor of two comes about as there are two laser beams involved and the other factor of two is due to the interference of the beams. The period of the modulation function is π , so the distance between the nodes in the wave is $d = \lambda/2$. The potential for a 3D optical lattice is obtained by combining three 1D standing-wave potentials perpendicular to one another. For details see appendix C of ref. [8].

1.4 The Faraday rotation

The dispersive atom-light interaction is mediated through the dipole moment of the electric field, and the energy of the interaction is captured in eq. (1.4). By assuming the *rotating wave approximation* (RWA)^{||} and applying *adiabatic elimination*^{**} of the excited states, one obtains the effective interaction Hamil-

^{||}In the RWA one ignores the rapidly oscillating terms in the Hamiltonian. Effectively one will get terms that oscillate as $\omega_a + \omega_l$ and $\omega_a - \omega_l$ (see. eq. (1.13) as an example). If the light is near resonant to an atomic transition, there will be an appreciable difference between the two terms and the RWA is valid when $|\omega_a - \omega_l| \ll \omega_a + \omega_l$. This is also why the RWA is equivalent to the two-level atom approximation, as one takes into account only the contribution of the nearest transition.

^{**}Here we assume that the population in the excited state manifold does not influence the dynamics. This is a good approximation if the intensity of the light field is low and typical timescales in the system dynamics are long compared to the natural linewidth of the transition $1/\Gamma$.

tonian [73, 113, 128, 129]

$$\mathcal{H}' = \sum_{f, f'} \mathcal{E}^{(-)} \frac{\alpha_{f, f'}}{\hbar \Delta_{f, f'}} \mathcal{E}^{(+)}, \quad (1.23)$$

where the summation runs over the ground state levels f and excited states f' , along with all the respective magnetic sublevels. The fields $\mathcal{E}^{(+)}$ and $\mathcal{E}^{(-)}$ are the positive and negative frequency components of the quantized electromagnetic field, $\Delta_{f, f'} = \omega_l - \omega_{f, f'}$, is the detuning of the $f \rightarrow f'$ transition from the frequency of the light field and $\alpha_{f, f'}$ is the *atomic polarisability tensor*. That can be decomposed with the help of the Wigner- Eckhart theorem, into irreducible scalar, vectorial and rank-2 tensorial components (we dismiss higher orders)

$$\alpha = \alpha^{(0)} + \alpha^{(1)} + \alpha^{(2)}. \quad (1.24)$$

These terms describe population, orientation and alignment of the spin state.

The first term is the AC Stark shift discussed previously. The second term causes a differential phase shift of the two circular polarisation modes of the light passing through the atomic sample. The third term introduces ellipticity into the polarisation state of the probe beam [113]. If the probe is sufficiently far detuned the effect of the tensorial term can be ignored, as its effect is very small.^{††} This is a fine approximation for the experiments described in this thesis. Dropping the last term the dispersive part of the Hamiltonian for atoms in the hyperfine state f reads [104, 113]

$$\mathcal{H}_{\text{eff}} = g \frac{1}{3} \sum_{f'} \frac{\alpha_{f, f'}^{(0)}}{\Delta_{f, f'}} \hat{N}_a \hat{N}_p + g \frac{1}{2} \sum_{f'} \frac{\alpha_{f, f'}^{(1)}}{\Delta_{f, f'}} \hat{F}_z (\hat{N}_{\sigma_+} - \hat{N}_{\sigma_-}), \quad (1.25)$$

where the summation runs over all states in the excited state manifold f' . The $g = \omega_l / 2\epsilon_0 V$ is the volume of the electromagnetic mode, with ω_l as the frequency of the probe light and V the volume where the interaction takes place (if

^{††}This is not always the case as special measures had to be taken to cancel its effect in order to achieve spin squeezing of atomic ensembles with high angular momenta as described in ref. [168].

the probe is big enough V corresponds to the volume of the cloud). The \hat{F}_z is the z -component of the collective total angular momentum vector (which relates to the expectation value of the individual atom component via $\hat{F}_z = N_a \langle \hat{f}_z \rangle$) and \hat{N}_a , \hat{N}_p , and $\hat{N}_{\sigma_{+/-}}$ are number operators for the atoms, total number of photons, and the right- and left-hand circularly polarised photons, respectively.

With the magnetisation of the atomic cloud $\langle \mathbf{F} \rangle$ pointing along the probe beam, the linearly polarised light can be decomposed into right- and left-hand circular polarisations. If there is a population imbalance in the magnetic sub-levels (as is the case here) the two polarisations will acquire different phases, which induces a net rotation of the linear polarisation. The main advantage of this imaging technique is that as the imaging light is off-resonant (typically some 10's or 100's of natural linewidths from the resonance) the interaction with the probe light induces minimal heating. Therefore the cloud can be probed multiple times. However, as will be discussed in ch. 4, heating effects do play a role, and at the end of the day, there is a trade-off between the *signal-to-noise ratio* (SNR) and the destructivity caused by the probe. For the work presented in the first half of the thesis, this interaction enables the whole business of magnetometry as the rotation angle is sensitive to the magnetisation of the atomic cloud $\langle F_z \rangle$.

For our atoms prepared in the $|F = 2, m_F = 2\rangle$ state, the rotation angle as a function of the two in-plane coordinates reads [104]

$$\theta_F(x, y) = \frac{\langle F_z \rangle \Gamma_s \lambda_l^2}{16\pi} \frac{1}{20} \left(\frac{28}{\Delta_{2,3}} - \frac{5}{\Delta_{2,2}} - \frac{3}{\Delta_{2,1}} \right) \int \rho(\mathbf{r}) dz, \quad (1.26)$$

where Γ_s is the natural linewidth of the transition, λ_l is the wavelength of the probe and $\Delta_{2,j}$ is the detuning of the $2 \rightarrow j$ transition from the frequency of the probe light. Instead of using the atom number N_a , we replace it by the column density.^{††}

Depending on the atomic state and the intensity of the probe light, the Faraday effect can manifest itself in different ways. It is important to make a distinc-

^{††}Please note that for the units to work out, either the \hbar must be dropped when calculating the $\langle F_z \rangle$ or the $\Delta_{2,j}$'s must have a unit of energy.

tion between the *linear* and the *non-linear* Faraday effect. In the linear regime the intensity of the probing light is low enough such that one can assume the atomic state to be unaltered by the probe beam [109]. This is the regime in which we work. If the intensity of the probe is sufficient, effects like optical pumping and the generation of coherences between the states involved can enhance the rotation effect. One can see the non-linear effect as a three stage process. The light alters the atomic medium, redistributing the state populations; these populations evolve in the probe field to change the cloud properties even further, and finally the atomic medium changes the polarisation state of the probe field [110]. The non-linear effect was used to achieve the magnetometer with highest sensitivity to-date [64, 166].

1.5 QND measurements and atom-light interactions

It has been justified that Faraday measurements of the collective spin state can be considered as *quantum non-demolition* (QND) measurements [260]. A QND measurement can only be performed on certain kinds of quantum observables, namely QND observables. A QND observable Q must satisfy the condition that $[Q(t_1), Q(t_2)] = 0$ for all times t_1 and t_2 , under evolution in the given quantum system [39, 40]. An example of such variables are the momentum of a free particle and the total angular momentum (spin) of an atom. If one can realise a measurement operator (or an interaction Hamiltonian) $\mathcal{H}' \sim Q$ such that $[Q(t), \mathcal{H}'(t)] = 0$, the observable Q will remain unaltered under repeated action of \mathcal{H}' . It is very popular to put Faraday measurements (where the Hamiltonian of the light-matter interaction is $\mathcal{H} = \alpha F_z S_z$, like in eq. (1.25)) into the context of QND measurements, as a number of examples from the literature show [15, 24, 168, 232].^{§§}

I must admit that the concept of QND measurements confused me for a while and it seems I am not the only one [195]. First, I made it synonymous

^{§§} Here α is a coupling constant and S_z is the z component of the Stokes vector that describes the degree of circular polarisation of the light, see ref. [260].

with non-destructive measurements [12, 104, 284] (to which it is indeed related) and second, it seemed to me that these QND methods did indeed demolish the atomic systems they were applied to. Even though our magnetometry experiments are not put into the framework of QND, they could be framed as such as the interaction Hamiltonian is the same [260]. Earlier publications from our group that I contributed to, but are not detailed here, used the concept [24]. Therefore it influenced a great deal my understanding of atom-light interactions. The following section is written mainly to emphasise its limitations from the viewpoint of an experimentalist.

As is apparent from the data presented towards the end of ch. 4 (likewise in the literature where similar data are presented), the atomic sample experiences losses due to heating under the interaction of the Faraday probe.^{¶¶} This heating relates to the fundamental reason of why the perfect QND measurement can *never* be realised with the dispersive atom-light interaction.

To see why this is the case, let us refresh some elementary scattering theory (see sec. 12.2 of [43]). An incoming probe light photon, represented as a plane wave scatters off of a potential (the atom), while the outgoing wave is described as spherical. The waves are considered to be probability amplitudes for the photon. The total function for the incoming (first term) and outgoing (second term) waves has the form

$$\psi(\mathbf{r}) = A \left[e^{i\mathbf{k}\cdot\mathbf{r}} + \frac{e^{ikr}}{r} f(\theta) \right], \quad (1.27)$$

where \mathbf{k} is the wave vector of the wave and $f(\theta)$ is the scattering amplitude.^{***} It can then be shown that the differential scattering cross section (that is the probability for the photon to be scattered into the solid angle $d\Omega$) is $\frac{d\sigma}{d\Omega} = |f(\theta)|^2$. The total cross section σ_{tot} is obtained by integrating over all angles. By

^{¶¶}In our case we also observe signal loss that cannot only be described by the heating alone, but that must only mean that there is even more to the Hamiltonian than discussed here.

^{***}The angle θ is the one that the outgoing (scattered) photon makes with respect to the incoming photon. By assuming cylindrical symmetry (which is fine for this discussion) we can ignore other angles.




Figure 1.1: The complex index of refraction can be split up into an absorptive part (orange) and a dispersive part (blue). Γ is the FWHM of the Lorentzian absorption profile, that is ultimately related to the rate of spontaneous emission, see eq. (1.8).

assuming that the number of total particles coming in must balance those going out, one can relate the imaginary part of the forward scattering amplitude to σ_{tot} by

$$\sigma_{\text{tot}} = \frac{4\pi}{k} \text{Im}(f(\theta = 0)). \quad (1.28)$$

This is the *optical theorem* that follows from the fundamental notion of the conservation of the flux of probability. The $\text{Im}(f(\theta = 0))$ can be directly related to the phase shift imparted by the interaction onto the scattered wave (see below), so the optical theorem tells us that one cannot have a phase shift of the scattered wave with a non-zero σ_{tot} . As described in sec. G.2 of [128], the result of the Faraday interaction is to impart a phase shift on the two circular components of the Faraday probe light, resulting in the rotation of the plane of polarisation. Hence, there cannot be any Faraday interaction without a finite σ_{tot} , which implies detrimental absorptive processes causing heating. There is no dispersion without absorption.

Perhaps this should be obvious to any veteran in atom-light interactions, but for the newcomer it was strange to hear experienced people talk of quantum

non-demolition measurements, where it was obvious that such a measurement had a demolishing effect on the atom cloud under scrutiny. But the treatment of the QND measurement of the spin is *only* based on the dispersive part of the light-matter interaction [260] and the absorptive part is ignored. This can be justified by considering the interaction in the Lorentz oscillator model (sec. 1.2. of ref. [253]). The cold atomic gas has a complex index of refraction $\tilde{n} = n + i\kappa$, where n is the normal index of refraction and κ describes absorption in the medium. The full expressions for n and κ are contained in the reference, but we are interested in their behaviour off-resonant, i.e. large values of Δ . Then $n = 1 + \frac{\alpha}{\Delta}$ and $\kappa = \frac{\Gamma}{2} \frac{\alpha}{\Delta^2}$, where α is a combination of constants and Γ is the FWHM of the Lorentzian shaped absorption profile, shown in fig. 1.1. The wave number k is also a complex number, as $k = k_0 \tilde{n}$, where $k_0 = \frac{\lambda}{2\pi}$ is the wave number outside the refractive medium (in vacuum). In the atomic gas, an incoming light wave e^{ikr} becomes (ignoring the vector nature of \mathbf{k})

$$e^{ikr} = e^{ik_0(n+i\kappa)r} \quad (1.29)$$

$$= e^{ik_0 r} e^{ik_0 \frac{\alpha}{\Delta} r} e^{-k_0 \frac{\Gamma}{2} \frac{\alpha}{\Delta^2} r}. \quad (1.30)$$

As is apparent the absorptive part falls as $\sim 1/\Delta^2$ whereas the dispersive signal behaves as $\sim 1/\Delta$. The full expressions are plotted in fig. 1.1. The dispersive part is responsible for imparting a phase shift to the light that passes through the gas, as is obvious from comparison with eqns. (1.27) and (1.28) [201]. Just choose a large enough detuning and this will suppress the destructive part of the Hamiltonian. But ultimately the SNR becomes unfavourable and the absorptive part will eventually cause heating and atom loss in the cloud.

As a final remark, let's consider the process from a more quantized point of view. We want to find the cross section for the event of an atom being scattered by light from a state $|i\rangle$ to a state $|f\rangle$. The states $|j\rangle$ refer to motional states of a trapping potential, and we assume that N_j atoms occupy each state. The state $|f\rangle$ can lie in the continuum. Following [159] (sec. 3.3) for cold bosons where N_j atoms populate the state $|j\rangle$, the cross section can be divided into a coherent

part and an incoherent part,^{†††}

$$\frac{d\sigma}{d\Omega} \sim \left| \sum_f N_i \langle i | e^{i\Delta\mathbf{k}\cdot\mathbf{r}} | i \rangle \right|^2 + \sum_{f \neq i} N_i |\langle i | e^{i\Delta\mathbf{k}\cdot\mathbf{r}} | f \rangle|^2. \quad (1.31)$$

The matrix element connects the state with initial momentum \mathbf{k}_i to the one with the final momentum \mathbf{k}_f , where $\Delta\mathbf{k} = \mathbf{k}_f - \mathbf{k}_i$ (see [222] and sec. 1.7 of [233]). The term on the left represents coherent scattering where the atom remains in the same state. This term gives rise to the dispersive signal that can be recorded e.g. by Faraday imaging. As the momentum transfer to the atom is small, this light is forward scattered into small angles [12]. The term on the right describes traditional Rayleigh scattering which is incoherent, and the light is isotropically scattered into space. The color of scattered photons is the same as that of the incoming ones. In the process the atomic sample is heated by the recoil imparted by the photons. The photons that are scattered into large angles give rise to the missing light in absorption imaging, and in fluorescence detection these are in fact the photons collected by the imaging system.

This picture is valid for low intensities of light where higher order effects do not play a role. By turning up the intensity higher order transitions can occur and the fluorescence spectrum exhibits the famous Mollow triplet which was very recently observed (unambiguously) in cold atoms for the first time [209]. This effect usually plays no role in typical cold-atom experiments.

In conclusion, all atom-light interactions have a fundamental destructive nature and the mechanism that causes the heating is the same one that provides us with the blue sky, at least when it doesn't rain.

^{†††}In ref. [222], the author discusses a third term due to bosonic stimulation which is relevant for the BEC. I leave that out here as the BEC as such plays no role in this discussion.

CHAPTER 2

Cooling and imaging dilute and degenerate atomic gases

The invention of the laser in 1960 [186] is doubtless the single most important technological step that enabled research of cold-atom physics. The narrow distribution in the frequency of the light emanating from a laser eased the addressing of individual atomic transitions. In the mid-1980s the field gained momentum as the basic machinery like the *magneto-optical trap* (MOT) was developed [228]. Nowadays nobody builds a cold-atom experiment without a MOT. It is the workhorse of the field, a device capable of bringing atoms from room temperature down to tens of μK .

In this chapter a brief account is given of the basic physics of Doppler and sub-Doppler cooling in section one and of the MOT in section two. That is followed by a discussion of the methods used for forced evaporative cooling in the third section. Section four covers the basic physics of Bose-Einstein condensation, and in the fifth section absorption imaging is explained.

2.1 Laser cooling

A two-level atom is free to move in one dimension between counter-propagating laser beams. From the reference frame of the atom, the frequencies of the lasers will depend on the velocity of the atom itself due to the Doppler effect. If the atom is moving at a velocity v , the frequencies of the laser beams with wave vectors k are shifted by $\pm kv$ for the atom moving towards (+) and away from (−) the laser beam. The atom can be made to preferentially scatter light from the laser it moves toward, by adjusting the rest-frame frequency of both lasers red of the atomic transition. The scattering rate Γ is given by eq. (1.9). An atom-light scattering event changes the momentum of the atom by $p = \hbar k$, and thus the light exerts a force on the atom $F = \frac{dp}{dt} = \hbar k \Gamma$. The net force on the atom due to the light from the two counter-propagating beams is

$$F_D = \hbar k (\Gamma(\omega_l + kv) - \Gamma(\omega_l - kv)) \simeq -\beta v, \quad (2.1)$$

where β is an effective damping coefficient.* In the approximation terms on the order of $(kv/\Gamma)^4$ have been ignored. As a result the lasers cause a velocity dependent restoring force on the atom. This is the effect of *optical molasses*.

The random nature of spontaneous emission is also responsible for a heating mechanism, as the scattered photons are emitted into a random direction, causing a diffusive motion of the atoms. The atom will on average gain two units of the *recoil energy* E_r per scattering event, where

$$E_r = \frac{\hbar^2 k^2}{2m}. \quad (2.2)$$

By balancing the cooling rate due to the optical molasses and the recoil heating rate†, the atoms will reach the *Doppler temperature* in steady-state

$$T_D = \frac{\hbar \Gamma}{2k_B}. \quad (2.3)$$

*The exact value of the β is irrelevant for this discussion. For its value see sec. 7.2 of ref. [191].

†The cooling rate is $(\frac{dE}{dt})_c = \frac{d}{dt} \frac{1}{2} m v^2 = F_D v$, and the heating rate is $(\frac{dE}{dt})_h = 4 E_r \Gamma$.

For ^{87}Rb , $T_D = 146 \mu\text{K}$ [252]. For further discussion see [191], ch. 3 and 7.

When the ground state energy levels are degenerate in the magnetic quantum number m_F , an atom moving in laser light where a polarisation gradient is present will experience a friction force. Early experiments with atoms in optical molasses reported much lower temperatures than expected from simple Doppler cooling [181]. The cooling mechanisms responsible are said to achieve *sub-Doppler cooling*. The reasons for how such a force arises for two counter-propagating beams with circular polarisations, traditionally named the $\sigma_+ - \sigma_-$ configuration, are discussed below. The interference of the polarisations of the two beams result in a linear polarisation at every point along the beam. This linear polarisation rotates around the propagation direction, forming a helix.

Sub-Doppler cooling requires a minimal angular momentum of 1 in the ground state.[‡] In ^{87}Rb we can envision a $F = 2 \rightarrow F' = 3$ level scheme, with a 5-fold degenerate ground state, and a 7-fold degenerate excited state. If the atom is at rest the linear polarisation will drive a π -transition, and due to Clebsch-Gordan coefficients[§] the population will be symmetrically distributed in the ground state manifold with the greatest portion in the $m_F = 0$ state.

As the atom begins to move with a velocity v , the local orientation of the linear polarisation is changed, and this changes the quantisation axis to which the spin system at rest is referenced to. The movement induces optical pumping between the sublevels, but as shown in ref. [63] the steady state population of the ground state is not symmetric as the atom starts moving. If the atom moves in the positive (negative) direction in a red-detuned laser field, the $m_F < 0$ ($m_F > 0$) sublevels will be more populated than $m_F > 0$ ($m_F < 0$), and the difference in the population of the states is proportional to the quantity $\frac{kv}{U_{AC}}$, where U_{AC} is the AC Stark shift due to the molasses beams themselves.[¶] The

[‡]Our experiment operates cooling on the ^{87}Rb , D2 line, so the total angular momentum of the ground state is $F = 2$.

[§]The Clebsch-Gordan coefficients are numbers that relate different electronic states of the total angular momentum, see. [233].

[¶]Equation (1.13) is still applicable, but here it would be fine to make the RWA as we only have a contribution from the $F = 2 \rightarrow F' = 3$ transition.

Clebsch-Gordan coefficients favour scattering of the σ_- (σ_+) beam it is moving towards, and therefore the atom is cooled.[‡]

In a steady state the equilibrium temperature in the $\sigma_+ - \sigma_-$ configuration is

$$T_\sigma = \frac{\hbar s_0 \Gamma_s^2}{2k_B \Delta} \left(\frac{29}{300} + \frac{254}{300} \frac{\Gamma_s^2}{4\Delta^2 + \Gamma_s^2} \right), \quad (2.4)$$

where s_0 is the saturation parameter, Δ is the detuning between the laser light and the atomic resonance and Γ_s is the rate of spontaneous emission, see eq. (1.8). For typical values of Δ and s_0 the cloud temperatures are $\sim 10 \mu\text{K}$. For details the reader is referred to ref. [63].

2.2 Magneto-optical trapping

Atoms in optical molasses will eventually diffuse out of the laser beams (even though they are confined in 3D), as there is no spatially dependent variation in the scattering force. One way to achieve trapping is to add an inhomogeneous magnetic field to the light field of the molasses. In a magnetic quadrupole trap made of a pair of coils in an anti-Helmholtz configuration, a magnetic field gradient is created between the coils. In the centre the magnitude of the field is zero and increases outwards in any direction. This spatially varying field creates a spatially dependent Zeeman shift of the magnetic sublevels, according to eq. (1.2). In this situation the scattering force does not only depend on the velocity of the atoms in the molasses, but also on their position. As they move away from the centre of the trap, the magnetic field shifts the energy levels closer to resonance and the atoms will scatter more light that pushes them back to the centre. This is the *magneto-optical trap* (MOT) first realised in 1987 [228].

[‡]This cooling mechanism is fundamentally different from the Sisyphus mechanism that cools atoms in the $\text{lin} - \perp - \text{lin}$ polarisation gradient. The main difference is that the AC Stark shift varies periodically in that case, whereas in the $\sigma_+ - \sigma_-$ case it is constant along the beam.

The workings of a MOT are indeed more complicated and depend on a variety of parameters. But it is good to keep in mind that the temperature distribution of atoms in a MOT is typically quite broad, extending over 10's of μK . A 3D MOT with large atom numbers like the one in our laboratory has two different temperature regimes. The inner regime around the centre of the trap is governed by the sub-Doppler mechanism explained above. The outer region takes over where the Zeeman shift becomes too great to allow for sub-Doppler cooling, as that mechanism relies on the degeneracy of the magnetic sublevels. This happens around the radius where the Zeeman shift equals the AC Stark shift caused by the molasses light. For a detailed analysis of the different phases of the MOT see ref. [272].

2.3 Forced evaporative cooling

There are other ways to cool atom clouds than by laser cooling. A thermal cloud of atoms is described by a Maxwell-Boltzmann velocity distribution (MBD) (see ch. 14 of ref. [162]). Forced evaporative cooling removes the hottest atoms from an ensemble, and this corresponds to cutting off the exponential tail of the MBD, allowing for a subsequent re-thermalisation of the atoms in the cloud. This yields an atom cloud with a lower average temperature, so the cloud is cooled. To ensure good thermalisation after the removal of hot atoms the rate of elastic two-body collisions has to be high enough. The rate of such collisions is $\gamma_2 \sim n$, where n is the density of the cloud.** Therefore it is bad for the efficiency of the evaporation to remove too many atoms at a time as the density will drop and so will the collision rate for the necessary thermalisation process.

In a standard experimental sequence (see sec. 3.1) we use two types of *forced evaporative cooling*. The first method is *microwave cooling* carried out in a quadru-

**One can imagine being a lonely melancholic atom thinking about the whereabouts of fellow atoms: "I wonder the chance of meeting another atom in this dilute vapour." It pauses—and then the epiphany: "It must be proportional to the density of the vapour!" For a proper treatment see sec. 4.6 of [219].

pole magnetic trap. To be susceptible to magnetic trapping the atom must be in a *weak-field seeking state*. For an atom to be in a weak-field seeking state, the Zeeman splitting of eq. (1.2) must increase the energy of the particular m_F state. In the ^{87}Rb ground state this is true for the $|F = 2, m_F = (1, 2)\rangle$ and $|1, -1\rangle$ hyperfine sublevels. The trapping potential is given by (see sec. 4.1 of ref. [219])

$$B(x, y, z) = \frac{m_F g_F \mu_B}{4} \frac{dB}{dz} \sqrt{x^2 + y^2 + (2z)^2} \quad (2.5)$$

where $\frac{dB}{dz}$ is the magnetic field gradient along the vertical direction (through the center of the coils). This generates a V-shaped trapping potential for our atoms which are in the $|2, 2\rangle$ state. By bathing the atom cloud in microwave radiation of the right frequency, the atoms with the highest velocities (that are able to reach the outskirts of the trap where the Zeeman splitting is the greatest) are transferred to the $|1, 1\rangle$ state, which is a *strong-field seeking state*. As a result they are expelled from the trap.^{††} As this transition requires the change of the nuclear spin i it is not an electric dipole transition but is coupled through the magnetic dipole moment. The microwave frequencies are typically swept from about 6.9 GHz towards the $F = 1 \rightarrow 2$ splitting of 6.83 GHz (see appendix A). As the clouds are rather dilute at this stage, the re-thermalisation process takes a long time so this is a slow method,^{‡‡} reaching temperatures of $\sim 1 \mu\text{K}$. For details see sec. 4.6 of ref. [219].

The second method is evaporative cooling in an optical potential (see sec. 1.3) which we perform in a crossed optical dipole trap (CDT). In this case the atoms are trapped regardless of their m_F state. As in the previous case, the hottest atoms reside high up in the potential, and are removed simply by lowering its brim. This is achieved by turning down the intensities of the laser beams providing the trap light. Optical evaporation is also different from the microwave evaporation in the respect that the trap shape changes during evaporation, whereas

^{††}The same effect could also be achieved by internal transitions in the $F = 2$ manifold, with RF frequencies exactly matching the Zeeman splitting.

^{‡‡}The whole process takes about 15 s in our experiment.

$\frac{dB}{dz}$ is held at a fixed value in method explained above. A detailed analysis of how optical evaporation is optimally achieved is presented in ref. [212]. Inspired by the central result presented in the article we vary the intensity of our dipole lasers like

$$I(t) = I_0 \left(1 + \frac{t}{\tau}\right)^{-\beta} \quad (2.6)$$

where I_0 is the initial value of the laser intensity and τ and β are constants that we determine experimentally. This method allows us to reach Bose-Einstein condensation at a temperature ~ 100 nK. Figure 2.1 presents a series of absorption images taken in our experiment where a BEC arises from a thermal cloud as the cloud is cooled by evaporative cooling. The implementation of this method of cooling is discussed in greater detail in ch. 3 of [131].

2.4 Bose-Einstein condensation

A very special state of matter appears as dilute vapours of bosonic atoms are cooled to temperatures well below the μK regime. This is the *Bose-Einstein condensate* (BEC), a macroscopic quantum-wave made of matter.^{§§} To get an intuitive idea about how this state of matter forms, one should consider the individual atoms as small wavepackets with a wavelength equal to the *thermal de Broglie wavelength*,

$$\lambda_T = \sqrt{\frac{2\pi\hbar^2}{mk_B T}}, \quad (2.7)$$

for particles at a temperature T and of mass m . As T decreases the density of the cloud increases and the inter-particle distance $n^{-1/3}$ shrinks. When λ_T and $n^{-1/3}$ are comparable in size ($\lambda_T n^{1/3} \simeq 1$) the individual matter-waves begin to synchronise and form one big matter-wave—the BEC. The quantity $\lambda_T n^{1/3}$ is essentially the *phase space density*

^{§§}The theoretical parts in this section are based on ref. [219], chs. 1, 2, 5 and 6, respectively. The reader is directed to that reference for more details.




Figure 2.1: A Bose-Einstein condensate emerges from the thermal cloud of cold atoms. From top left to bottom right the end point of the evaporation in the crossed optical dipole trap is decreased. This particular measurement was done with a microtrap superimposed to the crossed dipole trap (see sec. 4.2). The colour scale is the same for all frames.

$$\text{PSD} = n \lambda_T^3 = n \left(\frac{2\pi\hbar^2}{mk_B T} \right)^{\frac{3}{2}}, \quad (2.8)$$

which is a useful quantity to evaluate the efficiency of a cooling process, where typically one desires as many atoms as possible at as low temperatures as possible.

Bosons are particles with integer spin and due to the spin-statistics theorem [217], they obey Bose-statistics. The Bose-distribution function applies for non-interacting bosons at a thermal equilibrium and describes the probability of finding a particle in the state ν , given its energy E_ν

$$f(E_\nu) = \frac{1}{e^{(E_\nu - \mu)/k_B T} - 1}. \quad (2.9)$$

Here the *chemical potential* μ is a key quantity. For warm clouds the occupation of the low energy states is small, and the probability distribution resembles that of a MDB, and as $T \rightarrow \infty$, $\mu \rightarrow -\infty$. When T is lowered, μ becomes less negative and the occupation the low lying energy states, starts to rise. The form of the Bose distribution enforces the condition that $\mu \leq E_0$. The Bose-condensation of the atom cloud happens at the critical temperature T_c , and as $T \rightarrow T_c$, $\mu \rightarrow E_0$, and a macroscopic fraction of the particles in the system occupies the lowest lying energy state. As we see from eq. (2.9), the probability of occupation will rise significantly as $\mu \rightarrow E_0$.

The critical temperature for the onset of the BEC trapped in a 3D harmonic potential is

$$k_B T_c \simeq 0.94 \hbar \bar{\omega} N^{1/3} \quad (2.10)$$

where $\bar{\omega}$ is the geometric mean of the trapping frequencies along the three spatial coordinates, and N is the total number of particles in the system. The ratio of the number of condensed atoms to the total number of atoms is

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c} \right)^3. \quad (2.11)$$

This formula indicates that the onset of the phase transition from a thermal cloud to a BEC is sharp. Precision measurements of the phase transition [173], reveal that this is not true, as a Bose gas near T_c is not non-interacting as is assumed in the case of the Bose-distribution. Repulsive interactions play an important role. The effect of the interactions is both to smooth out the transition and slow down the rise of the condensate fraction below T_c .

The dominating interaction between two cold alkali atoms in a dilute cloud is the electric dipole-dipole interaction, also known as the van der Waals interaction. The interaction varies as $\sim r^{-6}$, where r is the interatomic separation. A basic quantum mechanical treatment of scattering theory reveals that when two cold atoms scatter, the scattered wave is essentially spherical. This is so-called s-wave scattering. That simplifies the process a great deal and enables

us to quantify the scattering process with a single quantity called *the scattering length*, denoted by a . Inter-particle separations in cold dilute gases are large, especially when they are compared with typical distances at which two atoms interact. Due to this, and the simple form of the atom-atom scattering, it is well justified to describe the interaction between two particles at coordinates \mathbf{r} and \mathbf{r}' with a contact potential

$$U(\mathbf{r}, \mathbf{r}') = \frac{4\pi\hbar^2 a}{m} \delta(\mathbf{r} - \mathbf{r}'). \quad (2.12)$$

We call the front factor U_0 for convenience.

To describe the condensed state we assume that all the different atoms in the cloud occupy the same state $|\phi(\mathbf{r})\rangle$ and the total wavefunction of the system is the product state of all the single particle states, $|\Psi\rangle = \prod_i |\phi(\mathbf{r}_i)\rangle$, where \mathbf{r}_i is the position vector of atom i . The Hamiltonian of an interacting gas of N particles is

$$\mathcal{H} = \sum_{i=1}^N \left[\frac{\mathbf{p}_i^2}{2m} + V(\mathbf{r}_i) \right] + U_0 \sum_{i<j} \delta(\mathbf{r}_i - \mathbf{r}_j). \quad (2.13)$$

The energy of the gas is simply the expectation value of the Hamiltonian, $E = \langle \Psi | \mathcal{H} | \Psi \rangle$. By introducing the concept of the condensate wavefunction $|\psi\rangle$ that satisfies the normalisation condition $\int |\psi(\mathbf{r})|^2 d\mathbf{r} = N$ and by applying variational calculus to minimise the energy E , one will arrive at the Gross-Pitaevskii equation (GPE)

$$\left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + U_0 |\psi(\mathbf{r})|^2 \right] |\psi(\mathbf{r})\rangle = \mu |\psi(\mathbf{r})\rangle. \quad (2.14)$$

This is the time-independent version of the GPE, and we see that the chemical potential μ has popped up again. The GPE resembles the time-independent Schrödinger equation with a non-linear term $U_0 |\psi(\mathbf{r})|^2$. The GPE is a mean-field theory, and the non-linear term accounts for the mean interactions of all the bosonic atoms in the ensemble.

For large BECs the interaction term dominates the kinetic term, which can as a result, be neglected. This is the so-called Thomas-Fermi approximation. By dropping the term out of eq. (2.14), one can solve for the atomic density $n(\mathbf{r}) = |\psi(\mathbf{r})|^2$ to yield

$$n(\mathbf{r}) = \frac{\mu - V(\mathbf{r})}{U_0}, \quad (2.15)$$

showing that BECs take the shape of their trapping potential. By assuming a harmonic trap $V(x, y, z) = \frac{1}{2}m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2)$, the radius of the condensate R_i (along dimension i) is determined by the condition $n(\mathbf{r}) = 0$, so

$$R_i = \sqrt{\frac{2\mu}{m\omega_i^2}}. \quad (2.16)$$

The chemical potential is also uniquely determined and has the form

$$\mu = \frac{15^{2/5}}{2} \left(\frac{Na}{\bar{a}} \right)^{2/5} \hbar\bar{\omega} \quad (2.17)$$

where $\bar{a} = \sqrt{\frac{\hbar}{m\bar{\omega}}}$ is the characteristic length scale in a harmonic trap.

2.5 Absorption imaging and time-of-flight measurements

It is safe to say that resonant absorption imaging is the most widely used technique for measuring cold atom clouds. The method is precise and can also be made accurate, although that requires a great deal of experimental work [173]. Resonant absorption imaging is preferable for dilute gases with ODs on the order of 1 [159]. The signals have high contrast but due to resonant scattering the atom cloud is heated and destroyed, so this method allows for only one good image of the cloud. At higher OD's it can be difficult to obtain the correct column densities, but methods exist to extract correct atom numbers for OD's at least up

to 10 [230]. Ultimately one is limited by the dynamic range of the camera [229]. In comparison, dispersive imaging techniques (see secs. 1.4 and 4.3) work well for samples with high OD, but their central advantage lies in the ability to take multiple images of the cloud in a single realisation of the experiment, enabling single-shot measurements of cloud dynamics. All the different dispersive techniques have the same SNR for small phase shifts [104]. Minimally-destructive absorption techniques also exist where a portion of the cloud is transferred to a different internal state and subsequently imaged absorptively [229]. This technique has a similar SNR to the dispersive techniques.

To understand how atom numbers can be acquired through light absorption, let us assume that a laser field of intensity I propagating along the z direction, is incident on an atom cloud. Its intensity will now be reduced at the rate of $\frac{dI}{dz}$ which equals the product of the incoming light intensity I , the density of scatterers $n(x, y, z)$, and the cross section of the light-matter interaction σ , of eq. (1.11). This results in Beer's law

$$\frac{dI}{dz} = -\sigma I n(x, y, z). \quad (2.18)$$

The atom cloud is illuminated and images of the laser beam with ($I_a(x, y)$) and without ($I_b(x, y)$) the atomic cloud present are recorded.[¶] The density is integrated along the line of sight and the intensity correspondingly drops from I_b to I_a , so one obtains the expression

$$\ln \left(\frac{I_a}{I_b} \right) + \frac{I_a - I_b}{I_{\text{sat}}} + \left(\frac{2\Delta}{\Gamma} \right)^2 (I_a - I_b) = -\sigma_0 \int_0^z n(x, y, z) dz. \quad (2.19)$$

The term on the right hand (apart from the minus sign) is the so-called *optical density*,

$$\text{OD}(x, y) = \sigma_0 \int_0^z n(x, y, z) dz \quad (2.20)$$

[¶]A third image $I_c(x, y)$ with no light present is also taken. That measures the technical noise level on the camera. When processing the data, I_c is subtracted from both I_a and I_b .

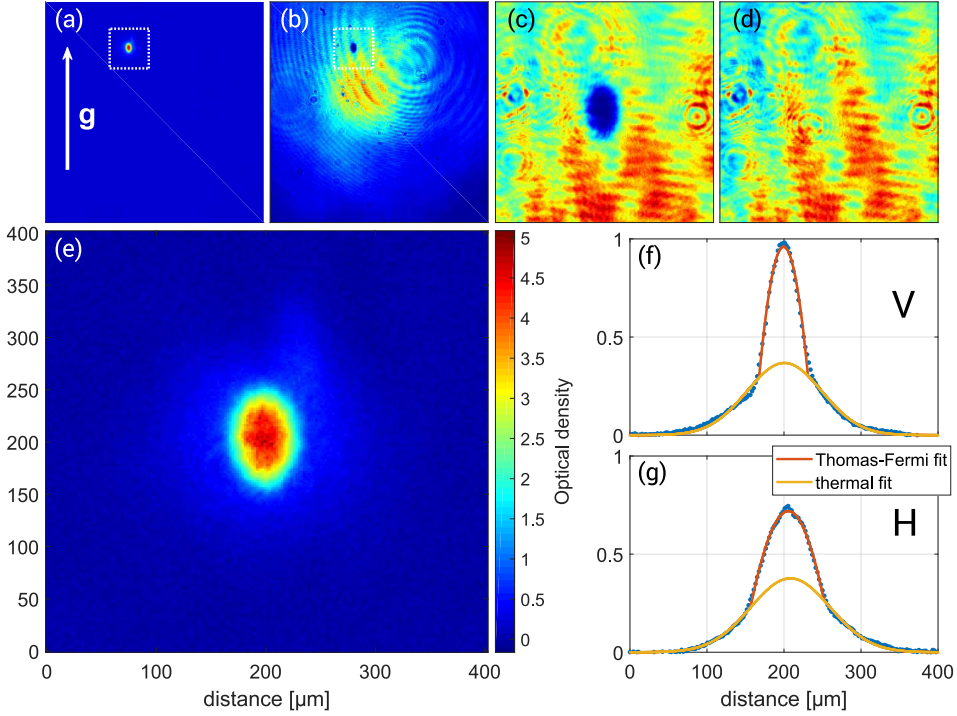


Figure 2.2: Reconstructing the atom density with absorption imaging. (a) The OD of the full field of view of the camera. The arrow indicates the direction of gravity. (b) The laser field with the atoms present, I_a . The white square demarcates the enlarged area shown in figure (c) where the absorption due to the atom cloud is apparent. (d) shows a reference image of the light field I_b . (e) is the enlarged image of the atom cloud, where the colour of the image is determined by its OD. The BEC is the high peak in the centre. (f) and (g) show the mean OD per pixel row, along the vertical (V) and horizontal (H) directions. The data are fit with a bimodal distribution.

which is routinely used in the field, as it comes out naturally from an absorption measurement.

Figure 2.2 (e) shows an absorption image (and its decomposition in frames (a)–(d)) of a BEC created with evaporation in the CDT, on top of a background of thermal atoms. The cloud was released from the trap and allowed to fall freely

under gravity for 15 ms, before the first image was taken. The optical density is obtained by putting the images I_a and I_b into eq. (2.19).^{***} Here $\Delta = 0$ so the third term simply drops out. The graphs in (f) and (g) represent the average OD along the vertical and horizontal directions. The data are fit by a function consisting of an inverse parabola (representing the Thomas-Fermi approximation for the density of the BEC in a harmonic potential) and a Gaussian (representing a cloud expanded according to a Maxwell-Boltzmann velocity distribution).^{†††}

The total number of atoms can be evaluated by summing the OD pixel by pixel,^{‡‡‡} giving

$$N_{\text{sum}} = \frac{1}{\sigma_0} \sum_{x,y} \text{OD}(x,y). \quad (2.22)$$

The value of σ_0 depends on the polarisation state of the light used for imaging (for details see ref. [252]).

Another important property that we frequently extract from absorption imaging is the temperature. By turning off the trapping potential the cloud will start to fall under the effect of gravity. As the confining potential is no longer present, the cloud will expand according to its velocity distribution. The in-trap momentum distribution of the cloud becomes clear as it falls. This is a time-of-flight (TOF) measurement, that effectively yields the Fourier transformation

^{***}Care has to be taken as the raw images are in units of bit/pixel, which indeed is a measure of the intensity of the incoming light. This is not an issue for the first term in eq. (2.19) but in the second term, I_{sat} has to be converted to the same unit. For that, a measurement of how much energy a bit count in the camera E_{bit} corresponds to is needed, as well as the area of one pixel in the image plane, A_{pix} . The saturation intensity in units of bits/pixel is then

$$I_{\text{bit}} = I_{\text{sat}} \frac{A_{\text{pix}}}{E_{\text{bit}}/\Delta t_{\text{img}}} \quad (2.21)$$

where Δt_{img} , is the pulse length of the imaging light.

^{†††}At these low temperatures, the cloud distributions are better described by the Bose function (see sec. 2.3 of ref. [219]) rather than a Gaussian. In our experiment we do not require high accuracy for those numbers, so this has been the convention.

^{‡‡‡}Remember to scale the σ_0 to the dimension of a pixel!




Figure 2.3: Extracting quantities by scanning the TOF. (a) The temperature of a thermal cloud expands linearly with the TOF. (b) A calibration measurement of the corresponding size of a pixel in the image plane.

of the in-situ momentum distribution. Assuming a Gaussian-shaped thermal cloud, the 1σ -width of the cloud along a direction i is given as

$$r_{\sigma,i} = \sqrt{\frac{k_B T}{m} \left(t_{\text{TOF}}^2 + \frac{1}{\omega_i^2} \right)} \quad (2.23)$$

where ω_i is the trap frequency along the i -th direction (see sec 4.5 of [159]). For a long t_{TOF} , the width $r_{\sigma,i}$ becomes linear in that variable, and grows independent of the initial trapping frequencies. An example of such a measurement is shown in fig. 2.3 (a). Here a cold thermal cloud of atoms was held in an optical trap about $80 \mu\text{K}$ deep and dropped into a TOF measurement. The TOF is scanned (changed stepwise in multiple realisations of the experiment), and the cloud expands as the TOF is increased. By fitting eq. (2.23) one obtains a temperature of $T = 421(8) \text{ nK}$.

Figure 2.3 (b) shows how to extract the corresponding size of an individual pixel in the image plane. This is a necessary calibration measurement for correct determination of both atom numbers and temperatures. The measurement procedure is identical to the one for obtaining the temperature, but here




Figure 2.4: An overnight stability test of the BEC creation of the machine. (a) The extracted BEC atom number as a function of the run of the experiment. (b) The BEC radius plotted against the BEC atom number. (c) The condensate fraction is plotted with the temperature of the thermal component. The critical temperature of the BEC is around 170 nK, as indicated by the rise of the condensate fraction.

one uses the fact that the distance Δs the cloud has fallen follows from elementary kinematics $\Delta s = \frac{1}{2}gt^2$, where g is the acceleration due to gravity. By a quadratic equation with one free parameter, the size of a camera pixel is mapped to the image plane.

To test the stability of our BEC machine, we occasionally run it overnight and observe the fluctuations in the fitted BEC number and the thermal part of the cloud. Figure 2.4 shows an example of such a measurement. In the (a) panel we see that when running the experiment continuously, it exhibits quite some fluctuations, especially in the BEC numbers. The fluctuations in the thermal part are small. For reference the mean atom number and 1σ fluctuations in the region between run 500 and 1000 are $N_{\text{th}} = 85(2) \times 10^3$ and $N_{\text{BEC}} = 76(9) \times 10^3$. There could be many reasons for these fluctuations like thermal drifts (although at this point the experiment had active temperature stabilisation described in sec. 6.5), drifts in the frequencies of the master cooling laser etc. We suspect fluctuations in the laser intensity controller of the dipole laser to be the culprit. Until now these fluctuations have not been a big issue for our experiments, but this is why some experiments choose to actively stabilise their atom numbers [105] or devise methods for post-selection of data [24]. How-

ever, as the BEC sizes fluctuate they can be used to (weakly) test some of the BEC theory presented in sec. 2.4. In (b) the BEC radius extracted from the fit is plotted as a function of the BEC atom number. According to eqns. (2.16) and (2.17) we expect that $R \sim N^{\frac{1}{5}}$. The data are fit by such a function, and the agreement is excellent given such a rough treatment. The final image (c) shows how the BEC arises as the temperature sinks below the critical temperature T_c .

CHAPTER 3

The experimental machinery for cold-atom magnetometry

Sections 3.4 and 3.5 are based on my progress report [83].

The experiment has seen many changes during my time in the Hires lab. It is not straightforward to simply list all the different parts of the machine necessary for the context of the thesis at once. The content of the whole chapter, apart from that of sections four and five, which are exclusively for the magnetometry experiments, is general for *all* the experimental work described in the thesis.

In the first section the typical experimental sequence carried out to produce a thermal cloud of atoms in the science chamber is described. The second section briefly discusses the basis of microscopy as that is important for all our experiments. In the third section an account is given of the two main technologies used for the spatial shaping of laser beams: the acousto-optical deflector and the digital micromirror device. In the fourth and fifth sections, technical details of the microtrap system and the spatially resolved Faraday polarimeter setup are




Figure 3.1: A schematic setup of the vacuum components of the experiment. Details are given in the text.

discussed. The sixth section discusses general noise properties of the cameras we use both for the Faraday measurements and the high-resolution experiment, and finally, the seventh section is devoted to the control system used in the lab, which is a home-built solution written in LabView.

The contents of this chapter rely on my knowledge of the experiment and on earlier reports of it, as detailed in chs. 2 and 3 of both [199] and [131].

3.1 The cold atom machine

The components of the vacuum setup: chambers, pumps and valves are all depicted in fig. 3.1. The experiment consists of four vacuum chambers. Our *Magneto-Optical Trap* (MOT) section, consists of a glass cell (HELLMA OPTICS) where we make a 2D MOT, which is attached to an octagonal chamber (KIMBALL PHYSICS) where a 3D MOT is realised. At a right angle to the MOT configuration a couple of tubes connect the 3D MOT chamber to a smaller vacuum

chamber cubic in size, called the *cube chamber* (KIMBALL PHYSICS). A specially made connection piece attaches the *science chamber* (KIMBALL PHYSICS) to the cube chamber. All the different chambers, viewports, tubes and flanges are made of 316LN steel, which is a non-magnetisable alloy, with some exceptions (the viewports in the MOT section and the gate valve connecting the 3D MOT and the cube). The science chamber is made out of titanium, for even better non-magnetic performance.

There are three vacuum pumps depicted in fig. 3.1. The ion getter pump (VARIAN, VACION PLUS) maintains the 3D MOT in the *ultra-high vacuum* (UHV) regime, at about $2 \cdot 10^{-10}$ mbar. A differential pumping stage between the 2D and 3D MOT chambers keeps a pressure difference of about two orders of magnitude. To reach *extremely-high vacuum* (XHV) an ion pump with a NEG coated getter material (SAES, NEXTORR D200) is situated in an angled tube between the cube chamber and the science chamber. This pump brought the science part of the experiment to the mid 10^{-11} mbar.* During a bakeout of the experiment in the fall 2016, we added another SAES pump (SAES, NEXTORR D500) to the science chamber side of the experiment, along with the necessary tube pieces and a valve (as a link to a turbomolecular pump to use during bakeout). The D500 brought us safely into the XHV regime with pressures $< 10^{-11}$ mbar. During the early days of the experiment, our team found out that the SAES NexTorr pumps don't work well with the titanium sublimation pump that was also built in. Therefore it was never used.

Let us link these different parts of the experiment together by introducing a typical experimental cycle. Every sequence starts off in the 2D MOT cell. Atoms are collected from a background gas of ^{87}Rb that evaporate from Rb dispensers (SAES) inside the vacuum system when current is run through them. The atoms

*It is a bit hard to tell exactly what is the level of pressure, as there is no dedicated pressure gauge on the setup. The display monitor for the NexTorr D200 pump, always flickered some low numbers in the nA regime, but in private communication with technical staff at SAES we were told rather to trust the reading of a dedicated pressure gauge than the gauge of the pump. Later on, our experience was that the SAES pump gauges read typically slightly higher values than a proper ion gauge.

collected in the trap are simultaneously pushed by a dedicated *push beam*, via light pressure through a differential pumping stage (a nozzle with an opening of 2 mm) into the 3D MOT chamber, where they are collected in a six-beam 3D MOT. The 3D MOT works at a detuning of $\Delta_{3D} \simeq 4.5 \Gamma_{D2}$, and a total optical power of about 200 mW is distributed such that there is twice as much power in the two beam pairs propagating in the horizontal plane compared to the beam pair that propagates vertically. This stage takes in total 5 s. As discussed in sec. 2.2 the temperature of the entire cloud varies from $\sim 10 \mu\text{K}$ in the centre to somewhere above the Doppler limit (at $146 \mu\text{K}$) at the outskirts of the MOT.

Subsequently the quadrupole magnetic field of the 3D MOT is turned off and a pure optical molasses phase is initiated. This phase lasts for 8 ms in order to reach sub-Doppler temperatures in the entire cloud. During that phase the detuning is increased from $7.1 \Gamma_{D2} \rightarrow 13.5 \Gamma_{D2}$, in a linear fashion. As a result of the Doppler cooling cycle, all of the atoms in the ensemble are now in the $F = 2$ state of the $5^2S_{1/2}$ manifold of the ^{87}Rb ground state (see level scheme in appendix A). The magnetic quantum number m_F is on the other hand not well defined for the whole cloud. As we need to magnetically trap the cloud later on, it must be pumped to a magnetically trappable state (see discussion in sec. 2.3). By applying a magnetic field along the vertical axis of the setup in the 3D MOT chamber, a quantization axis is now chosen for the quantum number F , and a circularly polarised beam (with σ_+ polarisation, see sec. 2.2 of ref. [95]), also propagating vertically, serves as a pump for the magnetic state for a period of 1.5 ms. As a result, the atoms are pumped into the maximally stretched state, $|F = 2, m_F = 2\rangle$.

The quadrupole MOT coils are mounted to a rail (PARKER, DAEDAL 401XR 600), such that they can be physically moved. Due to this they are typically referred to as the *transport coils*. After optical pumping the current in the transport coils is ramped up in 150 ms, to produce a magnetic field gradient of $\nabla B_{\text{trans}} = 150 \text{ G/cm}$. Within the next 2 s, the coils are moved over to the cube chamber. The cube is equipped with a pair of high-current quadrupole coils, made of hollow wire for efficient water cooling, shown in dark orange (coils) and light blue

(epoxy mount) colours in fig. 3.1. The atoms are loaded from the magnetic trap of the transport coils into the trap generated by the cube coils by relaxing one and ramping up the other to a gradient of about $\nabla B_{\text{cube}} = 130 \text{ G/cm}$. After loading into the cube trap, a loop antenna mounted inside the chamber delivers microwaves in three distinct linear ramps, ranging in frequency from about 6.92 GHz to about 6.85 GHz. This part of the sequence is the bottleneck when it comes to time efficiency, as it takes in total 15 s for it to run through.

At this stage the atom cloud counts about half a billion atoms, at a temperature $\sim 1 \mu\text{K}$, which is difficult to measure precisely with the imaging tools available. At this point a tightly focussed single-beam optical dipole trap propagating through the whole vacuum system along the negative longitudinal direction (see fig. 3.1), is ramped to a high power, giving a very deep trap. The light originates from a NUFERN NUAMP fibre amplifier operating at a wavelength of $\lambda_{\text{DT}} = 1064 \text{ nm}$. The *longitudinal dipole trap* (LDT) has a waist of $w_{\text{LDT}} \simeq 45 \mu\text{m}$, and at the loading power of 2.9 W the trap depth amounts to about 140 μK . The focus of the LDT can be shifted with the help of a remotely-controllable translation stage (THORLABS DDS220/M) to the science chamber. The position profile of the transport stage is determined by two parameters, a constant acceleration a_{ts} and a maximum velocity v_{ts} . The total transport distance is 228.4 mm and it takes 2.2 s using $a_{\text{ts}} = 100 \text{ mm/s}^2$ and $v_{\text{ts}} = 150 \text{ mm/s}$. As the cloud arrives to the science chamber, another dipole trap propagating transversal to the LDT, the *transverse dipole trap* (TDT), is ramped to a value of 6 W (70 μK), generating a *crossed dipole trap* (CDT).[†] The trap power of the LDT is subsequently ramped down to 1.75 W (80 μK), giving roughly equal trap depths in the two beams. The total depth is now around 140 μK , taking into the account the effect of the gravitational sag. Up until this point the cloud preparation takes around 25 s and the atom cloud contains 20–30 million atoms at a temperature of 10–20 μK . This can be considered to be the starting point of the all the different experiments described in this thesis, unless otherwise specified.

[†]A detailed description of this part of our system can be found in ch. 3 of [131].

3.2 The (very) basics of microscopy

Microscope objectives are central to both our implementation of magnetometry and the high-resolution setup. Although the quality of the optics and their light-gathering capacity is very different, the same basic physics applies.[‡]

The apparent size of an object depends on the angle it subtends in the field of view. By placing a convergent lens with a focal length f between the viewer and the object, the object can be brought closer to the viewer. If the object is placed within a distance of f from the lens, it will create an enlarged virtual image of the object that appears to be further away than the distance to the object itself. This is the entity seen through a looking glass, rather than the object itself.

The closer one can place the convergent lens to the object, the more light one can gather from it. The *numerical aperture* (NA) of an imaging system is precisely a measure of that, as portrayed in fig. 3.2 (a). It is defined for the most extreme ray that can enter a lens of a focal length f that has a free aperture d , at an angle of θ as

$$\text{NA} = n \sin(\theta) = n \sin \left(\arctan \left(\frac{d}{2f} \right) \right) \simeq \frac{nd}{2f} \quad (3.1)$$

where n is the refractive index of the material the ray travels in. In our case this material is air, so $n = 1$. Upon the assumption that θ is small, the approximation in the equation above is valid.

The numerical aperture of a microscope is also related to its resolving power, which is the level of detail that can be distinguished with the optic. If the system is free of optical aberrations, its performance is limited only by the diffraction of the light. To understand this in more detail, we must introduce the concept of the *point-spread function* (PSF). The PSF is an imaging system's response to a point source—or simply the image of the point source after propagation through the system (sec. 3.2. in ref. [124] covers the concept in details). It is

[‡]The section is written with the help of ref. [34].




Figure 3.2: Resolving an object by a lens. (a) The definition for numerical aperture of a lens of focal length f with a free aperture d . The NA is the sine of the angle θ . Source: Wikimedia Commons. (b) The Rayleigh resolution limit is the distance between the two vertical black dashed lines. See details in the text. (c) Fitting a Gaussian to an Airy disk.

convenient to represent it in coordinates (ρ, ζ) that relate to the cylindrical coordinates (r, z) by

$$\rho = \frac{2\pi}{\lambda} \text{NA } r, \quad (3.2)$$

$$\zeta = \frac{2\pi}{\lambda} \text{NA}^2 z, \quad (3.3)$$

where the image is formed at $z = 0$, and r is the radial coordinate in the image plane. The PSF is the function $I_{\text{PSF}}(\rho, \zeta)$. The radial intensity distribution of the PSF is identical to that of an Airy disk,

$$I_{\text{PSF}}(\rho, \zeta = 0) = I_0 \left(\frac{J_1(\rho)}{\rho} \right)^2, \quad (3.4)$$

where I_0 is the maximum intensity and $J_1(\rho)$ is the Bessel function of the first kind of order one (see sec. 5.2 of ref. [179]). The extent of the function axially is given by a sinc function[§]

$$I_{\text{PSF}}(\rho = 0, \zeta) = \frac{I_0}{\pi^2} \left(\frac{\sin(\zeta/4)}{\zeta/4} \right)^2. \quad (3.5)$$

[§]The extra factor of $\frac{1}{\pi^2}$ that appears in the axial part of the function is hidden in the $J_1(\rho)$.

Imagine that we have two such point sources side by side that we want to image and distinguish in the image plane ($\zeta = 0$). How closely can we space them? The Rayleigh criterion is the distance between the images of the two point sources where the first diffraction minima of one image lies in the maxima of the other, as illustrated in fig. 3.2 (b). That distance is equal to

$$r_{\min} = 0.61 \frac{\lambda}{\text{NA}}. \quad (3.6)$$

Other criteria for resolution limits exist, but in this text we plead to the Rayleigh criterion as its usage is widespread in our field of research.

To determine the resolution limit of an imaging system, one ideally acquires an image of a point source. In practice this can be approximated by an aperture that is smaller than what can be resolved by the system. It can be problematic to fit the Airy function directly, as the side lobes of the PSF are often not clearly visible due to residual aberrations still present in the imaging system. Thus it is quite common practice to fit a Gaussian function (eq. 1.15), see fig. 3.2 (c). The Rayleigh limit can be related to the fitted waist w_0 by the relation

$$r_{\min} = 1.434 w_0. \quad (3.7)$$

3.3 Shaping light

Light is central to the field of cold-atom physics. The laser provides means of trapping, cooling and manipulating cold atoms. Technical developments in the generation, handling and detection of light have in the past decades expanded the experimental horizon and brought new types of experiments within the reach of cold-atom labs around the world. One of those aspects, on which we put heavy emphasis in our laboratory, are some of the different ways of shaping laser light spatially. We rely on two distinct pieces of technology, *acousto-optical deflectors* (AODs) and *spatial-light modulators* (SLMs). Out of the different SLM technologies available we use *digital micromirror devices*. Generally speaking SLMs are more flexible than AODs, as they enable spatial control over

both the amplitude and the phase of a laser beam. Due to that they may be employed to correct for aberrations that arise in optical systems, as they give full control over the wavefront of the laser. However, their power efficiency is typically worse than that of AODs, which are also more easily implemented to optical systems.

The usage of acousto-optic devices is ubiquitous in modern optics labs as a means of controlling and regulating intensities of laser light as well as for fast switching of light fields, typically on the scale of 100's or even 10's of ns. In recent years they have also become popular for spatially shaping light beams, that are used to trap cold atoms. An AOD deflects a laser according to the frequency of an applied radio-wave. Often the AOD is operated with a single tone, only deflecting one beam at a time. This means that complex patterns of light have to be "painted". The first attempts to use AODs to make flexible trapping potentials for neutral atoms were reported almost 20 years ago [98]. Later they were employed to make simple static potentials for BECs to realise atom interferometers [245], however heating was an issue for making arbitrary trapping potentials. The first successful application for BECs was achieved in 2009 [133]. Since then this technology has been especially popular for making circular trapping potentials for BECs, an optimal platform for interferometric applications [26, 77, 188]. AODs have also been used to realise a cold-atom cloud collider, a flexible system with a high degree of spatial control for studies of scattering [231, 268]. Furthermore, they have been employed in a different context for a multiplexed quantum memory device for writing and reading out quantum states [227]. As discussed in the introduction, recent experiments relied on AOD controlled optical tweezers to generated arrays with unity atom filling [22, 87].

In the early 00's the power of SLMs for optical trapping was first properly demonstrated. Since they have been used to manipulate everything from microbeads to proteins to individual atoms [119]. The first successful realisation of an SLM based neutral atom trap dates from 2004 [29]. Here the traps were static, and intended for single atoms, but only a couple of years later a similar device was used to split and transport a BEC [36]. These first implementa-

tions relied on holographic beam shaping with a liquid crystal phase modulator (LCPM). An alternative method is to manipulate the amplitude directly by imaging the pattern on a DMD [42, 198]. This offers a very flexible way of creating arbitrary static potentials [108]. Recent reports also confirm that DMDs can be used in a dynamic mode to transport single atoms [256]. Another showcase of their usefulness, important for our particular application in the laboratory, is their combination with quantum gas microscopes to generate many-body states of ultra-low entropy [189].

The following subsections will briefly describe the technical aspects of the two technologies for light shaping.

Acousto-optical deflectors

The first realisation of acousto-optical deflection in 1932, was due to the experimental work of P. Debye and F. W. Sears [69]. A ray of light was passed through a container filled with liquid. A quartz crystal connected to an RF source, was immersed in the liquid, and by turning on the RF source, supersonic waves were excited in the liquid. The light ray was then observed to diffract into multiple orders due to the interaction with the sound waves in the liquid.

Present day AODs don't rely any more on liquids. Instead laser beams pass through transparent crystals and acoustic transducers create vibrational excitations (sound waves) that propagate in the crystal transversely to the laser beam, see fig. 3.3. A photon with a momentum k_p in the light ray passing through the crystal is diffracted by phonons, quantized vibrational excitations that have a momentum k_s . Momentum must be conserved so the diffracted photon momentum is ([65], sec. 19.11)

$$k_p' = k_p \pm k_s, \quad (3.8)$$

where the sign depends on if the photon-phonon scattering process increases or decreases the energy of the outgoing photon. The momentum conservation leads to a deflection in the angle of the outgoing photon with respect to the




Figure 3.3: A schematic figure of the function of an AOD. A laser beam enters the crystal (in light blue) from the left. As an RF source is coupled to a transducer it excites sound waves in the crystal propagating in the vertical direction in the image. The laser is diffracted into multiple orders by the density modulations generated by the sound waves.

angle of the incoming photon. The photon is diffracted into discrete diffraction orders as shown in fig. 3.3. The frequencies of the outgoing photons are also shifted discretely by integer multiples of the phonon frequency f_s in order to conserve energy. As a result, AODs can also be used to shift the frequency of laser light, which makes them extremely useful in laser locking applications for detuning the light compared to a given atomic transition.

The specific model used in our experiments (DTD-274HD6M by INTRAAC-TION CORP.) consists of two TeO_2 crystals. The laser beam passes through both crystals, which are mounted at right angles with respect to each other. One crystal deflects the light horizontally and the other vertically. This device enables 2D control over the angle of the outgoing laser beam.

Digital micromirror devices

A digital micromirror device is an electronic chip covered with tiny mirrors. For all the different applications described in the thesis, we use the same DMD model (DLP LightCrafter 6500 from TEXAS INSTRUMENTS). The active area




Figure 3.4: A schematic picture of a DMD. An incoming light beam (yellow) is reflected of the micromirrors on the DMD chip. Depending on the state of the mirrors, that are tilted by an angle $\theta = \pm 12^\circ$, the reflected beam is guided into or away from the imaging system. The inset is a microscope image of a DMD chip borrowed from [265].

of the chip hosts a 1080×1920 array of square mirrors, each with a side length of $7.56 \mu\text{m}$ [264]. When the device is off or in standby mode, the mirrors lie flat with respect to the substrate, but once turned on each mirror is tilted $\pm 12^\circ$, depending on whether the state of the mirror is *on* or *off*, as shown in fig. 3.4. The angle of the mirror is controlled by electrodes placed underneath it. The images that can be uploaded and displayed on the DMD are binary black-and-white images.

The DMD chip is interfaced with a computer via a USB connection to an evaluation board that comes with the device. A DMD from our setup is shown in fig. 3.5 (a), where the chip itself is the silvery square in the centre of the image. It can be operated with a GUI provided by Texas Instruments or directly from our experimental control software (see sec. 3.7), where the relevant portion of the GUI code has been modified to fit into that framework [223].

Until now we have only used the DMD in *pattern on-the-fly* mode. In this mode either a static image can be uploaded for display on the DMD, or a series of images may be uploaded and displayed as a movie. The memory buffer on




Figure 3.5: The DMD model in our experiment. (a) The DMD chip of a DLP LightCrafter 6500, containing about 2 million tiny mirrors. This one was used for the magnetometry experiments. (b) The DLP LightCrafter 6500 evaluation board, with the clock thief circuit mounted on top.

the evaluation board limits the total number of images to 400, and they can be swapped at a maximum rate of 9.523 kHz, corresponding to a minimum display time of a single image of 105 μs . Transport of individual atoms has been demonstrated using this feature [256].

Due to the risk of surface adhesion or deformation of the mirrors, that can be caused by the mirrors remaining too long in a tilted state, the manufacturer implements a *mirror clocking pulse* (MCP). This pulse interrupts the state of the mirrors and brings them momentarily to the flat state [140]. Within 3–5 μs it returns to the tilted state. The MCP operates at the maximum rate of the DMD, 9.523 kHz. These are unfortunately typical trapping frequencies one finds in tight lattice traps or even the tight optical tweezer potentials generated in the high-resolution setup (see ch. 6 and onwards). The MCP might then turn out to be an obvious source of heating of atoms residing in these traps, as the intensity of the light from the DMD flickers at this rate (see fig. 4.1 of ref. [223]). The authors of ref. [140] developed an electronic circuit (nicknamed *the clock thief*), that upon a TTL signal grounds the MCP and turns it momentarily off. As a

result, they report a 47-fold increase in the lifetime of their ultracold Fermi gas that has an energy E/\hbar close to the flickering frequency. Having read about their results we decided to implement the clock thief on all the DMDs that were built in to the high-resolution setup, as shown in fig. 3.5 (b). However, the clock thief was not implemented for the DMD controlling the Faraday signal (see sec. 3.5) as the total interaction time with the atoms is at most $80 \mu\text{s}$ during an entire experimental cycle, rendering any heating effects due to the flickering negligible in comparison to heating from the Faraday light itself.

3.4 Making microtraps

To enable imaging with good resolution and to make microtraps[‡] with waists of a few μm , the *science chamber* was equipped with two oppositely-facing re-entrant viewports. Two identical objectives were placed into each viewport, one for imaging and the other for creating microtraps, see fig. 3.6. The distance from the inner side of the window of either re-entrant viewport, to the centre of the chamber was measured to be 30.4 mm. The microscope objectives (LENSOPTICS) were custom designed to have an NA of 0.27 and a focal length of $f_{\text{obj}} = 36.8 \text{ mm}$ for 780 nm light. Their design is described in [5], and they have proven useful in other experiments [2, 6]. Figure 3.6 also shows the positioning of the Helmholtz coil pairs used for stray field compensation, and for applying the necessary fields for the magnetometry experiments.

A TOPTICA DLC PRO operating at a wavelength of 912 nm provided laser light for the microtraps. The light was amplified with a home-built tapered amplifier (TA) and brought to the experimental table through an optical fibre, after which the light entered the setup depicted in fig. 3.7. The first lens after the

[‡]The AOD-generated optical potentials in the magnetometry experiment will be referred to as *microtraps*, and the DMD-generated ones in the high-resolution experiment as *optical tweezers*. Both are obviously microtraps and optical tweezers, but this is done solely to distinguish the two setups. I will routinely refer to both potentials as *dimple traps* or simply *dimples* where they are used for BEC creation via the dimple trick.




Figure 3.6: A schematic setup figure of the science chamber. Identical microscope objectives are mounted into facing re-entrant viewports to enable microtrapping and spatially-selective probing of those. The relevant laser paths are shown: 1064 nm is for the CDT, 912 nm for microtraps and 780 nm for the probing. The positioning of the compensation coils used for background field compensation, and for sweeping the magnetic fields for the magnetometry experiments is also shown. The figure was first printed in [84].

AOD was placed exactly at the distance that minimised the convergence or divergence of an array of beams emerging from it. This beam was expanded to create a large beam that was focussed to a tight trap after the first imaging objective.

The heart of the setup was the AOD, that enables flexible two-axis control of the laser beam. To control the beam we needed a good way of controlling RF signals. For that purpose we used a commercially available DDS chip-based pulse generator (PulseBlasterDDS-II-300-AWG, by SPINCORE TECHNOLOGIES). The signals were amplified by RF amplifiers made by the electronic workshop in-house, and the RF tone was delivered to each axis of the AOD with a power of $1.00(2)$ W. Despite the center operational frequency of the AOD being specified at 27 MHz, we got better diffraction efficiencies around $\eta \simeq 40\%$ by running it at $f_{\text{AOD}} = 35.0$ MHz.

One of the advantages of using the AOD in combination with a flexible RF source is the ability to create arbitrary potentials. To achieve this the $(1, 1)$ diffraction order emerging from the AOD was aligned to the atomic cloud, and then the PulseBlaster was programmed to cycle through a given set of frequen-




Figure 3.7: A schematic setup of the microtrap optics. Details are given in the text. Figure borrowed from [131], and slightly modified.

cies, painting a pattern of light in the atomic plane. At every instance in time, only one laser beam was actually present at the location of the atoms, but by cycling through the frequencies quickly enough, the atoms would experience a time-averaged potential. For the time-averaging to work, one needs the shift frequency f_{sh} of the AOD, to be at least an order of magnitude faster than the typical trapping frequencies in the microtraps f_{mt} . For an array of microtraps with N traps we required $f_{\text{sh}} \geq 10 N f_{\text{mt}}$. The PulseBlaster itself can deliver pulses with a temporal resolution of 13.3 ns, which was more than sufficient for time averaging of these tight traps (which typically have $f_{\text{mt}} \sim 1$ kHz). An example of atoms trapped in time-averaged potentials recorded with absorption imaging is shown in fig. 3.8.

The limiting factor in terms of cycling speed in our system was actually the AOD crystal itself. In order to change the pointing of the laser beam, the sound wave has to travel through the diameter of the beam. This was why we chose a beam waist of only 800 μm , even though the AOD has a full aperture of 4 mm. The access time of the AOD is $t_{\text{acc}} = 1.6 \mu\text{s}/\text{mm}$, yielding a maximal shift rate of about $f_{\text{sh,max}} = 500$ kHz. We observed the same timescale when recording the rise time of the beam intensity on a photodiode. Another problem arose




Figure 3.8: A 3-by-3 array of atoms in microtraps in absorption imaging.

due to the finite propagation time of the sound wave through the crystal. If the difference between two adjacent frequencies $\Delta f < 300$ kHz, we observed a modulation in the beam intensity due to beating of the two frequencies in the crystal, an effect also reported elsewhere [87].

3.5 A setup for spatially-resolved Faraday imaging

As described in sec. 1.4, the Faraday effect manifests itself in the rotation of the plane of polarisation of an off-resonant light beam incident on an atomic sample. The setup described here was capable of shaping the probe light, in an arbitrary manner with a DMD. The surface of the DMD was imaged directly onto the atoms. The beam incident on the cloud featured clean linear polarisation and the amount of the polarisation rotation was detected after the chamber with a balanced homodyne polarimeter.

The probing light was provided by a TOPTICA DL100 780 nm laser. To lock the laser off-resonance, we used a home built beat-lock system. The signal from the Faraday laser was mixed with a reference signal from our master laser source (used for atom cooling and absorption imaging) which was spectroscopically locked to the $F = 2 \rightarrow F' = (2, 3)$ crossover transition in ^{87}Rb . The Faraday




Figure 3.9: A schematic setup of the probe light optics. Details are given in the text. Figure borrowed from [131], and slightly modified.

light was always detuned blue of the atomic transition. Details on the implementation can be found in ch. 7 of [192].

The DMD setup is schematically depicted in fig. 3.9. For the initial alignment of the DMD system we made use of the periodic structure of the chip. The chip works as a grating and by sending a beam backwards through the system onto the DMD with all the mirrors in the on-state, the brightest diffraction order was aligned into the outcoupler that brought light to the system. In this way the grating conditions for the incoming light were matched, maximising the efficiency of the system.

When small structures are uploaded to the DMD, it can be considered as a point source. To align the optical system, we used the light emerging from a single illuminated DMD pixel. All lenses were aligned with the help of a shear plate interferometer (THORLABS, SI035).[†] In between the outcoupling lens and the DMD we installed a lens precisely at the distance of 1 focal length

[†]Doing this in practice is still very challenging as the light level coming from a single DMD pixel is very low. Mounting a camera (IDS, UI-1240SE) directly onto the output of the interferometer enabled a good measurement.

away from the collimation lens after the DMD. This was done in order to provide a curved wavefront for the collimation lens to reduce the effect of non-isoplanatism (discussed below). The system demagnified the DMD image by a factor of $M_1 \simeq 1/22$. Before the light entered the objective it was sent through a *polarising beam splitter* cube (PBS) with a high nominal extinction ratio of $10^4:1$ (from LINOS). This was to ensure the linearity of the light polarisation of the Faraday imaging light. The cube rested on a tiltable mount that allowed for a precise alignment of the balanced polarimeter.

On the outgoing port of the chamber, the second objective re-collimated the Faraday light. Its polarisation was rotated by 45° with a $\lambda/2$ waveplate (mounted in a high precision rotation mount, THORLABS, PRM1/M), before it passed through a Wollaston prism (custom made by ФОСТЕК). The prism was designed to separate the orthogonal polarisations of the outgoing beams, by $\theta_W \simeq 0.3^\circ$. The co-propagating beams were imaged onto an *electron multiplying charge-coupled device* (EMCCD) camera. The magnification of the atom image was $M_2 \simeq 26$.

It is important that the intensity distribution of the light in the Faraday images is homogeneous. As mentioned above, we took precautions to minimise the effect of non-isoplanatism in our setup. This effect arises when a flat surface is imaged onto another flat surface. The optical distances between the object plane and the image plane will be slightly different for the two point sources, so interference effects may arise (see refs. [41, 270] and ch. 6 of [197]). To remedy distortions that may arise from this effect, the DMD was illuminated with a curved wavefront such that the wavefront emerging from the collimation lens was flat. Even though this effect was not pronounced in our system, we could still improve the flatness of our images with the 200 mm lens. We managed to decrease the sum of the squared residuals (subtracting the mean intensity from the image) by 40% with the lens in place. For detailed analysis see ch. 4 of ref. [170].

This lens also resolved another issue. When we imaged a single DMD pixel for the first time, a background field was prominent in the picture. The intensity of the field was constant regardless of how many DMD pixels were turned on




Figure 3.10: Images showing the versatility of light shaping with the DMD. The images shown are taken from one of the ports of the polarimeter. (a) An image of a checker-board pattern. (b) An image of a square. (c) A sample of randomly placed spots. The compilation is borrowed from [131].

or off, and it formed the same image as was formed when the whole chip was turned on, albeit at lower intensity. We could also focus this background light by moving the final imaging lens a few mm from the image plane of the DMD. We believe that this residual background light was due to the MCP discussed in sec. 3.3. We hypothesised that the light was focussed because the DMD chip had a slight curvature such that it acted as a concave mirror. With the lens in place this focusing effect was no longer observed.

In figs. 3.10 (a)–(c) some examples of DMD-forged light patterns imaged through the system are shown. The patterns are rendered from the image recorded of a single port of the polarimeter,** and they show the flexibility of the DMD in creating light patterns. The imaging system is notably not free from aberrations, nor is the illumination perfectly homogeneous. However, this did not limit the capabilities of the setup for the magnetometry experiment. We attribute the imaging imperfections primarily to a long total beam path (~ 3 m) of the probing light, necessitated by the high density of optics around the chambers. To eliminate stray light from sources at other wavelengths than 780 nm in the detection setup, we mounted two interference filters (SEMROCK) at an angle with respect to the propagation of the beam into the camera, to avoid interference effects.

**The intensity of balanced detection is cancelled in the homodyne procedure.

3.6 Technical sources of noise

All means of detecting atoms in our laboratory are based on measuring light. Both the magnetometry measurements and the fluorescence detection in the high-resolution setup rely on an EMCCD camera (ANDOR IXON ULTRA 897). The active area of the camera is a CCD chip with 512×512 pixels, and each pixel is $16 \times 16 \mu\text{m}^2$. Our model is also equipped with another CCD chip connected directly to the sensing chip that can be used as a buffer for storing images.

Figure 3.11 depicts the process of converting a detected photon into an image, see ref. [11]. A photon that strikes a pixel on the chip, which is made of a semiconductor material, is converted to an electron (1). The conversion is not perfect and the ratio between the number of incoming photons N_{ph} and the number of generated electrons N_e is called the *quantum efficiency*

$$\varepsilon = \frac{N_e}{N_{\text{ph}}}, \quad (3.9)$$

where ε is a number between 0 and 1. ε is a characteristic of the semiconductor material used in the chip, and it depends on both the incoming wavelength of the light and the operating temperature of the chip T_{CCD} . For the relevant regime of temperatures, ε is a decreasing function of T_{CCD} . Measurements of ε for the camera used for the Faraday measurements are presented in ref. [192]. At the typical operating temperature of $T_{\text{CCD}} = -40^\circ\text{C}$, $\varepsilon = 0.725(2)$, a value we also adopt for the camera used for the fluorescence detection as the models are identical.

As the exposure time for the image ends, it is read out. That is done by first shifting the electrons that have accumulated in the sensor chip vertically into the storage chip (2), and subsequently each row is shifted horizontally into the next stage (3). Stage (3) could be the optional step of *electron multiplying* (EM) amplification where every electron is multiplied by an EM gain factor G_{EM} (4) or it is alternatively shifted through the EM register without gain and then directed to the *analogue-to-digital converter* (ADC) that converts the electrons to ADC




Figure 3.II: How an EMCCD camera works. Image borrowed from [170], and modified.

counts with a gain factor G (5). The gain $G = 1/S$, where S is the sensitivity, which is the number specified by the manufacturer.

As mentioned, the camera can both be operated in a *conventional mode*, without EM gain, and in *EM mode* with EM gain. In the former mode the number of ADC counts becomes

$$N_{c,con} = \frac{\epsilon N_{ph}}{S}, \quad (3.10)$$

and in the latter mode

$$N_{c,EM} = \frac{\epsilon N_{ph} G_{EM}}{S}. \quad (3.11)$$

This process induces unavoidable noise in the resulting image. This noise stems mainly from three sources, thermally activated electrons; the process of shifting the electrons through the registers; or from the readout process in the ADC.

Thermally activated electrons accumulate over time and induce what is called *dark noise* δ_{dark} or *dark current* (if viewed per unit of time). This noise term is negligible compared to other noise sources at the operating temperature of

$T_{\text{CCD}} = -40^\circ\text{C}$. The shifting process induces what are called *clock-induced charges* (CICs) and contribute to a noise term δ_{CIC} . This can be minimised by adjusting the shift speed and the shift amplitude^{††} in the vertical shift processes, (2) in fig. 3.II. Generally speaking, slower shifts and larger amplitudes yield more noise. Both δ_{dark} and δ_{CIC} are typically very small for normal operation, but their effect is amplified when operating in EM mode. The readout process gives rise to a *readout noise* term δ_{read} that depends on the settings of the ADC converter, its sensitivity and the readout rate (i.e. the horizontal shift rate).

To account for the full noise budget, one must not forget the noisy nature of the photons themselves. If the photons induce N_e electrons, the noise associated due to the Poissonian nature of the photons^{‡‡} is $\delta_e = \sqrt{\varepsilon N_{\text{ph}}}$. The measured signal using the EM register is given by eq. (3.II). As the different sources of noise are all independent, their contributions are added in quadrature. Keeping in mind that all noise is amplified in the same way as the usual signal (apart from the readout noise term which is only amplified by $1/S$) the SNR in the image is [10]

$$\text{SNR} = \frac{\varepsilon N_{\text{ph}}}{\sqrt{F^2(\delta_{\text{dark}}^2 + \delta_{\text{CIC}}^2 + \delta_e^2) + (\delta_{\text{read}}/G_{\text{EM}})^2}}. \quad (3.12)$$

In the equation, F is a multiplication factor that depends on the operational mode of the camera. $F = 1$ in conventional mode, but due to the stochastic nature of the EM process the noise is enhanced by a factor of $F = \sqrt{2}$. The usefulness of the EM register is apparent from the equation as it can be used to effectively reduce δ_{read} . By equating the expressions for the SNR in the EM and conventional modes, assuming that δ_{dark}^2 , δ_{CIC}^2 and $\delta_{\text{read}}/G_{\text{EM}}$ can be made arbitrarily small, and the EM gain high enough to minimise the readout noise in EM mode, the break even point of the two modes is $N_{\text{ph,even}} = \frac{\delta_{\text{read}}^2}{\varepsilon} \simeq 10$ photons/pixel. This is obtained using the readout noise value ($\delta_{\text{read}} = 2.8 e^-$)

^{††}The amplitude controls the depth of the potential in which the electron is stored in each pixel.

^{‡‡}Noise of a signal from a source of a Poissonian nature is routinely called *shot noise*.

for the slowest readout rate available, as specified by the manufacturer. The use of EM gain is then favourable for fewer photons than $N_{\text{ph,even}}$. As this operation setting is horrendously slow, we typically use the fastest mode with $\delta_{\text{read}} = 9.7 \text{ e}^-$, giving $N_{\text{ph,even}} = 130$ photons/pixel. In the ideal case, all noise sources are small compared to δ_e where we retrieve the shot noise limit

$$\text{SNR}_{\text{SN}} = \frac{1}{\sqrt{\epsilon N_{\text{ph}}/F}}. \quad (3.13)$$

3.7 Alice: The experimental control system

Alice, the experimental control system was written in LabView by a former Masters student in the group^{§§}, see ref. [269]. It is a modular system where different parts of the experimental sequence can be easily grouped together in physically relevant groups. The fundamental entity in Alice is a *wave*. A wave is a command to some hardware that lasts for a defined period of time. The control system is connected to a set of analogue and digital output channels, but special waves can be made for special experimental hardware. For example, we have waves to communicate with the DMDs, the objective scanner of the high-resolution microscope, the RF source for the AOD, VISA interfaced instruments and the transport stage used for optical transport of the atom cloud.

Waves can be combined into *blocks* in order to control many channels at once and/or sequentially. Every wave has variables that define its action and duration. These variables can either be local within the wave, or defined as global variables for the block. In this way one may construct a block out of many waves, where the block performs a certain task but only some of the variables remain accessible to the user.

^{§§}The program got its name due to the developer's fondness for Lewis Carroll's *Alice in Wonderland*. The experiment control system is the rabbit hole that Alice tumbles down to enter Quantumland.

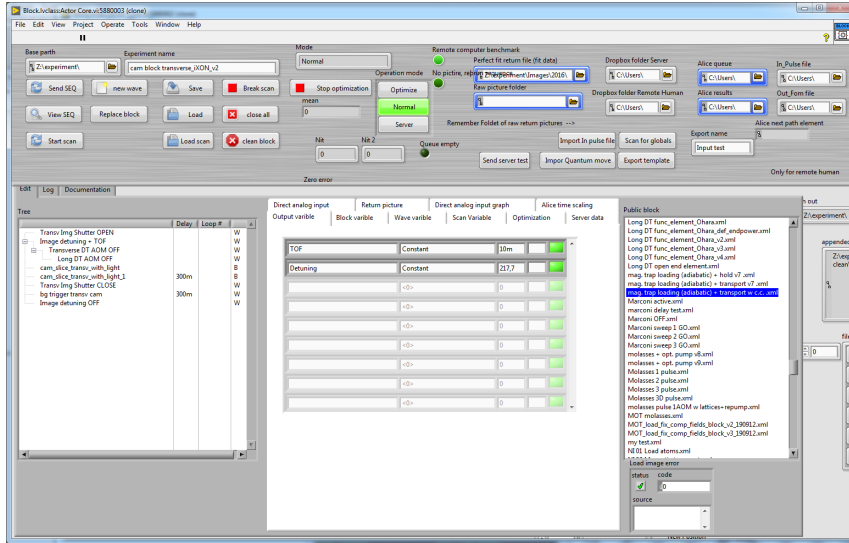


Figure 3.12: A screenshot of the experimental control system Alice. The particular block shown is the one used to acquire absorption images. Two variables in this block are global, the TOF and the frequency of the imaging light.

Let's take an example of the task of acquiring an absorption image in time-of-flight. As explained in sec. 2.5, one needs 3 different images under 3 different conditions. The procedure is roughly as follows:

- Open the mechanical shutter for the imaging light and set the detuning of the imaging light with the AOM.
- Wait for the time defined in the time-of-flight measurement.
- Pulse the AOM that enables fast control of the light and simultaneously trigger the camera.
- Wait until the atom cloud falls out of the field of view of the image system.
- Pulse the AOM that enables fast control of the light and simultaneously trigger the camera.

- Close the mechanical shutter.
- Trigger the camera.
- Set the detuning of the imaging light back to where it was.

The procedure is always the same, apart from the time-of-flight and the frequency of the detuning that controls the image light. This set of waves can be saved as a block with only these two knobs as global variables, encapsulating the rest of the code to make it more user-friendly.

The program features a common pool for frequently used blocks, known as *public blocks* (the list on the right hand side in fig. 3.12), where the imaging block can be saved. With a simple drag-and-drop feature a public block (the list on the left hand side in fig. 3.12) can be pulled into any sequence, which is merely one big block in the eyes of Alice. The interface features an intuitive tree structure, where waves/blocks at the leftmost level are executed sequentially, but an indented wave/block is executed in parallel with the one above. Figure 3.12 features an example where two waves have been indented for execution in parallel with the third wave. By construction, this design enforces modular thought, which is closer to the physical intuition of the operator.

All the output variables of a block can be accessed and scanned in a series of experimental runs. The variables are scanned in a linear manner and any number of variables can be set to scan either simultaneously or in a different dimension of the scan. Only the patience of the operator limits the number of dimensions in a scan, and multiple scans can be executed in series.

As reported in [132] (and in a greater detail in [131]), Alice has also been used to interface a remote optimisation task and a citizen science experiment, where the general public gained control over laser intensities and quadrupole coil currents via a game interface developed in collaboration with the game developers of ScienceAtHome. For a long time we operated Alice via LabView 2013, and when performing long scans the program would occasionally crash. After updating to LabView 2017 we experienced no longer this issue.

CHAPTER 4

Dispersive probing of atomic clouds

Sections 4.1, 4.3 and 4.4 are based on my progress report [83], but have been edited and expanded.

This chapter describes the characterisation of the microtraps and the Faraday probe used for the magnetometry experiments. In the first section we discuss the general properties of the microtraps, and in the second section measurements of the loading from a reservoir into a single microtrap are accounted for. In the third section we shift focus to the polarimeter and describe general properties of the dual-port Faraday imaging technique employed here, and in the fourth section the destructive nature of the Faraday probe is evaluated. The fifth section discusses the signal-to-noise ratios of the Faraday signals and the technical noise level of the detection system is analysed.

Before diving into the details I want to convey the general idea of the spatially-selective probing technique. Figure. 4.1 features a raw Faraday image of five atom clouds trapped in microtraps, in a configuration shaped like the number five on a die. The fluctuations in the image give an idea of the measurement




Figure 4.1: A raw Faraday image of atoms in five microtraps. The red cylinder demarcates the minimal size of a probe beam, demonstrating the spatially-selective nature of the setup. The image was first printed in [84].

noise of a single shot. One atom cloud is encircled with a red cylinder that represents the smallest Faraday probe beam we could make with the DMD driven imaging system. This allows parts of a system to be weakly measured whilst the remainder remains untouched. To our knowledge this was the only cold-atom experiment so far, capable of making local weak measurements.

Recently a similar idea was carried out experimentally in a different context, where AOD controlled, counter-propagating laser beams, were utilised to read and write atomic states in an implementation of a multimode quantum memory [227]. In that case the atom cloud was freely falling, having been released directly from a MOT, which limited coherence times τ_{coh} only to a few μs . We envision that by adding flexible trapping potentials to such an experiment one could extend τ_{coh} greatly.

The method also resembles single-atom spin addressing schemes [277, 281]. In those cases the spin addressing beam represents truly a local interaction between the light and a single atom residing in a single well of an optical lattice,

but it is only used to alter the atomic state, not to image the atom.* In contrast, our setup enables measurements of the collective spin state of an ensemble of atoms.

4.1 Characterising a microtrap

As a first step in the characterisation of the microtrap potentials, we measured their radial trap frequency. In order to excite the atom cloud into oscillation we loaded a single trap from the reservoir cloud of the CDT (see sec. 4.2), and then quickly (in 1 ms) shifted the position of the potential by $5\text{ }\mu\text{m}$ (about one waist of the trap). Afterwards a series of 20 Faraday images was recorded. To capture the in-situ dynamics the images were spaced by $T_{\text{aq}} = 100\text{ }\mu\text{s}$. For a detailed discussion of how the probing was carried out, see sec. 4.3. From each picture we extracted the cloud position by fitting a Gaussian function to it, and to each series we fitted a damped sinusoidal extracting the trapping frequency as shown in fig. 4.2 (a). By varying the power in the microtrap we could measure the waist of the beam using a fit with a simple square root function (see eqns. (1.14), (1.15) and (1.19) that yield $f \sim \sqrt{P}$). The fit is shown in fig. 4.2 (b) as the orange curve and it yields a waist of $w_{\text{d}} = 4.3(1)\text{ }\mu\text{m}$, in good correspondence with what we expected.

Another standard way of characterising the trap is to perform a lifetime measurement. Here, the given trap configuration was loaded and held for a variable amount of time, followed by an absorption image. The atoms slowly leave the trap mainly due to collisions with the background gas and recoil heating of the atoms due to off-resonant scattering with the trap light. First we measured the lifetime of the reservoir cloud. We usually evaporated to powers of $P_{\text{LDT}} = 150\text{ mW}$ and $P_{\text{TDT}} = 630\text{ mW}$ for the individual beams in the CDT,

*One must be careful with language here, as the interaction between the atom and the laser would be a measurement of its state in the quantum measurement sense, as it definitely collapses the local wavefunction to a given state. The only difference here is that the observer does not directly look on an image of the entity being probed.

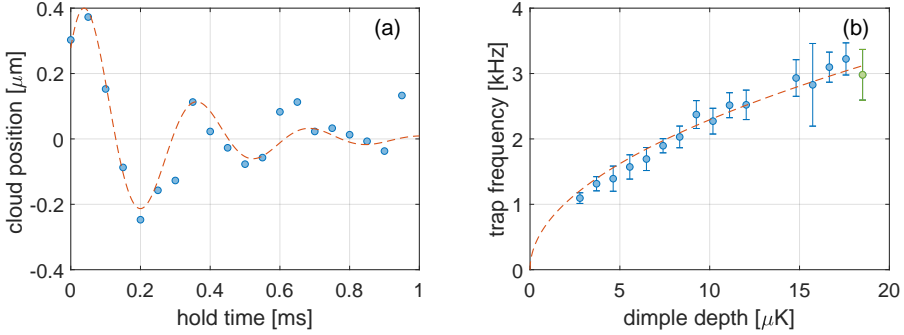


Figure 4.2: Characterising a microtrap. (a) An oscillating atom cloud probed with multiple Faraday images within a single experimental realisation. The frequency is extracted by fitting an exponentially damped sinusoidal to the data. The frequency of this particular trace is $f = 3.0(3)$ kHz. (b) The trapping frequency extracted as a function of the depth of the microtrap. The fitting function is a square root function. The trace in (a) is represented by the green datapoint.

yielding a trap depth close to $5 \mu\text{K}$, taking into account the gravitational sag. This provided an atom cloud of about $2 \cdot 10^6$ atoms at a temperature of about 700 nK , which will be our *reservoir cloud* in what follows. As shown in fig. 4.3 (a) the hold time was varied while the atom number was recorded by absorption imaging. To account for free evaporation (where hot atoms simply escape the trapping potential), a double exponential function was fit to the data. The decay constant of the slower process was taken as the measurement of the $1/e$ trap lifetime yielding $\tau_{\text{res}} = 16.4(3) \text{ s}$. Superimposing a single dimple onto the reservoir with a power of $P_d = 200 \mu\text{W}$ gives a dent in the potential of about $1.5 \mu\text{K}$ in depth. This effectively halved the lifetime to $\tau_{\text{r+d}} = 8.7(2) \text{ s}$, likely due to the enhancement of atom density in the dimple and thus an increased rate of destructive three-body collisions, see fig. 4.3 (a). In light of the measurements presented in sec. 4.2 a small BEC must have been present in the dimple potential.

We also characterised the lifetime in the pure dimple. Loading a single dim-

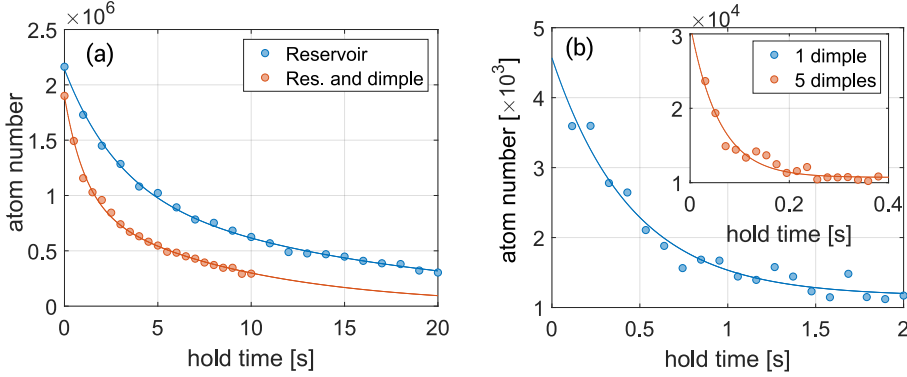


Figure 4.3: Measuring lifetime in different configurations. (a) The reservoir (blue) has $\tau_{\text{res}} = 16.4(3)$ s, and a dimple superimposed in the reservoir (orange) has $\tau_{\text{r+d}} = 8.7(2)$ s. (b) A single dimple, without the reservoir (blue) has $\tau_{1\text{d}} = 450(60)$ ms. The inset shows the lifetime of atoms loaded into a 5 microtrap array with $\tau_{5\text{d}} = 59(8)$ ms. The array configuration resembles the one depicted in fig. 4.1.

ple at a depth of about 800 nK yields the lifetime measurement in fig. 4.3 (b), at a $\tau_{1\text{d}} = 450(60)$ ms. Loading 5 dimples (inset of fig. 4.3 (b)) cuts the lifetime almost by a factor of 10, down to $\tau_{5\text{d}} = 59(8)$ ms. To get good loading of so many dimples and to be able to hold them for a reasonable time, the individual intensities of the AOD generated laser beams were individually adjusted. When using many dimples we did not have a controlled way of measuring the power in the individual traps, but the loading was roughly homogeneous. We estimate that each potential should have been around $8 \mu\text{K}$.

In hindsight we could have put more effort into ensuring the homogeneity of the loading of an array of microtraps of a given size. There are reports in the literature of BECs held in time-averaged AOD potentials with a lifetime exceeding 20 s [26], so there was surely room for improvement. The diffraction efficiency of the AOD varies quite strongly with the applied frequency so effort has to be put into controlling the intensity of the applied RF field to at least homogenise the depth of the individual traps. Such work was in fact implemented in the lab (see ref. [82]), but unfortunately it was not utilised for our

experiments, as we expected to become power-limited due to the reduction of the maximum attainable power in the microtraps.

4.2 Loading a microtrap from a reservoir

A trap of a small volume inside a large reservoir of cold atoms will locally increase the PSD of the trapped cloud. To understand why this is the case one should recall the expression for the PSD, eq. (2.8). A local deformation of the potential will increase the density of atoms n , but as this part of the cloud remains in good thermal contact with the reservoir (being much bigger with a higher thermal capacity) its temperature remains unchanged. This is called the *dimple trick* and it was used in the early days of BEC experiments to adiabatically cross the phase-transition in a reversible manner [221], and later to realise multiple crossings in a single realisation of an experiment [251]. The concept is illustrated in fig. 4.4. The concept of multiple crossings was revisited and supplemented by monitoring the crossing of the phase transition with Faraday imaging [24]. This trick was essential for the realisation of the first cesium BEC [279], because all other traditional methods were limited by heating due to three-body recombination processes. Since then the dimple trick has become one of the standard tools for making BEC's [147, 254]. Recently the loading dynamics of a dimple from a reservoir have been studied both experimentally [106] and theoretically [76].

Our goal was in particular not to create a BEC with the dimple, but merely to maximize our Faraday signal. Thus we set out to characterise the loading of a dimple trap from a reservoir. As a first step we characterised the reservoir by scanning the end point of evaporation and performing TOF measurements to obtain atom numbers and temperatures. The result of that scan is depicted in fig. 4.5. As expected, we see in plots (a) and (b) that the atom number and temperature decrease as the trap is made shallower. In fig. 4.5 (c) there is a $\log \text{PSD} - \log N$ plot typically used to evaluate the quality of an evaporation




Figure 4.4: Enhancing the PSD with the dimple trick. (a) A cold atom cloud resides in a trap of large volume. (b) A local deformation of the potential increases the PSD in the dimple as the atom density rises, but due to thermal contact with the reservoir the temperature stays the same. If the reservoir cloud is cold enough this may bring its energy below the chemical potential μ and condense the cloud.

ramp. The PSD of the reservoir is calculated according to the formula[†]

$$\text{PSD} = N \left(\frac{\hbar \bar{\omega}}{k_b T} \right)^3, \quad (4.1)$$

and at the green point on the graph, we have in fact crossed the phase transition to a BEC. The fit function is to a power law aN^b , and gives a measure of the efficiency of the ramp. The exponent obtained from the fit is $b = -1.04(1)$, which is low compared to [207], but the ramp still does its job.

Later the characterisation was extended and we performed large 2D parameter scans. We varied the depth of the reservoir during the loading of the dimple and the dimple depth. We performed in total three individual scans for different parameter ranges. The sequence was as follows: we evaporated to a variable end point in the CDT, ramped the dimple to the loading power in 80 ms, turned off the CDT and held for 20 ms, then took an absorption image after 5 ms in TOF. In all cases we also took a Faraday image prior to the absorption image.

[†]We use eq. (2.8) and plug in the expression for the thermal deBroglie wavelength, assuming the trap volume through $\frac{1}{2} m \omega_i^2 \sigma_i^2 = \frac{1}{2} k_b T$.




Figure 4.5: A characterisation of the reservoir potential. (a) TOF measurements provide the atom number and (b) the temperature of the cloud. These can be combined to yield the PSD (c), measuring how far we are from the transition to a BEC. The green data point represents a condensed cloud.

The results are depicted in fig. 4.6 and the Faraday results are only shown for one of the scans.

Panels (a) and (b) show the scan with the widest range. The depth of the reservoir endpoint was varied from $0 \rightarrow 100 \mu\text{K}$ and the microtrap depth was scanned from $8 \rightarrow 110 \mu\text{K}$. In (a) the number of atoms obtained from absorption imaging is shown, whereas (b) depicts the mean Faraday rotation averaged over a $10 \times 10 \mu\text{m}^2$ area, around the position of the dimple cloud. As can be seen there is a good correspondence between the two that was also confirmed by other measurements. The maximal atom number loaded in the dimple was obtained when the two were equal in depth at $30 \mu\text{K}$, yielding $N = 3 \cdot 10^5$. The area of good loading is more sensitive as a function of the dimple depth than as a function of the reservoir depth. For dimple powers that were too low, there was no good loading as the volume of the deformation was simply too small. For deeper (and deeper) dimples the effective dimple volume grows and more atoms can be loaded. However, at some point the effective volume becomes too large and the temperature of the whole sample will start to increase. That increase in temperature is caused by adiabatic compression, making the dimple trick less effective.

In fig. 4.6 (c)–(d) we have cut down the scan range of the dimple by more




Figure 4.6: Scanning the reservoir and dimple depths for optimum loading. (a)–(b) The thermal atom number from absorption imaging and the corresponding Faraday signal. (c)–(d) Narrowing the scan range of the dimple depth reveals a region where BECs are created. (e)–(f) By scanning the reservoir depth more finely we see that BECs arise for a large range of trap parameters. Outliers (in the top row of (a), (d) and (f)) are due to poor fits to the TOF data.

than an order of magnitude. Here we only focus on the absorption imaging data, and apply fits of bimodal distributions in order to obtain the numbers of thermal and condensed atoms. No BEC is formed for higher CDT depths as the PSD is simply not high enough for the dimple to cross the transition. As before, if the dimple is ramped too deep, the dimple trick becomes ineffective, so there exists an optimum in the BEC loading landscape.

In fig. 4.6 (e)–(f) the scan range of the CDT has been narrowed. According to our previous characterisation of the reservoir in fig. 4.5, we know that the phase transition to a BEC has not been crossed at a trap depth of $3\ \mu\text{K}$. This means that the vast majority of the BECs here were created with the dimple trick. The resolution of the reservoir scan in fig. 4.5 close to the condensation was not fine enough in order to tell the position of the threshold precisely, but from about $3\ \mu\text{K}$ depth and lower in the CDT, the thermal and BEC numbers decrease towards lower CDT depths. The biggest BECs are obtained for a dimple depth from $2\text{--}4\ \mu\text{K}$ and a CDT depth of $5\text{--}10\ \mu\text{K}$. The optimal ratio of the dimple depth to the CDT depths is $2:5$ for maximizing the BEC size, $1:1$ for maximal loading of atoms.

There are many routes one could take to further analyse this dataset to gain better understanding of the loading dynamics. These would require more focus on the temperature of the cloud and on the PSD. As simple thermodynamic models of the dimple trick exist [61, 244] (see also [199], sec. 5.1), a comparison with theory should be straightforward. I will however, save this as an opportunity for future work.

4.3 Dual-port Faraday imaging

The detection of Faraday rotation, which emerges from the asymmetric response of the magnetic sublevels to near-resonant light, has commonly been realised with photodiodes both in applications for magnetometry [145, 166, 213], spin-squeezing of hot [175] and cold ensembles [60, 240, 261], and in applications involving teleportation of states on light onto an atomic ensemble [242]. In these

cases spatial information about the cloud distribution is lost entirely. Dispersive imaging methods like phase-contrast imaging have been used in the past to acquire spatially resolved images of BECs [12], and their minimally-destructive nature has been used to probe dynamics in them [13, 134]. The first successful implementation of spatially-resolved Faraday imaging I am aware of, dates from 1997. There the method was used to detect the first lithium BEC [38]. Since then, the method has been used to study density-dependent effects in the interaction of light and matter [153], in-trap dynamics of thermal clouds [104], the BEC phase transition [24], and to realise feedback-driven atom number stabilisation with sub-Poissonian number fluctuations [105]. In addition, the technique was recently employed for single-atom detection in a quantum gas microscope [291].

In our experiment, we set up a dual port Faraday imaging (DPFI) system with balanced homodyne detection. The configuration of the detection system was essentially the same as that described in [152, 153], with one important difference. In ref. [153] the light beam was split on a regular PBS after the chamber, and the beams were imaged through separate optical systems on the camera. In our case a Wollaston prism splitted the beam, introducing only a small angle between the different polarisations, and the beams were propagated through the *same* imaging system onto the camera. Our hope was that this configuration would minimise the effects of fringing the other system suffered from, and ease the alignment as beam path differences are circumvented

In our system, a linearly polarised laser beam was guided through the chamber. Afterwards, the polarisation of the light was rotated by 45° with a half-wave plate before it was split into two light beams with orthogonal linear polarisations by the Wollaston prism. Thereafter the two light beams were imaged onto separate regions of the camera, see fig. 4.7. The system was adjusted such that equal amounts of light emerged from the two ports of the Wollaston prism. In this way, any rotation of the plane of polarisation by the atom cloud yielded an increase in the light intensity of one port of the polarimeter and a decrease in the other. As a result, the setup was sensitive to the sign of the rotation angle. In comparison, the method of dark field Faraday imaging (DFFI), measures only




Figure 4.7: A schematic setup of the balanced homodyne polarimeter. Figure borrowed from [131] and slightly modified.

the light that is rotated by the atoms [104]. Note that the SNR of the two methods (and in fact other dispersive probing methods like phase contrast imaging) are the same for small rotation angles, as shown in [104].

A balanced homodyne Faraday detection system is considered ideal for the detection of spin-squeezing [260], which was the primary motivation for building a DPFI system, even though we realised neither projection-noise-limited measurements [168] nor spin-squeezing of our spin ensembles.[‡] We were nonetheless successful in applying the system for magnetometry as detailed in ch. 5.

The light intensity measured in either of the ports of the DPFI system, given an incident intensity I_0 is (see [153] or [152], ch. 5 for the derivation)

$$I_{H,V} = \frac{1}{2} I_0 (1 \pm \sin(2\theta_F)), \quad (4.2)$$

where θ_F is the Faraday rotation angle given by eq. (1.26), and I_H and I_V are the intensities measured coming from the horizontal (H) and vertical (V) ports of the polarimeter. It is straightforward to solve eq. (4.2) above to get an expression for the measured angle

$$\theta_F = \frac{1}{2} \arcsin \left(\frac{I_H - I_V}{I_H + I_V} \right), \quad (4.3)$$

which is the expression applied to generate a Faraday image, as the one displayed in fig. 4.8.

[‡]The former point is obviously a necessary criterion for even attempting to realise the latter.




Figure 4.8: Generating a Faraday image. (a) A raw image of the two polarisation fields as recorded on the Andor iXon camera. (b)–(c) Enlarged versions of the regions demarcated by the white boxes in (a). (d) An image composed according to eq. (4.3) with no atoms present in the field. It is subtracted from an identical image with the atoms present to yield the final image shown in (e). This subtraction is done in order to account for any spatial inhomogeneities present in the imaging light. This figure has been published in [84].

The DPFI system works best when the light intensities in the two ports, I_H and I_V , are equal, i.e. when the polarimeter is *balanced*. To produce the best Faraday image the two light fields recorded on the camera, $I_H(x, y)$ and $I_V(x, y)$, must be subtracted with respect to the actual centre of the light beam in the post-processing of the raw data. A simple algorithm minimises the noise in the subtracted field $I_H - I_V$ with respect to the centre position of the subtraction of the two fields. The subtraction was done to the precision of a single camera pixel. If the centres are not well determined, the resulting image will be very noisy, especially on the edges of the light field, which will also lead to an

artificial shape of image of the actual atom distribution. A detailed account of the balancing procedure can be found in ch. 6 of [192].

Figure 4.8 (a) depicts the two light fields I_H and I_V as recorded on the camera. When speed was not a requirement we typically used a flat top beam as seen in the image. The intensity profile of the beam varies somewhat mainly due to imperfections of the imaging setup. In the raw images, some fringes with a rather large spacing (~ 5 fringes across the beam) are visible. This does not affect the image of atoms in a single microtrap, but we see a spatial modulation in images of bigger clouds. These fringes must somehow have been encoded in the polarisation profile of the beam before it was split on the Wollaston. To minimise the effect of spatial inhomogeneities of the probe beam, a Faraday image without the atoms was also recorded as shown in fig. 4.8 (d). This image was then subtracted from the atom image to produce the final image in fig. 4.8 (e). Here an image of atoms trapped in a single microtrap is shown. Peak Faraday rotation angles of $\theta_F \simeq 15^\circ$, are reached using a probe pulse length of $T_p = 2 \mu\text{s}$ and a probe detuning of $\Delta_p = 1.130 \text{ GHz}$. The choice of these values is explained in sec. 4.4.

Before we built in the Faraday imaging system and the microtrap setup, we pre-aligned the dual microscope system. This was done by overlapping reference beams coming through the opposite side of the chamber. These served as a reference for the optical axis of the imaging system. During this procedure we faced stability issues originating from the mounts we used for the objectives. The objectives were mounted to a five-axis translation stage (x, y, z , tip and tilt), which in turn rested on the main optical breadboard. In order to reach the viewports they were about 35–40 cm in height. At the time there was no dedicated temperature stabilisation system on the experiment (see sec. 6.5), and the alignment drifted slightly on a day-to-day basis due to thermal instabilities. If we were to build in such objectives again, a mount attaching them physically to the chamber allowing only for the freedom of the z -translation, would be much more preferable.

In the alignment process we found that the performance of the dual microscope system would always be sub-optimal. The most severe limitation posed

the high-resolution viewport of the chamber. Due to a manufacturing flaw it extended almost 2 mm too far into the chamber.[§] Another limitation was that we were unsuccessful in re-collimating a collimated beam that was passed through the microscope system. Here the main reason was that the focal length of the objectives was around 2 mm shorter than indicated by the specification of $f = 36.8$ mm. Due to the fixed depth of the re-entrant viewports, we simply could not place them far enough into the chamber. This resulted in an effective $\text{NA} = 0.11$ (compared to the design NA of 0.27). With the correctly manufactured high-resolution viewport we should be able to achieve an NA of 0.19.

According to eq. (3.6) an NA of 0.11 gives a Rayleigh resolution limit of $r_{\min} = 4.3 \mu\text{m}$. To quantify the performance of the imaging system, a 2×2 pixel array (much smaller than the PSF of the imaging system) was displayed on the DMD and imaged on the camera. Fitting the PSF to transversal cuts of the recorded Airy pattern formed on the camera yielded a size of $r_x = 5.3(1) \mu\text{m}$ and $r_y = 4.3(1) \mu\text{m}$, corresponding to diffraction limited performance for the usable aperture. These measurements are described in detail in [170].

4.4 The destructive nature of the Faraday probe

The Faraday interaction is a dispersive effect, and as such it does not cause any heating or loss in the atom cloud. All our measurements however, reveal a substantial signal loss during probing, but the heating of the cloud remains low. We find that it is important to understand the limitations of our dispersive probing scheme as a function of detuning from the relevant atomic transition and the intensity of the probe light. This section will describe our explorations to that end.

To begin with we discuss the effect of the Faraday probing on the reservoir cloud according to data obtained by absorption imaging after TOF. We performed a lifetime measurement (varied the hold time in the reservoir trap) which

[§]The high-resolution viewport was eventually replaced with a correctly manufactured one during a bakeout in the fall 2016.

was followed by a train of 40 Faraday pulses (taken at the rate of 50 Hz, yielding a total acquisition time of 800 ms and an interaction time of 80 μ s), complemented by an absorption image taken after TOF in the end. The Faraday pulses were $T_p = 2 \mu$ s long, and the detuning was set to $\Delta = 1.130$ GHz. The data is presented in fig. 4.9 (a). The blue circles represent the lifetime measurement shown earlier in fig. 4.3 (a) and the orange circles represent the data of the measurement explained here. It is obvious that the Faraday light causes a signal loss. There are immediately two possible causes: recoil heating or state transfer.

As the linearly polarised Faraday light passes through the sample, it can drive σ_+ and σ_- transitions between the $F = 2$ and $F' = 3$ manifolds due to the absorptive part of the light-matter interaction. In line with the loss analysis in ch. 5 of ref. [103], we can start by estimating the number of recoil events per photon. We estimate typical pulse intensities to be about 30 mW/cm² (corresponding to 750 photons per pixel of the camera). Adopting a value of the saturation intensity for the $\sigma_{+/-}$ transition of $I_{\text{sat},\sigma} = 1.67$ mW/cm² [252], we can calculate the scattering rate from eq. (1.9) to be $\Gamma = 2\pi \times 400$ Hz. The intensity of the σ_+ component alone is half of the total intensity and during the total interaction time of $T_{\text{tot}} = 80 \mu$ s, the number of recoil events per atom are only $N_{\text{rec}} = T_{\text{tot}} \Gamma = 0.1$. Comparing the reservoir depth of 5 μ K to the energy gain from a recoil event of $2E_{\text{rec}} \simeq 700$ nK, we see immediately that this mechanism cannot cause all the loss in the signal. In fig. 4.9 (b) we plot the temperature of the atom cloud. The inset (green circles) shows that the heating of the cloud due to the Faraday light is 100–150 nK.[¶] This matches the expected recoil heating which we estimate as $2E_{\text{rec}} N_{\text{rec}} \simeq 140$ nK. For comparison, it is plotted as the purple dashed line.

By comparing the atom numbers of the first shots of the two lifetime graphs in (a), we see that the signal is halved due to the probe light even though the temperature of the cloud does not rise significantly. In ref. [103] state transfer induced by the σ_- component is identified to be the main cause for atom

[¶]It is important to note that the cloud is released into TOF directly after the pulse train is finished. There is no subsequent time to allow for thermalisation.

loss from their magnetic trap, due to optical pumping into an untrapped state. That atom loss channel is not present in an optical trap, but it could still lead to a signal loss as a portion of the atoms might be pumped into the $F = 1$ manifold. A simulation, based on a master equation approach[†] (see sec 2.3.3 of [167] and 7.8 of [253]) reveals that under typical probing conditions^{**} only 1% of the population ends in the $|F = 1, m_F = 1\rangle$ state. Even if one were left with some population distribution in the magnetic sublevels of the $F = 2$ manifold after the Faraday pulse trains, the absorption imaging should be insensitive to that. The population very rapidly goes into steady state under the resonant linearly polarised pulse, bearing no memory of the state it was in after the Faraday pulse train.

This atom loss is really puzzling as the absorptive part of the Hamiltonian seems in no way sufficient to explain the observations. Ideally we would revisit the experiment and set out to measure the heating due to one pulse (instead of a train) and try to measure any fraction of atoms that could have been scattered to the $F = 1$ manifold, by adding repumping to the sequence. Unfortunately the Faraday setup was removed to make way for the high-resolution setup so this exercise must be left to other experiments. As a final remark, there is probably some density dependence on the losses. This is in any case true for the densest clouds, as experiments have revealed enhanced losses under dispersive probing up to a factor of three higher for a BEC compared to a thermal cloud [33]. But as our cloud is thermal at this stage, this enhancement should not play a vital role.

As a next step we wanted to investigate the destructivity of the probe as a function of its detuning Δ and the length of the probe pulse T_p . The signal decays exponentially during the pulse train and as an attempt to quantify the loss, we assume that a cross section σ_F can be associated to this mechanism. We

[†]The simulation is made in a program that was written by a postdoc that was in our group at the time when we did the Faraday work. As the program was not my work, its functionality will not be explained here. Thank you Mario for a nice tool!

^{**}Probe detuning of $\Delta = 1.130$ GHz, total interaction time of $T_{\text{tot}} = 80$ μ s, probe intensity of 30 mW/cm² and a magnetic field component $B_x = 20$ mG (along the probe direction).

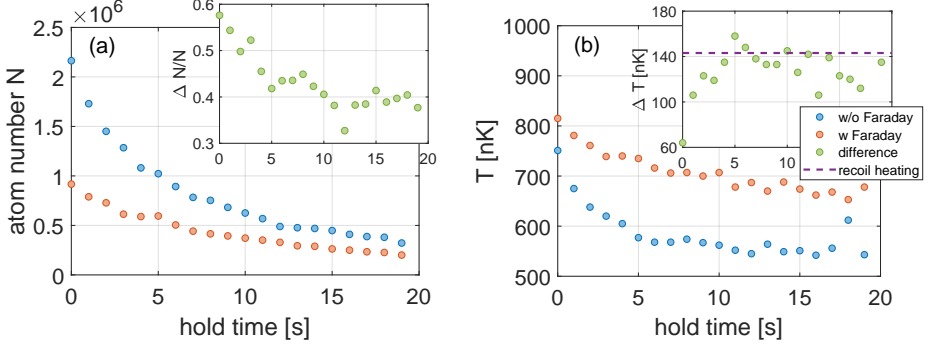


Figure 4.9: Atom loss and heating due to the Faraday probe as measured by absorption imaging. (a) Atom loss in the reservoir cloud, with (orange circles) and without (blue circles) the Faraday pulse train. The inset shows the loss due to the Faraday pulse as a function of time, normalised to the atom number without the probe. (b) The temperature of the reservoir cloud with (orange circles) and without (blue circles) the Faraday pulse train. The inset shows the temperature difference of the two measurements, and the purple dashed line is the estimated recoil heating.

take the loss rate to be

$$R = \dot{N}_{\text{ph}} n_{\text{at}}(x, y) \sigma_F, \quad (4.4)$$

where \dot{N}_{ph} is the number of photons per unit time and $n_{\text{at}}(x, y)$ is the column density of the atoms. Both quantities are evaluated within the same ROI. We take the loss rate as $R = d\theta_F$, the decrement in the Faraday signal, and recognise that according to eq. (1.26) $\theta = \alpha \frac{n_{\text{at}}(x, y)}{\Delta}$, where α is a proportionality constant. From the data we naturally extract the relative drop in the signal of the Faraday angle $\frac{d\theta_F}{\theta_F}$, scaled with the total number of photons in a pulse of a length T_p , so $\dot{N}_{\text{ph}} = N_{\text{ph}}/T_p$. The photon number N_{ph} can be obtained from the raw polarisation fields I_H and I_V registered on the camera. We combine these expressions and define $\alpha \Delta \sigma_F$, as the *destructivity*

$$\eta = \frac{T_p}{N_{\text{ph}}} \left| \frac{d\theta_F}{\theta_F} \right|. \quad (4.5)$$




Figure 4.10: Quantifying the destructivity of the Faraday probe. (a)–(c) Measurements with fixed $T_p = 2 \mu\text{s}$, varying Δ . (d)–(f) Measurements with fixed $\Delta = 1.130 \text{ GHz}$, varying T_p . (a) and (d) show the results for the reservoir potential, (b) and (e) that of a microtrap superimposed to it, and (c) and (f) the measurements for the atoms trapped in the pure dipole potential. The lines between the points are guides to the eyes.

Assuming that θ_F is well described by an exponential function of the form $ae^{-t/\tau}$, yields the ratio $\left| \frac{d\theta_F}{\theta_F} \right| = \left| \frac{1}{\tau} \right|$, so in our case the destructivity becomes

$$\eta = \frac{T_p}{N_{ph}} \left| \frac{1}{\tau} \right|. \quad (4.6)$$

The quantity η has a unit of [events/photon], which is reasonable. As the effect of a single photon is very small, the plots are given in units of [events/ 10^3 photons].

For the measurements we chose to use the top-hat beam shape displayed in fig. 4.8 that encompasses the entire reservoir cloud. The cloud was probed our with 40 Faraday pulses and Δ and T_p were varied. These measurements were conducted in three different trap configurations: in the CDT reservoir prepared at a depth of 5 μ K and containing about $2 \cdot 10^6$ atoms at a temperature 700 nK; with a dimple potential superimposed at a depth of about 1.5 μ K; and with the dimple alone at the same depth. Keeping the pulse duration fixed at $T_p = 2 \mu$ s (which is the shortest possible pulse in our system) we varied Δ in five steps from 0.756 to 1.503 GHz, as the beat-lock allowed. At the central value of $\Delta = 1.130$ GHz, the pulse width was varied from 2 μ s to 20 μ s. Each measurement was repeated five times.

The averaged data for each setting is displayed in fig. 4.10. The traces decay with two apparent timescales. The first decay is about 100 ms and the latter one is slower, exceeding 1 s. The fast decay might indicate a density-dependent loss mechanism. To quantify the losses in the Faraday signal (without knowing their physical causes) we extract decay constants from the traces. We choose to focus on the longer timescale so the data are fit by a single exponential, where we exclude the first 10 points of each trace. This yields good results in terms of fit quality. In the following analysis we leave out the system where the single tweezer was loaded, mainly because the SNR is so poor, as is apparent from fig. 4.10 (c) and (f).

In fig. 4.11 (a) we plot the mean rotation angle in a ROI of $12 \times 8 \mu\text{m}^2$ around the position of the microtrap as a function of the probe detuning using the first

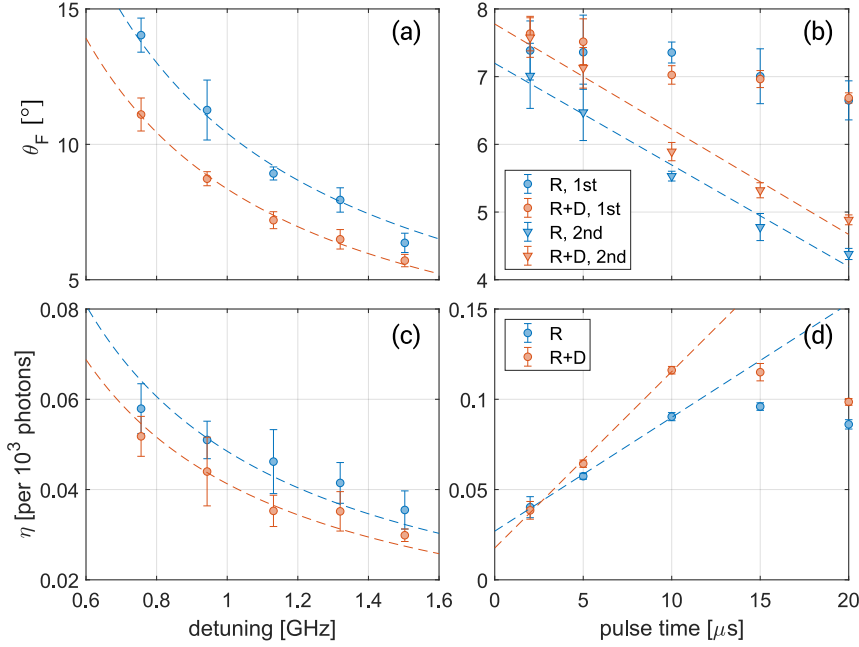


Figure 4.11: Destructivity of the Faraday probe. (a) The Faraday angle varies as $\sim \frac{1}{\Delta}$ as expected. (b) The mean Faraday angle for the first and second images in a pulse train. (c) Evaluating the destructivity η as a function of the detuning. A clear reduction in η is observed. (d) Saturations manifest themselves as more photons interact with the sample per pulse. The legends apply for figures standing side by side.

measurement in the pulse train. The data are well fit by a function varying as $\sim \frac{1}{\Delta}$, as expected from eq. (1.26). In fig. 4.11 (b) we plot the angle as a function of the pulse length for two different points in the pulse train. The circles are rendered from the first image in the train and the triangles from the second image. From eq. (4.3), we would not expect any variation in the rotation angle with higher incident photon numbers, so the reduction in the value of θ_F for the circles is slightly surprising. This must mean that signal loss effects are already visible on the timescale of a long pulse ($\sim 10 \mu$ s), slightly reducing the signal. That effect is even more clear for the triangles, where one pulse has already caused a

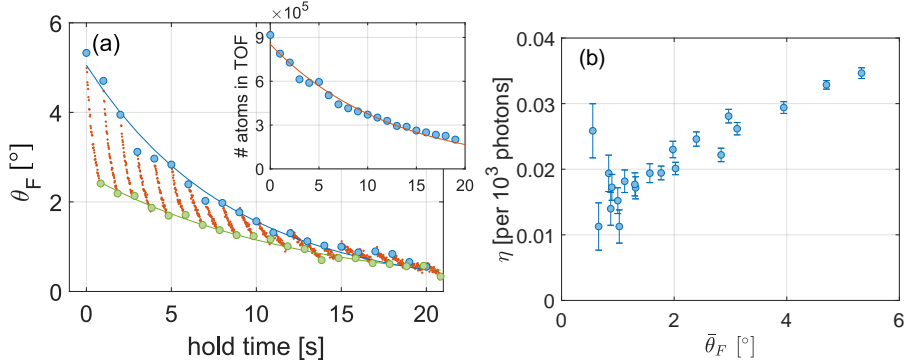


Figure 4.12: Destructivity and sample density. (a) Faraday pulse trains shown in red, taken for after a variable hold time in the reservoir trap. (b) Destructivity η extracted from the plot on the left is seen to fall as the cloud becomes more dilute. Details are given in the text.

significant signal loss. The reduction in the signal hints at a linear dependence of the loss mechanism on the light intensity.

In fig. 4.11 (c), η is plotted as a function of detuning. The dashed lines are linear fits of the form $\sim \frac{1}{\Delta^{\alpha}}$, where α is a free parameter. The fits result in an exponent around 0.8 for both traces. In the discussion above we designated $\eta = \alpha \Delta \sigma_F$, indicating that $\sigma_F \sim \frac{1}{\Delta^{\alpha'}}$, where α' amounts to about 1.8. From the perspective of a simple minded experimentalist it is most likely that $\alpha' = 2$. Finally in fig. 4.11 (d), η is plotted as a function of the pulse length T_p . The dashed lines are linear fits to the first three points, to emphasise how the data deviates from linearity for longer pulse durations. These are clear saturation effects. As the pulses grow too long there are simply more photons than the atoms can interact with. The effect is also clear in the raw data, fig. 4.10 (d) and (e), where the traces for $T_p \geq 10 \mu\text{s}$ all converge to the same level. The value of η is slightly lower when the dimple is superimposed, indicating greater robustness to the probe. This might be caused by the higher densities in the dimple, so saturation effects might manifest themselves sooner compared to a more dilute reservoir.

Let us now revisit the lifetime measurement discussed earlier where a train

of Faraday pulses was set off after a variable hold time in the reservoir. The results are depicted in fig. 4.12 (a). Here the selected ROI encapsulated the whole reservoir cloud, about $35 \times 50 \mu\text{m}^2$. The pulse trains are shown in red, with the first point in blue and the last point in green. Single exponentials are fitted to the first and last points in the Faraday pulse train, giving a $1/e$ lifetime of $\tau_{\theta,i} = 8.5(3) \text{ s}$ and $\tau_{\theta,f} = 12.2(3) \text{ s}$. The inset figure is the lifetime obtained from the TOF measurements, yielding $\tau_{\text{TOF}} = 12.2(4) \text{ s}$, exactly matching $\tau_{\theta,f}$. The lifetime of $\tau_{\theta,i}$ can be compared to the one obtained in sec. 4.1, and we see a lifetime that is almost half of the value reported there. If the trace is refit with a double exponential like the TOF data, the second exponential gives $11.0(1.5) \text{ s}$, which is still shorter than the value of sec. 4.1.

Shifting the focus back to the orange pulse trains of fig. 4.12 (a), one may observe that the slope of the individual traces decreases with the hold time. We attribute this effect to changes in the density of the atomic sample. To each trace a single exponential is fitted and the individual decay constants extracted, enabling the calculation of η . In fig. 4.12 (b) η is depicted as a function of the rotation angle obtained from the first Faraday image of each trace. In a dense sample with higher values of θ_F the destructivity is greater than in a more dilute sample.

4.5 Signal-to-noise ratio of the Faraday images

There are other means of judging the quality of ones signal, than those explored in sec. 4.4. In ref. [249] a careful study of the SNR for a varying detuning of the probe was carried out for Larmor precession signals. Using the dataset discussed in the section above, we can also extract the SNR of the Faraday signal. The noise is estimated by calculating the variance of θ_F on the level of a single pixel. As the density profile of the cloud provides an inhomogeneous background, we cannot consider $\text{Var}(\theta_F)$ to be a pure measure of the noise in the signal. Therefore we use the method described in ref. [143], that should be insensitive to curvatures in

the underlying structure.^{††} The noise in the image is taken as $\sigma_\theta = \sqrt{\text{Var}(\theta_F)}$, yielding

$$\text{SNR}_\theta = \frac{\theta_F}{\sigma_\theta}. \quad (4.7)$$

From the dataset of the previous section we extract SNR_θ and plot it in fig. 4.13 as a function of various quantities. The errorbars on SNR_θ show the standard error obtained from the five measurement repetitions. As a reminder we have two data sets: one where $T_p = 2 \mu\text{s}$ and Δ is varied, and the other where $\Delta = 1.130 \text{ GHz}$ and T_p is varied. In all graphs the quantities are plotted for the first image in the trace with and without the dimple superimposed (blue and orange circles respectively), likewise for the last image in each trace (yellow and purple triangles). In (a) the SNR_θ is plotted as a function of detuning. Similarly to fig. 4.11 (a) it seems to vary as $\sim 1/\Delta$, albeit it levels more clearly off for $\Delta > 1.130 \text{ GHz}$, indicating that there is not so much loss in SNR if we choose to work at a higher detuning. The relative reduction between the 1st and last shot is the least for the dimple system at a detuning of 1.503 GHz. In (b) the SNR exhibits a clear linear relation, but as we want to use the DPFI setup to probe dynamics, the high SNR values are unfavourable as they come at the cost of η . This is apparent from comparing the traces for the first and last images.

This is why it is instructive to compare the two quantities η and SNR_θ . In (c) they are plotted for the dataset where Δ is varied. The interdependence is apparent. Even though the datasets for the reservoir (blue points) and the reservoir with the dimple added (orange) are clearly separated in (a), they still have the same η - SNR_θ dependence. This plot shows the SNR-destructivity trade-off. Finally in (d) we look for an optimum in the ratio of the two quantities. The propagated error on the ratio becomes quite large, making it hard to distinguish an optimum. However, by comparing the first and last images in the traces, the least reduction is yet again observed for the dimple system.

^{††}The method is based on the convolution of the image with a mask that represents the Laplacian operator, which measures curvature. In the process I noticed edge effects in the convoluted image, so a two pixel wide frame around the image was left out in the analysis.

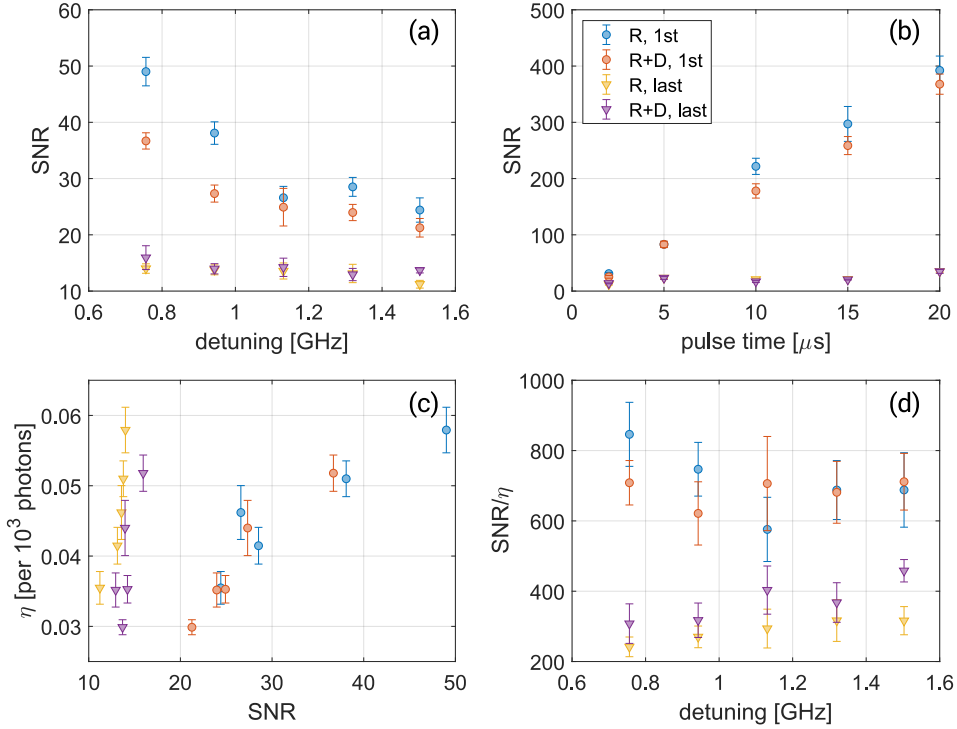


Figure 4.13: SNR for a pulse train of Faraday images. (a) The SNR is seen to decrease with the detuning as expected. It is very different for the first and last image of the trace, and the relative change in the SNR is lower for higher detuning. (b) The SNR increases linearly with the pulse length and drops dramatically for the last image in each trace at longer pulses. In the legend R stands for reservoir and R+D stands for reservoir and dimple. (c) The interdependence of η and the SNR is clear. (d) The ratio the SNR to η is plotted in order to find an optimum in that variable. In fact the least reduction in SNR/ η is observed at the highest value of Δ . The legend in (b) applies for all images.

Ref. [104] gives a good analysis of the SNR for various dispersive methods. All the methods have the same SNR at low angles, but DPFI in fact shows the strongest dependence, and for the typical peak rotation angles of $\sim 15^\circ$, DPFI has 10–15% worse performance than the other methods. Therefore in order to maintain a good SNR, it is preferable to work below $\theta_F \sim 5^\circ$.

The preceding analysis was only done during the writing of this thesis. At the time of the experiments, we chose to work with $T_p = 2 \mu\text{s}$ long pulses as those were the shortest we could create, having least effect on the sample. In terms of the detuning, we continued to work with $\Delta = 1.130 \text{ GHz}$. Ultimately the choice of parameters depends on the application. In our case we should have looked directly at the magnetometer signals and optimised their operating conditions in terms of Δ and T_p , using the measurement precision of the magnetic field.

In the concluding part of this section we set out to determine the detection noise limit in our DPFI system. The Faraday angle is calculated according to eq. (4.3). The argument of the arcsin function, $S = \frac{I_H - I_V}{I_H + I_V}$, is directly proportional to the intensity of the incoming light. The noise properties of S are thus in line with the results of sec. 3.6, yielding $\text{Var}(S) = \frac{F^2}{\epsilon N_{\text{ph}}}$. Now as $S = \sin(2\theta_F) \simeq 2\theta_F$ for small angles, we get^{††}

$$\text{Var}_{\text{SN}}(\theta_F) = \frac{F^2}{4\epsilon N_{\text{ph}}}. \quad (4.8)$$

Here F is a factor depending on the operational mode of the camera.

To measure the detection noise limit, we acquired a set of 100 images for various settings of the camera (shown in the legend of fig. 4.14 (b)), and for each setting we altered the level of illumination by varying the pulse length T_p . From a region of 20×20 pixels, free of obvious imaging distortions, we calculate the variance in the noise pixel by pixel across each set of 100 images. In fig. 4.14 (a), this is done for the homodyne signal $I_H - I_V$ as a function of the number of photons in the pulse. Each dot represents the variance calculated from single set

^{††}Keep in mind that $\text{Var}(aX) = a^2 \text{Var}(X)$.

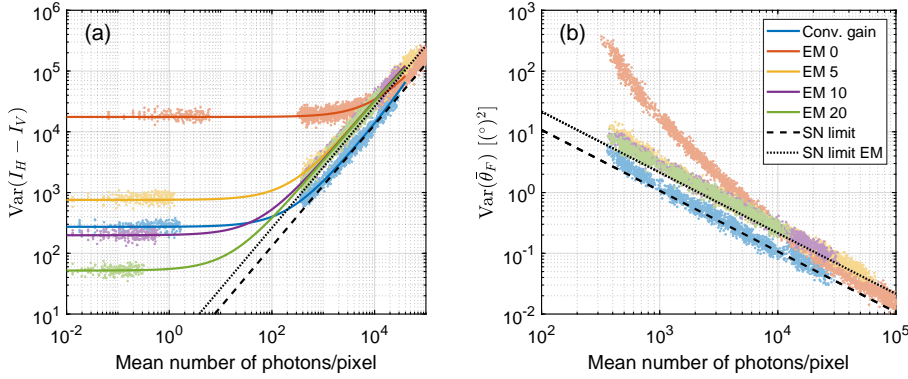


Figure 4.14: Measurements of the detection noise limit of the DPFI system. (a) The variance of $I_H - I_V$ strongly depends on the number of photons that hit a camera pixel. The technical noise level of the camera is indicated by the position where the curves flatten out. (b) The focus is on the relevant region of the photons/pixel for our mode of operation. Here the variance of θ_F is directly plotted.

of pixels. The shot-noise-limited performances with and without EM gain are graphed as a dotted line and a dashed line respectively. The coloured lines are second-order polynomial fits to the data. For high enough photon numbers, the detection limit is well described by photonic shot noise. For intensities of 500–1000 photons/pixel, the conventional register (no EM gain) is clearly the best choice. The chosen setting of the conventional register, is the one with the fastest readout rate and lowest sensitivity. It is also apparent from fig. 4.14 (a) that it never pays off to choose a low EM gain ($G_{EM} < 10$), as those curves always lie above the conventional gain. The break-even point between $G_{EM} = 20$ (the highest measured) and the conventional gain lies around 100 photons/pixel. At around 200 photons/pixel the conventional gain register will always perform better than the EM gain register. This means that it is always beneficial to work without EM gain for DPFI, as that method typically involve intensities above that limit.

In fig. 4.14 (b) the $\text{Var}_{SN}(\theta_F)$ is shown, now in the interesting region from

100 photons/pixel and upwards. The variance in the detected angle decreases as a function of the number of photons used, and it is obvious that the conventional gain register is always the best in terms of noise performance.

Unfortunately these noise measurements were only conducted as the very last thing before we stopped the operation of the magnetometry experiments. As noted earlier, we did in fact run all of our magnetometry measurements with $G_{EM} = 5$, which was a rather bad life choice. Looking closely at fig. 4.14 (a) that setting is still close to the asymptote of the EM shot noise limit at the relevant intensity. All in all, it seems as we could have lowered the noise by a factor of 2–4, depending on ones level of optimism, by choosing the correct mode of operation for the camera.

Spatially-selective optical magnetometry with ultracold atoms

The bulk of the measurements presented in this chapter have been published previously in [84]. The text is also based on chapter 3 of my progress report [83] but those sections have in most cases been rewritten.

This chapter describes two implementations of spatially-resolved, single-shot, in-situ magnetometers. The first is a *vectorial magnetometer* first described in ref. [104]. In this mode of operation we can determine with high *precision* the magnetic field components parallel and transversal to the direction of the probe beam. A schematic figure of the situation is shown in fig. 5.1 (a). This magnetometer is realised by scanning the orientation of an external magnetic field (indicated by the blue arrows) aligned with the probe beam, from positive to negative. This field sweep is performed slowly enough, such that the polarisation of the spin ensemble can follow. The homodyne signal will exhibit a zero crossing, indicating the field component along the beam, and the width of the crossing is a measure of the transversal components.




Figure 5.1: Two cold-atom magnetometers. A signal (shown in green) is recorded by a homodyne detector in two situations. (a) The orientation of a magnetic field component aligned with the probe beam is swept from positive to negative (blue arrows). The magnetisation of the atom cloud (F) follows (red arrows). The zero crossing of the signal (top) is a measure of the field along the beam, and the width of the transition measures the transversal components. (b) The magnetic field is aligned perpendicular to the probe axis and the magnetisation of the atom cloud. Due to that, F will precess around the direction of B at the Larmor frequency ω . The recorded signal (bottom) oscillates at the same frequency.

Our second implementation is that of a *Larmor magnetometer*, based on the coherent precession of cold atoms in an external magnetic field. The response of the spin ensemble is in this case proportional to the modulus of the magnetic field component perpendicular to the probe axis.* In this mode of operation, the field can be measured with high *accuracy*. The spin ensemble precesses around the magnetic field axis with the Larmor frequency, as shown in fig. 5.1 (b). The rate of the precession is directly proportional to the strength of the field. The homodyne signal will oscillate at the same frequency.

In the article this chapter is based on, an error in calculations of magnetic fields in our experiment rendered all fields quoted in ref. [84] a factor of two too small. This has been remedied in the following chapter.

*In ref. [84], a similar description of the Larmor case (p. 2, par. 1) is downright wrong. The keyword *perpendicular*, somehow went fishing.

The first section is devoted to a discussion of the fundamental limits of optical magnetometry. In the second and third sections the operation of the vectorial and Larmor magnetometry are presented in the respective order. In the fourth section, the technique of spatially selective probing is discussed and an outlook is given in the fifth and final section.

5.1 The limits of optical magnetometry

To compare the performance of different magnetometers, the figure of merit is the projection-noise-limited *sensitivity* [47]. The projection noise is the fundamental uncertainty in a measurement of the collective angular momentum of our atomic sample, caused by the quantum mechanical nature of the measured quantity. A coherent spin state is created by optically pumping a sample into a certain hyperfine state. Due to decoherence that arises e.g. from spin-exchange collisions or fluctuations in the background magnetic fields, the lifetime of such states is limited to its *coherence time*. Assuming that we can probe the sample for the full coherence time τ_{coh} and achieve an uncertainty in the phase determination of a single Larmor cycle to the level of $\delta\phi \simeq 1$,[†] the uncertainty in the Larmor frequency ω becomes $\delta\omega = \delta\phi/\tau_{\text{coh}}$. Using N atoms one can improve this by a factor of \sqrt{N} , corresponding to the atomic shot noise limit, and by repeating the measurement (or by measuring over many coherence times which can be achieved by continuous optical pumping of the sample), one can effectively improve the uncertainty by another factor of $\sqrt{T_{\text{exp}}/\tau_{\text{coh}}}$, where T_{exp} is the total experimental run time. The Larmor frequency relates to the Larmor frequency via the gyromagnetic ratio γ

$$\omega = \gamma B, \tag{5.1}$$

[†]This corresponds to the experimentalist's ability to count the number of precession cycles during the measurement time, that is a reasonable bound in our case.

where $\gamma = \frac{g_F \mu_B}{\hbar}$. According to the considerations given above the sensitivity is

$$\delta B \sqrt{T} = \frac{1}{\gamma} \frac{1}{\sqrt{N \tau_{\text{coh}} T_{\text{exp}}}}. \quad (5.2)$$

This quantity has the dimension[‡] of $\text{T Hz}^{-\frac{1}{2}}$, analogous to conventional noise spectra. This makes very good sense for measurements that can be carried out continuously for many coherence times with a high sampling rate. Such measurements can be typically obtained at a high rate and with a *duty cycle* (of measurement time vs. experiment time) close to unity. Examples of these are refs. [64, 166], and the concept of *bandwidth* becomes appropriate. To incorporate the volume in the measure of sensitivity some quote the *volumetric sensitivity* that is normalised to the volume of the sensor (see e.g. ref. [64]), and has a unit of $\text{T Hz}^{-\frac{1}{2}} \text{cm}^{-\frac{3}{2}}$.[§]

In the case of ultracold atomic magnetometers like the ones presented here or in refs. [196, 274], the preparation time of the atomic sample is much longer than the actual measurement time. These magnetometers have a very low duty cycle rendering the concept of sensitivity not particularly useful. Therefore cold-atom based magnetometers will hardly compete with other techniques like *spin-exchange relaxation free* (SERF) magnetometers or *superconducting quantum-interference devices* (SQUIDs) when it comes to practical implementation and commercialisation. The strength of cold-atom magnetometers lies inside laboratory walls, where they can be used to monitor and even correct for fluctuations

[‡]It is a convention in the field of ultracold quantum gases (where my heart beats) to quote magnetic fields in the non-SI unit of Gauss. This is not the conventional choice for the field of optical magnetometry, where researchers tend to stick to the Tesla. Thus all magnetic fields are quoted in units of Gauss, but magnetic field sensitivities are based on the Tesla in order to ease direct comparison with the literature. For reference 1 Gauss = 10^{-4} Tesla.

[§]Having now read quite a few articles in the literature, I sometimes get the feeling that the art of finding the unit that makes ones measurement look best, is widely practised. This work contributes to the large pool of measures by refusing to use the commonly quoted sensitivity as the most relevant metric.

of stray magnetic fields (see refs. [172, 248]). In such applications the absolute precision is more important than bandwidth.

On those grounds it is well justified to distinguish between three different concepts: the *single-shot precision*, which is simply the precision of a single magnetic field measurement, the *single-shot sensitivity*, where one takes the time of a single measurement into account T_{meas} (not the full experimental cycle time T_{exp}), and the conventional *sensitivity*, which takes the full experimental cycle time into account, enabling comparison of different types of magnetometers. This distinction has been made before in the literature, see refs. [274] and [196], although the terminology might be slightly different from what is presented here.[‡]

To understand the full capacity of one's cold-atom magnetometer it is useful to calculate the single-shot sensitivity. If we assume as before that we can measure the accumulated angle of the Larmor signal to 1 rad within in the measurement time T_{meas} , then the error in the frequency is $\delta\omega = 1/T_{\text{meas}}$. For a sample of N atoms one gets, in analogy with eq. (5.2),

$$\delta B \sqrt{T_{\text{ss}}} = \frac{1}{\gamma} \frac{1}{\sqrt{N T_{\text{meas}}}}. \quad (5.3)$$

This underlines a fundamental difference in the operation of cold atom magnetometers to those that use hot vapours, as typical pump-probe setups simply cannot be realised with trapped cold atomic gases due to heating effects. For such schemes to work out the cloud would have to be cooled simultaneously, either via the interaction of the probe itself or by other means. It is rare for cold atoms system to have interrogation times exceeding 1/10 of a typical spin coher-

[‡]As an example of the necessity to impose this distinction one could read ref. [88]. The authors claim a field sensitivity (which I would call single-shot precision) of 94 pT for an interrogation time of 15 ms yielding a “field sensitivity per second” of $12 \text{ pT Hz}^{-\frac{1}{2}}$. This number is flashed in the abstract and compared to the result of [274] that report $8 \text{ pT Hz}^{-\frac{1}{2}}$. That result was however obtained at the full cycle time, so for a fair comparison the number reported in [88] should be larger by a factor of 30–50 (for reasonable BEC preparation times).

ence time (~ 100 's of ms), as detrimental effects of the probe manifest themselves on much shorter timescales.

To calculate the sensitivities of our magnetometers (and this is also widespread practice in the field), we simply multiply the obtained single-shot precision δB either by $\sqrt{T_{\text{meas}}}$ or $\sqrt{T_{\text{exp}}}$ providing $\delta B\sqrt{T_{\text{ss}}}$ and $\delta B\sqrt{T}$ respectively. This can be contrasted with arrangements that allow for acquisition of many data points in a single run where the sensitivity is often determined from the Fourier spectrum of the recorded oscillatory signal. Denoting ΔB as the width of the Fourier peak corresponding to the precession frequency and its height above the noise floor as a measure of the SNR we obtain (see [239])

$$\delta B\sqrt{T_{\text{ss,FFT}}} = \frac{\Delta B}{\text{SNR}}. \quad (5.4)$$

In the case of SERF magnetometers $\delta B\sqrt{T_{\text{FFT}}} \rightarrow \delta B\sqrt{T}$ as there is no preparation time involved.

Aside from the measure of sensitivity, atomic magnetometers can be built to measure DC [166] or AC [278] components of a magnetic field; scalar or vectorial fields [104]; or even gradients [169, 285], where the spatial resolution of the gradiometer plays a role. With industrial demands of miniaturisation the effective area or volume of the magnetometer is also important [196] (supp. mat.).

5.2 Vectorial magnetometry

The method we implemented for vectorial magnetometry was first presented in ref. [104], inspired by the Hanle-type magnetometers described in the introduction. In this scheme a number of Faraday images are acquired as an external magnetic field parallel to the probe direction is swept from positive to negative. The operation requires the speed of the sweep to be slow compared to the instantaneous Larmor frequency in the applied field, otherwise the macroscopic spin polarisation $\langle f_z \rangle$ will not follow the sweep. According to eq. (1.26) the sign

of θ_F depends on whether the $\langle f_z \rangle$ is parallel (positive) or anti-parallel to the probe direction (negative). In contrast to the implementation in ref. [104] that relied on dark field Faraday imaging and was not sensitive to the sign of θ_F , our DPFI setup does distinguish between the different orientations of the spin ensemble along the probe beam.

Technical implementation

The magnetometer sequence was started by preparing a cloud of atoms in the CDT, that had been evaporated to around $2 \cdot 10^6$ atoms at a temperature of 700 nK. Three pairs of orthogonally mounted coils could provide bias fields, enabling full control over the magnetic field at the position of the atoms. Those were used both to cancel external magnetic fields transversal to the probing direction, and to maintain a field along z that kept the atoms oriented. This field was turned on within 200 ms of the end of the evaporation in the CDT. The magnetometer was realised with a train of 2 μ s long Faraday pulses at a detuning of 1.130 GHz where the intensity of each light pulse was about 200 pW/ μ m². Meanwhile the magnitude of the B_z magnetic field was swept from positive to negative. In the process the resulting Faraday angle θ_F changed sign. The point where it changes sign measures the B_z field and the sharpness of the sign change was a measure of how well one had compensated the transversal external magnetic fields.

The compensation coils in our system were configured as shown in fig. 3.6. The vertical coils were very close to being in a Helmholtz configuration, whereas the two coil pairs in the horizontal plane were quite far away from one another relative to their radius. The direction of the bias field they produce was also at a 45° angle to the probe light, making the compensation of the horizontal stray fields and the control of the field parallel to the Faraday light slightly non-trivial. To ease the problem we parametrised the control of those coil pairs in terms of the sum and difference of the fields generated. Assuming a linear dependence of the magnetic field B with the applied coil current I such that $B = \alpha I$ (also

confirmed by measurements), where α is the gain of the coil pair,[‡] we put

$$B_{\text{sum}} = \frac{\alpha}{\sqrt{2}}(I_1 \alpha + I_2) \quad (5.5)$$

$$B_{\text{diff}} = \frac{\alpha}{\sqrt{2}}(I_1 \alpha - I_2). \quad (5.6)$$

The factor of $\sqrt{2}$ comes due to the geometry of the coil setup and in the control program we added a factor α that we determined experimentally. It accounted for inhomogeneities in the coil setup and differences due to inconsistencies in the gain of the different coil pairs. We found a value of $\alpha = 1.10(1)$.

To find suitable compensation fields, we applied the magnetometer sequence where the B_{sum} (B_z) component was swept, scanning the B_y (vertical) field, the B_{diff} (B_x) field and then the α parameter in order to achieve the sharpest slope on the vectorial magnetometer. As an example of the compensation procedure, measurements from a scan of B_x are presented in fig. 5.2 (a), where the mean rotation angle $\bar{\theta}_F$ measured in a ROI of about $20 \times 40 \mu\text{m}^2$ was used. In this measurement B_z was swept by $\Delta B_z = 400 \text{ mG}$ in 800 ms for values of B_x ranging over a similar magnitude. The best value of the B_x compensation coil current (I_{diff}) was marked with red circles, whereas worse values were indicated by the blue traces. As B_x was moved away from the best value, the slope at the sign change decreased.

During our first trials of background field compensation the Faraday images were obtained at a rate close to 100 Hz. As is obvious from fig. 5.2 (b) we observed that the zero-crossing measured by sequential realisations of the magnetometer fluctuated. Note that even though the horizontal axis is different from the one shown in fig. 5.2 (a), they are interchangeable as the magnetic field was swept at a constant rate during the acquisition. By changing the rate of acquisition to 50 Hz and locking the onset of the pulse train to the phase of

[‡]To begin with, the values of α_1 and α_2 were determined independent of the atom signal and were in fact found to be almost identical (as expected), so we adopted a single value α for both. Their values were later made more accurate by a cross calibration with the Larmor magnetometer.

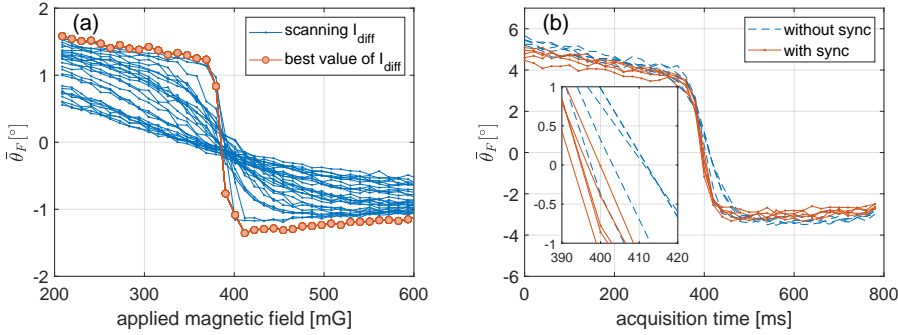


Figure 5.2: Adjusting the vectorial magnetometer. (a) Compensating the transversal fields in the operation of the vectorial magnetometer. (b) As the operation of the magnetometer was triggered on the same phase of the mains line, it reduced the shot-to-shot scatter in the determined fields.

the power line, we could eliminate these fluctuations arising from 50 Hz magnetic field noise. The inset of fig. 5.2 (b) is enlarged around the crossing point of the field. The spread in the blue measurements (without sync to powerline) is considerably greater than the spread in the red curves (with sync to powerline).

Performance of the vectorial magnetometer

An optimised operation of the vectorial magnetometer is depicted in fig. 5.3 (a), where the same ROI is used as before. In each measurement the field was swept by $\Delta B_z = 80$ mG, at a rate of 100 mG/s. The linear relation between the applied magnetic field and the acquisition time is emphasised by the dual horizontal axis. Focusing on the blue circles, we observe an initial decay explained by the interaction of the Faraday probe and the atom cloud (see secs. 4.4 and 4.5). At a given value of B_z , the effect of the perpendicular field components $B_\perp = \sqrt{B_x^2 + B_y^2}$ begin to influence the dynamics of the macroscopic spin polarisation. Eventually the signal drops quickly and changes sign when the magnitude of the applied field equals the background field component $B_z = B_{z,0}$.

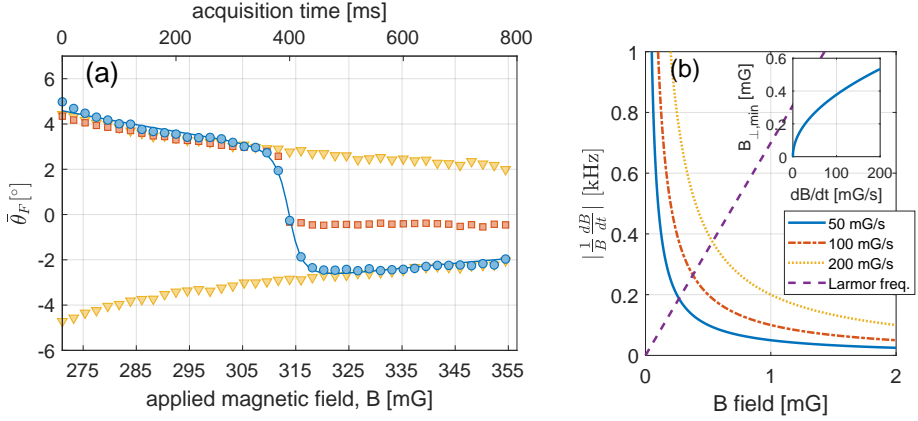


Figure 5.3: The vectorial magnetometer. As the acquisition time is linearly related to the sweep time the horizontal axes are dual, acquisition time on top and applied magnetic field on the bottom. (a) An optimised realisation of the vectorial magnetometer is shown with blue circles and a fit to that curve as the blue solid line. The yellow triangle traces are sweeps taken at a fixed magnetic field showing the decay of the signal due to the interaction of the probe and the atoms (the bottom horizontal axis does not apply here). The orange squares are a realisation of the magnetometer where the background fields were compensated so well that the spin polarisation of the cloud was lost. (b) The relative rate of change of the magnetic field for various sweep rates are shown with the solid lines. The Larmor frequency in the corresponding field is shown as the broken purple line. The inset shows the minimal transverse field needed to maintain the magnetisation of the cloud. Figure (a) was printed previously in ref. [84].

The background and the applied fields point in opposite directions. As B_z is scanned further it reaches a maximum and the decrement can again be explained by detrimental effects in the atom-light interaction.

Based on the assumption that the Faraday signal θ_F follows the projection of the z-component of the field onto the total magnetic field, we expect that $\theta_F \sim (B_z - B_{z,0})/|\mathbf{B}|$, where $\mathbf{B} = (B_z - B_{z,0}, B_\perp)$.** Assuming an exponential

**For more details see ref. [104], sec. III B.

Table 5.1: Precision and sensitivities of the vectorial magnetometer. See details in the text.

	δB	$\delta B \sqrt{T}$	$\delta B \sqrt{T_{ss}}$
$B_{z,0}$	100 μG	55 $\text{nT Hz}^{-\frac{1}{2}}$	9.0 $\text{nT Hz}^{-\frac{1}{2}}$
B_{\perp}	200 μG	110 $\text{nT Hz}^{-\frac{1}{2}}$	18.0 $\text{nT Hz}^{-\frac{1}{2}}$

decay caused by the destructive nature of the probe we arrive at a function

$$\bar{\theta}_F = -A e^{-k B_z} \frac{(B_z - B_{z,0})}{\sqrt{(B_z - B_{z,0})^2 + B_{\perp}^2}}, \quad (5.7)$$

which is fit to the trace and drawn as the solid line in fig. 5.3 (a). The fit parameters are: A the amplitude of the function, k the destructivity decay constant, and the field components $B_{z,0}$ and B_{\perp} . This particular trace yields $B_{z,0} = 313.76(10)$ mG and $B_{\perp} = 2.6(2)$ mG. The precision and sensitivities of the two components of the magnetometer are listed in table 5.1. These results are a significant improvement over the earlier implementation of ref. [104], by more than two orders of magnitude both in terms of the obtained precision and absolute sensitivity.^{††} The absolute sensitivity is obtained using the full experimental cycle time of $T_{\text{exp}} = 30$ s, and the single-shot sensitivities take into account the measurement time $T_{\text{meas}} = 800$ ms.

The precision of our vectorial magnetometer is comparable to other results. Ref. [266] routinely reports precision on the level of $\delta B = 50 \mu\text{G}$, which is the same as the level reached in ref. [248]. That is also true for ref. [25], reporting shot-noise limited single shot precision of $\delta B = 100 \mu\text{G}$ at the 1 kHz bandwidth of the measurement. These experiments all have their different traits, but our results are close to a typical performance in an unshielded environment.

^{††}The sensitivities reported in ref. [104] are $\delta B \sqrt{T_{ss}} = 60 \text{ nT Hz}^{-\frac{1}{2}}$, which are calculated using only the measurement time, not the full experimental cycle time. This statement is based on the reported uncertainty in the zero crossing $\delta B_{z,0} = 13$ mG. Private communication the authors, reveals a cycle time of $T_{\text{exp}} = 90$ s, and thus our implementation improves δB by a factor of 100.

The two traces with yellow triangles in fig. 5.3 (a) are obtained by keeping B_z at a fixed value, and depending on the orientation of the magnetic field the atom signal is negative or positive. As expected, the blue traces asymptote to the yellow ones, indicating that there is no other signal loss other than the one due to the interaction of the atoms and the Faraday probe.

However if the background fields are compensated carefully enough, a sudden reduction in the Faraday signal is observed, as represented with the orange squares in fig. 5.3 (a). This magnetometer relies on the condition that the atomic spins can follow the magnetic field. If the relative rate of change of the magnetic field $|\frac{1}{B} \frac{dB}{dt}|$ becomes greater than the Larmor frequency in the corresponding field, this condition is not met any more and the atomic sample depolarises. This is illustrated in fig. 5.3 (b). The solid, dotted dashed and dotted lines show $|\frac{1}{B} \frac{dB}{dt}|$ for different sweep rates, and the purple dashed line the Larmor frequency in the field. If, for the rates presented here (100 mG/s), the fields are compensated better than about 0.4 mG (that is, the magnetic field corresponding to the crossing point of the orange solid line and the purple broken line) one might expect to lose the spin polarisation. This condition can also be viewed as defining the minimal transversal magnetic field remaining when the Faraday signal changes sign. Equating the Larmor frequency at $B_{\perp, \min}$ to $|\frac{1}{B_{\perp, \min}} \frac{dB}{dt}|$ one gets the condition that

$$B_{\perp, \min} = \sqrt{\frac{1}{\gamma} \frac{dB}{dt}}. \quad (5.8)$$

This is plotted in the inset of fig. 5.3 (b).

The vectorial magnetometer was used to monitor drifts and fluctuations in the magnetic field environment in our experiment. This was done by continually running the sequence over a period of around eight hours. The results are shown in fig. 5.4 where the instantaneous values of the measurement are plotted as a function of the run number of the experiment, referenced to the mean value of each trace $\Delta B = B(t) - \bar{B}$. The shaded area around the traces are 1σ errors of the fits. The histograms to the right are rendered from the values obtained and




Figure 5.4: Monitoring the local magnetic field environment. The timescale of the measurement is close to eight hours. The histograms to the right are rendered from the whole trace. Figure printed previously in ref. [84].

represent the drifts in the magnetic field background. The standard deviation of the histograms amount to 0.32 mG and 0.34 mG for the longitudinal and transversal field components respectively.

5.3 Larmor magnetometry

The most sensitive magnetometers are based on measurements of Larmor precession. Such a magnetometer can only measure scalar fields. To get an intuitive understanding of the basic physics of Larmor precession, let us consider a two level spin system placed in a static external magnetic field.^{††} For a quantum mechanical treatment, we need to find the Hamiltonian that describes the energy

^{††}Examples of how to treat such a problem can be found in refs. [120, 233].

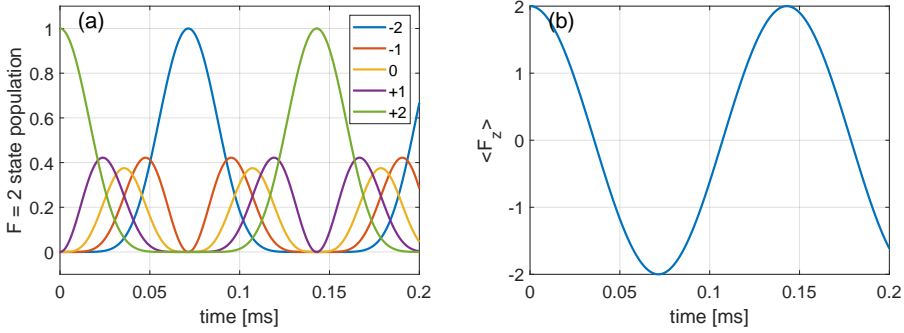


Figure 5.5: (a) The population dynamics of the $F = 2$ ground state manifold of ^{87}Rb , undergoing precession in a magnetic field of 10 mG. The numbers in the legend stand for the m_F number of the magnetic sublevel. (b) The expectation value of the $\langle F_z \rangle$ operator amounts to a sine wave.

in the system. We take the spin state to be fully characterized by the spin operator $\mathbf{S} = \frac{\hbar}{2}(\sigma_x, \sigma_y, \sigma_z)$, where the σ_i are the Pauli spin operators. The magnetic moment that captures the response of the spin system to an external magnetic field \mathbf{B} , relates to \mathbf{S} via the gyromagnetic ratio γ , so $\boldsymbol{\mu} = \gamma\mathbf{S}$.

The energy of the particle in an external magnetic field is given by the Hamiltonian $\mathcal{H} = -\boldsymbol{\mu} \cdot \mathbf{B} = -\gamma\mathbf{S} \cdot \mathbf{B}$. Taking the magnetic field to point in the z direction, $\mathcal{H} = -\gamma\frac{\hbar}{2}\sigma_z B_z$. By solving the time-dependent Schrödinger equation one can calculate the expectation values of the spin components (which is the relevant atomic quantity, according to eq. (1.26)) to be

$$\langle \mathbf{S} \rangle = \begin{pmatrix} \langle S_x \rangle \\ \langle S_y \rangle \\ \langle S_z \rangle \end{pmatrix} = \frac{\hbar}{2} \begin{pmatrix} \sin(\alpha) \cos(\gamma B_z t) \\ -\sin(\alpha) \sin(\gamma B_z t) \\ \cos(\alpha) \end{pmatrix} \quad (5.9)$$

where the α is the angle between the spin vector and the magnetic field. If we apply a field along z and measure $\langle S_x \rangle$ or $\langle S_y \rangle$ those signals will be modulated at ω_L . It is evident from eq. (5.9) that ω_L is related to the magnitude of the

magnetic field via eq. (5.1), so a measurement of the Larmor frequency is also a measurement of the magnetic field.

The atom cloud is initialised in the $|F = 2, m_F = 2\rangle$ stretched state. The gyromagnetic ratio in the $5^2S_{1/2}$ state of ^{87}Rb is approximately 0.7 MHz/G (see [252]), which gives us an idea of the typical fields that can be measured in this way. In the $F = 2$ ground state manifold, the electronic state is 5-fold degenerate in the m_F quantum number. The actual dynamics of the m_F -populations during precession are in fact more complex than in the simple example outlined above. Figure 5.5 (a) shows a simulation of the ground state population in the $F = 2$ manifold undergoing Larmor precession in a magnetic field of $B_z = 10$ mG. The program numerically solves a master equation (see sec. 7.8 of [253]) for the $F = 2 \rightarrow F' = 3$ transition in ^{87}Rb . Provided the populations of the different sublevels P_{m_F} ,

$$\langle F_z \rangle = \sum_{m_F} m_F P_{m_F} \quad (5.10)$$

where $m_F = \{-2, \dots, 2\}$. This quantity is graphed in fig. 5.5 (b), and is in fact a perfect sine wave. This is why the two level image of Larmor precession is so useful.

Technical implementation

Our implementation of the Larmor magnetometer was as follows. After the preparation of an atomic cloud in the CDT as described in sec. 3.1, we ramped up a single dimple to a typical power of $P_d = 200 \mu\text{W}$, giving a trap depth of around 1.5 μK . This locally enhanced the OD, giving a stronger Faraday signal. The magnetometer was capable of measuring an applied magnetic field along the y-direction (according to fig. 3.6). The transversal components to the probe direction B_y and B_z , were compensated to the precision of the vectorial magnetometer (about 0.2 mG). At this point we maintained the magnetisation of the cloud by a small field of 10 mG along the probe direction. The precession was achieved by quickly ramping B_y to 600 mG at an initial rate of 10 mG/ μs . The

current drivers were made by the Institute's electronic workshop with speed and low noise in mind. The ramp of the B_y field was so fast that the collective spin polarisation did not follow and began to precess around the y -axis. The small applied B_z component was subsequently reduced to zero and the B_y field was settled to the value intended to be measured. Exactly 1 cycle of the 50 Hz mains frequency after the B_y field was turned on, the cloud was imaged with a train of 40 Faraday pulses^{§§} at a rate of $f_{\text{aq}} = 44.31(2)$ kHz yielding a total probing time of $T_{\text{meas}} = 925 \mu\text{s}$. Our procedure of acquisition was similar to the one reported in ref. [134].

This fast rate of acquisition was obtained by operating the camera in *fast kinetic mode*. We defined an area of 25×512 (height \times width) camera pixels at the top of the CCD chip onto which we imaged a very small Faraday light patch that fitted into the crop region. This was easy to achieve using the DMD. When an image is taken, the charges that have accumulated in the CCD device are shifted downwards from the crop region to clear it and make it ready for the next pulse. As explained in sec. 3.6, our camera has a buffer chip below the image chip that is equal in size. As a result we can fill up 1024 pixels giving 40 images in total. The Faraday signal θ_F was obtained in the same manner as before where we typically found the mean rotation in a ROI of about $5 \times 5 \mu\text{m}^2$ (10×10 pixels).

One drawback of the method of initiating the Larmor precession with a sudden change in the magnetic field was that the phase of the oscillation was different from shot to shot. A better procedure would have been to use optical pumping instead, which would fix the phase of the initial Larmor precession. An even greater improvement would be to couple this optical pumping light in a spatially-selective way as we do with the probe light. Then the precession could be initiated locally, a feature which is highly desirable especially if the magnetic sensing is to be done in conjunction with other measurements as explained in the outlook section of this chapter.

^{§§}The properties of the Faraday pulses were the same as for the vectorial implementation.

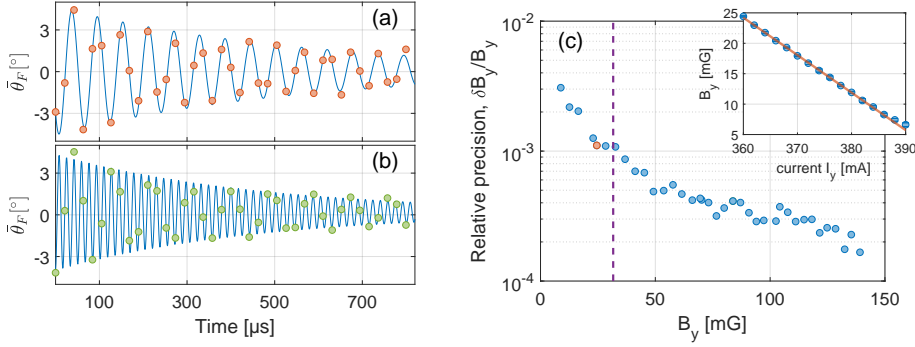


Figure 5.6: Measuring Larmor oscillations. (a)–(b) Examples of the Larmor signals obtained in the experiment. Damped sinusoidal fits to the data provide the Larmor frequency and the application of eq. (5.1) provides the magnetic field. For the (a) $B_y = 24.51(3)$ mG and for (b) $B_y = 83.79(3)$ mG. (c) The fractional uncertainty of B_y as B_y is varied from 9 to 140 mG. The red and green points in the data represent the realisations in (a) and (b). The purple dashed line marks the Nyquist aliasing criterion. The inset shows the calibration of the coils used to provide the B_y field, using the atomic signal. Parts of the figure were printed previously in ref. [84].

Performance of the Larmor magnetometer

The experimental sequence described previously provides signals as shown in fig. 5.6 (a). The data are fit with an exponentially damped sinusoidal function, where the amplitude, damping coefficient of the exponential, frequency and phase of the sine and an overall offset are free parameters. The magnetic field corresponding to the frequency is calculated via eq. (5.1). As a result we measure the B_y component to a precision of $\delta B = 30 \mu\text{G}$. Given the total measurement time $T_{\text{meas}} = 925 \mu\text{s}$ we get $\delta B \sqrt{T_{\text{ss}}} = 900 \text{ pT Hz}^{-\frac{1}{2}}$. Taking into account the full experimental cycle we obtain the absolute sensitivity of $\delta B \sqrt{T} = 16 \text{ nT Hz}^{-\frac{1}{2}}$.

To investigate the operational range of our implementation the applied field was scanned from about 3 mG to 70 mG. The relative single-shot precision $\delta B_y / B_y$ is plotted as a function of the applied field in fig. 5.6 (b). The absolute precision is more or less constant over the range of frequencies measured so $\delta B_y / B_y$ im-

proves with increasing B_y . The Nyquist–Shannon sampling criterion is plotted as the dashed purple line. It is the magnetic field equivalent to $\frac{f_{\text{aq}}}{2}$ (see ref. [241]). The measured fields above it will suffer from aliasing, but as we can follow the precession frequency from below $\frac{f_{\text{aq}}}{2}$, there is no doubt about the absolute values of the magnetic fields above $\frac{f_{\text{aq}}}{2}$.

The magnetic field is obtained directly from the Larmor frequency so this magnetometer is not only precise but also accurate. With an absolute ruler for magnetic field strength at hand we cross-calibrated the vectorial magnetometer in order to improve its accuracy. The inset of fig. 5.6 shows the linearity of the measured field as a function of the applied current to the coils I_y . A linear fit describes the data very accurately, apart from the two points corresponding to the lowest applied fields we measured. As the applied field is low, the transversal components are compensated relatively poorly and so they are excluded in the fit. The calibration yields a gain of 0.6161(6) G/A, and by virtue of this calibration we can measure accurately the B_{\perp} component in the operation of the vectorial magnetometer. By scanning B_y we obtain different values of B_{\perp} from the fits and these values fix the calibration of B_z , the field component which is swept in the vectorial magnetometer.

It is important to test the stability of the Larmor signal within a single trace and see if any signs of frequency chirping are present in the oscillation. Such features could arise from high frequency components in the magnetic field background that cannot be compensated passively with DC fields. As the total probe time is close to 1 ms we should be sensitive to components with frequencies > 1 kHz. To test this hypothesis 80 Larmor traces were obtained with a constant applied field of $B_y = 24.2$ mG. The data are fit with a damped sinusoidal as before to the full trace, to the first 20 points (left hand side), and to the last 20 points (right hand side). A scatter plot of the obtained ω_{left} and ω_{right} normalised to the full trace frequency ω_0 is shown in fig. 5.7 (a). The circle is the mean of the 1σ errors of the fits, so statistically speaking, we would expect approximately 68% of the points to lie within that circle. This is in fact the case, with 18 points lying outside the circle. There seems to be no systematic variation

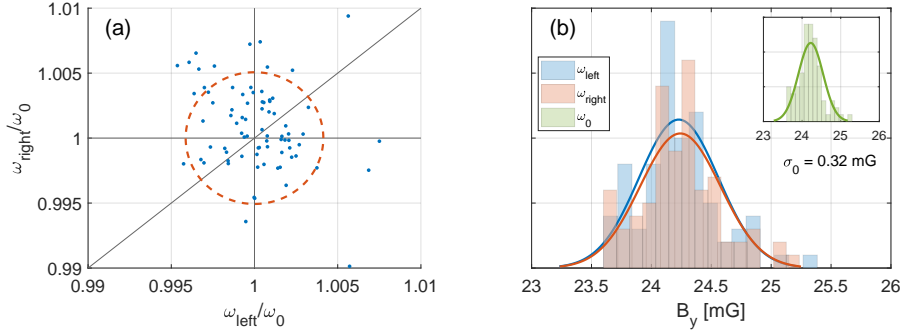


Figure 5.7: Testing the stability of the Larmor oscillation within a single run. (a) A scatter plot of the left and right frequencies obtained from the fits. The circle represents the average 1σ error of the fits. (b) A histogram of the frequencies fitted to the left and right side of the traces. The inset shows the result of fitting the full trace. The fits are Gaussian used to extract the mean and standard deviation of the distributions.

in the frequency during the probing time.

Figure 5.7 (b) shows a histogram of the magnetic fields from the individual runs for the left and right side of the trace. As is apparent from (a) there is no difference between the left or right distributions. The solid lines are Gaussian fits giving the mean value (μ) and the standard deviation (σ) of the distribution. The inset shows the histogram of fits using the full traces. The values of μ for all distributions are identical as to the 4-th decimal place. The value of σ is a measure of the shot-to-shot fluctuations, the residual jitter in the stray fields. For the full trace it amounts to $\sigma_0 = 0.32$ mG, compatible with the 0.32–0.34 mG obtained from the eight hour operation of the vectorial magnetometer, shown in fig. 5.4. As the single-shot precision is around an order of magnitude better than the shot-to-shot fluctuations, the magnetometer could be used in-sequence with other magnetic-field-sensitive measurements to improve their precision, either by active feedback stabilisation [248] or by post-selection of data conditioned on the field value obtained from the magnetometer [24, 172].

The single-shot precision of our Larmor magnetometer is among the best

cold-atom optical magnetometers that have been realised. It is on par with the results of refs. [134, 145, 285] and exceeds that of refs. [90, 169] by an order of magnitude. The results of the squeezed lattice magnetometer of ref. [196] at $\delta B = 3.1 \mu\text{G}$, the squeezed BEC interferometer of ref. [205] at $\delta B = 150 \text{ nG}$, and the best cold-atom magnetometer to date in ref. [274] at $\delta B = 9 \text{ nG}$, remain unprecedented.

A recent realisation of an RF-driven cold-atom magnetometer claims to hold the record for non-degenerate, non-squeezed atom ensembles reaching single-shot sensitivities of $\delta B \sqrt{T_{\text{ss,FFT}}} = 330 \text{ pT Hz}^{-\frac{1}{2}}$, and single-shot determination of magnetic fields down to $\delta B = 7 \mu\text{G}$ [59]. The operation of this magnetometer permits a bandwidth-based definition of the sensitivity (due to high sampling rate), making an exact comparison to our single-shot sensitivity of $900 \text{ pT Hz}^{-\frac{1}{2}}$ not straightforward. In the paper they compare their result to ours (among others) and claim that our sensitivities lie in the several $\text{nT Hz}^{-\frac{1}{2}}$ regime, which is true if one takes into account our absolute sensitivity. But their notion of bandwidth does not take their full experimental cycle (around 7 s) into account. Thus their claim of exceeding other results by an order of magnitude is not justified.

The supplemental material of ref. [196] gives a great overview of the performance of magnetometer sensitivities for a great variety of platforms and plots them as a function of volume. To place ourselves on such a plot we need to estimate the volume of the atom cloud in the microtrap. Assuming a thermal distribution of atoms, the volume can be taken to be $V = (2\pi)^{\frac{3}{2}} \sigma_x \sigma_y \sigma_z$ where $\sigma_i = \sqrt{\frac{k_b T}{m \omega_i^2}}$ (see sec. 2.3 of [219]). Using the expressions for the trapping frequencies ω_i in eqns. (1.19) and (1.20), we get that for a single microtrap

$$V = \frac{\pi}{2} \left(\frac{\pi k_b T}{U} \right)^{\frac{3}{2}} \frac{w_d^4}{\lambda}, \quad (5.11)$$

where T is the temperature of the atoms, U is the trap depth and w_d is the beam waist. Assuming that the kinetic energy $\pi k_b T = U$, which is a good

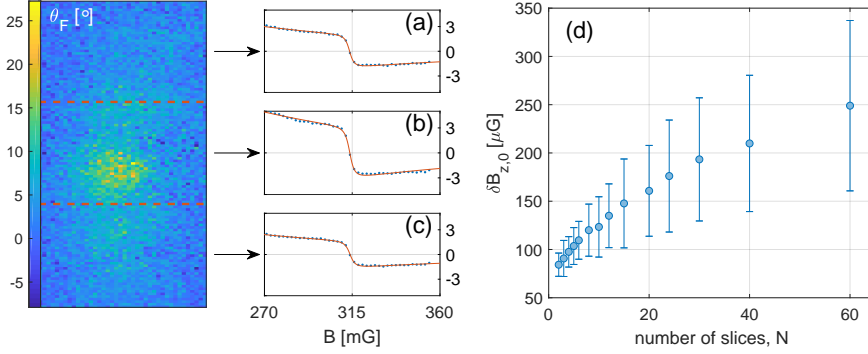


Figure 5.8: Precision and area of evaluation. The graphs (a)–(c) represent the magnetometry traces as obtained from slicing up the Faraday image of the cloud as indicated by the red dashed lines to the left of the curves. The plot in (d) shows the change in the precision as the number of slices is increased. Figure printed previously in ref. [84].

approximation for the atoms in the microtrap,^{¶¶} and using the measured value of $w_d = 4.3 \mu\text{m}$ we arrive at $V_{\text{dim}} \simeq 600 \mu\text{m}^3$, and that places us very close to ref. [286].

As a concluding point of this section it is worth estimating how far we are from shot-noise-limited performance according to eq. (5.3). Given that ~ 1000 atoms reside in the dimple potential^{***} and the measurement time is ~ 1 ms, we get $\delta B \sqrt{T_{\text{ss}}} = \frac{1}{\gamma T_{\text{meas}} \sqrt{N}} = 20 \mu\text{G}$. This indicates that we cannot be very far from shot-noise-limited performance, given the atom number and the probing time of the realisation. During operation, there may well be more atoms along the line of sight than ~ 1000 , which would pull down the limit. But as stated in the concluding remarks of sec. 4.5 we had tried and failed to confirm projection-noise-limited performance of the detection system. In any case we could have lowered the technical noise level by using the proper mode of the camera.

^{¶¶}In any case the usual rule of thumb that the energy is about 1/10 of the depth hardly holds here, as new reservoir atoms constantly flow in as old ones are removed from the potential.

^{***}This is justified according to fig. 4.6 (bottom) with a $U_{\text{res}} = 5 \mu\text{K}$ and $U_{\text{dim}} = 1.5 \mu\text{K}$.

5.4 Spatially-selective probing and detection

There are two *spatial* aspects of our measurement setup. Due to the CCD camera we have *spatially-resolved* detection and by using the DMD to shape the probe it can be made *spatially-selective*.

The operation of our magnetometers does not rely on the spatial resolution of the camera. As one reduces the ROI for integration of the signal the SNR is usually reduced. To test this trade-off we analyse a single realisation of the vectorial magnetometer, a set containing in total 40 Faraday images. The ROI is chosen to be $20 \times 60 \mu\text{m}^2$ in size (40×120 pixels) capturing the entire reservoir cloud, as shown in the leftmost image of fig. 5.8. The image is now split into N slices along its longer side. For each slice the average Faraday angle $\bar{\theta}_F$ is calculated for the whole series of measurements, and a magnetometer traces like those in the figs. 5.8 (a)–(c) are fitted. The precision of the zero crossing of the magnetic field is obtained by fitting eq. (5.7) for each of the N slices giving $\delta B_{z,0,i}$ (where $i = 1, \dots, N$) and so the mean value and the standard deviation of the $\delta B_{z,0,i}$'s can be determined for each N . The result is graphed in fig. 5.8 (d). As expected, the precision of $B_{z,0}$ is reduced for smaller ROIs.

By dissecting the image into N regions, the setup can measure N -th order variations in the magnetic field. This is perhaps not so useful in the middle of a vacuum chamber, far from all sources of magnetic fields, but it is much more relevant close to the surface of microscopic circuits. Such an ability has been demonstrated before, where a BEC on an atom chip was used as a sensor of the complex structure of electric and magnetic fields close to its surface [286]. It is however realistic to have a gradient present at the location of the atoms in a large vacuum chamber, and the vectorial magnetometer achieves gradiometric precision of $3.0 \mu\text{G}/\mu\text{m}$.

The main advantage of the flexible control over the shape of the Faraday light provided by the DMD is that parts of the atomic system can be measured, leaving the rest unaffected by the probe light. To test if this is indeed the case we prepared an atom cloud in a single trap and obtained two Faraday images. The size of the probe was as small as we could make it: a patch with a waist of

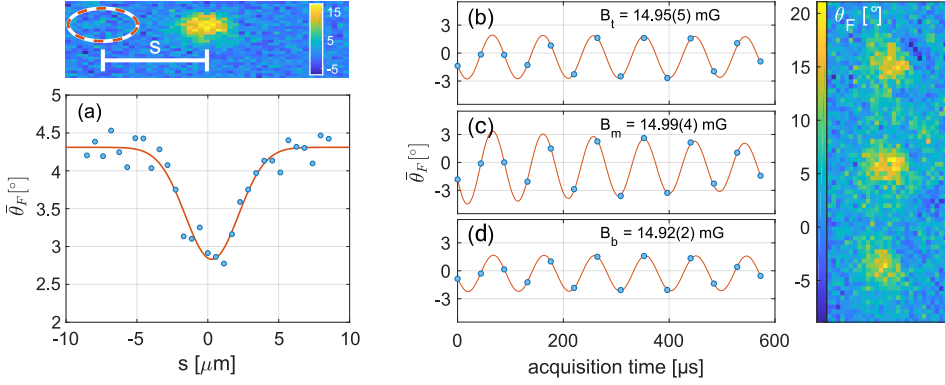


Figure 5.9: Spatially-selective probing. (a) Establishing the locality of the Faraday probe. (b)—(d) Simultaneous Larmor oscillation in a triple microtrap system. Figure printed previously in ref. [84].

3.8 μm . In the first image of each run the probe was at a distance s from the position of the atoms in the trap (marked with a red broken circle in the top panel of fig. 5.9 (a)) and in the second image the probe measures the cloud. The distance s was scanned across the cloud in multiple realisations of the experiment. The Faraday angle obtained from the second image is plotted as a function of the distance s in fig. 5.9 (a). The first probe clearly has a detrimental effect on the signal if the beam overlaps with the cloud. The solid line in the figure is a Gaussian fit, from which we extract the $1/e^2$ radius of 3.8(3) μm in accordance with what is expected from the size of the probe beam. This eliminates the need to screen stray light which might originate from the probe source, which is often achieved with a razor blade properly placed in the intermediate imaging plane.

The final measurement utilises most of the capabilities of the experiment. Three microtraps are generated in a row spaced by 13 μm , and superimposed to the reservoir. The system is shown in the image to the far right in fig. 5.9. The atoms were put into Larmor precession and probed with a train of 14 Faraday pulses spanning 570 μs . The size of the probe encompassed all three clouds. The spatially-resolved detection was employed to extract the angle from a $10 \times$

$10\text{ }\mu\text{m}^2$ ROI around each atom cloud and the result is shown in figs. 5.9 (b)–(d). This experiment shows the potential of the spatially-selective probing scheme, where the two outer clouds could be used for magnetic field measurements while some magnetic-field-sensitive process takes place in the centre cloud.

As a result we can put an upper bound on any residual magnetic field gradient, that must be $\frac{dB_y}{dy} < 3.0\text{ }\mu\text{G}/\mu\text{m}$. This number could be improved by a factor of 2–3 by separating them only by one waist of the microtrap. State-of-the-art cold atom gradient magnetometers are at least one order of magnitude better, reaching gradient precision of about $100\text{ nG}/\mu\text{m}$ [169, 196, 285]. The best result I found in the literature is that of ref. [289] where two BECs were used in a Ramsey type interferometer to reach gradiometric precision of $3\text{ nG}/\mu\text{m}$.

5.5 Outlook

In this chapter, two high-precision cold-atom optical magnetometers have been presented. Their sensitivity is orders of magnitude from competing with state-of-the-art cell-based hot-vapour magnetometers, but their real potential lies elsewhere. The small diffusion of cold-atom clouds makes them excellent for spatial imaging of magnetic field environments. We have shown our method’s potential to measure gradients in 1D, but by trapping atoms in a grid of microtraps, gradients could be measured in 2D. With improvements of the machinery enabling micro-trapping in 3D, as was achieved recently in ref. [53], the magnetic field environment could be mapped in full 3D.

An even more interesting route to take is to use the magnetometers in conjunction with other measurements of quantities that depend strongly on the local magnetic field. Such a methodology was demonstrated recently in our experiment, where an in-sequence Faraday measurement of the phase space density was used to reduce shot-to-shot fluctuations in the determination of the transition from a thermal cloud to a BEC [24]. In another recent example a microwave-based magnetometry scheme was used to reduce measurement fluc-




Figure 5.10: Spatially-selective Faraday measurements could enable magnetometry in simultaneity with other magnetic field sensitive measurements. Here an atom ensemble has been separated into a sensor region, in the four corners of the illustration (two of which are being probed), and an interaction region in the centre where quantum spin dynamics take place. That process is prone to error in the presence of fluctuating magnetic fields. The figure was printed previously in ref. [84].

tuations of magnetic field sensitive variables [172], so the research community has interest in such techniques.

The spatially-selective nature of our probing scheme enables a refinement of this idea as the fluctuating magnetic fields could be measured in portions of the system, whereas the rest remains unharmed by the measurement. This concept is captured in fig. 5.10 where small atomic clouds in the four corners of the illustration form *sensor regions* and the atoms in the egg-tray potential make up an *interaction region* where magnetic field sensitive dynamics takes place, such as quantum simulation of spin physics [246]. This interaction region is intentionally drawn as such for the reader to make mental connections to atoms in optical lattices which are the topic of the second half of the thesis.

Apart from these avenues, the spatially-selective probing holds other promises. The original purpose of the experiment was to realise spin-squeezing in the style of ref. [240]. Due to high levels of technical noise, we did not reach the atom

shot noise limit. I am sure the story would be different had we operated the detector in the right regime and worked some more on the matter. With high enough ODs in the microtraps the achievement of spin-squeezing should be feasible. Shaping of the probe light could allow for local squeezing of atoms in a microtrap array.

Weak measurements also allow for experiments with active feedback. Such implementations require fast acquisition, readout and analysis that can be achieved on an FPGA on a timescale close to 1 ms. One could imagine extracting the atom number and temperature on the fly to get an estimate of the PSD as was done in [24] and feed back on the system to stabilise that quantity. That could prove useful for experiments that require stable atom numbers and temperature, e.g. for deterministic preparation of the filling fraction in an optical lattice for a subsequent quantum simulation. Recently Faraday imaging was used to achieve number stability of atom clouds at the atom shot noise limit [104, 173].

Finally spatially-selective probing could be used to measure parts of a coupled atomic system like a Bose-Josephson junction [3, 247]. Such a two-level tunnel-coupled system exhibits two distinct phases; an oscillatory phase where the atom population oscillates between the wells and a self-trapping phase, where the self-energy (the interaction energy due to the $U_0|\Psi|^2$ term in the Gross-Pitaevskii equation) is so high that the oscillatory phase is prohibited by energy conservation. With the help of a local probe one might remove atoms from one of the wells, either to cross from the self-trapping to the oscillatory phase or vice-versa. Due to heating effects such an experiment might be difficult. I still believe that local probing would be of interest to extract information of parts of a quantum system [232, 276], or to induce dynamics in portions of a system.

This technique has a great potential in field of cold-atom physics; both if employed for the measurement of typical error-inducing quantities in-sequence with other measurements; and for the purpose of inducing and controlling dynamics and quantum states [259].

New machinery for the high-resolution experiment

The vision for the high-resolution experiment at Århus University is to be a state-of-the-art quantum simulator for bosons. The machine will be able to prepare atoms in a single layer of a 3D optical lattice with unity filling and image them through a microscope with a resolution to the level of a single site of the optical lattice [243]. The novelty of the machine will lie in the combination of arbitrary far off-resonant optical potentials and local spin addressing of individual atoms. Static and dynamic potential landscapes can be realised with the ability of the DMD for cycling through a multitude of images that could be used for the purpose of atom transport and quantum optimal control. The addressing scheme is realised on the $F = 1 \rightarrow 2$ hyperfine transition in the ground state manifold, using a combination of near-resonant light and microwaves [281].

* * *

In October 2016 we stopped doing experiments with the magnetometry setup and started to incorporate parts for the high-resolution experiment. As explained below this required a bakeout of the science chamber. Naturally we had to strip

the setup of all optics in the vicinity of the parts to be baked out. The 912 nm microtrap system and the DPFI imaging setup were permanently removed and so was the transverse dipole beam in the cube chamber we used to create a crossed dipole trap for the Alice Challenge [132].

The chapter below describes all modifications that have been made to our experiment for the high-resolution experiment and supplements the description provided in refs. [131, 199]. The first section covers changes made to the science chamber, and the second section describes our microwave system in that chamber. The third section describes the laser setup around the science chamber, and the fourth section that of the high-resolution optical breadboard. The fifth section describes the climate system constructed around the experiment table to stabilise the temperature on the experiment table.

6.1 The viewports of the science chamber

A top view of the science chamber is depicted in fig. 6.1 (a) and its viewports are marked with letters A–E. The original design of the science chamber had to include access for many different optical setups: the high-resolution part from below (A), the re-entrant viewports used in the magnetometry experiments (B), and the viewports on the lattice axes (C and D). To begin with the experiment was designed to support superlattices for qubit manipulation, like is described in ref. [151]. Soon after we started building the optical lattices we quickly diverted from the plan of leaving space for those optics. This was both because space is really scarce around our main chamber, and it seemed like the DMDs could be used instead of superlattices to realise collisional qubit gates by shuttling atoms around [282]. This is however, the reason why one of the lattice axes is equipped with a re-entrant viewport for internal reflection (D) that reflects 1064 nm light (and transmits 780 nm), ensuring the stability required for precise phase control of the different coloured standing waves. As explained in sec. 7.5 this design is a limiting factor for the achievable depths of our lattices. All lattice viewports, apart from the retro-reflections at (A) and (D), have an an-




Figure 6.1: The science chamber. (a) The different viewports are marked with letters A-E. (b) A cross sectional view of the science chamber. The objective enters on the bottom. The red blob marks the position of an atom cloud.

gled window (seen to the left in fig. 6.1 (b)) tilted by 15° to avoid standing waves between the windows. The longitudinal dipole beam enters at (E) and is a normal flat AR coated window.* A cross section of the science chamber (along the internally reflected lattice axis) is shown in fig. 6.1 (b). Here the top viewport, with its 15° tilt, is also visible.

The old high-resolution viewport was flawed in two ways. As discussed in sec. 4.3 it was more than 2 mm too high, which is not a problem for the high-resolution imaging itself but limited the achievable NA of the objectives used for magnetometry. The other flaw was more serious. The optical coating on the viewport is in fact quite complex. It is designed to be reflective in the near-infrared (from 700–1100 nm) apart from two transmission windows at $\Delta\lambda_1 = 760\text{--}850\text{ nm}$ and $\Delta\lambda_2 = 930\text{--}1020\text{ nm}$ (see appendix D of ref. [199]). In this way the viewport is transmissive at 780 nm (for the fluorescence light from the atoms), and at 940 nm (for creation of arbitrary potentials). It is highly reflecting

*As seen in fig. 6.5 the window itself is a bit further away, as there is a T-piece connector that enables access for a vacuum pump.

at 912 nm (for superlattice purposes), and at 1064 nm (for internal reflection of the vertical lattice beam).

As the high-resolution objective has very high NA, it can accept light from high angles of incidence. The issue with the old coating was that the transmission windows only worked well at 0° angle of incidence, severely limiting the achievable NA. This had to be remedied with a new viewport and a new coating. The new piece (MPF, CUSTOM DESIGN) was also 2.4 mm lower than the old one so the surface of the viewport is now about 7 mm below the centre of the vacuum chamber. In other respects it is identical to the old one. It has a glass window (CORNING 7980, GRADE A) with a diameter of about 56 mm, and a thickness of 5.923 mm. Its flatness is greater than $\lambda/4$, good enough for diffraction-limited performance.

6.2 Microwave system and quadrupole coils

To perform spin addressing, a source for microwave radiation must be present in the science chamber. As microwave evaporation is routinely performed in the cube chamber there was a system at hand we could build upon. The microwave source is a MARCONI 2024 controlled via a GPIB interface. The source is operated in the vicinity of 980 MHz. To reach the 6.8 GHz required for driving transitions between the $F = 1$ and $F = 2$ hyperfine states, the signal is sent through a diode (HEROTEK, GC1000RC) that creates harmonics of the main frequency. With low-pass (MINI-CIRCUITS, VLF-6400+, DC–6.4 GHz) and high-pass (MINI-CIRCUITS, VHF-7150+, 7.8–11 GHz) filters we select out the 7th harmonic. Afterwards the signal is amplified in several stages before it is guided to the atoms.

The spin addressing scheme requires a greater degree of control over the microwave radiation than the current system has to offer. To meet the demands the old system was upgraded with a versatile RF signal generator (PHOTONICS TECHNOLOGIES, VFG-150). The RF signal from the VFG is mixed (MINI-CIRCUITS, ZMX-7GR) with the microwave signal from the Marconi to pull it

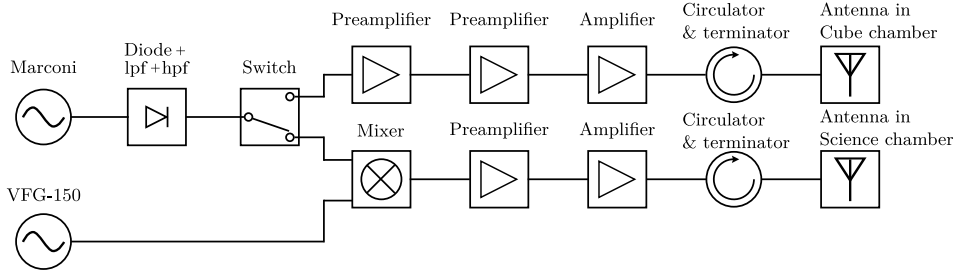


Figure 6.2: A schematic setup of the microwave system used in the cube and the science chamber. Figure borrowed from [178].

into the right frequency range. The schematic of the microwave setup is shown in fig. 6.2. With a switch (MINI-CIRCUITS, ZFSWA2-63DR+) we select which branch (the cube or the science chamber) the signal is sent to. In both stages the signal goes through pre-amplification (MINI-CIRCUITS, ZVE-8G+)[†] and amplification (KUHNE ELECTRONIC, KU PA 640720-10A) delivering more than 1.5 W of power to the antenna. Due to imperfect impedance matching with the antennas, a big portion of the wave is reflected. To protect the final amplifier, its outgoing port is protected with a circulator (MTC, H119FFF) guiding the reflected wave to the terminator.

One of the central goals of the new machine is to achieve the ability to load atoms into a single plane of an optical lattice. This is quite challenging as the spacing between the nodes in our lattices is only 532 nm. A discussion of our plans to that end are included in the outlook. One of those options is the method of magnetic field slicing [280]. For that to be viable we built in a pair of gradient coils of the same design as for the coils attached to the cube chamber. Each coil has 4×4 windings of a square shaped copper wire (obtained from DAN-FYSIK with a side length of 2 mm) with a hollow core enabling efficient water cooling (for details on the coils and the water cooling see sec. 2.1.3 and app. A of ref. [199]). The coils are cast in epoxy and mounted directly on the flanges

[†]On the cube branch there are two pre-amplifiers, first a MINI-CIRCUITS, ZX60-8008E-S+ and the second mentioned in the text.

on the top and bottom of the chamber. Due to the design of the system, the bottom coils had to be built in at the same time as the high-resolution objective.

6.3 Optical systems for the science chamber

The high-resolution experiment required new optical setups. Most notably we built in the high-resolution breadboard (covered in the next section), a 3D optical lattice system, and optical molasses around the science chamber. A schematic picture of the laser setup in the horizontal plane around the chamber is shown in fig. 6.3.

The crossed dipole trap is depicted in green and is similar to the earlier setup. The transverse dipole trap had to be redesigned as the dual objective system was removed, but the present setup still provides a beam with a waist of about $90\text{ }\mu\text{m}$. The last mirror of the TDT, and the second to last of the LDT are mounted in a controllable piezo driven mirror mounts (RADIANT DYES, N-470 PIEZOMIKE), such that day-to-day adjustments of the position of the CDT can be done without opening the compartments around the experiment. We do especially see drifts of the LDT beam (a few μm per day), which is possibly due to both its very long beam path ($\sim 3\text{ m}$), and some slight hysteresis in the transport stage (see sec. 3.1). The stability of the whole setup is greatly improved when the climate system is operating (see sec. 6.5).

The horizontal lattice beams are depicted in blue. The main difference between the horizontal 1 (H_1) and horizontal 2 (H_2) lattice axes (see fig. 6.3), is that H_1 is externally reflected and H_2 is internally reflected on a re-entrant viewport (D in fig. 6.1 (a)). The distance from the centre of the chamber to the surface of the re-entrant viewport is 30 mm , which constrains the achievable waist of the H_2 lattice beam. Naturally the focus of the beam must be placed on the viewport itself leaving the size at the atoms up to Gaussian beam propagation. By minimising the beam size at a distance of 30 mm from the focus using eq. (1.16), one can obtain a waist of about $150\text{ }\mu\text{m}$ at the atoms, given that one has a $100\text{ }\mu\text{m}$ beam at the viewport. As a result both lattice beams are designed

to give 100 μm waists where the focus of H_1 is placed at the atoms, and the focus of H_2 placed at the re-entrant viewport. The beam sizes were confirmed by beam profiling.[‡] The size of the H_2 beam 30 mm from the focus was measured to be 140 μm . Both beams have pure linear polarisation that is cleaned by a polarising beam splitter cube after the last mirror before it enters the chamber.

The light for the horizontal lattice axes is provided by a 50 W fibre amplifier at 1064 nm wavelength (AZUR LIGHT SYSTEMS). The light distribution is done on a separate optical breadboard (shown in appendix B) and the light is passed through photonic crystal fibres (NKT PHOTONICS, AEROGUIDE-POWER) to the setup. The H_1 axis is externally reflected by a specially coated dichroic (see appendix C of ref. [199]) that is highly transmissive for 780 nm light and reflective at 1064 nm at a 0° angle of incidence. The retro-reflection of all lattice axes is achieved with the reflective element in a cat's eye configuration, where the beam is focussed onto the retro-reflector. This configuration enhances the stability of the setup and minimises distortions of the reflected beam due to any surface deviations on the reflector.

Absorption images of the cloud can be acquired on both lattice axes (path shown in orange in fig. 6.3). This provides a good starting point for the alignment of the optical lattices as the lattice beams are also visible on the cameras. The alignment of the H_1 axis is considerably easier than that of the H_2 axis. For the purpose of the alignment of the H_1 we have installed a flip mirror before the retro-dichroic. This mirror sends the beam to a dump, turning our lattice into a single-beam dipole trap. To be sure that the beam hits the cloud centrally, the CDT is turned off and the H_1 beam is pulsed on at high power for ~ 1 ms, followed by a TOF measurement. If the beam hits the cloud on the side, the

[‡]When aligning the lattices we routinely measured the beam size of the laser at various points in the setup. In the beginning we often did those measurements only with the laser in seed mode or with the fibre amplifier running at low current, simply because more precautions have to be taken when operating at high power. This turned out to give incorrect results. The profile of the laser is radically different at lower power mainly because a relatively big portion of the seed light propagates in the cladding of the fibre. *Always measure the properties of the beam under the settings in which it will be operated.*

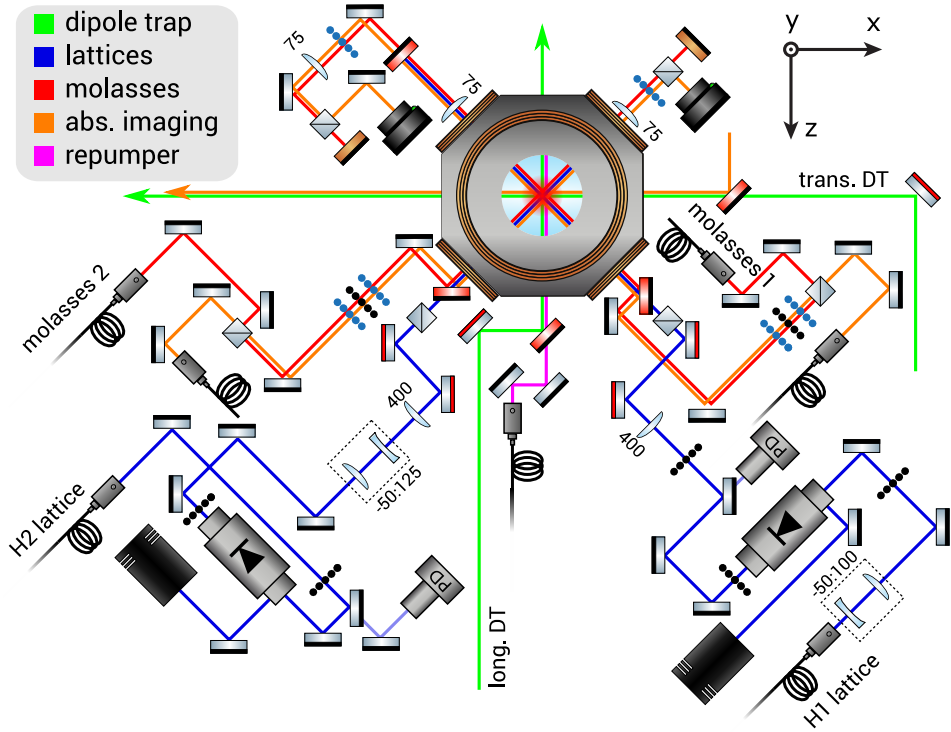


Figure 6.3: The optical setup around the science chamber. Lens sizes are given in mm and a legend for the optical elements is available in appendix B. Details are given in the text. Figure borrowed from [178] and edited.

cloud will accelerate towards the region of highest intensity. By scanning the beam position across the cloud one obtains a dispersion-like curve in the cloud position and can deduce the centre of the beam. The retro-reflected beam is overlapped to an iris in the beam path (before the optical isolator) and the lattice depth is maximised either by Kapitza-Dirac measurements or by modulation spectroscopy (both methods are explained in ch. 7). For the H2 axis we rely on the camera for initial positioning, but as the retro-reflector is fixed, the beam must be walked. In the end that must be done by using the atoms as a measure




Figure 6.4: The optical setup for the vertical lattice. Lens sizes are given in mm and a legend for the optical elements is available in appendix B. Details are given in the text. Figure borrowed from [178] and edited.

of the lattice depth, which is considerably more cumbersome than for the H_1 . In our experience the H_2 axis was generally more stable than the H_1 axis, but because of the limitation of the beam size it requires twice the optical power to achieve the same depth.

The beam shaping optics for the vertical lattice beam are placed on a bread-board mounted above the science chamber. The path is shown in fig. 6.4. Most notably it is designed to provide a narrow waist of $55\text{ }\mu\text{m}$ at its focus, and this corresponds to a waist of about $60\text{ }\mu\text{m}$ at a distance of $3.5\text{--}4.0\text{ mm}$ above the viewport where the atomic cloud is held in the CDT. As the last lens in the optical setup is mounted inside the top viewport, we never confirmed the size of the beam directly. The choice for the narrow waist was mainly because of power considerations. The light is provided by our Nufern amplifier that also gives light to the CDT. In fact the 0th order (un-diffracted) beam emerging from the AOM for the TDT is used for the vertical lattice. For details the reader may study the schematics of the high power table given in fig. B.3 (bottom) of appendix B. This axis is aligned in the same way as the H_2 . For further details on the particularities of our lattice setup, and lessons learned, see appendix D.

The beams for the optical molasses in the horizontal direction are depicted

in red in fig. 6.3. The optics are designed to provide beams with a waist of around 0.7 mm and estimations using the lattice cameras reveal them to be about 0.5 mm on both axes at the position of the atoms. The paths feature a combination of waveplates (quarter-half-quarter), with which one can produce any given polarisation of the laser beam. They are adjusted to give beams with circular polarisation. The horizontal molasses beams are retro-reflected by mirrors glued on top of ring-shaped piezo actuators (PIEZOMECHANIK, HPST150/4-10/12). By modulating them with an amplitude of several λ at a frequency much greater than the exposure time of a fluorescence image, one avoids that standing waves in the molasses light appear on the images (see fig. 4.11 of ref. [280]).

The light collected by the objective from the atoms is the scattered molasses light. As a result we cannot shine a molasses beam in from above, as that would completely dominate the faint fluorescence light from the atoms. To provide some cooling in the vertical direction, it turns out to be enough to pass a single molasses beam through the objective from below (its path is sketched in blue in fig. 6.5 (c)). This so-called third molasses beam is overlapped with the imaging path via an uncoated glass plate that ensures about 5% reflection of the incoming light. A portion of the light is unavoidably back-reflected into the imaging system giving rise to speckle patterns on the camera. This is minimised by walking the beam such that the main portion of the retro-reflection misses the imaging system. To attain a reasonable waist (one that covers the whole cloud) the optical system must focus the beam right before the objective. Our estimates were that in this way we could get a beam with a waist of 100 μm at the atoms, but this turned out not to be the case. The beam only illuminated a portion of the cloud and suffered heavily from fringing. However, it is still necessary for cooling as is discussed in sec. 9.4.

The molasses light comes directly from the master laser and is brought near resonance by AOMs in double pass configuration. A schematic of the distribution system is given in fig. B.3 (top). A necessary repumper beam is shone into the chamber through the same viewport as the LDT beam (magenta in fig. 6.3). Its size is similar to that of the molasses beams.

6.4 The high-resolution optical breadboard

The centrepiece in the machinery of our quantum gas microscope is the optical breadboard that hosts the high-resolution objective (d. *den gule pølse*) and other high-resolution optics. The objective (ASE OPTICS, CUSTOM DESIGN) has an NA of 0.69, an effective focal length of 5 mm and a working distance of 12.95 mm. It is designed to perform at the diffraction limit for 780–790 nm light without any chromatic shifts. For details see sec. 7.3.2 of ref. [199]. It is mounted on a piezo-driven objective scanner (PI, PI-FOC P-725.4CA) with a scan range of 400 μm which enables control of its vertical position to nm precision. The objective and the scanner are shown in fig. 6.5 (a).

The alignment of the objective to the window of the high-resolution viewport is highly critical as any residual tilt will lead to an aberrated image. Therefore the setup must allow for some means of moving and tilting the objective irrespective of the viewport. To achieve this without affecting the relative alignment of the objective with respect to other optics the whole breadboard is mobile. A schematic top view of the breadboard is given in fig. 6.5 (b). The board is laid on top of a metal block that is identical in shape to it and three micrometer screws (NEWPORT, BM30.10) are the sole points of contact between the breadboard and the bottom plate. Their positions are marked by red circles in the sketch. To enable high precision positioning on top of the experimental table, micrometer screws can be placed around it as marked by the green arrows in the sketch. The yellow dot marks the position of the objective. The breadboard is custom made by the Institute's mechanical workshop, machined out of aluminium. Both plates are 4 cm thick and given the dimensions from the sketch, the whole construction with all optics weighs a few hundred kilos.

There are four optical systems on the breadboard and they are all sketched in fig. 6.5 (c). The sketch is simplified and marks only how the different systems are combined. A detailed description of all the separate optical paths and alignment strategies can be found in ref. [130], but below the most important aspects will be discussed. The imaging path is marked in orange and together with a 750 mm achromatic doublet the system has a magnification around $M = 152.3(9)$ (see

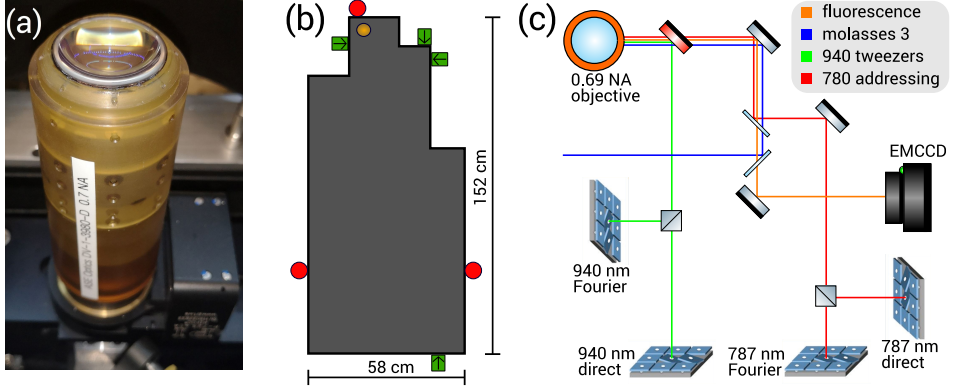


Figure 6.5: The high-resolution optics. (a) An image of the objective mounted to the objective scanner. (b) A sketch of the breadboard itself and the means to move it. (c) A coarse schematic setup of the optics on the hires board. Details are given in the text.

measurements in sec. 9.3). As a result a single camera pixel corresponds to $d_{\text{pix}} = 105.5(2)$ nm and so the field of view of the camera spans a region of $54 \times 54 \mu\text{m}^2$. As the periodicity of the lattice is 532 nm, a single atom covers an area of 5×5 pixels. The molasses beam from below, discussed in the previous section, is marked in blue.

In addition there are two optical systems capable of projecting light onto the atoms spatially shaped by DMDs. One is operated with light at a 940 nm wavelength intended for the creation of far off-resonant potentials and the other uses near-resonant light at a wavelength close to 787.5 nm. That light is used for local spin-addressing that has been realised previously both in a quantum gas microscope [281] and in a large spacing 3D optical lattice [277]. At this wavelength the light shifts generated by σ_+ polarised laser beams, with respect to the D1 and D2 lines, cancel each other for the $|F = 1, m_F = 1\rangle$ state, but there is still a significant light shift for the $|F = 2, m_F = 2\rangle$ state. This shift can be used to pull individual atoms locally into resonance with microwave radiation used to transfer atoms between the two spin states.

Both of those systems have two paths. On one arm the light is shaped by

a DMD in a direct imaging configuration, and on the other arm the DMD is placed in the Fourier plane of the imaging system. These arms are combined by a non-polarising beam splitter cube. Both methods of imaging have their pros and cons. The direct imaging method is conceptually simpler and is preferable for big structures as the intensity that can be obtained is independent of the structure. However, any residual optical aberrations inherent in the system will affect the pattern. When the DMD is placed in the Fourier plane of the imaging system one gains control both over the amplitude and the phase of the light. Here one obtains a way of correcting for residual aberrations. The drawback is that large and complex patterns will become power limited, as they require high-frequency components in Fourier space that effectively requires one to turn off a large fraction of the mirrors on the DMD.

As explained in ref. [130] much work has been done in the group on mapping out the aberrations efficiently. Briefly, to project a certain image onto the atoms one would naively upload its Fourier transform to the DMD, but as it turns out this will not yield a good result. The generated image will be highly aberrated, mostly due to the curvature of the DMD itself, which amounts to several λ across the whole chip. These aberrations (and others) can be measured and corrected for by an interferometric measurement, where the phase of the light across the chip is mapped out with respect to its centre. In this procedure one acquires a *phasemap* that can be used to correct the Fourier transform of the desired image. This procedure of measuring the phasemap can be time consuming. However, as is explained in ref. [130], this time can be cut down to about half an hour. Until now, we have not reached the point of using the Fourier DMDs in the experiment and so their function won't be detailed any further. Ref. [294] offers a fine account of Fourier imaging with DMDs.

Each of the direct imaging DMDs can be illuminated either with a large (more uniform, less intense) or a small (less uniform, more intense) beam. The two beams are shone in from different sides to the DMD. As explained in sec. 3.3 the mirrors can be tilted $\pm 12^\circ$. Therefore either orientation can be used as the on or off state, depending on the direction of the incoming beam. One can project the same onto the atoms simply by inverting the image uploaded to the

DMD (black \leftrightarrow white). A smaller beam (with higher intensity) leads to deeper traps, at the cost of a smaller usable area on the DMD. The laser systems that provide light to the 787 and 940 nm setups are depicted in fig. B.2 of appendix B.

The whole breadboard with DMDs was pre-aligned and characterised before instalment to the system [130]. The final alignment of the imaging system had to be done in-situ with the correct viewport before the chamber was closed and baked out. A specially made test plate was placed on top of the viewport and imaged (with 780 nm light) to guide the alignment of the system. The plate was made of glass, coated with a thin gold layer on one side. In the gold, 100 nm holes had been etched out constituting tiny pinholes that were well below the resolution limit of the imaging system. The positions and tilts of the board were adjusted by the micrometer screws. Having obtained a satisfying result the board was removed again as it could not be in place for the bakeout. Afterwards it was placed in again and the alignment could be reproduced by resetting the micrometer screws to their optimal positions. This procedure is described in details in ch. 7 of ref. [130].

As a final comment let us examine the 940 nm direct system, which we have used the most, in more detail. The demagnification to the image plane is $1/M_{940} = 92.7(5)$. The other systems have a similar demagnification. This means that each DMD pixel corresponds to about 80 nm at the position of the atoms. To make tight tweezers we have routinely used circles with a radius of 7 DMD pixels (R7) that yields tweezers with a waist of 780 nm.[§] This is limited by diffraction (and some residual aberrations in the system) so smaller patches will only result in lower intensity.

Before the DMD we have maximally about 500 mW of optical power. Out of that we are able to regulate 21 μ W in a single R7 tweezer, as measured in the

[§]The term Rn describes a tweezer with a radius of n DMD pixels. A single DMD pixel corresponds to about 80 nm in the image plane. Although it might seem excessive to include the size of the spot used on the DMD when the actual waist can be quoted, one must realise that the correspondence between the object size and the image size is not 1 : 1, as ultimately the image will be limited by diffraction. Thus this number is quoted for completeness, where it is applicable.

intermediate image plane. The efficiency of the optical system from the intermediate image plane to the atom plane is about 45%. The tweezers have a waist of $w_{940} = 780$ nm (see sec. 8.2), corresponding to tweezer potentials of 1.5–2.0 μ K maximum depth. Using the small outcoupler we should be able to gain a factor of five in the depth.

6.5 The climate system

Drifts of ambient temperatures on the timescale of hours are an inevitable part of life on the surface of the pale blue dot. The in-house climate system provided by the University maintains the temperature in the laboratory between 22–24°C, depending on outside weather conditions. On very hot summer days[‡] lab temperatures rise some degrees above that level. Such drifts cause the pointing of laser beams to change by some μ m as most opto-mechanical components are made of metal that is prone to thermal expansion and contraction. The optical lattices and the high-resolution optics require a long term stability in the position of atoms and lasers on the level of 1 μ m.

To achieve stable temperatures on the experiment table we designed and built a dedicated climate system for precise regulation of the temperature down to the level of 0.1°C. The principle of the system is simple and relies on the fact that its easy and fast to heat things and hard and slow to cool. The system consists of two stages: a cooling stage that brings the temperature of the air to T_c a few degrees below the desired temperature T_d and a heating stage that heats the air to T_d . Any drifts in T_c are easily compensated by the heating element.

The schematic is shown in fig. 6.6 (a). A fan-coil unit (AERMEC TDA) sucks in air from the lab. The University provides cooling water typically around 13°C, which serves as the coolant in the fan coil unit. A controllable valve (SIEMENS SAS61) mounted to the inlet line of cooling water steers the temperature at this stage. After the fan coil the air is blown through a resistive heating element (SYSTEMAIR CB 250-3), from where the air is guided into the experimental compart-

[‡]Which are blessedly few here in Denmark!

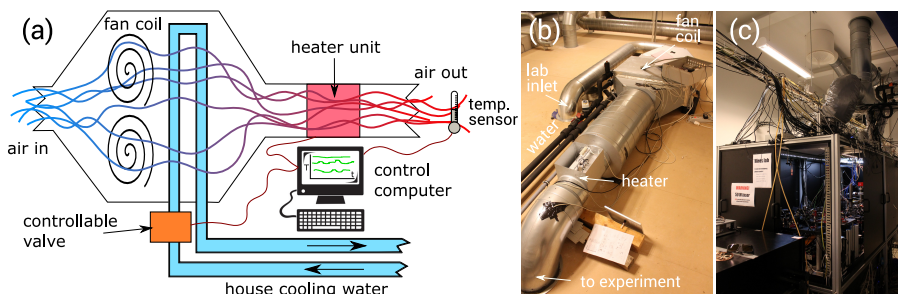


Figure 6.6: The climate system. (a) A schematic figure of the climate system. (b) An image of the cooling and heating units. (c) The experimental compartments. The outlet from the climate system is guided through the ceiling in the lab into two separate HEPA filters that rest on top of the experiment cage.

ments. The air from the climate system flows through HEPA filters (VENFILTER BTSR244805/250H14) that filters dust and ensures homogeneous and laminar air flow onto the experiment table, which is closed off in a box with sliding doors (that can easily be unmounted for access) and a cover on the top.

The central parts of the climate system (fan-coil and heating unit) sit on top of the ceiling of the laboratory that is accessible from outside. In fig. 6.6 (b) these parts are shown. The black patch on the tube that goes towards the experiment table marks the position of the NTC temperature sensor (TEXAS INSTRUMENTS LMT70) used for the regulation of the temperature. The air is guided down to the lab through the tube extending from the ceiling towards the experiment in fig. 6.6 (c). The temperature sensors are read out by a NATIONAL INSTRUMENTS myRIO device that runs LabView. Software PID loops running on the micro-controller regulate the pre-cooling of air via the adjustable valve, as well as the resistive heating element that warms up the air to the desired point.

To estimate the quality of the regulation data are evaluated for a period of 1 week. This particular week was not a busy lab week, but at the time of writing there was no access to data where the experiment was running continuously. The temperature on the experiment table during that week ranged from 20.3°C to 20.4°C, meeting our demands of 0.1°C stability.

CHAPTER 7

Optical lattices

*Ef gárur ganga dal í dal
gjarnan lýsist rúmið,
en falli þar sem fjall í sal
fýkur yfir húmið.*

An optical lattice is an optical potential created by two counter-propagating laser beams. As the light is coherent it will interfere to create a standing wave along the beam with a periodicity of $d = \lambda/2$. We use 1064 nm light resulting in a lattice period of $d = 532$ nm, and as the laser is red-detuned with respect to the atomic transitions the high intensity regions form an attractive potential. Along a single spatial dimension x , the potential is

$$V(x) = V_0 \cos^2(kx), \quad (7.1)$$

where $k = \frac{\lambda}{2\pi}$ is the wave vector and V_0 the depth of the lattice which can be calculated from eqns. (1.14), (1.15) and (1.22). At the origin, $V_0 = \frac{8P}{\pi w_0^2} U_0$ where P is the optical power of the beam, w_0 is its waist and U_0 is the atom-light coupling




Figure 7.1: A 1D optical lattice. (*top*) Two focussed counter-propagating laser beams interfere and a standing wave is created along the beam. (*bottom*) For a red-detuned laser an attractive potential is created at high intensity regions. The Gaussian nature of the focussed beam affects the depth of the potential wells along the beam.

constant of eq. (1.14). Typically the Gaussian shape of the interfering beams has to be considered, as is evident from eq. (1.22). This feature is illustrated in fig. 7.1, and as a result the depth of the individual potential wells varies along the beam. However, as the Rayleigh range of our lattices (~ 30 mm) is long compared to the periodicity, the system can typically be considered as homogeneous over relevant length scales and eq. (7.1) is a good approximation.

Periodic potentials are ubiquitous in nature as all crystals are essentially periodic structures of atoms on a lattice. The analogy between electrons travelling in a crystal of ions and neutral atoms in optical lattices is apparent. It is however much easier to control individual atoms in optical lattices than single electrons in solid state materials, which renders atoms in lattices a perfect test bed for theories in solid state physics.

The depth of an optical lattice is often cited in terms of the recoil energy of the trapping light, that is half of the energy an atom obtains from scattering a lattice photon. In our case, such events are rare as the light we use is very far off-resonant from the D-line doublet in ^{87}Rb . To relate it to other units of energy,

one lattice recoil at 1064 nm amounts to*

$$E_r = \frac{(\hbar k)^2}{2m} = \frac{h^2}{2m\lambda^2} = 97.3 \text{ nK} \times k_B = 2.03 \text{ kHz} \times h. \quad (7.2)$$

The purpose of our optical lattices is twofold: first, they provide a periodic potential landscape for physics investigations, and second, they provide a tightly confining deep potential for site-resolved imaging of atoms. A BEC immersed in an optical lattice exhibits *superfluid* properties at low lattice depths $< 10 E_r$, where the location of each atom is ill defined but their phases are in sync. When ramped deep the atoms become localised to a given site of the lattice potential and the phase coherence vanishes. This is the *Mott insulating* state and all remnants of superfluidity vanishes around $20 E_r$. For the purpose of imaging single atoms (see ch. 9) the lattices are ramped *deep* to 1000's of E_r . This dual purpose of the optical lattices where several Watts of laser light are required to reach the imaging regime, and only some 10's of mW for the physics regime, makes the technical implementation challenging.

This chapter discusses different aspects of our optical lattice setup and experiments we have performed to characterise its performance. The first section contains a theoretical account of how band structure arises for atoms in optical lattices. The second and third sections cover the different ways we use to calibrate the depth of our optical lattice. The fourth section is devoted to our observation of the superfluid to Mott insulator phase transition. The fifth and final section delves into heating issues caused by Brillouin scattering in optical fibres.

*Please note that there is a slight ambiguity in definitions here. The recoil temperature of ^{87}Rb is quoted as $T_r = 362 \text{ nK}$ [252], corresponding to *two* units of the E_r at 780 nm, whereas in the optical lattice literature [118] the depth is quoted in units of *one* lattice recoil. In the case of a scattering event between the atom and a lattice photon it still gains *two* units of recoil energy [187].

7.1 Band structure in periodic potentials

Atoms in optical lattices are quantum mechanical particles in periodic potentials.[†] The energy of such a particle is readily obtained from the time-independent Schrödinger equation, $\mathcal{H}|\psi\rangle = E|\psi\rangle$. We assume an infinite lattice, and so ignore the harmonic confinement of the potential. This is a fine approximation for our BECs that span a few 10's of μm whereas the periodicity of the lattice is only $0.5\ \mu\text{m}$ [72]. The Hamiltonian has a kinetic and a potential term, and for a 1D lattice of a depth V_0 Schrödinger's equation becomes,

$$\left(\frac{\hat{p}^2}{2m} + V_0 \cos^2(kx) \right) |\psi\rangle = E |\psi\rangle. \quad (7.3)$$

In this situation Bloch's theorem for quantum waves in periodic potentials applies (see ch. 8 in ref. [17]). The state vector $|\psi\rangle$ is decomposed into a product of a plane wave and a function $|u\rangle$ that is modulated at the frequency of the lattice. Each function is labelled by the band number (the energy level) n and the number q , the so-called *quasimomentum* of the lattice. The Bloch functions are

$$|\psi_{n,q}\rangle = e^{iqx} |u_{n,q}\rangle. \quad (7.4)$$

The naming of q is apparent if we substitute the Bloch ansatz into the Schrödinger equation, arriving at (mainly by applying the momentum operator $\hat{p}^2 = -\hbar^2 \frac{\partial^2}{\partial x^2}$)

$$\left(\frac{(\hat{p} + q)^2}{2m} + V_0 \cos^2(kx) \right) |u_{n,q}\rangle = E_{n,q} |u_{n,q}\rangle. \quad (7.5)$$

As $|u\rangle$ is periodic in the lattice potential it may be written as a Fourier sum of plane waves

$$|u_{n,q}\rangle = \sum_{\alpha} c_{\alpha,n,q} |\alpha\rangle, \quad (7.6)$$

[†]The derivation in this section is based on refs. [117] and [233].

where $\alpha \in \mathbb{Z}$ and $|\alpha\rangle = e^{2ikx\alpha}$.[‡] In addition the plane waves $|\alpha\rangle$ form a closed basis. The energies are the matrix elements of the Hamiltonian $\langle u' | \mathcal{H} | u \rangle$, and by focusing on one term at a time, we obtain for the kinetic part

$$\left\langle u' \left| \frac{(\hat{p} + q)^2}{2m} \right| u \right\rangle = \sum_{\alpha'} \sum_{\alpha} c_{\alpha'}^* c_{\alpha} \frac{\hbar^2}{2m} (2k\alpha + q)^2 \langle \alpha' | \alpha \rangle, \quad (7.7)$$

and the potential part

$$\begin{aligned} \langle u' | V_0 \cos^2(kx) | u \rangle &= \frac{V_0}{2} \sum_{\alpha'} \sum_{\alpha} c_{\alpha'}^* c_{\alpha} (e^{2ikx} \langle \alpha' | \alpha \rangle + e^{-2ikx} \langle \alpha' | \alpha \rangle) \\ &= \frac{V_0}{2} \sum_{\alpha'} \sum_{\alpha} c_{\alpha'}^* c_{\alpha+1} \langle \alpha' | \alpha + 1 \rangle + c_{\alpha'}^* c_{\alpha-1} \langle \alpha' | \alpha - 1 \rangle. \end{aligned} \quad (7.8)$$

(7.9)

As a result, the Hamiltonian only has non-zero kinetic terms on the diagonal and the potential terms on the off-diagonal

$$\mathcal{H}_{\alpha\alpha'} = \begin{cases} \frac{\hbar^2}{2m} (2k\alpha + q)^2, & \text{if } \alpha = \alpha' \\ \frac{V_0}{2}, & \text{if } \alpha = \alpha \pm 1. \end{cases} \quad (7.10)$$

To solve the system numerically, the sum must be truncated at some reasonable value of α . For rendering fig. 7.2 we have $-10 \leq \alpha \leq 10$, and we plot the dispersion relation in the first Brillouin zone ($-\hbar k \leq q \leq \hbar k$). The green dotted line in each panel signifies the lattice depth. Odd bands are coloured blue, and even bands are orange. In (a) there is no lattice present and the dispersion relation is parabolic ($E \sim k^2$), identical that of a free particle. In the presence of a lattice a band gap opens up between the first and the second band, and a depth

[‡]The exponent $2kx$ comes about as $\cos^2(kx) = \frac{1}{2}(\cos(2kx) + 1)$. For clarity the coefficients $c_{\alpha,n,q}$ will only be written with the subscript relevant for the given calculation. For the remainder of the section we write $c_{\alpha,n,q} = c_{\alpha}$.

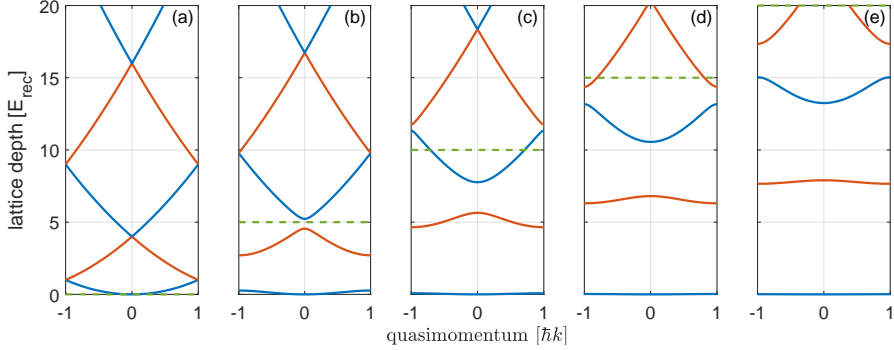


Figure 7.2: Band structure in a simple sinusoidal 1D lattice. The depth of the lattice is varied from 0 to $20 E_r$ from left (a) to right (e) in steps of $5 E_r$. In the absence of a lattice (a), the dispersion is that of a free particle and with increasing depth, band gaps open up on the edges $|\hbar k| = 1$ and in the centres $|\hbar k| = 0$ of the first Brillouin zone and the bands flatten out. The structure approaches that of a harmonic oscillator.

of $5 E_r$ is sufficient to open up a gap between the second and third band as seen in (b). The effect of an increasing lattice depth flattens out the bands and as the depth becomes sufficient the structure approaches that of a harmonic oscillator, with an even energy spacing of $\hbar\omega$ between each band, where ω is the (angular) trap frequency. This is also supported by the intuition that as the potential becomes deep, it can readily be approximated with a harmonic potential around the bottom of each lattice site.

7.2 Calibrating lattices: Kapitza-Dirac scattering

In 1933 P. L. Kapitza and P. A. M. Dirac proposed an elegant experiment that could demonstrate particle-wave duality in a clear manner. In this experiment a beam of electrons should traverse a standing wave of light, and be diffracted by it. The effect was first observed with atoms in 1986 [116] but the observation of the original proposal was only achieved in 2001 [96]. The effect is identical to the diffraction of laser light by a grating, where the roles of the matter and the

laser are interchanged. A BEC—a coherent matter-wave—can be diffracted by a pulsed optical lattice, that acts as the grating [211].

As the depth of the lattice affects the dynamics of the BEC, the Kapitza-Dirac effect can be used to calibrate it. This is the method we routinely use in the lab to monitor the lattice depth on a day to day basis. It works primarily in the low-power regime, as the dynamics become very complex for depths $> 20 E_r$ [8].

Let us imagine that a BEC is abruptly exposed to an optical lattice potential for a short period of time Δt . In the previous section we saw that the wavefunction of an atom in a lattice can be naturally expanded in Bloch waves (planewaves with different momenta),[§]

$$|\psi_{n,q}\rangle = \sum_{\alpha} e^{ix(2k\alpha+q)} c_{\alpha,n}. \quad (7.11)$$

As the momentum distribution of a BEC is very narrow we take it for a plane wave of the lowest order $|\phi_{\text{BEC}}\rangle = e^{iqx} |0\rangle$ [72]. Typically the BEC possesses no net momentum relative to the lattice, but in the following calculation we allow for $q \neq 0$. By non-adiabatically exposing the BEC to the lattice it is projected onto the Bloch states. We get that

$$|\Phi(t=0)\rangle = \sum_n |\psi_{n,q}\rangle \langle\psi_{n,q}|\phi_{\text{BEC}}\rangle \quad (7.12)$$

$$= \sum_n c_{0,n}^* e^{-iqx} |\psi_{n,q}\rangle. \quad (7.13)$$

During the time where the lattice is present each component gains a phase factor from the time-evolution operator,

$$|\Phi(t)\rangle = \sum_n c_{0,n}^* e^{-iqx} |\psi_{n,q}\rangle e^{-i\frac{E_{n,q}}{\hbar}t}. \quad (7.14)$$

[§]In the following calculation the subscript n is included on the coefficients $c_{\alpha,n}$.

After the lattice is turned off the state is projected onto the plane wave basis and the different components of that state are

$$|\Phi'(\Delta t)\rangle = \sum_{\alpha} |\alpha\rangle \langle\alpha|\Phi(\Delta t)\rangle \quad (7.15)$$

$$= \sum_{\alpha} |\alpha\rangle \sum_n c_{0,n}^* e^{-iqx} \langle\alpha|\psi_{n,q}\rangle e^{-i\frac{E_{n,q}}{\hbar}\Delta t} \quad (7.16)$$

$$= \sum_{\alpha} |\alpha\rangle \sum_n c_{0,n}^* c_{\alpha,n} e^{-i\frac{E_{n,q}}{\hbar}\Delta t}. \quad (7.17)$$

As an effect the BEC is projected into n energy bands with some components also mapped into separate momentum components (multiple of $2\hbar k$) depending on the lattice depth of the pulse. If $q = 0$, only even energy bands are populated due to symmetry. After the lattice is turned off the populations in the different bands will interfere within the different momentum orders. As an example, in the $+2\hbar k$ order there will be a population in bands 0 and 2 that give rise to an oscillatory signal according to eq. (7.17) [8]. We typically work in the weak pulse regime where we only need to take into account population in bands 0 and 2 which exhibit Rabi-like oscillations. The contribution from a population in the band 4 only becomes significant $\gtrsim 15 E_r$ (see for example fig. 6.5 (b) in ref. [8]).

In our case, a typical experimental sequence for a Kapitza-Direc measurement involves the creation of a pure BEC in the CDT by evaporation followed by a short pulse of the lattice. We achieve this by bypassing the intensity regulation of the laser power (PID controller) as its response is very slow, and adjust the lattice power by attenuating the output of the AOM driver itself with RF attenuators. The lattice signal is controlled with a TTL signal providing a true square pulse. The cloud is subsequently imaged in TOF.

An example of such a calibration measurement is depicted in fig. 7.3. Panel (a) shows a concatenated series of images obtained after 15 ms TOF, where the BEC was exposed to a $13 E_r$ deep potential for a time ranging from 0 to 150 μs . The data has been processed in panel (b) by determining the atom number

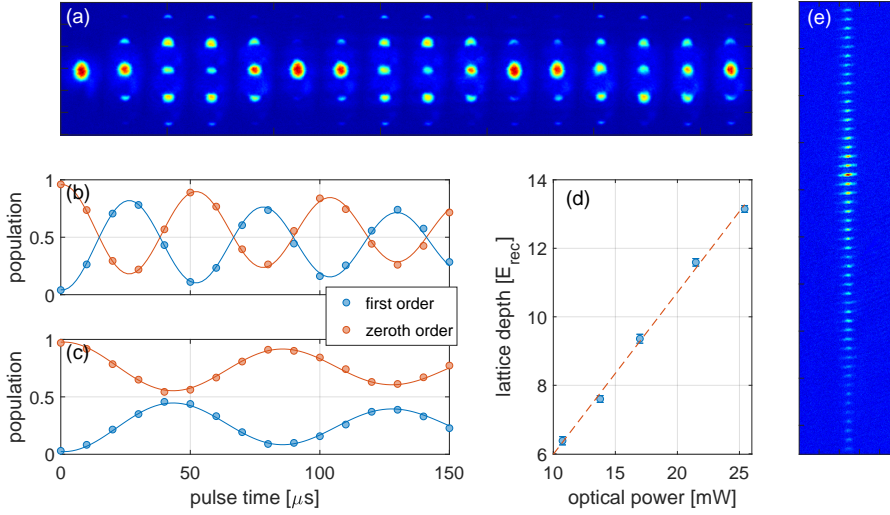


Figure 7.3: Diffracting matter with light. Panel (a) is a compilation of absorption images of a BEC diffracted by the vertical lattice beam, taken after 15 ms TOF. In the series the pulse length is scanned from 0 to 150 μs from left to right. In (b) the atom number in the centre peak and an adjacent first order has been extracted and normalised to the total atom number. The data are fit with a damped sinusoidal. The data in (c) are obtained with a shallower lattice. These measurements correspond to one point on the calibration graph in panel (d). In (e) a matter diffraction with the lattice at full power is shown. The diffraction orders extend out of the field of view of the camera.

in the centre peak and one of the side peaks (with a momentum of $\pm 2\hbar k$). The populations are normalised as we only require relative populations to measure the depth. They oscillate with a frequency $\omega_{\text{KD}} = (E_{2,0} - E_{0,0})/\hbar$ obtained from a fit of a damped sinusoidal.[¶] The data in panel (c) are taken at a lower lattice depth (around $6 E_r$), yielding a slower oscillation. By repeating these measurements for different depths, one obtains a calibration graph like

[¶]The average frequency of the two is used to calculate the depth from a graph like the one displayed in the inset of fig. 7.4 (a), which shows the calculated transition frequency for a given lattice depth at $q = 0$.

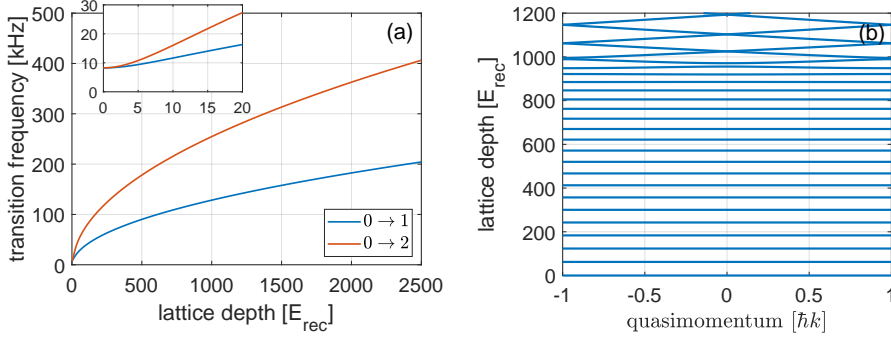


Figure 7.4: Deep optical lattices. (a) A calculation of the transition frequencies between the ground band and first and second bands at $q = 0$ as a function of the lattice depth. The inset is a zoom in on the low depth regime. (b) A band structure calculation for a $1000 E_r$ deep lattice.

the one displayed in panel (d). This particular measurement yielded a slope of $472(17) E_r/W$. It was our experience that this calibration could be extrapolated to much higher depths, something we confirmed with modulation spectroscopy, as explained in the next section. Panel (e) displays a Kapitza-Dirac signal taken with the vertical lattice at full power, far beyond the weak pulse regime. At least 30 momentum orders are visible in the image.

7.3 Calibrating lattices: Modulation spectroscopy

Another standard method for characterising the eigenfrequencies of an optical trap is to modulate its intensity at various frequencies and observe the response of the atom cloud. As trapping frequencies (or multiples of them) are matched, the modulation results in heating of the cloud. In fig. 7.4 (b) the band structure of a $1000 E_r$ deep lattice has been calculated. The even level spacing is obvious so the potential is well approximated by a harmonic one. Adding a sinusoidal

intensity modulation at a frequency ω_m gives a potential

$$V(t) = \frac{1}{2}m\omega^2\hat{x}^2 \cos(\omega_m t) \quad (7.18)$$

where \hat{x} is the position operator. The matrix element coupling states $|i\rangle$ and $|j\rangle$ yields

$$\langle i|V(t)|j\rangle \sim \langle i|\hat{x}^2|j\rangle \quad (7.19)$$

$$\sim \langle i|\hat{a}^2 + \hat{a}^\dagger\hat{a} + \hat{a}\hat{a}^\dagger + \hat{a}^\dagger\hat{a}^\dagger|j\rangle = \begin{cases} 1 & \text{if } |i-j| = 0, \pm 2 \\ 0 & \text{otherwise,} \end{cases} \quad (7.20)$$

where in the second step the expression for \hat{x}^2 in terms of the ladder operators of the harmonic oscillator (see ref. [233], p. 93) was inserted. Thus, an intensity modulated harmonic trapping potential can only couple transitions between levels separated by $2\hbar\omega$. In a more classical picture the modulation excites the breathing mode in the cloud, oscillating at twice the eigenfrequency of the trapping potential.

The transition frequencies for the $0 \rightarrow 1$ and $0 \rightarrow 2$ transitions are plotted in fig. 7.4 (a). The reason for the inclusion of the $0 \rightarrow 1$ process is discussed below. For high lattice depths we obtain that $f \sim \sqrt{V_0}$ as expected by eq. (1.19). The inset is a zoom to the low depth end of the spectrum, and as $V_0 \rightarrow 0$, $f \rightarrow 8.253$ kHz which is the energy between bands 0 and 1 or 0 and 2 in the case of the free particle, see fig. 7.2 (a).

To perform modulation spectroscopy on the atom clouds, they are first adiabatically loaded into the lattice beam to be measured. The intensity controller for the laser light is designed such that upon a TTL signal, it commands the AOM driver to fix the amplitude of the RF wave it outputs, and an external signal (coming through the intensity controller) is used to modulate the amplitude of that wave. This results in an amplitude modulation of the laser intensity as coming from the AOM. We use a function generator (AGILENT, 33220A) to supply the modulation signal. A modulation typically consists of 500–1000 periods at a modulation depth of a few % of the total amplitude. The experimental sequence ends with an absorption image in TOF.

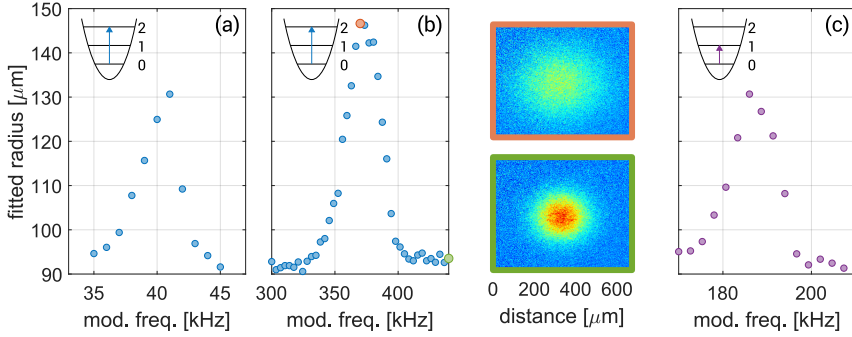


Figure 7.5: Modulation spectroscopy. (a) and (b) are modulation spectroscopy signals of the $0 \rightarrow 2$ transition, taken at two different values of the power in the optical lattice beam. The absorption images show the corresponding clouds of the orange and green data points in (b) after 3 ms TOF. The data in panel (c) are taken at the same power as the one in (b), but most likely corresponds to a $0 \rightarrow 1$ transition.

From the images we extract the width of the cloud, and this can be directly related to its temperature by eq. (2.23). By scanning the modulation frequency, traces like the ones displayed in panels (a) and (b) of fig. 7.5 were obtained. Here, modulation spectroscopy was performed in the vertical lattice beam for beam powers of 45 mW and 4 W respectively. The images to the right of (b) are TOF images. In these cases we loaded a BEC into the lattice, but in the process the cloud was heated out of the BEC phase so the falling cloud was thermal. The image on the top belongs to the orange data point and the bottom figure to the green point. The effect of the modulation is apparent from an increasing width of the cloud accompanied by a depletion in the density (the colour scale of the images is the same).

The shape of the profiles is asymmetric. Towards the lower frequencies the rise of the signal is slower, resulting in a longer tail compared to the abrupt fall of the signal after the peak has been reached. As the lattice potential is provided by a Gaussian beam, the trapping frequencies in each well will vary along the beam with the highest frequencies at its focus, as is indicated by fig 7.1 (bottom). As a result, the response of the atoms to the modulation frequency will also vary,

Table 7.1: The properties of the different lattice axes. The first two columns contain the measured waist and the estimated depth at 1 W of optical power. The next two columns list the typical lattice depths achieved on each axis during the period when the microscope was functional (June–August 2018), and the alignment efficiencies as compared to the theoretical estimation.

	w_0	V_0 at 1 W	calibration	efficiency
H1	100 μm	400 E_r	450 E_r/W	> 100%
H2	140 μm	200 E_r	160 E_r/W	80%
V	60 μm	1100 E_r	450 E_r/W	40%

with a clear cut-off at the highest frequency. The depth according to fig. 7.4 (a) is 36 E_r for (a) and 2300 E_r for (b), as determined by the highest frequency where the modulation clearly affects the cloud.

The data in panel (c) are taken with the same beam power as in (b). However, we also observe a resonance signal at frequencies corresponding to the $0 \rightarrow 1$ transition, which according to the reasoning above should not be visible. The most likely explanation is that the trap is not entirely harmonic and odd anharmonic components ($\sim \hat{x}^3$ or higher), do couple the $0 \rightarrow 1$ transition.

We have found the results of the two different methods of Kapitza-Dirac scattering and modulation spectroscopy to coincide. Table 7.1 lists the properties of the different lattice axes. The beam waists of the horizontal axes are obtained from a beam-profile measurement taken during the alignment procedure and the vertical waist is estimated from the optics design. The results of efficiencies > 100% indicate that we probably overestimate the waist of the H1 beam. After August 2018 the alignment of the vertical lattice was improved significantly to yield around 900 E_r/W (80%).

7.4 The transition from a superfluid to a Mott insulator

A superfluid can flow without friction. A BEC exhibits superfluid properties. Let's imagine we have means of stirring up our BEC and can induce a net flow in

it. For a superfluid there exists a critical velocity below which there are no means of dissipating heat, so nothing can hamper its flow. The Landau critical velocity is the point at which it becomes energetically possible to create excitations in the fluid, which can dissipate heat (see ch. 10 of ref. [219]).

This property is also present when a BEC is immersed in a weak optical lattice potential given that the immersion process itself is adiabatic. Envision a 3D optical lattice potential. In the superfluid state the atoms flow freely between sites of the lattice (by tunnelling), so the number of particles on each site fluctuates considerably, as a result of a projective measurement. The atom ensemble is coherent and that property is maintained as the atoms can tunnel in the shallow lattice. If, however, the lattice potential is made deep, the atoms become localised on a given site. This entails that the atom number on a given site becomes well determined as a result of the on-site interactions between atoms, and the system loses its phase coherence. This is called a Mott insulating state, in analogy with an electrical insulator with a filled band. The transition from the superfluid to the Mott insulating state is not driven by thermal fluctuations. The physics of the transition is governed by a competition between the kinetic energy and local interaction energy, a significant property of a quantum phase transition that this system exemplifies [118].

As a theoretical basis for the description of the system we use the Bose-Hubbard Hamiltonian. This is the bosonic counterpart of the seminal model introduced by J. Hubbard in 1963 to describe electron behaviour in solid state materials [138]. The model assumes all particles to be contained in the lowest energy band of the lattice, so no excitations are allowed. The Hamiltonian is presented in the framework of second quantisation (where occupation numbers of quantum states are quantised and altered by operators) and reads [118, 295]

$$\mathcal{H}_{\text{BH}} = -J \sum_{\langle i,j \rangle} \hat{a}_i^\dagger \hat{a}_j + \frac{1}{2} U \sum_i \hat{n}_i (\hat{n}_i - 1) + \sum_i \varepsilon_i \hat{n}_i. \quad (7.21)$$

Here \hat{a}_i and \hat{a}_i^\dagger are annihilation and creation operators for a particle on site i and \hat{n}_i is the number operator that counts the number of particles on site i .

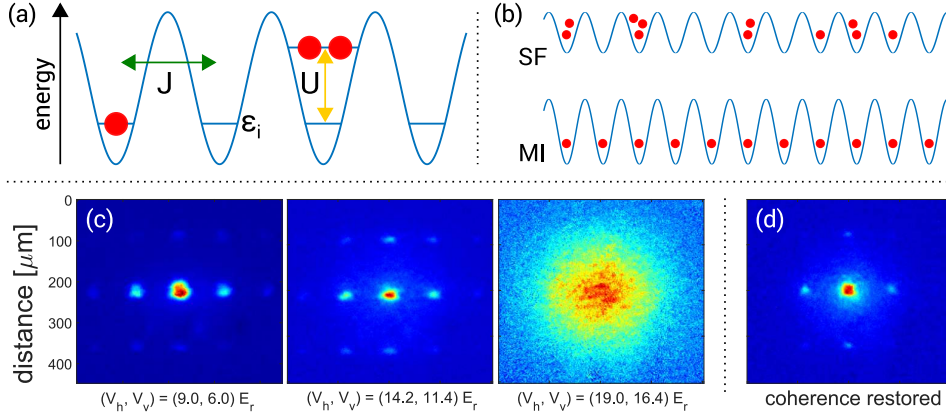


Figure 7.6: The Bose-Hubbard model and the superfluid to Mott insulator transition. (a) An illustration of the components of the Bose-Hubbard model, the tunnel coupling between neighbouring sites J , the on-site interaction U , and the energy shift per site ϵ_i due to the underlying potential. Inspiration for this figure comes from the artwork of Jesper H. M. Jensen. (b) Snapshots of the superfluid and Mott insulator phases. (c) Absorption images of the superfluid to Mott insulator phase transition. For low lattice depths the phase coherence of the condensate results in the characteristic interference peaks separated at momentum $2\hbar k$. As the lattice is ramped deep the phase coherence vanishes. (d) By ramping the lattice into the Mott regime and back down to the superfluid regime, the coherence of the state can be restored.

There are three terms in the Hamiltonian, and fig. 7.6 (a) illustrates the physics of each term. The first term describes the kinetic energy in the system and is characterised by the hopping amplitude J . The sum runs over neighbouring sites only. The second term captures the on-site interaction. If there are particles present on a given site, the energy increases by U . The third term captures energy shift ϵ_i due to the overall shape of the potential. In a typical experimental setting the optical lattice is in fact enveloped in a harmonic function which is a weak change compared to the fast spatial modulation of the lattice potential. As a result this term is often omitted.

To calculate the energies in the system we must know how to calculate J and U . To describe atoms that are localised to a given lattice site i , we use the

so-called Wannier functions $w(x - x_i)$ which have that property (see sec. 3.1.4 of ref. [117]). The hopping amplitude is the matrix element of the kinetic and potential part of the GPE Hamiltonian (first two terms of eq. (2.14)), calculated with $w(x - x_i)$ between two neighbouring lattice sites i and j , $J = - \int w(x - x_i)(T + V(x))w(x - x_j)dx$. The interaction constant is the matrix element of the interaction part (third term of eq. (2.14)) on site i so $U = U_0 \int |w(x - x_i)|^4 dx$. In the limit of deep lattices ($V_0 \gg E_r$) J and U can be calculated exactly [295] where

$$J = \frac{4}{\sqrt{\pi}} E_r \left(\frac{V_0}{E_r} \right)^{\frac{3}{4}} e^{-2\sqrt{\frac{V_0}{E_r}}} \quad (7.22)$$

and

$$U = \sqrt{\frac{8}{\pi}} \kappa a E_r \left(\frac{V_0}{E_r} \right)^{\frac{3}{4}}. \quad (7.23)$$

where a is the scattering length. The hopping element is exponentially damped with increasing lattice depth, whereas the on-site interaction term increases as the particles are squeezed together in deeper traps. When the J term dominates, the ground state is a product of single-particle superposition states that are spread all over the lattice. The occupation per lattice site follows Poissonian statistics. On the other hand when U dominates, the occupation numbers are Fock states with no number fluctuations, as fig. 7.6 (b) illustrates.

The boundary between the superfluid and the Mott state lies at $U/J = z \times 5.83$ [273], where z is the number of nearest neighbours in the lattice. In a cubic lattice $z = 6$, so we expect a transition around $13 E_r$. This quantum phase transition was first observed in 2002 and reported on in a seminal paper that ignited the field of quantum simulation [118].

As another benchmark for the functionality of our lattices, we have observed this phase transition in our laboratory. The experimental parameters were never optimised to perform the measurement to high accuracy, but still what is reported below contains all the qualities of the physics involved. We

considered it to be an important milestone, that confirms the ability of our machine to work with a truly quantum gas.

The experiment starts by creating a BEC. Then the lattices are ramped linearly to a given value in 100 ms, ensuring an adiabatic turn on.^{||} Immediately after the depth is reached the lattice potential is turned off and the cloud is dropped into TOF. The result of such a measurement where the end point of the lattice ramp is varied from the superfluid regime into the Mott regime is shown in fig. 7.6 (c). As the lattice potential is abruptly turned off in the superfluid regime, coherent matter waves at each site expand and interfere to create the distinct momentum peaks that are visible in the two images to the left. In this particular realisation (which was still the best dataset at hand) we did not keep all lattice axes at equal depth, but still all the qualities of the phase transition are visible. For this reason the interference peaks are much more pronounced vertically than horizontally. When interpreting the images, one must also keep in mind that our lattice axes are diagonal to the imaging axis. This is why five peaks are seen horizontally across the image, and the ratio of the distances from the centre peak to the top and to the side is about $\sqrt{2}$. As the depth of the lattice is increased an incoherent background arises, but the width of the superfluid peaks does not change. In the end the peaks vanish and the Mott state takes over as an entirely incoherent cloud.

The phase coherence of the superfluid state can be regained by ramping down the lattice back into the superfluid regime. An example of such a measurement is shown in fig. 7.6 (d). Here the sequence resembles the one before, but after the Mott regime is reached (here the lattice was ramped to $20 E_r$) the lattice depth is held constant for 20 ms and ramped down to $6 E_r$. We measured that 20 ms was sufficient to restore the coherence of the superfluid state.

Lets focus more closely on the scan in fig. 7.6 (c). The endpoint of the lattice was scanned from $9 E_r \rightarrow 19 E_r$ for both of the horizontal axes and $6 E_r \rightarrow 16.4 E_r$ for the vertical axis in 20 points with three repetitions. In

^{||}The criterion for an adiabatic ramp of an optical lattice is discussed in ref. [72]. Effectively the criterion is $\frac{dV}{dt} < 16 \frac{E_r^2}{\hbar}$.




Figure 7.7: The visibility of a Mott insulator. (a) The visibility as a function of the lattice depth. (b) The scaling of the visibility in the Mott regime. The broken line is a power-law fit to the last 10 points (14–19 E_r).

ref. [112] the phase coherence through the transition, and especially its reminiscence in the Mott phase is analysed in detail. The following discussion is based on that article. Similar to a classical optics experiment, where the fringe visibility of an interferometer is a figure of merit, we define the visibility of momentum peaks in the TOF image as

$$\mathcal{V} = \frac{n_{\text{max}} - n_{\text{min}}}{n_{\text{max}} + n_{\text{min}}}. \quad (7.24)$$

The n_{max} is chosen as the sum of the signal in a ROI around one of the interference peaks in the first momentum order (marked by a white rectangle in the inset of fig. 7.7 (b)). This is referenced to an empty region within the realms of the atom distribution (also marked in white). The data in fig. 7.7 (a) follow two slopes, and the intersection of the two lies between 13 and 14 E_r . This signifies the transition point from the superfluid phase to the Mott insulator phase and our results are both in accordance with theory [273] and earlier experiments [118]. The same data are graphed in fig. 7.7 (b) on a log-log plot with the ratio $\frac{U}{6J}$ on the horizontal axis. This quantity is calculated by eqns. (7.22) and (7.23). In the latter part of ref. [112] the scaling of \mathcal{V} in the Mott regime is

discussed. According to their theoretical results the integrated visibility in the Mott regime should scale as

$$\mathcal{V} = \frac{4}{3}(n_0 + 1) \frac{6J}{U}, \quad (7.25)$$

where n_0 is the filling factor of the lattice. The data from $14 E_r$ onwards are fit with a power law $\mathcal{V} = \alpha \left(\frac{6J}{U}\right)^b$, and from the fit we obtain $b = -1.01(7)$, in perfect agreement with the prediction. The pre-factor $\alpha = 2.4(3)$, yields a filling fraction of $n_0 = 0.8(1)$ in the lattice, which is also in line with our expectations.

7.5 Noise-inducing Brillouin scattering in optical fibres

It is well known that optical traps can induce heating of an atom ensemble by multiple means. These mainly include recoil heating due to spontaneous emission, intensity noise [235] and pointing noise. Heating due to spontaneous emission is minimised by using far off-resonant light. For a large detuning Δ , the scattering rate is linked to the potential U_{AC} of eq. (1.14) by $\Gamma = \frac{\Gamma_s U_{AC}}{\hbar \Delta}$ [121]. For deep lattices at $1000 E_r$ the heating rate due to this process amounts to about 150 nK/s for each lattice axis, which is very low compared to the depth, but this is an unavoidable heating mechanism.

When we first began to work with optical lattices we saw heating/instabilities caused by pointing drifts of the horizontal lattices. As discussed in appendix D those effects were minimised by coupling the lattice light through a fibre and by stabilising the temperature in the experiment to combat thermal drifts.

Soon after we fixed the pointing issue, we encountered a severe (and disproportionate) reduction of the sample lifetime in the optical lattices at high depth. Figure 7.8 (a) presents an example of such measurements. For the blue and yellow data points the H₂ lattice was operated at 13 W (2350 E_r) yielding a lifetime of about $\tau = 0.4$ s, but for the orange trace at 11 W (2000 E_r) the lifetime is improved by a factor of five to about $\tau = 2.2$ s. By changing the depth

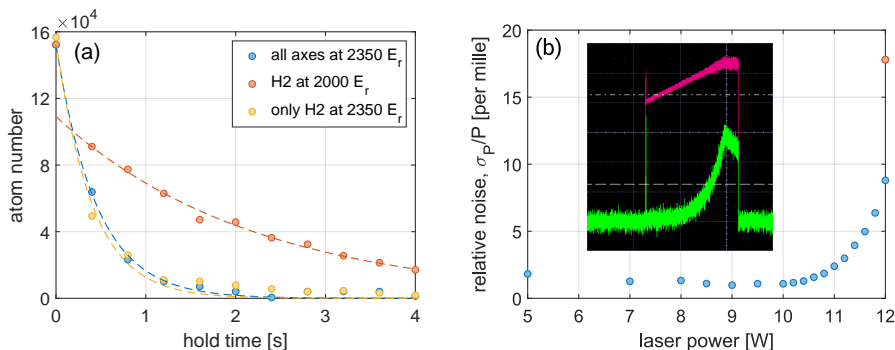


Figure 7.8: Brillouin scattering causes heating. (a) Lifetime measurements of atom clouds in the deep optical lattices, where the depth of the different axes is limited to different levels (see legend). (b) The relative noise as measured on an oscilloscope. The orange point is taken with the lattice in an unregulated mode. The inset is a snapshot of a linear ramp of the H₂ lattice (magenta), and the green trace is the light level as measured in the reflected port of an optical isolator in the beam path. The temporal spacing between the grey horizontal lines is 500 ms.

of the other axes (H₁ and vertical) regardless of the H₂, we excluded them as the source of this issue. The inset of fig. 7.8 (b) is a snapshot from an oscilloscope that monitors the laser power (in magenta) and the light level in the reflected port of an optical isolator that was later added to the laser setup. As the regulated power reaches 11 W (14–15 W before the fibre), noise is visible on the signal. We quantified this noise by obtaining the mean and standard deviation of the photodiode traces on an oscilloscope, and plot the relative noise σ_P/P as a function of P , as is seen in fig. 7.8 (b). There is a clear and a steep increase in the relative noise after we reach 11 W.

We ascribe this effect to Brillouin scattering that occurs in optical fibres subject to high intensities of light. In short, at some threshold intensity the light begins to alter the structure of the material in the fibre through electrostriction, generating an acoustic wave in the fibre. A portion of the light is scattered by the acoustic wave back from where it came (for details see ref. [164]). This is a non-linear process that limits the amount of light that can be coupled through

an optical fibre, but more importantly the backscattered light induces instabilities and enhanced noise in the laser source itself. The effect of Brillouin scattering is proportional to the length of the optical fibre and inversely proportional to its area. By using a shorter fibre or one with a larger mode field diameter we should be able to pass more light through it. To avoid getting into that detrimental regime before replacing the fibre, we simply limited the output power of the AOM for the H₂ lattice axis. Consequently this limited the depth of our lattices to 2000 E_r, which turned out still to be enough for high contrast fluorescence imaging.

Despite taking these precautions, it was most likely due to light backscatter from this effect that the ALS fibre amplifier broke down in the beginning of September 2018. After we got the laser back from repair, we added an extra stage of optical isolation after the output (as shown in the schematic of the setup in fig. B.3) which we previously thought to be unnecessary (having consulted with ALS staff) as there was an isolator in the laser head itself. Furthermore, after the fire incident discussed in the outlook, we decided to remove the internally reflective viewport on the H₂ axis in the science chamber and place a normal viewport instead. This enables us to design it in a similar way to the H₁ lattice axis so in the future the depth of our lattices should only be limited by the maximum power of the fibre amplifier.

Optical tweezer arrays

This chapter covers measurements that have been made to characterise the far off-resonant tweezer arrays generated with the 940 nm light. We have so far only used the direct imaging DMD with the beam that illuminates the full chip, i.e. the big outcoupler arm (see sec. 6.4). Our trials with the system were in hindsight coloured by the fact that the sizes of the tweezers and our power budget had been poorly characterised. The tweezers were both broader by 30–40% and the power that arrived to the atoms was lower than expected, resulting in tweezers that were at least an order of magnitude shallower than initially thought.

The structure of the chapter is as follows. The first section covers the first measurements conducted to overlap the tweezer array with the reservoir cloud. The second section accounts for general characterisation measurements, where various parameters, like the depth, spacing and size are scanned. Some images of the tweezer arrays taken with the fluorescence imaging system (then in an immature state) are shown and discussed, as well as results from modulation spectroscopy. The third section discusses characterisation measurements of the light potentials that we made after the fire, and those results are compared to earlier measurements. The fourth section discusses the role of light-assisted collisions in the tight tweezers and the optical lattices. In the fifth and final section

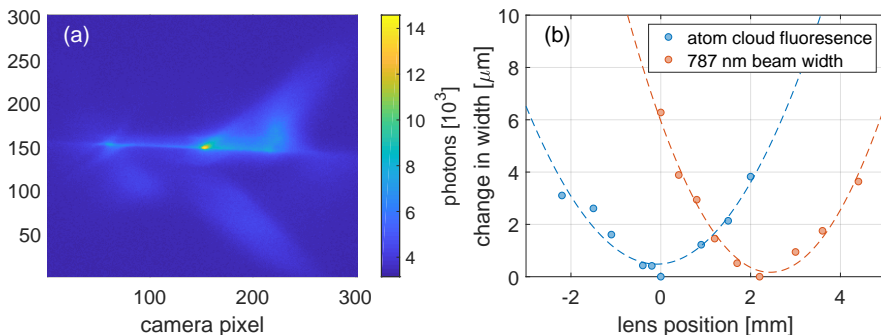


Figure 8.1: Searching for the tweezer in the haystack. (a) A fluorescence image of an atom cloud in early stage resonant optical molasses, taken from above. (b) A comparison of the width of the atom fluorescence and the 787 nm beam size as a function of the position of the last lens in the imaging system before the camera.

preliminary results for transport of atoms/atom clouds with the tweezers are shown.

8.1 Loading optical tweezers

Searching for the first signal of the tweezers is a bit like looking for the needle in a haystack. One can only align the system to a certain degree in advance, but when the chamber is closed off and the size of the reservoir is only a fraction of a mm this can become troublesome. To guide the search we built up a preliminary fluorescence imaging system that imaged the cloud from above. The light is collected by a 50 mm lens that is mounted inside the top viewport and focuses also the vertical lattice beam (see schematics of fig. 6.4). The fluorescent light passes through a dichroic and is then imaged by a second lens onto a EMCCD type camera. An example of such an image is shown in fig. 8.1 (a). A thermal cloud in the CDT is illuminated with resonant molasses beams. The LDT is visible as the horizontal line across the image, and the traces left by the molasses beams are apparent at 45° angles. With such images at hand we could easily

guide a DMD generated tweezer on top of the CDT and observe it from above.

To determine the correct focus position we took images of the atom cloud and of a tweezer generated with the 787 direct imaging DMD (to minimize any possible chromatic shift that could arise from comparing the atom position to the 940 light), while varying position of the last lens of the imaging system. The change in the width of the atom cloud and the laser beam were determined by Gaussian fits. The results are shown in fig. 8.1 (b), where the focus position of the 787 tweezer was found to be roughly 2 mm above the atom cloud. As a result we raised the high-resolution breadboard by that amount. After this manoeuvre we were successful in observing an effect of the tweezers on the reservoir cloud via standard absorption imaging from the side.

Even though we saw a detrimental effect on the reservoir cloud in the first shots (for large tweezers), we could not hold any atoms against gravity. The Rayleigh range for those traps is only several micrometers and our method of the lens focusing is not precise to that level. Our high-resolution objective is mounted to a scanner which can be translated vertically, thereby moving the tweezers up and down. To determine the height we conducted lifetime measurements of the atom cloud with a single tweezer at a waist of about $6\text{ }\mu\text{m}$ (R70) superimposed to the reservoir for various positions of the scanner. The reservoir is prepared and the tweezer is ramped adiabatically to about 5 mW. Afterwards the system is held for 4 s. This measurement is repeated at various positions of the objective, and the results are shown in figs. 8.2 (a) and (b). The atom number and cloud width are determined from an absorption image where each measurement point was acquired twice. The TOF was set too short for a reliable determination of the temperature, but the width is related to the cloud temperature. A hump is visible in the atom number of the graph (a) accompanied by a dip in temperature (b). For the dips on either side around the hump the temperature of the cloud increases. We also observed BECs arising from the reservoir at positions around 120 and 260 μm .

To aid the interpretation the real potential depth is plotted (as the blue solid line) for different positions of the tweezer in fig. 8.2 (c). Here we assume Gaussian beam propagation, but as this tweezer is quite large the potential resembles




Figure 8.2: Scanning the focus position of the tweezers. A tweezer with a waist of about $6\text{ }\mu\text{m}$ was used in all these measurements. (a) Atom number after a fixed hold time of 4 s with the tweezer superimposed to the reservoir. (b) The cloud width from the same dataset as of the one in (b). (c) The calculated potential depth along gravity for a varying position of the tweezer. The orange curve is the depth of the reservoir in the absence of the tweezer. The inset shows the radial potential for the two positions marked in the figure.

to some extent a flat-top. When the tweezer is below the reservoir it lowers the reservoir depth locally, which will aid evaporation and support the creation of a BEC. The presence of the tweezer itself also enhances the PSD via the dimple trick. But we observe that as the tweezer is moved higher into the cloud the temperature is increased and the atom number decreases (as is apparent in (a) and (b) around an objective position of $175\text{ }\mu\text{m}$). We interpret this as a result of increased density and three-body losses. In those regimes we would also expect an increase in the temperature of the cloud, as it is mainly the coldest atoms that are removed from the sample. The hump in the centre is however harder to explain. One explanation might be that when the tweezer is accurately overlapped with the cloud the volume it encompasses is in fact smaller than when it is slightly offset. Another hypothesis could be that at some point the potential depth in the vertical direction is unaffected by the tweezer (see the crossing of the blue and red curve). The effect of the tweezer on the cloud is therefore less than when the two are offset. Note that the range of the displayed tweezer displacement in (c) is in fact twice that of the objective scan range, to emphasize the dispersive shape of the curve. We never confirmed if the location of the hump

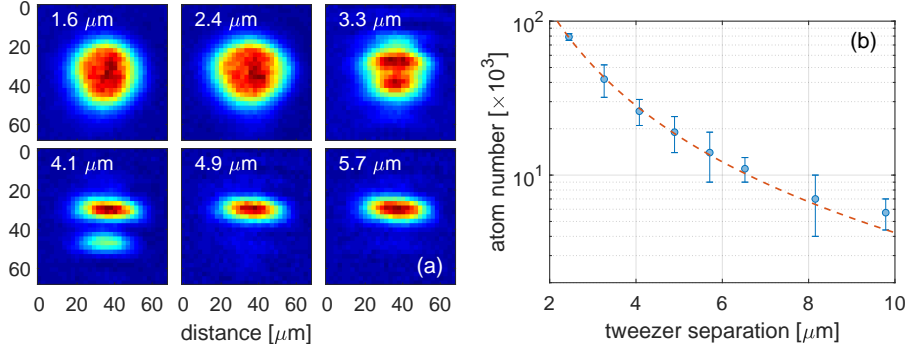


Figure 8.3: Varying the spacing in the tweezer grid. (a) In-situ absorption images for a varying spacing in the grid. The tweezer separation is indicated in the top left corner of each panel. (b) The atom number as a function of array separation. The fit is a power-law, giving an exponent of around -2 as expected.

coincided exactly with the objective position that was later on routinely used to overlap the tweezers with the reservoir.

8.2 Characterising tweezer arrays

Having determined the rough position of the tweezer we could then load and hold atoms in a tweezer array. In fig. 8.3 (a) atoms are loaded to an array of tweezers with a variable spacing and a fixed waist of about $1.0 \mu\text{m}$ (R10) and an absorption image is taken in-situ. The tweezer array covers the entire cloud, so for a tight spacing there are more tweezers that overlap with the cloud than in the case of a large spacing. For tightly spaced tweezers one big cloud is trapped, with an extent much larger than the expected Rayleigh range of $3\text{--}4 \mu\text{m}$. As the tweezers are separated the clouds start to separate and eventually only one elongated cloud remains, which is the shape one would intuitively expect.

Such periodic structures of light exhibit some non-trivial near-field diffraction as one moves away from the focus position [49]. One expects the pattern to be repeated one *Talbot length*, d_T , away (then shifted in phase by a factor of

π)

$$d_T = \frac{a^2}{\lambda}. \quad (8.1)$$

Here a is the periodicity in the structure and λ is the wavelength of the light. We see separated planes emerging at a spacing of $3.3 \mu\text{m}$ where $d_T = 12 \mu\text{m}$, and for $4.1 \mu\text{m}$ the separation is clear ($d_T = 18 \mu\text{m}$). For that case atoms are only in the Talbot plane below the main plane due to the pull of gravity. In both cases the distance to the adjacent planes matches d_T . To avoid trapping in the nearest Talbot planes we must increase the spacing of the array. In the following we work with a spacing of $5.7 \mu\text{m}$.

Figure 8.3 (b) displays the atom number after TOF. The errorbar is the standard deviation of five repetitions. The atom number is fitted with a power law in the spacing D , $N = aD^{-b}$, yielding an exponent of $b = 2.07(4)$ as expected from the fact that the number of tweezers overlapping with the reservoir and being loaded should drop quadratically with D .

To find the point of optimal loading the reservoir depth and tweezer depth are scanned for an array of both $1.0 \mu\text{m}$ (R10) and $0.78 \mu\text{m}$ (R7) tweezers. The results of those measurements are shown in fig. 8.4 (a) and (b). The optimal loading conditions in terms of depth are about double for the $0.78 \mu\text{m}$ tweezer to that of the $1.0 \mu\text{m}$ tweezer. When the tweezer array is dropped into TOF we see that the expansion is far too small for a thermal cloud. By scanning the TOF we can fit eq. (2.23) to acquire a temperature. For the vertical expansion, this yields only 20 nK . We therefore argue that as we load the tweezers multiple BECs are created in the traps by the dimple trick. This is also expected from our previous experience with the dimple trick, see fig. 4.6 (e)–(f). Along these lines we wanted to compare if there was any visible difference between loading a tweezer array adiabatically at a given value of the depth U_{tw} , or to load it at the optimal value and subsequently evaporate to U_{tw} . The results of such a measurement are shown in fig. 8.4 (b) for $0.78 \mu\text{m}$ traps. The evaporation was exponential with a time constant of 2 ms and lasted for a total of 10 ms . There is no discernible difference between the two in terms of atom number. The widths

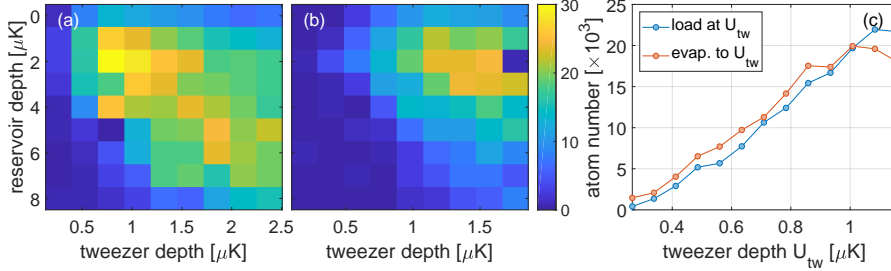


Figure 8.4: Optimal loading of a tweezer array for an array of 1.0 μ m (a) and 0.78 μ m tweezers (b). (c) A 0.78 μ m tweezer array that covers the entire cloud is either loaded directly at U_{tw} or loaded at the optimal point in (b) and subsequently evaporated to U_{tw} .

of the clouds (not plotted) were also identical. That leaves us to believe that the BECs are in fact created by the dimple trick.

As a measure of how many tweezers the reservoir cloud can accommodate, at the fixed spacing of $D = 5.7 \mu\text{m}$ the number of tweezers in a square array is varied in fig. 8.5 (a). The atom number increases linearly up to about a 5×5 array and after a 7×7 array there is no increase in atom number. Like in the case presented above, where the spacing was increased one would perhaps expect a quadratic increase in atom number, but that is not so obvious here. The first three points in the graph hint at such a behaviour, but as the array size is increased beyond a 3×3 array, this is no longer the case probably because the cloud grows thinner towards its edges. All in all, this hints at that between 30–40 tweezers are loaded with atoms from the cloud. Comparing to fluorescence images taken later on (see fig. 8.10 (d)), we confirm that we load atoms into about 40 tweezers. This means that the atom number in each tweezer varies from about 500–1000 atoms depending on where in the array it sits.

In a different measurement the radius of the tweezers in a 4×4 square array was varied and the available power scanned for each setting. The depth of the tweezer grows with the optical power as long as the size of the trap is not limited by the PSF of the imaging system. For shallow tweezers the effect of gravitational sagging is also important. The results are shown in fig. 8.5 (b). We

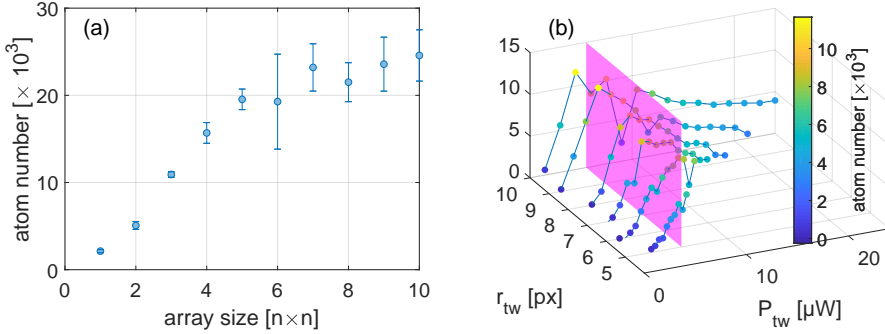


Figure 8.5: Characterising the optical tweezers. (a) Varying the size of a square array. The loading is saturated for a 7×7 array. (b) The atom number achieved at the available optical power for different patch sizes on the DMD. The pink patch is set at a power of $5.5 \mu W$ as a guide to the eye.

can achieve a state of maximal loading (where we see a maximum in the atom as a function of P_{tw}) for tweezer radii down to a radius of 6 DMD pixels, but comparing the curves to the pink patch (at $P_{tw} = 5.5 \mu W$), we can see that the peak is reached considerably later in the R6 curve compared to the others. We believe this to indicate that up to this point we are not limited by the PSF of the imaging system. For maximum power efficiency it is best to use the largest possible tweezer size that just reaches that limit. We believe that the R7 tweezer patches will serve that purpose.

As mentioned in the introduction to the chapter, we were under the impression that our tweezers should be smaller than they later turned out to be. The characterisation measurements reported in ref. [130] indicated that we should have a diffraction limited optical system for the tweezers at $NA = 0.69$. At 940 nm that would correspond to waist of about 580 nm . However, after having worked with them for a period of time we were starting to suspect that they were bigger than we initially thought. This suspicion arose mainly during attempts of using the tweezers to load a single lattice plane, discussed in ch. 10.

Another reason to believe that the tweezers were bigger than expected were

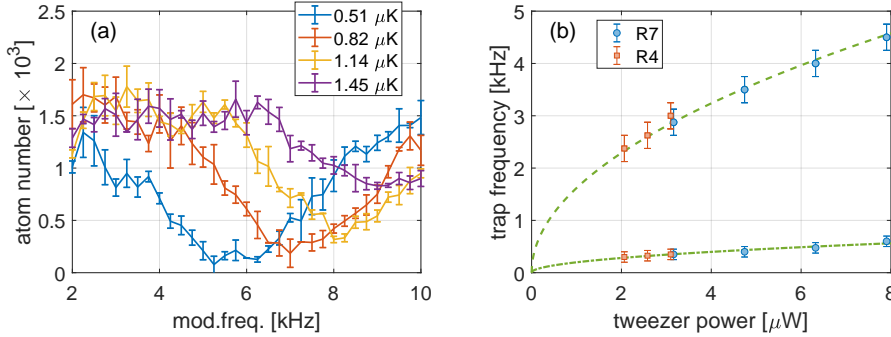


Figure 8.6: Modulation spectroscopy in optical tweezers. (a) The atom number in a tweezer array is depleted as the modulation frequency is varied. Each trace contains three repetitions. (b) The extracted trap frequencies vary as \sqrt{P} as expected, from which we can determine the tweezer waist.

independent experiments with modulation spectroscopy of atoms in the tweezers. An example of data from such measurements is presented in fig. 8.6 (a). The sequence is as follows. The reservoir is evaporated to the usual loading depth containing about 600,000 atoms at a temperature of 0.5 μ K. The tweezers are ramped to a depth of a few 100 nK, enough to hold the atoms against gravity. In the following, the reservoir is also ramped down. Afterwards the tweezers are ramped to the depth intended for the measurement. Subsequently the amplitude is modulated at a considerable fraction (10–20%) of the total depth for 50 cycles, at a given frequency f_{mod} which is scanned in multiple realisations.

Ideally one would find the highest frequency which offers some disturbance of the atom number (see discussion in sec. 7.3), and calculate the corresponding trap frequency to be the half of that value, as the intensity modulation drives the breathing mode at twice the trap frequency. However, both the breadth of the resonances and the scan range of the data shown fig. 8.6 (a) make it difficult to determine that edge, so instead we take the centre of the dip as determined from Gaussian fits as the measure. This approach will overestimate the waist.

The data in (a) corresponds to the blue points in the upper curve in fig. 8.6 (b).

They were taken with an R7 spot uploaded to the DMD. We were also able to locate another resonance at a lower frequency, indicated by the blue circles on the lower branch.* We attribute the higher branch to the radial trap frequency, and the lower branch to the axial one. To see if we could make tighter tweezers in the system we carried out the same measurements with an R4 spot on the DMD. As it turned out, the results from the two measurements yielded the same result, indicating that the R7 tweezer is at the limit of the optical system. The waist of the tweezer is extracted by fitting eq. (1.19) to the data. This yields a waist of $0.85(2) \mu\text{m}$. The uncertainty of the fit however does not capture the uncertainty in the determination of the actual waist. From the blue and the orange data curves fig. 8.6 (b), one can deduce that the edge is about 30% higher than the dip, at least for this dataset. Assuming that is the case the waist shrinks to $0.75(2) \mu\text{m}$. Based on the results presented in sec. 8.3 we adopt a value of $0.78(2) \mu\text{m}$, that corresponds to taking the resonance frequencies to be about 20% higher than presented in the plot.

8.3 Post-fire characterisation of the optical tweezers

As has become clear we had started to doubt the characterisation measurements of the DMD systems reported in ref. [130]. On the sad occasion of the fire incident discussed in the outlook, we took the opportunity to re-characterise the DMD systems in-situ. We constructed a re-imaging system out of a 0.75 NA commercial microscope objective (OLYMPUS, UPLSAPO 20X) with an effective focal length of 9 mm and a 300 mm lens. The re-imaging objective was mounted into a gimbal mount (THORLABS, GM100/M) that enabled tilting of the optic without severe translation of the image. The measured magnification of the imaging system was $M = 32.0(1)$.

As an example of an image that we acquired of a tweezer array from the 940 nm direct system is shown in the inset of fig. 8.7 (a). It was taken with an IDS UI-1240SE, a CMOS camera with an 8 bit dynamic range, which we fre-

*The ratio of the two branches is $\omega_r/\omega_z \simeq 8$.

quently use for absorption imaging. Looking closely, one may notice that the background is suspiciously flat. The camera automatically subtracts a certain reference voltage called *black level*, which can change from image to image as it depends on the temperature of the camera chip. In addition there is a constant *offset* value added before the readout. This is designed to suppress effects of dark noise, but if the settings are incorrect details will be removed from the background [142], which for our purposes must be included. We found out by digging through the old characterisation data, that it had been acquired with the wrong settings and thus heavily underestimated all beam waists.

This is shown in fig. 8.7 (a) where the mean count in a 3×3 pixel region around the peak of each spot in the inset figure is plotted against the fitted waist. The measured waist should not depend on the intensity of the laser, but it does. To understand this setting we obtained a series of images of a circular patch that was scanned over the DMD chip (effectively varying the level of illumination) and imaged it onto an identical camera in an intermediate imaging plane. The results are shown in fig. 8.7 (b), where the data were acquired with the black level compensation on and off (blue and orange data points). The size on the vertical axis represents the spot size in the intermediate imaging plane. When *off*, the fitted data depend on the illumination level comparably to the data in (a). In contrast, when switched *on*, the fitted waist is constant, albeit with an increased noise for low light levels. For this particular setting the fitted waist is still 20% too small at the highest illumination, when the black level compensation is set off. The images displayed with (c) and (d) show exemplifying spots with and without the black level compensation, together with fits to transversal cuts.

With the right setting of the camera we obtain proper images of the light fields for both the 787 and 940 nm direct imaging DMD systems. The results are shown in fig. 8.8. Examples of spots are given as well as Gaussian fits to transversal cuts through the data. The PSF as provided by ASE Optics is shown as a reference. To get a value for the waist we take the mean and standard deviation of the fits to the cuts with an average peak count above 70% of the dynamic range (~ 180 counts). For the 787 system this yields for horizontal (h) and vertical (v)

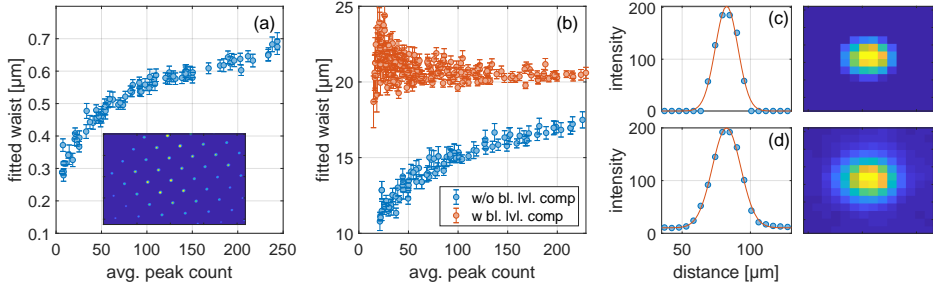


Figure 8.7: The importance of knowing your camera. (a) Using the wrong setting, the fitted waist depends on the average peak count in the image shown in the inset. (b) A set of images obtained with (orange) and without (blue) the black level compensation setting in the intermediate imaging plane. For a correct measurement, the camera's black level compensation should be switched on. (c) and (d) show exemplary spots with and without black level compensation, together with fits to transversal cuts.

cuts

$$w_{787,h} = 0.77(2) \mu\text{m} \quad (8.2)$$

$$w_{787,v} = 0.79(2) \mu\text{m}, \quad (8.3)$$

and for the 940 system

$$w_{940,h} = 0.78(2) \mu\text{m} \quad (8.4)$$

$$w_{940,v} = 0.92(2) \mu\text{m}. \quad (8.5)$$

The reason for the large aspect ratio of the 940 system is not clear. We inspected the beam paths and improved slightly the alignment of both systems, but could not get rid of the tail-like structures. We had hoped that this was somehow due to bad alignment of the re-imaging optics, which are in itself tedious to align, but as the 787 data looks much better, it is hard to blame that part of the system. In light of the modulation spectroscopy measurements reported in sec. 8.2, we will use the $w_{940,h}$ when modelling the tweezers. The power efficiency of the 940 setup was also observed to drop by 55% between the intermediate imaging

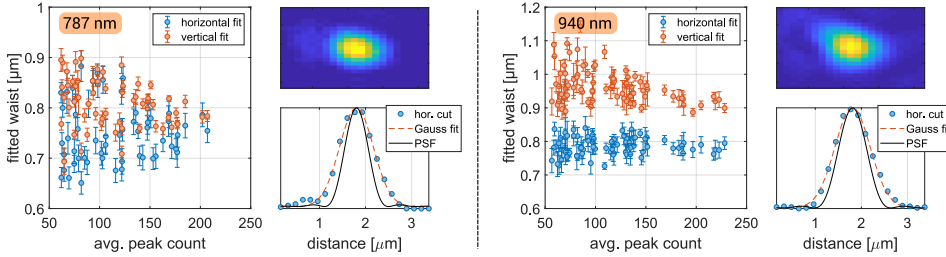


Figure 8.8: Spot sizes of spots generated by the 787 nm and 940 nm direct DMDs. Details are given in the text.

plane and the atoms, a quantity that had not been previously measured when the chamber was initially built in.

8.4 Towards few atom loading with light-assisted collisions

A central objective of our experiment is the loading and detection of individual atoms. There are essentially two routes to achieve few atoms in optical tweezers. Either one can load the trap at a very low depth, such that it only holds very few atoms (above we saw the we can maximally load up to 500–1000 atoms in each tweezer) or one relies on light-assisted two-body collisions. In the following discussion we focus on the second option.

Elastic two-body collisions are necessary to ensure the thermalisation of atomic samples. At high densities in intense light fields close to an atomic transition, these collisions can become inelastic and have detrimental effects on the atom number in the sample [92, 174]. This process is due to *light-assisted two-body collisions* (LACs). LACs are a key to understanding recent measurements of quantum phases in optical lattices where the filling fraction exceeds unity [21, 243]. They have also been used to near-deterministically prepare single atoms in tight optical tweezers [123].

The process is schematically described in fig. 8.9 (a), where the molecular potential for two interacting atoms is sketched. The blue curve represents the

potential when both atoms are in the ground (S+S) state and the orange denotes the situation when one is in the ground state and the other in the first excited (S+P). The process can be described as follows [137, 257]: (1) Two atoms in the ground state coincidentally approach one another due to thermal movement. One atom is excited by absorbing a photon from a light field which is present (2). The slope of the S+P potential is however greater than that of the S+S potential and the atoms accelerate towards each other (3). After a while the excited atom spontaneously emits a photon (4), thereby gaining a kinetic energy ΔE . If ΔE is large enough the atoms can escape the trap (5).

The loss rate is obtained from a rate equation [174]

$$\frac{dN}{dt} = -\beta' N^2(t), \quad (8.6)$$

where β' is an atom loss constant. The equation can be integrated to yield the time it takes to reach an atom number N , starting from N_0 atoms

$$t = \frac{2\sqrt{2}}{\beta} \left(1 - \frac{N_0}{N} \right) \frac{1}{n}. \quad (8.7)$$

Here β is the two-body loss rate normalised to the volume V occupied by the atoms [99], the number typically quoted in the literature. This indicates that if the densities can be maintained at a high enough level, one will end either with 0 or 1 atom in the trap, with 50% probability [237].

We have observed these effects mainly in our optical tweezers where densities reach 10^{16} cm^{-3} .[†] In our experiment the LACs are driven by near-resonant light. At first we observed the effect using the molasses light that is 40 MHz detuned from the $F = 2 \rightarrow F' = 3$ transition (about $7 \Gamma_s$). Later we also observed LACs using the repumper light which is 6.8 GHz detuned. Such far-detuned light was also used in ref. [280] for the same purpose.

[†]We obtain this number by assuming the validity of eq. (5.11), an atom number of 1000, and the ratio $\frac{\pi k_B T}{U} = \frac{1}{10}$.

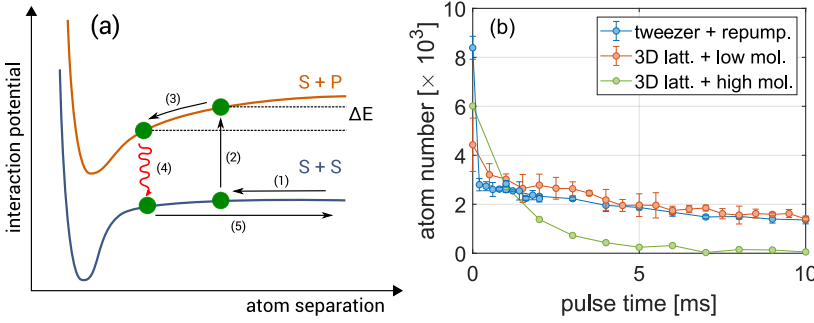


Figure 8.9: Light-assisted collisions. (a) A schematic that explains the microscopic picture of an LAC. (b) Atom number as a function of the LAC pulse duration. LACs are induced in our tweezer arrays (blue), but there is no clear sign of such a process when atoms are trapped in deep 3D optical lattices (orange and green).

In the following, an investigation of this effect in our system is presented. Three different data sets are plotted in fig. 8.9 (b). The blue data points are atoms counted in TOF that were loaded into an array of $0.78 \mu\text{m}$ tweezers, where they were held without any reservoir present at a depth of about $1 \mu\text{K}$, and exposed to light from our repumper laser. The intensity of the beam is estimated to be about $10 \text{ mW}/\text{cm}^2$. There is a clear and immediate drop in the atom number to about 1/3rd of its initial value. Estimating the timescale of the losses from eq. (8.7) by using a value of $\beta = 10^{-12} \text{ cm}^3/\text{s}$ [92][‡] we obtain that considerable fractions of the atoms are lost on the timescale of $t_{\text{LAC}} = 10 \mu\text{s}$. As the repumper laser is only controlled with a mechanical shutter (rise-time of several μs) but not an AOM, we cannot resolve the timescale to a greater precision. It is also clear that there are still atoms remaining after the LAC pulse. Assuming about 40 loaded tweezers, this amounts to about 50 atoms in each. The obvious reason for why the LACs halt is that they become ineffective as the densities drop. Our shallow tweezers are not able to maintain the density needed to reach a single atom (with 50% probability), as is routinely done in

[‡]This is the value measured at 900 MHz red-detuning, but it is the one closest to 6.8 GHz I could find in the literature.




Figure 8.10: Examples of early-stage fluorescence images of atom clouds and tweezer arrays taken in deep 3D optical lattices. (a) The reservoir cloud from which we load our tweezers, averaged from five exposures. (b) Atoms in a tweezer array, averaged from five exposures. Atoms are loaded in about 50 individual tweezers. (c) A single exposure where the lattice pinning was preceded by a 50 ms repumper pulse. (d) An average of 20 exposures from the same sequence as used in (c). The individual tweezers are more distinct, as compared to (b).

very tight tweezers [237]. By applying longer pulses we can still remove atoms, just not as effectively.

To compare the densities in the tweezer to the densities in the optical lattices, data for atoms pinned into a deep 3D optical lattice is presented. At that point we used the molasses beams to induce the LACs, but they are much closer detuned. Data are shown for two cases where the total intensity is about 10 mW/cm^2 (orange) and 30 mW/cm^2 (green). No abrupt losses are observed in either case, only a much slower decay due to recoil heating of the near-resonant light is visible. At this time the molasses beam provided no cooling effect. We believe the reason for the absence of LACs simply to be our densities. The atom number loaded into the lattice barely suffices for unity filling. As a result there are very low chances of having two atoms or more per site, which is necessary for the LAC process.

Figure 8.10 shows early-stage fluorescence images of atoms in optical tweezers. Figs. (a) and (b) show the reservoir and atoms loaded to an array of $0.5 \text{ } \mu\text{K}$ deep tweezers with a waist of $0.78 \text{ } \mu\text{m}$. The images are averages of five exposures each 500 ms long. The atom cloud is illuminated with (non-cooling) molasses beams in the horizontal plane. After loading the tweezer array, the atoms are

pinned in very deep lattices within 10 ms. Figure (d) is an average of 20 images (100 ms exposure each). The main difference between (b) and (d) is that in (d) a 50 ms repumper pulse was applied *before* the lattice pinning, whereas in (b) a 100 ms pulse was applied *after* the pinning. The difference is obvious. There are more atoms in (b) and some are clearly scattered around in the directions of the molasses beams (along the diagonal). However, these streaks are not present in (d), of which fig. (c) is a single exposure. By applying a 50 ms pulse we certainly get further into the few atoms regime as indicated by the data in fig. 8.9 (a). Looking at (c), one is tempted to assume that some tweezers even have single atoms.[§]

8.5 Transporting atoms in optical tweezers

With the high magnification of the fluorescence imaging system, we also conducted some trials of using the DMDs' dynamic abilities to display movies and transport atoms. The DMD controller board can hold up to 400 images and display them at a rate close to 10 kHz (see sec. 3.3). The exposures in fig. 8.11 show atoms in tweezers before and after they have been transported. The distance profile is constructed such that the acceleration is constant and positive in the first half of the move, and constant and negative in the second half, resulting in an S-shaped transport profile. The images in (a) and (b) show the start and end point of transport.

The atoms are loaded into a single tweezer with a waist of $0.78\text{ }\mu\text{m}$ at a depth of $0.7\text{ }\mu\text{K}$. Afterwards the tweezer was moved by a distance of $12\text{ }\mu\text{m}$ in a total time of 23 ms to (b). The movie was split up into 75 images and generated such that the image of the tweezer was moved by two pixels (corresponding to 160 nm in the atom plane) between each frame. In order to achieve the required transport profile, the display rate was varied up to a maximum of 6.5 kHz in this

[§]At the time we were starting to look into few photon counting methods for single atom detection, in the style of ref. [30], but exactly a week later after the data were acquired our optical molasses started to work.




Figure 8.11: Transporting atoms in tweezers. Atoms in a single tweezer before (a) and after (b) being moved $12\ \mu\text{m}$ in 23 ms. From (c) to (d) the corners of a 4×4 array were moved outwards by $5.6\ \mu\text{m}$.

measurement. In a different setting atoms were trapped in a 4×4 array (c) and the corners were moved diagonally outwards by a distance of $5.6\ \mu\text{m}$ (d). This was done in 50 frames and a single pixel step size ($80\ \text{nm}$ in the atom plane), reaching a maximum frame rate of $4.5\ \text{kHz}$. The total time of the move was $21\ \text{ms}$. In both cases a part of the atoms are lost during the transport. However, these tests were preliminary and not optimised to yield good results.

Atoms are routinely shuttled around with AODs [22, 87] to create defect free tweezer arrays of single atoms, but to our knowledge there is only one report so far on atom transport using DMDs [256]. In that case $1.4\ \mu\text{m}$ tweezers were loaded directly from a ^{87}Rb MOT. The traps were created by a DMD in Fourier configuration using trapping light at $785\ \text{nm}$, reaching a trap depth of $400\ \mu\text{K}$ which is more than 500 times deeper than the depth of our traps. Due to the tight confinement, the process of LACs boils the population down to 0 or 1 atom. During transport, the atoms are cooled by optical molasses that carried away a lot of the transport-induced heating. In the experiments single atoms were moved a distances of $25\ \mu\text{m}$ in $132\ \text{ms}$ time, spilt up in 30 frames. The step size was to $0.89\ \mu\text{m}$, more than half of the tweezer waist. The smallest step size we can realise in our system is only about $1/10$ th of the tweezer waist, so we can perform moves in a smoother manner. Moreover, our atoms are also colder, as we start with BECs in our tweezers. All in all, this route is worth more detailed investigations as soon as the experiment is up and running again.

Site-resolved fluorescence detection

Leiðin að hjarta atómsins er stráð innrauðum ljóseindum.

Fluorescence imaging is a standard method for detecting single atoms in a quantum gas [210]. As an atom absorbs a resonant photon it becomes transparent to the light which makes absorption imaging of a single atom very difficult [255]. Even though high phase shifts have been obtained with trapped ions [148], dispersive imaging techniques of single atoms are still inferior to fluorescence methods in terms of contrast [291]. This chapter describes our route to fluorescent single-atom detection and offers an analysis of the performance of the high-resolution imaging system. Its structure is as follows. In the first section the main topics are the optical molasses and general aspects of our imaging system. The second section discusses the properties of our single-atom signals and the third section describes a measurement of the magnification of the high resolution imaging system and how to assign the grid of the optical lattice to an atom image. The fourth and last section accounts for heating issues we traced down to the H_I lattice axis, that remain unresolved.

9.1 Optical molasses and fluorescence imaging

To obtain a good fluorescence signal from an atom one needs to collect as many scattered photons from it as possible. The photon collection efficiency is maximised by placing a high NA lens very close to the atom cloud. In addition, tight confinement in deep potentials and simultaneous cooling allow for multiple scattering events per atom before it is lost from the trap. To estimate the attainable signal from a single atom, let us imagine that we expose it to optical molasses for a time T_{exp} , from which it scatters photons at the rate Γ . The atom is assumed not to be lost during the exposure time. The collection efficiency of the objective can be calculated by integrating over the emission pattern of the atoms in the portion of the sphere subtended by the objective. Assuming an isotropic emission pattern the collection efficiency is [187]

$$\Omega = \frac{1}{4\pi} \int_{\theta_-}^{\theta_+} \int_{-\phi_0}^{\phi_0} \sin(\theta) d\theta d\phi, \quad (9.1)$$

where the limits are $\theta_{\pm} = \frac{\pi}{2} \pm \arcsin(\text{NA})$ and $\phi_0 = \arcsin\left(\sqrt{\frac{\text{NA}^2}{\sin^2(\theta)} - \frac{1}{\tan^2(\theta)}}\right)$. A numerical evaluation for an $\text{NA} = 0.69$ yields $\Omega = 13.8\%$. The transmission of the objective itself is at least 90% according to the manufacturer. As shown in fig. 6.5 (c), the light collected by the objective passes through a specially coated dichroic with 99% transmission (see appendix E of ref. [199]), and two glass plates that are estimated to have 91% transmission.* To shield from stray light three interference filters (SEMROCK, FFOI-780/12-25) are put in the imaging path. The transmission at 780 nm is 96% for each filter. All in all we lose about 35% of the fluorescence light before it arrives to the camera, so the transmission of the optics is $\eta = 0.65$. In total

$$N_{\text{ph}} = \Gamma T_{\text{exp}} \Omega \eta \quad (9.2)$$

*The glass plate transmission is found according to Fresnel's equations (see sec. 1.5.2 of ref. [34]), assuming a 50/50 mixture of S and P polarised light.




Figure 9.1: Preliminary adjustments of optical molasses and the cooling environment. (a) Light pressure in free space molasses, with and without the retro-reflection. (b) Microwave spectroscopy on the $F = 2$ to $F = 1$ hyperfine transition. The data are fit to a Gaussian function to determine the centre of the transition. (c) Varying the current in the compensation coils reveals the background magnetic fields.

photons hit the EMCCD register of the camera. Its quantum efficiency reduces the number of photons detected, but this value is well characterised as discussed in sec. 3.6. The detuning of our optical molasses is set to $\Delta = 40$ MHz, and for a saturation intensity of $10 I_{\text{sat}}$ the scattering rate is about 5% of the resonant rate in steady state. For a 0.5 s exposure time this should yield a signal of about 10^5 photons per atom.[†]

First steps to molasses cooling

The optical molasses that are used for cooling in the lattices consist of two retro-reflected beam pairs in the horizontal plane that propagate at right angles to one another, and a single beam that enters through the objective from below (the third molasses beam). The polarisation of all beams is adjusted to be circular realising a $\sigma_+ - \sigma_-$ configuration in which sub-Doppler temperatures are achieved (see sec. 2.1). As a first step in the alignment of the optical molasses, we

[†]The design and setup of the imaging system is very similar to the one described in ref. [243]. Under similar imaging parameters to those discussed here they collect only about 1/10th of the expected number of photons, as is reported in ref. [280].

observe how a single beam pushes the atomic cloud. The back-reflection is carefully aligned onto an iris early in the beam path. Fine adjustment is made to the retro-reflectors such that the atom cloud stays put when exposed to the molasses beams, indicating that the intensities of the two beams are equal. The beam that passes through the chamber is focussed by a lens onto the retro-reflecting mirror (see fig. 6.3), so any losses in the optics along the path can be made up for by adjusting the position of the retro-mirror to slightly focus the retro-reflected beam, and increase the effective intensity at the atoms. Figure 9.1 (a) demonstrates a position measurement of a thermal cloud in single-beam molasses with (blue data) and without (orange data) the retro beam for a varying pulse length of the optical molasses. The cloud is pushed by light pressure in the absence of the retro-reflected beam, as there is no trapping potential present during the pulse.

For a well-functioning optical molasses the background magnetic fields must be compensated to within about 100 mG [182], otherwise Larmor precession compromises the optical pumping mechanism of sub-Doppler cooling. We use the microwave system in the science chamber (see sec. 6.2) to perform microwave spectroscopy on the $F = 1$ to $F = 2$ hyperfine transition. An example of such a measurement is shown in fig. 9.1 (b). For each point in the graph the microwave frequency is scanned discretely over 10 kHz in 500 points dwelling 10 ms at each frequency. Subsequently the cloud is dropped for some TOF and the atoms are imaged resonantly in absorption imaging. When the resonance is hit, a portion of the cloud is transferred into the $|F = 1, m_F = 1\rangle$ state rendering that part invisible to the absorption imaging beam. The dip is fit by a Gaussian function and its centre determined from the fit. An atom cloud typically fills up a trap to the level of its own temperature. The trapping potential naturally varies across the cloud, so the atoms experience a spatially-varying AC Stark shift equal to that temperature. Here the temperature of the cloud is about 700 nK, corresponding to about 14 kHz. This coincides well with the width of the dip.

In the presence of a bias magnetic field B the degeneracy of the hyperfine states is lifted according to the Zeeman shift, eq. (1.2), and the energy difference between the $|F = 2, m_F = 2\rangle$ and $|F = 1, m_F = 1\rangle$ states increases as $\Delta E =$




Figure 9.2: Minimising detection noise. Panels (a) and (b) are histograms of images taken with different light sources present in the experimental compartments, with 1 and 3 interference filters stacked along the path respectively. (c) A measurement of the camera noise for various settings of the vertical shift speed and amplitude.

$3\mu_B g_F B = 2.1 \text{ MHz/G}$. By varying the compensating magnetic field, the resonance is shifted accordingly as shown in fig. 9.1 (c), where the centre of the Gaussian fits is plotted against the applied current in compensation coils. The location of the true magnetic field zero is obtained from fitting a function of the form $\Delta f_{\mu W} = A|I - I_0|$. In fig. 9.1 (c) the transversal fields were already well compensated. For this particular set of coils (the vertical coil pair) the calibration factor amounts to $0.675(12) \text{ G/A}$. This value is similar to what we obtained by Larmor magnetometry (see the inset of fig. 5.6), but differs by 10%. That discrepancy is natural as the compensation coils had to be removed and remounted after the bakeout following the installation of the high-resolution equipment.

Minimising stray light and detection noise

To detect single atoms, the camera should be operated in a mode that minimises technical noise. As a first step we measured the noise from various light sources on the experiment table as shown in fig. 9.2 (a). The graph shows a histogram of the ADC counts (see sec. 3.6) obtained for each setting. The distribution is characteristic of EMCCD cameras and consists of a Gaussian shaped distribution and an exponential tail. The former part stems from the nature of the ADC

register which offsets the image by around 500 ADC counts (called a baseline) and has a width that is a measure of the readout noise. The latter part comes from the stochastic nature of the process of EM amplification which essentially follows a Poisson distribution, hence the exponential falloff. The data are accumulated for an exposure time of 1 s with the camera operating at maximal EM gain of $G_{\text{EM}} = 300$. The light sources are run at typical powers during the imaging procedure. The technical noise level (blue) is taken with the shutter of the camera open. For these measurements a single interference filter (bandpass, 780 ± 10 nm) was mounted on the camera. In addition the camera is shielded off with a cardboard box and a system of 1 inch lens tubes is built around the light path. As is obvious from (a) the largest contribution comes from the vertical lattice beam, which is shone directly into the path. Some of that light is bound to leak through the coating of the high-resolution viewport, as the lattice beam is very intense. This is remedied by adding two more filters (at an angle to avoid standing waves) into the path, resulting in fig. 9.2 (b). The purple data of the horizontal molasses beams stand out as they cannot be filtered away. At the time when we acquired the data the third molasses beam was not properly aligned and so it was not measured. That beam is however the largest source of noise we have in our images and will be discussed later.

Section 3.6 accounted for general aspects of the operation of an EMCCD camera. At the operational temperature of $T_{\text{CCD}} = -40^\circ\text{C}$ the dark current is measured to be $I_{\text{dark}} = 0.042(2) \text{ e}^-/\text{px}/\text{s}$.[‡] As indicated by the data-sheet of the camera the readout noise is dictated by the readout rate and the pre-amplifier setting. For typical settings in EM mode at 17 MHz readout rate and the pre-amplifier setting with sensitivity $S = 4.24$, this noise term amounts to $\delta_{\text{read}} = 74 \text{ e}^-/\text{px}$, but as it is suppressed by the EM gain, it only results in $\delta_{\text{read}} = 0.35 \text{ ph}/\text{px}$ (photon equivalent counts).

As the shifting process of the electrons from the CCD register to the ADC converter induces CICs (see sec. 3.6), we mapped out the noise for different settings of the vertical shift time Δt_v and amplitude A_v . The CICs are am-

[‡]Here px stands for pixel.

plified in the EM register so the data represented in fig. 9.2 (c) were obtained at $G_{\text{EM}} = 300$ with a readout rate 5 MHz to minimise the effect of δ_{read} in the measurement.[§] Each data point is the mean pixel value as measured out of 100 images taken at each setting. Unsurprisingly less amplitude and faster shift time minimises the CICs. However, for the two fastest settings we observed spilling of electrons from a hot pixel[¶] present in the chip leaving a trace on the images. To avoid this we work with $A_v = 0$ and $\Delta t_v = 0.9 \mu\text{s}$ where $\delta_{\text{CIC}} = 0.197(8) \text{ e}^-/\text{px} = 0.27(1) \text{ ph/px}$. Note that as the measured value is similar to δ_{read} , it represents the joint contribution of the two noise terms for this setting. For more details of the procedure for minimisation of the detection noise, see ref. [171].

The colour of the optical molasses

The first fluorescence signals that were acquired in the lab through the high-resolution imaging system were taken without confinement in deep lattices. The molasses did not provide any cooling, but still the structure of the trapping potentials projected by the 940 nm DMD was apparent in those images. The atoms were easily heated and diffused quickly out of the traps, which were also shallow, maximally about $2 \mu\text{K}$ deep.

After adding deep confinement by the optical lattices, it looked as if single-atom signals were present in some images, as is exemplified in fig. 9.3. In (a), atoms were transferred from optical tweezers to the vertical lattices, populating several lattice planes. The vertical lattice potential is ramped down to a level just above levitation, spilling most of them out of the trap, and the atom distribution is subsequently pinned and imaged. In (b) an atom cloud is trapped in a single tweezer and its population reduced by evaporation and immediately pinned and imaged. From small light patches, as those magnified in panels (c) and (d) we count up to 50–100 photons.

[§]For this setting $\delta_{\text{read}} = 0.17 \text{ ph/px}$.

[¶]A hot pixel is a defect camera pixel that accumulates an excessive amount of charges and appears very bright.




Figure 9.3: Low-photon single-atom signals. Panels (a) and (b) are exposures taken with non-cooling optical molasses, with only a few atoms trapped in the deep optical lattices. Panels (c) and (d) are enlarged regions of light patches from where we count around 50 photons.

This low photon collection efficiency was an indication that some part of the system was not performing as expected. We had had problems with heating in the H₂ lattice axis and suspected that the H₁ lattice was also not performing as it should. However, the photon count from the light patches roughly matched the estimated number we could expect for single atoms trapped in 2000 E_r deep lattices *without* cooling, and this was a clear indication that the molasses were not functioning at all. Such indications are also present in fig. 8.11 that were taken under similar conditions, where diffuse lines appear under imaging along the direction of the horizontal molasses axes (along the diagonals). The most natural explanation is that atoms are blown out of the traps and diffuse predominantly along the direction of the molasses beams.

The individual arms of the molasses were controlled by individual AOMs as the optical setup in fig. B.3 (top) indicates. The RF wave supplied to each AOM was provided by separate AOM drivers, adjusted to the same frequency to within 200 kHz corresponding to the jitter in the potentiometer used to adjust it. This was a crucial mistake. As a step in the investigation of the poor performance of the molasses, we built a small setup that distributed the light from a single AOM to the three molasses axes. Bringing all beams to the same




Figure 9.4: Atomic luv. Atoms trapped in a heart shaped potential generated by the 940 nm DMD. (a) Atoms are held by heart alone. (b) Atoms pinned in deep optical lattices before imaging. (c) Optical molasses cooling atoms pinned in deep optical lattices. Single atoms are clearly discernible. (d) An enlarged image of the region framed in white in (c). The signal from this atom amounts to 830 photons. (e) The image uploaded to the DMD to create the potential.

frequency had an enormous impact on the fluorescence signal. Finally, the optical molasses were cooling the atoms in the deep lattices.

Figure 9.4 shows the fluorescence signal of atoms trapped in a heart shaped potential, under different imaging conditions. Images (a) and (b) show atoms exposed to the horizontal molasses at different frequencies. In (a) the atoms were held by the heart potential alone, but in (b) they were pinned in deep lattices before imaging which kept the atom heart more uniform. Image (c) is taken with all three molasses at the same frequency. Individual atoms are clearly visible even though the imaging conditions were to be improved. A pure BEC was loaded into the heart-shaped light box, and the atoms were held by it for 50 ms before imaging. The white rectangle demarcates the region of panel (d), a zoom of an individual atom trapped and cooled in a single lattice site. To create the potential we uploaded the image shown in (e) to the DMD.[¶] The width of the line in the heart is about 20 DMD pixels, corresponding to roughly 1.5 μm at

[¶]The image was hand-drawn in Microsoft's finest program, MS Paint.

the position of the atoms.

The theory of sub-Doppler cooling (see sec. 2.1) in optical molasses is non-trivial and the standard theoretical treatment is limited to the 1D case [63, 191]. The intuition that one has from the 1D theory is not easily generalised to higher dimensions, but as we found out experimentally one cannot assume simply that 1D molasses along two perpendicular directions provide 2D optical molasses. The interplay between the beams is important and cooling is definitely enhanced as the beams are brought close together in frequency. This was something we failed to realise when designing the system. Perhaps it should have been clear that at least the third beam could not have provided any cooling on its own, and interference with one of the other molasses beams would be necessary to create a polarisation gradient along the vertical direction. In our setup for the science chamber we required individual control over each beam. This was mainly chosen because we knew of the necessity to offset their frequencies at least some 10's of Hz with respect to one another to avoid standing wave patterns between the different axes (see sec. 4.5 of [280]). However, we judged such fine control over the phase of the beams unnecessary in the first iteration, and saw it only as a second step in the realisation of good optical molasses.** Thus we were left with the unfortunate situation of three molasses beams that all had slightly different colours.

Strong signals by cooling

As we now held the key to a stronger fluorescence signal it was straightforward to improve it by scanning the intensities of the molasses beams. Figure 9.5 (a) shows the total number of photons scattered per atom from a BEC pinned in deep optical lattices, for a varying exposure time. The traces have been corrected for a varying atom number during the imaging procedure using the data pre-

**Since then we have built an RF control system (using a SynthHD from WINDFREAK TECHNOLOGIES) where the phases and frequencies of all three beams are set precisely with respect to one another.




Figure 9.5: Improving the atom signal. The graph in (a) shows the scattered number of photons per atom. The trace has been corrected to account for the varying atom number. The scattering rates as extracted from (b) as a function of the molasses intensity. (c) The atom number as measured in absorption imaging with and without molasses cooling.

sented in (c). Each trace is fit with a line (no offset) and the scattering rate is extracted according to eq. (9.2).

The horizontal molasses beams have a beam waist of about $w_H = 0.5$ mm and the third beam is close to $w_3 = 10$ μm . The relative intensities of the beams were adjusted such that the scattering rates were equal for the horizontal beams, and the scattering rate of the third beam was half of the rate for the horizontal beams. The peak intensity of the horizontal beams is calculated as $I_H = \frac{2P}{\pi w_H^2}$, where P is the optical power in the beam. The total intensity is assumed to be $I_{\text{tot}} = 5 I_H$, which is normalised to the saturation intensity for $\sigma_{+/-}$ polarisation, $I_{\text{sat},\sigma} = 1.67$ mW/cm² [252]. Simultaneously with the molasses beam, light from a linearly polarised repumping beam (resonant to the $F = 1$ to $F' = 2$ transition) with an estimated intensity of 50 mW/cm², is shone into the cloud in order to prevent the depumped atoms from remaining in the $F = 1$ ground state manifold. The measured rate Γ_{meas} is plotted as a function of the saturation intensity in fig. 9.5 (b). It rises until around $I = 30 I_{\text{sat},\sigma}$, where it bends off and afterwards decreases. In the most intense regions of the third beam we see a clear depletion in the signal. We believe that the equilibrium temperature of the molasses is simply getting too high at this point, and atoms are lost from

the lattice.^{††} For reference the expected scattering rate per atom is graphed as the purple dashed line, and 1/10th of that function is plotted as the green dashed line. The reduced rate compared to theoretical expectations could be explained to some degree by ellipticity of the molasses beams, where the effective I_{sat} is in fact greater than $I_{\text{sat},\sigma}$. In any case do our results match what is reported in ref. [280].

The atom number and cloud sizes were also measured by absorption imaging for a BEC loaded into a deep lattice with and without optical molasses. This is presented in fig. 9.5 (c). The particular dataset was acquired at a total intensity of about $15 I_{\text{sat},\sigma}$, but the behaviour was not notably different with double or half the intensity. The atom number equilibrates under the molasses to about 20×10^3 atoms, whereas the number drops below that level without any cooling present. The green dashed line is an exponential fit to the data, and was used to correct for the atom number in the calculation of the scattering rates in fig. 9.5 (a). The inset shows the cloud size after 1 ms TOF, which clearly reaches a steady state in the optical molasses.

As a final remark we should state that we have chosen to work with molasses at a red-detuning of $\Delta = 40$ MHz. Along with the high AC Stark shift experienced by the atoms in the optical lattices, about $\Delta_{\text{AC}} = 12$ MHz at $2000 E_r$,^{‡‡} the atoms are almost $9\Gamma_s$ away from the $F = 2 \rightarrow F' = 3$ transition. This is similar to the parameters reported in ref. [280]. For now we have not varied this parameter, and our reasoning was that when off-resonant the temperature scales as $\frac{I/I_{\text{sat}}}{\Delta}$, see eq. (2.4), so in this regime we should be able to adjust the temperature of the molasses as needed without sacrificing too much in the scattering rate which is also linear in the parameter I/I_{sat} . Our intention was ultimately to be able to adjust the optical molasses by observing how the photons coherently

^{††}Using eq. (2.4), the molasses temperature for $I = 30 I_{\text{sat},\sigma}$ is about $400 \mu\text{K}$. This cannot be as the lattices are only $200 \mu\text{K}$ deep. Reducing the effective saturation value by an order of magnitude, as is justified by fig. 9.5 (b), yields a temperature of $40 \mu\text{K}$, about 1/5th of the total depth.

^{‡‡}The contribution is about 4 MHz from each axis with respect to the field-free environment.




Figure 9.6: Bright fluorescence signals. Panel (a) is single fluorescence exposure of a BEC, and (b) is an average of 40 frames. In (c) and (d) atoms were trapped in a 4-by-4 tweezer array, with (c) and without (d) an LAC pulse preceding pinning in the deep lattices.

scatter off a Mott insulator state, as is described in ref. [283]. During the time when the microscope was working we did not get to that point, but eventually we will.

9.2 Bright fluorescence images of single atoms

Let us inspect examples of bright fluorescence images as displayed in fig. 9.6. Panel (a) shows a single exposure of a BEC. First a BEC is prepared in the CDT, then all lattices are ramped to a depth of $20 E_r$ in 100 ms using a ramp shaped like a \sin^2 function. Subsequently the lattices are pinned within 10 ms to about $2000 E_r$ and the distribution is exposed to optical molasses ($\sim 20 I_{\text{sat},\sigma}$) for 400 ms. The granular look of the image is due to single atoms scattering off of multiple sites in the lattice. Its extent along the line of sight should be similar to its vertical extent on the figure. The $1/e^2$ width is $34 \mu\text{m}$ according to a Gaussian fit to a vertical cut through the image in panel (b), which is an average of 40 exposures. Thereby we expect it to populate about 60 anti-nodes of the vertical lattice. A horizontal cut through the image (not shown) reveals a clear flat top structure of the distribution. We take this as an indication of atom depletion during the molasses phase due to LACs (see sec. 8.4) in the centre where the cloud is densest, hinting at a filling fraction $n > 1$. In panels (c) and (d) we see atoms trapped in a $0.5 \mu\text{K}$ deep 4-by-4 array of tight tweezers. Here the atoms are pinned directly




Figure 9.7: Purple rain. (a) An exposure of about 50 atoms. From the atom in the inset we collect 8000 photons. The 33 brightest signals are superimposed to produce (b). The image is fit by a 2D Gaussian function and the residual of the fit is shown in (c). Horizontal and vertical cuts through (b) are displayed in (d) and (e) along with Gaussian fits thereof and the PSF of the objective.

after about 100 ms of hold time in the tweezer potential. In (c) a 2 ms long pulse of repumper light that induces LACs precedes the pinning and imaging (same imaging conditions as for the BEC), but in (d) it is absent. Looking closely at panel (c), some single-atom signals are clear but mostly they are blurred by signals from atoms in adjacent lattice planes along the line of sight.

To attain a better view of the performance of the imaging system we must trap only a few atoms that are spread out so the individual signals become clear. The image displayed in fig. 9.7 (a) is produced by loading a single shallow tweezer such that the atoms barely levitate against gravity. After a hold time of 100 ms in the tweezer, the vertical lattice is ramped in 10 ms to a depth of $130 E_r$. The tweezers are turned off and the lattice potential is ramped down to $20 E_r$ in a 1 s long evaporation ramp. The two horizontal lattice axes are ramped to the same depth within 1 ms and afterwards the lattices are pinned. As explained in sec. 10.3, the focus position of the tweezers is about $1.5 \mu\text{m}$ away from the optimal imaging position. In the subsequent 200 ms after pinning, the objective

is moved by this amount and afterwards the distribution is imaged for 500 ms. This provides clear single-atom signals, but at variable intensities. This is most likely due to that atoms are trapped in a few planes of the vertical lattice.

To measure the quality of the imaging system, 33 of the brightest signals are superimposed (centre is determined by a Gaussian fit) to create an averaged image as shown in fig. 9.7 (b). This image is fit by two means. First a 2D Gaussian that is free to rotate yields waists of $w_{\min} = 502(5)$ nm and $w_{\max} = 548(5)$ nm, along the minor and major axes of the ellipse that are drawn by the dotted lines in (b). The residual of the fit is shown in (c), and in fact it exhibits some coma. In panels (d) and (e) a Gaussian is fit (blue dashed line) to a horizontal and a vertical cut through the data as marked by the dashed lines in (b). This results in $w_H = 519(8)$ nm and $w_V = 503(9)$ nm. The PSF (as provided by ASE Optics) is plotted as the orange solid line. This method of determining the resolution of the system gives a lower bound, because (i) the process of superimposing the images can only broaden the true PSF, (ii) the signals used in the averaging procedure do not all lie in the same vertical lattice plane, and (iii) atoms in optical lattices are not true point sources. By combining eqs. (3.6) and (3.7) we obtain a lower bound for the NA,

$$\text{NA}_{\text{eff}} = 0.66. \quad (9.3)$$

This is within 5% of the expected $\text{NA} = 0.69$.

9.3 Imaging optical lattices with atoms

Clear images that contain many distinguishable atoms can be used to determine the underlying lattice grid that traps them.^{§§} Such a procedure also determines precisely the magnification of the high-resolution imaging system. In this dataset used below, the experiment prepared atoms in the same manner as was done for fig. 9.7 (a). However, for the imaging procedure five fluorescence

^{§§}The methods described in this section have been applied earlier in the same context, see ref. [280].

Table 9.1: Angles and spacing of the lattice grid

	θ_{\min}	lattice spacing	magnification	pixel size
H ₁	45.37(2)°	5.08(2) px	152.8(6)	104.7(4) nm
H ₂	−44.93(1)°	5.05(1) px	151.9(3)	105.3(2) nm

exposures were obtained, each 200 ms long. To each of the five images we apply a peak finder that first filters out high frequency Fourier components and then applies a threshold to the image. In this way the positions of the atoms (x_i, y_i) are determined, and from those the mutual distances $(\Delta x_{ij}, \Delta y_{ij}) = (x_i - x_j, y_i - y_j)$ between each two atoms is calculated. From the data a histogram is generated for either Δx_{ij} or Δy_{ij} and Gaussian functions with a common height, width and separation are fit to the data as shown in fig. 9.8 (a). The image is rotated and the fitted spacing between the peaks and width are plotted in the (b) and (c) panels. The horizontal lattices propagate along the diagonals of the images, so we expect these angles to be close to 45°. To the widths of the Gaussian fits we fit a quadratic function to determine the angle at which the width is minimised. The histogram displayed (a) is the one where the angle minimises the width. The results are displayed in table 9.1. We adopt a value of $M_{hi} = 152.4(6)$ for the magnification of the imaging system. The data for H₁ (not shown) are significantly more noisy, which is partially reflected in the error of the numbers shown in the table.

Having determined the spacing and angles of the lattices we can overlay a grid onto each image. To do this correctly the phase of the standing wave must be determined according to[¶]

$$V = \frac{V_0}{2} \cos \left(2\pi \frac{x}{d} + \phi \right), \quad (9.4)$$

[¶]This formula is equivalent to eq. (9.4) apart from the addition of the phase, and a constant offset.

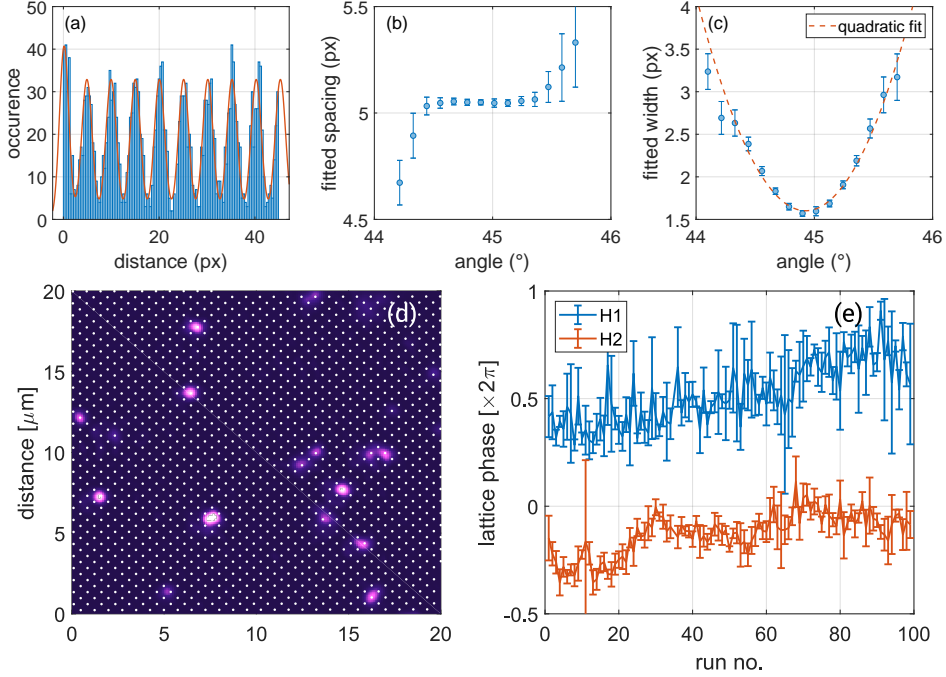


Figure 9.8: Fitting a lattice grid to single-atom signals. (a) A histogram of the mutual distances for the H2 lattice is shown. (b) From the histogram the spacing of the peaks can be determined as a function of the angle of the lattice. (c) The width of the peaks plotted along with a quadratic fit. (d) Determining the phase of the optical lattice. A Fourier filtered fluorescence exposure with the lattice grid superimposed as the white dots. (e) The phases of the H1 and H2 lattices are extracted for the entire dataset. The horizontal axis spans about 1 hour.

where d is the lattice spacing. The phase is determined by overlaying a grid $\{x_g\}$ with the angles and spacing found previously. The distance from the atom coordinate x_i to a grid-point is calculated modulo the lattice spacing

$$\Delta x_i = x_i - x_{g,j} \mod \left(\frac{\lambda}{2} \right), \quad (9.5)$$

and the corresponding phase is determined as

$$\phi = 2\pi \frac{\Delta x_i}{d}. \quad (9.6)$$

This is calculated for all the atoms in the image, and the mean and standard deviation of the sample are taken as the phase and the error in its determination.

This has been done in fig. 9.8 (d), where the lattice grid is marked by the white dots in the image. The atom image has been cleaned by filtering out high frequency components in its Fourier spectrum. This ensures a better stability of the phase determination. The graph in fig. 9.8 (e) shows the calculation of the lattice phase for the H1 and H2 horizontal lattice axes for the entire dataset used earlier. Temporally the graph spans about one hour. Each datapoint denotes the mean of the phases extracted for the five images taken in each experimental run and the errorbar the standard deviation of those numbers. The H2 trace has been shifted by $-\pi$ for clarity. The noise of the H1 lattice is notably higher and its overall drift about two times that of the H2 lattice.

9.4 Heating caused by lattices

As has been hinted at in the previous sections we did encounter heating issues in the system after we had a working microscope. We tracked this heating down to the H1 lattice axis. The heating effects both manifested themselves on the fluorescence images and in absorption imaging, and are described in this section.

Atom heating in fluorescence images

These issues became clear when we acquired multiple exposures of a few atoms prepared in the optical lattice. We observed clear atom loss and about 20% of the atoms hopped between consecutive images, primarily in the direction of the H_2 lattice beam. First the focus was set on the optical molasses. In a set of measurements we varied the overall intensity of the optical molasses. The atoms were prepared in the same way as described in the beginning of sec. 9.3. In each run five 200 ms exposures were obtained and in all cases the time between the images was short (5 ms) apart from one case where it was set to 200 ms. During that time the atoms were not cooled. To every image we apply a peak-finding algorithm. The number of atoms is counted in each image and the average fractional change in the atom number is calculated for all 100 images obtained at each setting. This is plotted in fig. 9.9 (a). We clearly see that more atoms are lost for higher molasses intensities. This is maybe no big surprise, but even at the lowest intensities there is still a notable loss.

In another set of measurements^{***} the intensity was fixed to around $15 I_{\text{sat},\sigma}$ and the molasses beams were blocked one at a time. The data are analysed as before, but the experiment was only repeated five times for each realisation, which at least partially explains larger fluctuations. The results are shown in fig. 9.9 (b). When all beams are cooling, the population is reduced by roughly 10%. We also observe that the atoms primarily hop along the direction of the H_2 lattice beam. As one of the horizontal beams is blocked the atoms start to hop around more, which is reflected in an increased fluctuation and noise in the trace. The most likely reason for the relative atom number rising above one, is that due to increased hopping atoms appear in later images as they spread out during imaging. When the third molasses beam is blocked, atoms disappear quickly from the image, probably because they start to hop along the vertical direction. This shows the necessity of using all three molasses beams.

We were however never fully happy with the third molasses beam. By shin-

^{***}This latter set was taken three weeks later and the molasses were re-aligned at least once during that period. This can explain the small differences in absolute numbers.

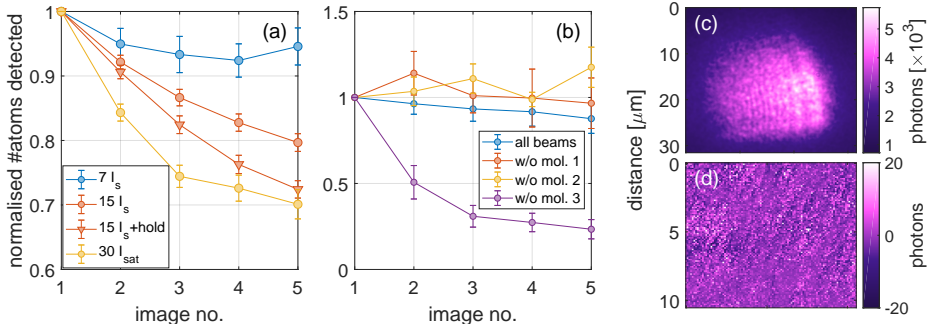


Figure 9.9: Atom loss due to hopping. (a) The fractional change in atom number was measured for different molasses powers. (b) Here the cooling beams are blocked one at a time. (c) The third molasses beam. (d) The noisy background of the third molasses beam.

ing it into a thermal cloud of atoms its intensity distribution may be imaged by averaging a multiple of exposures. Such an averaged image is shown in fig. 9.9 (c). The beam only covers an area of $20 \times 20 \mu\text{m}^2$. In addition the intensity distribution is not homogeneous and fringing is apparent, indicating a cut in the path.^{†††} But even though the third molasses beam is necessary to cool the atom distribution, it introduces speckles and enhances the noise in the images. This is seen in fig. 9.9 (d), that displays a single image of the background from which an averaged image of the background has been subtracted. Therefore it measures the noise present in a single shot. As the figure shows, there are fringes present that vary from shot to shot. The fringes likely stem from interference of the molasses beam with itself, as it is combined into the imaging path by a glass plate. It is difficult to get rid of this background pattern, but in future measurements we could try to wash it out by scanning the beam pointing with a piezo-driven mirror at a rate which is much faster than the exposure time.^{†††}

^{†††} At the time of writing we have in fact improved the beam quality.

^{†††} Such piezo modulators are already built in on the retro-reflecting mirrors of the horizontal molasses beams. However, we never saw any clear effect of using them. That was probably only caused by a lack of a clean enough atom signal (like a uni-layer Mott insulator state).

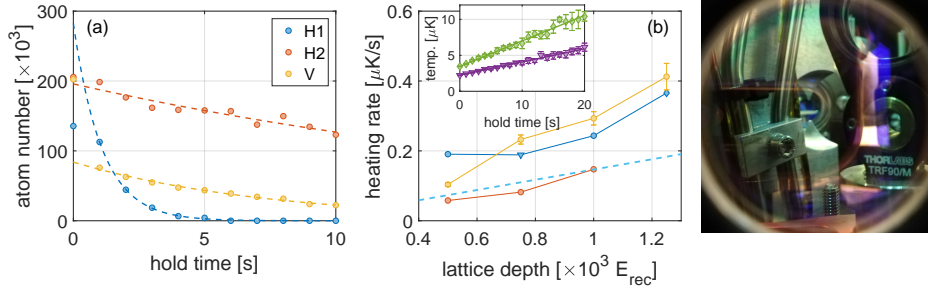


Figure 9.10: The erratic H1 lattice. (a) Lifetime measurements indicate a serious issue with the H1 lattice. (b) Measurements taken in a different setting, indicate that the heating rates are similar for all axes. (c) A potential cause for the erratic behaviour. The holder of the transport coils might have partially blocked the beam path of the retro-reflected H1 beam, depending on its alignment.

Atom heating in absorption images

As another step in these investigations we performed lifetime measurements in the lattices. A BEC was carefully loaded into a 1D lattice, pinned and held for a variable time. The results, shown in fig. 9.10 (a), reveal a particularly poor lifetime in the H1 lattice of only $\tau_{\text{H1}} = 1$ s. The H2 lattice accommodates most the atoms as it has the largest waist, and likewise the vertical lattice has the smallest waist and is expected to hold the fewest atoms. The lifetimes for both of those axes are $\tau_{\text{H2,V}} = 6$ s. In the time to come we carefully inspected the different aspects of the lattice beam, its alignment, the control electronics, optical elements in the beam path, and potential effects of thermal lensing that we found to be minimal. We changed around the optical fibres and inspected the retro-reflection but found no obvious reason for the poor performance. It was in this process where our Azur laser eventually broke down, probably due to back-reflections from the optical fibres due to Brillouin scattering as discussed towards the end of sec. 7.5.

In its absence we coupled light from the Nufern laser (used for the dipole traps and the vertical lattices) into the lattice paths. In this way we could have

low power lattices, or fairly high power in one axis at a time. Instead of looking at lifetimes, we focussed on thermometry to extract heating rates. The results are shown in fig. 9.10 (b). The inset is an example of heating rate measurements of the H₁ lattice. For comparison the expected heating rate $\frac{dE}{dt} = 2E_r\Gamma$ is plotted as the dashed line in light blue. This time we measured no excessive heating in H₁, and in fact it was comparable to the other axes. Thus we blamed the heating on the Azur amplifier.^{§§§} When the Azur arrived back, we measured again extremely poor lifetimes in the H₁ lattice so once more we diverted light from the Nufern into the H₁. Now the heating rates were just as bad with the Nufern, as we measured with the Azur, which lead us to believe the amplifier was in fact not the issue. We swapped an optical isolator in the beam path of the axis and measured all of a sudden much better lifetimes again. It was however difficult to believe that a static optical element could cause such heating.

In hindsight it was difficult to track down the heating issue. Although we managed to isolate the H₁ lattice as the source, the exact reason for it was never clear. When looking back on the lifetime and heating measurements the degree of the effect varied dramatically. On top we had the results of the fluorescence measurements in mind where the atoms seemed to primarily hop along direction of the H₂ lattice, which is hard to reconcile with the H₁ being the sole source of heating. Perhaps some part of the thermal hopping can also be traced to the molasses, but when many things are wrong simultaneously, it is very time-consuming to disentangle the issues.

During the reconstruction of the experiment after the fire we found a potential cause for the poor lifetime of the H₁ lattice. In all experimental sequences the atoms are transported from the 3D MOT to the cube chamber by movable coils. The coils are only sent back to the 3D MOT at the end of each sequence, so they stay at the cube for the whole time experiments are conducted in the science chamber. Figure 9.10 (c) is taken through the backside of the retro-reflecting dichroic of the H₁ lattice, looking along its path back into the science chamber.

^{§§§}We still found this hard to believe as the H₂ lattice shares light with the H₁ lattice, so any intensity noise arising in the amplifier itself should also appear on the H₂ axis.

The transport coil holder is visible and the tubes on it are for cooling water. The window of the science chamber is in the centre of the image. It is of course hard to fully realise the situation from a single image, but looking from this angle it seems quite probable that parts of the coil holder could block the beam path. This could also explain why the effect was so erratic, as a realignment procedure of the lattice axis could result in partial blocking of the retro beam. Looking back, we failed to inspect the retro-reflected beam with the transport coils at the cube chamber as all experimental sequences automatically drive it back. In future experiments we will make sure to drive the transport stage back to the 3D MOT before any lattice beams are turned on. That will hopefully resolve the issue permanently.

Atom cloud tomography

Our first strategy for loading a single lattice plane was to rely on the short Rayleigh range of the optical tweezers. However, as discussed in ch. 8, the tweezers turned out to be wider and shallower than we initially thought. The distribution of atoms loaded into the different planes of the vertical lattice is still of interest, especially as the high-resolution objective can be moved to image different planes of the vertical lattice. By acquiring multiple fluorescence exposures at different positions of the objective, one can even trace out the positions of individual atoms and tomographically reconstruct the 3D atom distribution in a sparsely loaded optical lattice.

Full 3D tomography has not been demonstrated before in wavelength scale optical lattices. Until now, the method has only been applied in large-spacing lattices [200] and in 3D arrays of optical tweezers [23], where the spacing between atom layers is minimally $5\text{ }\mu\text{m}$. One study, however, reports on methods to prepare and read out a bilayer Mott insulator system in a quantum gas microscope [225]. In their approach a superlattice facilitated the creation of a tunable double-well potential along the line of sight. By virtue of a tunable standing wave in the molasses light, the scattering rate could be adjusted between layers. At first, light was collected from both layers, with only one of them in focus.

The scattering rate was tuned to be lower for the plane out of focus. After imaging the atoms in the in-focus plane, they are subsequently heated out by intense molasses, that was tuned to cool the neighbouring layer. By removing a glass plate in the imaging path the focus of the system was shifted to that plane, allowing for a proper image to be taken. The distribution in the two planes was determined in post-processing.

The chapter is structured as follows. In the first section the translation of the objective by the scanner is characterised. In the second section, we theoretically analyse the expected distribution of atoms loaded from the tight tweezers into the vertical lattice. The third section covers an analysis of a dataset where (on average) 1–2 atoms were loaded into the vertical lattice, and the objective position was changed *between* experimental realisations. The fourth and final section accounts for the analysis of the loading of 20–30 atoms into the vertical lattice, where the objective focus is scanned *within* an experimental sequence, enabling atom cloud tomography.

10.1 Moving the high-resolution objective

The high-resolution objective is mounted on a piezo-driven scanner (see sec. 6.4) and can be positioned with nm resolution. As a result it can be translated between consecutive fluorescence exposures within the same sequence. To characterise its movement we recorded position profiles, and varied the amplitude of the movement and the *slew rate* v_{\max} , which is basically the maximum allowed velocity of the movement. Figure 10.1 (a) shows the results where the 90% settling time T_{90} (the time it takes to travel 90% of the desired distance) is plotted as a function of v_{\max} . For an increasing v_{\max} the travel time settles to about 70 ms, which eventually becomes limited by the PI-control of the scanner itself. The optimal control parameters depend on the weight of the objective, but the adjustment of the P and I coefficients of the control must be made with care. If one were to drive the piezo too hard it might induce resonant oscillations that can, in the worst case, cause permanent damage to the piezo-electric material.




Figure 10.1: Characterising the movement of the objective scanner. (a) Measurements of T_{90} are shown as a function of v_{\max} . (b) To improve those results the mechanical resonances of the objective scanners are measured. (c) The optimised position profile after suppression of the mechanical resonances.

The manual of the scanner has instructions on how to optimise the movement. This basically involves measuring the mechanical resonances in the system and subsequently applying so-called notch filters that damp the effect of those resonances when the scanner is moved. An example of such a measurement is displayed in fig. 10.1 (b), which shows the Fourier transform of the position profile (the inset) recorded after an abrupt shift of the scanner. The location of the highest two resonance peaks can now be typed into the control program, and the program subsequently suppresses the resonances. The improvement is demonstrated in fig. 10.1 (c). The blue curve shows the position profile, and T_{90} is marked as the yellow dashed line as the objective was shifted by $dz = 0.5 \mu\text{m}$. In comparison the old profile is plotted in orange and the corresponding T_{90} as the purple dashed line. The 90% settling time is improved by a factor of two, and residual oscillations during the move are no longer present. The inset shows the optimised position profile for $dz = 5 \mu\text{m}$, where the T_{90} of that curve is marked in green and lies just below 50 ms (the inset shares the horizontal axis with the main graph). The T_{90} of the $dz = 0.5 \mu\text{m}$ movement is also marked for reference.

These measurements were however conducted in the period after the Azur

amplifier broke down, so for all data reported below we worked with the old settings and the time between exposures was set to 200 ms. This ensured enough time for the objective scanner to settle before the next image was taken. In the future we can safely reduce this time to 50 ms and cut down the total imaging time by 600 ms.

10.2 Slicing a hypothetical BEC

Our goal is to understand the population distribution of the atoms in the vertical lattice, as loaded from the optical tweezers. Small BECs are created in our tweezer potentials by virtue of the dimple trick. Their density is approximated by the Thomas-Fermi approximation of the GPE equation, see eq. (2.15). We work in cylindrical coordinates and assume a harmonic potential of the form

$$V(r, z) = \frac{1}{2}m(\omega_r^2 r^2 + \omega_z^2 z^2), \quad (10.1)$$

where ω_r and ω_z are the radial and axial trapping frequencies. The radial boundary condition is obtained by setting the BEC density, given by eq. (2.15), to zero. That yields $R(z) = \sqrt{R_{\text{BEC}}^2 - \varepsilon^2 z^2}$ where $\varepsilon = \frac{\omega_z}{\omega_r}$ and R_{BEC} is the radial extent of the BEC given by eq. (2.16). By integrating out the radial direction, we get the density $\tilde{n}(z)$ as a function of the line-of-sight coordinate z

$$\tilde{n}(z) = \int_0^{2\pi} d\theta \int_0^{R(z)} r dr |\psi(r, z)|^2 \quad (10.2)$$

$$= \frac{2\pi}{U_0} \int_0^{\sqrt{R_{\text{BEC}}^2 - \varepsilon^2 z^2}} (\mu - V(r, z)) r dr \quad (10.3)$$

$$= \frac{\pi}{U_0} \frac{m\omega_r^2}{4} (R_{\text{BEC}}^2 - \varepsilon^2 z^2)^2, \quad (10.4)$$

where μ is the chemical potential given by eq. (2.17).

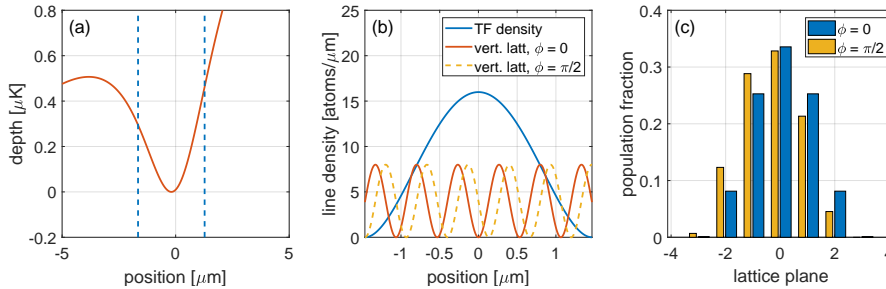


Figure 10.2: Slicing up a dimple BEC. (a) The tweezer potential along the line of sight. The dashed lines mark the extent of the atom density shown in blue in (b) panel, that shows the line density of the Thomas-Fermi approximation of the GPE for $N = 25$ atoms. The lattice potential is drawn when the center of an anti-node overlaps with the middle of the cloud ($\phi = 0$), and when its phase is shifted by $\phi = \pi/2$. (c) A histogram of the population fraction that is loaded into different lattice planes of the mode-matched and shifted lattices in (b).

With the atom distribution in hand we can slice up the density distribution and calculate the population fraction loaded into each lattice plane. Figure 10.2 shows an example of such a calculation. In panel (a) the tweezer potential is drawn in orange along the line of sight (which is also the direction of gravity), and the blue dotted lines represent the extent of the BEC, which is represented in panel (b) by the blue line. For large atom numbers and shallow traps, one might doubt the validity of the harmonic approximation as the potential is severely modified by gravity.* The atom distribution in (b) is rendered using the radial and axial trapping frequencies measured by modulation spectroscopy as reported in sec. 8.2 with $N = 25$ atoms. The atom number is chosen as these are typical numbers loaded into the tweezer in the scan described in sec. 10.4.† The lattice potential is drawn as the solid orange line, with an anti-node opti-

* Luckily the radial extent of the BEC only grows as $R \sim N^{\frac{1}{5}}$, so at least adding more atoms will not render the approximation invalid, but to be sure one should solve the GPE in the correct potential.

† For such low atom numbers one might rightly doubt the quality of such a simple treatment. This is first because the Thomas-Fermi approximation grows better with larger BECs (and 25 is not many) and second because the GPE is after all a mean-field approach (and 25 is




Figure 10.3: Population fraction in different vertical lattice planes rendered for (a) 25 atoms (b) and 100 atoms. In (c) we also have 100 atoms but the maximal power in the tweezer is 10 times higher than our current maximal optical power. The numbers in the legend stand for the plane of the vertical lattice with respect to the central plane.

mistically placed in the center of the distribution. The case when the phase of the lattice has been shifted by $\phi = \pi/2$ is drawn as the yellow dashed line. Panel (c) shows a histogram of the fractional population in each plane obtained from integrating from crest to crest, for the two lattices shown in (b). The main distribution is represented by the blue bars, and the one from the shifted potential by the yellow ones. The effect of a mismatch is to skew the loaded distribution, which otherwise populates about five adjacent planes.

It is easy to extend this calculation to higher atom numbers and greater trap depths, as shown in figs. 10.3 (a)–(c). Here the calculations were done in the unshifted lattice. The trap setting used to generate fig. 10.2 is marked by the green dashed line in panel (a). As is expected more atoms provide a broader distribution and deeper tweezers narrow it. If we were to rely on this method as a route towards single plane loading it would definitely be favourable to add more tweezers, each with few atoms, instead of adding more atoms to fewer tweezers. Instead of relying on the 940 nm tweezer system with the big outcoupler (that provides a more homogeneous illumination), we could either use the small out-

not many). For a better theoretical treatment of N interacting bosons one could consult a multiconfigurational time-dependent Hartree method for bosons (MCTDHB) [4], but such fine theory is outside the mental capacity of a humble experimentalist writing a PhD thesis.

coupler for illuminating the DMD, or use the holographic DMD system (see sec. 6.4). This should be favourable in terms of attainable depth, at least for few tweezers. The calculation in panel (c) was rendered with 100 atoms up to an optical power of 10 times the maximal power of the tweezer system that we currently use. For the highest depths only three planes of the vertical lattice are populated.

10.3 Counting a few atoms in multiple lattice planes

To measure the atom distribution in the vertical lattice we must scan the position of the objective. For the dataset covered in this section, five fluorescence exposures were acquired in each experimental run. The first image was taken with the objective at the same position as where the tweezer-trapped cloud was loaded to the vertical lattices z_{load} , then afterwards the objective was moved to another position z_i that was scanned between consecutive runs. For each image the atoms were exposed to the optical molasses for 100 ms and the time between images was set to 300 ms.[‡] Prior to imaging, a single tweezer was loaded at a depth of 0.5 μK . After turning off the CDT the atoms were held in the tweezer for 10 ms before a 4 ms long repumper pulse depleted the atom number by LACs. After another 100 ms of hold time, the distribution was pinned in deep lattices and imaged as described earlier. The objective position z_i was scanned over a range of 5 μm with a resolution of $dz = 0.5 \mu\text{m}$, just below the lattice spacing of $d = 0.532 \mu\text{m}$. This cycle was repeated 120 times to acquire statistics. In the following analysis the image obtained at z_{load} is omitted, and the other images are treated individually. The position of the objective will be referenced to the central position of the scan z_0 , by the use of the quantity $\Delta z = z_i - z_0$. This position happens to coincide with the z_i where atoms are detected.

[‡]This was the first experiment that we conducted where the objective was moved between exposures, so in hindsight we used a rather short exposure time and a long time between the images. The atom hopping was in fact not as pronounced in this scan as in later scans.

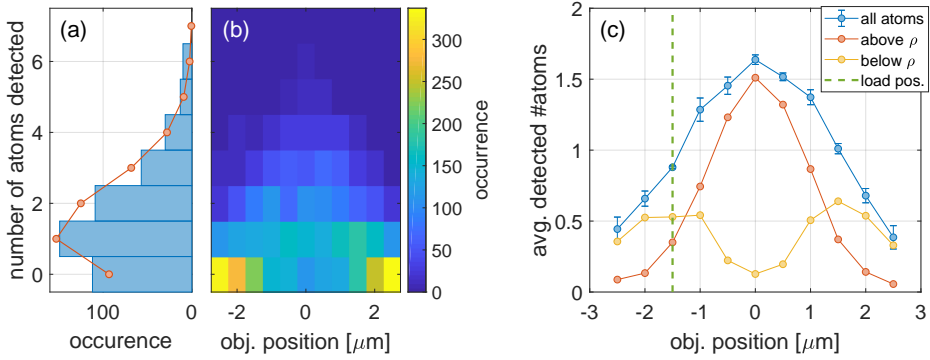


Figure 10.4: Atom counting statistics. (a) A histogram of the number of atoms detected at z_0 , fit by a Poisson distribution. (b) A histograms for all objective positions z_i . The image shares the vertical axis with (a). (c) Average atom number detected as a function of z_i (blue points). Atom signals above (orange circles) and below (yellow points) ρ , the mean PTR.

Atoms out of focus will appear fainter and cover a larger area on the image than those that are in focus. As a measure of this feature we adapt a quantity that we call the *peak-to-total ratio* (PTR). It is ratio of the *peak count* in a 3×3 pixel array around the centre of the atom image, to the *total count* in a 7×7 pixel array around the same centre,

$$\text{PTR} = \frac{N_{3 \times 3}}{N_{7 \times 7}}. \quad (10.5)$$

We apply a peak finding algorithm to all images, that locates the atoms and the centre of the signal to the precision of a camera pixel. The PTR is calculated from a Fourier-filtered version of the raw image.

We begin by looking at the number of atoms detected as a function of the position z_i . Figure 10.4 (a) is a histogram of the atom number at the centre position z_0 . The orange circles represent the Poisson distribution that best fits the data, yielding a mean number of $\bar{n} = 1.63(8)$ atoms. The image in fig. 10.4 (b) is the histogram compiled at all the different positions z_i . The distribution has a spear-like shape indicating that most atoms are loaded in the vicinity of z_0 , and

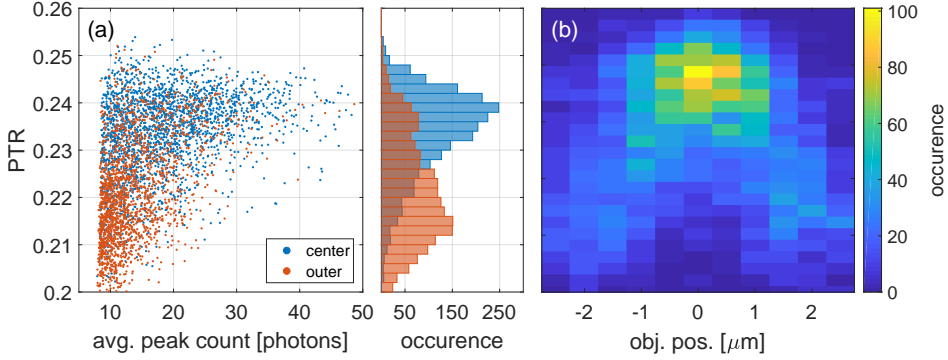


Figure 10.5: Inspecting the PTR of atoms detected. (a) A scatter plot of the PTR and $\bar{N}_{3 \times 3}$ for atoms detected close to z_0 (blue points) and far from z_0 (orange points). The histogram shows that the PTR distributions are notably different. (b) An image of the histograms as a function of z_i .

it is also clear that towards the edges no atoms are found in the vast majority of the images. In fig. 10.4 (c) we plot \bar{n} as a function of z_i . The maximum number of atoms lies $1.5 \mu\text{m}$ above z_{load} , as is indicated by the green dashed line. This is most likely the effect of a chromatic shift between the 940 nm light in which the atoms are held, and the 780 nm light they emit. This is an important observation as the objective must always be moved by this amount before imaging if the tweezer potentials have been used earlier in the sequence. By looking at the PTR values of the atoms detected we filter the \bar{n} that have $\text{PTR} > \rho$ (orange circles) and $\text{PTR} < \rho$ (yellow circles), where ρ is the mean value of the PTR for the whole dataset. We see that atoms with low PTR are detected away from the z_0 and signals with higher PTR values are typically found close to the centre.

To get a better idea of how distinguishable the atom distributions are in and out of focus, a scatter plot is presented of the PTR and the average peak count in the centre of the atom signal, $\bar{N}_{3 \times 3}$, also obtained from the filtered image. The dataset presented in fig. 10.5 (a) is split in two. The blue points contain atoms found in images within $|\Delta z| \leq 0.5 \mu\text{m}$ (three positions in total) and the orange points contain $|\Delta z| \geq 1.5 \mu\text{m}$ (six positions in total). The atoms

detected in the vicinity of z_0 generally have higher PTR, and a higher density of high values of $\tilde{N}_{3 \times 3}$. The histogram on the side shares a vertical axis with the scatter plot, and here it is clear that at least on the length scale of $1 \mu\text{m}$ there is a notable difference in the distributions. In fig. 10.5 (b), the histogram is presented with full resolution, so each vertical row corresponds to a given z_i . The Λ -like shape of the distribution is related to the underlying atom distribution, but the measurement result itself is convolved with the axial profile of the PSF as given by eq. (3.5). This is inherently the limiting factor in this approach, as there is no obvious way to determine with certainty if a given atom signal is in or out of focus and how far away it is from the focal plane.

10.4 Tomographic reconstruction of sparse atom clouds

To gain information about into which lattice plane a given atom was loaded, the position of the objective had to be changed between consecutive exposures within the same experimental run. This was achieved in a set of experiments where a single tweezer was loaded at low power and the cloud was subsequently moved over to a $130 E_r$ deep vertical lattice followed by an evaporation in that potential down to $20 E_r$. The sequence, apart from the imaging procedure, was the same as the one used to obtain fig. 9.7 (a), and is described in sec. 9.2. For the data presented here five images were obtained, each with an exposure time of 200 ms, separated by the same amount of time. During the time between the images the objective was translated by $dz = 0.5 \mu\text{m}$ through a set of positions $\{z_i\}$, spanning a total range of $2 \mu\text{m}$. This sequence was repeated 100 times to acquire statistics.

The dataset is not free of heating effects.[§] For further data processing we filter out atoms that are not found within a distance of one pixel (up, down, left, right and diagonally) from one image to the next. Essentially, the atoms we consider have to stay put in their lattice site in all five images for the total duration of the imaging procedure of 1.8 s. In this way we obtain traces $\{a_i\}$,

[§]Indeed the data were taken on the same day as the hopping data discussed in sec. 9.4.

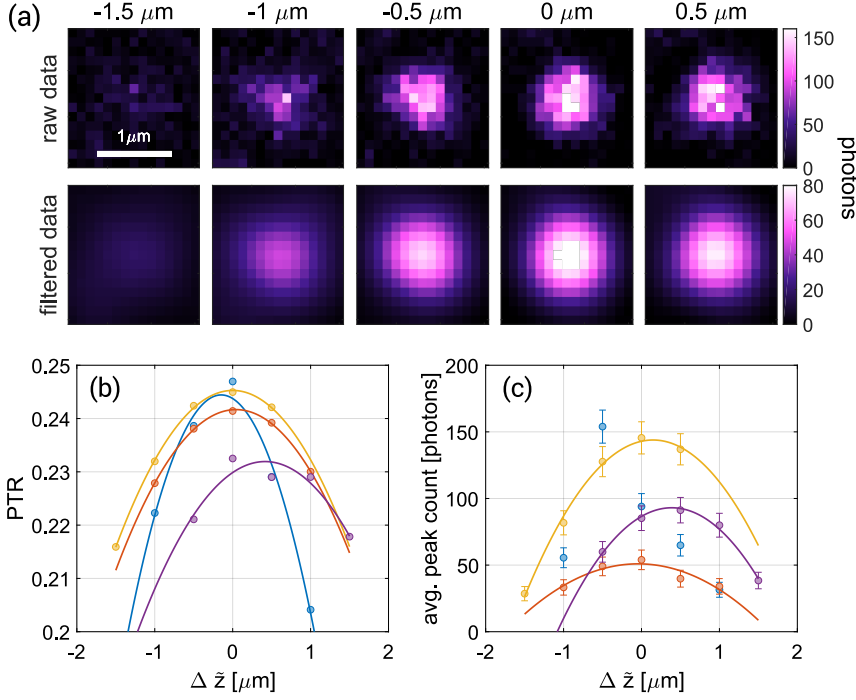


Figure 10.6: Tracing out the atom position. The top row in (a) shows an example trace of raw atom images and the bottom row contains the filtered versions. The numbers above the top row denote $\Delta \tilde{z}$. In (b) and (c) the PTR and $\bar{N}_{3 \times 3, \text{raw}}$ (as calculated by the raw image) are plotted as a function of $\Delta \tilde{z}$, respectively.

where $i = \{1, \dots, 5\}$. An example of such a single-atom trace and how the atom image changes as a function of z_i , is given in fig. 10.6 (a). The figure shows a series of five raw atom images in the top row and their filtered versions in the bottom row. For each image we find the PTR by eq. (10.5), and the distance from the maximum PTR to a given image $\Delta \tilde{z} = z_i - z_{\text{max}}$ is written above the top row. For this particular trace the maximum PTR was found in the 4th image of the series.

An example of some PTR traces is given in fig. 10.6 (b), and the particular

one in (a) is shown in yellow. Through each set of points a quadratic function is fitted, and this is also drawn as a guide to the eye. In fig. 10.6 (c) we plot $\bar{N}_{3 \times 3, \text{raw}}$ as a function of $\Delta \bar{z}$ for the same traces as shown in (b). The errorbars represent the photon-shot noise $\sqrt{\bar{N}_{3 \times 3, \text{raw}}}$. Even though the PTR values are similar, e.g. for the orange and yellow traces, the orange and the yellow $\bar{N}_{3 \times 3, \text{raw}}$ traces differ at least by a factor of two. There can be many reasons why the photon collection efficiency is so different for the two atoms, but it is likely due to non-homogeneous illumination in the optical molasses. In the blue data trace the second point stands out and has about twice as many counts as the equivalent point on the other side of the z_{max} . We attribute this to another atom that hopped into the line of sight in a different plane of the vertical lattice, but was gone again in the next image. This does not seem to affect the PTR value significantly, and shows the robustness of that measure to spurious events like this one.[¶] However, focusing on the purple trace we see that the correct focus position of the trace seems to be mis-identified when the quadratic fit is considered. As stated above the center is identified as the point with the maximum PTR value. As this case demonstrates, one should rather rely on the center of the fit as a measure of the plane into which the given atom was loaded, and this will be done in a future iteration of the analysis. All in all, this discussion shows the power of the technique. By the atom trace $\{a_i\}$ the correct plane of the atom along the line of sight can be determined with a greater certainty than in the approach described in sec. 10.3, where the objective was not moved during imaging.

As shown in the inset of fig. 10.7 (a) we find on average $n_1 = 30(9)$ atoms in the first image of each series. However after the filtering process we are left with 169 atom traces $\{a_i\}$ which corresponds to close to 6% of the total number of atoms detected in all of the initial images. The atom loss during imaging was thoroughly discussed in sec. 9.4, but one could point out especially that a portion of the atoms found in the first image lies outside the region of the third molasses beam, and so they are expected to be heated out quite fast. In

[¶]In future work we hope that such events will be suppressed with better cooling in the molasses, and minimal lattice heating.

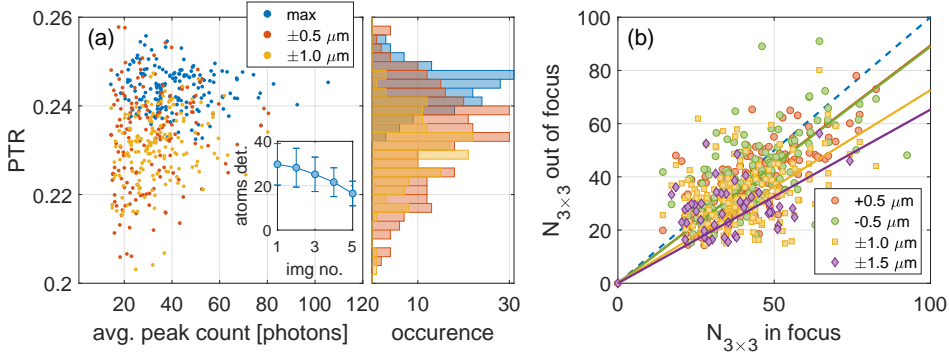


Figure 10.7: Correlations in PTR and peak count. (a) A scatter plot of the PTR and $\tilde{N}_{3 \times 3}$. The histogram of the PTR coordinate is shown to the right. The inset displays the average number of atoms detected each of the five images obtained, and the errorbar denotes the standard deviation. (b) The correlation of $\tilde{N}_{3 \times 3}$, both in and out of focus for the different planes.

fig. 10.7 (a) we show a scatter plot of the PTR and $\tilde{N}_{3 \times 3}$. Like in fig. 10.5 (a) all the points lie above the diagonal of the plot, but the bulk of the data are grouped around 1/3rd of the maximal value of $\tilde{N}_{3 \times 3}$. The blue points correspond to atoms identified at z_{max} , and the orange and yellow data points at $|\Delta \tilde{z}| = 0.5 \mu\text{m}$ or $|\Delta \tilde{z}| = 1.0 \mu\text{m}$ respectively. The histograms on the right to the scatter plot are not as easily distinguishable as those in fig. 10.5 (a), but the separation $\Delta \tilde{z}$ is also much less. The group of the z_{max} count is however considerably narrower than the other groups.

Another way to characterise the data is to find the correlation between the maximum peak count $\tilde{N}_{3 \times 3}$ in and out of focus. This is presented in fig. 10.5 (b). The 1 : 1 correlation is drawn for reference as the blue dashed line. Each data set is fit by a linear function (without an offset) and as expected the slope is reduced for planes further away from the location of z_{max} . In the graph we treat separately the data $0.5 \mu\text{m}$ above and below z_{max} , and as expected they yield exactly the same slope (the green and orange solid lines are practically indistinguishable in the graph).

As a next step we superimposed all the $169 \{a_i\}$ traces with respect to the

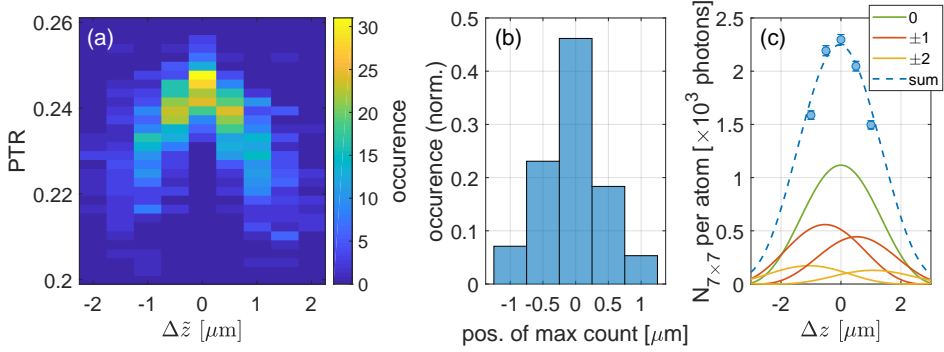


Figure 10.8: Shifting atom traces. (a) Histograms of the PTR for each coordinate $\Delta\bar{z}$. (b) A histogram of the location of the maximum PTR as a function of the objective position. (c) Using that data we model the intensity of the light along the axial coordinate, and compare to measurements.

$\Delta\bar{z}$ coordinate, and a histogram for the PTR value is compiled for each value of $\Delta\bar{z}$. This is represented in the image of fig. 10.8 (a). The image is equivalent to the one that we would obtain if all atoms had been loaded into the same lattice plane. One may now directly compare this image to the main result of the static objective approach, fig. 10.5 (b), which is much broader and more uncertain. We see that the tomographic approach reduces substantially the uncertainty in this distribution.

We are now equipped to answer the question of how the atoms loaded into the vertical lattice were distributed into the respective planes. This is done by attributing the location of the z_{max} as the plane into which it is loaded. The result is displayed in the histogram of fig. 10.8 (b), where the occurrence is equivalent to the population fraction loaded into a given plane. This result can be directly compared to a similar plot of fig. 10.2 (c), which was the result of the simple GPE based theory represented in that section. That particular figure was rendered for $N = 25$ atoms, which is the mean number of atoms detected in all images. The difference between the two situations is however that in the experiment, some evaporation was carried out in the vertical lattice after loading, but

still the simple theory predicts the total number of loaded planes. The asymmetry of the measured population fraction could stem from the phase of the lattice with respect to the center of the atom cloud as was demonstrated in fig. 10.2 (c), given that the phase drift was negligible over the one hour it took to acquire the data. The gravitational sag of the tweezer potential from which it was initially loaded could also play a role. The measurement yields a notably larger fraction in the central plane compared to the prediction, which could be the result of the evaporation.

To a good approximation, we expect each atom to emit light according to the PSF of the imaging system. Along the line of sight its intensity should follow eq. (3.5), the axial shape of the PSF. The total light intensity as measured from the atoms should be equal to the sum of all the individual PSFs, that we denote by $I_i(z)$, for an atom in lattice plane i . As we know the distribution of atoms in the planes along z from fig. 10.8 (b), we can use those numbers as population weights p_i to achieve a normalised total intensity $I_{\text{tot}}(z)$. That function is modelled by

$$I_{\text{tot}}(z) = \sum_i p_i I_i(z - id) \quad (10.6)$$

where $i \in \{-2, \dots, 2\}$ and d is the lattice spacing. The average total number of photons collected per atom is plotted in fig. 10.8 (c). The $I_{\text{tot}}(z)$ is the blue dashed line, and the individual components of eq. (10.6) are also drawn. As the function $I_{\text{tot}}(z)$ is normalised, its height is fit to match the data. This simple theory matches the measurements quite well. In future work, we could improve the comparison to the model by scanning the objective over a larger range.

As a final exercise, we can now tomographically reconstruct the 3D spatial distribution of atoms. This is shown in fig. 10.9 (a) and (b) from slightly different angles. The distribution is rendered by the z_{max} of each of the 169 atoms, and the horizontal coordinate is obtained from the image at which the PTR was found to be at a maximum. This represents an average distribution as it samples about 6% of the atoms present in the first image, and it is compiled using all 100 repetitions of the measurement. However, with improved molasses

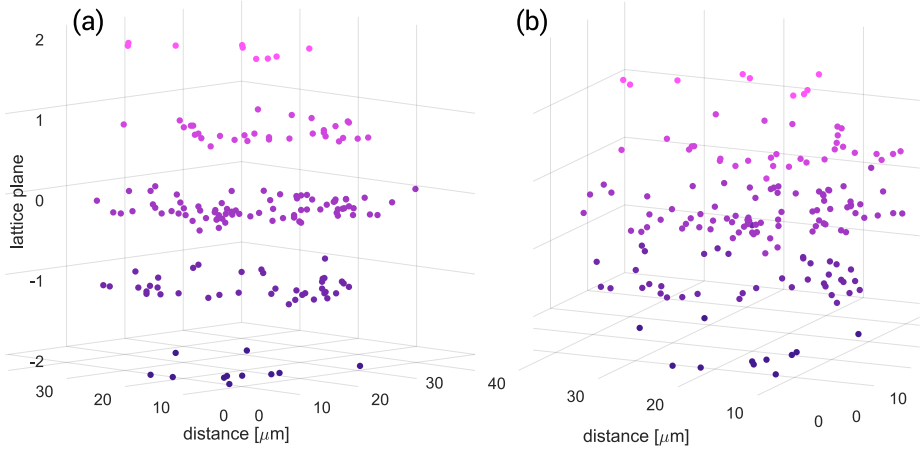


Figure 10.9: A tomographic reconstruction of the atom cloud loaded from a single tweezer into the vertical lattice. The frames (a) and (b) are taken from slightly different angles and show different aspects of the distribution. The colour of the atom depends on the lattice plane.

cooling, and better stability of the optical lattices nothing prohibits such images to be rendered from each single realisation of the experiment.

OUTLOOK

A brave new experiment

The outlook is split in four parts. The first section describes the fire accident, its causes and effects, and briefly accounts for the reconstruction of the experiment and its present status. The second section sketches out different strategies to achieve single plane loading and the third section discusses the feasibility of the single-site spin addressing scheme. The fourth section outlines the longer-term goals, and which promises the machine holds for future physics experiments.

0.1 The rise of the Phoenix

February the 11th 2019 was a particularly bad Monday in the lab. We forgot to turn on the water-cooling for the high-current coils (not for the first time), but this time the interlock that prevents the system from overheating in such situations did not work as intended. We had recently built in two high-current coils around the science chamber. They were run on two separate power supplies (DELTA ELEKTRONIKA, SM 15-400), identical to the one used for the cube (see fig. 6.5). At the time, the interlock channels of each power supply had been hooked up in series to the interlock system of the existing power supply, in a way



Figure O.1: Ashes to Ashes. (a) The author inspects the scene. Robert hides behind the optics. *(b)* The remnants of the cube chamber.

that seemed to work all right, at least when all the power supplies were turned on.* This morning however, we only needed one power supply.

The high-current coils used for microwave evaporation overheated and melted and burnt the epoxy they were mold in. There was probably fire in the lab for some time, as the epoxy went up in flames. Luckily it did not spread to any electric cables or other plastic material that is present in vast amounts above the experiment table. Water hoses broke in the heat and sprinkled the optics table. Vacuum broke and all chambers were filled with soot. Fire, water, and air filled chambers. The elemental forces of nature drove the experimentalist to the dust.†

*We came to learn that this was of course by no means the right way to do things. Both were the interlocks connected in the wrong way, but more importantly, the functionality of the interlock had not been properly tested after the alterations. Always test your fail-safes in *all* modes of operation. Else your machine might spontaneously combust.

†My grandfather's only lesson to me on physics became appropriate (in a very loose translation).

Hvert orð þitt kvelur mig,
hvar sjást þín þín gæði?
Andskotinn eigi þig
eðlisfræði.

Your words torture,
where are your qualities?
Go to hell
physics.




Figure O.2: The death of an experiment. The machine faithfully creates BECs until its last breath. The dark clouds on the horizon eventually reach the vacuum chambers and suddenly their insides turn optically thick.

The scene was frightful, as fig. O.1 shows so vividly, and it was immediately clear that this incident would set back our progress for a long time (at least on the timescale of a PhD student). The machine recorded images in the science chamber during the warmup routine, and it managed to make 22 BECs before things started to fail, as is shown in fig. O.2. In the 64th shot (after about 40 minutes of continuous operation) the OD in the chamber changes abruptly. That marks the point of the vacuum failure.

There was only one way ahead: to mend what was broken. Luckily all the high-resolution optics as well as the optics around the science chamber were intact, apart from the high-resolution viewport that had grown a thick soot-doped epoxy coating. All vacuum components had become extremely dirty on the inside, but with the help of dirt-eating bacteria, organic solvents and great assistance from the mechanical workshop, we could clean almost all of the vacuum parts to XUV grade. The optics and viewports were carefully rinsed with hundreds of cotton swabs and (m)ethanol. Even the SAES getter pumps were functional after several bakeouts and cycles of getter-activation. The cube however had to be replaced. After cleaning, the experiment was baked out in two parts outside of the lab. First we baked the MOT section, and later the cube and science chambers were baked together (see fig. 3.1). We back-filled the MOT section with argon and then the two parts were joined via a gate valve connected after the cube chamber. Afterwards we could pump the system down to same vacuum level as before. At the time of writing we have atoms anew, with a couple of months of work still ahead of us. The new experiment is appropriately nick-named *The Phoenix*.

O.2 Loading a single lattice plane

The ability to extract reliable results from a quantum gas microscope experiment depends crucially on the imaging fidelity of single atoms, which is in all cases above 95%. The depth of focus of our imaging objective is close to $1.5\text{ }\mu\text{m}$, which naturally renders the resolution of dense atom clouds along the line of

sight challenging. This is why all quantum gas microscope experiments have exclusively worked with atoms in a single plane, with minor exceptions [225]. Today, different approaches are used to achieve the loading of a single lattice plane. Methods include magnetic field slicing [243] (explained below); trapping based on the combination of an evanescent wave and a large spacing optical lattice [115]; compression of a degenerate cloud by means of an accordion lattice [139]; or loading of a large spacing lattice from a tightly focussed and highly elongated optical tweezer [208] (supp. mat.). The last method suffers from the fact that atoms do spill over into neighbouring lattice planes. They are however subsequently removed by selectively transferring them via a local RF addressing scheme under a magnetic field gradient to states where spin-changing collisions deplete the population.

Our tomographic approach presented in ch. 10 has obvious limitations when it comes to imaging dense samples in multiple planes. It seems for instance to be impossible to precisely determine the location of holes in a multi layer Mott insulator with a filling fraction $n < 1$. However, its strength lies in analysis of few-atom systems distributed over a handful of lattice planes along the line of sight. As an example of an experiment that fulfils such conditions would be the realisation of quantum random walks in 3D, something that has been investigated before in 1D [224]. The majority of the experiments that have been realised in quantum gas microscopes, also rely on rather small populations. With some further work, especially in terms of homogenising the number of photons collected from each atom, this method could be adapted to work for precise atom counting along the line of sight in a few layers.

Despite these prospects, it is still one of our immediate goals to achieve the loading of a single plane, and in what follows, different means to that end are outlined.

Magnetic field slicing

The method of magnetic field slicing has been successfully employed for single plane loading, as described in ref. [243]. Figure O.3 (a)–(f) explains the process.

An atom cloud in the $|F = 2, m_F = 2\rangle$ hyperfine state, is trapped in a 1D optical lattice with a lattice spacing d as shown in (a). By applying a magnetic field gradient $\frac{dB}{dz}$ (b), the energy levels are shifted dependent on the clouds position by the Zeeman effect. The atoms are exposed to microwave radiation (c) which is adjusted to be resonant only with a single lattice plane, transferring them to the $|1, 1\rangle$ state (shown in blue). To get rid of all but one atom layer a light pulse, resonant with the $F = 2 \rightarrow F' = 3$ transition, pushes all atoms in the $|2, 2\rangle$ state out (d). Now the single layer can be shifted back to $|2, 2\rangle$ state (e) and we end up an atom cloud in a single layer (f).

As we will see, this approach both requires the stabilisation of current in the gradient coils to a high degree and for a homogeneous loading the flatness of the magnetic field at the atoms is also important. To analyse this in more detail let us calculate the total magnetic field (in cylindrical coordinates) at a distance $(\Delta r, \Delta z)$ from the centre of the quadrupole field. By assuming only linear contributions (which is reasonable close to the centre of the trap), we have that

$$B(\Delta r, \Delta z) = \sqrt{\left(\frac{dB}{dz}\Delta z + B_z\right)^2 + \left(\frac{dB}{dr}\Delta r + B_r\right)^2}, \quad (\text{O.1})$$

where the $\frac{dB}{dz}$ and $\frac{dB}{dr}$ are gradients in the vertical and radial directions, and B_z and B_r are offset fields.

Now we impose the condition of a homogeneous field in the centre of the trap, which will enable loading of a Mott insulating state of a radius r_0 in a single lattice plane. We take the limiting condition to be the point where the field strength at a radius r_0 in plane number n is equal to the field strength in the centre of the neighbouring plane $n + 1$, i.e. $B(\Delta r + r_0, \Delta z) = B(\Delta r, \Delta z + z_0)$.[‡] Solving for B_z gives the vertical offset field required to achieve the condition as

[‡]It can of course be debated if this is exactly the right condition, but it is correct at least up to a constant front factor.




Figure O.3: Magnetic field slicing. (a)–(f) Preparation of atoms in a single layer of an optical lattice by magnetic field slicing. (g) Stability requirements for the magnetic field gradient used for slicing, and hence the current in the quadrupole coils. here are two requirements that must be met, first a given size of the Mott insulator must be reached, and second the background field fluctuations set a limit for the minimal gradient. Details are given in the text.

a function of the gradient. If we recall that $\frac{dB}{dz} = 2 \frac{dB}{dr}$ we obtain

$$B_z(\Delta r, \Delta z) = \frac{1}{2z_0} \frac{dB}{dz} \left(\frac{1}{4}(r_0^2 + 2r_0\Delta r) - (z_0^2 + 2z_0\Delta z) \right)^2 + r_0 B_r. \quad (\text{O.2})$$

The fixed point in our experiment is the location of the objective, so all other entities must be aligned to that point. The quadrupole coils around the science chamber are placed directly on the flanges of the vacuum chamber. Their alignment with respect to the centre of the chamber is achieved by mechanical means only (not involving the atoms) and as the atom cloud is located slightly below the centre, it will not be in the physical centre of the anti-Helmholtz system. We carried out some preliminary tests of the quadrupole coils while the experiment was still intact, and we found that the coils were offset radially by about 1 mm and vertically by 2.5 mm.[§] By controlling the coils independently using

[§]We have now altered the coil holders such that their mechanical alignment will improve. They are also now closer to the atoms which will allow for higher gradients.

separate power supplies, we can shift the location of the zero field vertically to coincide with the location of the tweezer-trapped clouds above the objective, but the radial displacement is hard to fix.

From the first half of the thesis we found shot-to-shot fluctuations of stray magnetic fields to be on the level of $\sigma_B = 0.5$ mG. In fig. O.3 (g) we plot the ratio of B_z/σ_B , assuming $\Delta r = 1$ mm, $\Delta z = 0$ mm and $B_r = 0$ mG. If we demand that the difference in the magnitude of the magnetic field between adjacent planes (ΔB), is an order of magnitude greater than σ_B , this will require a gradient

$$\frac{dB}{dz} = 10 \frac{\sigma_B}{d} = 100 \text{ G/cm.} \quad (\text{O.7})$$

For comparison, ref. [85] reports 50 G/cm. For a homogeneous Mott insulator of only $r_0 = 5$ μm this puts us close to a fractional current stability of 10^{-5} . To build such a power-supply will require some electrical engineering. A recent approach reaches a fractional stability of high currents to the level of 10^{-6} [267], which would definitely fit our needs.

Large-spacing lattice from the side

Interference patterns of lattice beams with a larger spacing than $\lambda/2$ can be created by introducing an angle 2θ between the beams.[‡] This will result in an optical lattice with a spacing of

$$d = \frac{\lambda}{2 \sin(\theta)}, \quad (\text{O.3})$$

as exemplified in fig. O.4 (a). We could foresee the implementation of such a system using a dual microscope system from the side (see fig. 3.6). To minimise the lattice spacing, θ must be large and will ultimately be limited by the NA

[‡]We prevent this from happening in our current lattice beams by shifting their optical frequencies by 10's of MHz relative to one another.

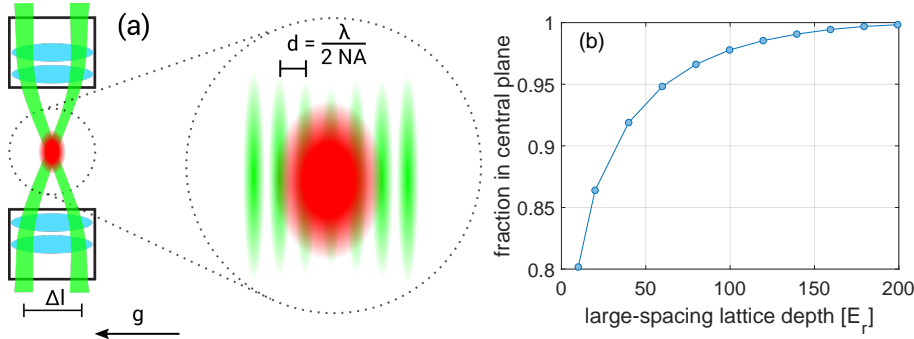


Figure O.4: Loading a large spacing optical lattice. (a) A schematic figure of the creation of a large spacing lattice using the dual microscope setup. The g -arrow points in the direction of gravity. (b) The fraction of atoms loaded from a BEC in the large spacing lattice to a single node of the short spacing lattice as a function of the depth of the large lattice spacing. Parameters for the calculation are given in the text.

of the objective. For the most extreme rays $\sin(\theta) = \text{NA}$. The objectives are estimated to have an effective $\text{NA} = 0.19$ (see sec. 4), but as the atom cloud is located below the centre let us assume an $\text{NA} = 0.15$. This will produce a lattice with a spacing of $d = 3.5 \mu\text{m}$.

Such an optical lattice could be used as an intermediate step in a preparation sequence for a single plane. With that tool at hand we have two options.

- (i) If the system is designed such that the distance between the incoming beams (Δl), can be varied (this can be achieved e.g. with an AOM), the spacing in the lattice can also be varied. By reducing Δl to 1/10th of its maximal value, the spacing becomes $d \simeq 30 \mu\text{m}$, which would suffice to engulf the BECs we create in our CDT (see fig. 9.6). The BEC could then be compressed by increasing Δl . This could maximally yield about 10^5 atoms in a single plane.^{||}

^{||}There is some flexibility in the length scales here. First if $\Delta l = 30 \mu\text{m}$ is not easily realised, the BEC could both be made smaller by further evaporation, and subsequently compressed in the CDT before loading over to the lattice.

- (ii) We could rely on the short Rayleigh range of the optical tweezers and load atoms from an array into a single anti-node of the large-spacing lattice. According to fig. 8.5 (a), this method should yield $2 \cdot 10^4$ atoms in a single plane.

Both of these schemes might introduce heating in the cloud such that the BEC might be lost, or at least the condensate fraction somewhat reduced. To reliably load from the large-spacing lattice over to the short-spacing lattice, the extent of the atom cloud must not exceed 532 nm. Using GPELab [14], we calculated the spatial extent of a BEC of 10^4 atoms in a harmonic potential assuming a beam waist of $w_0 = 100 \mu\text{m}$ and a lattice spacing of $d = 3.5 \mu\text{m}$, where the aspect ratio of the trapping frequencies is $\frac{\omega_z}{\omega_r} = \frac{\pi}{\sqrt{2}} \frac{w_0}{d} = 63$. Figure O.4 (b) shows the calculation of the atom fraction that ends up in the central plane of a short-spacing lattice (assuming that an anti-node of the lattice coincides with the BEC), as a function of the depth of the large-spacing lattice. The remaining atoms are distributed equally to the nearest adjacent planes. We see that already at the depth of $60 E_r$, around 95% of the atoms populate the central plane of the short spacing lattice.

Light-sheet-aided microwave transfer

The method described below is in many respects the same as the magnetic field slicing scheme discussed above. In fact, the schematics presented in fig. O.3 (a)–(f) is applicable, but the spatially dependent energy shift is not created by a magnetic field. Instead the shift is due to an optical potential. For this to work out, two conditions must be met. (i) The light field must be *differential* in the states of the ground state manifold, $F = 1$ and $F = 2$, which means that in its presence the energy levels of the two states are shifted by different amounts. (ii) The difference in the light shift between the main plane and the neighbouring plane must be so large that the transfer pulse (in terms of frequency), acts exclusively only in one plane.




Figure O.5: Spin flips in a light-sheet. (a) The light shift for the $|F = 1, m_F = 1\rangle$ and $|F = 2, m_F = 2\rangle$ hyperfine states due to σ_+ and σ_- polarised light. (b) The differential light shift between $|F = 1, m_F = 1\rangle$ and $|F = 2, m_F = 2\rangle$ states as produced by a σ_+ polarised beam at a wavelength of 787.6 nm made with a $NA = 0.15$ objective. The blue solid line corresponds to a transversal cut through the light-sheet intensity, as a function of the distance from its center. Parameters for the calculation are given in the text. The vertical lattice potential is drawn as the solid orange line. Its depth is irrelevant in this context. The inset shows the difference in the light shift between the green filled circles as a function of the NA.

The ideal light color to achieve a high differential light shift U_{diff} , in ^{87}Rb lies between the D1 and D2 lines. Figure O.5 (a) shows the AC stark shifts for the two hyperfine states $|F = 1, m_F = 1\rangle$ and $|F = 2, m_F = 2\rangle$ in the presence of σ_+ and σ_- polarised light. The expression for the shift is similar to eq. (I.14) but takes into account the polarisation states of the light. That variant can be found in sec. II.B of ref. [121]. Of particular interest here is the “magic” wavelength for the $|F = 1, m_F = 1\rangle$ state where the contributions from the D1 and D2 lines cancel one another, that is roughly at a wavelength of 787.6 nm in a σ_+ polarised light field.

Let us now assume that we can generate a light-sheet from the side using the AOD and the dual microscope system of the magnetometry experiments. As in the previous section, the beam is modelled as a diffraction-limited Gaussian beam at $NA = 0.15$, and calculated according to the criteria discussed in sec. 3.2. That calculation is shown as the blue solid line in fig. O.5 (b), where we have as-

sumed a power of 10 mW in the beam (comparable to what we used for the microtraps in the magnetometry experiments), and a wavelength of 787.6 nm. The horizontal axis denotes distance along the vertical direction and its origin is chosen in the centre of the light-sheet potential. The vertical lattice potential (short-spacing) is also drawn as the orange solid line, and the positions that correspond to an anti-node are also drawn with filled circles on the solid blue line. We are interested in the difference in the light shift between the central plane and the adjacent one ΔU , marked as the green circles in fig. O.5 (b). For this particular calculation we get that $\Delta U = 3.5 \text{ MHz} \times \hbar$. The inset shows how this result changes when we relax the condition of the NA. Even at the old NA of 0.11, the $U_{\text{diff}} = 1.0 \text{ MHz} \times \hbar$. In ref. [280], the author used a so-called HS1 pulse (details in ref. [107]) with a width of about 60 kHz. Assuming that we can use similar parameters, the scheme as outlined above is promising. The main challenge of this method will be to stabilise the light-sheet to a particular lattice node with $\sim 100 \text{ nm}$ precision.

Evaporation strategies

Regardless of the strategy we choose for single plane loading we will need to perform subsequent forced evaporation in the lattice. This will serve both to cool the atoms and stabilise the atom number, as well as to get rid of any atoms that might still be in neighbouring planes. As discussed in ch. 9 we found that evaporation in the vertical lattice enabled us to remove some atoms from neighbouring planes and clearly image individual atoms. If the population difference between planes is great enough, one's intuition is that an optimal evaporation ramp for the main plane is not at all optimal for the adjacent planes, so atoms there will be predominantly lost. Such a purification effect is also hinted at in ref. [280]. However, evaporation in the lattice alone will only cause the hottest atoms to fall downwards, risking that the falling atoms would be trapped in the lattice potential again further downstream.

As reported in appendix A of ref. [85], the magnetic slicing scheme is followed by evaporative cooling in the vertical lattice. The cooling process is fa-




Figure O.6: The differential light shift of a 787.6 nm beam through the high-resolution objective, intended for site-resolved spin-addressing.

cilitated by a magnetic field gradient in the horizontal direction that, together with an optical trap with about $10 \mu\text{m}$ waist at the plane of the atoms, creates a controllable potential barrier. Any atoms that remain in adjacent planes of the vertical lattice should also be removed in this process. We could implement such a gradient either by means of a horizontal gradient field or via an optical gradient imaged onto the atoms through the objective with the DMD.

Another way of using the DMD to evaporate the atoms in the lattice, would be to project tweezer potentials in a ring around the optical lattice beam to enable a flow of the hottest atoms radially outwards into them, in order to realise so-called ghost-beam evaporation in the style of ref. [68]. Atoms that flow into the tweezer are then pulled downwards by gravity, and away from the vertical lattice.

0.3 Single-site spin addressing

Let us briefly discuss the feasibility of our single-site spin addressing scheme. Figure O.6 shows a similar calculation as was presented in fig O.5 (b). Here, however, the blue solid line represents the light coming from the 787 tweezer

system that was shown in sec. 8.3, to have a waist of $w_{787} = 780$ nm, used in this calculation. The curve in the figure is rendered for a power of $6 \mu\text{W}$ in the beam. For reference, the horizontal lattice is drawn as the solid orange line.

After the fire, we took the opportunity to measure the optical powers we could get in the DMD-generated 787 beams using the laser system shown in fig. B.2 (left). The highest power we measured was $0.3 \mu\text{W}$ in the image plane, using the Fourier DMD to create a single tweezer. This number seems horrifically low, but it is due to both low efficiency of the DMD systems in general, and the fact that the 787 arm is combined with the imaging path using a glass plate where 95% of the light is lost. The direct DMDs were worse, simply because fewer mirrors reflect light. To attain reasonable depths, we can boost the optical power by incorporating a tapered amplifier in the system. Comparing it to the 940 laser system, this should give us at least a factor of 20 in the maximum optical power or about $6 \mu\text{W}$ in the Fourier arm. That should bring us safely into the regime where the spin of a single atom can be selectively flipped with high fidelity.

0.4 Future experiments

In ch. 10 of this thesis we demonstrated how to reconstruct the 3D distribution of atoms trapped in deep optical lattices, as measured in a quantum gas microscope experiment. By scanning the position of the high-resolution imaging objective we can measure which plane of the optical lattice an atom was loaded into along the line of sight. This is an important step towards the realisation of 3D quantum many-body physics with quantum gas microscopes. These novel experimental tools have in the past decade also proven to be extremely powerful for the purpose of quantum simulation, as was detailed in the introduction of this thesis.

There are many paths that can be taken with an experiment such as ours. With a deterministic way of preparing a single plane, the limits of the tomographic approach can first of all be further investigated. As mentioned earlier

in the outlook, one application would be to track atoms in a 3D quantum random walk. We could prepare a 2D structure (such as a 3-by-3 array of atoms), and allow them to diffuse in the 3D lattice. To access non-trivial dynamics in this setting we could tilt the lattice, by a magnetic field gradient, or even modulate it to study diffusion of higher motional states. For detection we would rely on the tomographic imaging.

The topic of quantum optimal control lies close to the heart of our research group. The flexibility of the experiment will enable studies in protocols for quantum information processing. Joining those two fields is an essential part in making quantum technologies tenable in the 21st century [1, 67]. By relying on DMD aided state preparation [158], both in terms of occupation and spin, we could realise spin networks [177]. Along these lines one could study the transport of a spin in a ring (router) or in a chain (link) and apply methods of optimal control to the process by shaping potentials via the DMDs, optimising transport fidelity, speed, etc. We could also study transport of single atoms directly using the dynamic capabilities of the DMD, as we demonstrated with small atom clouds in sec. 8.5. The transport of single atoms is an essential ingredient in certain realisations of digital quantum computations with neutral atoms in optical lattices [282]. The elementary quantum logic gates used in that proposal are also applicable for quantum optimal control studies [149].

Apart from the ideas discussed in the outlook at the end of ch. 5, there are other routes that involve weak measurements. Until now, weak measurements have been employed in an optical lattice system to suppress tunnelling in a quantum gas by the quantum Zeno effect [216], or to study the dissipative Bose-Hubbard model [271]. To observe such dynamics in combination with single-site resolution would be fascinating. This seems viable with the flexible light-shaping tools at hand in our system, that could even allow for control the phases of the system via local measurements [259].

As a final point I would like to mention the idea of opening a state-of-the-art quantum simulation experiment to other experts in the field, or even to the public. This idea is not entirely unfamiliar, as was briefly discussed in the preface. In the Alice Challenge, we enabled remote access to our experiment through a

gamified control interface [132]. In that case people were allowed to optimise the BEC size in our experiment, but there is nothing that hinders similar studies to be carried out with more complex and intricate quantum systems.

APPENDIX A

The ^{87}Rb level scheme

All near-resonant interaction between the ^{87}Rb atoms and coherent light, used in our laboratory happens on the D2 line. In the level scheme reprinted on the next page, the interaction of the nuclear angular momentum \mathbf{i} and the total electronic angular momentum \mathbf{j} is taken into account. The energy splitting due to the linear Zeeman effect is marked in parentheses at each hyperfine state.

The frequencies of the our near-resonant lasers, that we use for cooling and imaging are drawn in the figure. In particular, the master laser is locked to the $F' = 2$ to $F' = 3$ crossover, and detuned by 300 MHz. That light is redistributed for various purposes and pulled to the right frequency.

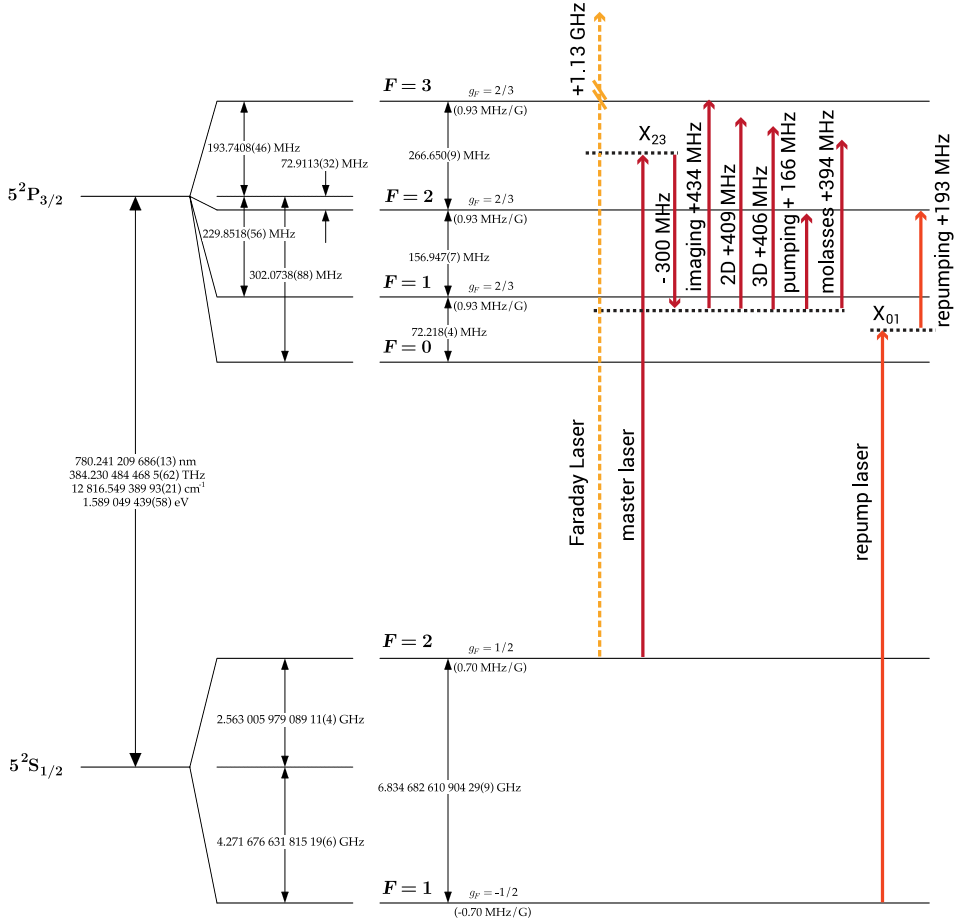


Figure A.1: The atomic level structure of ^{87}Rb , on the D2 transition line. The frequencies of the relevant lasers and laser control systems have been drawn on top. Reprinted with permission from [252], and modified.

APPENDIX B

Laser systems

This appendix has schematic images of four laser systems. A legend for the different symbols is shown in fig. B.1. In fig. B.2 the systems that distribute the light to the 787 (left) and 940 nm (right) laser systems are shown. Each system has three different outcouplers that lead to the high-resolution breadboard, two for the direct imaging DMDs and one for the Fourier DMD. To start with we both had a home-built laser and less optical isolation (36 dB) between the laser source and the tapered amplifier in the 940 system. This gave rise to great instabilities (both in intensity and frequency) in the laser system so we doubled the isolation. In the end we also swapped the home-built source with the Topica laser that was previously used for the 912 nm microtraps in the magnetometry. The laser was easily tuned to 940 nm.

The molasses setup is shown in fig. B.3 (top). The AOMs operate around 200 MHz, and are set up in a double pass configuration. A part of the light is also used for the absorption imaging on the lattice axes. The high-power laser setup is shown in fig. B.3 (bottom). The Nufern amplifier provides light for the two beams of the crossed dipole trap and the vertical lattice beam. The Azur amplifier provides light for the horizontal lattice axes.




Figure B.1: A legend for the laser setups found in ch. 6 and in this appendix.

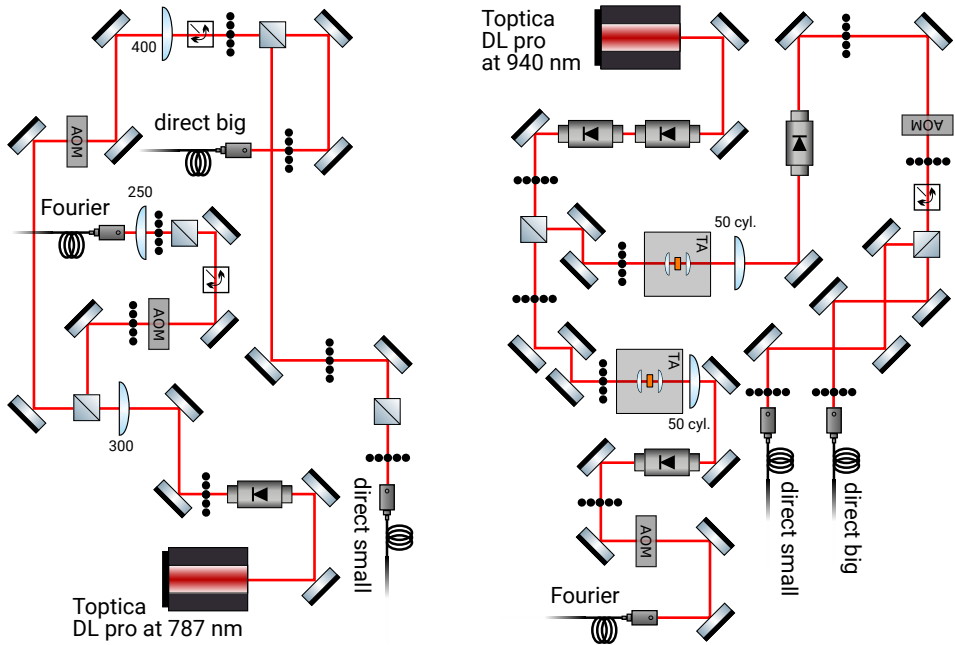


Figure B.2: (left) The 787 nm laser distribution system. (right) The 940 nm laser distribution system.

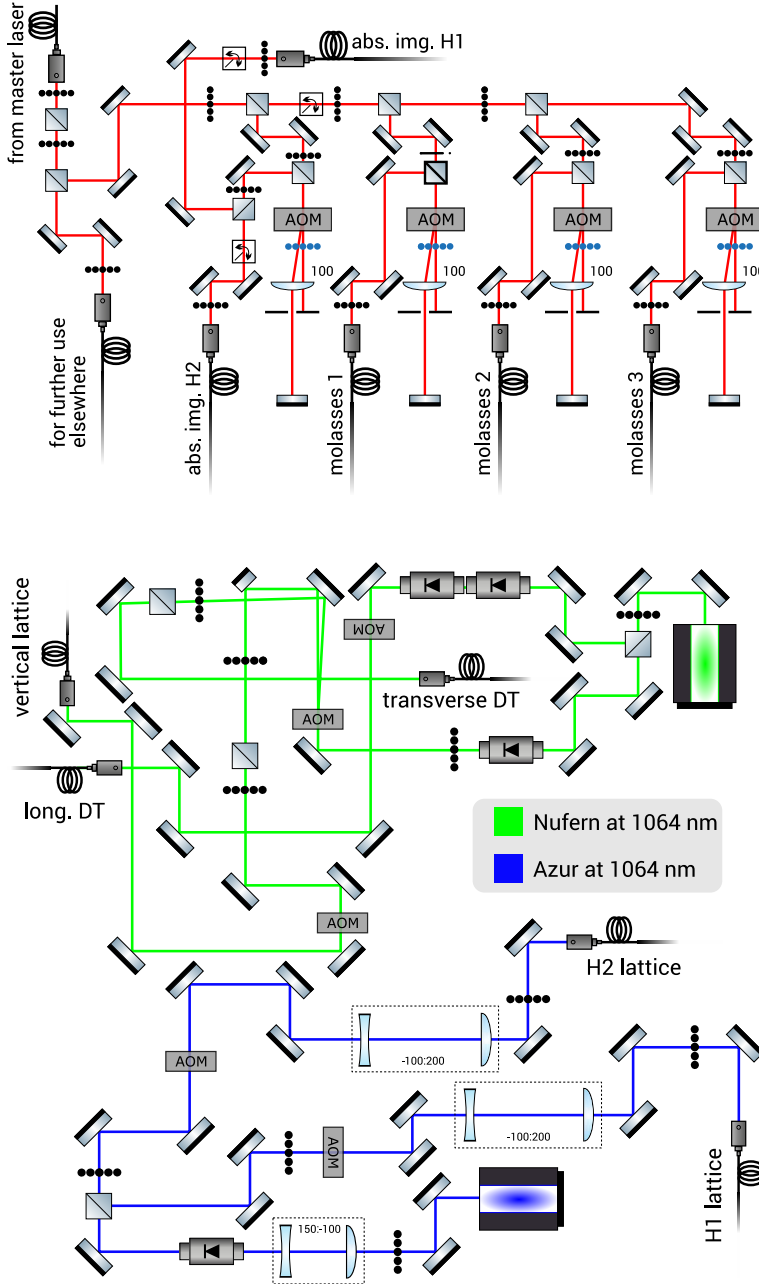


Figure B.3: (top) The molasses laser distribution system at 780 nm. A part of the light is also used for imaging. (bottom) The high-power 1064 nm laser system.

The logging server

All laboratories continually measure a multitude of quantities signifying the health status of their machinery. In our case these are laser powers, temperatures around the lab, the temperatures of the climate system sensors, the water flow and temperature of the high power magnetic field coils and even the magnetic field background. To ease the monitoring of all those numbers we run an online monitoring system which is available from the internet through a web browser.

The system is run from a server (**Ubuntu Linux**, 16.04 LTS) hosted by the IT department of Århus University. On the logging server we have an SQL style time series database (**InfluxDB**, v. 1.7.4) for writing, reading and the manipulation of data. There is also a web application (**Chronograph**, v. 1.7.4) that graphs data uploaded to the database. Both are freely available from influxdata.com, see documentation [144]. The database automatically coarse-grains old data sets as the amount of data can become quite excessive. The data can be downloaded from the server for further analysis if necessary.

The logging of all log-able quantities is handled in LabView and via its native LabViewHTTPClient we can directly upload data using the InfluxDB API



Figure C.1: A snapshot of the online Chronograph server, used for logging in our experiment. The graph displayed is the readout from the temperature sensors of the climate system over a period of one week. The vertical axis is not calibrated to temperature. The bottom line monitors the cooling part of the system, the middle one the heating part and the top one is the regulated temperature measured on the experiment table.

protocol. A snapshot for the page we use to monitor the climate system is shown in fig. C.1. Please note that the vertical axis is not calibrated to temperature.

Notes on our optical lattice setup

The optical lattice setup has seen quite some changes since it was first constructed. Here I make a couple of comments on the importance of fibre coupling high-power lattices and on the intensity regulation system.

To image the fluorescence from single atoms we need deep lattices. In the first iteration of constructing the optical lattices we decided to place the laser head of the Azur laser (that provides light for the two horizontal lattice axes), directly on the experiment table to lose as little power as possible. As a result the laser setup with optical isolators, AOMs for fast switching and all the beam shaping optics were fitted around the science chamber. The laser head had to be placed on one side of the laser table, resulting in quite some difference in the lengths of the beam paths of the H₁ and H₂ lattices, where the H₁ path was close to being 1 m longer. We worked with this lattice design for 7 months (learnt how to make Kapitza-Dirac scattering and modulation spectroscopy, observed the superfluid to Mott insulator transition and even saw faint single-atom signals) but due to constant drifts in the alignment and obvious shot-to-shot jitter in the depth of the horizontal lattices (especially on the H₁ axis), we decided to couple the output of the laser through optical fibres. That did indeed remedy all pointing issues and shot-to-shot depth fluctuations we had with the lattices.

Always fibre couple your lattices.

To regulate and control the laser powers we use an AOM driver (providing RF signals to the lattice AOM) coupled to a PI-circuit both made by the Institute's electronic workshop. The transmitted part of the lattice light through one of the mirrors in the setup, is measured on a photodiode (THORLABS, PDA36). Soon after we started operating the optical lattices we realised that there was something rotten in the state of Denmark. Troubles arose when we tried to regulate the lattice at low powers. The system only responded to a control signal with an excessive delay of several 100's of μs and the following ramp up was also limited to such a timescale for low regulation signals. For high signals the response was much better, but still with a delay. We found out that output of the RF amplifier in the AOM driver had a logarithmic response in the control voltage, rendering it difficult to control low voltages. The solution was to use a different amplifier that had a linear response in the control voltage. That amplifier did however not provide enough RF power for the AOMs (they require up to 2 W to reach high diffraction efficiency), so the internal amplifier was bypassed and a constant gain amplifier (MINI-CIRCUITS, ZHL-I-2W-S+) was added afterwards. After this fix we could well regulate the low-power end of the lattice. However, the regulation circuit still failed to turn on the lattices smoothly. Every time the TTL signal to the AOM driver is opened the intensity spikes. To prevent this from affecting the delicate ultracold cloud we open the TTL much earlier in the sequence, before we start the evaporative cooling process in the CDT. The downside is that this results in a small amount of leak light through the fibre, that give maximally a depth of $0.5 E_r$ on the vertical lattice axis. Adding a D part to the regulation loop might remedy the issue.

At the current stage there is only one step in the regulation. We can either decide to be able to regulate the lattices at high or low power, by switching the gain of the photodiode manually. This means that in a single experiment, we cannot both regulate low powers precisely and the high powers for ramping really deep and imaging. Our hope was that under normal circumstances we would always work in the low power regulation regime and when we needed to go deep for imaging, we would simply ramp the lattices to saturation level of the photodi-

ode. At that stage the light intensity controller maximises the output intensity so the lattices are ramped to maximum power. This methodology has resulted in a few broken photodiodes that probably do not tolerate the high laser intensities at high gain. We never reached the point of validating if this method would work or if it would induce too much heating. Most certainly it is a simple approach compared to other high-power optical lattice designs recently implemented in experiments performing state-of-the-art quantum simulation [190], but time will tell if this will be sufficient.

The atomic Mach pendulum swing

During our first attempts to load the tight tweezers from below, we encountered something we did not quite understand at first. We had been experimenting with imposing the tweezer potential into the reservoir, and afterwards we allowed it to expand in a waveguide formed by the LDT beam. When we dropped our cloud into TOF we saw it wiggling over the computer screen like a snake, as is shown in fig. E.1 (left). After some consultation with more experienced people, we figured that we had excited the trapped cloud into oscillation. Due to the propagation of a Gaussian beam, different portions of the cloud will experience different trapping frequencies. This is very similar to the so-called Mach pendulum swing, where a series of pendulums that hang from a bar in strings of different lengths, will oscillate at different frequencies. The number of wiggles along the pendulum swing will vary as a function of the time after the onset of the oscillation. If one is patient enough the initial configuration (all pendulum in a straight line) will revive.

To study this a bit we released cold atoms from the CDT into the waveguide. Immediately after, the intensity in the LDT was doubled for 1 ms and then put back to the former value. This sharp intensity step put the cloud into oscillation, and as the trap frequency varies along the beam the atoms will os-




Figure E.1: The formation of an atomic snake. (*left*) Two snakes slither in our vacuum chamber. The two different images are taken after TOF, for two different hold times in the waveguide. (*right*) The frequency as measured as a function of the position along the LDT beam. The blue shaded region is the 1σ error of the trap frequency.

cillate at different frequencies. We scanned this hold time, and fit a sinusoid to the trace formed by the vertical position of the cloud (in a narrow region of interest shifted along the LDT beam) and the time coordinate, and extracted the trap frequency. The result is shown in fig. E.1 (right). The trap frequency clearly changes along the beam, but on top we see additional wiggles in the trace. We believe that is due to coupled motion of the atoms along both the vertical direction, and the direction into the figure (so to speak).

Bibliography

- [1] Acín, A., I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, and F. K. Wilhelm (2018). The quantum technologies roadmap: A European community view. *New Journal of Physics*, 20(8):080201.
- [2] Alberti, A., C. Robens, W. Alt, S. Brakhane, M. Karski, R. Reimann, A. Widera, and D. Meschede (2016). Super-resolution microscopy of single atoms in optical lattices. *New Journal of Physics*, 18(5):053010.
- [3] Albiez, M., R. Gati, J. Fölling, S. Hunsmann, M. Cristiani, and M. K. Oberthaler (2005). Direct Observation of Tunneling and Nonlinear Self-Trapping in a Single Bosonic Josephson Junction. *Phys. Rev. Lett.*, 95(1):010402.
- [4] Alon, O. E., A. I. Streltsov, and L. S. Cederbaum (2008). Multiconfigurational time-dependent Hartree method for bosons: Many-body dynamics of bosonic systems. *Physical Review A*, 77(3).
- [5] Alt, W. (2002). An objective lens for efficient fluorescence detection of single atoms. *Optik*, 113(3):142–144.
- [6] Alt, W., D. Schrader, S. Kuhr, M. Müller, V. Gomer, and D. Meschede (2003). Single atoms in a standing-wave dipole trap. *Physical Review A*, 67(3).
- [7] Amarel, S. (1966). On the mechanization of creative processes. *IEEE Spectrum*, 3(4): 112–114.
- [8] Andersen, H. K. (2008). *Bose-Einstein Condensates in Optical Lattices*. PhD thesis, Aarhus University, Aarhus.
- [9] Anderson, M. H., J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell (1995). Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, 269 (5221):198–201.

- [10] Andor (2019). Comparing EMCCD, ICCD and CCD Cameras | Andor Learning Centre. Link: <https://andor.oxinst.com/learning/view/article/ccd,-emccd-and-iccd-comparisons>. Accessed: 2019-08-08 12:36:40.
- [11] Andor (2019). What is an Electron Multiplying CCD (EMCCD) Camera? | Andor Learning Centre. Link: <https://andor.oxinst.com/learning/view/article/electron-multiplying-ccd-cameras>. Accessed: 2019-08-08 12:25:55.
- [12] Andrews, M. R., M.-O. Mewes, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle (1996). Direct, Nondestructive Observation of a Bose Condensate. *Science*, 273(5271):84–87.
- [13] Andrews, M. R., D. M. Kurn, H.-J. Miesner, D. S. Durfee, C. G. Townsend, S. Inouye, and W. Ketterle (1997). Propagation of Sound in a Bose-Einstein Condensate. *Physical Review Letters*, 79(4):553–556.
- [14] Antoine, X. and R. Duboscq (2014). GPELab, a Matlab toolbox to solve Gross–Pitaevskii equations I: Computation of stationary solutions. *Computer Physics Communications*, 185(11):2969–2991.
- [15] Appel, J., P. J. Windpassinger, D. Oblak, U. B. Hoff, N. Kjærgaard, and E. S. Polzik (2009). Mesoscopic atomic entanglement for precision measurements beyond the standard quantum limit. *Proceedings of the National Academy of Sciences*, 106(27):10960–10965.
- [16] Arute, F. et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510.
- [17] Ashcroft, N. W. and N. D. Mermin (1976). *Solid State Physics*. Brooks/Cole, Belmont, CA.
- [18] Ashkin, A. (1970). Acceleration and Trapping of Particles by Radiation Pressure. *Physical Review Letters*, 24(4):156–159.
- [19] Ashkin, A. (2000). History of optical trapping and manipulation of small-neutral particle, atoms, and molecules. *IEEE Journal of Selected Topics in Quantum Electronics*, 6(6):841–856.
- [20] Bakr, W. S., J. I. Gillen, A. Peng, S. Fölling, and M. Greiner (2009). A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice. *Nature*, 462(7269):74–77.
- [21] Bakr, W. S., A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner (2010). Probing the Superfluid–to–Mott Insulator Transition at the Single-Atom Level. *Science*, 329(5991):547–550.

- [22] Barredo, D., S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys (2016). An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science*, 354 (6315):1021–1023.
- [23] Barredo, D., V. Lienhard, S. de Léséleuc, T. Lahaye, and A. Browaeys (2018). Synthetic three-dimensional atomic structures assembled atom by atom. *Nature*, 561(7721):79–82.
- [24] Bason, M. G., R. Heck, M. Napolitano, O. Elfasson, R. Müller, A. Thorsen, W.-Z. Zhang, J. J. Arlt, and J. F. Sherson (2018). Measurement-enhanced determination of BEC phase transitions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 51 (17):175301.
- [25] Behbood, N., F. Martin Ciurana, G. Colangelo, M. Napolitano, M. W. Mitchell, and R. J. Sewell (2013). Real-time vector field tracking with a cold-atom magnetometer. *Applied Physics Letters*, 102(17):173504.
- [26] Bell, T. A., J. A. P. Glidden, L. Humbert, M. W. J. Bromley, S. A. Haine, M. J. Davis, T. W. Neely, M. A. Baker, and H. Rubinsztein-Dunlop (2016). Bose–Einstein condensation in large time-averaged optical ring potentials. *New Journal of Physics*, 18(3):035003.
- [27] Bell, W. E. and A. L. Bloom (1957). Optical Detection of Magnetic Resonance in Alkali Metal Vapor. *Physical Review*, 107(6):1559–1565.
- [28] Bell, W. E. and A. L. Bloom (1961). Optically Driven Spin Precession. *Physical Review Letters*, 6(6):280–281.
- [29] Bergamini, S., B. Darquié, M. Jones, L. Jacubowicz, A. Browaeys, and P. Grangier (2004). Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator. *JOSA B*, 21(11):1889–1894.
- [30] Bergschneider, A., V. M. Klinkhamer, J. H. Becher, R. Klemt, G. Zürn, P. M. Preiss, and S. Jochim (2018). Spin-resolved single-atom imaging of ${}^6\text{Li}$ in free space. *Physical Review A*, 97(6):063613.
- [31] Bernien, H., S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin (2017). Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584.
- [32] Boll, M., T. A. Hilker, G. Salomon, A. Omran, J. Nespolo, L. Pollet, I. Bloch, and C. Gross (2016). Spin- and density-resolved microscopy of antiferromagnetic correlations in Fermi-Hubbard chains. *Science*, 353(6305):1257–1260.
- [33] Bons, P. C., R. de Haas, D. de Jong, A. Groot, and P. van der Straten (2016). Quantum Enhancement of the Index of Refraction in a Bose-Einstein Condensate. *Physical Review Letters*, 116(17).

- [34] Born, M. and E. Wolf (2001). *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, Cambridge, 7th edition.
- [35] Bose (1924). Plancks Gesetz und Lichtquantenhypothese. *Zeitschrift für Physik*, 26(1):178–181. English translation: “Planck’s Law and Light Quantum Hypothesis” by H. Dekant, *Observations on Quantum Computing and Physics*. Link: <https://wavewatching.net/lost-papers/>.
- [36] Boyer, V., R. M. Godun, G. Smirne, D. Cassettari, C. M. Chandrashekar, A. B. Deb, Z. J. Laczik, and C. J. Foot (2006). Dynamic manipulation of Bose-Einstein condensates with a spatial light modulator. *Physical Review A*, 73(3).
- [37] Bradley, C. C., C. A. Sackett, J. J. Tollett, and R. G. Hulet (1995). Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions. *Physical Review Letters*, 75(9):1687–1690.
- [38] Bradley, C. C., C. A. Sackett, and R. G. Hulet (1997). Bose-Einstein Condensation of Lithium: Observation of Limited Condensate Number. *Physical Review Letters*, 78(6): 985–989.
- [39] Braginsky, V. B. and F. Y. Khalili (1996). Quantum nondemolition measurements: The route from toys to tools. *Reviews of Modern Physics*, 68(1):I–II.
- [40] Braginsky, V. B., Y. I. Vorontsov, and K. S. Thorne (1980). Quantum Nondemolition Measurements. *Science*, 209(4456):547–557.
- [41] Brainis, E., C. Muldoon, L. Brandt, and A. Kuhn (2009). Coherent imaging of extended objects. *Optics Communications*, 282(4):465–472.
- [42] Brandt, L., C. Muldoon, T. Thiele, J. Dong, E. Brainis, and A. Kuhn (2011). Spatial light modulators for the manipulation of individual atoms. *Applied Physics B*, 102(3): 443–450.
- [43] Bransden, B. H. and C. J. Joachain (2003). *Physics of Atoms and Molecules*. Prentice Hall, Harlow, England, 2nd edition.
- [44] Brown, P. T., D. Mitra, E. Guardado-Sanchez, P. Schauß, S. S. Kondov, E. Khatami, T. Paiva, N. Trivedi, D. A. Huse, and W. S. Bakr (2017). Spin-imbalance in a 2D Fermi-Hubbard system. *Science*, 357(6358):1385–1388.
- [45] Brown, P. T., D. Mitra, E. Guardado-Sanchez, R. Nourafkan, A. Reymbaut, C.-D. Hébert, S. Bergeron, A.-M. S. Tremblay, J. Kokalj, D. A. Huse, P. Schauß, and W. S. Bakr (2019). Bad metallic transport in a cold atom Fermi-Hubbard system. *Science*, 363(6425):379–382.

- [46] Budker, D. and Derek F. Jackson Kimball (2013). *Optical Magnetometry*. Cambridge University Press, New York.
- [47] Budker, D. and M. Romalis (2007). Optical magnetometry. *Nature Physics*, 3(4):227–234.
- [48] Carpentier, A. V., Y. H. Fung, P. Sompet, A. J. Hilliard, T. G. Walker, and M. F. Andersen (2013). Preparation of a single atom in an optical microtrap. *Laser Physics Letters*, 10(12):125501.
- [49] Case, W. B., M. Tomandl, S. Deachapunya, and M. Arndt (2009). Realization of optical carpets in the Talbot and Talbot-Lau configurations. *Optics Express*, 17(23):20966.
- [50] Cheneau, M., P. Barmettler, D. Poletti, M. Endres, P. Schauß, T. Fukuhara, C. Gross, I. Bloch, C. Kollath, and S. Kuhr (2012). Light-cone-like spreading of correlations in a quantum many-body system. *Nature*, 481(7382):484–487.
- [51] Cheuk, L. W., M. A. Nichols, M. Okan, T. Gersdorf, V. V. Ramasesh, W. S. Bakr, T. Lompe, and M. W. Zwierlein (2015). Quantum-Gas Microscope for Fermionic Atoms. *Physical Review Letters*, 114(19):193001.
- [52] Cheuk, L. W., M. A. Nichols, K. R. Lawrence, M. Okan, H. Zhang, E. Khatami, N. Trivedi, T. Paiva, M. Rigol, and M. W. Zwierlein (2016). Observation of spatial charge and spin correlations in the 2D Fermi-Hubbard model. *Science*, 353(6305):1260–1264.
- [53] Chisholm, C. S., R. Thomas, A. B. Deb, and N. Kjærgaard (2018). A three-dimensional steerable optical tweezer system for ultracold atoms. *Review of Scientific Instruments*, 89(10):103105.
- [54] Chiu, C. S., G. Ji, A. Mazurenko, D. Greif, and M. Greiner (2018). Quantum State Engineering of a Hubbard System with Ultracold Fermions. *Physical Review Letters*, 120(24):243201.
- [55] Chiu, C. S., G. Ji, A. Bohrdt, M. Xu, M. Knap, E. Demler, F. Grusdt, M. Greiner, and D. Greif (2019). String patterns in the doped Hubbard model. *Science*, 365(6450):251–256.
- [56] Choi, J.-y., S. Hild, J. Zeiher, P. Schauß, A. Rubio-Abadal, T. Yefsah, V. Khemani, D. A. Huse, I. Bloch, and C. Gross (2016). Exploring the many-body localization transition in two dimensions. *Science*, 352(6293):1547–1552.
- [57] Chu, S., L. Hollberg, J. E. Bjorkholm, A. Cable, and A. Ashkin (1985). Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure. *Phys. Rev. Lett.*, 55(1):48–51.
- [58] Chu, S., J. E. Bjorkholm, A. Ashkin, and A. Cable (1986). Experimental Observation of

- Optically Trapped Atoms. *Physical Review Letters*, 57(3):314–317.
- [59] Cohen, Y., K. Jadeja, S. Sula, M. Venturelli, C. Deans, L. Marmugi, and F. Renzoni (2019). A cold atom radio-frequency magnetometer. *Applied Physics Letters*, 114(7):073505.
 - [60] Colangelo, G., F. M. Ciurana, L. C. Bianchet, R. J. Sewell, and M. W. Mitchell (2017). Simultaneous tracking of spin angle and amplitude beyond classical limits. *Nature*, 543(7646):525–528.
 - [61] Comparat, D., A. Fioretti, G. Stern, E. Dimova, B. L. Tolra, and P. Pillet (2006). Optimized production of large Bose-Einstein condensates. *Phys. Rev. A*, 73(4):043410.
 - [62] Cooper, A., J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres (2018). Alkaline-Earth Atoms in Optical Tweezers. *Physical Review X*, 8(4).
 - [63] Dalibard, J. and C. Cohen-Tannoudji (1989). Laser cooling below the Doppler limit by polarization gradients: Simple theoretical models. *Journal of the Optical Society of America B*, 6(11):2023.
 - [64] Dang, H. B., A. C. Maloof, and M. V. Romalis (2010). Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer. *Applied Physics Letters*, 97(15):151110.
 - [65] Davis, C. C. (1996). *Lasers and Electro-Optics: Fundamentals and Engineering*. Cambridge University Press, Cambridge.
 - [66] Davis, K. B., M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle (1995). Bose-Einstein Condensation in a Gas of Sodium Atoms. *Phys. Rev. Lett.*, 75(22):3969–3973.
 - [67] de Touzalin, A., C. Marcus, F. Heijman, I. Cirac, R. Murray, and T. Calarco (2016). Quantum Manifesto.
 - [68] Deb, A. B., T. McKellar, and N. Kjærgaard (2014). Optical runaway evaporation for the parallel production of multiple Bose-Einstein condensates. *Physical Review A*, 90(5):051401.
 - [69] Debye, P. and F. W. Sears (1932). On the Scattering of Light by Supersonic Waves. *Proceedings of the National Academy of Sciences*, 18(6):409–414.
 - [70] Degen, C. L., F. Reinhard, and P. Cappellaro (2017). Quantum sensing. *Reviews of Modern Physics*, 89(3):035002.
 - [71] Dehmelt, H. G. (1957). Modulation of a Light Beam by Precessing Absorbing Atoms. *Physical Review*, 105(6):1924–1925.
 - [72] Denschlag, J. H., J. E. Simsarian, H. Häffner, C. McKenzie, A. Browaeys, D. Cho,

- K. Helmerson, S. L. Rolston, and W. D. Phillips (2002). A Bose-Einstein condensate in an optical lattice. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 35(14): 3095–3110.
- [73] Deutsch, I. H. and P. S. Jessen (2010). Quantum control and measurement of atomic spins in polarization spectroscopy. *Optics Communications*, 283(5):681–694.
- [74] Duarte, P. M., R. A. Hart, J. M. Hitchcock, T. A. Corcovilos, T.-L. Yang, A. Reed, and R. G. Hulet (2011). All-optical production of a lithium quantum gas using narrow-line laser cooling. *Physical Review A*, 84(6).
- [75] Dupont-Roc, J., S. Haroche, and C. Cohen-Tannoudji (1969). Detection of very weak magnetic fields (10^{-9} Gauss) by ^{87}Rb zero-field level crossing resonances. *Physics Letters A*, 28(9):638–639.
- [76] Dutta, S. and E. J. Mueller (2015). Kinetics of Bose-Einstein condensation in a dimple potential. *Phys. Rev. A*, 91(1):013601.
- [77] Eckel, S., F. Jendrzejewski, A. Kumar, C. J. Lobb, and G. K. Campbell (2014). Interferometric Measurement of the Current-Phase Relationship of a Superfluid Weak Link. *Physical Review X*, 4(3).
- [78] Edge, G. J. A., R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky, and J. H. Thywissen (2015). Imaging and addressing of individual fermionic atoms in an optical lattice. *Physical Review A*, 92(6):063406.
- [79] Einstein, A. (1917). Zur Quantentheorie der Strahlung. *Physikalische Zeitschrift*, 18:121–128. English translation: “On the Quantum Theory of Radiation,” by D. ter Haar, *The Old Quantum Theory*, Pergamon Press, New York (1967), p. 167. Link: [http://hermes.ffn.ub.es/luisnavarro/nuevo_maletin/Einstein%20\(1917\)_Quantum%20theory%20of%20radiation.pdf](http://hermes.ffn.ub.es/luisnavarro/nuevo_maletin/Einstein%20(1917)_Quantum%20theory%20of%20radiation.pdf).
- [80] Einstein, A. (1924). Quantentheorie des einatomigen idealen Gases. *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, pages 261–267. English translation: “Quantum Theory of the Ideal Monoatomic Gas.” by L. Amendola. Link: <https://www.thphys.uni-heidelberg.de/~amendola/otherstuff/einstein-paper-v2.pdf>.
- [81] Einstein, A. (1925). Quantentheorie des einatomigen idealen Gases. Zweite Abhandlung. *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, pages 3–14. English translation: “Quantum Theory of the Ideal Monoatomic Gas. Second Treatise.” by H. Dekant, *Observations on Quantum Computing and Physics*. Link: <https://wavewatching.net/lost-papers/>.
- [82] Ejlersen, P. (2016). *Investigation of Complex Control Landscapes of an Evaporative Cooling Process*. Master’s thesis, Aarhus University, Aarhus.

- [83] Elíasson, O. (2017). *Dual Port Faraday Imaging for Magnetometry in Tweezer Arrays*. Progress report towards PhD, Aarhus University, Aarhus.
- [84] Elíasson, O., R. Heck, J. S. Laustsen, M. Napolitano, R. Müller, M. G. Bason, J. J. Arlt, and J. F. Sherson (2019). Spatially-selective *in situ* magnetometry of ultracold atomic clouds. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 52(7):075003.
- [85] Endres, M. (2013). *Probing Correlated Quantum Many-Body Systems at the Single-Particle Level*. PhD thesis, Ludwig-Maximilians-Universität München, München.
- [86] Endres, M., M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schauß, C. Gross, L. Mazza, M. C. Bañuls, L. Pollet, I. Bloch, and S. Kuhr (2011). Observation of Correlated Particle-Hole Pairs and String Order in Low-Dimensional Mott Insulators. *Science*, 334(6053):200–203.
- [87] Endres, M., H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin (2016). Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science*, 354(6315):1024–1027.
- [88] Eto, Y., H. Ikeda, H. Suzuki, S. Hasegawa, Y. Tomiyama, S. Sekine, M. Sadgrove, and T. Hirano (2013). Spin-echo-based magnetometry with spinor Bose-Einstein condensates. *Physical Review A*, 88(3).
- [89] Faraday, M. (1933). *Faraday's Diary: Being the Various Philosophical Notes of Experimental Investigation*, volume 4. G. Bell and Sons, Ltd, London. Text available at https://archive.org/stream/faradaysdiarybeioofara_2/faradaysdiarybeioofara_2_djvu.txt and the relevant entries are in fact available in a preview copy of the diary <http://faradaysdiary.com/ws3/faraday.pdf>.
- [90] Fatemi, F. K. and M. Bashkansky (2010). Spatially resolved magnetometry using cold atoms in dark optical tweezers. *Optics Express*, 18(3):2190–2196.
- [91] Feldon, D. F., K. Litson, S. Jeong, J. M. Blaney, J. Kang, C. Miller, K. Griffin, and J. Roksa (2019). Postdocs' lab engagement predicts trajectories of PhD students' skill development. *Proceedings of the National Academy of Sciences*, page 201912488.
- [92] Feng, P., D. Hoffmann, and T. Walker (1993). Comparison of trap-loss collision spectra for ^{85}Rb and ^{87}Rb . *Physical Review A*, 47(5):R3495–R3498.
- [93] Fernholz, T., H. Krauter, K. Jensen, J. F. Sherson, A. S. Sørensen, and E. S. Polzik (2008). Spin Squeezing of Atomic Ensembles via Nuclear-Electronic Spin Entanglement. *Physical Review Letters*, 101(7).
- [94] Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488.

- [95] Foot, C. J. (2007). *Atomic Physics*. Oxford Master Series in Physics ; 7 : Atomic, Optical, and Laser Physics ; Oxford Master Series in Physics. Oxford University Press, Oxford.
- [96] Freimund, D. L., K. Aflatooni, and H. Batelaan (2001). Observation of the Kapitza–Dirac effect. *Nature*, 413(6852):142–143.
- [97] Frese, D., B. Ueberholz, S. Kuhr, W. Alt, D. Schrader, V. Gomer, and D. Meschede (2000). Single Atoms in an Optical Dipole Trap: Towards a Deterministic Source of Cold Atoms. *Physical Review Letters*, 85(18):3777–3780.
- [98] Friedman, N., L. Khaykovich, R. Ozeri, and N. Davidson (2000). Compression of cold atoms to very high densities in a rotating-beam blue-detuned optical trap. *Physical Review A*, 61(3).
- [99] Fuhrmanek, A., R. Bourgain, Y. R. P. Sortais, and A. Browaeys (2012). Light-assisted collisions between a few cold atoms in a microscopic dipole trap. *Physical Review A*, 85(6).
- [100] Fukuhara, T., A. Kantian, M. Endres, M. Cheneau, P. Schauß, S. Hild, D. Bellem, U. Schollwöck, T. Giamarchi, C. Gross, I. Bloch, and S. Kuhr (2013). Quantum dynamics of a mobile spin impurity. *Nature Physics*, 9(4):235–241.
- [101] Fukuhara, T., P. Schauß, M. Endres, S. Hild, M. Cheneau, I. Bloch, and C. Gross (2013). Microscopic observation of magnon bound states and their dynamics. *Nature*, 502(7469):76–79.
- [102] Fukuhara, T., S. Hild, J. Zeiher, P. Schauß, I. Bloch, M. Endres, and C. Gross (2015). Spatially Resolved Detection of a Spin-Entanglement Wave in a Bose-Hubbard Chain. *Physical Review Letters*, 115(3).
- [103] Gajdacz, M. (2015). *Non-Destructive Imaging and Feedback with Ultracold Gases*. PhD thesis, Aarhus University, Aarhus.
- [104] Gajdacz, M., P. L. Pedersen, T. Mørch, A. J. Hilliard, J. Arlt, and J. F. Sherson (2013). Non-destructive Faraday imaging of dynamically controlled ultracold atoms. *Rev. Sci. Instrum.*, 84(8).
- [105] Gajdacz, M., A. J. Hilliard, M. A. Kristensen, P. L. Pedersen, C. Klempt, J. J. Arlt, and J. F. Sherson (2016). Preparation of Ultracold Atom Clouds at the Shot Noise Level. *Physical Review Letters*, 117(7).
- [106] Garrett, M. C., A. Ratnapala, E. D. van Ooijen, C. J. Vale, K. Weegink, S. K. Schnelle, O. Vainio, N. R. Heckenberg, H. Rubinsztein-Dunlop, and M. J. Davis (2011). Growth dynamics of a Bose-Einstein condensate in a dimple trap without cooling. *Physical Review A*, 83(1):013630.

- [107] Garwood, M. and L. DelaBarre (2001). The Return of the Frequency Sweep: Designing Adiabatic Pulses for Contemporary NMR. *Journal of Magnetic Resonance*, 153(2):155–177.
- [108] Gauthier, G., I. Lenton, N. M. Parry, M. Baker, M. J. Davis, H. Rubinsztein-Dunlop, and T. W. Neely (2016). Direct imaging of a digital-micromirror device for configurable microscopic optical potentials. *Optica*, 3(10):1136–1143.
- [109] Gawlik, W. and J. M. Higbie (2013). Magnetometry with cold atoms. In Budker, D. and D. F. J. Kimball, editors, *Optical Magnetometry*, pages 167–189. Cambridge University Press, New York.
- [110] Gawlik, W. and S. Pustelny (2009). Nonlinear Faraday effect and its applications. In Drampyan, R., editor, *New Trends in Quantum Coherence and Nonlinear Optics*, volume 263 of *Horizons in World Physics*, pages 47–84. Nova Science Publishers, New York.
- [111] Georgescu, I. M., S. Ashhab, and F. Nori (2014). Quantum simulation. *Reviews of Modern Physics*, 86(1):153–185.
- [112] Gerbier, F., A. Widera, S. Fölling, O. Mandel, T. Gericke, and I. Bloch (2005). Phase Coherence of an Atomic Mott Insulator. *Physical Review Letters*, 95(5).
- [113] Geremia, J. M., J. K. Stockton, and H. Mabuchi (2006). Tensor polarizability and dispersive quantum measurement of multilevel atoms. *Phys. Rev. A*, 73(4):042112.
- [114] Gericke, T., P. Würtz, D. Reitz, T. Langen, and H. Ott (2008). High-resolution scanning electron microscopy of an ultracold quantum gas. *Nature Physics*, 4(12):949–953.
- [115] Gillen, J. I., W. S. Bakr, A. Peng, P. Unterwadtzer, S. Fölling, and M. Greiner (2009). Two-dimensional quantum gas in a hybrid surface trap. *Physical Review A*, 80(2):021602.
- [116] Gould, P. L., G. A. Ruff, and D. E. Pritchard (1986). Diffraction of atoms by light: The near-resonant Kapitza-Dirac effect. *Physical Review Letters*, 56(8):827–830.
- [117] Greiner, M. (2003). *Ultracold Quantum Gases in Three-Dimensional Optical Lattice*. PhD thesis, Ludwig-Maximilians-Universität München, München.
- [118] Greiner, M., O. Mandel, T. Esslinger, T. W. Hansch, and I. Bloch (2002). Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415(6867):39–44.
- [119] Grier, D. G. (2003). A revolution in optical manipulation. *Nature*, 424(6950):810–816.
- [120] Griffiths, D. J. (2005). *Introduction to Quantum Mechanics*. Pearson Education, New Jersey, 2nd edition.
- [121] Grimm, R., M. Weidemüller, and Y. B. Ovchinnikov (2000). Optical Dipole Traps for

- Neutral Atoms. In *Advances In Atomic, Molecular, and Optical Physics*, volume 42, pages 95–170. Elsevier.
- [122] Gross, C. and I. Bloch (2017). Quantum simulations with ultracold atoms in optical lattices. *Science*, 357(6355):995–1001.
- [123] Grünzweig, T., A. Hilliard, M. McGovern, and M. F. Andersen (2010). Near-deterministic preparation of a single atom in an optical microtrap. *Nature Physics*, 6(12):951–954.
- [124] Gu, M. (1999). *Advanced Optical Imaging Theory*. Springer Series in Optical Sciences. Springer, Berlin, Heidelberg.
- [125] Guardado-Sanchez, E., P. T. Brown, D. Mitra, T. Devakul, D. A. Huse, P. Schauß, and W. S. Bakr (2018). Probing the Quench Dynamics of Antiferromagnetic Correlations in a 2D Quantum Ising Spin System. *Physical Review X*, 8(2).
- [126] Haas, M., U. D. Jentschura, and C. H. Keitel (2006). Comparison of classical and second quantized description of the dynamic Stark shift. *American Journal of Physics*, 74(1):77–81.
- [127] Haller, E., J. Hudson, A. Kelly, D. A. Cotta, B. Peaudecerf, G. D. Bruce, and S. Kuhr (2015). Single-atom imaging of fermions in a quantum-gas microscope. *Nature Physics*, 11(9):738–742.
- [128] Hammerer, K., A. S. Sørensen, and E. S. Polzik (2010). Quantum interface between light and atomic ensembles. *Reviews of Modern Physics*, 82(2):1041–1093.
- [129] Happer, W. (1972). Optical Pumping. *Reviews of Modern Physics*, 44(2):169–249.
- [130] Haugaard, N. A. (2017). *Building a High Resolution Imaging System for Detecting, Addressing and Manipulating Single, Optically Trapped Atoms*. Master’s thesis, Aarhus University, Aarhus.
- [131] Heck, R. (2016). *A Toolbox for Optimized and Stable Experiments with Ultracold Quantum Gases*. PhD thesis, Aarhus University, Aarhus.
- [132] Heck, R., O. Vuculescu, J. J. Sørensen, J. Zoller, M. G. Andreassen, M. G. Bason, P. Ejlertsen, O. Eliasson, P. Haikka, J. S. Laustsen, L. L. Nielsen, A. Mao, R. Müller, M. Napolitano, M. K. Pedersen, A. R. Thorsen, C. Bergenholtz, T. Calarco, S. Montangero, and J. F. Sherson (2018). Remote optimization of an ultracold atoms experiment by experts and citizen scientists. *Proceedings of the National Academy of Sciences*, 115(48):E11231–E11237.
- [133] Henderson, K., C. Ryu, C. MacCormick, and M. G. Boshier (2009). Experimental demonstration of painting arbitrary and dynamic potentials for Bose–Einstein conden-

- sates. *New Journal of Physics*, 11(4):043030.
- [134] Higbie, J. M., L. E. Sadler, S. Inouye, A. P. Chikkatur, S. R. Leslie, K. L. Moore, V. Savalli, and D. M. Stamper-Kurn (2005). Direct nondestructive imaging of magnetization in a spin-1 bose-einstein gas. *Physical Review Letters*, 95(5).
 - [135] Hild, S., T. Fukuhara, P. Schauß, J. Zeiher, M. Knap, E. Demler, I. Bloch, and C. Gross (2014). Far-from-Equilibrium Spin Transport in Heisenberg Quantum Magnets. *Physical Review Letters*, 113(14):147205.
 - [136] Hilker, T. A., G. Salomon, F. Grusdt, A. Omran, M. Boll, E. Demler, I. Bloch, and C. Gross (2017). Revealing hidden antiferromagnetic correlations in doped Hubbard chains via string correlators. *Science*, 357(6350):484–487.
 - [137] Hoffmann, D. (1996). *Light-Induced Collisions of Ultracold Rubidium Atoms*. PhD thesis, University of Wisconsin–Madison, Madison.
 - [138] Hubbard, J. (1963). Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*.
 - [139] Huber, F. G. (2014). *Site-Resolved Imaging with the Fermi Gas Microscope*. PhD thesis, Harvard, Cambridge, Massachusetts.
 - [140] Hueck, K., A. Mazurenko, N. Luick, T. Lompe, and H. Moritz (2017). Note: Suppression of kHz-frequency switching noise in digital micro-mirror devices. *Review of Scientific Instruments*, 88(1):016103.
 - [141] Hänsch, T. W. and A. L. Schawlow (1975). Cooling of gases by laser radiation. *Optics Communications*, 13(1):68–69.
 - [142] IDS Imaging Development Systems GmbH (2016). The key to contrast and detail: Adjusting the black level. Link: https://en.ids-imaging.com/techtipps-detail/en_techtip-black-level.html. Accessed: 2019-09-20 14:39:54.
 - [143] Immerkjær, J. (1996). Note: Fast Noise Variance Estimation. *Computer Vision and Image Understanding*, 64(2):300–302.
 - [144] InfluxData (2019). InfluxDB 1.7 documentation | InfluxData Documentation. Link: <https://docs.influxdata.com/>. Accessed: 2019-08-27 13:40:31.
 - [145] Isayama, T., Y. Takahashi, N. Tanaka, K. Toyoda, K. Ishikawa, and T. Yabuzaki (1999). Observation of larmor spin precession of laser-cooled Rb atoms via paramagnetic faraday rotation. *Physical Review A*, 59(6):4836–4839.
 - [146] Islam, R., R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner (2015). Measuring entanglement entropy in a quantum many-body system. *Nature*, 528(7580):77–83.

- [147] Jacob, D., E. Mimoun, L. D. Sarlo, M. Weitz, J. Dalibard, and F. Gerbier (2011). Production of sodium Bose–Einstein condensates in an optical dimple trap. *New Journal of Physics*, 13(6):065022.
- [148] Jechow, A., B. G. Norton, S. Händel, V. Blüms, E. W. Streed, and D. Kielpinski (2013). Controllable Optical Phase Shift Over One Radian from a Single Isolated Atom. *Physical Review Letters*, 110(11).
- [149] Jensen, J. H. M., J. J. Sørensen, K. Mølmer, and J. F. Sherson (2019). Time-optimal control of collisional $\sqrt{\text{swap}}$ gates in ultracold atomic systems. *arXiv:1907.08504 [quant-ph]*.
- [150] Jo, G.-B., J. Guzman, C. K. Thomas, P. Hosur, A. Vishwanath, and D. M. Stamper-Kurn (2012). Ultracold Atoms in a Tunable Optical Kagome Lattice. *Physical Review Letters*, 108(4):045305.
- [151] Jørgensen, N. B., M. G. Bason, and J. F. Sherson (2014). One- and two-qubit quantum gates using superimposed optical-lattice potentials. *Physical Review A*, 89(3):032306.
- [152] Kaminski, F. (2012). *Coherent Atom-Light Interaction in an Ultracold Atomic Gas Experimental Study of Faraday Rotation Imaging and Matter-Wave Superradiance*. PhD thesis, University of Copenhagen, Copenhagen.
- [153] Kaminski, F., N. S. Kampel, M. P. H. Steenstrup, A. Griesmaier, E. S. Polzik, and J. H. Müller (2012). In-situ dual-port polarization contrast imaging of faraday rotation in a high optical depth ultracold ^{87}Rb atomic ensemble. *The European Physical Journal D*, 66(9).
- [154] Kastler, A. (1973). The Hanle effect and its use for the measurements of very small magnetic fields. *Nuclear Instruments and Methods*, 110:259–265.
- [155] Kaufman, A. M., B. J. Lester, and C. A. Regal (2012). Cooling a Single Atom in an Optical Tweezer to Its Quantum Ground State. *Physical Review X*, 2(4):041014.
- [156] Kaufman, A. M., B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazard, A. M. Rey, and C. A. Regal (2014). Two-particle quantum interference in tunnel-coupled optical tweezers. *Science*, 345(6194):306–309.
- [157] Kaufman, A. M., B. J. Lester, M. Foss-Feig, M. L. Wall, A. M. Rey, and C. A. Regal (2015). Entangling two transportable neutral atoms via local spin exchange. *Nature*, 527(7577):208–211.
- [158] Kaufman, A. M., M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner (2016). Quantum thermalization through entanglement in an isolated many-body system. *Science*, 353(6301):794–800.
- [159] Ketterle, W., D. S. Durfee, and D. M. Stamper-Kurn (1999). Making, probing and un-

- derstanding bose-einstein condensates. *arXiv:9904034 [cond-mat]*.
- [160] Kim, H., W. Lee, H.-g. Lee, H. Jo, Y. Song, and J. Ahn (2016). In situ single-atom array synthesis using dynamic holographic optical tweezers. *Nature Communications*, 7(1).
 - [161] Kim, H., Y. Park, K. Kim, H.-S. Sim, and J. Ahn (2018). Detailed Balance of Thermalization Dynamics in Rydberg-Atom Quantum Simulators. *Physical Review Letters*, 120(18).
 - [162] Kittel, C. and H. Kroemer (1980). *Thermal Physics*. W. H. Freeman and Company, New York, 2nd edition.
 - [163] Knappe, S. and J. Kitching (2013). Microfabricated atomic magnetometers. In *Optical Magnetometry*, pages 125–141. Cambridge University Press, New York.
 - [164] Kobayakov, A., M. Sauer, and D. Chowdhury (2010). Stimulated Brillouin scattering in optical fibers. *Advances in Optics and Photonics*, 2(1):1–59.
 - [165] Koepsell, J., J. Vijayan, P. Sompet, F. Grusdt, T. A. Hilker, E. Demler, G. Salomon, I. Bloch, and C. Gross (2019). Imaging magnetic polarons in the doped Fermi–Hubbard model. *Nature*, 572(7769):358–362.
 - [166] Kominis, I. K., T. W. Kornack, J. C. Allred, and M. V. Romalis (2003). A subfemtotesla multichannel atomic magnetometer. *Nature*, 422(6932):596–599.
 - [167] Koschorreck, M. (2010). *Generation of Spin Squeezing in an Ensemble of Cold Rubidium* 87. PhD thesis, Univeristat Politècnica de Catalunya, Barcelona.
 - [168] Koschorreck, M., M. Napolitano, B. Dubost, and M. W. Mitchell (2010). Quantum Nondemolition Measurement of Large-Spin Ensembles by Dynamical Decoupling. *Physical Review Letters*, 105(9):093602.
 - [169] Koschorreck, M., M. Napolitano, B. Dubost, and M. W. Mitchell (2011). High resolution magnetic vector-field imaging with cold atomic ensembles. *Applied Physics Letters*, 98(7):074101.
 - [170] Krarup, O. (2016). *Spatial Light Modulation with Digital Micromirror Devices*. BSc thesis, Aarhus University, Aarhus.
 - [171] Krarup, O. (2018). *Imaging Single Atoms in an Optical Lattice*. Master’s thesis, University of Aarhus, Aarhus.
 - [172] Krinner, L., M. Stewart, A. Pazmiño, and D. Schneble (2018). In Situ Magnetometry for Experiments with Atomic Quantum Gases. *Review of Scientific Instruments*, 89(1):013108.
 - [173] Kristensen, M. A., M. B. Christensen, M. Gajdacz, M. Iglicki, K. Pawłowski, C. Klempt,

- J. F. Sherson, K. Rzażewski, A. J. Hilliard, and J. J. Arlt (2019). Observation of Atom Number Fluctuations in a Bose-Einstein Condensate. *Physical Review Letters*, 122(16).
- [174] Kuppens, S. J. M., K. L. Corwin, K. W. Miller, T. E. Chupp, and C. E. Wieman (2000). Loading an optical dipole trap. *Phys. Rev. A*, 62(013406):13.
- [175] Kuzmich, A., L. Mandel, and N. P. Bigelow (2000). Generation of Spin Squeezing via Continuous Quantum Nondemolition Measurement. *Physical Review Letters*, 85(8): 1594–1597.
- [176] Labuhn, H., D. Barredo, S. Ravets, S. de Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys (2016). Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models. *Nature*, 534(7609):667–670.
- [177] Langbein, F. C., S. Schirmer, and E. Jonckheere (2015). Time optimal information transfer in spintronics networks. pages 6454–6459. IEEE (2015).
- [178] Laustsen, J. S. (2018). *Single Site Resolved Imaging: Towards a Quantum Gas Microscope*. Progress report towards PhD, Aarhus University, Aarhus.
- [179] Lebedev, N. N. (1972). *Special Functions & Their Applications*. Dover, New York, revised english edition.
- [180] Lester, B. J., N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal (2015). Rapid Production of Uniformly Filled Arrays of Neutral Atoms. *Physical Review Letters*, 115 (7).
- [181] Lett, P. D., R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf (1988). Observation of Atoms Laser Cooled below the Doppler Limit. *Physical Review Letters*, 61(2):169–172.
- [182] Lett, P. D., W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook (1989). Optical molasses. *JOSA B*, 6(11):2084–2107.
- [183] Lewenstein, M., A. Sanpera, and V. Ahufinger (2012). *Ultracold Atoms in Optical Lattices: Simulating Quantum Many-Body Systems*. Oxford University Press.
- [184] Li, J.-R., J. Lee, W. Huang, S. Burchesky, B. Shteynas, F. C. Top, A. O. Jamison, and W. Ketterle (2017). A stripe phase with supersolid properties in spin-orbit-coupled Bose-Einstein condensates. *Nature*, 543(7643):91–94.
- [185] Léonard, J., A. Morales, P. Zupancic, T. Esslinger, and T. Donner (2017). Supersolid formation in a quantum gas breaking a continuous translational symmetry. *Nature*, 543 (7643):87–90.
- [186] Maiman, T. H. (1960). Stimulated Optical Radiation in Ruby. *Nature*, 187(4736):493–494.

- [187] Martinez-Dorantes, M., W. Alt, J. Gallego, S. Ghosh, L. Ratschbacher, and D. Meschede (2018). State-dependent fluorescence of neutral atoms in optical potentials. *Physical Review A*, 97(2).
- [188] Mathew, R., A. Kumar, S. Eckel, F. Jendrzejewski, G. K. Campbell, M. Edwards, and E. Tiesinga (2015). Self-heterodyne detection of the in situ phase of an atomic superconducting quantum interference device. *Physical Review A*, 92(3):033602.
- [189] Mazurenko, A., C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner (2017). A cold-atom Fermi–Hubbard antiferromagnet. *Nature*, 545(7655):462–466.
- [190] Mazurenko, A., S. Blatt, F. Huber, M. F. Parsons, C. S. Chiu, G. Ji, D. Greif, and M. Greiner (2019). Implementation of a stable, high-power optical lattice for quantum gas microscopy. *Review of Scientific Instruments*, 90(3):033101.
- [191] Metcalf, H. J. and P. van der Starten (1999). *Laser Cooling and Trapping*. Springer New York, New York.
- [192] Mikkelsen, S. (2016). *Homodyne Detection of Ultracold Atomic Clouds*. Master’s thesis, Aarhus University, Aarhus.
- [193] Miranda, M., R. Inoue, Y. Okuyama, A. Nakamoto, and M. Kozuma (2015). Site-resolved imaging of ytterbium atoms in a two-dimensional optical lattice. *Physical Review A*, 91(6).
- [194] Mitchell, M. W. and S. P. Alvarez (2019). Quantum limits to the energy resolution of magnetic field sensors. *arXiv:1905.00618 [quant-ph]*.
- [195] Monroe, C. (2011). Demolishing quantum nondemolition. *Physics Today*, 64(1):8.
- [196] Muessel, W., H. Strobel, D. Linnemann, D. Hume, B., and M. Oberthaler, K. (2014). Scalable Spin Squeezing for Quantum-Enhanced Magnetometry with Bose-Einstein Condensates. *Phys. Rev. Lett.*, 113(10):103004.
- [197] Muldoon, C. (2012). *Control and Manipulation of Cold Atoms Trapped in Optical Tweezers*. PhD thesis, Oxford University, Oxford.
- [198] Muldoon, C., L. Brandt, J. Dong, D. Stuart, E. Brainis, M. Himsworth, and A. Kuhn (2012). Control and manipulation of cold atoms in optical tweezers. *New Journal of Physics*, 14(7):073051.
- [199] Müller, R. C. G. F. (2015). *A New Experiment to Probe and Manipulate Quantum Systems*. PhD thesis, Aarhus University, Aarhus.
- [200] Nelson, K. D., X. Li, and D. S. Weiss (2007). Imaging single atoms in a three-dimensional array. *Nature Physics*, 3(8):556–560.

- [201] Newton, R. G. (1976). Optical theorem and beyond. *American Journal of Physics*, 44(7):639–642.
- [202] Nichols, M. A., L. W. Cheuk, M. Okan, T. R. Hartke, E. Mendez, T. Senthil, E. Khatami, H. Zhang, and M. W. Zwierlein (2019). Spin transport in a Mott insulator of ultracold fermions. *Science*, 363(6425):383–387.
- [203] Nogrette, F., H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys (2014). Single-Atom Trapping in Holographic 2D Arrays of Microtraps with Arbitrary Geometries. *Physical Review X*, 4(2).
- [204] Norcia, M. A., A. W. Young, and A. M. Kaufman (2018). Microscopic Control and Detection of Ultracold Strontium in Optical-Tweezer Arrays. *Physical Review X*, 8(4).
- [205] Ockeloen, C. F., R. Schmied, M. F. Riedel, and P. Treutlein (2013). Quantum Metrology with a Scanning Probe Atom Interferometer. *Physical Review Letters*, 111(14):143001.
- [206] Ohl de Mello, D., D. Schäffner, J. Werkmann, T. Preuschoff, L. Kohfahl, M. Schlosser, and G. Birkl (2019). Defect-Free Assembly of 2D Clusters of More Than 100 Single-Atom Quantum Systems. *Physical Review Letters*, 122(20):203601.
- [207] Olson, A. J., R. J. Niffenegger, and Y. P. Chen (2013). Optimizing the efficiency of evaporative cooling in optical dipole traps. *Physical Review A*, 87(5):053613.
- [208] Omran, A., M. Boll, T. A. Hilker, K. Kleinlein, G. Salomon, I. Bloch, and C. Gross (2015). Microscopic Observation of Pauli Blocking in Degenerate Fermionic Lattice Gases. *Physical Review Letters*, 115(26):263001.
- [209] Ortiz-Gutiérrez, L., R. C. Teixeira, A. Eloy, D. F. da Silva, R. Kaiser, R. Bachelard, and M. Fouché (2019). Mollow triplet in cold atoms. *New Journal of Physics*, 21(9):093019.
- [210] Ott, H. (2016). Single atom detection in ultracold quantum gases: A review of current progress. *Reports on Progress in Physics*, 79(5):054401.
- [211] Ovchinnikov, Y. B., J. H. Müller, M. R. Doery, E. J. D. Vredenburg, K. Helmerson, S. L. Rolston, and W. D. Phillips (1999). Diffraction of a Released Bose-Einstein Condensate by a Pulsed Standing Light Wave. *Physical Review Letters*, 83(2):284–287.
- [212] O’Hara, K. M., M. E. Gehm, S. R. Granade, and J. E. Thomas (2001). Scaling laws for evaporative cooling in time-dependent optical traps. *Physical Review A*, 64(5).
- [213] Palacios, S., S. Coop, P. Gomez, T. Vanderbruggen, Y. N. M. de Escobar, Martijn Jasperse, and M. W. Mitchell (2018). Multi-second magnetic coherence in a single domain spinor Bose–Einstein condensate. *New Journal of Physics*, 20(5):053008.
- [214] Parsons, M. F., F. Huber, A. Mazurenko, C. S. Chiu, W. Setiawan, K. Wooley-Brown, S. Blatt, and M. Greiner (2015). Site-resolved imaging of fermionic ${}^6\text{Li}$ in an optical

- lattice. *Physical Review Letters*, 114(21):213002.
- [215] Parsons, M. F., A. Mazurenko, C. S. Chiu, G. Ji, D. Greif, and M. Greiner (2016). Site-resolved measurement of the spin-correlation function in the Fermi-Hubbard model. *Science*, 353(6305):1253–1256.
 - [216] Patil, Y. S., S. Chakram, and M. Vengalattore (2015). Measurement-Induced Localization of an Ultracold Lattice Gas. *Physical Review Letters*, 115(14):140402.
 - [217] Pauli, W. (1940). The Connection Between Spin and Statistics. *Physical Review*, 58(8):716–722.
 - [218] Pednault, E., J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff (2019). Leveraging secondary storage to simulate deep 54-qubit sycamore circuits. *arXiv:1910.09534 [quant-ph]*.
 - [219] Pethick, C. J. and H. Smith (2008). *Bose-Einstein Condensation in Dilute Gases*. Cambridge University Press, Cambridge.
 - [220] Pezzè, L., A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein (2018). Quantum metrology with nonclassical states of atomic ensembles. *Reviews of Modern Physics*, 90(3):035005.
 - [221] Pinkse, P. W. H., A. Mosk, M. Weidemüller, M. W. Reynolds, T. W. Hijmans, and J. T. M. Walraven (1997). Adiabatically Changing the Phase-Space Density of a Trapped Bose Gas. *Phys. Rev. Lett.*, 78(6):990–993.
 - [222] Politzer, H. D. (1997). Bose-stimulated scattering off a cold atom trap. *Physical Review A*, 55(2):1140–1146.
 - [223] Poulsen, J. E. (2018). *Trapping and Manipulation of Ultra Cold Atoms in Tailorpmade Light Fields*. Master’s thesis, Aarhus University, Aarhus.
 - [224] Preiss, P. M., R. Ma, M. E. Tai, A. Lukin, M. Rispoli, P. Zupancic, Y. Lahini, R. Islam, and M. Greiner (2015). Strongly correlated quantum walks in optical lattices. *Science*, 347(6227):1229–1233.
 - [225] Preiss, P. M., R. Ma, M. E. Tai, J. Simon, and M. Greiner (2015). Quantum gas microscopy with spin, atom-number, and multilayer readout. *Physical Review A*, 91(4):041602.
 - [226] Prodan, J., A. Migdall, W. D. Phillips, I. So, H. Metcalf, and J. Dalibard (1985). Stopping Atoms with Laser Light. *Physical Review Letters*, 54(10):992–995.
 - [227] Pu, Y.-F., N. Jiang, W. Chang, H.-X. Yang, C. Li, and L.-M. Duan (2017). Experimental realization of a multiplexed quantum memory with 225 individually accessible memory cells. *Nature Communications*, 8:15359.

- [228] Raab, E. L., M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard (1987). Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59(23):2631–2634.
- [229] Ramanathan, A., S. R. Muniz, K. C. Wright, R. P. Anderson, W. D. Phillips, K. Helmer-son, and G. K. Campbell (2012). Partial-transfer absorption imaging: A versatile technique for optimal imaging of ultracold gases. *Review of Scientific Instruments*, 83(8):083119.
- [230] Reinaudi, G., T. Lahaye, Z. Wang, and D. Guéry-Odelin (2007). Strong saturation absorption imaging of dense clouds of ultracold atoms. *Optics Letters*, 32(21):3143–3145.
- [231] Roberts, K. O., T. McKellar, J. Fekete, A. Rakonjac, A. B. Deb, and N. Kjærgaard (2014). Steerable optical tweezers for ultracold atom studies. *Optics Letters*, 39(7):2012–2015.
- [232] Rogers, B., M. Paternostro, J. F. Sherson, and G. De Chiara (2014). Characterization of Bose-Hubbard models with quantum nondemolition measurements. *Physical Review A*, 90(4):043618.
- [233] Sakurai, J. J. and J. Napolitano (2011). *Modern Quantum Mechanics*. Pearson, San Francisco, 2nd edition.
- [234] Salomon, G., J. Koepsell, J. Vijayan, T. A. Hilker, J. Nespolo, L. Pollet, I. Bloch, and C. Gross (2019). Direct observation of incommensurate magnetism in Hubbard chains. *Nature*, 565(7737):56–60.
- [235] Savard, T. A., K. M. O’Hara, and J. E. Thomas (1997). Laser-noise-induced heating in far-off resonance optical traps. *Physical Review A*, 56(2):R1095–R1098.
- [236] Schauß, P., M. Cheneau, M. Endres, T. Fukuhara, S. Hild, A. Omran, T. Pohl, C. Gross, S. Kuhr, and I. Bloch (2012). Observation of spatially ordered structures in a two-dimensional Rydberg gas. *Nature*, 491(7422):87–91.
- [237] Schlosser, N., G. Reymond, I. Protchenko, and P. Grangier (2001). Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature*, 411(6841):1024–1027.
- [238] Schlosser, N., G. Reymond, and P. Grangier (2002). Collisional Blockade in Microscopic Optical Dipole Traps. *Physical Review Letters*, 89(2).
- [239] Seltzer, S. J. and M. V. Romalis (2004). Unshielded three-axis vector operation of a spin-exchange-relaxation-free atomic magnetometer. *Applied Physics Letters*, 85(20):4804–4806.
- [240] Sewell, R. J., M. Koschorreck, M. Napolitano, B. Dubost, N. Behbood, and M. W. Mitchell (2012). Magnetic Sensitivity Beyond the Projection Noise Limit by Spin Squeezing. *Physical Review Letters*, 109(25).

- [241] Shannon, C. (1949). Communication in the Presence of Noise. *Proceedings of the IRE*, 37(1):10–21.
- [242] Sherson, J. F., H. Krauter, R. K. Olsson, B. Julsgaard, K. Hammerer, I. Cirac, and E. S. Polzik (2006). Quantum teleportation between light and matter. *Nature*, 443(7111):557.
- [243] Sherson, J. F., C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr (2010). Single-atom-resolved fluorescence imaging of an atomic Mott insulator. *Nature*, 467(7311):68–72.
- [244] Shevchenko, A., M. Kaivola, and J. Javanainen (2004). Thermodynamics of a multicomponent-atom sample in a tightly compressed atom trap. *Physical Review A*, 70(1):011403.
- [245] Shin, Y., M. Saba, T. A. Pasquini, W. Ketterle, D. E. Pritchard, and A. E. Leanhardt (2004). Atom Interferometry with Bose-Einstein Condensates in a Double-Well Potential. *Physical Review Letters*, 92(5).
- [246] Simon, J., W. S. Bakr, R. Ma, M. E. Tai, P. M. Preiss, and M. Greiner (2011). Quantum simulation of antiferromagnetic spin chains in an optical lattice. *Nature*, 472(7343):307–312.
- [247] Smerzi, A., S. Fantoni, S. Giovanazzi, and S. R. Shenoy (1997). Quantum Coherent Atomic Tunneling between Two Trapped Bose-Einstein Condensates. *Phys. Rev. Lett.*, 79(25):4950–4953.
- [248] Smith, A., B. E. Anderson, S. Chaudhury, and P. S. Jessen (2011). Three-axis measurement and cancellation of background magnetic fields to less than 50 μG in a cold atom experiment. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(20):205002.
- [249] Smith, G. A., S. Chaudhury, and P. S. Jessen (2003). Faraday spectroscopy in an optical lattice: A continuous probe of atom dynamics. *Journal of Optics B: Quantum and Semiclassical Optics*, 5(4):323–329.
- [250] Sompet, P., Y. H. Fung, E. Schwartz, M. D. J. Hunter, J. Phrompao, and M. F. Andersen (2017). Zeeman-insensitive cooling of a single atom to its two-dimensional motional ground state in tightly focused optical tweezers. *Physical Review A*, 95(3):031403.
- [251] Stamper-Kurn, D. M., H.-J. Miesner, A. P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle (1998). Reversible Formation of a Bose-Einstein Condensate. *Phys. Rev. Lett.*, 81(11):2194–2197.
- [252] Steck, D. (2010). *Rubidium 87 D Line Data*. 2.1.4 edition. Available at: <http://steck.us/alkalidata>.
- [253] Steck, D. (2017). *Quantum and Atom Optics*. 0.11.6 edition. Available at:

- <http://steck.us/teaching>.
- [254] Stellmer, S., B. Pasquiou, R. Grimm, and F. Schreck (2013). Laser Cooling to Quantum Degeneracy. *Physical Review Letters*, 110(26):263003.
 - [255] Streed, E. W., A. Jechow, B. G. Norton, and D. Kielpinski (2012). Absorption imaging of a single atom. *Nature Communications*, 3(1).
 - [256] Stuart, D. and A. Kuhn (2018). Single-atom trapping and transport in DMD-controlled optical tweezers. *New Journal of Physics*, 20(2):023013.
 - [257] Suominen, K.-A. (1996). Theories for cold atomic collisions in light fields. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 29(24):5981–6007.
 - [258] Sørensen, J. J. W. H., M. K. Pedersen, M. Munch, P. Haikka, J. H. Jensen, T. Planke, M. G. Andreasen, M. Gajdacz, K. Mølmer, A. Lieberoth, and J. F. Sherson (2016). Exploring the quantum speed limit with computer games. *Nature*, 532(7598):210–213.
 - [259] Sørensen, J. J. W. H., M. Dalggaard, A. H. Kiilerich, K. Mølmer, and J. F. Sherson (2018). Quantum control with measurements and quantum Zeno dynamics. *Physical Review A*, 98(6):062317.
 - [260] Takahashi, Y., K. Honda, N. Tanaka, K. Toyoda, K. Ishikawa, and T. Yabuzaki (1999). Quantum nondemolition measurement of spin via the paramagnetic Faraday rotation. *Physical Review A*, 60(6):4974–4979.
 - [261] Takano, T., M. Fuyama, R. Namiki, and Y. Takahashi (2009). Spin Squeezing of a Cold Atomic Ensemble with the Nuclear Spin of One-Half. *Physical Review Letters*, 102(3):033601.
 - [262] Tamosiune, A. (2016). ScienceAtHome.org | The Unforgettable April. Link: <https://www.scienceathome.org/community/blog/the-unforgettable-april/>. Accessed: 2019-10-16 17:21:04.
 - [263] Tarruell, L., D. Greif, T. Uehlinger, G. Jotzu, and T. Esslinger (2012). Creating, moving and merging Dirac points with a Fermi gas in a tunable honeycomb lattice. *Nature*, 483(7389):302–305.
 - [264] Texas Instruments (2016). DLP65000.65 1080p MVSP Type A DMD. Link: <http://www.ti.com/lit/ds/symlink/dlp6500flq.pdf>.
 - [265] Texas Instruments (2018). Getting Started With TI DLP®DisplayTechnology. Application report DLPA059C, Texas Instruments. Link: <http://www.ti.com/lit/an/dlpao59c/dlpao59c.pdf>.
 - [266] Thiele, T., Y. Lin, M. O. Brown, and C. A. Regal (2018). Self-Calibrating Vector Atomic Magnetometry through Microwave Polarization Reconstruction. *Physical Review Let-*

- ters, 121(15):153202.
- [267] Thomas, R. and N. Kjærgaard (2019). A digital pid controller for stabilizing large electric currents to the ppm level for feshbach resonance studies. *arXiv:1909.11257 [cond-mat, physics:physics]*.
 - [268] Thomas, R., M. Chilcott, C. Chisholm, A. B. Deb, M. Horvath, B. J. Sawyer, and N. Kjærgaard (2017). Quantum Scattering in an Optical Collider for Ultracold Atoms. *Journal of Physics: Conference Series*, 875:012007.
 - [269] Thorsen, A. R. (2015). *A Modular Control System for Cold Atom Experiments*. Master's thesis, Aarhus University, Aarhus.
 - [270] Tichenor, D. A. and J. W. Goodman (1972). Coherent Transfer Function*. *Journal of the Optical Society of America*, 62(2):293.
 - [271] Tomita, T., S. Nakajima, I. Danshita, Y. Takasu, and Y. Takahashi (2017). Observation of the Mott insulator to superfluid crossover of a driven-dissipative Bose-Hubbard system. *Science Advances*, 3(12):e1701513.
 - [272] Townsend, C. G., N. H. Edwards, C. J. Cooper, K. P. Zetie, C. J. Foot, A. M. Steane, P. Szriftgiser, H. Perrin, and J. Dalibard (1995). Phase-space density in the magneto-optical trap. *Physical Review A*, 52(2):1423–1440.
 - [273] van Oosten, D., P. van der Straten, and H. T. C. Stoof (2001). Quantum phases in an optical lattice. *Physical Review A*, 63(5):053601.
 - [274] Vengalattore, M., J. M. Higbie, S. R. Leslie, J. Guzman, L. E. Sadler, and D. M. Stamper-Kurn (2007). High-Resolution Magnetometry with a Spinor Bose-Einstein Condensate. *Physical Review Letters*, 98(20).
 - [275] Vijayan, J., P. Sompet, G. Salomon, J. Koepsell, S. Hirthe, A. Bohrdt, F. Grusdt, I. Bloch, and C. Gross (2019). Time-resolved observation of spin-charge deconfinement in fermionic hubbard chains. *arXiv:1905.13638 [cond-mat]*.
 - [276] Wade, A. C. J., J. F. Sherson, and K. Mølmer (2016). Manipulation of collective quantum states in Bose-Einstein condensates by continuous imaging. *Physical Review A*, 93(2):023610.
 - [277] Wang, Y., X. Zhang, T. A. Corcovilos, A. Kumar, and D. S. Weiss (2015). Coherent Addressing of Individual Neutral Atoms in a 3D Optical Lattice. *Physical Review Letters*, 115(4).
 - [278] Wasilewski, W., K. Jensen, H. Krauter, J. J. Renema, M. V. Balabas, and E. S. Polzik (2010). Quantum Noise Limited and Entanglement-Assisted Magnetometry. *Physical Review Letters*, 104(13).

- [279] Weber, T., J. Herbig, M. Mark, H.-C. Nägerl, and R. Grimm (2003). Bose-Einstein Condensation of Cesium. *Science*, 299(5604):232–235.
- [280] Weitenberg, C. (2011). *Single-Atom Resolved Imaging and Manipulation in an Atomic Mott Insulator*. PhD thesis, Ludwig-Maximilians-Universität München, München.
- [281] Weitenberg, C., M. Endres, J. F. Sherson, M. Cheneau, P. Schausz, T. Fukuhara, I. Bloch, and S. Kuhr (2011). Single-spin addressing in an atomic Mott insulator. *Nature*, 471(7338):319–324.
- [282] Weitenberg, C., S. Kuhr, K. Mølmer, and J. F. Sherson (2011). Quantum computation architecture using optical tweezers. *Physical Review A*, 84(3):032322.
- [283] Weitenberg, C., P. Schauß, T. Fukuhara, M. Cheneau, M. Endres, I. Bloch, and S. Kuhr (2011). Coherent Light Scattering from a Two-Dimensional Mott Insulator. *Physical Review Letters*, 106(21):215301.
- [284] Wigley, P. B., P. J. Everitt, K. S. Hardman, M. R. Hush, C. H. Wei, M. A. Sooriyabandara, P. Manju, J. D. Close, N. P. Robins, and C. C. N. Kuhn (2016). Non-destructive shadowgraph imaging of ultra-cold atoms. *Optics Letters*, 41(20):4795–4798.
- [285] Wildermuth, S., S. Hofferberth, I. Lesanovsky, E. Haller, L. M. Andersson, S. Groth, I. Bar-Joseph, P. Krüger, and J. Schmiedmayer (2005). Bose–Einstein condensates: Microscopic magnetic-field imaging. *Nature*, 435(7041):440–440.
- [286] Wildermuth, S., S. Hofferberth, I. Lesanovsky, S. Groth, P. Krüger, J. Schmiedmayer, and I. Bar-Joseph (2006). Sensing electric and magnetic fields with Bose-Einstein condensates. *Applied Physics Letters*, 88(26):264103.
- [287] Wineland, D. and H. G. Dehmelt (1975). Proposed $10^{14} \delta\nu/\nu$ laser fluorescence spectroscopy on Ti^+ mono-ion oscillator. *Bull. Am Phys. Soc.*, 20(637).
- [288] Wineland, D. J., R. E. Drullinger, and F. L. Walls (1978). Radiation-Pressure Cooling of Bound Resonant Absorbers. *Physical Review Letters*, 40(25):1639–1642.
- [289] Wood, A. A., L. M. Bennie, A. Duong, M. Jasperse, L. D. Turner, and R. P. Anderson (2015). Magnetic tensor gradiometry using Ramsey interferometry of spinor condensates. *Physical Review A*, 92(5).
- [290] Yamamoto, R., J. Kobayashi, T. Kuno, K. Kato, and Y. Takahashi (2016). An ytterbium quantum gas microscope with narrow-line laser cooling. *New Journal of Physics*, 18(2):023016.
- [291] Yamamoto, R., J. Kobayashi, K. Kato, T. Kuno, Y. Sakura, and Y. Takahashi (2017). Site-resolved imaging of single atoms with a Faraday quantum gas microscope. *Physical Review A*, 96(3):033610.

- [292] Yang, F., A. J. Kollár, S. F. Taylor, R. W. Turner, and B. L. Lev (2017). Scanning Quantum Cryogenic Atom Microscope. *Physical Review Applied*, 7(3):034026.
- [293] Zeiher, J., J.-y. Choi, A. Rubio-Abadal, T. Pohl, R. van Bijnen, I. Bloch, and C. Gross (2017). Coherent Many-Body Spin Dynamics in a Long-Range Interacting Ising Chain. *Physical Review X*, 7(4).
- [294] Zupancic, P. P. J. (2013). *Dynamic Holography and Beamshaping Using Digital Micromirror Devices*. Master's thesis, Ludwig-Maximilians-Universität München, München.
- [295] Zwerger, W. (2003). Mott Hubbard transition of cold atoms in optical lattices. *Journal of Optics B: Quantum and Semiclassical Optics*, 5(2):S9–S16.



This thesis covers experimental work conducted within two distinct fields of cold-atom physics. The first part describes an experiment capable of spatially-selective dispersive measurements of a cold-atomic system in part or in whole, by virtue of a digital micromirror device (DMD). Atom clouds are trapped in optical tweezers made and controlled by an acousto-optical deflector. Two types of high-precision magnetometers are realised. One is vectorial and reaches a single-shot precision of $\delta B = (100, 200) \mu\text{G}$ for the field components parallel and transversal to the probe light, exceeding a previous realisation by two orders of magnitude. The other is a scalar magnetometer based on Larmor precession and reaches $\delta B = 30 \mu\text{G}$, putting it on equal footing with other state-of-the-art cold-atom magnetometers.

The second part is devoted to the characterisation and the first results of a new quantum gas microscope experiment. It features a 0.69 NA microscope objective for high-resolution fluorescence imaging of individual atoms trapped in deep optical lattices. The imaging system is shown to perform near the diffraction limit. By means of DMD-generated off-resonant tight optical tweezers, projected through the high-resolution optics, we can load only a few planes of a co-propagating 1D lattice. In a single realisation of the experiment we acquire multiple fluorescence images, where the objective is translated between images, bringing different planes of the optical lattices in focus. In this way we can tomographically reconstruct the atom distribution in 3D.



AARHUS UNIVERSITY