Bose-Einstein condensates in optical lattices

Henrik Kjær Andersen

QUANTOP
Department of Physics and Astronomy
University of Aarhus, Denmark

PhD Thesis
September 2008

This thesis is submitted to the Faculty of Science at the University of Aarhus, Denmark, in order to fulfill the requirements for obtaining the PhD degree in Physics.
The studies have been carried out at the Department of Physics and Astronomy under supervision of Associate Research Professor Michael Budde and Associate Professor Michael Drewsen, from August 2004 to September 2008.

An electronic version of this thesis is available at:
`http://www.phys.au.dk/quantumgas/PhD_hka.pdf`

Acronyms

ADC	Analog-to-Digital Converter
AOM	Acousto Optical Modulator
AWG	Arbitrary Waveform Generator
BEC	Bose-Einstein Condensate
CCD	Charge-Coupled Device
CCS	Camera Control System
CMOT	Compressed Magneto-Optical Trap
DAC	Digital-to-Analog converter
DDS	Direct Digital Synthesis
ECDL	Extended Cavity Diode Laser
ECS	Experimental Control System
EMCCD	Electron-Multiplying Charge Coupled Device
EOM	Electro-Optical Modulator
FMS	Frequency Modulation Spectroscopy
FPGA	Field-programmable Gate Array
LED	Light-Emitting Diode
MPS	Mechanical Positioning System
MOT	Magneto-Optical Trap
OD	Optical Density
OI	Optical Isolator
PBS	Polarizing Beamsplitter
PDH	Pound-Drever-Hall
PGR	Point Grey Research
PID	Proportional-Integral-Derivative
QP	Quadrupole
QUIC	Quadrupole-Ioffe Configuration
RF	Radio Frequency
TA	Tapered Amplifier
TF	Thomas-Fermi
TOF	Time-Of-Flight
TSP	Titanium Sublimation Pump
VCO	Voltage-Controlled Oscillator

Contents

Publications

[I] J. F. Bertelsen, H. K. Andersen, S. Mai, and M. Budde, *Mixing of ultracold atomic clouds by merging of two magnetic traps*, Phys. Rev. A 75, 013404 (2007).

We are currently doing experiments that will lead to more papers—see chapter 9.

Abstract

Dansk resumé. *I min eksperimentelle ph.d.-uddannelse har jeg arbejdet med at opbygge et laboratorium til at køle en gas af atomer fra stuetemperatur ned til det absolutte nulpunkt, hvorved man opnår et såkaldt Bose-Einstein kondensat. Et Bose-Einstein kondensat er interessant, fordi det kan bruges til at studere fundamentale fysiske (kvante) egenskaber, som ikke kan undersøges ved stuetemperatur. Derudover har jeg undersøgt, hvordan man kan ændre egenskaberne af kondensatet vha. laserstråler. I fremtiden håber man at kunne bruge den grundforskning der i øjeblikket bliver lavet indenfor dette område i forskellige teknologiske opfindelser.*

English abstract. During my PhD studies I have participated in building a laboratory for cooling a gas of atoms from room-temperature down to the absolute zero, where a so-called Bose-Einstein condensate is formed. A Bose-Einstein condensate can be used to study fundamental properties of quantum physics, which cannot be examined in a gas of atoms at room-temperature. I have also studied how it is possible to change the properties of the condensate with laser beams. In the future Bose-Einstein condensates will perhaps be used in different technological inventions.

Preface and outline

Group history

In the beginning of 2003 Michael Budde started *The Aarhus Quantum Gas Laboratory*. The goal was to build a modern quantum gas laboratory with the opportunity to make many different types of experiments. However, the first step on the way was to make a Bose-Einstein Condensate (BEC) of rubidium 87—we managed to do this in March 2006. The second subsidiary goal was to build a 3D optical dipole/lattice setup. This has now been finished and both the general BEC setup and the optical lattice setup are well-characterized and optimized. Another subsidiary goal was to add a second species (lithium 6 and/or 7) to the setup, but so far this project has been put on standby (see below).

From the beginning Jesper Fevre Bertelsen and I joined the group and together with Michael we started to build the experiment. Jesper Fevre Bertelsen became a PhD student in the summer of 2003 and I finished my bachelor report about the laser systems later that year. Jesper handed in his partly experimental and partly theoretical PhD-thesis in the summer of 2007. Until the summer of 2004 I continued to work in the laboratory on different projects. In the summer of 2004 I became a PhD-student, and I started working on the computer control systems.

In February 2005 Christian Schori joined the group as a PostDoc working on the optical lattices, but in the summer of 2005 he was offered another job and decided to leave us. Up until this point the plan was that I should have built the lithium part of the setup, but instead we now decided that I should work on the optical lattices after finishing my work on the BEC setup. In the autumn of 2005 Sune Mai joined the group as a MSc-student and he started working on the experimental sequence for getting a BEC together with Michael. In the beginning of 2007 Sune handed in his theoretical master thesis and became a PhD-student. Since then Sune has been working mostly with theory. In the beginning of 2006 Martin Larsen joined the group as a bachelor student and later that year he handed in a report about our microwave equipment and experiments performed with it. Also Otto Højager Attermann Nielsen has been a BSc-student in the group, and he has done measurements on BECs close to the critical temperature T_c, and compared the results to theoretical simulations. Otto handed in his BSc-thesis in the beginning of 2008.

In the end of 2007 Sung Jong Park joined the group as a PostDoc. Sung Jong had been working with atomic clocks and the plan was that he should have built the lithium setup. However in November 2007 Michael Budde decided to leave the group, and he is now working for a private company. For this reason Michael did not want us to start another project, so Sung Jong and I have since then been working close together on different experiments with the optical lattice setup. Currently we are still without a group leader and will be until sometime in 2009.

Thesis outline

This thesis presents the work I have done during my PhD-studies under daily supervision of Michael Budde and formal supervision of Michael Drewsen. My contribution to the laboratory has mainly been in building different parts of the experiment. I have spend almost all my time working in the laboratory on three different projects:

- Building the BEC-setup—this part was mainly done together with Jesper and Michael. Later when we had made our first BEC I worked on optimizing and improving the setup in order to get everything stable.
- Working on the computer control systems for the laboratory.
- Building the 3D optical lattice setup. Christian Schori worked on the initial phases of this project. Sung Jong and I have done all the experiments with the optical lattices together.

My work has been divided approximately equally between these three projects—each represents approximately 1½ years of work (including attending courses and teaching).

This thesis is divided into two main parts. The first part will describe the setup for making a BEC including the computer systems and some experiments performed with the setup (chapter 2–5). The second part will contain information about the optical lattice setup, and some of the measurements we have conducted using the optical lattices (chapter 6–8). A slightly more detailed outline is given in the following:

In chapter 1 a short introduction to the field of quantum gases (and quantum gases in optical lattices) is given.

In chapter 2 the experimental setup for making a BEC of rubidium 87 is described.

In chapter 3 the Experimental Control System (ECS) is described. This system controls all the equipment in the laboratory during the experiments.

In chapter 4 the Camera Control System (CCS) is described. This computer system is used for grabbing images (which is the outcome of our experiments).

In chapter 5 it is described how we analyze the grabbed images, and different experiments performed with the basic BEC-setup are presented.

In chapter 6 some general theory of BECs in optical lattices is introduced. This theory will mainly be used in chapter 8.

In chapter 7 the 3D optical lattice setup is described.

In chapter 8 we will look at different experiments done with the optical lattice setup.

In chapter 9 a conclusion to the thesis will be given and future plans will be discussed.

Because of the lack of a group leader, I have been thinking much about the next generation of people in the laboratory during the writing of this thesis. This means that I have provided a little more details on the setup, and what we have been working on, than I otherwise would have done. Also, all the data I present in this thesis will contain information about the runs it is based on. For this reason it will be easy for future group members to go back and find the data analysis and see in details how the experiments were performed. I hope in this way we can pass on as much of our knowledge and experience with the system as possible. Fortunately Sune was a big part of making the first BEC in our group, and he is very experienced in running the system. Also, Sung Jong is becoming very experienced with the system, and he will hopefully continue for one more year in the group.

Acknowledgments

First of all, I would like to deeply thank my two supervisors, Michael Drewsen and especially Michael Budde. I have never regretted that I decided to work for you on building a completely new and very complex laboratory. I have learned *very* much from you and from my work in general—both regarding physics but also regarding working methods, leadership and much more.

Next, I would like to thank Jesper. Jesper and I worked together on the BEC experiment, and we developed a close friendship. I wish him all the best of luck in the future, and I hope we can meet once in a while and go to a concert or in the cinema.

Also, I want to thank Sune. I hope (and think) he will do many great experiments in the laboratory during the next couple of years, and I am sad that we did not have more time to play with FPGAs together.

Next, I will give my most profound thanks to Sung Jong for all the work he has done on the optical lattice experiments. It has been a distinct pleasure for me to work close together with him for the last many months. I wish we could have more time together to do new great experiments. I also hope he and his wife and daughter will have a good time in Denmark or wherever they decide to live. I also want to thank him very much for proofreading this thesis.

Next, I want to thank Christian who started the optical lattice experiment. I am very sad that he decided to leave the group since he is a very sympathetic and skilled person.

I also want to mention the other two former bachelor students Martin and Otto. They are both very pleasant (and fun) company, and I hope they will find the right places to continue their studies (perhaps in our group) and reach their goals.

Next, I want to mention the theoretical group (especially I thank Nicolai Nygaard very much). We have discussed different issues regarding the optical lattices after Michael left. This has been very helpful, and I hope we can continue the collaboration in the future.

It is also very important for me to thank the technical staff at the University of Aarhus—especially Erik Søndergaard and Frank Mikkelsen from the electronics department for the enormous amount of work they have done on different equipment and helping us in general. Also, I want to thank the staff in the workshop and especially Torben Hyltoft Thomsen and Uffe Simonsen.

Finally I would like to thank my friends (both inside and outside the university) and my family. In particular, I want to thank Vivi Petersen very much for her support and patience, and also for proofreading this thesis.

Henrik Kjær Andersen
September, 2008

One

Introduction

1.1 Bose-Einstein condensation

One of the most dynamic research areas within atomic physics today is the field of quantum gases. An atomic quantum gas is a gas of atoms cooled so much that it can no longer be described by Maxwell-Boltzmann statistics. Instead the gas can be described by either Bose-Einstein statistics (if the atoms are bosons) or Fermi-Dirac statistics (if the atoms are fermions). This means that by cooling a gas of atoms down to temperatures very close to the absolute zero a completely new world opens. In this world quantum physics rules and for this reason the wave-nature of the atoms becomes visible. In 1924 Louis de Broglies postulated that matter has a wave-like nature and that the characteristic wavelength is $\lambda_{db} = h/p$, where h is Planck's constant and p is the momentum of the matter. If this wavelength is calculated for an atom with a large (thermal) velocity it will be very small and the corresponding wave packet (or extension of the atom) will also be very small. However when the temperature drops this wavelength becomes larger and larger and close to absolute zero λ_{db} becomes comparable to the average distance between the atoms. For bosons the atomic waves of all the different, indistinguishable atoms will start to overlap and a Bose-Einstein Condensate (BEC) will be formed [1]. In this quantum phase all the atoms are in the same quantum state and share coherently the same macroscopic wavefunction.

The condition for getting a BEC (or for having the different wave packets overlapping) can be written in terms of the so-called phase-space-density (the number of atoms within a volume λ_{dB}^3) [1]

$$n\lambda_{\mathrm{dB}}^3 = n\left(\frac{2\pi\hbar^2}{mk_BT}\right)^{3/2} \gtrsim 1, \tag{1.1}$$

where $\lambda_{\mathrm{dB}} = (2\pi\hbar^2/mk_BT)^{1/2}$ is the de Broglie wavelength, n the density of the atoms, m the mass of an atom, k_B the Boltzmann constant and T the temperature. When the phase-space-density goes above a critical value around unity the atomic wave packets will overlap and a large fraction of the atoms will form a BEC. If we isolate the temperature we get a condition for the critical temperature T_c, at which the BEC is formed[1]

---

[1] Notice that eq. (1.1) is only an approximate expression based on simple arguments. A more thorough analysis is made in [2], where the critical temperature is calculated for different trapping potentials.

$$T_c \lesssim \frac{2\pi\hbar^2}{mk_B} n^{2/3}. \tag{1.2}$$

A BEC is typically formed in clean, dilute gases with densities of about $n = 10^{14}\text{cm}^{-3}$ at a critical temperature around $T_c = 1\,\mu\text{K}$.

Another way to look at Bose-Einstein condensation is given in the following. The way we make a BEC is by trapping and cooling a gas of neutral atoms in an external potential $V(\mathbf{r})$ (e.g. formed by a magnetic field). In this potential the atoms are occupying the different single-particle motional states (with energy ϵ_i) of the trap. The mean occupation number $\overline{n}_i$ of atoms in the i-th state is given by the Bose-Einstein distribution [3]

$$\overline{n}_i = \frac{1}{e^{(\epsilon_i - \mu)/(k_B T)} - 1}, \tag{1.3}$$

here μ is the chemical potential, k_B is the Boltzmann constant, and T is the temperature. One way to think about the chemical potential μ is that it just determines the number of atoms N for a given temperature T since

$$N = \sum_i \overline{n}_i. \tag{1.4}$$

At high temperatures (where μ is large and negative [3]) the Bose-Einstein distribution reduces to the Maxwell-Boltzmann distribution where $\overline{n}_i \ll 1$ for all i. When we cool the gas the mean occupation number in the lowest states becomes larger and larger and the chemical potential approaches the ground state energy ϵ_0. In the limit where T goes to zero all the atoms will end up in the ground state and μ will be equal to ϵ_0 (so the maximum possible value of the chemical potential is the ground state energy—otherwise the mean occupation number could be negative). This phenomenon is called Bose-Einstein condensation and the cloud of atoms (all being in the ground state) is called a Bose-Einstein condensate.

Bose-Einstein condensation is a phase transition which only happens in bosonic systems. When cooled sufficiently fermions will form a degenerate Fermi gas since they cannot share the same quantum state due to the Pauli exclusion principle. This effect becomes significant when the temperature of the fermion gas goes below the Fermi temperature T_F. At this temperature the state occupation of the lowest energy levels (ϵ_i) has increased from essentially zero to almost one. For electrons in metals this happens already below many thousands of kelvins, but for a dilute atomic gas it does not happen until the temperature is lowered to around $1\,\mu\text{K}$. Currently we are not studying Fermi gases in the laboratory, but we will perhaps do it in the future.

Since Bose-Einstein condensation is a quantum phenomenon we need to use quantum mechanics when we want to describe our experiments. This normally includes solving the Schrödinger equation and finding wave functions of the system (in this case the condensate). However the wave equation that is used to describe BECs is a non-linear Schrödinger equation called the Gross-Pitaevskii equation [4]

$$\left[-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + gN|\psi(\mathbf{r})|^2 \right]\psi(\mathbf{r}) = \mu\psi(\mathbf{r}) \tag{1.5}$$

where N is the atom number (it is assumed that the atom number is large so that $N \approx N - 1$), m is the mass of an atom and $g = 4\pi\hbar^2 a/m$ is an interaction parameter where a is the s-wave scattering length.

This equation differs from the normal Schrödinger equation by the fact that there is a third term that describes the interaction between the different particles. It is still a one-particle equation but each particle experiences a mean interaction from all the other particles given by the interaction parameter g.

One can think about this equation in the following way. The first (kinetic energy) term becomes small if the wave function spreads out (becomes wide). The second (potential energy) term becomes small if the wave function becomes more narrow. The last term can be both positive and negative (depending on a). For positive a this term tries to spread out the wave function and make the condensate bigger (like the kinetics energy term). In order for a BEC not to collapse/implode during the cooling g (and therefore a) has to be positive[2].

1.1.1 BEC experiments

Cornell, Wiemann and co-workers at JILA were the first to make a BEC of a dilute gas of ^{87}Rb atoms in 1995 [6] (70 years after Bose and Einstein made the theoretical prediction in 1925). A few months later Ketterle and co-workers at MIT made a condensate of ^{23}Na [7]. Cornell, Wiemann and Ketterle won the 2001 Nobel Prize in physics for their work [8, 9]. Since then the field has grown enormously and today there are many BEC laboratories around the world.

A large number of species have been Bose condensed including H, ^{4}He, ^{7}Li, ^{23}Na, ^{41}K, ^{52}Cr, ^{85}Rb, ^{87}Rb, ^{131}Cs, ^{174}Yb and latest ^{170}Yb [10] and ^{39}K [11] ([10] also contains references to all the previous mentioned species). Degenerate Fermi gases have been made with e.g. ^{40}K [12] and ^{6}Li [13]. The reason why these lists are not longer (all though they are already long) is that it is very hard to do laser cooling of non-alkali metals.

Since the first BEC was created in a dilute, neutral gas many exciting experiments have been performed. The wavelike nature (or wave-particle duality) of a BEC has e.g. been studied in experiments where two BECs are released from two different traps and allowed to interfere with each other. In these experiments interference fringes are seen (just like laser beams can form interference patterns) [14]. Vortices and superfluidity have been investigated extensively in BECs (e.g. [15, 16]). Quantized vortices are a signature of the superfluidity of a BEC and they can e.g. be formed by stirring a condensate with a laser beam. Observation of Feshbach resonances in BECs have proved to be a very effective tool in tuning the interaction strength (or scattering length) between the atoms (e.g. [17]). By sweeping across a Feshbach resonance it is e.g. possible to form a molecular BEC from paired fermions [18–20]. It has also been demonstrated that it is possible to slow down the speed of light using a BEC [21]. Recently it has been shown how to transform light into matter and back into light again using BECs [22]. These studies show how optical information can be stored in (or imprinted on) a BEC, altered and finally written back onto propagating optical fields. This might lead to dynamical controllable optical delay lines which e.g. can be used in optical communication systems.

Today, many groups are going for two or three species (bosons and/or fermions) and many other advanced extensions to these experiments in order to be able to make (and manipulate) new quantum systems of ultracold atoms. With more than one species it becomes possible to study mixtures and form heteronuclear molecules. Also, the pursuit of a continuous-wave atom laser and a quantum computer (probably far into the future) is some of the big goals of the field. However there are also room for more alternative (but likewise great) experiments like e.g. a “drop tower BEC experiment” where BECs can be studied in micro-gravity [23]. In general the BEC-field is slowly moving towards commercial products.

---

[2] Cesium 133 is an example where there are problems with a negative scattering length—however a can be tuned using a Feshbach resonance and a stable BEC of cesium atoms can still be achieved [5].

Above I have mentioned a few examples of some of the fields within the BEC area. This is by no means an attempt to make a complete list of all the interesting experiments that have been performed within this area—this would be almost impossible.

1.2 Optical lattices

Another very hot field within this area is to study ultracold quantum gases in optical lattices. An optical lattice is a standing wave of monochromatic light—a periodic intensity pattern formed by the interference of two or more laser beams. Depending on the wavelength of the lattice light compared to the electronic transition frequencies within an atom, the atoms can be trapped in regions of either maximum or minimum intensity. The simplest possible optical lattice consists of two counter propagating beams with the same wavelength. In this light field a one-dimensional array of microscopic trapping potentials are formed. If more laser beams are used it is possible to form a periodic 3D crystal-like structure that very much resembles the crystal structure of a material (known from solid state physics). For this reason atoms loaded into a 3D optical lattice will in many ways resemble electrons moving in the periodic potential of a solid state crystal. However optical lattices provide many advantages—for example[3]

- The atoms move in a potential which can be dynamically controlled. The depth of the potential can be changed just by changing the intensities of the different optical lattice beams. By doing this the interactions between the atoms can also be varied. Furthermore the potential can be switched on and off and modulated. Also, by choosing slightly different frequencies of two counter propagating beams, it is possible to make a moving lattice where the different sites (or traps) move with a certain controllable velocity[4].
- The lattice geometry and the lattice constant can be chosen. It is e.g. possible to make sc, fcc and bcc crystal structures. The lattice constant ($\lambda/2$) can (at least to some extend) be chosen by choosing the wavelength λ of the optical lattice laser. Furthermore, if two laser beams are used with a mutual angle less than 180° the period of the optical lattice will be longer than $\lambda/2$. Also, it is e.g. possible to change the angle of the different beams and thereby change the crystal structure dynamically (e.g. [25]). Another opportunity is to make so-called super-lattices, where two (or more) lattice beams with different wavelengths on top of each other are used (e.g. [26]).
- In an optical lattice the dimension of the quantum gas can be changed from 3D to 1D or 2D.

It is really remarkable how much physics we can learn by exploring the huge amount of parameters, which are experimentally controllable in systems with quantum gases in optical lattices. An example of how an optical lattice can change the properties of a quantum gas will be given in the following.

If a BEC is loaded adiabatically into a periodic potential (only populating the lowest Bloch band) it can be described by the so-called Bose-Hubbard model. This model is known from solid state physics, but the link to BECs in optical lattices was made in [27]. The model is capable of describing a quantum phase transition that can be seen when a BEC (superfluid) is placed in an optical lattice. The phase transition is a consequence of

---

[3] A longer list can be found in [24].

[4] The velocity of the moving lattice is given by $v = \Delta\nu\lambda/2$, where λ is the wavelength of the optical lattice and $\Delta\nu$ is the relative frequency detuning between the two beams.

two competing interactions, namely the tendency of particles to jump into neighboring lattice sites, and the interparticle force that tries to separate the atoms from each other and distribute them across the entire lattice. These interactions are accounted for in the Bose-Hubbard Hamiltonian (and therefore it can also describe the phase transition [24, 27–29])

$$H = -J \sum_{\langle i,j \rangle} b_i^\dagger b_j + \frac{1}{2} U \sum_i \hat{n}_i (\hat{n}_i - 1) \tag{1.6}$$

where b_i ($b_i^\dagger$) is the annihilation (creation) operators of atoms on the i-th site, and $\hat{n}_i = b_i^\dagger b_i$ is the number of atoms on site i. $\langle i,j \rangle$ indicates a sum over nearest neighbors.

The first (kinetic energy) term describes the tunneling/hopping coupling J of bosons between nearest neighbor sites. The second term represents the onsite, repulsive interaction U between bosons.

In the superfluid phase ($U/J \ll 1$) no exact number of atoms are located on each lattice site, but instead each atom is delocalized over the entire lattice. In this case there is a phase coherence between the atomic wavefunctions on the different lattice sites, and this leads to interference patterns when the condensate is released from the lattice. In the Mott insulator phase ($U/J \gg 1$) there is a fixed number of atoms on each lattice site, and no phase coherence (no interference) is seen. The phase transition was first observed in the group of I. Bloch [29] in 3D (in a 3D optical lattice) and later in the group of T. Esslinger in 1D (also in a 3D optical lattice but with a high intensity along two axes allowing only for tunneling in one-dimension) [30]. The transition can be seen just by changing the lattice laser intensity. By doing this the tunneling to neighboring sites (dominating in the superfluid phase, where the lattice is shallow) as well as the strength of the on-site repulsive interactions (dominating in the Mott insulator phase, where the lattice is deep and the tunneling rate is low) can be tuned. This experiment is one of the most famous experiments performed with BECs in optical lattices. However within this field there have also been conducted many other intriguing experiments. A few of them are described in the following.

Bose-Einstein condensation can be studied in fewer than three dimensions e.g. in a so called Tonks-Girardeau gas. A Tonks-Girardeau (TG) gas is made in a 1D Bose gas in an optical lattice [31–33]. In the TG regime the repulsive interactions between the atoms dominate the kinetic energy. This means that the atoms (although strongly interacting) in some ways begin to act as non-interacting fermions, since the atoms are prevented from occupying the same position in space (in order to minimize their mutual repulsion). Bose/Fermi mixtures in optical lattices have also been realized (e.g. [34, 35]). In [35] it is studied how the superfluid to Mott-insulator transition depends of the concentration of the fermionic (impurity) species[5].

In the more "commercial direction" an optical lattice (atomic) clock has been realized in [36]. Here the frequency of an optical transition within atoms loaded into an optical lattice is measured. Since this atomic clock is based on an optical transition (rather than a microwave transition, which is used e.g. in the current cesium fountain clocks), it is a strong candidate to the atomic clocks of the future. Quantum computing with neutral atoms trapped in an optical lattice is also a field of great interest. The atoms in the different lattice sites can act as qubits, and one of the big goals within this field is currently to achieve single site addressability in optical lattices [37, 38].

A more detailed review of quantum gases in optical lattices can be found in [24, 39, 40].

---

[5] Being able to make these kind of advanced experiments, where different species are mixed in an optical lattice, was what we had in mind when we started building the laboratory. Except the second species (which has been put on hold), we have most of the needed features today.

Two

BEC setup

2.1 Introduction

In this chapter the experimental setup for producing BECs in our laboratory is described. I will start with an overview of the experimental steps we go through to make a BEC. After this I will go into more details about e.g. how the laser system works, how the vacuum chamber is designed and how we have implemented some of the different experimental sequences.

I have contributed to building most of the setup described in this chapter together with Jesper and Michael. Additional information about the setup can be found in Jesper's PhD thesis [41].

2.2 Overview of the experimental steps

The experiment is placed on two optics tables. One of them is used for all the lasers and most of the optics. It will throughout this thesis be referred to as the *laser table*. On this table all the light used in the experiment is generated. Detailed information about the setup on this table can be found in section 2.3. The other table is used for the vacuum chamber, optics and much more. This table will be referred to as the *vacuum table* throughout this thesis and here the BEC is created. Details about the setup on this table will be given in section 2.4. The two optics tables are optically connected through single mode, polarization maintaining fibers.

In this short overview section we will focus on the setup on the vacuum table which can be seen in figure 2.1.

There are four main experimental steps to create a BEC in our laboratory

1. MOT, CMOT, state preparation and catch
2. Transport and transfer
3. Evaporative cooling
4. Absorption imaging

First the atoms are cooled from room temperature to about 150 µK in the MOT-chamber (see figure 2.1) using laser cooling (step 1). The atoms are then transported to

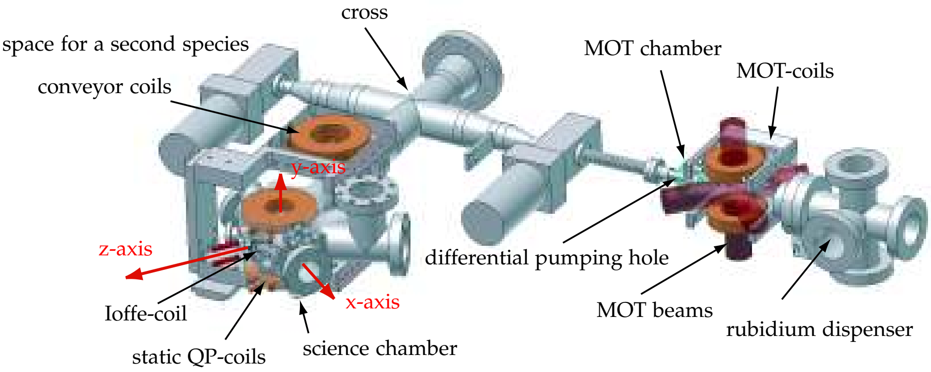

Figure 2.1: The vacuum chamber placed on the vacuum table. On this drawing only the most important parts are shown. At the cross different arms meet: the rubidium MOT arm, the not yet created second species MOT arm and the arm going to the science chamber (the last arm of the vacuum chamber is used for an ion-pump). See the text for further comments.

the science chamber (step 2) and cooled down to the BEC regime (<500 nK) using evaporative cooling (step 3). At the end images of the BEC are grabbed (step 4). After the BEC is formed and before the images are grabbed (between step 3 and 4) various experiments can be done with the BEC—e.g. the BEC can be loaded into an optical lattice.

1. MOT, CMOT, state preparation and catch

The first step towards a BEC in our laboratory is to cool and trap about 2×10^9 rubidium 87 atoms in a so-called Magneto-Optical Trap (MOT) located in one of the arms of the vacuum chamber. By sending current through a rubidium dispenser atoms are released out into the MOT chamber (a cylindrical glass-cell) where they can be cooled and trapped using laser beams and a magnetic Quadrupole (QP)-field (formed by two magnetic coils in an anti-Helmholtz configuration). More details about the MOT will be given in section 2.4.3.

In the MOT-phase we are running 17 A through the MOT-coils which gives a field gradient of about 11 G/cm along the axial direction. The field gradients in the MOT-phase are not steep enough to keep the atoms suspended when the laser beams are turned off. However if the current in the MOT-coils is cranked up to e.g. 250 A the atoms are caught in a purely magnetic trap—this procedure is referred to as *catch* and it ends the first step towards a BEC. The reason why we want to trap the atoms in a purely magnetic trap is that we want to move them to another part of the vacuum chamber.

Before catch some preparation is needed. The MOT-cloud is about 5 mm in diameter so a fast increase of the current in the MOT-coils would result in a large, unwanted increase in potential energy of the outermost atoms in the trap resulting in heating. To avoid this a so-called Compressed Magneto-Optical Trap (CMOT) is created in the last $\sim$ 35 ms of the MOT period [42]. Here the atomic cloud is spatially compressed to minimize this effect. The CMOT-phase needs to be short otherwise we will lose all the atoms (more details in section 2.4.4).

In the MOT-phase the atoms are (most of the time) distributed in the different m_F states of the $F = 2$ ground state (see figure 2.2). In the end of the CMOT-phase (just before catch) we put the atoms in the same F, m_F state. Normally we choose the $F = 2$,

$m_F = +2$ state. This state is a low-field seeking state which means that the atoms can be trapped in the magnetic field minimum set by the QP-coils (details are given in section 2.4.5).

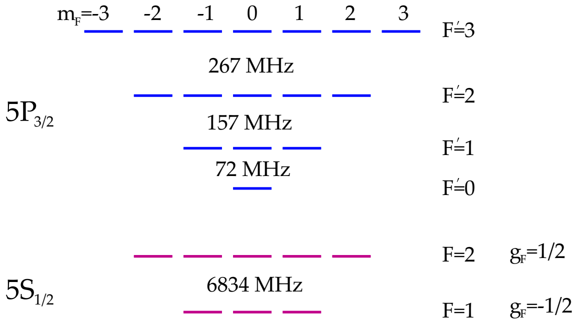

Figure 2.2: The rubidium 87 D2 line ($5^2S_{1/2}$ to $5^2P_{3/2}$—780 nm)—not to scale. Rubidium has two naturally occurring isotopes—Rb-85 (72%) and Rb-87 (28%).

2. Transport and transfer

For the MOT to be efficient a relatively high partial pressure of rubidium is needed in the MOT-chamber (about 1×10^{-10} torr)—otherwise we will simply not trap enough atoms to be able to make a BEC. However, we also need a long lifetime of the trapped atomic cloud since the upcoming evaporative cooling step takes a long time ($\sim$ 30 s). With a pressure of 1×10^{-10} torr there will be a lot of collisions with background atoms at room temperature and this will destroy the cloud before the BEC is created. For this reason we move the trapped atoms to better vacuum. The vacuum chamber is divided into two parts connected by a 4 mm in diameter *differential pumping hole* (see figure 2.1). The pressure in the low pressure part of the vacuum chamber is extremely low (about 5×10^{-12} torr) and this ensures a lifetime long enough to be able to create a BEC.

The atomic cloud (trapped in the QP-trap formed by the MOT-coils) can be moved through the differential pumping hole just by moving the coils mechanically. About 50% of the atoms are lost in this procedure but when the atoms reach the good pressure part they can be moved all the way to the science chamber practically without any loss or heating. The MOT-coils can move all the way from the MOT-chamber to the *cross* (see figure 2.1). After this the conveyor coils (also forming a QP-trap) take over and move the atoms from the cross to the science chamber. The moving system used in this transport phase is referred to as the Mechanical Positioning System (MPS) and it is based on large motorized translation stages (it will be described in details in section 3.5). At the science chamber the atomic cloud is transfered from the conveyor coils to static QP-coils mounted on the science chamber.

3. Evaporative cooling

The next event is the evaporative cooling step [1, 42, 43]. The idea in this procedure is to remove the most energetic atoms from the cloud and thereby increase the phase-space-density to get a BEC.

The stationary QP-coils mounted on the science chamber are not sufficient to keep the atoms trapped when further cooling towards the critical temperature is done. The problem is that the QP-field has a zero-point (where no magnetic field is present) and even though this is not a problem with $\sim 150\,\mu$K hot atoms it becomes a problem when we want to lower the temperature further. The reason is that the atoms can make a non-adiabatic spin-flip transition when they cross the magnetic zero-point at the center of the trap. This means that they can change state from a trapped sub-state (e.g. the $m_F = +2$ state) to an un-trapped sub-state and disappear from the trap (Majorana spin flip [44]). One of the ways to solve this problem is to use a third coil—a so-called Ioffe-coil placed on the z-axis (see figure 2.1). This coil together with the QP-coils form a so-called Quadrupole-Ioffe Configuration (QUIC)-trap. By choosing the right current ratio between the QP-coils and the Ioffe-coil a harmonic potential with no zero-point can be obtained (instead this potential has a bias field at the center of the trap). When we ramp up the current in the Ioffe-coil (i.e. when we go from the QP-trap to the QUIC-trap) the trap minimum and thereby the atoms are moved about 9 mm towards the Ioffe-coil.

The most energetic atoms in the QUIC-trap are removed by a Radio Frequency (RF) "knife" that transfers atoms that can reach a certain height (set by the RF-field) in the harmonic, magnetic trap to an un-trapped (high field seeking) sub-state ($F = 2$, $m_F = +2$ atoms are transfered to the $F = 2$, $m_F = -2$ state). The RF-field is provided by an antenna inside the science chamber and the RF-frequency must be ramped down slowly so only the most energetic atoms are lost and so that the remaining atoms have time to thermally equilibrate to a lower temperature through elastic collisions (only the tail of the Maxwell Boltzmann distribution should be cut away [43]). At the same time it must be done fast compared to the lifetime of the atoms in the science chamber. In practice this means that cooling times of 25 s are normal in our setup. It is not obvious how this frequency ramp should be created but an exponentially decreasing ramp is a good starting point [45]. We use an exponentially decreasing ramp with a time constant of 15 seconds ramping down from 17 MHz to about 700 kHz in 25 s. However in our case the evaporative cooling is first done in the QP-trap before the atoms are transfered to the QUIC-trap and then continued. This is done to remove the most energetic atoms (making the cloud smaller) and thereby avoid losses when the atoms are transfered to the QUIC-trap. During the evaporative cooling step the magnetic field must be very stable—otherwise the trap minimum will bounce and the RF-cooling will not work properly. This means that there are very high demands on the current controllers for the QUIC-coils.

If all these steps are performed correctly we end up with a BEC with about 3×10^5 atoms with a temperature below 500 nK. The BEC can now be used in further experiments.

4. Absorption imaging

The outcome of a quantum gas experiment is normally images. We send imaging light (resonant with the cyclic $F = 2$ to $F' = 3$ transition) through the atoms and onto a Charge-Coupled Device (CCD) camera—if we use $F = 1$ atoms the repump laser is first used to bring the atoms resonant with the imaging light (details in section 2.3.1). Since the atoms scatter photons out of the imaging beam the atoms are seen as a shadow in the imaging light. The imaging pulse needs to be short because it heats up the atoms and because the atoms move on a millisecond time scale. For these reasons we typically use 50 µs long imaging pulses with an intensity far below the saturation intensity of the transition.

Instead of looking at shadows in a bright image we introduce the Optical Density (OD) to describe the light absorption due to the atoms. The OD is defined by Beer's law [42]

$$I = I_0 e^{-\mathrm{OD}} \tag{2.1}$$

where I is the intensity hitting the CCD camera and I_0 is the intensity sent towards the atoms[1]. The OD (calculated for each pixel) is given by

$$\mathrm{OD} = \ln\left[\frac{I_{ref} - I_{bck}}{I_{raw} - I_{bck}}\right] \tag{2.2}$$

Here I_{raw} is a RAW image of both the atoms and the laser beam (corresponding to I) and I_{ref} is a reference image of the laser beam alone (corresponding to I_0). To account for the background light I_{bck} (with no light and no atoms) is introduced. If no atoms are seen then $I_{ref} = I_{raw}$ and $\mathrm{OD} = 0$. However if atoms are seen then $\mathrm{OD} > 0$.

The in-trap size of a BEC is typically only a few tens of micrometers in diameter. This means that the BEC will often be broadened by the experimental resolution of the imaging system. To circumvent this problem the trap is often turned off and the cloud is allowed to expand for a certain Time-Of-Flight (TOF) before the RAW image is grabbed. When all the images are grabbed the OD can be used to calculate the integrated density along the line of sight and parameters like the number of atoms and the temperature of the cloud can be found (this is described in chapter 5).

Time line

A time line of the experiment can be seen in figure 2.3. It may seem to be a relatively limited number of events we go through to make a BEC. However this is really only a rough overview. In reality most of these events involves many sub-events that all have been optimized carefully.

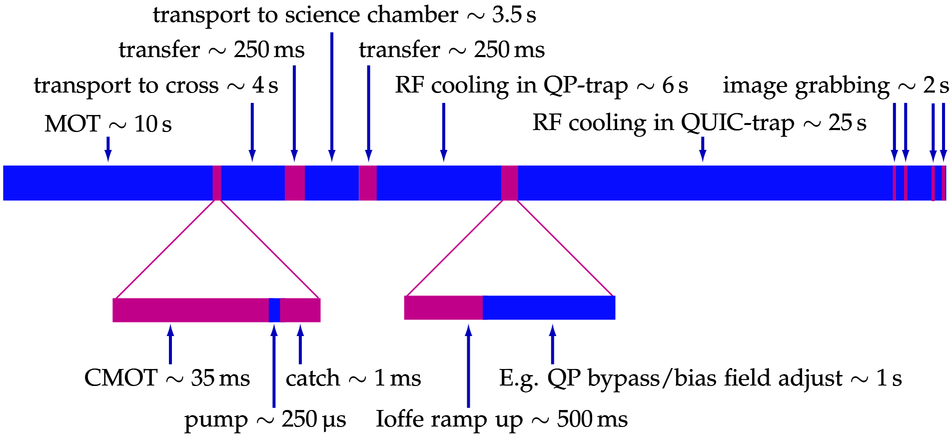

Figure 2.3: A time line showing the most important events in our experimental sequence to make a BEC (not to scale). This figure is meant as an overview of the different procedures described in this chapter. The whole sequence lasts for about 1 minute. After this the BEC is destroyed by the imaging beam and we run the sequence again.

This finishes the short overview of the experiment. In the following sections we will go into more details.

---

[1] Another often used parameter is the transmission T ($OD = -\ln(T)$).

2.3 Setup on the laser table

In this section the setup for producing all the light used in the MOT and later in the imaging sequence will be described.

2.3.1 The rubidium lasers

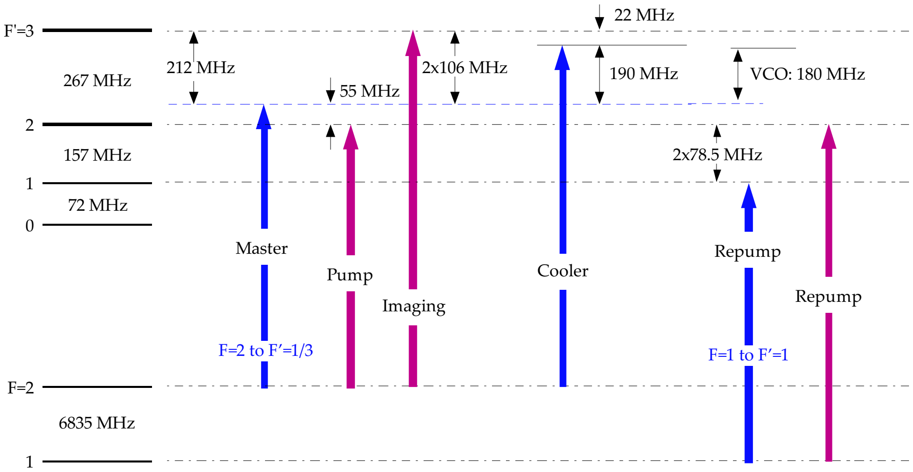

Figure 2.4: Rubidium D2 line (780 nm)—not to scale. The blue arrows indicate where the three lasers (Master, Cooler and Repumper) are locked. The red arrows indicate the different beams taken from the three lasers.

The three rubidium lasers are named Master, Cooler and Repumper

- **Master:** The light from the Master laser is used for three purposes: optical pumping (the state preparation in the last few milliseconds of the MOT-sequence), imaging and it provides the locking signal for the Cooler laser[2]. The Master diode laser is locked directly to the $F = 2$, $F' = 1/3$ crossover (see figure 2.4) using Doppler-free saturated absorption spectroscopy [46]. For optical pumping the light from the Master is shifted 55 MHz down by a single pass Acousto Optical Modulator (AOM) setup. This brings it on resonance with the $F = 2$ to $F' = 2$ transition. For the imaging light the Master is shifted 212 MHz up by a double-pass AOM setup (2×106 MHz). This provides an easy way to do off-resonant imaging since we can just change the AOM frequency.

- **Cooler:** The light from the Cooler laser is used to laser cool and trap the atoms in the MOT (described in section 2.4.3). The Cooler is offset-locked 190 MHz above the Master. This corresponds to 22 MHz below the $F = 2$ to $F' = 3$ transition (see figure 2.4).

---

[2] Actually it is also used for locking the length of a cavity, which again is used to lock the optical lattice lasers. This will be described in more details in chapter 7.

- **Repumper:** The light from the Repumper is used to bring atoms from the $F = 1$ state back to the $F = 2$ to $F' = 3$ cyclic transition. The Repumper is directly locked to the $F = 1$ to $F' = 1$ transition using absorption spectroscopy. It is then shifted $2 \times 78.5\,\mathrm{MHz}$ up to be resonant with the $F = 1$ to $F' = 2$ transition by a double-pass AOM setup.

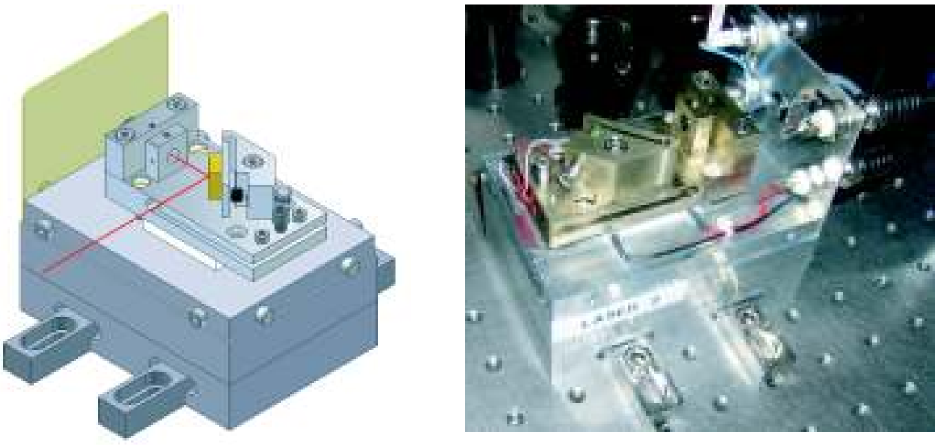

Figure 2.5: Our laser setup. All the mechanical parts and the electronics are home-made.

2.3.2 The diode laser setup

All our rubidium lasers are based on laser diodes [47]. Our home-made laser setup can be seen in figure 2.5. We use a standard Littrow configuration Extended Cavity Diode Laser (ECDL) setup [48, 49]. The parts are machined in an alloy[3] which is flexible and has a good thermal conductivity. Between the different parts in the setup we use a thin layer of heat compound[4] to ensure a good heat transfer. A thermistor is placed close to the diode laser and a peltier element transport heat from the setup to the aluminum base-block (or in the opposite direction). A temperature Proportional-Integral-Derivative (PID)-controller keeps the temperature of the setup constant within 1 mK. The laser diodes[5] are especially chosen to be close to 780 nm at room temperature. All the lasers provide about 50 mW of power and have a linewidth below 500 kHz.

The stability of these lasers are extremely good—they can survive severe treatment and they can be stored for half a year and still be resonant with the rubidium D2 line when they are turned on. Part of the reason why they are so stable is that we use uncoated laser diodes which means that we have both an internal (inside the diode) and an external cavity (formed between the diode and the grating). We have a few times used coated diodes, but we have always decided to replace them with uncoated ones.

A problem with our design (this problem only turns up when we are using coated diodes) is that the vertical screw that controls the vertical feedback has a tendency of

---

[3] $CuNi_7Zn_{39}Pb_3Mn_2$.

[4] AOS Thermal Compounds—HTC-60.

[5] Sanyo DL-7140-201W "non-coated" laser diodes (chosen from the DL-7140-201S series).

settle into the metal/alloy. We could fix this problem with a ceramic contact surface for the screw.

2.3.3 Locking methods

The Master and the Repump lasers are frequency locked to a saturated absorption signal using Frequency Modulation Spectroscopy (FMS) [50–52]. I will not go into details (these can be found in [50]), but just mention that we use a 10 MHz modulation signal that is sent to the current controller for the laser diodes and to the 10 MHz (saturated absorption spectroscopy) detector as a reference. We use a servo-amplifier to convert the error signal from the detector to a dc-signal (up to 30 Hz) for feedback to the piezo that controls the angle of the grating and an ac-signal (up to 100 kHz) for feedback to the current controller for the laser diodes. Initially we used only the feedback to the grating, however the lock becomes much better if we also use feedback to the current controller.

The Cooler laser is locked to the Master laser using an offset lock [53]. In short we have a fast (up to 300 MHz) detector that can measure the beat frequency between the Master and the Cooler laser (the beat frequency should be around 190 MHz—see figure 2.4). By mixing the beat signal from the Master laser with a signal from a Voltage-Controlled Oscillator (VCO), dividing the new signal into two, and introducing a phase delay on one of the signals before they are mixed again, a locking signal will be produced. This signal is sent to a servo-amplifier (identical to the one we use for the FMS-locks). Also in this case we use both dc- and ac-feedback. The offset between the Master and the Cooler can be adjusted just by changing the VCO frequency. A schematics showing how to build this kind of lock and the relevant formulas can be found in appendix A.

The Cooler laser could also be locked to an atomic resonance and then shifted in frequency using an AOM. We do not do this because we need as much power as possible in the cooling beams, and if we used an AOM we would probably lose at least 20% power.

Both locking methods are able to correct for a significant amount of noise and the lasers will typically stay locked all day long. The locking signals for all the lasers can be seen in appendix A.

2.3.4 Optical table

A schematics of the laser table setup can be seen in figure 2.6. The Cooler light is amplified by a Tapered Amplifier (TA) that provides 500 mW of optical power (it is seeded with approximately 15 mW). A TA is roughly speaking a laser diode with no cavity. This means that when light is sent through the TA we get an amplification of the seed frequencies due to stimulated emission. The TA setup will be described in more details in chapter 7 (we also use this setup for the optical lattices). The Cooler light is mixed with the repump light and sent into three MOT fibers (however the repump light is primarily sent into only one of the three MOT fibers—this is due to the polarization of the repump light compared to the polarization of the Cooler light).

The Master is used for either pump light or imaging light. Since we never use both pump and imaging light at the same time, we use an electronic flipper mounted with a half-wave plate, followed by a Polarizing Beamsplitter (PBS) to switch between having light from the Master going into the pump fiber or into the three imaging fibers. The same way we can switch the repump light between being used in the pump fiber or in the imaging fibers (the repump light in the imaging fibers is only used if we use $F = 1$ atoms). This arrangement is used to get as much power in each beam as possible, and it is not a problem in the runs, since we always have much time to switch the flippers.

To be able to control where all the light is going we have different shutters in the setup—e.g. we have a shutter in front of all three imaging fibers, and the AOM setups are typically also followed by a shutter so that we can turn off the light completely. For diagnostics all the lasers can easily be coupled into a Fabry-Perot interferometer which is also used in the lock-chain for the optical lattices (described in chapter 7).

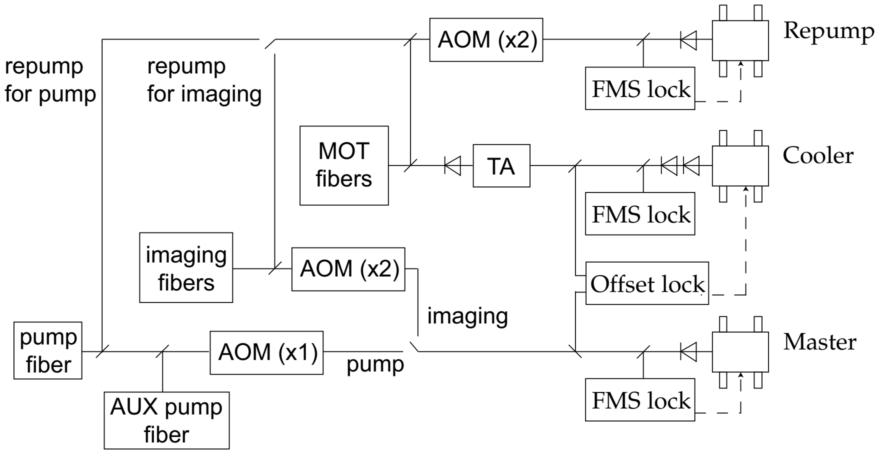

Figure 2.6: Schematics showing the rubidium part of the setup on the laser table. The Cooler is locked using the offset lock, but we also have a saturated absorption spectroscopy setup (FMS-lock) on the Cooler (mostly used for diagnostics). The different AOMs are used to shift the frequencies of the light in accordance with the description in section 2.3.1 (x1 means single pass, x2 means double-pass). Two places in the setup we can switch between using the light for pumping and for imaging. More information is given in the text.

2.4 Setup on the vacuum table

In this section the setup on the vacuum table is described and at the same time the different experimental procedures are presented and discussed in some details. The different parts of the setup on the vacuum table are introduced in the order that they are used in the experiment sequence.

2.4.1 Vacuum chamber

The vacuum chamber is made of stainless steel except for the MOT-chamber, which is a custom-made cylindrical glass-cell. After the vacuum chamber was assembled two turbo pumps were used to pump the vacuum chamber down to 1×10^{-8} torr. After this two ion-pumps (one placed in each part of the chamber) were turned on[6]. The vacuum chamber (both parts) was then baked out at up to 190 °C for two weeks. After this the turbo-pumps were turned off and the valves in front of them were closed. Since then we have not opened the vacuum chamber[7].

---

[6] The gauge on the ion-pump in the high pressure part of the science chamber has risen a factor of 10 in 2008. We cannot see any changes in the MOT or in the general behavior of the experiment so we think it is just an effect of the ion-pump (and not really there), but it is something that should be checked when the vacuum chamber has to be opened.

[7] During the baking it was very important to have water cooling (and not just air cooling) on the turbo-pumps—otherwise they will get too hot and shut down (we tried this—this was a smaller disaster).

Besides the ion-pumps we have two Titanium Sublimation Pumps (TSPs) in the good vacuum part. We run them perhaps ones every half a year when the lifetime in the science chamber becomes too short (the normal lifetime is about 150 s).

2.4.2 Dispensers

We have four rubidium dispensers in the vacuum chamber. If necessary, we can easily change to a new dispenser just by moving two wires connected to pins on the outside of the vacuum-chamber. Right now, we are running on the fourth dispenser, however the others are perhaps not completely empty—we change when the needed current (and the resulting heating) gets too high. The dispensers last longer and longer, probably because the vacuum chamber has now been "saturated" with rubidium atoms. The previous dispenser lasted for more than one year. When all four dispensers are empty we need to open the vacuum chamber and replace them. However we can close a valve between the MOT-chamber and the science chamber, so we only need to bake out the MOT-chamber. The dispensers are placed in such a way that the atoms have to pass a corner before they can reach the glass-cell. When the MOT-chamber has to be opened, we might place the new dispensers in such a way that they point directly towards the glass-cell to try to minimize the consumption of rubidium. Also, today it is possible to buy really big rubidium dispensers. Perhaps these initiatives could make it possible to run for many years without having to open the vacuum chamber.

We run 3.5 A–5.5 A through the dispenser depending on the amount of rubidium left in the dispenser and on whether we are starting up or have been running for a while. We have a 6 A fuse on the wires to the dispenser to avoid a potential malfunctioning current supply from coating the vacuum chamber with rubidium. In the beginning of the day we run the dispenser for 5–10 minutes at a high current (e.g. 4.5 A) and then we turn down the current to a level that will keep the loading time constant throughout the day (e.g. 3.5 A). If the loading time of the MOT is different from approximately 10 s we adjust the current by 0.1 A (up or down). This way the loading time will stay constant just by slightly adjusting the current to the dispenser a couple of times during a day.

2.4.3 MOT

Our 3D MOT setup consists of counter propagating laser beams along three mutually orthogonal directions (6 beams in total) and a magnetic QP-field produced by the MOT-coils. The counter propagating laser beams have opposite circular polarizations. More information about how a MOT works can be found in e.g. [41, 50, 54–57].

In the following the reason for locking the lasers as described in section 2.3.1 will become clear. The cooling effect of a MOT is analogous to normal laser (Doppler) cooling. We cool on the $5\mathrm{S}_{1/2}$ $F = 2$ to $5\mathrm{P}_{3/2}$ $F' = 3$ cyclic transition giving us an effective two level system (since selection rules demands $\Delta F = 0, \pm 1$, and disregarding m_F sub-states). We have optimized the red-detuning of the Cooler laser to 22 MHz below the F=2 to F'=3 transition, which corresponds to 3.7 natural linewidths (3.7 Γ). If an atom ends up in the $F = 1$ state (this happens in about 1‰ of the excitations) it needs to be put back onto the cyclic transition. This is done with the repump beam (tuned to the $F = 1$ to $F' = 2$ transition). If the atom afterward falls back to the $F = 2$ state it becomes resonant with the Cooler laser again—if not it gets one more trip.

The MOT beams are about 20 mm in diameter and we have about 25 mW of power in each beam. Since we have a cylindrical glass-cell we use telescopes for all the beams instead of just having retro-reflectors. The laser beams are disturbed through the curved glass-cell, so we would have troubles getting the same amount of light back through the glass-cell (with a reasonable beam profile) if we used retro-reflectors. Also, since we are using telescopes we have almost no MOT optics in "free air" on the vacuum table. The MOT light is coupled into three polarization maintaining fiber splitters[8]. Every fiber has two output fiber couplers that can be placed directly in the telescopes. The splitters are specified to split 50/50 at 780 nm. This means that even though the input intensity is drifting the MOT beams are always balanced within 1%.

The number of atoms we load in the MOT is not very critical for getting a BEC in the end of a run sequence. We can load only 60% of the atoms, we normally load in the MOT, and still get a BEC with our normal settings, so we are far from the limit (actually we can also load twice the normal amount of atoms if we push the system). The final number of atoms in the BEC will depend on the number of atoms collected in the MOT, but this dependence is not linear—we do not get twice the amount of atoms in the BEC if we load twice the amount of atoms in the MOT[9].

2.4.4 CMOT

The CMOT is obtained by increasing the detuning of the Cooler light from about 3.7 to 6.5 natural linewidths below the resonance. At the same time the repump power is reduced. Together this reduces the radiation pressure in the trap and the spatial extent of the MOT cloud becomes smaller. This results in a higher phase-space-density of the cloud [42].

The obvious way to detune the Cooler laser would be to change the VCO-frequency of the offset lock and thereby change the offset to the Master. However the servo-amplifier is not fast enough to keep the laser locked in this procedure. Instead we have a module that unlocks the Cooler laser, change the injected current to the diode laser slightly (and thereby also the frequency) and afterwards it takes the frequency back and locks the laser again.

2.4.5 State preparation

In the end of the MOT-sequence we either put the atoms in the $F = 1$, $m_F = -1$ state or in the $F = 2$, $m_F = +2$ state. These two states can both be trapped in a magnetic field minimum, which we will see in section 2.4.6.

To get the atoms in the $F = 1$, $m_F = -1$ sub-state we just turn off the repump laser $\sim 500\,\mu\text{s}$ before the cooling light is turned off and the atoms are caught in the magnetic field. If an atom ends up in the $F = 1$ state it will stay here since there is no repump laser to put it back into the cyclic $F = 2$ to $F' = 3$ transition. When the current in the MOT-coils is cranked up (almost) all the atoms are in the $F = 1$ state. About 1/3 (actually we measure closer to 40%) of these atoms will be in the $m_F = -1$ sub-state and therefore trapped in the magnetic field.

If we want to get the atoms into the $F = 2$, $m_F = +2$ sub-state we use optical pumping. A constant magnetic field of about 1 G, provided by two extra coils mounted outside the MOT-coils in a Helmholtz configuration, defines a quantization axis and σ^+ pump-light (coming from the vertical, upper telescope) on the $F = 2$ to $F' = 2$ transition is turned on. This puts almost all of the atoms in the $F = 2$, $m_F = +2$ sub-state. During this procedure the repump light is still on. All the experiments presented in this thesis is done with $F = 2$, $m_F = +2$ atoms.

---

[8] From Canadian Instrumentation and Research.

[9] This might be due to 3-body collisions—see e.g. [45] and the comments in chapter 5.

2.4.6 The magnetic QP-trap

The magnetic field of a QP-trap has a zero-point at the center of the trap. Around this zero-point ($x = y = z = 0$) the magnetic field B increases approximately linearly in all directions. The gradient of the magnetic field along the axis of the coils (here defined to be parallel to the z-axis) is twice the gradient in the radial direction $r = \sqrt{x^2 + y^2}$ [56]

$$B \approx A\sqrt{r^2 + (2z)^2}, \qquad A = \left.\frac{\partial B}{\partial r}\right|_{r=0^+}. \tag{2.3}$$

For our coils this is a good approximation out to at least 1 cm from the trap center. The atoms move in a potential V that depends on the magnitude of the magnetic field B

$$V(x,y,z) = m_F g_F \mu_B |B(x,y,z)| \tag{2.4}$$

here g_F is the g-factor, m_F is the magnetic quantum number and μ_B is the Bohr magneton. Atoms can be trapped in this potential if B is large and if the temperature of the atoms is very low.

It can be proved that in a region without charges and currents a magnetic field can only have a local minimum (like in eq. (2.3)) but not a local maximum [58]. This means that the sub-states where $g_\mathrm{F} \cdot m_\mathrm{F} < 0$ cannot be trapped—only week field seeking states where $g_\mathrm{F} \cdot m_\mathrm{F} > 0$ can. So eq. (2.4) explains why the atoms in the following sub-states $|F = 2, m_F = +1\rangle$ ($g_\mathrm{F} = 1/2$), $|F = 2, m_F = +2\rangle$ ($g_\mathrm{F} = 1/2$) and $|F = 1, m_F = -1\rangle$ ($g_\mathrm{F} = -1/2$) can be trapped and why they are trapped in the B-field minimum[10].

2.4.7 Catch

In the end of the MOT-phase we turn off all the light, and at the same time we need to turn on the magnetic QP trapping potential (introduced in the previous subsection) fast to catch the atoms in the field of the MOT-coils. If this turn-on time is too slow, the atoms will have time to move away from the center of the trap before it is loaded (resulting in fewer atoms loaded into the trap at a higher temperature). The current controller for the MOT-coils has a feature that makes it possible to charge a big capacitor and "fire" it when we want to turn on the current fast (e.g. during catch). A transistor controls the current until the power supply can take over. This setup makes it possible to ramp up the current to e.g. 250 A in around 100 µs. Initially we used this feature, however our power supply[11] can ramp up in a few milliseconds, and when we compare the two different loading techniques, we see no significant difference in the atom number or in the temperature of the clouds, indicating that a turn-on time of a few milliseconds is fast enough. For this reason we normally just use the power supply directly to control the current in the MOT-coils during catch.

2.4.8 Transport and transfer

Even though it is possible to make a BEC in a purely optical trap in just a few seconds [59, 60] (the same as the lifetime in our MOT-chamber), it is generally more common to move the atomic cloud from a MOT chamber to a low pressure science chamber, where the lifetime makes it possible to make a BEC using RF-evaporative cooling in a magnetic trap.

---

[10] Actually $|F = 2, m_F = 0\rangle$ ($g_\mathrm{F} = 1/2$) can also be trapped.

[11] Delta Elektronika SM15-400.

Different schemes can be chosen for moving the atoms into the low pressure part. Setups with two MOTs are widespread [61]. Here a 2D MOT is placed in the high pressure part and from this MOT the atoms are transfered/loaded through a differential pumping hole into a 3D MOT located in the low pressure part of the vacuum chamber. The transfer between the two MOTs can be done using a push beam.

Another approach is to catch the atoms in a purely magnetic trap and then move this trap and thereby the atoms. This can be done by making a chain of overlapping QP-coils and running suitable currents through them—in this way the trapping geometry of the potential can be maintained during the transport process [62].

An alternative approach is just to move the magnetic coils mechanically [42]. We use this method because it is an easier and more elegant solution in many aspects. E.g. we save a lot of space since we do not need to have extra MOT beams or coils mounted all the way from the MOT-chamber to the science chamber (which also needs a lot of control electronics). This gives us much more space for other things around the science chamber. The total distance between the rubidium MOT-chamber and the science chamber is about 90 cm—but we could easily move the atoms over a larger distance.

The transfer of atoms from the MOT-coils to the conveyor coils and later from the conveyor coils to the static QP-coils mounted on the science chamber is not very critical. As an example we overlap the MOT/conveyor coils with 250 A running in the MOT-coils and 0 A in the conveyor coils. The transfer can afterwards be done almost adiabatically by two linear ramps—the MOT current is ramped from 250 A to 0 A in 250 ms, and at the same time the conveyor current is ramped from 0 A to 250 A. A similar transfer is done from the conveyor coils to the stationary QP-coils mounted on the science chamber.

2.4.9 Science chamber

The custom-made science chamber can be seen in figure 2.7. The science chamber is tilted 55 degrees with respect to the vacuum tube that connects the science chamber to the rest of the vacuum chamber. This gives us three mutually perpendicular directions through the science chamber that can be used for the optical lattices. The sizes of the windows along these three axes through the science chamber are chosen in such a way that we can get a good imaging resolution on all the axes. On top of the three main axes we have two auxiliary axes (four windows) diagonally through the science chamber. Currently we have only been using the three main axes, but the auxiliary axes could e.g. be used for extra laser beams. All the windows are anti-reflection coated, which especially is a nice feature considering that we are using optical lattices. If we did not have anti-reflection coated windows, we would most likely have to avoid using the lattice beams at normal incidence (in order to prevent reflections from destroying the standing wave pattern).

QUIC-trap

The magnetic field produced by the QUIC-coils can be calculated using the Biot-Savart law [63]. This is necessary to do in the design process, but these calculations take a lot of computation time and often they depend critically on the geometrical parameters. When the trap is in use one can instead use an analytical approximation to describe the magnetic field from the QUIC-trap [1]

$$\mathbf{B} = B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + B' \begin{pmatrix} x \\ -y \\ 0 \end{pmatrix} + \frac{B''}{2} \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{1}{2}(x^2 + y^2) \end{pmatrix}. \tag{2.5}$$

A configuration that gives this field is often referred to as a Ioffe-Pritchard trap (and the QUIC-trap is one of them—different variants of this trap can be found in [64]). This kind

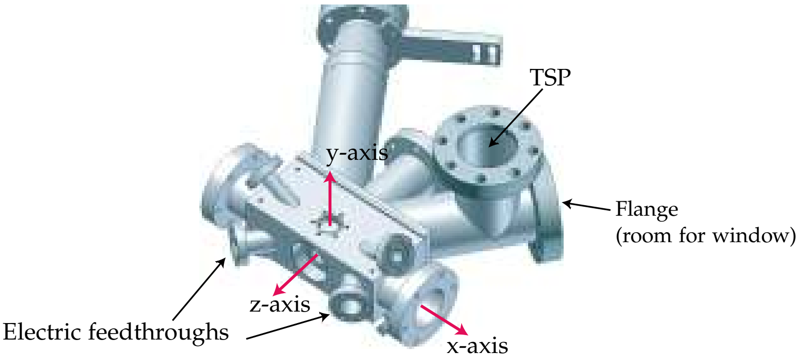

(a) A drawing of the science chamber. The science chamber is made of stainless steel and have three main axes and two auxiliary axes.

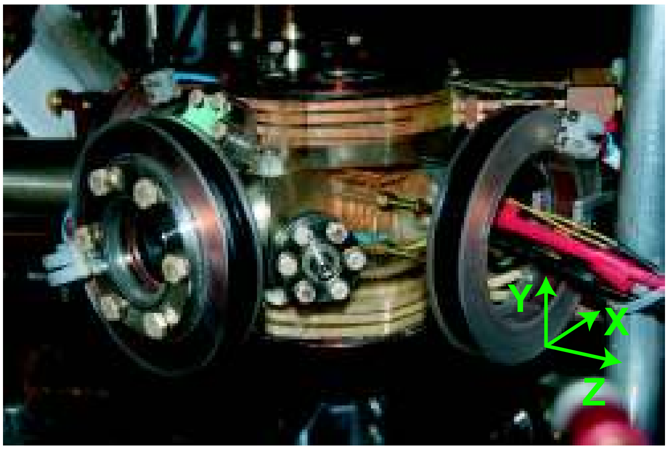

(b) A picture showing the science chamber—here the different magnetic coils can also be seen. Slightly to the right of the center of the picture is the Ioffe-coil. Above and below this coil the big QP-coils can be seen. The other coils (further away from the center of the science chamber) is the shim-coils. The y-shim coils are very thin and glued onto the QP-coils. Characteristic coil parameters for all our coils are listed in the end of this section.

Figure 2.7: The science chamber. Notice that the viewpoint is different in the two figures.

of trap has the wanted bias field B_0, which is the reason why the Ioffe-coil is used. The different constants B_0, B' and B'' can be measured for a given trap of this type—we will return to this in chapter 5. The potential obtained with this magnetic field is found by inserting eq. (2.5) into eq. (2.4).

The QUIC-coils are mounted firmly with screws to the science chamber (see figure 2.7). It is very important that these coils cannot move at all, since if they move just a few micrometers it will change the magnetic trap significantly. The Ioffe-coil was initially designed in such a way that the windings were placed in an aluminum/plastic housing. Since water was pumped through the housing the wires moved slightly from run to run. This resulted in severe problems in the evaporative cooling sequence, since the trap bottom would change all the time. It turned out that a 5 µm movement of the Ioffe-coil could explain the few tens of kHz drift in the trap bottom we observed. For this reason a new Ioffe-coil was designed and tested. After a couple of retries we got a working Ioffe-coil made in the same way as the other coils (casted into epoxy—more about this in section 2.4.10). The demands to the stability of the QUIC-trap also explains why we cannot use the moving conveyor coils as a part of the QUIC-trap, since the position of these coils changes a lot (up to 50 µm) from run to run.

Along all three main axes we have shim coils mounted on the science chamber (see figure 2.7b). These coils can, beside being used to shim for different magnetic fields (e.g. the earth magnetic field), also be used to shape the magnetic trap in the science chamber—this is e.g. done before the optical lattice is loaded. We can also easily bypass the different shim coils or change from a Helmholtz to an anti-Helmholtz configuration for the different coil-pairs. For this reason we can e.g. use the shim-coils for Stern-Gerlach experiments, where an inhomogeneous magnetic field is used to separate the different m_F components (this can also be done with a current pulse on the Ioffe-coil).

The conveyor coils can also be used to provide (or shape) a magnetic field in the science chamber. With the conveyor coils we can change from an anti-Helmholtz to a Helmholtz configuration electronically during a run. This functionality was implemented to be able to set a high, homogeneous magnetic field in the science chamber, which e.g. can be used to sweep across Feshbach resonances (however so far we have not done any experiments with Feshbach resonances).

QUIC controller

The QUIC controller (see figure 2.8) is designed in such a way that we can start out with only the QP-trap on, and then slowly turn on the Ioffe coil and load the atoms into the QUIC-trap. Also, we can bypass current around the QP-coils and thereby adjust the current ratio between the QP-coils and the Ioffe-coil. With the new (and stronger) Ioffe-coil we are able to place it (with micrometers precision) in such a way that we can run the same current in all the QUIC-coils (e.g. 300 A) and get a bias field of around 700 kHz. However, currently we are using the z-shim coils together with the QUIC-trap (with 300 A in all three coils) to get the right trap minimum (around 700 kHz or $B_0 \sim 1\,\mathrm{G}$). The only drawback of doing this is that the turn-off time of the current in the shim coils is around 5 ms (this could of course be improved with some control electronics). This is significantly slower than the turn off of the QUIC-trap, which is only about 100 µs.

RF-evaporative cooling

Let us look at the condition that needs to be fulfilled in order for the RF-evaporative cooling to transfers atoms from a trapped state in the QUIC-trap into an un-trapped state. When the RF-frequency ν_{RF} (or more precisely the energy $h\nu_{RF}$) is equal to the energy difference between two subsequent m_F numbers in the magnetic trap the transfer

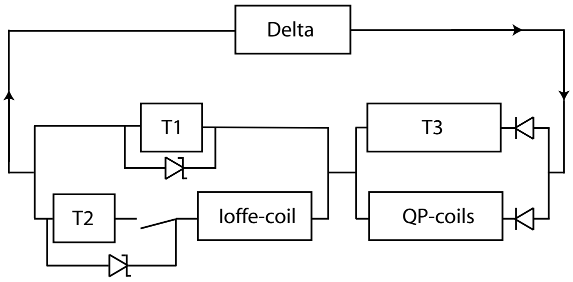

Figure 2.8: A rough schematic showing the principle of operation of the QUIC-controller. "Delta" refers to our power supply (a Delta Elektronika SM15-400), and T1, T2 and T3 refer to transistors capable of controlling currents up to 400 A. Initially, all the current runs through the QP-coils and through T1. When we want to turn on the Ioffe-coil, we "open up" T2 and "close down" T1 in order to get all the current through the Ioffe branch. T3 (the QP-bypass transistor) can be used to bypass *some* of the current around the QP-coils, which effectively makes the Ioffe-coil stronger. The QUIC-controller adjusts the current in the different branches actively. Initially, the different currents were measured with resistors with a really small temperature coefficient (e.g. Zeranin resistors), but we have replaced most of them with current transducers (Danfysik Ultrastab 867-400I), which are much easier to implement technically and also much more precise.

will happen (see e.g. eq. (2.4))

$$|B(x,y,z)| = \frac{\hbar\omega_{RF}}{g_F\mu_B} \approx \frac{\nu_{RF}}{g_F \cdot 1.4\,\mathrm{MHz/G}} \tag{2.6}$$

here g_F is the g-factor, μ_B is the Bohr magneton and $\omega_{RF} = 2\pi\nu_{RF}$ is the angular RF-frequency. This condition will be fulfilled at a certain "height" (or "shell") in the QUIC-trap. When the RF-frequency is swept down the resonance shell becomes smaller and smaller. To specify the location of this shell for an atom in state m_F, we will define a truncation energy

$$\epsilon_t = m_F\hbar(\omega_{RF} - \omega_0), \tag{2.7}$$

where $\omega_0 = \mu_B g_F |B(0,0,0)|/\hbar$ is the resonance ("bias") frequency is the center of the trap. We are going to use these results in chapter 5.

The RF-signal for the evaporative cooling is generated using a Direct Digital Synthesis (DDS) (10 kHz to 70 MHz) and a 75 W RF-amplifier[12] is used to amplify the signal from the DDS (details on how we program the DDS is given in chapter 3). The signal from the amplifier is sent to the RF-antenna inside the science chamber[13].

2.4.10 The magnetic coils

The MOT, conveyor and QUIC-coils are made of squared copper-wires with a 2 mm hole in the middle for direct water cooling. A thin insulation material was twisted around

---

[12] Amplifier Research 75A250—this amplifier needs good cooling (cold air circulation) otherwise it will malfunction.

[13] We have a spare RF-antenna placed inside the science chamber if the first coil is destroyed (has not happened yet).

the wires before they were shaped and casted into epoxy under vacuum[14]. Our power supplies and electronics are specified for 400 A and the cooling is dimensioned for this. We normally use only up to 300 A and at this current we can run continuously without having to worry about the coils getting to hot. However some of the control electronics will be very hot after some minutes.

A 10 bar pump is used to pump the water through the small hole in the wires. We use demineralized water in a closed system and this water is cooled by the cooling water in the house (through a heat exchanger). We were afraid that all the parallel connected coils would result in different flows in the different tubes (and wires). For this reason a fancy system was designed where it is possible to measure and adjust the water flow individually in all the different tubes. However, it turned out that the water flow in all the different tubes was approximately the same, but it might come in handy later on if we get new equipment that also needs water cooling.

A list of the different coils can be seen in table 2.1. The radial field gradient of the MOT trap is 0.36 (G/cm)/A and for the conveyor trap it is 0.20 (G/cm)/A [41]. The characteristic parameters of the QUIC-trap will be given in chapter 5.

Name	Windings	Average radius (mm)
MOT coils	13	32
MOT pump coils	25	45
Conveyor coils	31	42
QP-coils	24	36
Ioffe-coil	10	13
x-shim coils	208	45
y-shim coils		35
z-shim coils	208	45

Table 2.1: The different magnetic coils used in the experiment. Precise measurements and simulations of the magnetic field from the MOT and conveyor coils can be found in Jesper's PhD thesis [41]. In chapter 5 precise measurements of the parameters describing the magnetic field of the QUIC-trap can be found.

2.5 Two species

Our setup supports two species since a second species part can be added opposite to the rubidium MOT (see e.g. figure 2.1)—this part will use another set of magnetic coils and a third translation stage. The setup will be able to bring second species atoms to the cross, where they can be mixed with rubidium atoms, before the cloud is transfered to the science chamber using the conveyor coils.

Our setup stands out with this transport sequence that makes it possible to mix two species (prepared in two different MOT-chambers) at the cross or in the science chamber. Jesper has done a lot of mixing of two rubidium clouds trapped in the MOT-coils and in the conveyor coils, respectively [41], [I]. These studies showed that it was possible to mix two rubidium clouds virtually without losing atoms (so we would end up with twice the normal amount of atoms) and with an acceptable heating of the atomic cloud(s). Jesper also simulated mixtures between rubidium and lithium and also in this case it looked

---

[14] Loctite Hysol (Resin epoxy RE2038 and Hardener HD3475) was used. The vacuum was pumped down to 1×10^{-5} mbar before the epoxy was injected/sucked into the vacuum chamber. This caused a pressure rise to 1 mbar and after this the pressure was actively varied around 1 mbar to remove air bubbles from the casting.

promising indicating that this "separated MOT-setup" in fact is an alternative to the current two species MOTs (e.g. [65, 66], or three [67]—also in the group of R. Grimm/F. Schreck a quantum gas experiment with a three species MOT has been built), where all the species are loaded in the same MOT. The nice thing about our setup is that the MOTs can be optimized independently and there is no need for e.g. dichromatic optics. However the experimental parameters (mixing velocity and currents in the different coils) have to be chosen carefully for the mixing to work, so it is not trivial but simulations (that have been checked against experimental measurements) can tell us how to do it. Everything we have designed and built so far have been prepared for this second species. If needed one could in principle also think about extending the setup in such a way that it would be possible to add a third species.

2.6 Final notes

In this section I will just make a few concluding comments to the BEC apparatus. Pictures of the setup on the laser and the vacuum table can be seen in figure 2.9. The BEC experiment has been running for the last 1½ year without any serious problems (it took us half a year to get everything stable after we got the first BEC). Most of the days we have a BEC in the first run and the entire setup can be turned on in 15–30 minutes. There is no (measurable) "heating" of the system so we do not need to run the experiment over and over again to make it stable—we can take data from the first run and the experiment can easily be controlled by one person.

The entire experimental setup has been built in such a way that everything can be controlled by analog (0 to 5 volts) or digital (0 or 5 volts) channels. This has been easy to achieve since almost all of the electronics in our laboratory is home-built by the electronics workshop in collaboration with us. The analog and digital channels are connected to a computer, which is able to control all the equipment in the laboratory with a precision of 0.1 µs. This kind of computer system is very important in a BEC laboratory because, as seen in this chapter, there are many different techniques involved in a quantum gas experiment. The computer systems we use to control our laboratory are described in the following chapters.

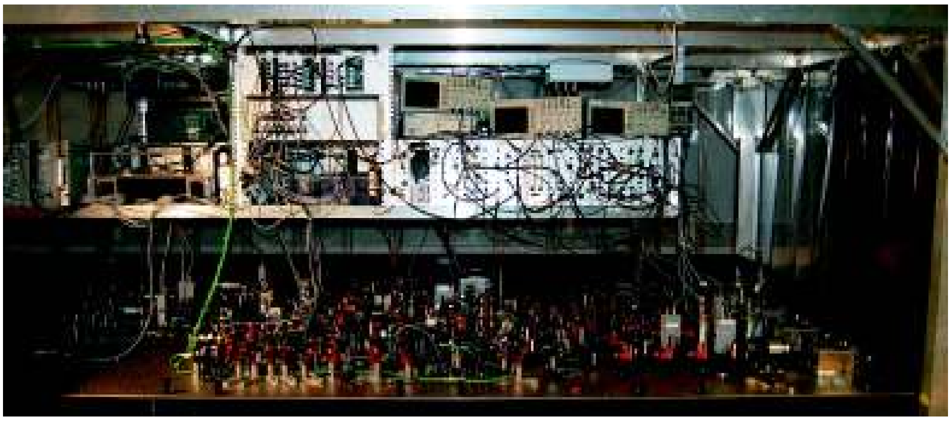

(a) Laser table setup.

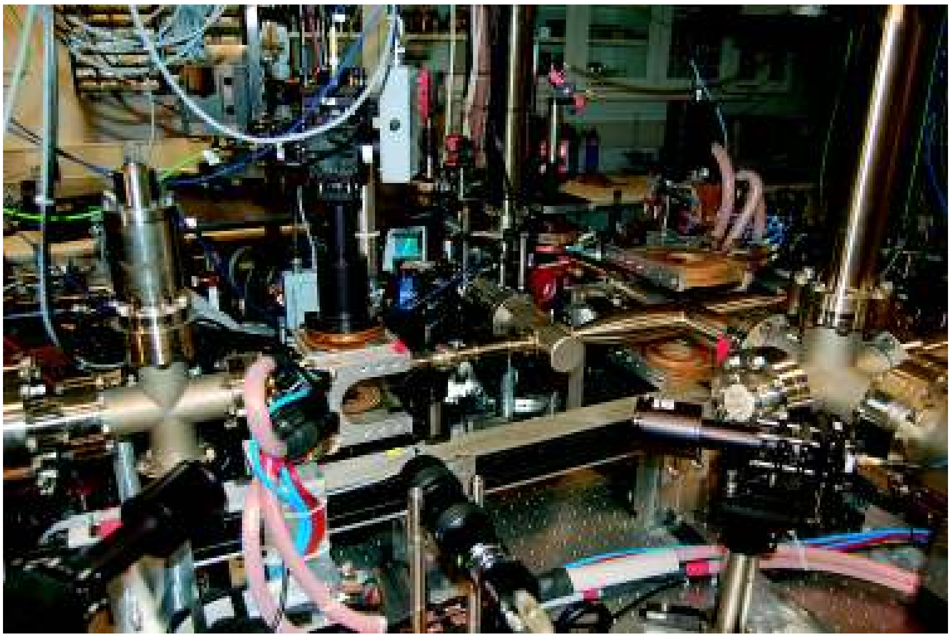

(b) Vacuum table setup.

Figure 2.9: Pictures of the experimental setup on the laser table and the vacuum table.

Three

Experimental control system

3.1 Introduction

All the different experimental steps presented in the previous chapter is controlled by two computer systems. The Experimental Control System (ECS) (described in this chapter) is able to control all the equipment in the laboratory and it makes sure that everything is synchronized. ECS is the heart of the experiment—everything of importance can be controlled from this system. At the same time a Camera Control System (CCS) (described in chapter 4) is needed to be able to grab the necessary images at the end of a run. With these two systems our daily work in the laboratory takes place in front of the computers almost all of the time. For the data analysis Michael has developed a fit-program which is able to analyze the grabbed images. More information about the data analysis can be found in chapter 5.

ECS and CCS consist of partly hardware and partly software and they are very extensive so the aim in this and the following chapter is not to provide a complete, detailed description of all the different parts of the systems, but instead to provide an overview of the most important parts and features.

I will start this chapter with a few general comments on ECS. After this I will go through the hardware part of ECS in section 3.3 and the software part in section 3.4. The Mechanical Positioning System (MPS) will be described in section 3.5 and comments on a big, future update of ECS will be given in section 3.6. This update is going to replace the only commercial product in the system with our own.

Besides putting up all the hardware and writing all the software described in this chapter I have spent a lot of time testing a lot of equipment in the laboratory and getting everything stable in order to make nice, reproduceable BECs with the system.

3.2 Background

Since this kind of control system is really an essential part of all quantum gas experiments, we decided from the beginning that we would not compromise on this system. Having a system that is not properly dimensioned can be a limiting factor or at least annoying and time consuming in the everyday work in the laboratory. Our focus on this area has already paid off and we are sure that ECS (and CCS) will form the base of the experiment for many years to come.

All of the software is written in Borland Delphi 7 professional and it is all written from scratch. Delphi is an object oriented programming language that in many aspects resemble C++. Almost all of the drivers provided by the companies behind the different equipment where written in C or C++ which meant, that the headers for these dll-drivers had to be translated from C to Delphi. We could also have chosen to write everything in C++ instead of Delphi, but since I was most experienced in Delphi we choose this programming language.

Focus has been on making a user-friendly, highly-flexible system with many advanced features. ECS is designed in such a way that it could also be used to run other kinds of experiments where a good time resolution is needed and where many things need to be computer controlled. For this reason we have made a few attempts to try to find a company that was interested in making the system into a commercial product. However one of the problems with this idea (besides that the market perhaps is not big enough) was that we did not have time to help in the development process. Different groups have also been very interested in buying ECS and CCS, but again it stranded because we did not have time to help them implementing the system(s). Perhaps sometime in the future the systems will be adopted in other laboratories—meanwhile they are continuously updated and new features are added.

3.3 ECS hardware

An overview over the whole system is given in figure 3.1. Even though the schematics in this figure looks quite complicated the strong feature about this system is that it relies on relatively simple ideas. The ECS computer takes care of all the timing and provides Arbitrary Waveform Generator (AWG) channels (in the following called analog channels) and digital channels that can be used to control equipment in the laboratory.

3.3.1 DIO64 cards and channel setup

ECS is based on two DIO64 (Digital Input/Output) PCI boards (each with 64 channels) placed in the ECS computer. These boards generate the output that controls the different equipment in the laboratory. ECS consists of software to control the two DIO64 boards and a variety of electronic modules to process the output generated by these boards. The electronic modules are the connecting link between the computer and the equipment that we want to control. In figure 3.2 the hierarchy of ECS is shown.

The DIO64 board is a commercial product manufactured by *ViewPoint Systems*, and it is the only part of the system which is not home-made—and actually this is going to change in the future since these boards will be replaced with our own Field-programmable Gate Array (FPGA) based circuit boards developed at the University. However, the basic principles of the new cards are the same as for the DIO64 cards (in fact the DIO64 cards are also based on FPGAs), so here it is described how the DIO64 cards work. In section 3.6 it is described how we can improve the functionality of the system significantly by using our own cards.

We clock the DIO64 cards at 10 MHz which give us a precision of 0.1 µs in the timing of the different events in the laboratory. To ensure that the two DIO64 cards are exactly synchronized the cards are linked together inside the computer in such a way that both cards use the same clock. The 64 digital channels in each card can be configured as either inputs or outputs (in blocks of 16 channels). We have configured the two cards in the following way:

- The first card is programmed to have 48 digital output channels and 16 input channels. These channels provides 48 digital output channels and 16 digital input channels for ECS.

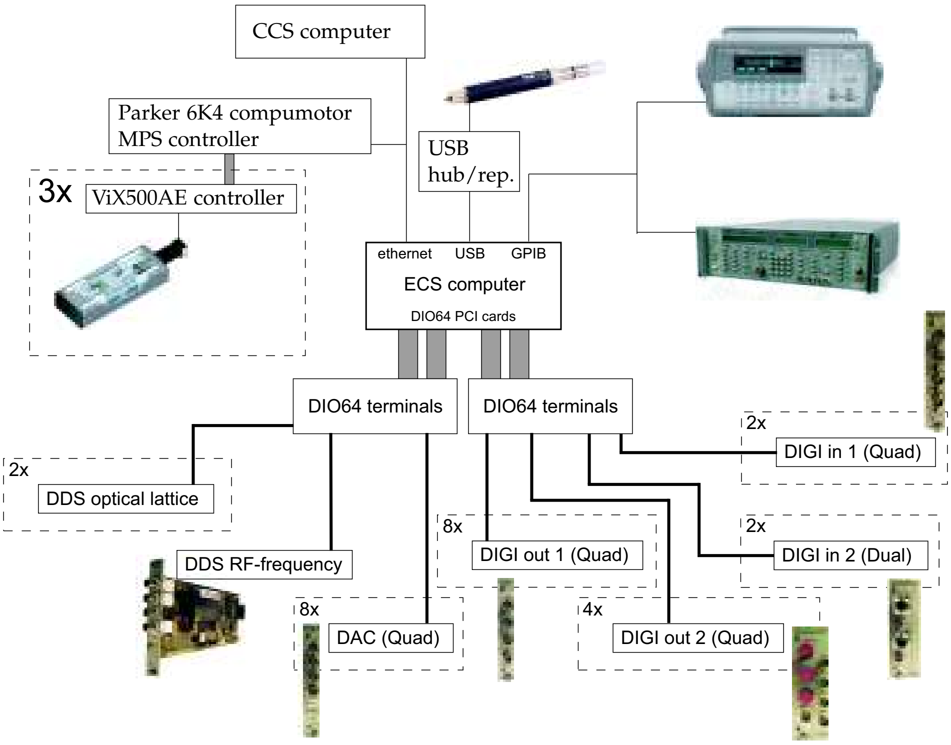

Figure 3.1: A schematic overview of ECS. ECS uses two DIO64 cards for all the analog and digital channels. Besides this MPS is connected to ECS together with a function generator, a microwave generator and a motorized translation stage (used to move an imaging lens when we grab images). These instruments cannot be controlled with digital and analog channels and therefore they are directly programmed by ECS. Further comments to this figure will be given in the text.

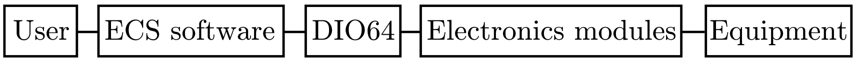

Figure 3.2: The ECS hierarchy.

- The second card is programmed to have 64 digital output channels. These channels provides 32 analog channels for ECS.

So, the second card generates the analog channels. However since the DIO64 card can only create digital pulses we need to be able to translate a series of digital pulses into a voltage. This is done by Digital-to-Analog converter (DAC) chips placed inside the electronic modules that provide the analog channels. Normally the analog channels are used to set a 0–5 V voltage, but they can also be used to control Direct Digital Synthesis (DDS) based electronics[1]. A DDS-chip generates an oscillating signal—a sine wave with a certain frequency. It can (like the DAC-chips) be programmed by a series of digital pulses. We use these chips to generate the RF-frequencies used for the evaporative cooling sequence and for generating the frequencies for two of the AOMs used for the optical lattices.

Channel type	Time resolution	Precision	Number of channels
Digital output (DIGIs)	1 µs	-	48
Digital Input	-	-	16
Analog (DACs)	down to 50 µs	16 bit	32
DDS (DDSs)	down to 50 µs	32 bit	32

Table 3.1: The different ECS channels. The analog and DDS channels have to share the same 32 channels. This means that since we at this moment use 3 DDS channels we have only 29 analog (DAC) channels available.

A great feature about the DIO64 cards is that only *changes* in the status of the channels have to be specified—however when the output on one of the 64 channels needs to be changed the output on all the channels has to be specified again. In table 3.2 an example of the DIO64 data format is given. The numbers in this table are all presented as a four digit hexadecimal number which means that all numbers between 0x0000 and 0xFFFF (16 bits) can be placed in a given cell (0x indicates hexadecimal numbers). The two first columns (32 bits) are reserved for a *time stamp* and the next four columns are reserved for output channel 1 to 64. The time stamp (that is the same as the "clock cycle number") that corresponds to a specific time, will depend on the clock frequency (in this case 10 MHz). For example, if we want channel number 19 to go high after 1 second the time stamp should be $1\,\mathrm{s}/0.1\,\mu\mathrm{s} = 10^7$. However this number has to be converted into two 16 bit numbers

$$(10^7 \bmod 65536) = 38528 \qquad \text{(or 0x9680)},$$
$$(10^7 \operatorname{div} 65536) = 152 \qquad \text{(or 0x0098)}.$$

Setting channel 19 high is done by writing 0x0004 in column 17-32 since

$$2^{((19-1) \bmod 16)} = 4 \qquad \text{(or 0x0004)}.$$

The last two rows in table 3.2 are made the same way. After 1 s, 2 ms and 500 µs channel 15 and 25 are set high and channel 19 is still kept high (row number 3). 1 µs later all channels are pulled low except channel 64 which is set high (row number 4). In

---

[1] We use AD9954 DDSs manufactured by Analog Devices.

principle all that the ECS software does, is to create this table and then send it to the DIO64 card which then generates the output. This table could also be made manually and then send to the DIO64 card—this would however be *very* complicated! For the run giving the first BEC ECS had more than 3.8×10^6 rows in the table for the DIO64 card handling the analog channels. It is however only the analog channels that needs these massive amount of data—in comparison, the table for the DIO64 card handling the digital channels had less than 100 rows in the first BEC run.

TS	TS	1-16	17-32	33-48	49-64
0x0000	0x0000	0x0000	0x0000	0x0000	0x0000
0x9680	0x0098	0x0000	0x0004	0x0000	0x0000
0xF828	0x0098	0x4000	0x0104	0x0000	0x0000
0xF832	0x0098	0x0000	0x0000	0x0000	0x8000

Table 3.2: The DIO64 data format. The first two columns in the array contains the Time Stamp (TS) of the event. The last four columns contain the output on channel 1 to 64 divided into groups with 16 channels in each. Hexadecimal numbers are indicated by 0x. A table like this one is generated and sent to the DIO64 card before every run.

3.3.2 Electronic modules

In the case of the analog channels the electronic modules are needed to be able to translate between digital pulses and analog voltages. The equipment that needs digital signals could in principle be connected directly to the DIO64 terminals. However all the DIO64 channels are connected to custom-made electronic modules. This is done for the following three reasons:

- Usability: The electronic modules provide a really nice interface to ECS since the modules are distributed across the laboratory. This means that every time new equipment needs to be connected to ECS there is always a module/channel close by. We just have to connect the equipment to the module(s) (using BNC-cables) and then tell the ECS-software that a new machine has been connected.
- Ground issues: All the electronic modules are electrically isolated by means of optocouplers from the DIO64 boards. This is done to protect the equipment from the noisy environment of the computer and avoid ground-loops.
- Noise issues: The signals from the DIO64 cards are pretty noisy because of the relatively high clock frequency and the long cables (up to 15 meters) connecting the DIO64 cards and the modules. All of the signals are therefore cleaned up in the electronic modules.

In table 3.3 an overview over the different electronic modules can be seen. They are all made as EURO-modules with BNC-connectors on the front for output, and input from the DIO64 cards on the back (BNC or D-sub (9 or 15 pin)). The modules all have a response time of less than 100 ns and most have four channels per module. From the table it is seen that there are seven different types of modules. The racks for the modules are located five different places in the laboratory.

The two types of digital output modules are capable of controlling equipment that can be "on or off" such as shutters, flippers, fans, camera triggers, triggers for the mechanical positioning system and the different transistors in the current controllers for the magnetic

Electronic module	Channels per module	Description
Analog output (DAC)	4	The analog channels can be used to set a 0 to 5 V voltage (actually the region can be set arbitrary—we also have a couple of channels with a -1 to $+1$ V interval but we try to restrict ourselves to 0–5 V).
Digital output—type 1	4	General TTL trigger (0 or 5 V).
Digital output—type 2	4	Like type 1 except that the high voltage can be adjusted continuously from 0 to 5 V (3 channels: high-voltage set by potentiometer, 1 channel: high-voltage set by external voltage—e.g. from an analog module).
Digital input—type 1	4	Can detect a low or high (>2.5 V) signal. Used for status check of the system, trigger or interrupt of the ECS software.
Digital input—type 2	2	Like type 1 except that the trigger level can be varied by a potentiometer or by a reference BNC-input voltage.
DDS (RF generator)	1	10 kHz to 70 MHz RF-synthesizer
DDS (lattice control)	2	10 kHz to 160 MHz RF-synthesizer

Table 3.3: The different electronic modules. In figure 3.1 images of the different modules can be seen.

coils. The analog channels can be used to control equipment such as AOMs, Electro-Optical Modulators (EOMs), current in the different magnetic coils, the power of the microwave- and RF-signals and so on. The two types of digital input channels are used as triggers and interrupts.

3.3.3 DAC and DDS programming

We use serial digital data transmission for the data to the DAC and the DDS modules. To ensure that the DAC modules and the DIO64 cards are synchronized the DIO64 card is used to create the clock frequency for the analog modules. This clock does not run continuously because it is generated from rows in the DIO64 table and we want to make this table as small as possible (the clock is only needed when data is sent). Alternatively we could use a central clock (a commercial frequency reference that just runs continuously) and distribute it throughout the laboratory to all electronic modules, DIO64 cards and so on. This would perhaps also be a good solution, however it is nice to generate all the signals in the software since we can then also very easily change to e.g. a new DAC-chip.

Slightly after the data bits are received the analog module receives a load pulse from another DIO64 channel. The load pulse tells the multiplying 16-bit DAC-chip to set a voltage on the analog channel that corresponds to the last 16 bits received

$$V_{out} = \frac{\text{word}}{\text{0xFFFF}} \cdot 5\,\text{V} \tag{3.1}$$

here "word" refers to the 16 bit word (numbers between 0x0000 and 0xFFFF) sent to each analog channel and converted to a corresponding output voltage V_{out} on the front of the module. The module will output this voltage until a new load pulse is received.

The four channels in each DAC-module share the same clock and load. This may sound strange since we then have to clock out the voltages for all four channels even though we only want to change the value of one of them. However if the voltage on one of the four channels needs to be changed the necessary time stamps in the DIO64 table are already reserved. It does not matter whether one channel, four channels or all of the DAC-channels are updated—the same number of rows in the DIO64 tables are needed. This also means that nothing is saved by using different clocks or loads for the different channels in an analog module. On the contrary by using the same clock and load only six wires (or DIO64 channels) have to be used for each analog module—one wire for load, one for the clock and four wires for the data transmission to the four analog channels. In principle we could use the same clock and load for all the modules, but we do not do this because of noise considerations.

The system uses the "nonreturn to zero" (NTZ) technique where the duration of each digital pulse is equal to the period of the clock. Therefore the 16 bits for the DAC module actually take up about twice the number of rows (namely 34) in the DIO64 table because the clock pulse has to go high and low for each bit. The bits are read out on falling clock signals—this means in the middle of a data pulse (see figure 3.3). The length of a half clock period is set to $5 \times 0.1\,\mu\text{s} = 0.5\,\mu\text{s}$. We began to see some errors in the transmission at $0.1\,\mu\text{s}$ so we chose $0.5\,\mu\text{s}$ to be well above this limit. With this value it takes $34 \times 0.5\,\mu\text{s} = 17\,\mu\text{s}$ to clock out the 16 bits. The DDS-chip (described in the next section) needs up to 48 data bits for each transmission. This corresponds to 98 rows in the DIO64 table and this sets the limit to our update rate since it takes $98 \times 0.5\,\mu\text{s} = 49\,\mu\text{s}$ to update a DDS module. Therefore the general update interval of ECS was chosen to be $50\,\mu\text{s}$ for the analog channels. However running a DDS with a $50\,\mu\text{s}$ resolution for just 1 s typically take up around 1.7×10^6 rows in the DIO64 table so the default time resolution is only 1 ms (the resolution of each ramp can be set in the ECS software).

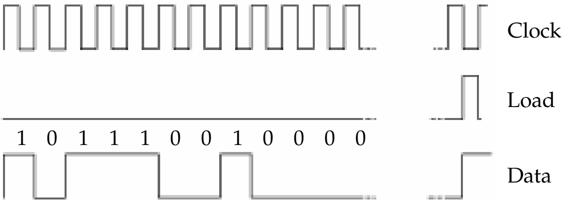

Figure 3.3: This sketch shows how the serial data transmission is made. The clock wire sets the clock frequency for the multiplying DAC. The load wire indicates when the 16 bits have been sent. The data wires (here only one is shown but there are four for each analog module) contain the 16 bit word that sets the voltage on the DAC channel.

DDS

The principle behind the DDS programming is slightly different from the way the DACs are programmed since we use a build in linear ramp function in the DDS-chips. Initially we were using the same principle as used for the DACs where we clock out different *static* values that are used until the next value is received. However the problem is that ECS only has a $50\,\mu\text{s}$ time resolution on the analog channels. This means that if we e.g. want to make a fast ramp in frequency with a DDS it is forced to "jump" the frequency in large steps every time it is updated. Sune has been doing experiments with RF-transfers and here it is very important that we can sweep the frequency of the DDS with a step size that is smaller than the Rabi-frequency of the transition. For this reason we had to change to the linear sweep mode of the DDS. This gives a lot of complications but the good thing is that the frequency is updated at least once every 100 ns.

One problem with this solution is that there is a finite resolution on the frequency step size (32 bit) and on the number of clock cycles (or time) between the frequency ν is stepped (8 bit). This means that the slope of the linear sweep

$$\frac{d\nu}{dt} = \frac{N\Delta\nu}{M\Delta t} \tag{3.2}$$

can almost never be chosen exactly right (here N is the number (0 to $2^{32}-1$) of frequency steps $\Delta\nu$ and M is the number of clock cycles (1 to 256) between the frequency is updated. We use a 400 MHz clock for all our DDSs but the frequency is only stepped once every four clock cycles (at a frequency of 100 MHz) which means that $\Delta t = 10\,\text{ns}$ and $\Delta\nu \approx 23.3\,\text{mHz}$).

By calculating the *continued fraction* of a given real number (the frequency slope) one can calculate the "best" fraction N/M representing this number with certain restrictions to the size of the denominator and/or numerator [68]. These mathematical models also contains possibilities to always choose a fraction that is smaller/bigger than the real number. ECS uses this model and in our case the denominator M is restricted to be a number between 1 and 10—because of this the output frequency is updated every 10 ns to 100 ns. However, if the sweep is several seconds long it can still be a really bad approximation. Therefore the slope of the sweep is changed/updated after 1 ms to 500 ms (can be set in the software where the ramp is generated). This means that if a linear sweep is chosen in ECS it will perhaps not be completely linear but have small bends. Of course ECS takes this into account and makes sure that we end up at the right spot in the end, but still it can be annoying if we want to know the exact sweep rate in some experiment—for this reason we will implement functionality in ECS that makes it possible to lock the sweep rate.

The DDS programming is very difficult to do really good. We have made all kind of strange functionality to be able to change the sweep slope and the sweep direction so that we can create completely arbitrary waveforms. The main problem is that ECS does not use the same clock as the DDS-chip. Because of this we get a bit out of sync as time goes by. This might sound as a small effect but after a 30 s sweep it can be significant. This means that we have to "reset" the frequency continuously during the ramping (and this is what we do every 1 ms to 500 ms). We do this by exploiting that the DDS has a max/min limit on the frequency sweep which we can use to force the frequency to the right value. If we always slightly overestimate the slope of the sweep we will hit this max/min value some nano/microseconds too early. This way we are sure that we have the right frequency before the next sweep (or slope-change) is performed. However this means that we have to update the min/max frequency all the time and this together with the fact that we want to be able to make completely arbitrary waveforms (i.e. be able to change the sweep direction and rate arbitrary many times) makes it hard. If the DDS and ECS were perfectly synchronized we would know the exact output value of the DDS in all of the ECS clock cycles. In this case we could make sure that the load pulse that e.g. change the sweep rate or direction would be received in the right clock cycle and everything would be fine. We have thought about using the same clock for the DIO64 cards and the DDSs (of course multiplied by some integer) and we think this would ease the programming significantly. However now we have made this solution and it works.

3.4 Software

The software part of ECS consists of more than 18.500 lines of source code to be able to provide all the functionality that we want. All the sequences, settings and calibrations created by the user in ECS is saved in a *.ecs file. From these informations the output on all the different channels can be generated. ECS always backup the last 5 runs no matter what. Furthermore we normally run ECS with autosave on. This means that a new ECS file is created for each run. Because of this we can always reload a run and see exactly how the settings where in this specific run and then modify it or run it again. With autosave on, each file is automatically given a time stamp containing the date, time and run-number and saved in a folder with the right week-number.

The program is designed to use two monitors. The first monitor contains the `Main Window` which is the primary interface to the software. The `Main Window` is divided into two panels (see figure 3.4). The digital output can be edited directly on the left panel but this panel also gives access to the `Digital Window` where more advanced settings for each digital channel can be set. The right panel provides access to the `Analog Window` where the different ramps can be generated. The different channels are ordered in *groups* to increase the clarity in the `Main Window` and also other places in ECS. A screenshot of the entire `Main Window` can be seen in figure B.1 (in appendix B).

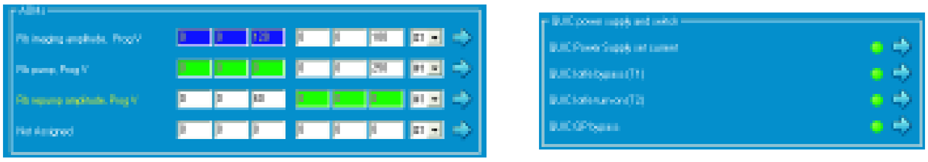

a) Digital group

b) Analog group

Figure 3.4: Examples of a digital and an analog group. (a) The digital group “AOMs” taken from the left panel located in the `Main Window`. This group contains four channels. The three first edit boxes contain the start time of a pulse, given in s, ms and µs respectively. The next three boxes contain the stop time or duration of the pulse according to the settings the user has made for the specific channel. The drop down box gives the opportunity to switch between the different pulses. The arrow opens the `Digital Window` for the particular channel. (b) An analog group also containing four channels. The green traffic light indicates that the channel is enabled and the arrow opens the `Analog Window` for the particular channel.

The second monitor is reserved for the `Event Browser`. In this window all events on all the different channels are represented as time-lines (see figure 3.5). The current maximum run-time in ECS is two minutes which should be enough for the experiments we are performing (however if necessary it can easily be increased). Each time-line either contains the complete 2 minutes run sequence or a selected part of it. If the complete run sequence is shown the left and right side of the time-line corresponds to 0 s and 120 s, respectively. However the user can zoom in on events (all the way down to the µs regime) and in this case the left and right time-values changes. A time marker can be placed to see if different channels are correctly synchronized and we can also e.g. measure the time between two different events in this window. The `Event Window` is an important tool in the debugging of a run sequence and in general it provides a great overview of what is going on in the current run. A screenshot of the `Event Browser` can be seen in figure B.2.

Not all of the 48 digital and 32 analog channels are shown by default in the `Main Window` and in the `Event Browser`. The channels that are shown is a consequence of how the channels are defined in `Options` (see figure 3.6). The channel structure is saved in every ECS file. This means that every time an old file is loaded the structure in the `Main Window` (and many other places in the program) changes. In this way we can in a

a) Digital channel

b) Analog channel

Figure 3.5: Time-lines of a digital and analog channel. (a) shows two trigger pulses on a digital channel controlling one of the CCD cameras. White means high (5 V), black low (0 V). Since this camera responds on a falling edge in this case the channel is inverted (this can be set in the `Digital Window`) and the pulses are therefore black. (b) shows an analog channel. Here the colors are not only white and black but represents all the different voltages between 0 and 5 V—or between a subset of voltages used in the current view. The complete `Event Browser` can be seen in figure B.2.

few seconds switch between completely different experiments that use different channels and have completely different run sequences. E.g. if we are doing mixtures between two rubidium clouds we do not need to have any of the optical lattice channels.

ECS uses two different sets of channel numbers: *software numbers* and *laboratory numbers*. There is a one to one correspondence between these numbers. Every channel on all the electronic modules has its own laboratory number that stays fixed and is shown on the module. However, in the software we want to be able to group different channels and change the way the channels are ordered. Instead of switching the cables of the different modules we assign a software number to the laboratory number (see figure 3.6). In this way we can change the order of the channels in the software without affecting the output. We plan to expand this functionality with a third set of numbers—see section 3.4.9.

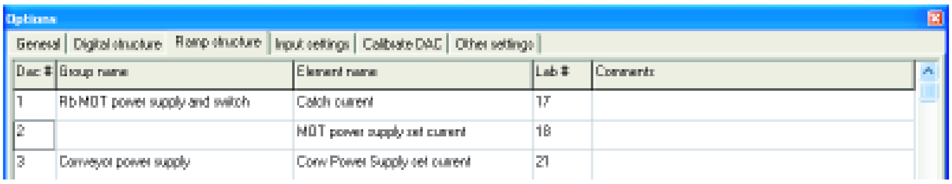

Dac #	Group name	Element name	Lab #	Comments
1	Rb MOT power supply and switch	Catch current	17	
2		MOT power supply set current	18	
3	Conveyor power supply	Conv Power Supply set current	21	

Figure 3.6: A segment of the `Options Window`. Here the Ramp Structure tab is shown. This tab contains the current structure of the analog channels. Here the different channels can be named and joined in different groups which makes the `Main Window` more structured. Also the one to one correspondence between software channels and laboratory channels is set here. When this correspondence is made the user can change the order of the software channels and insert new channels wherever needed.

3.4.1 Digital Window

In the `Digital Window` and in the `Main Window` the digital output is created by the user. In these windows we can make digital pulses with a duration down to 1 µs. In the `Digital Window` a number of settings for each digital channel and for each pulse can be set. For example we can account for dead time in the equipment by specifying a time offset for each digital channel. With this offset the synchronization of different equipment is easier because the user does not have to account for dead time in the equipment—ECS does it. The user can decide if the dead time should be shown in the time lines or not.

3.4.2 Analog Window

A lot of thoughts have been put into the design of the way the analog channels should be programmed by the user and actually this is one of the only parts of ECS that has been redesigned. We ended up with the principle of creating the ramps as shown in figure 3.8. The idea is that the user defines so-called *time nodes* (points defined by their time and value/voltage). The line that connects two time-nodes is called a *time segment*. For each time segment a ramp type can be chosen—e.g. linear, exponential, s-curve or sine. ECS automatically calculates the intermediate points between all time nodes and shows the result on a graph in the `Analog Window` (see e.g. figure B.3) and in the `Event Browser`. In the case of the linear ramp no extra information is needed to construct the ramp whereas in the case of e.g. the exponential ramp a time constant τ needs to be specified. Currently we have these four different line types but any imaginable ramp-type could easily be implemented. The expressions for the line types can be seen in table 3.4 and the corresponding curves can be seen in figure 3.7. For each time segment we also specify a time resolution (for the normal analog channels it can be 50 µs, 100 µs, 500 µs, 1 ms or 5 ms). This time resolution determines how often a ramp is updated. The default value is 1 ms but for a particularly critical event better resolution can be chosen. If 50 µs is chosen ECS can e.g. represent a 1 kHz sine and this can e.g. be used to measure trapping frequencies in the magnetic trap as described in chapter 5.

Ramp type	Expression	Needed inputs
Linear	$V(t) = V_0 + a(t - t_0)$, where $a = \frac{V_1 - V_0}{t_1 - t_0}$	None
Exponential	$V(t) = a_0 + a_1 \exp((t - t_0)/\tau)$, where $a_0 = V_0 - a_1$ and $a_1 = (V_0 - V_1)/(1 - \exp((t_1 - t_0)/\tau)$	τ can be pos/neg.
S-curve [69]	$V(t) = V_0 + \frac{V_1 - V_0}{1+\exp(-\alpha(t-(t_1+t_0)/2)/(t_1-t_0))}$	α (> 0)
Sine	$V(t) = \min(V_0, V_1) + \frac{\lvert V_1 - V_0\rvert}{2}(1 + cos(2\pi(t - t_0)\nu + \delta))$, where $\delta = 0$ if $V_0 > V_1$ else $\delta = \pi$	ν (> 0)

Table 3.4: The different ramp types in ECS. The ramps are defined between two time nodes (t_0,V_0) and (t_1,V_1). It is assumed that t_1 is larger than t_0 but no restrictions is put on V_0 compared to V_1 (they can even be negative which can make sense if a calibration is used for the channel—this will be described later). Notice that the sine behaves slightly different that the other curves since it does not necessarily go through the end time node (but always through the start time node). This is because the user specifies the frequency ν (not the number of cycles between the time nodes) and for this reason the curve will not go through the end time node unless $t_1 - t_0$ is chosen right with respect to ν. Otherwise ECS will just jump to V_1 in the end. However the graphs showing the output (and the `Event Browser`) always show exactly what is sent out so the user can immediately see it and adjust the duration of the time segment if wanted.

An example of the way we define a ramp can be seen in figure 3.8. Here we start after 4.012 150 s at 0 V. During the next 2.322 s the channel is ramped from 0 V to 4.3234 V using a linear ramp. The time resolution of this time segment is set to be 50 µs. For the next 5.587 s the voltage is kept at 4.3234 V. After this an exponentially decreasing ramp with a time constant of -1.754 s and a resolution of 1 ms ramping down to 0 V in 12.5 s is defined. In a similar way the rest of the output can be analyzed.

It can also be seen in figure 3.8 that the ramps are ordered in groups. The first time node in a group is an *absolute* time T whereas the rest are durations Δt. This means that by changing the time of the first time node the whole group is shifted. This also means that by changing the duration of one of the time segments somewhere in a group all

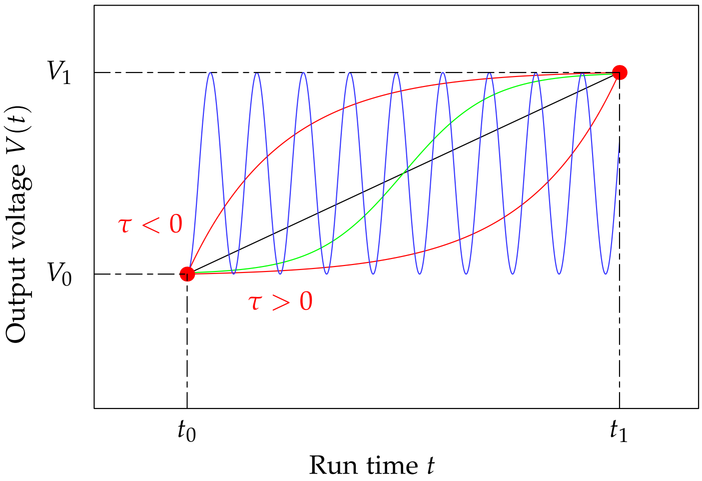

Figure 3.7: The different ramp types currently implemented in ECS: Linear (black), Exponential (red), S-curve (green), Sine (blue). Notice that the Exponential ramp can "bend" both ways depending on the sign of τ. Two time nodes ($[t_0, V_0]$ and $[t_1, V_1]$) specifies the end points of each curve. ECS calculates the intermediate points—sometimes an extra parameter is needed to specify the curve (see table 3.4). When these different ramp types are lined up after each other really complicated ramps can be generated. However ECS also has a feature that makes it possible to load data points generated manually or in another program. This way any imaginable ramp type can be used.

the following time segments are moved as a response to this change (but other groups are not effected by this). This is an essential feature because often we are interested in changing the duration of a certain event in the run sequence. If we e.g. want to make the transport sequence from the MOT-chamber to the science chamber faster, then everything that we do after the atoms arrive to the science chamber needs to be moved (also on all the other channels). In practice this is almost impossible or at least very complicated to do manually because there are often *many* events that need to be moved! Exactly how ECS is able to move the output on all channels will be described in the next section.

An important feature is that the graph showing the output in the Analog Window really shows *exactly* what is sent out. This means that one can zoom in and see every point in a given ramp. When we zoom in on a few points we can also see how the output in reality is stepping the voltage. If a time resolution of 50 µs is used we will see small steps—if 1 ms is chosen the steps will be larger.

In the example shown in figure 3.8 a voltage with 5 significant digits is set. 16 bit corresponds to a precision better than 0.0001 volt and the reference voltage on the analog modules support this precision but it is not necessarily equal to 5.0000 V. Therefore a calibration of each analog module (made with a very precise multimeter) can be specified in `Options`. This is however not the only calibration that can be made. A calibration between the output voltage and a *new unit* can also be set. The new unit could e.g. be an AOM-frequency (in MHz) or a current in the magnetic coils (in A). This means that instead of typing a voltage corresponding to e.g. 350 A we can just type 350 into the value field in the `Analog Window`. This unit change is based on a cubic-fit. This means that even though the equipment does not have a linear response a good representation can still be made. This unit change is made in the `Analog window`. Here also a minimum

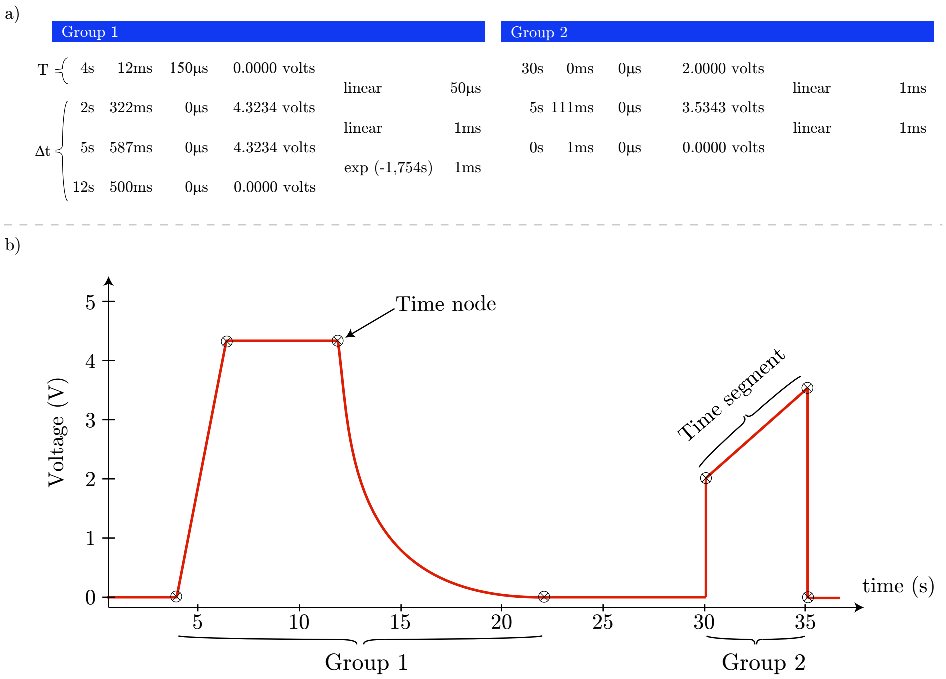

Figure 3.8: The principle behind the way analog ramps are generated in the `Analog window`. The ramps are defined by specifying so-called *time nodes*. For each *time segment* the ramp type and the resolution can be set. See text for further comments.

and a maximum value for each channel can be set (to avoid setting a damaging voltage to sensitive equipment) and the channel type (Analog, RF DDS or lattice AOM DDS) can be chosen. If one of the DDS types is chosen the unit automatically changes to frequency (kHz). The `Analog Window` can be seen in figure B.3.

When building the DIO64 table for the analog channels ECS makes sure that only the absolutely necessary number of time stamps are used to generate the output. This means that if nothing changes on any of the channels in a certain time period, none of the channels are updated. ECS also analyzes whether or not a change is made on one of the DDS-channels every time something is clocked out. If not, the DDS channels are not updated and less time stamps are needed (since up to 48 bits are used for a DDS but only 16 bits are used for an analog channel). These features dramatically reduce the amount of data needed to make the output on all the analog channels. Typically it takes about 3–6 seconds to build the tables for the DIO64 cards.

3.4.3 Global offsets

The *global offsets* are used to shift/change events taking place on a number of different ECS-channels simultaneously. ECS contains 40 time- and 40 value-offsets (can be increased) all specified in the `Global Offset Window` (see figure 3.9). The global offsets can be used everywhere a time or a value (voltage, frequency, current, . . .) is specified.

When e.g. a time-offset is defined for a time field somewhere in ECS the offset is added to the time already specified in the field. In figure 3.4 three time offsets on the digital channels are set (indicated by the blue and green colors of three of the time-boxes). All global offsets have different colors and when a global offset is added to a field the color of the field changes from white to the color of the offset. This makes it easier to get an overview of the offsets.

MOT	0	5	0	2/55	I_catch	100	2
CMOT	0	35	0	1/55	I_MOT_transport	250	2

Figure 3.9: An extract of the `Global Offset Window` where the global offsets are defined. At the left the time-offsets are shown, at the right the value offsets. In this window 40 time and 40 value offsets can be specified. Sums of time offsets can also be defined in this window. The first time offset specified here is called "MOT"—by changing the value of this offset the MOT period is changed (this offset is normally set to zero because we use a trigger-input to control the loading time of the MOT). This means that the following outputs on all the ECS-channels are shifted in time.

The opportunity to add one time-offset to a field is not always enough because often an event in the end of a run sequence depends on many different events earlier in the run. It is therefore necessary to be able to specify more than one time-offset for each field. However after many considerations we decided that only one offset should be allowed for each field, but instead new offsets containing sums of other offsets can be defined. This provides the same functionality.

In the run resulting in the first BEC 30 time-offsets and 13 value-offsets were used. The time-offsets were used 66 places in the program and indirectly (through other offsets) 894 times! This might sound complicated but it is "complicated" for ECS—not for the user and that is the point.

Now it is also more clear why the ideas described in the last section involving time nodes were important. By defining a global offset on the first time node in a group the whole group is moved when the offset is changed. By defining a global offset on one of the other time nodes in a group the duration of the associated time segment will be changed and the rest of the group will be moved but otherwise unchanged. Since a global offset can be assigned to all fields (on all channels) and since sums of offsets can be defined this provides all the functionality one could think of.

The global offsets also makes it possible for ECS to handle a series of runs where one of the global offset is automatically changed from run to run—thereby we can measure how some parameter (e.g. the atom number) changes as a function of the global offset. Before looking into different run-modes a couple of other features should be mentioned.

3.4.4 Static output

From the `Static Output Window` a static output on all the channels can be set. This feature is very useful when setting up equipment and for general testing purposes. However the output for making a MOT is also a static procedure (we just turn on the laser light and the magnetic field). In every ECS run five default settings for the outputs can be set (e.g. the MOT settings). Besides these five static output settings we also have a (static) initialization sequence which contains non-damaging settings for all the channels. This initialization is called when ECS opens, closes, before and after each run and if a run is aborted. This ensures e.g. that if one of the current controllers is left at 400 A after a run sequence by a user-mistake it will be pulled low when the initialization is made.

3.4.5 Input trigger/interrupt

As mentioned earlier ECS has 16 input channels that can be used to interrupt a run in case of problems or to trigger a new run.

ECS has a table in `Options` where the expected inputs on some/all of the 16 channels can be set. If one or more of the inputs are wrong when a new run sequence is executed the run will be aborted. Also during a run sequence ECS checks every 3 seconds to see if all the inputs are still right. If not, ECS will immediately abort the run and execute the initialization sequence. So far only the temperature sensors on the magnetic coils are monitored by ECS. If a coil gets too hot (e.g. if no cooling water is turned on) the temperature sensor connected to one of the input channels will change status.

An example of how an input channel can be used as a trigger is described in the following. When executing a run sequence ECS can be configured to run one of the 5 static output sequences before the actual (dynamic) run sequence is started. This mode is associated with the trigger function because it is a change in the status of one of the input channels that determines when ECS goes from the static output to the real run mode. We use this function to ensure that the number of atoms collected in the MOT is approximately the same from run to run. By choosing the static output sequence containing the MOT-settings and connecting a MOT-detector to an input channel the number of atoms loaded in the MOT can be kept constant. The detector is connected to an input module type 2 (see table 3.3) which means that the input status changes when the detector signal goes above a pre-defined voltage. This pre-defined voltage can be set manually on the module or it can be provided through a BNC-input connector. This pre-defined voltage can therefore be provided by ECS itself since an analog channel can be connected to the BNC-input connector on the input channel. The pre-defined (reference) voltage is set in the static MOT-sequence (which is used until ECS is triggered).

3.4.6 Run control window

In the `Run Control Window` all settings concerning a run execution can be defined. It is e.g. here we chose which cameras we want to use, what the trigger-mode should be and much more. Three different run-modes can be chosen: normal, cyclic and macro (all of these modes can be combined with the input trigger as described in the previous section). The normal mode just executes a run. The cyclic mode executes the same run sequence over and over again, whereas the macro mode also executes a number of run sequences but here the value of a global offset is updated from run to run. In the future we will implement a feature that makes it possible to step through two global offsets at the same time (the system is prepared for this functionality). It will be possible either to change both offsets before each run or to step through all the possible combinations of the two (within a specified range and with a specified step size).

3.4.7 Axis settings

All the parameters used in connection with the cameras and the general imaging sequences are set in axis-settings in ECS. Here 150 parameters can be defined for each of our camera axes (x-axis, y-axis, z-axis and aux-axis). The idea with these parameters is to have all information about a given run saved in the *.ecs file—also camera/imaging settings which are not directly used in ECS but just passed on to CCS and later to the fit-program. By having this kind of system we always know exactly what all the settings were when the run was made—and we only have to reload the corresponding *.ecs file to see it. The axis settings contains information (in 4 columns for x-, y-, z- and aux-axis) about the camera (e.g. gain, exposure or shutter time, mode, temperature, offset, EMCCD gain, pixel shift speed and fan settings) and the imaging conditions (e.g. calibrated pixel size, polarization, imaging AOM frequency and absorption (Lorentz) line width).

The 150 parameters is defined in a text-file (but the values of the parameters are set in ECS). Each parameter can be defined as an integer, a real or a string. When a parameter is defined it should (for backward compatibility) never be changed into something else. If a parameter is outdated we just leave it. If we have to add additional information we use a new parameter that has never been used before (also because the meaning of each parameter is hard-coded in CCS and in the fit program).

Global Offsets can also be used in the `Axis Settings Window` and this means that the values here are automatically updated e.g. in the case where we are running a macro run and we are changing a parameter that is defined in the `Axis Settings Window` (e.g. the imaging AOM frequency—which later enters in the atom number calculation in the fit-program).

3.4.8 Other features

Some additional software modules are added to ECS which are specific for our laboratory. We have e.g. a module for a frequency generator (Agilent 22330A), a microwave generator (Wiltron 6722B) and motorized translation stage (Thorlabs Z625B) which is used in the imaging sequence. These modules are an integrated part of ECS and this means that we can e.g. use global offsets in the programming of these machines. If we are using the frequency generator we can e.g. define global offsets for the frequencies that the generator should use in the given run—this is especially useful in connection with macro-runs. Since all the Global Offsets are also sent to CCS and later imported in the fit-program we can choose to view these global offsets in the data-analysis. This means that there is no manual work (and no chance of possible mistakes) since the data entering the data-analysis is never manually typed in by the user. The frequency generator is e.g. used when we measure trapping frequencies of the optical lattices as described in chapter 8.

In the case of the microwave generator ECS can program it in such a way that we e.g. can transfers the atoms between the different hyperfine states. The user interface is made in such a way that the user can just select the start state (e.g. $F = 2$, $m_F = +2$) and the wanted end state. ECS will then calculate the needed sweeps and the user has to use a normal digital channel to provide the necessary trigger pulses for the generator (and thereby control when the transfers should occur). Even though this transfer works fine we would like to have an even more precise frequency control over the microwave generator during sweeps. For this reason we have bought a mixer and we will perhaps in the future use a (not yet created) DDS-module together with the locked microwave generator to get a frequency precision in the Hz regime (during sweeps the (unlocked) microwave generator has a large uncertainty in the frequency—up to one or two MHz).

ECS checks *everything* the user does extensively. If something is wrong it will be handled. If the user makes an error on one of the output channels (e.g. tries to make a time-overlap between two different events on the same channel) ECS will indicate this by showing an "error" on the relevant channel(s). This is a strong feature because it means that the user is reminded that something is wrong but without showing an error message in a pop-up window or likewise. Furthermore everything is updated in ECS without the user has to press an "ok" button or something similarly. For every keystroke everything of importance is updated (including the `Event Browser` and the other graphical representations of the output on the different channels). E.g. if a global offset is changed it will often affect the output on many channels at the same time but for each keystroke the user will see the effects and be able to check that everything is still fine. ECS will of course not be able to execute a run before all errors are corrected.

This ends the overview of the most important features in ECS. In the following section a few possible improvements will be described which we plan to implement soon.

3.4.9 Possible improvements

Queue System

A nice improvement to ECS would be to implement a run queue system. Currently we are often using macro-runs where the system automatically is varying one parameter. However one could imagine to extend this idea and make it possible for ECS to control a very long sequence of runs automatically. This way we do not have to be in the laboratory if we are just measuring something where we do not have to take action upon what we see. This feature can be implemented by having a queue list in ECS where we can specify the ECS-files that we want to execute. In the beginning of a day we would have to prepare different ECS-files and make a schedule for the execution (put them in the queue). ECS would then just load the first file in the queue, execute it, load the next file, execute it and so on. The ECS-files could be normal runs, cyclic- or macro-runs. If the MOT does not load or if one of the digital inputs (or interrupts) are wrong ECS should be able to tell us that we have a problem. This could be done easily by allowing ECS to send e-mails. This way one could also get a SMS-message on the mobile phone using one of the free e-mail to SMS services that exists on the Internet. This queue system (including e-mail service) could be implemented in a day or two since nothing in the existing code would have to be changed—it would just be a small module on top of the exiting system.

Instrument numbering

We already have two sets of numbers in the system: software numbers and laboratory numbers (as described earlier). A nice way to do this even better would be to give all the different instruments a unique *instrument number* (or more than one if it has different inputs/outputs). This way one would have to specify three different numbers—a software number (it defines in which order the channels should be shown in ECS), a laboratory number (it specifies a specific channel in an electronic module placed somewhere in the laboratory) and an instrument number (it is related to the instrument). Every time an old instrument is replaced by a new one new numbers should be taken into use even though the new instrument does the same as the old one. This means that even though we e.g. only have 32 analog channels there can be much more than 32 instrument numbers using analog inputs. This way every unit calibration in ECS should be made with respect to the instrument number and not with respect to the laboratory number. The calibration of the "5.0000 V" voltages should still be related to the laboratory numbers. A table should relate software number to instrument number—this table should be saved in the *.ecs file. Another table should contain information about which laboratory number a certain instrument (number) is connected to—this table should be saved globally (not in the *.ecs-files and therefore be used in all runs). This way we can move an instrument to another channel and all we have to do is to change the laboratory/instrument number relation in the ECS software. Since this is stored globally we can also load old runs and everything will still work (as long as all the instruments that were used in this run are still connected somewhere in the laboratory) since the ECS run only contains the relation between the software number and the instrument number. Since all the laboratory numbers are calibrated (to 5.0000 V) the calibration made for the specific instrument will also still work since what is sent out takes into account both calibrations (both the 5 V calibration and the unit-calibration—if used). If we implement this scheme ECS would also know if we loaded an old run which was based on a different instrument and it could therefore also warn us about changes in the experimental setup since the run was executed (which could be a long time ago). It would be a few days of work to implement this feature.

3.5 Mechanical Positioning System

The Mechanical Positioning System (MPS) is based on commercial products manufactured by *Parker Compumotor*. The system consists of three motorized translation stages about 75 cm in length. The translation stages are heavily loaded since besides supporting the magnetic coils they also have to move the associated cables and the tubes for the water cooling of the magnetic coils. The translation stages are specified to have a precision of 5 μm but from run to run it is closer to 50 μm. Each of the translation stages have their own controller (a ViX500AE) which is controlled by a so-called 6K4 controller. The 6K4 controller is connected to the ECS computer by ethernet and it contains *many* possibilities for programming the stages. We use a built-in so-called S-curve motion (a smooth motion) that can be programmed by setting only the distance, maximum velocity, average acceleration, maximum acceleration, average deceleration and maximum deceleration of the motion. However the units used for these parameters have to be given in "counts" and "revolutions". Therefore two calibrations were made to translate to SI units[2]. With these calibrations the position, velocity and acceleration to all times are known. ECS also uses these calibrations to calculate the travel time of the individual motions. In a typical run the MOT-coils are moved 493 mm to the vacuum chamber cross in about 3.5 s and obtain a maximum velocity of 200 mm/s and a maximum acceleration of 1200 mm/s^2. After that the conveyor coils are moved 369 mm in about 3.2 s with the same maximum velocity but the maximum acceleration is only 300 mm/s^2. The difference is due to the fact that we want to get out of the MOT-chamber fast (because of the low lifetime). We always use half the value of the maximum acceleration for the average acceleration since this gives the most smooth curve.

MPS is also completely integrated in ECS. This means that the user just has to specify the parameters for the S-curve motion, the stage number and the start time for the motion (which can be global offsets). Before each run ECS creates a 6K4 program and uploads it to the 6K4 controller. This program contains (aside from different settings) the S-curves for the different stages. Each motion is initiated by an external trigger input on the 6K4 controller from one of the digital modules. Currently we have implemented the system in such a way that it can only move one translation stage at a time (to save trigger pulses), but the system is prepared to support motion on the different axes at the same time. This will be necessary if a second species is implemented.

The translation stages also have so-called end-stop sensors that can be programmed to stop the motion if a coil moves too far. However if this software safety mechanism for some reason does not work we have mounted solid physical stops on the stages to ensure that the coils will not hit the vacuum chamber. A special reset sequence that uses the end-stops and is configured in ECS can be executed. This sequence is not a built-in function of the MPS controller so it had to be created. The sequence is able to place the stages at a given position (also with an uncertainty of 50 μm). This reset sequence is e.g. used if a run is aborted. In this case MPS will shutdown the motion on all the different translation stages and they are therefore left at unknown positions.

MPS induces vibrations in the vacuum table and this means that e.g. if we have atoms trapped only in the optical lattice the lifetime is reduced from many seconds to only a

---

[2] The velocity and acceleration is measured in rev/s and rev/s^2 respectively ("rev" means revolutions). Distances are measured in counts. We have made two calibrations: 409.7±0.1 counts/mm and 10.00±0.01 mm/rev. These two calibrations are used in ECS to calculate distances, velocities and accelerations in SI units.

few if one of the translation stages are moving during this “hold” time.

3.6 Future expansion

Currently we have 32 analog channels. After implementing the 3D optical lattice we have used around 25 of these to control the entire laboratory. This means that in the future we will run out of channels if e.g. a second species is implemented. Currently we are not using all the DIO64 channels on the PCI-card controlling the analog channels so we could get a few more channels this way. Also we could buy more DIO64-cards and place them in the ECS-computer together with the two cards we already have. However the more channels we use the more calculations ECS has to perform to create all the necessary digital pulses for the DACs.

3.6.1 FPGA

Another solution that will give us many more channels and take away much of the computation time from ECS is to use Field-programmable Gate Arrays (FPGAs) to program the analog channels. A FPGA is a semiconductor device that can be used to make complex digital systems (relatively) easy. The FPGAs we are using have 1 million logical gates that can be connected almost arbitrary and this means that very complicated logical hardware can be made just by sitting in front of a computer. We use an “all-in-one” circuit development card[3]. This kind of card is normally used when designing a circuit and it has many nice features: e.g. it has its own 50 MHz clock, switches and Light-Emitting Diode (LED)s for testing purposes and 16 mb flash memory. These cards are cheap (around 100$) and the FPGA can be programmed and reprogrammed with a computer using the free Xilinx ISE WebPACK software. These boards are mounted on home-made printed circuit boards placed in EURO-racks (just as all our other ECS electronic modules)[4]. The FPGA-modules can be connected directly to the computer using USB connectors and to the analog or DDS modules using D-sub connectors. So far we have not changed to the FPGA-control but we will soon be ready to do the update.

Sune has been the pioneer in this project and he has made the *really* extensive programming of the FPGAs—together we have discussed the design and worked on implementing the controls in ECS. Detailed information about the programming of the FPGA and much more information can be found in a progress report written by Sune [70]. Here I will instead try to put it into an ECS context.

Currently we have made 10 FPGA-modules. Each of these modules are able to control up to 24 analog channels[5]—so we will be able to control up to 240 analog channels! Both the DACs and the DDSs are programmed with the same FPGA-modules. The really nice thing about using our own modules is that we can generate the digital pulses for the analog modules directly in the FPGAs. This way ECS only needs to tell the FPGAs what the output should be like—then the FPGAs generate the digital pulses for the analog modules directly in the hardware. This means that many of the calculations done in ECS are now moved to the FPGAs. This cannot be done with the DIO64-cards since they are designed for general purposes. The FPGA-modules are especially designed for our purposes (and they can do them really well) but they cannot do anything else.

---

[3] Nexys cards from Digilent. These cards are based on a XC3S1000 Spartan3 FPGA.

[4] We could also build the circuit completely from scratch but the Digilent test-boards makes it a bit easier for the electronics department and since they are very cheap we choose this solution.

[5] A DDS takes up the space of four DACs so if e.g. one DDS is used there are only 20 DAC-channels left on the FPGA.

Besides giving us a lot of channels (and releasing ECS from a lot of work) the FPGAs give us a better time resolution since they *always* use a 50 µs update time. The reason is that the FPGA-modules are designed to make linear sweeps (just like the DDS does) and since it is all done in the hardware it does not really matter how often the output should be updated. ECS just sends information about the start and stop values and what the slope of the ramp should be. After this the FPGA takes care of the rest. ECS calculates the "best" fraction N/M representing the slope (like described in section 3.3.3) and the numbers N and M are sent to the FPGAs together with a start time and start and stop values. ECS also splits the different ramps up into small sections and the system is designed in such a way that the ramping time will never be more than some milliseconds before the slope is updated and because of this (and many other small tricks) the output will always represent a given arbitrary ramp almost perfect.

Since we now will have FPGA-modules placed around the laboratory we need a global clock to synchronize all the output. We are going to use a 10 MHz external clock which we currently take from our microwave generator. Perhaps we will buy a dedicated external 10 MHz frequency standard/oscillator in the future to get an even better reference. The 10 MHz clock will be distributed throughout the whole laboratory through cables. The clock is multiplied by a factor of four in the FPGAs to provide a 40 MHz system clock for the FPGAs. Currently we do not know if we want to make our own FPGA-modules for the digital channels. It will be much easier to make FPGA-modules for the digital channels however here we cannot really gain anything with our own modules so we will probably just continue to use the DIO64 cards to control the digital channels (these cards should of course also use the external clock). We will then get 112 digital channels instead of the current 48 since we now can use both DIO64-cards for the digital channels.

Looking further into the future we have prepared the FPGA-modules with BNC-input connectors (and possibilities for putting in Analog-to-Digital Converters (ADCs) on these input channels). This way one could imagine to make e.g. PID-circuits directly in the FPGAs. We could e.g. connect the detector signal from the optical lattice detectors controlling the power in the lattices to one of these inputs. The FPGA could then receive the wanted detector signal from the computer and try to get this signal on the detector using the PID-circuit and providing output to the AOM controlling the power. This would allow us to change and adjust the circuits really fast and we could make exceptions and other things which can only be done using digital electronics.

Four

Camera control system

4.1 Introduction

The Camera Control System (CCS) is the other main part of the computer system that controls the laboratory. In this chapter I will start by presenting the main features of CCS in section 4.2. In section 4.3 I will comment on the hardware connected to CCS. After this I will describe the CCS software in section 4.4 and make a few comments on how to get a good image quality in section 4.5. In the end of this chapter I will comment on how to focus the imaging systems (section 4.6).

I have tested and implemented the equipment described in this chapter and written the CCS software to control it.

4.2 Background

The camera system used in quantum gas experiments needs to be able to grab series of images (RAW, REF and BCK images) and calculate the OD from these images as described in chapter 2. The main challenge related to this is that the timing of the cameras has to be really good since we normally use 50 µs imaging pulses and a short exposure time (a few hundred microseconds) or/and a short shutter time (a few milliseconds) of the cameras. The reason why it is important to use a short exposure/shutter time is that we want to limit the amount of background light in the images but of course we still want to capture the full imaging pulse. Besides having a good timing of the cameras we also want the cameras to be able to grab several images within a small time window (e.g. two images within 50 ms). This normally limits the amount of interference in the images which we will return to in section 4.5. Besides these demands (which is needed during the real BEC-experiments) it is also very useful if the cameras are able to stream video images. This functionality is e.g. useful for looking at the laser beams (imaging and optical lattice beams) or looking at the atoms in the MOT. These demands may seem very simple but it can really take a long time to get the software and hardware working the exact right way.

An important question is why we bother to write our own software since these cameras are normally shipped with some standard software. Even though it might be possible to use this software it would be very inconvenient since the software provided with the cameras is not designed for BEC-experiments and also it would not be compatible with ECS. For these reasons we would have to use all our time on setting up the cameras, grabbing and saving images manually.

The main features of CCS

1. **Simultaneous imaging:** CCS is able to grab images on four different cameras at the same time
2. **Stream video-images:** All the cameras connected to CCS can be used in a video-mode
3. **Image analysis:** We have different basic analysis tools which we can use in the analysis of the grabbed images
4. **TIFF-compatible:** CCS supports 16-bit gray scale TIFF-images
5. **Remote controlled:** Normally CCS is remote controlled by ECS
6. **Optical lattice alignment:** CCS is able to control motorized mirrors which are used in the alignment procedure of the 3D optical lattices

We wanted CCS to be able to grab images on different cameras at the same time. Normally we only grab images on one axis at a time because each of the imaging beams heat up the atomic cloud. However the ability to grab images on different axes at the same time can e.g. be used when a second species is implemented. By using two imaging beams (resonant with rubidium and the second species, respectively) it will be possible to grab an image of both kind of atoms at the same time (by using two different axes/cameras). In a test setup CCS was able to grab images of the same 25 µs imaging pulse on four different cameras at the same time (still with a very short exposure time on all the cameras).

We have decided not to implement the image fitting software in CCS (the image analysis is instead done in the fitting program). This could of course be done but it would be like having two different programs put together in one. We would still have to make it in such a way that we would be able to analyze old images without being disturbed when new images arrive from the cameras. However as indicated in point three CCS has some basic features which we use when we are just looking at the images. We can e.g. zoom in on the clouds, change the OD scale, view and compare the last 20 grabbed image series, measure distances in the images (in pixels) and see the exposure amount. Besides that we have some features that can be used in connection with the optical lattices—e.g. a feature that can be used to calculate the so-called visibility.

4.3 Hardware

A schematics of the CCS hardware setup can be seen in figure 4.1. Currently we have up to 7 cameras connected to the CCS computer at the same time. They are from three different manufacturers—Andor Technology, DTA and Point Grey Research (PGR). The different cameras (including three identical PGR Scorpion cameras) are listed in table 4.1.

The Andor cameras are by far the best cameras in our setup. They are both Electron-Multiplying Charge Coupled Device (EMCCD) cameras and they have a high quantum efficiency. The DU-888 is a back-illuminated camera with a quantum efficiency of more than 70% and the DU-885 has a quantum effciency of around 50% (both at 780 nm). The quantum efficiency is even better (up to more than 90% for the DU-888) for shorter wavelengths which would be the case if we implemented lithium as the second species[1]. The

---

[1] Because of the great sensitivity of the Andor cameras we are able to grab fluorescence images. With twice the amount of light we normally use in absorption imaging we are able to grab nice fluorescence images of thermal clouds and BECs.

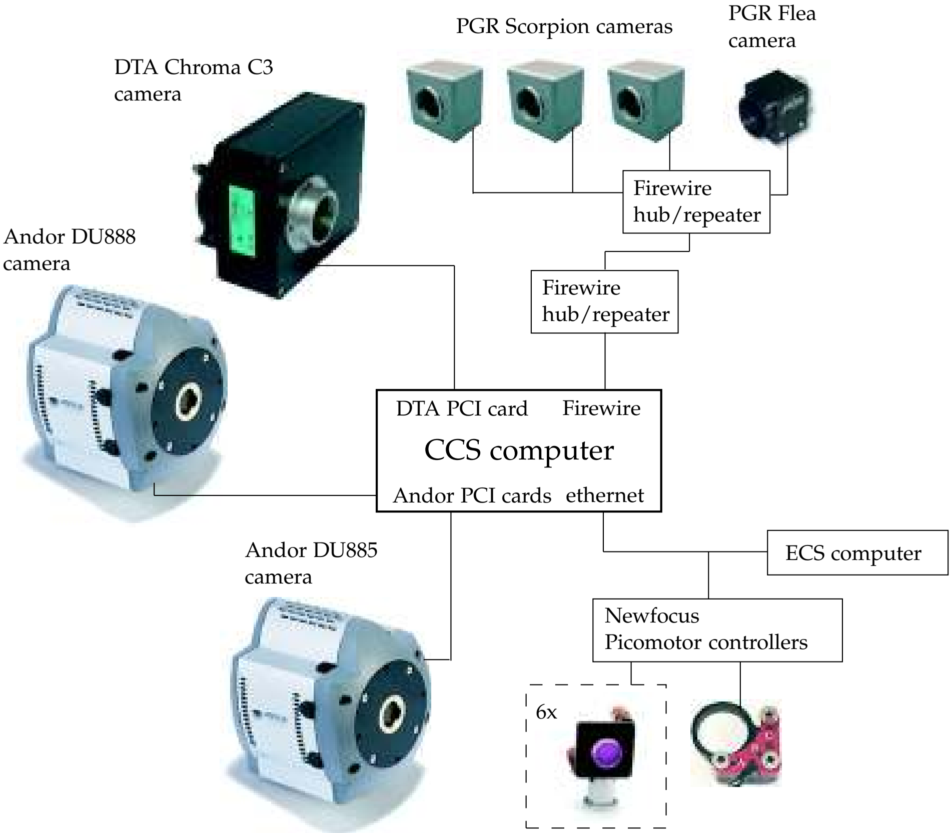

Figure 4.1: A schematic overview of CCS. Besides controlling all the cameras CCS also controls motorized mirrors. These mirrors are used in the alignment procedure of the optical lattices. The CCS computer is connected to the ECS computer through an ethernet connection.

camera name	pixel size (µm)	chip-size (pixels)	ADC
Andor iXon^{EM}+ DU-885 (x-axis)	8x8	1004x1002	14 bit
Andor iXon^{EM}+ DU-888 (y-axis)	13x13	1024x1024	14 bit
DTA (z-axis)	6.8	2184x1472	14 bit
PGR Flea (MOT camera)	4.65	1024x768	8 bit
PGR Scorpion (x-axis and misc)	4.8	1600x1200	12 bit

Table 4.1: The different cameras connected to the CCS computer.

two Andor cameras are our primary cameras—currently they are mounted as our x- and y-axis cameras. They are both cooled to $-80\,°\mathrm{C}$ to reduce dark noise/current and they have a 14-bit ADC. Because of these cameras we decided to buy a fast SCSI-harddrive which can be used if many images have to be grabbed and stored fast—however so far we have not used this feature.

The DTA camera (manufactured by DTA) is our old primary camera. It has now been moved to the z-axis. It has a 14 bit ADC and the CCD-chip is cooled to about $-10\,°\mathrm{C}$ to reduce dark current. It should be mentioned here that Jesper has done much of the initial testing and programming of the DTA camera.

CCS is also able to control up to 9 PGR cameras (manufactured by *Point Grey Research*). These cameras are not cooled but have an electronic shutter which makes shutter times down to 50 µs possible. The signal-to-noise ratio and the quantum efficiency is much better for the DTA and especially for the Andor cameras than for the PGR cameras but the PGR cameras are fine for grabbing images of the atoms in the MOT (the Flea camera is used for this purpose) or for large thermal clouds in the science chamber (one of the scorpion cameras is used for grabbing images of thermal clouds in the science chamber along the x-axis).

We found severe bugs in the drivers from all three manufacturers that provided us from setting up and running the cameras exactly the way we wanted. In the end we managed to convince all companies that they should fix them (and we got new drivers) but we have really used a lot of time dealing with annoying problems and proving to the companies that their software/hardware was not working the way it was intended.

Motorized mirrors from New Focus are connected to CCS through an ethernet connection. With the help of the CCS software these mirrors make it easy to align the 3D optical lattice. We have 6 Model 8807 1" mirror mounts and 1 Model 8852 2" mirror mount. They are connected to 5 Model 8753 driver module which again is connected to a Model 8752 ethernet controller module. The functionality of this setup will be described in the following section.

4.4 CCS software

The CCS software consists of just below 10,000 lines of source code and it is written from scratch in Borland Delphi 7 professional.

CCS is a “stand-alone” program but during an ECS run CCS is remote controlled by ECS through an ethernet connection and thereby the two systems more or less act as one in some respects. However even though the CCS computer is controlled by ECS it does not mean that the user cannot work in CCS while waiting for the new images. The user can e.g. look through old images in CCS and in general use most of the other features in CCS in between the runs (only a small blinking lamp indicates that a run is in progress). However when the run is finished the new images arrives to the screen. CCS receives a start signal from the ECS computer and detailed information about the current run (as already described in some details in chapter 3). This information is among other things used to initialize the different cameras. It is possible to overrule the commands from ECS but we never use this feature. ECS and CCS automatically check if the other program is open and connect to each other if possible. This means that the only thing the user has to do when starting up the experiment is to open ECS and CCS, load an *.ecs file in ECS and press “go” in either ECS or CCS. Then everything (including the cameras) are initialized and the run is executed.

When CCS finishes a run all the grabbed images are stored on the harddrive together with a *.ccs file containing information about the given run. Normally we grab four images in each run because we grab two background images—one for the RAW-image and one for the REF-image. However in principle we could grab up to 4×4 images in every run since we have four different axes (namely the x-, y-, z- and aux-axis). CCS stores the grabbed images in a 16 bit gray scale Tagged Image File Format (TIFF). Initially we thought about making our own file format but it is nice to be able to open the images in a normal graphics editing program (even though many of these program does not support 16 bit gray scale images). Since we have 8, 12 and 14 bit cameras and since the image format is 16 bit the data bits are shifted so only the *most* significant 8, 12 or 14 bits are used instead of the least significant bits. In this way the images can be loaded from a graphics editing program without being very dark and for the data analysis it does not matter where the bits are located—the OD density gives the same if all the images are shifted/multiplied with a constant.

On a typical day in the laboratory we run 100-300 runs corresponding to 400-1200 images (each of these images take up a few megabytes of space). All the images and the *.ccs file are automatically saved with the same time stamp as the *.ecs file. Reload functions are implemented in ECS and CCS that make it easy to reload old runs just by knowing the date and the run number.

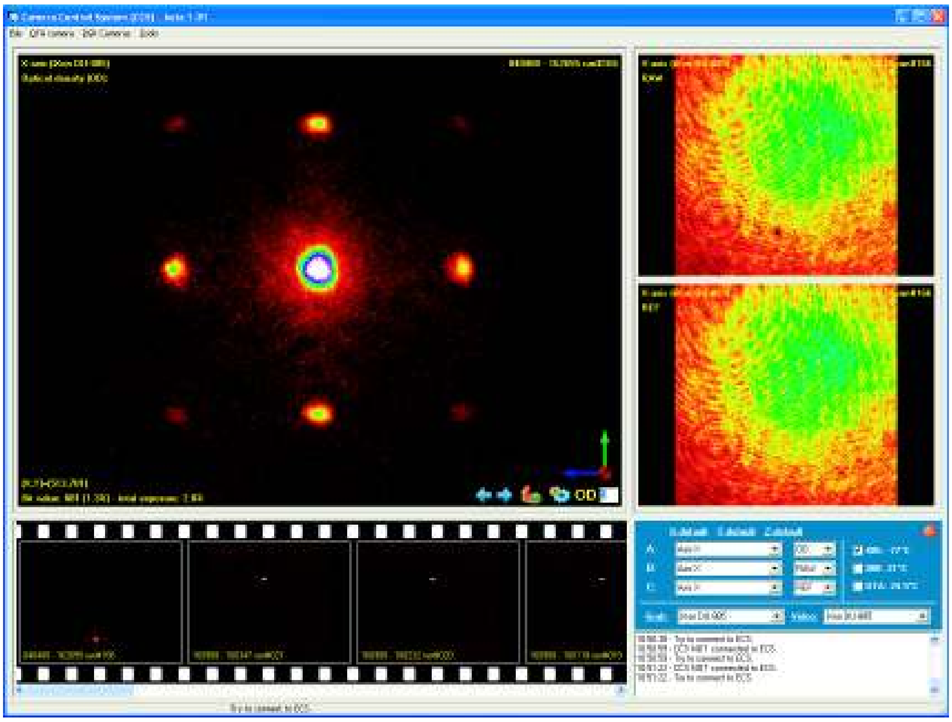

Figure 4.2: The main window of CCS. Here three images are shown at the same time. Normally we show the OD, the RAW and the REF image for the current camera. However if more cameras are used at the same time we can also chose to see e.g. the OD for each of the cameras. In the film-strip we can see the last 20 image series grabbed (or loaded) and if we press one of these images the complete run is brought back into the three main images with the same e.g. zoom-settings as used in the current viewing. When CCS is used for aligning the optical lattice the two right images and the film-strip is replaced with aligning features (see e.g. figure B.4 and figure B.5).

The cameras can be used in three different modes: normal mode, capture mode, and video mode. In the *normal mode* CCS (and the cameras) is triggered by the ECS-computer. This mode is always (and only) used in connection with an ECS run. In this mode ECS tells CCS which cameras should be used, how many images that should be grabbed and so on. From these informations CCS is able to set up all the cameras in a “wait for trigger” mode (and able to do it again during a run). All the cameras are triggered by normal ECS digital channels. In *capture mode* CCS just grabs an image on the selected camera—in this mode the settings used for the camera are defined in CCS. In *video mode* the cameras stream video images. Also here the camera settings are defined locally in CCS.

When the different cameras are awaiting a trigger pulse they normally lock the software, however CCS uses different so-called “threads” (CCS is a multi-threaded application). We have one thread for each of the four axes/cameras and a thread for the rest of the CCS software. If a camera blocks one of the threads it will not affect the others. Moreover, threads are also used in the video mode which makes it possible to stream video with a high frame rate and at the same time use the other features in CCS. If not all of the trigger pulses are received (or other errors happen during a run) CCS just terminates the relevant threads and tells the user than an error has occurred.

4.4.1 Motorized mirror mounts

The seven motorized mirrors connected to CCS are used to move the optical lattice beams. These mirrors can be controlled by CCS in two different modes:

- The first mode is used when we stream video images of one of the optical lattice beams and want to move this beam on the screen.

- The second mode is used when we are looking at an atomic cloud trapped in the lattice beams and want to move this cloud to ensure a good overlap with the magnetic trap center.

The first mode is simple. If we e.g. are using the x-axis camera CCS knows what mirror (and motor) it has to move in order to move the x-lattice beam left, right, up or down on the screen. The movement can be done just by pressing the associated arrows in CCS (see e.g. figure B.4). This method is used in the rough alignment procedure of the optical lattice beams.

The second mode is used when we want to do the fine alignment of the optical lattices based on the position of the atomic cloud (if the optical lattice beams are right aligned the atomic cloud should be located at the center of the magnetic trap also when the cloud is only trapped in the optical trap). In this case we cannot move the atomic cloud up/down and left/right just by moving one mirror. Instead the up/down movement is controlled by one mirror and the left/right movement is controlled by another. Again CCS knows which axis we are using and we can just press the right arrow if we want the atomic cloud to move right. If we are aligning a retro-reflected beam we have to tell CCS that we are aligning the retro-reflector but otherwise it is the same. CCS also makes a 1D fit (in both directions) to the beam or the atomic cloud. In the case where we are looking at the atomic cloud CCS can tell us in which direction(s) we should move the cloud since it knows where the center of the magnetic trap is located along the different axes. For this reason we can align more degrees of freedom at the same time. Detailed information about the aligning procedure of the optical lattices can be found in appendix E, and a screen dump of this feature can be seen in figure B.5.

4.5 Image quality and kinetics mode

Since we are using monochromatic imaging light, interference can be a big issue. We can live with an interference pattern if it does not move between the RAW and the REF image because then the OD calculation will "remove" it. However, the interference pattern (and the laser intensity) will often change in time—e.g. movable dust and moving/vibrating optics can change the interference pattern in time. Of course we try to keep everything as clean as possible but it is our experience that it is more important to address the following issues to get a good image quality:

- Stable setup
- Fast imaging

To avoid vibrations we have designed good mounting for the cameras and generally we try to mount everything as good as possible. Also we turn off the air cooling (fans) on the cameras a few seconds before we grab the images. Better mounting and reducing the vibrations by turning off fans have improved our image quality dramatically.

Another way to improve the image quality can be to use a fast camera. If the interference moves on a large time scale compared to the time interval between the different images this will help. We normally run the cameras in the so-called kinetics mode (also sometimes referred to as "frame transfer"). The kinetics mode makes it possible to grab images with a few ms time interval since the camera does not have to read out the data before all the images are grabbed (of course there is a limit to how big and how many images we can grab in this mode in one sequence). We are typically using a 30 ms–50 ms time interval between the RAW- and the REF-image which reduces the interference "noise" in the OD images much compared to when we use a longer time interval. We could make the time interval between the RAW- and the REF-image even smaller. However the problem is that we need to wait for all the atoms to disappear from the field of view before we grab the REF-image. On the y-axis the atoms hit the bottom of the vacuum chamber after approximately 60 ms. However since we are heating up the atomic cloud with the RAW-beam the atoms will be gone earlier.

The Andor and the DTA cameras can be used in kinetics mode. The idea is that only a part of the CCD-chip is exposed when an image is grabbed—the rest of the chip is blocked behind a mask. In our case half of the CCD-chip is blocked. This means that when we grab an image this image (or these pixels) can be shifted fast to the region behind the mask and the pixels originally located behind the mask are just dumped. This procedure only takes a few milliseconds since no data needs to be read out. Now another image can be grabbed and both images can be read out (which is a slow process)—see figure 4.3. Since the REF-image is not shifted into the mask area fast during the read out the REF-image will typically contain more background light. For this reason we use two different background images and in general it is very important that the sequences are *exactly* the same when we grab the RAW/REF images as when we grab the two background images except that the laser light should (of course) not be on when the background images are grabbed. However the imaging shutters should e.g. also be opened and closed (as normal) during the background imaging to account for scattered light through the imaging AOM when it is off.

The Andor cameras are built for the kinetics mode—half of the approximately 2000 × 1000 pixels CCD-chip is blocked from the factory leaving an approximately 1000 × 1000 pixels area for imaging. The DTA camera support the kinetics feature but here we have to make the mask ourselves. Instead of placing the mask directly on the CCD-chip we place it in an image plane (see e.g. figure 7.7 in chapter 7).

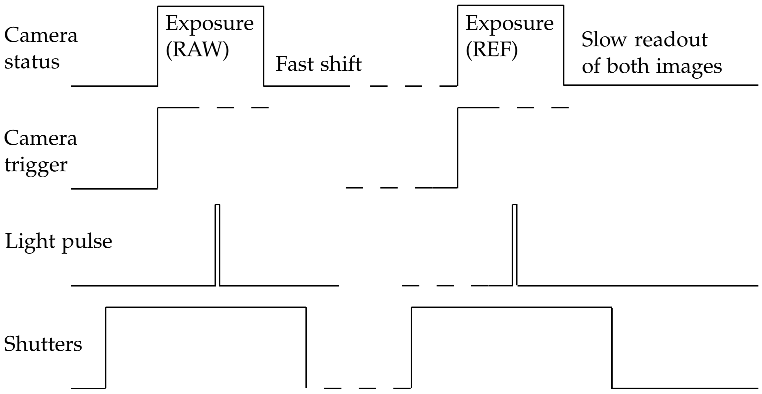

Figure 4.3: A sketch showing the imaging sequence when we grab the RAW and the REF images. When we grab the two background images the sequence is exactly the same except that the light pulse is off (using the AOM). The second trigger pulse will be missed if it is placed too close to the first trigger pulse because the camera needs to have time to shift the pixels before the next image can be grabbed. The shutters are slow and need a few ms to open and close so they are effectively opened for much more time than the duration of the light pulse (50 µs). The exposure time is normally chosen to be a few hundred microseconds to be sure that the light pulse hits the exposure interval (before the pixels start to shift). After the exposure time has elapsed the camera shifts (or reads out) the images. Up until the first trigger pulse the camera should perform "keep clean cycles" to make sure that there is no "noise" on the CCD-chip before it is exposed.

4.6 Focusing the imaging system

Since we want to have a high resolution of the atomic clouds the depth of focus is only a few hundred micrometers so we need to have a good way to focus the imaging system. One way to do this is by using the fact that the refractive index changes across an atomic resonance [71]. If the imaging beam is red detuned the refractive index is above unity. However if the imaging beam is blue detuned the refractive index goes below unity.

Figure 4.4 gives a qualitative idea of the consequences of this. In the case of red detuning ($n > 1$) the atomic cloud will act as a converging lens for the light—the light that hits the atomic cloud will be focused onto a single spot (see figure 4.4). If the camera is focused on a plane in front of the atoms (between the atomic cloud and the camera) there will be a bright spot on the RAW-image and correspondingly a dark spot in the OD-image (if the imaging beam is red detuned). If the camera is focused on a plane behind the atomic cloud it will look like all the light that hits the atoms is coming from a larger area—so we get a dark spot on the RAW-image and a bright spot on the OD-image. The same analysis can be made in the case of blue detuning. The precision of this focusing method is below 100 µm.

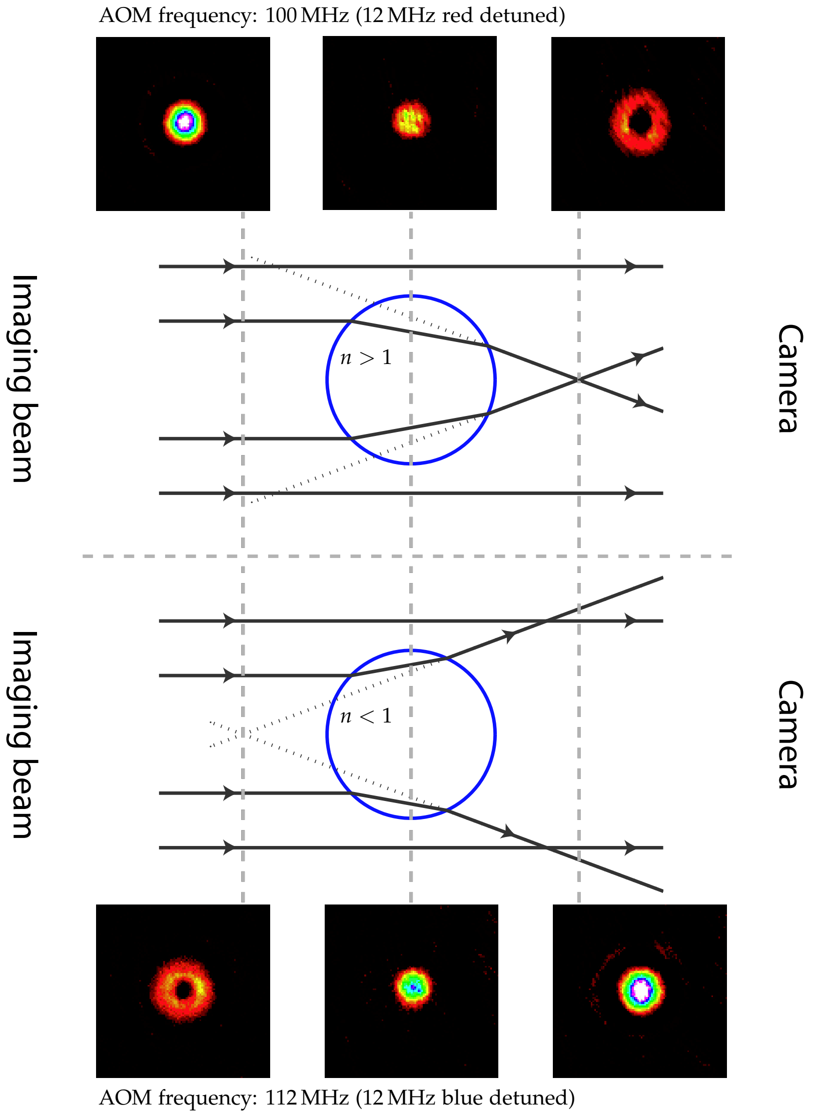

Figure 4.4: A sketch showing the different images we get when the imaging system is focused on different planes. The drawings indicate the influence of an atomic cloud in the beam path. The atomic cloud (blue circle) can be assigned a refractive index that depends on the detuning. Red detuned ($n > 1$) gives a bright spot in the center of the RAW image (corresponding to a dark spot in the center of the OD-image shown here) when the camera is focused on a plane in front of the the atomic cloud. The other case can be analyzed the same way. The images are 2 mm defocused (in both directions). More information is given in the text. Based on 200508 run 6–24.

Five

BEC results

5.1 Introduction

In this chapter I will present different experiments performed with the BEC-setup. In section 5.2 I will go through some basic theory which is needed in the rest of this chapter. After this I will present different techniques for doing fundamental characterization of the setup. In section 5.3 I will present different methods used for calibrating the pixel size of the imaging system. Section 5.4 contains three different ways of measuring the magnetic trap frequencies of the QUIC-trap. The pixels size and the magnetic trap frequencies can e.g. be used in the determination of the atom number which is discussed in section 5.5. In the end of this chapter (section 5.6 and 5.7) different BEC images are analyzed.

Most of the measurements presented in this chapter was conducted together with Sung Jong.

5.2 Background

5.2.1 Optical density and absorption cross-section

We saw in section 2.2 (eq. (2.2)) how the OD can be calculated from three grabbed images (or four images in the case where we have two different backgrounds as described in the previous chapter)

$$\mathrm{OD} = \ln\left[\frac{I_{ref} - I_{bck}}{I_{raw} - I_{bck}}\right]. \tag{2.2}$$

However the optical density is also given by

$$\mathrm{OD}(x,y) = \sigma \int_{-\infty}^{\infty} n(x,y,z)dz, \tag{5.1}$$

where $n(x,y,z)$ is the particle density, σ is the absorption cross-section (which is assumed constant—independent of x, y and z here) and z is the propagation direction of the imaging beam. Throughout this chapter it is important to notice that the x, y and z directions *not* necessarily are the same as the x, y and z axis defined for our QUIC-trap. Here z is just the imaging axis and x and y are the horizontal and vertical directions in the grabbed images.

The OD can easily be integrated to give the atom number N

$$N = \frac{1}{\sigma} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathrm{OD}(x,y) dx dy. \tag{5.2}$$

In a real experiment the integral in eq. (5.2) will be a sum over the camera pixels

$$N = \frac{\Delta_x \Delta_y}{\sigma} \sum_i \sum_j \mathrm{OD}(x_i, y_i) \tag{5.3}$$

where Δ_x and Δ_y are the pixel size of the imaging system (typically a pixel corresponds to a few µm).

In general σ depends on the detuning, the light intensity and the light polarization [72]

$$\sigma = \frac{\sigma_0}{1 + 4(\Delta/\Gamma)^2 + (I/I_{sat})} \tag{5.4}$$

where σ_0 is the on-resonance cross section (which e.g. depends on the light polarization), Δ is the detuning of the imaging light, Γ is the natural linewidth of the optical transition, I is the intensity of the imaging light and I_{sat} is the saturation intensity. Normally we try to be far below the saturation intensity which means that we can neglect the I/I_{sat} term.

Our fitting program allows us to make the sum in eq. (5.3) directly in a pre-defined area and thereby determine the atom number if we know σ, Δ_x and Δ_y. However more often we fit the OD image to different models (or fit-functions) depending on the temperature of the cloud. This way we can extract the atom number but also other quantities from the fit parameters as we will see below.

5.2.2 Thermal cloud

When the magnetic (and/or optical) trap is turned off (to time $t = 0$) the atomic cloud will expand and fall in the field of gravity. If collisions and gravity are neglected the position of each individual atom will evolve in time t as

$$\mathbf{r}(t) = \mathbf{r}_0 + \mathbf{v}t \tag{5.5}$$

where $\mathbf{r}_0$ is the initial position of the atom in the trap and $\mathbf{v}$ is the velocity of the atom. This means that after a long TOF t we are actually measuring the momentum distribution of the atoms (and not the spacial distribution).

In the following we will look at an ideal Bose gas trapped in a harmonic potential (e.g. the QUIC-trap)

$$V(x,y,z) = \tfrac{1}{2} m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2) \tag{5.6}$$

which is turned off to $t = 0$. Here m is the mass of an atom and ω_i are the trapping frequencies. For our QUIC-trap we have a radial trap frequency $\omega_r = \omega_x = \omega_y$ and an axial trap frequency $\omega_a = \omega_z$.

When the temperature approaches the critical temperature T_c the effects of quantum (Bose) statistics on the density distribution sets in. In this case the density distribution (as a function of TOF t) is given by [1]

$$n_{bose}(x,y,z,t) = \frac{1}{\lambda_{dB}^3}\sqrt{\frac{1}{1+\omega_x^2t^2}}\sqrt{\frac{1}{1+\omega_y^2t^2}}\sqrt{\frac{1}{1+\omega_z^2t^2}} \\ g_{3/2}\left[\exp\left(\frac{\mu}{k_BT}-\frac{x^2}{w_x(t)^2}-\frac{y^2}{w_y(t)^2}-\frac{z^2}{w_z(t)^2}\right)\right] \quad (5.7)$$

where $\lambda_{dB} = (2\pi\hbar^2/mk_BT)^{1/2}$ is the De Broglie wavelength, μ is the chemical potential, k_B is the Boltzmann constant, T is the temperature, g is the Bose function (also called the polylogarithm function)

$$g_j(\chi) = \sum_{i=1}^{\infty}\frac{\chi^i}{i^j} \quad (5.8)$$

and $w_i(t)$ are the widths of the cloud given by

$$w_i(t) = \sqrt{\frac{2k_BT}{m\omega_i^2}+\frac{2k_BT}{m}t^2} = \sqrt{\frac{2k_BT}{m\omega_i^2}\left(1+\omega_i^2t^2\right)}. \quad (5.9)$$

We can use eq. (5.1) to get the OD (this integral can directly be made using the definition of g—eq. (5.8))

$$\mathrm{OD}_{bose}(x,y,t) = \frac{\sigma}{\lambda_{dB}^3}\sqrt{\frac{1}{1+\omega_x^2t^2}}\sqrt{\frac{1}{1+\omega_y^2t^2}}\sqrt{\frac{1}{1+\omega_z^2t^2}} \\ w_z(t)\sqrt{\pi}g_2\left[\exp\left(\frac{\mu}{k_BT}-\frac{x^2}{w_x(t)^2}-\frac{y^2}{w_y(t)^2}\right)\right]. \quad (5.10)$$

Notice here that we get something with g_2 instead of $g_{3/2}$ as for the density distribution (eq. (5.7))[1]. The number of atoms can be calculated using eq. (5.2)

$$N_{bose} = \frac{w_x(0)w_y(0)w_z(0)}{\lambda_{dB}^3}\pi^{3/2}g_3\left[\tilde{z}\right] \quad (5.11)$$

where $\tilde{z}$ is the fugacity ($\tilde{z} = \exp(\mu/k_BT)$). Looking at eq. (5.10) we are inspired to define a fit-function of the following form

$$f_{bose}(x,y) \equiv a_1 + a_2g_2\left[a_7\exp\left(-\left(\frac{x-a_3}{a_4}\right)^2-\left(\frac{y-a_5}{a_6}\right)^2\right)\right] \quad (5.12)$$

where a_1 is an offset, a_3 and a_5 are the center positions of the cloud, a_4 and a_6 are the widths of the cloud (given by eq. (5.9)—the widths depend on t), a_7 is the fugacity $\tilde{z}$ and a_2 is the amplitude given by

$$a_2 = \frac{\sigma}{\lambda_{dB}^3}\sqrt{\frac{2k_BT}{m\omega_z^2}}\sqrt{\pi}\sqrt{\frac{1}{1+\omega_x^2t^2}}\sqrt{\frac{1}{1+\omega_y^2t^2}}. \quad (5.13)$$

---

[1] The g_2 function can just be evaluated numerically with a small script summing e.g. the first 10.000 terms of eq. (5.8) (it converges nicely when χ is between 0 and 1) or it can also be calculated with the built in PolyLog[j=2,z] function in Mathematica. In the fit-program we use a lookup table of g_2 to decrease the computation time.

The number of atoms can generally be expressed in terms of a_2, a_4, a_6 and a_7 as

$$N = \frac{a_2 a_4 a_6}{\sigma} \pi g_3\big[a_7\big]. \tag{5.14}$$

The atom number is proportional to the amplitude and the two widths. Also note that this result is time independent (even though a_2, a_4 and a_6 depend on time the product does not) which of course should be the case since we do not lose atoms in this model—the cloud is just reshaped in time.

The temperature can be found by grabbing images to different TOFs t. If the width squared (w_a^2 or w_r^2) is plotted as a function of $2k_B t^2/m$ it gives a straight line with the temperature T as the slope (see eq. (5.9)).

Instead of basing the temperature calculation on several runs we can also base it on a single run if we have a large TOF. If we isolate the temperature in eq. (5.9) we get

$$T = \frac{m w_i(t)^2}{2k_B}\left(\frac{1}{1/\omega_i^2 + t^2}\right) \approx \frac{m w_i(t)^2}{2k_B t^2}. \tag{5.15}$$

For large t ($t \gg 1/\omega_i^2$) the atomic cloud has expanded so much that it is much bigger than the in-trap size. In this case the initial in-trap size (for t=0) can be neglected and the approximation made in eq. (5.15) is valid. The radial width (the smallest) should be used in the calculation of T.

Even though the theory presented above is valid for temperatures all the way down to T_c we currently do not fit to eq. (5.12) directly since it can be hard to get stable values for a_2 and a_7. Instead we look at the low and high temperature limits.

Low temperature limit

Below (or near) T_c the chemical potential μ can be set equal to zero. In this case we can set $\tilde{z} = 1$ and we fit to eq. (5.12) with a_7 locked to 1. This means that the Bose function implemented in the fit-program (with $a_7 = 1$) should only be used in the low temperature region (e.g. in the bimodal fits where the cloud is partly condensed).

In this case the atom number is given by

$$N_{bose} = \frac{a_2 a_4 a_6}{\sigma} \pi \mathrm{Zeta}[3] \quad \text{(for } a_7\text{=1)} \tag{5.16}$$

where Zeta$[x]$ is the Riemann zeta-function (Zeta$(3) = g_3(1) \approx 1.2021$).

High temperature limit

In the high temperature regime where μ is large and negative ($|\mu| \gg k_B T$ and $\mu < 0$) we have that $\tilde{z} \ll 1$. This means that $g_j(\tilde{z}) \approx \tilde{z}$. In this case the density distribution in (5.7) reduces to a Gaussian distribution

$$n_{gauss}(x,y,z,t) = n_0(t)\exp\left(-\frac{x^2}{w_x(t)^2} - \frac{y^2}{w_y(t)^2} - \frac{z^2}{w_z(t)^2}\right) \tag{5.17}$$

where the central density $n_0(t)$ is given by

$$n_0(t) = \frac{1}{\lambda_{dB}^3}\sqrt{\frac{1}{1+\omega_x^2 t^2}}\sqrt{\frac{1}{1+\omega_y^2 t^2}}\sqrt{\frac{1}{1+\omega_z^2 t^2}}\exp\left(\frac{\mu}{k_B T}\right). \tag{5.18}$$

It is expected that the distribution in eq. (5.17) should be Gaussian since the Maxwell-Boltzmann velocity distribution (which is valid for $T \gg T_c$) is Gaussian. This means that no matter what the in-trap spatial distribution is the cloud will approach a Gaussian shape as it expands. However for a harmonic potential (e.g. the QUIC-trap) the in-trap spatial distribution is also Gaussian (which also can be seen from eq. (5.17))[2]. The OD can be calculated using eq. (5.10)

$$\mathrm{OD}_{gauss}(x,y,t) = \sigma n_0(t)\sqrt{\pi} w_z(t) \exp\bigg(-\frac{x^2}{w_x(t)^2} - \frac{y^2}{w_y(t)^2} \bigg). \tag{5.19}$$

From this we are inspired to define a fit-function for hot thermal clouds ($T \gg T_c$)

$$f_{gauss}(x,y) \equiv a_1 + a_2 a_7 \exp\bigg[-\Big(\frac{x-a_3}{a_4}\Big)^2 - \Big(\frac{y-a_5}{a_6}\Big)^2 \bigg]. \tag{5.20}$$

We cannot distinguish a_2 and a_7 in this case so $a_2 a_7$ should be considered as a single fit parameter here. In this case the atom number is given by

$$N_{gauss} = \frac{a_2 a_4 a_6}{\sigma} \pi a_7. \tag{5.21}$$

5.2.3 BEC region

In the very low temperature limit a BEC is formed and in this case we need to solve the Gross-Pitaevskii equation to find the density distribution of the BEC. Let us write up the Gross-Pitaevskii equation presented in chapter 1 again [4]

$$\bigg(-\frac{\hbar^2 \nabla^2}{2m} + V(\mathbf{r}) + gN|\psi(\mathbf{r})|^2 \bigg)\psi(\mathbf{r}) = \mu\psi(\mathbf{r}) \tag{5.22}$$

where $N-1 \approx N$ and the interaction parameter g is given by $g = \frac{4\pi\hbar^2 a}{m}$, where a is the s-wave scattering length and m is the mass of an atom.

When the interactions are repulsive and the number of atoms is large (in the Thomas-Fermi (TF) approximation [4]) we can neglect the kinetic energy term and we are left with a Gross-Pitaevskii equation which is easy to solve. Since it is a non-linear equation and since we have removed the kinetic energy term (which normally acts on the wave function) we can simply "divide" with the wave function and the result is

$$n_{BEC}(x,y,z) = N|\psi(x,y,z)|^2 = \max\bigg(\frac{\mu - V(x,y,z)}{g}, 0 \bigg). \tag{5.23}$$

Again we will assumed that the BEC is made in a harmonic potential like eq. (5.6) which is turned off at $t=0$. We can now calculate the OD using eq. (5.1)

$$\begin{aligned} \mathrm{OD}_{BEC}(x,y,t) &= \sigma \int_{-\infty}^{\infty} n_{BEC}(x,y,z)dz \\ &= \frac{2}{3}\frac{\sigma}{g}\frac{m}{\omega_z}\Big(\frac{2\mu}{m}\Big)^{3/2}\Big(1 - \frac{x^2}{w_x(t)^2} - \frac{y^2}{w_y(t)^2}\Big)^{3/2}. \end{aligned} \tag{5.24}$$

Here the widths $w_i(t)$ are made time-dependent. The reason is that it can be shown that a BEC will retain the same TF (inverted parabola) shape of the density distribution during

---

[2] For the QP-trap the in-trap distribution is *not* Gaussian—however at least for hot thermal clouds trapped in our QP-traps it is still a very good approximation to use the formulas presented in this high temperature limit [41].

ballistic expansion (Castin and Dum [73]). The widths of a BEC will change as a function of TOF t in accordance with[3]

$$w_i(t) = \lambda_i(t)\sqrt{\frac{2\mu}{m\omega_i^2}} \tag{5.25}$$

where $\lambda_i(t)$ are the so-called scaling parameters. The scaling parameters depend on the trapping frequencies of the trap. However in the elongated QUIC-trap where $\epsilon = \omega_a/\omega_r \ll 1$ the radial scaling parameter is given by [73]

$$\lambda_r(\tau) = \sqrt{1+\tau^2} \tag{5.26}$$

and the axial scaling parameter is

$$\lambda_z(\tau) = 1 + \epsilon^2(\tau \arctan\tau - \ln\sqrt{1+\tau^2}) + O(\epsilon^4) \tag{5.27}$$

where $\tau = \omega_r t$.

We are now inspired to define a fit-function

$$f_{BEC}(x,y) \equiv a_1 + \max\left[a_2\left(1 - \left(\frac{x-a_3}{a_4}\right)^2 - \left(\frac{y-a_5}{a_6}\right)^2\right)^{3/2}, 0\right]. \tag{5.28}$$

Here a_1 is an offset, a_3 and a_5 are the center positions of the cloud, a_4 and a_6 are widths given by eq. (5.25) and a_2 is

$$a_2 = \frac{4}{3}\sigma\frac{\mu}{g}\sqrt{\frac{\mu}{\frac{1}{2}m\omega_z^2}}\frac{1}{\lambda_x(\tau)}\frac{1}{\lambda_y(\tau)} \tag{5.29}$$

In the (elongated) QUIC-trap the scaling parameters λ_x and λ_y are simply given by the eq. (5.26) and eq. (5.27). However in general they should be calculated for the used trap [73]. Also notice that in order to use these scaling parameters for the QUIC-trap it is important *not* to use the axial direction in the QUIC-trap (the Ioffe-axis) for imaging. Along this axis the in-trap shape of the BEC is not elongated (but round)—both the horizontal and vertical direction of the images grabbed along this axis corresponds to the radial direction of the trap. The number of atoms can be calculated using eq. (5.2)

$$N_{BEC} = \frac{2}{5}\pi\frac{a_2 a_4 a_6}{\sigma}. \tag{5.30}$$

5.2.4 Fitting

In table 5.1 the different fit-types implemented in our fitting-program can be seen. We do real 2D fitting of the clouds as indicated by the definitions of the fitting functions. Alternatively one can make 1D fits instead of real 2D fits to reduce the computation time. The data can be reduced to two 1D fits by summing all rows and all columns in an image—notice that if this is done other expressions are needed in the analysis of the images).

---

[3] This expression can be derived from eq. (18) in [73]—using that the width of a parabola ($y = c - ax^2$) is $\sqrt{c/a}$.

Fit type	Application	Equation
2D Gaussian	Thermal ($T \gg T_c$)	(5.20)
2D Bose	Thermal ($T \gtrsim T_c$)	(5.12) with $a_7 = 1$
2D Gauss/TF (bimodal)	Thermal and BEC ($T \sim T_c$)	(5.20) + (5.28)
2D Bose/TF (bimodal)	Thermal and BEC ($T \sim T_c$)	(5.12) with $a_7 = 1$ + (5.28)
2D TF	Pure BEC ($T \ll T_c$)	(5.28)

Table 5.1: The different fit-functions implemented in our fitting program and their applications. Because of the characteristic "bimodal" shape of a partly condensed cloud we can fit to a sum of the expressions for a thermal cloud and a BEC. Besides the 2D fit types presented in this table we can also do 1D Gauss, 1D TF and 1D Gauss/TF fits in the fitting program but we do not really use this functionality since the 2D fits does not require much computation time (only a few seconds).

5.3 Pixel size calibration

It is important to know the pixel size of the cameras since it e.g. enters in the atom number (eq. (5.3)) and in the temperature (through the widths—eq. (5.9)) of the cloud. We do the pixel size calibration in two different ways:

1. Gravity—free fall calibration
2. Optical lattices—interference peaks calibration

Gravity—free fall calibration

Normally we make the pixel calibration by releasing the atoms from the magnetic trap and letting them fall in the field of gravity. We can then track the position of the atomic cloud as a function of TOF and since we know the acceleration due to gravity g the pixel size can be determined. This (of course) only works with cameras mounted in such a way that they look through the science chamber horizontally (perpendicular to the gravity potential). However the calibration can be transfered to a camera looking through the science chamber vertically since we can compare the size of the atomic clouds using different cameras.

Since the magnetic trap does not turn off instantaneously, we have investigated if it influences the free fall of the atomic cloud. Also stray magnetic fields could influence the free fall of the atoms. However it turns out (which will become clear below) that all these effects can be neglected if we just assign a starting velocity b to the cloud—so fitting to an expression looking like this

$$z = \frac{g}{2a}t^2 + bt + c \tag{5.31}$$

where z is the vertical position of the atomic cloud in pixels, $g = 9.82\,\mu\text{m}/\text{ms}^2$ is the acceleration due to gravity and a is the pixel size (earlier named Δ_x and Δ_y—if the pixels (on the camera) are squared then $\Delta_x = \Delta_y$).

Optical lattice—interference peaks calibration

The good thing about using the optical lattices in the pixel calibration is that this method can be used for all directions/cameras because it does not rely on gravity. When a BEC is loaded into an optical lattice adiabatically and the lattice potential is switched off abruptly a matter wave will be emitted from each lattice site. After TOF the phase coherent matter waves will form interference peaks (corresponding to the Fourier transform of the wave function in the lattice). These discrete momentum peaks move out with a constant velocity that depends on the wavelength of the lattice laser (this is described in more details in chapter 6).

The pixel size can be found if the position of a momentum peak is plotted as a function of TOF t and fitted to the following formula

$$z = \frac{2\hbar k}{am} t \tag{5.32}$$

where z is the position of the $2\hbar k$ momentum peak (relative to the center $0\hbar k$ peak—in pixels), $k = \frac{2\pi}{\lambda}$, λ is the wavelength of the lattice laser, m is the mass of a rubidium atom and a is the (unknown) pixel size.

Results

We have made three different measurements based on the "free fall method"—free fall from the QUIC-trap, free fall from an optical trap, and free fall from an optical trap in $m_F = 0$ (a small bias field was applied in this case). Also we have calibrated the pixel size using the interference peaks. In table 5.2 the results can be found. All four measurement are performed with the x-axis imaging system (the x-axis camera is looking through the science chamber horizontally).

We expect the free fall measurement performed with the optical trap in the $m_F = 0$ state to be the most precise measurement since all the magnetic coils were off a long time before the optical potential was turned off in this measurement, and on top of that the atoms were transfered to the $m_F = 0$ state (which to first order is not influenced by stray magnetic fields).

Method	Pixel size (µm/pixels)
Free fall from QUIC-trap	1.947±0.017
Free fall from z-lattice trap	1.948±0.005
Free fall from 2D yz-dipole trap ($m_F = 0$)	1.948±0.005
Interference peaks (1D z-lattice)	1.91±0.05

Table 5.2: Comparison of different methods for determining the pixel size of the x-axis imaging system. Especially in the "free fall from the QUIC-trap" measurement it is important to assign a starting velocity to the atomic cloud because of the trap turn-off—otherwise it gives a value almost 10% off compared to the other results. There is no specific reason why we have used a 1D lattice in one of the measurements and a 2D dipole trap in another measurement—the point is just that they are both optical traps (where the magnetic field has been turned off). In the "interference peaks" measurement the atoms were released from a combined QUIC/1D optical lattice trap. It would be a better strategy to release the atoms from a combined dipole/1D optical lattice potential. Based on 140408 run 9–50, 115–125 and 100408 run 21–37.

Conclusion

From these measurements we now know that it is not very critical which method we use. All the methods are in good agreement and they are all easy to perform.

5.4 Magnetic trap frequencies

We have used three different ways to determine the trapping frequencies of the QUIC-trap:

1. Expansion of a thermal cloud.
2. Parametric heating [74].
3. Oscillating cloud [74].

1. Expansion of a thermal cloud

A very easy way to determine the trapping frequencies of the magnetic trap can be done using a thermal cloud. Using eq. (5.9) we can find the aspect ratio w_r/w_a of a thermal cloud as a function of TOF t

$$\frac{w_r}{w_a} = \frac{\omega_a}{\omega_r}\sqrt{\frac{1+\omega_r^2 t^2}{1+\omega_a^2 t^2}} \tag{5.33}$$

The two unknowns ω_a and ω_r can be found by measuring the aspect ratio as a function of t for a thermal cloud and make a fit to eq. (5.33). The result for the elongated trap where the BEC is created can be seen in figure 5.1. Especially the radial trapping frequency is not determined very precisely this way.

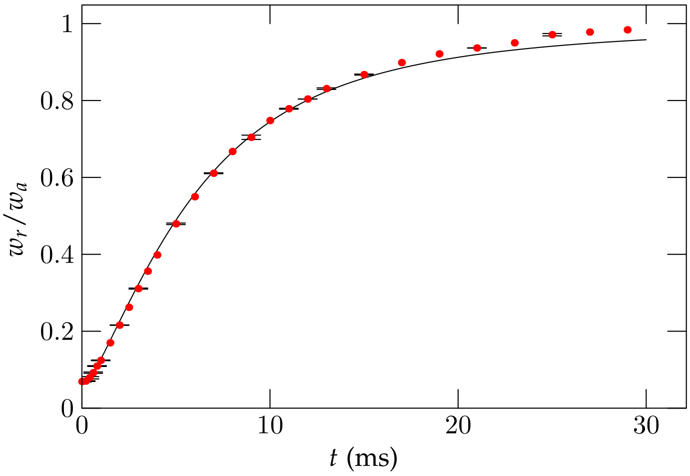

Figure 5.1: Expansion of a thermal cloud. Here the aspect ratio w_r/w_a is plotted as a function of TOF t. The data is fitted to eq. (5.33) and the result is that $\omega_r = 2\pi \times (282\,\text{Hz} \pm 33\,\text{Hz})$ and $\omega_a = 2\pi \times (17.7\,\text{Hz} \pm 0.2\,\text{Hz})$. It can be seen that the aspect ratio of a thermal cloud approaches unity for long TOFs. For a BEC this aspect ratio would go above one. The RF-end frequency for the evaporative cooling set was set to 1000 kHz in these measurements (the RF-bottom was around 757 kHz). When using this method it is important also to have points at small TOFs to make the uncertainty in the determination of ω_r and ω_a as small as possible. Based on 050208 run 147–167 and 180208 run 142–176.

2. Parametric heating

In this scenario the atoms are trapped in the QUIC-trap. By using an auxiliary ac magnetic field we can do resonant excitation (or parametric heating) when the frequency of the ac magnetic field is equal to one of the trapping frequencies—we just shake the trap with a certain frequency and when we hit the resonances the atoms heat up. Since the magnetic field from the QUIC-trap is really deep we use a "RF-knife" (placed at 850 kHz) to remove the heated atoms from the trap. Without the RF-knife we do not see any loss of atoms but only a change in the width of the cloud. This change in width can also be used to find the resonance frequencies but it often gives better results to use the RF-knife and look at the number of atoms instead.

In figure 5.2 an example of a measurement of the radial trapping frequency of the QUIC-trap is shown. Actually we have not tried to measure the axial trapping frequency using this method but it should also work (see also comment about this in the description of the third method).

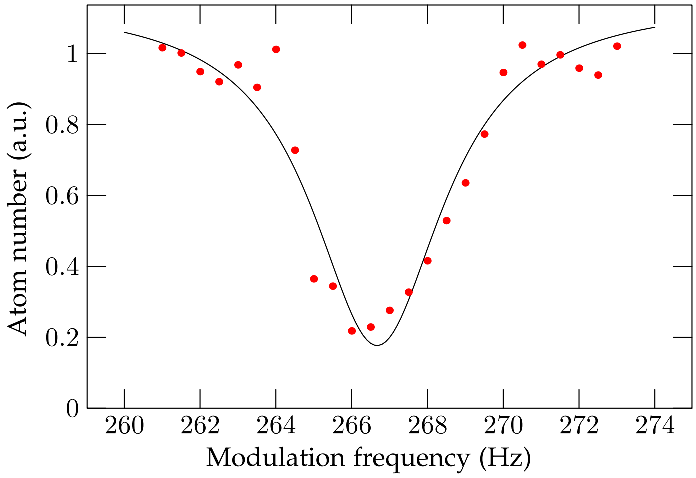

Figure 5.2: A parametric heating experiment. Here we have fitted to an inverted Lorentzian and we get a radial trap frequency of 266.7 Hz ± 0.2 Hz. In this case we used the y-shim coils to provide the trap modulation. When the modulation frequency was equal to the radial trapping frequency atoms were lost from the trap. Based on 280108 run 125–149.

3. Oscillating cloud

Another very precise method and perhaps also a more direct way to do it is by using the method also described in e.g. [74]. Here the axial frequency is measured by monitoring the oscillation of an atomic cloud in the QUIC-trap. The oscillation is achieved by "placing" the atoms off-center with respect to the magnetic field minimum. This means that when the atoms are released they will start to slosh back and forth in the trap. If this movement is large compared to the resolution of the imaging system the trap frequency can directly be monitored just by grabbing images after different "sloshing" times and looking at the position of the atoms. Different approaches can be used to pull the atoms away from the trap minimum. We could use a dipole trap aligned off-center compared to the magnetic field minimum. However since we need to move an optical lattice beam and realign it again it is much easier to use the magnetic coils directly. For the axial trapping frequency we make a linear ramp with the QP-bypass feature which decrease the current in the QP-coils (making the Ioffe-coil effectively stronger). This pulls the atoms closer to the Ioffe coil. After the slow ramp we abruptly change the QP-bypass value back to the initial value (zero) and watch the sloshing. An example where this method has been used can be seen in figure 5.3a.

The radial frequency is much higher than the axial frequency and it is often stated that this cannot be measured the same way. This explains why some are using method 2 for the radial trap frequency and method 3 for the axial trapping frequency measurement. However we tried to do it anyhow and (at least in our setup) it is not a problem to resolve the movement at least up to around 300 Hz. Of course we need to be careful not to pull the atoms too far "up hill" since this would resolve in large oscillations which is a problem since the trap becomes more and more anharmonic the further we move away from the trap center. However in figure 5.3b it can be seen that we can easily resolve the movement even though the oscillation amplitude is only around 4 pixels (or $< 15\,\mu\text{m}$). The perturbing ramp was performed with the x-shim coils in this case.

Conclusion

All the methods give results in good agreement. However method three stands out—it is just very nice and easy to do. After checking that we were able to get the same results using these different methods we decided always to use the oscillating cloud method for this kind of measurements.

In table 5.3 it can be seen how the trapping frequencies changes as a function of the current in the z-shim coils (I_{z-shim}) and the current in the QUIC-coils (I_{QUIC}). It can be seen that the trap becomes more "round" when I_{z-shim} is increased or I_{QUIC} is decreased. The top row of table 5.3 ($I_{z-shim} = 0.35\,\text{A}$ and $I_{QUIC} = 300\,\text{A}$) is the trap we use when the BEC is created. Here the aspect ratio of the trap/cloud is 1:15. So, this trap is relatively elongated but it insures at the same time a high confinement which is necessary if we want to get a BEC. Because of three-body collisions the lifetime of the BEC in this trap is low (below 1 s). However when the BEC is created we can "open up" the trap which reduces the density of the cloud and makes the lifetime of the BEC much longer. The lowest row of table 5.3 ($I_{z-shim} = 4\,\text{A}$ and $I_{QUIC} = 150\,\text{A}$) is the trap we change to before we load the optical lattice. More details on the loading procedure of the optical lattices can be found in chapter 8.

5.4.1 QUIC-trap characterization

The trapping frequencies can be used to determine the atom number of a BEC which we will see in section 5.5. However we can e.g. also use the measured trapping frequencies to make a model for the magnetic field created by the QUIC-trap. In chapter 2 we saw an analytical (approximated) model for the magnetic field of the QUIC-trap

$$\mathbf{B} = B_0 \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + B' \begin{pmatrix} x \\ -y \\ 0 \end{pmatrix} + \frac{B''}{2} \begin{pmatrix} -xz \\ -yz \\ z^2 - \frac{1}{2}(x^2 + y^2) \end{pmatrix}. \tag{2.5}$$

We can now find the three constants B_0, B' and B''. B_0 is the bias field which can be found because we can measure where the bottom of the trap is located using the RF-field

$$B_0 = \frac{\hbar \omega_{RF}}{g_F \mu_B}. \tag{5.34}$$

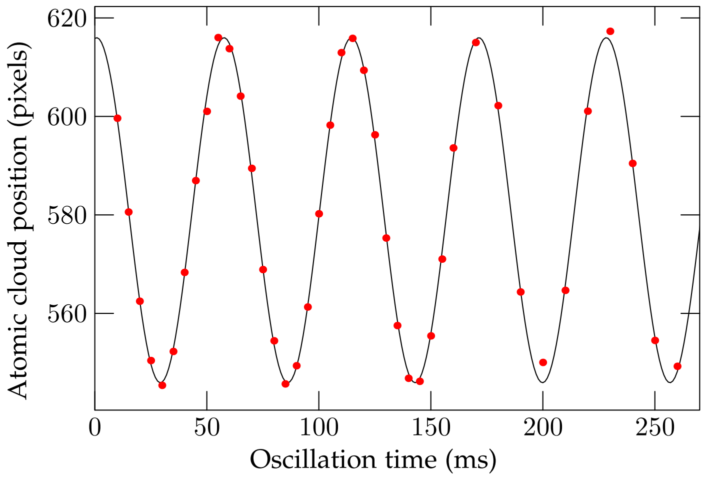

(a) Axial oscillation frequency 17.6 Hz ± 0.2 Hz. The trap was perturbed using the QP-bypass functionality. Based on 280108 run 047–089.

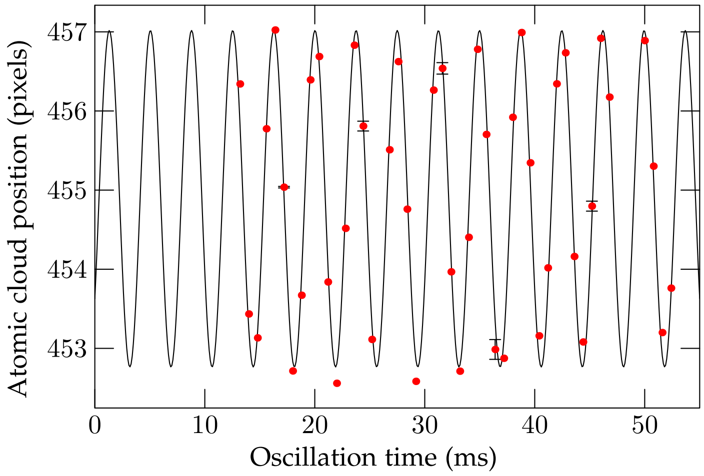

(b) Radial oscillation frequency 267.3 Hz ± 0.4 Hz. Notice that we are not using points with an oscillation time less than around 15 ms. This is because it takes some time before the trap is stable after the perturbation. In this case the x-shim coils were used together with the Agilent function generator to perturb the trap. Based on 050208 run 12–60.

Figure 5.3: Trapping/sloshing frequencies of atomic clouds trapped in the QUIC-trap. In the axial direction we can make large oscillations since the potential is harmonic across a large area. However in the radial direction we have to be careful not to choose too large modulations because here the trap is only harmonic in a small region.

Here ω_{RF} is the angular frequency corresponding to the RF-frequency where all the atoms are lost from the QUIC-trap.

Putting the expression for the magnetic field (eq. (2.5)) into the expression for the magnetic potential (eq. (2.4)) we arrive with an expression that can be compared to a harmonic potential with the following form $V = \frac{1}{2}m(\omega_r^2 r^2 + \omega_z^2 z^2)$. Setting r equal to zero we immediately get B'' in terms of the trapping frequency ω_z

$$B'' = \frac{m\omega_z^2}{g_F m_F \mu_B}. \tag{5.35}$$

If we set $z = 0$, make a Taylor expansion of eq. (2.4) for r around zero and put in the expression for B'' we arrive (after a few calculations) at

$$B' = \sqrt{\frac{B_0 m(2\omega_r^2 + \omega_z^2)}{2g_F m_F \mu_B}}. \tag{5.36}$$

These formulas can be used to translate the measured trapping frequencies into the parameters (B_0, B' and B'') describing the QUIC-trap. Along the axial direction (for $x = y = 0$) the field is harmonic and goes like $B_z = B_0 + B''z^2/2$. Along the radial direction the field is harmonic just in center but otherwise it increases linear like $B_x = B'x$.

We can use the measured trapping frequencies to try to adjust the geometrical parameters in Biot-Savart calculations of the magnetic trap. The difference between these simulations and the analytical model (eq. (2.5)) is very small. If we compare the two methods for a QUIC-trap approximately equal to the first row in table 5.3 (we use the exact same trapping frequencies in the two methods) we get a maximum deviation of about 2.5% in the axial direction and 2% in the radial direction when we go 1 mm away from the trap center. This deviation is so small that we can just use the analytical expression to represent the trap (still Biot-Savart calculations are necessary in the design process).

Values for B_0, B' and B'' as a function of I_{z-shim} can be seen in table 5.4. B'' is almost constant as a function of I_{z-shim} whereas B' changes a lot.

I_{z-shim} (A)	I_{QUIC} (A)	Axial trap freq (Hz)	Radial trap freq (Hz)
0.35	300	17.58±0.02	267.3±0.4
0.75	300		172.6±0.6
1	300	17.72±0.01	147.7±0.4
1.5	300		120.2±0.8
2	300	17.71±0.03	102.1±0.3
3	300	17.70±0.03	82.6±0.2
4	300		78.2±0.3
4	150	12.33±0.01	38.4±0.2

Table 5.3: Trap frequencies for the QUIC-trap as a function of I_{z-shim} and I_{QUIC}. All frequencies are measured using the "oscillating cloud" method. Since the axial trap frequency is almost constant (when $I_{QUIC} = 300$ A) we only measured it for a few settings. The first row shows the parameters of the trap we use when we make a BEC. Before the optical lattice is loaded the QUIC-trap is reshaped into the settings seen in the last row. This means that the aspect ratio of the trap is changed from 1:15 to 1:3. Based on 280108, 290108, 310108, 040208, 050208.

I_{z-shim} (A)	RF bot. (kHz)	ν_r(Hz)	ν_z(Hz)	B_0 (G)	B' (G/cm)	B'' (G/cm^2)
0.35	757	267.3	17.58	1.08	217	190
0.75	1625	172.6		2.31	206	
1.00	2168	147.7	17.72	3.08	204	193
1.50	3253	120.2		4.63	203	
2.00	4339	102.1	17.71	6.17	200	193
3.00	6509	82.6	17.70	9.26	200	192
4.00	8680	78.2		12.3	215	

Table 5.4: Trap parameters as a function of z-shim current. I_{QUIC} is kept constant at 300 A. B_0, B' and B'' are calculated using eq. (5.34), (5.36) and (5.35). Typical values of B_0, B' and B'' can be found in [64].

5.5 Atom number

We have used two different methods to determine the atom number N

1. Cross section calibration
2. Expansion of a BEC

1. Cross section calibration

In [72] the on-resonant cross section for circular polarized light ($\sigma^{\pm}$) on the $|F=2, m_F=\pm 2\rangle$ to $|F=3, m_F=\pm 3\rangle$ cyclic transition is given theoretically ($\sigma_0 = 2.907 \times 10^{-9}\,\text{cm}^2$)[4]. This means that if we use circular polarized light and have a well defined quantization axis (a magnetic field parallel to the propagation direction) we know (theoretically) what the cross section should be. By switching between using circular polarized imaging light (and a bias field) and our normal (linear) polarized imaging system (with no bias field) we can determine the on-resonant cross section for our imaging system. We transfer this "calibration" by comparing the total sum of the OD of the clouds grabbed in the two different situations, and assume that the atom number is constant from run to run (see e.g. eq. (5.3)—in practice we average over many runs). Since we now know σ_0 we can determine the atom number using eq. (5.21), (5.16) or (5.30) depending on the fit-function we are using (in case of the bimodal fits the total atom number is (of course) given by a sum of the number of atoms in the thermal cloud and the number of atoms in the BEC).

Off-resonant imaging

If we want to determine the atom number when we have detuned the imaging light we need to know the (Lorentz) linewidth of the imaging transition. The natural line-width of the imaging transition is $\Gamma = 6.0\,\text{MHz}$. However this transition can be broadened and shifted e.g. because of stray magnetic fields. In figure 5.4 an absorption profile measured in our setup can be seen. The width of this profile is broadened to 7.5 MHz and this value should be used in the calculation of σ. If the profile is constant then it is easy to correct for but if it changes (and shifts) as a function of e.g. TOF it becomes more tricky because the fitting program needs to know the profiles for different TOF in order to determine the atom number.

---

[4] In principle this method can also be used with linear polarized light but it is done in a more controlled fashion (and works better) if we use circular polarized light.

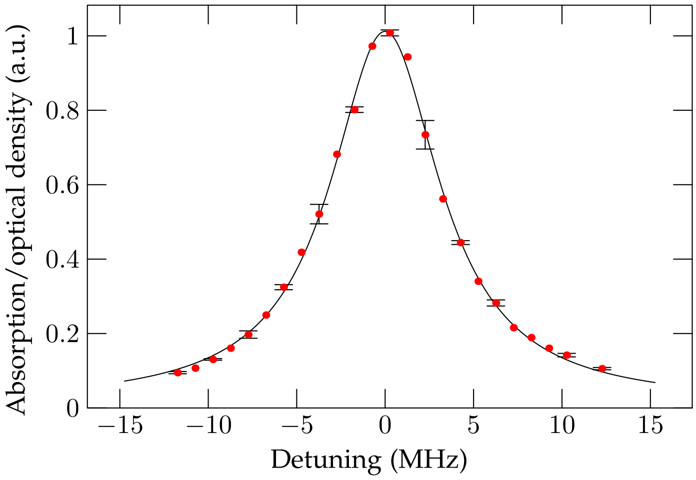

Figure 5.4: Lorentz absorption profile. The points are fitted to a Lorentz with a FWHM $= \Gamma = 7.5\,\mathrm{MHz}$. The AOM center frequency (corresponding to a detuning equal to zero) is 105.9 MHz. In this measurement the z-shim coils were turned off to $t = 0\,\mathrm{ms}$ but it takes several milliseconds before the field is actually off and mostly because of this the linewidth depends on TOF in our experiment. Based on 140208 run 6–43.

2. Expansion of a BEC

Another way to determine the atom number (of a BEC) is by using the theory for an expanding BEC introduced in section 5.2.3. When a BEC is released (in this case from an elongated harmonic trap) the expansion depends on the number of atoms in the BEC and on the trapping frequencies of the harmonic trap. Since we know the trapping frequencies very precise we can determine the atom number just by looking at the shape of the cloud.

In the expression for the widths of an expanding BEC (eq. (5.25)) the chemical potential μ is the only unknown (assuming that we know the trapping frequencies and that we have done a TF-fit to determine the widths). This means that we can determine the number of atoms N through the following relation (derived in Castin and Dum [73])

$$N_{BEC} = \frac{1}{15a} a_{ho} \left(\frac{2\mu}{\hbar\bar{\omega}} \right)^{5/2} \tag{5.37}$$

where $a_{ho} = \sqrt{\frac{\hbar}{m\bar{\omega}}}$ is the harmonic oscillator length, $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$ and a is the s-wave scattering length. The scattering length a is known with a reasonable precision (within $\sim 5\%$) and it can e.g. be found in [4]—we could also go the other way and find the scattering length if we knew the atom number.

In order to determine the atom number we have isolated μ in eq. (5.37) and pasted it into the expressions for the widths given by eq. (5.25) (with different scaling parameters λ_r or λ_z depending on the axis). The resulting expression can be used to make a fit to the widths as a function of TOF where the only unknown is the atom number. This is done in figure 5.5 for both the axial and the radial widths. Notice that in this method we do not need to know σ—instead we can use this method to find σ (or σ_0).

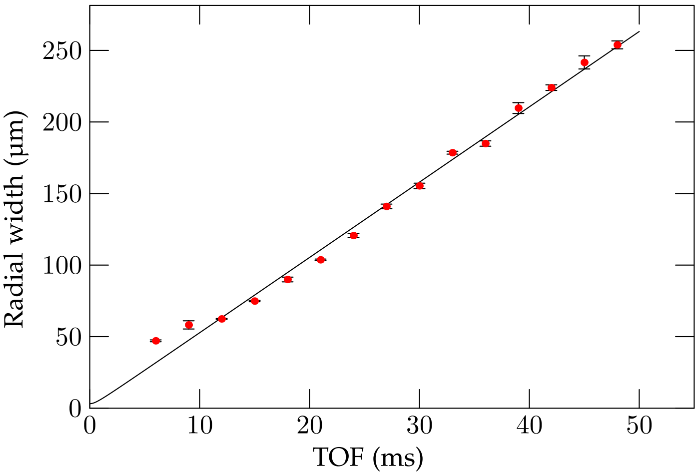

(a) The atom number based on the radial expansion is $N = 290 \times 10^3 \pm 11 \times 10^3$.

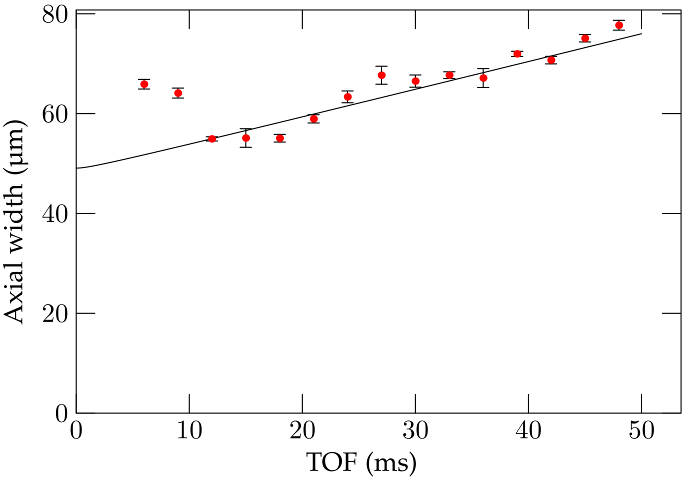

(b) The atom number based on the axial expansion is $N = 336 \times 10^3 \pm 16 \times 10^3$.

Figure 5.5: Axial and radial widths of a BEC in free expansion from the QUIC-trap (which has been turned off) as a function of TOF. Only runs with a peak optical density between 0.2 and 0.5 are used to reduce imaging effects of having a high OD (this means that the data is taken with different detunings on the imaging light, which also can be a problem because we have seen that the widths depend slightly on the detunings—however using only images with a low OD gives the best result). The data is fitted to the theoretical expressions for the widths where only the atom number is unknown. We assign the deviation present at low TOFs to imaging effects—for this reason the first two points have not been included in the fits. The trap frequencies of the QUIC-trap was in this example 267.3 Hz and 17.58 Hz ($I_{z-shim} = 0.35\,\mathrm{A}$ and $I_{QUIC} = 300\,\mathrm{A}$). These measurements were performed with the y-axis imaging system. Based on 140208 run 138–230.

Conclusion

In the first method it is difficult to determine the atom number completely independent of e.g. I_{z-shim} and TOF since we need to know the Lorentz profile for the different settings we are using. For this reason we often use some "reference" settings where we apply the same well-defined turn-off of the magnetic field, wait for the same TOF and so on. Under these conditions we measure the Lorentz absorption profile and this way we can always go back to these settings if we want to know what the atom number is.

Alternatively we have use method number two where we directly get the atom number (but it only works for a BEC). Instead of measuring a lot of points (to different TOF) we can just base the atom number on one run (i.e. one point in figure 5.5 for long TOFs).

Using these methods we can perhaps determine the atom number within 10–20%. However if it becomes necessary to determine the atom number more precisely we should really put more work into this and explore other ways. However so far we have not been so interested in the atom number (except that we want it to be stable).

5.6 Expansion of a BEC

In the previous section we looked at the expansion of a BEC and used this to determine the atom number. In this section we will look a bit more into this expansion.

When a BEC is released from an (elongated) harmonic trap the widths of the BEC will change (as already studied) in accordance with eq. (5.25) with the scaling parameters given by (5.26) and (5.27). Now, let us look at the aspect ratio instead

$$\frac{w_r(t)}{w_a(t)} = \frac{\lambda_r(t)\sqrt{2\mu/m\omega_r^2}}{\lambda_a(t)\sqrt{2\mu/m\omega_a^2}} = \epsilon\frac{\lambda_r(t)}{\lambda_a(t)} \tag{5.38}$$

where $\epsilon = \omega_a/\omega_r \ll 1$.

In figure 5.6 images of a BEC after different TOF can be seen and the aspect ratio is plotted in figure 5.7 also as a function of TOF. This aspect ratio goes above unity and this is one of the most characteristic signatures of having a BEC and not a thermal cloud.

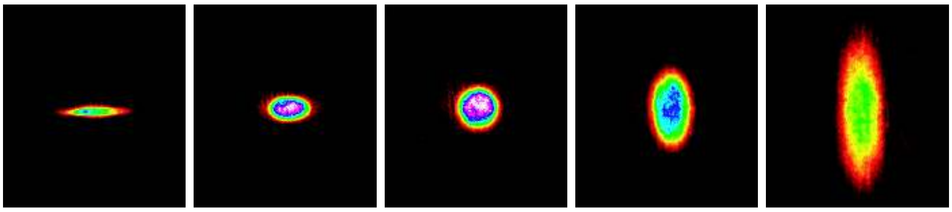

Figure 5.6: Typical absorption images of an (almost) pure BEC after different TOF. Y-axis imaging. TOF = 0 ms, 5 ms, 10 ms, 20 ms and 40 ms. This release was done from the elongated trap where we make the BEC (I_{z-shim} = 0.35 A, I_{QUIC} = 300 A). The imaging area is around 400 µm × 450 µm. The elliptical shape (for long TOFs) indicates that this is a highly non-thermal distribution. Here we are in fact imaging a single, macroscopically occupied quantum wave function [6]. Based on 120208 run 80, 81, 82, 84 and 88.

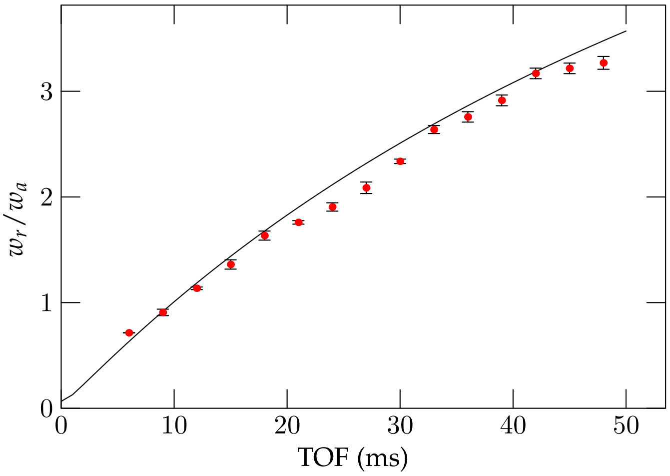

Figure 5.7: The aspect ratio of a BEC released from an (elongated) harmonic trap as a function of TOF. The theory curve is plotted using eq. (5.38) with the (measured) trapping frequencies taken from the first row in table 5.3 (no fitting is performed). Here it can be seen how the fundamental behavior of a BEC is different from the behavior of a thermal cloud (see e.g. figure 5.1). The aspect ratio goes above unity indicating a non-thermal distribution. Same data as used in figure 5.5.

5.7 Formation of a BEC

When we get close to the critical temperature T_c a BEC will start to form in the middle of the thermal cloud. In figure 5.8 it can be seen how the number of atoms in the thermal cloud and in the BEC depends on the RF-end frequency. We like this graph because it tells us something about the stability of our setup. It can e.g. be seen that we have a span of 60 kHz where some part of the cloud is a BEC. It can also be seen that the edge where we cut away all the atoms is narrow (1 kHz–2 kHz) indicating that the trap is very stable. Also it can be seen that we have a wide (perhaps 20 kHz) plateau where the atom number in the BEC is almost constant[5]. Normally we sweep the RF-frequency down to 5 kHz–10 kHz from the edge since here we have an almost pure BEC and the fluctuations in the atom numbers are small due to the presence of the plateau. If we want to check that the BEC-system is working properly we measure this profile.

It can also be seen in figure 5.8 that we typically form BECs with around 3×10^5 atoms. However we can push the system and get up to 1×10^6 atoms in the BEC. To do this we load more atoms in the MOT and use a fast linear ramp in the end of the evaporative cooling sequence when three-body loses set in. When we do this it seems like the atom number fluctuates more, however it is not something we have investigated a lot so it might be possible to optimize this procedure and get a high, stable atom number.

In figure 5.9 it can be seen what the formation of a BEC out of a thermal cloud looks like in the absorption images. The images clearly show that we only have a thermal cloud for high RF-end frequencies. After more cooling the BEC can be seen as a dense, elliptical peak in the middle of the thermal cloud (a bimodal distribution). Finally, the

---

[5] This might be due to three-body collisions which can have an effect of keeping the atom number constant—see e.g. [45].

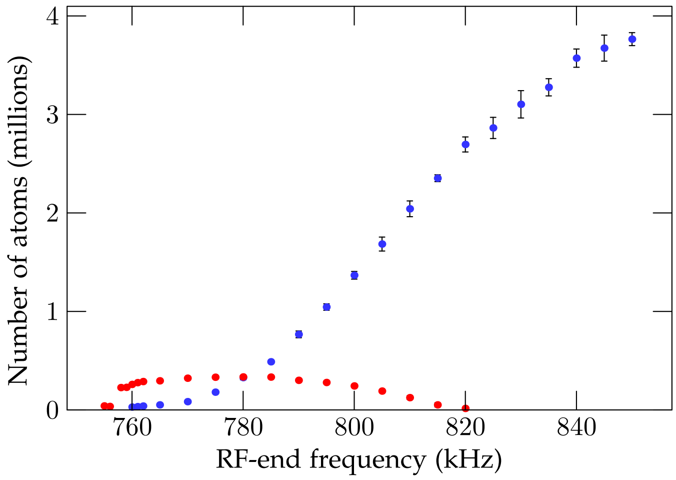

(a) Here the thermal component (blue) and the BEC component (red).

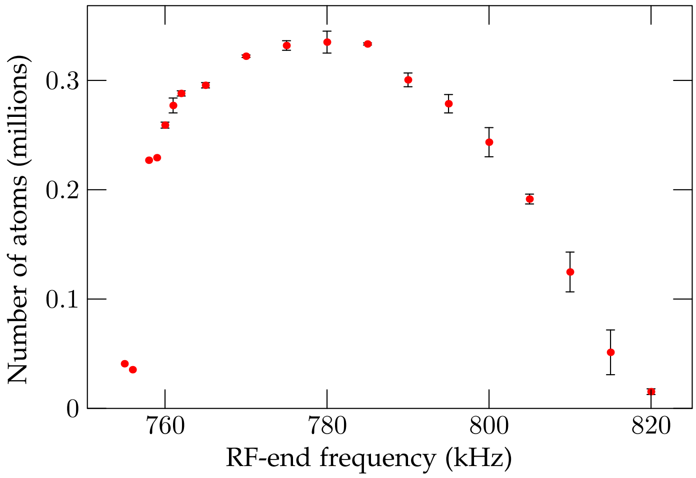

(b) Here only the BEC component.

Figure 5.8: The number of atoms in the thermal component and in the BEC component as a function of the RF-end frequency. We cut away all the atoms just by changing the RF-end frequency 1 kHz–2 kHz indicating that the magnetic trap is really stable. A pure thermal cloud exists for high RF-end frequencies. When the RF-end frequency is lowered a BEC is formed. An almost pure BEC is formed in the last 10 kHz before the edge where all the atoms are lost. A wide plateau with an almost constant atom number in the BEC is also seen. Shot to shot fluctuations in atom number is less than 10%. If something in the setup is changed (or if e.g. a tool is forgotten close to the science chamber) this profile will shift because of the changed magnetic bias field. Based on 080208 run 9–79.

thermal cloud almost goes away and we are left with an almost pure BEC. All the images in figure 5.9 are analyzed using a bimodal Bose/TF-fit (see e.g. table 5.1)—examples of these fits can be seen in figure 5.10. The 2D fits are represented in 1D by placing all (pixel) rows of the image after each other in a one dimensional array. In figure 5.10 it can be seen that the data is really nicely represented by the fits. Close to T_c (but still a thermal cloud) the data is represented by a Bose fit (figure 5.10a and 5.10b). In the intermediate region the bimodal shape is clearly seen in figure 5.10c. In figure 5.10d the TF-shape indicating a BEC is seen. The number of atoms in the thermal part and in the BEC part can be extracted from these fits (since we have calibrated σ_0). The temperature can be determined from the thermal part of the fits using eq. (5.15).

This ends the first part of this thesis about the basic BEC setup and results made with this setup. In the next chapter I will change subject to BECs in optical lattices.

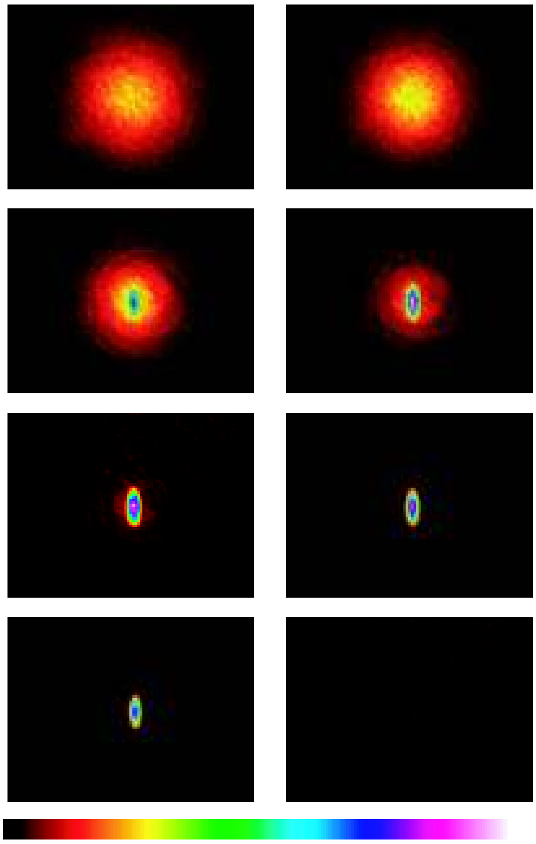

#1
$\nu_{RF} = 840\,\mathrm{kHz}$
$N_{sum} = 3.6 \times 10^6$
$N_{bose} = 3.6 \times 10^6$
$N_{BEC} = 0$
$\epsilon_t = 8.0\,\mu\mathrm{K}$
$T = 1.1\,\mu\mathrm{K}$

#2
$\nu_{RF} = 830\,\mathrm{kHz}$
$N_{sum} = 3.2 \times 10^6$
$N_{bose} = 3.2 \times 10^6$
$N_{BEC} = 0$
$\epsilon_t = 7.0\,\mu\mathrm{K}$
$T = 0.83\,\mu\mathrm{K}$

#3
$\nu_{RF} = 810\,\mathrm{kHz}$
$N_{sum} = 2.2 \times 10^6$
$N_{bose} = 2.1 \times 10^6$
$N_{BEC} = 0.10 \times 10^6$
$\epsilon_t = 5.1\,\mu\mathrm{K}$
$T = 0.63\,\mu\mathrm{K}$

#4
$\nu_{RF} = 795\,\mathrm{kHz}$
$N_{sum} = 1.2 \times 10^6$
$N_{bose} = 0.98 \times 10^6$
$N_{BEC} = 0.24$
$\epsilon_t = 3.7\,\mu\mathrm{K}$
$T = 0.46\,\mu\mathrm{K}$

#5
$\nu_{RF} = 780\,\mathrm{kHz}$
$N_{sum} = 0.56 \times 10^6$
$N_{bose} = 0.25 \times 10^6$
$N_0 = 0.31 \times 10^6$
$\epsilon_t = 2.2\,\mu\mathrm{K}$
$T = 0.27\,\mu\mathrm{K}$

#6
$\nu_{RF} = 765\,\mathrm{kHz}$
$N_{sum} = 0.35 \times 10^6$
$N_{bose} = 0.05 \times 10^6$
$N_0 = 0.30 \times 10^6$
$\epsilon_t = 0.78\,\mu\mathrm{K}$
$T < 0.25\,\mu\mathrm{K}$

#7
$\nu_{RF} = 758\,\mathrm{kHz}$
$N_{sum} = 0.24 \times 10^6$
$N_{bose} = 0.01 \times 10^6$
$N_{BEC} = 0.23 \times 10^6$
$\epsilon_t = 0.11\,\mu\mathrm{K}$
$T < 0.25\,\mu\mathrm{K}$

#8
$\nu_{RF} = 755\,\mathrm{kHz}$
$N_{sum} = 0$
$N_{bose} = 0$
$N_{BEC} = 0$
$\epsilon_t = -0.18\,\mu\mathrm{K}$
$T < 0.25\,\mu\mathrm{K}$

Figure 5.9: These images show how a BEC is formed out of a thermal cloud. The only difference between the runs is the RF-end frequency ν_{RF}. All the images are fitted to a bimodal Bose/TF-fit. N_{bose} is the number of thermal atoms, N_{BEC} is the number of condensed atoms and N_{sum} is the total number of atoms. The truncation energy ϵ_t is calculated using eq. (2.7) and the temperature T is calculated from the width of the thermal component using eq. (5.15). In this case we have used y-axis imaging with TOF equal to 30 ms. The imaging area is approximately 1.55 mm × 1.35 mm and white corresponds to OD $= 2$ on the colorscale. The magnetic trap parameters are equal to the ones listed in the first row in table 5.3 and the trap bottom is located around 756.5 kHz. The axial direction of the trap is horizontal in these images and the radial direction is vertical. Based on 080208 run 51, 13, 57, 40, 43, 66, 73, 48 (ordered by RF-end frequencies descending).

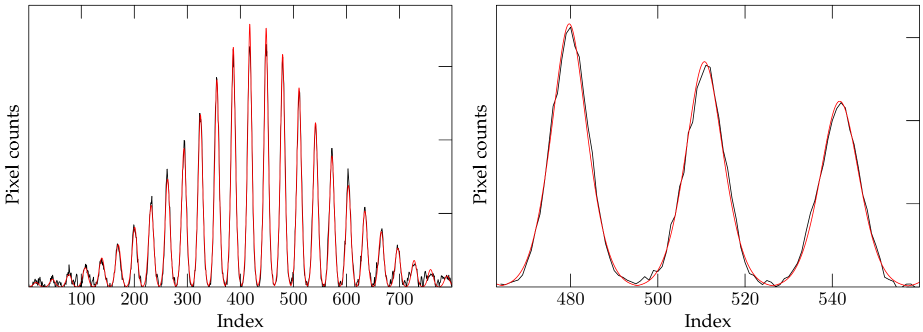

(a) Image #1 from figure 5.9. Here the TF part is negligible so it is a pure Bose fit. The bin-factor is set to 19 to make it more easy to see the different loops.

(b) Image #1 from figure 5.9. Same as (a) but here we have zoomed in on three loops.

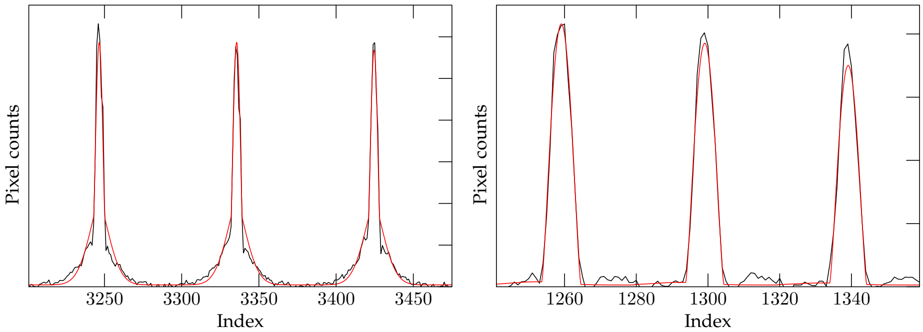

(c) Image #4 from figure 5.9. Here the bimodal shape is clearly seen. A bin-factor of three is used and we have zoomed in on three loops.

(d) Image #7 from figure 5.9. An almost pure BEC (and an almost a pure TF fit). Here a bin-factor of three is used and we have zoomed in on three loops.

Figure 5.10: These graphs show how the shape of the cloud changes when we go from a thermal cloud to a BEC. Data points (black lines) and 2D bimodal Bose/TF fits (red lines). The 2D-fits (and the data) are stretched out to be shown in one dimension by placing all horizontal rows in the image after each other in a one-dimensional array. Image #1, #4 and #7 from figure 5.9 are fitted. Normally we do not use bin-factors unless the cloud is really big but here we have done it to make it easier to see what is going on (especially in figure (a)).

Six

Introduction to optical lattices

6.1 Introduction

In this chapter I will introduce some basic theory of Bose-Einstein condensates in optical lattices. In section 6.2 I will show how the Stark effect makes it possible to trap neutral atoms in a light field. I will also introduce a (Gaussian) light field and calculate trap depths and trapping frequencies for different optical traps. In section 6.3 I will present band-structure calculations which proves very useful in the understanding of what happens when a BEC is loaded into an optical lattice. In the end of this chapter I will introduce a new kind of hybrid optical trap that we have implemented in the setup (section 6.7).

6.2 Optical trap

6.2.1 AC Stark shift

When atoms are placed in a light field the electric field of this light will induce an oscillating electric dipole moment in the atoms. The electric field of the light will interact with this induced dipole moment and this alters the energy levels of the atoms and makes it possible to trap neutral atoms using only laser beams. In the case of far-detuned laser light the shift of the energy levels can be treated as a second order perturbation of the electric field [75]. The generally well-known second-order time-independent perturbation theory for non-degenerate states gives an energy correction ΔE_i of the i-th state[1]

$$\Delta E_i = \sum_{j \neq i} \frac{|\langle j|\mathcal{H}|i\rangle|^2}{\mathcal{E}_i - \mathcal{E}_j}. \tag{6.1}$$

In this case the Hamiltonian $\mathcal{H} = -\boldsymbol{\mu} \cdot \mathbf{E}$ ($\boldsymbol{\mu} = -e\mathbf{r}$) describes the atom-light interaction and $\mathcal{E}_i$ is the unperturbed energy of the i-th state. The sum is carried out over all other states j with unperturbed energies $\mathcal{E}_j$. For far detuned laser light the atoms will spend (by far) most of their time in the ground state. This means that it is the correction to the ground state energy of the atoms that we need to calculate and this correction depends in

---

[1] Non-degenerate perturbation theory can be used because there is no coupling between degenerate ground states for pure linear or pure circular polarizations.

principle on all coupled excited states. However normally the main contribution comes only from a few excited levels so let us start by looking at the correction to the ground state energy of a two level atom[2]

$$\Delta E = \frac{3\pi c^2}{2\omega_0^3}\frac{\Gamma}{\Delta}I, \tag{6.2}$$

where ω_0 is the angular (optical) transition frequency between the ground state level and the excited level, ω is the angular frequency of the laser beam, Γ is the spontaneous decay rate of the excited state (also equal to the natural linewidth), $\Delta = \omega - \omega_0$ is the detuning and I is the light intensity.

Some interesting features are seen in eq. (6.2). First of all for red-detuned lattice light ($\omega < \omega_0$) the detuning Δ is negative resulting in a negative energy correction ΔE to the ground state energy. This means that in this case atoms will seek towards high light intensity (there will be a "well" or "dip" in the ground state energy proportional to the light intensity I). This is opposite for blue-detuning. Another important tendency is that the potential scales as I/Δ. For the optical lattices to work it is important that the rate at which photons are absorbed by the atoms is very low. Fortunately the scattering rate [75]

$$\Gamma_{sc} = \frac{3\pi c^2}{2\hbar\omega_0^3}\left(\frac{\Gamma}{\Delta}\right)^2 I \tag{6.3}$$

scales as I/Δ^2 so this means that by choosing a large detuning and a high intensity the scattering rate can be made small and the potential large.

For our alkali atoms we will use a slightly more refined expression than eq. (6.2) since the main contribution to the shift of the ground state energy comes from the D1 line (${}^2\mathrm{S}_{1/2} \rightarrow {}^2\mathrm{P}_{1/2}$) and the D2 line (${}^2\mathrm{S}_{1/2} \rightarrow {}^2\mathrm{P}_{3/2}$) [75]

$$V(x,y,z) = \frac{\pi c^2\Gamma}{2\omega_0^3}\left(\frac{2+\mathcal{P}g_\mathrm{F}m_\mathrm{F}}{\Delta_{2,F}} + \frac{1-\mathcal{P}g_\mathrm{F}m_\mathrm{F}}{\Delta_{1,F}}\right)I(x,y,z), \tag{6.4}$$

where $\mathcal{P}$ is the light polarization ($\mathcal{P} = 0$ for linear and ± 1 for circular $\sigma^\pm$ polarized light, respectively), $\Delta_{1,F}$ and $\Delta_{2,F}$ is the detuning of the lattice light from the transition frequency between the ground state ${}^2\mathrm{S}_{1/2}$ (with total angular momentum F) and the center of the hyperfinesplitting ${}^2\mathrm{P}_{1/2}$ and ${}^2\mathrm{P}_{3/2}$, respectively (so it is still assumed that the excited-state hyperfine structure is unresolved).

6.2.2 Trap depth and trapping frequencies

Eq. (6.4) can be used to calculate the trap depths and the trapping frequencies of different optical traps. Here we will look at the case where we have three mutually perpendicular linear polarized Gaussian laser beams with perpendicular polarizations. If the beams are retro-reflected we will get a simple-cubic (sc) optical lattice. If the beams are *not* retro-reflected we will not have a standing wave and I will throughout this thesis refer to this kind of trap as a "dipole trap". In appendix C the trap depths and the trapping frequencies are calculated for 1D, 2D and 3D optical dipole traps and lattices of this type. In the following a few of the main results will be presented.

---

[2] The calculation is done in details in [75].

In a 1D optical lattice (with lattice light propagating in the +z and -z directions) the intensity of the Gaussian light field is given by

$$I(x,y,z) = \frac{8P}{\pi w^2(z)} e^{-2(x^2+y^2)/w^2(z)} \cdot \cos^2(kz), \tag{6.5}$$

where P is the beam power, $w(z) = w_0\sqrt{1 + z^2/z_0^2}$ is the spot size of the beam, $z_0 = \frac{\pi w_0^2}{\lambda}$ is the Rayleigh length, $k = 2\pi/\lambda$, where λ is the wavelength of the lattice beams, and w_0 is the beam waist (all these parameters are assumed equal for the two counter propagating beams). The main difference between a 1D dipole trap and a 1D lattice is that in the case of a 1D dipole trap there is no $\cos^2$ term—no standing wave[3]. The trapping potential for a 1D optical lattice $V_{\text{1D-lat}}(x,y,z)$ can be calculated for a given position (x,y,z) by inserting eq. (6.5) into eq. (6.4). The trap depth $V_0 \equiv V_{\text{1D-lat}}(x=0, y=0, z=0)$ is given by

$$\begin{aligned} V_0 &\equiv |V_{\text{1D-lat}}(x=0, y=0, z=0)| \\ &= \frac{\pi c^2 \Gamma}{2\omega_0^3}\Big(\frac{2 + \mathcal{P} g_F m_F}{\Delta_{2,F}} + \frac{1 - \mathcal{P} g_F m_F}{\Delta_{1,F}}\Big) \cdot \frac{8P}{\pi w_0^2} \end{aligned} \tag{6.6}$$

which means that we can express the potential as

$$V_{\text{1D-lat}}(x,y,z) = V_0 \frac{1}{(1 + z^2/z_0^2)} e^{-2(x^2+y^2)/w^2(z)} \cos^2(kz). \tag{6.7}$$

Normally we express the depth of the (1D) optical lattice in terms of the dimensionless parameter s, which is the number of recoil energies E_r

$$s \equiv \frac{V_0}{E_r}. \tag{6.8}$$

Here $E_r = \hbar^2 k^2/2m$, $k = 2\pi/\lambda$, where λ is the wavelength of the optical lattice laser, and m is the mass of an atom.

Notice that even though we call V_0 (or s) the depth of the lattice it is really only the maximum depth of a 1D optical lattice. In 2D and 3D we get one or two more terms on eq. (6.7) but we still use the same V_0. This means that e.g. in 3D the maximum depth is $3V_0$. Still we say that the depth of the 3D optical lattice is V_0.

The trapping frequencies can be found by making a Taylor expansion of eq. (6.7) around $x = y = z = 0$ (this is done in appendix C). The result is

$$\omega_x = \omega_y = \sqrt{\frac{4V_0}{m w_0^2}}, \qquad \omega_z = \sqrt{\frac{2V_0}{m} \cdot k^2}. \tag{6.9}$$

In this 1D case the axial trapping frequency is set by the wave vector k. This means that a trapping frequency of several kHz in the axial direction is obtained. The radial trapping frequencies are however only several Hz. Similar calculations can be done in 2D and 3D (see appendix C). The results can be seen in table 6.1.

Examples of the trap depth and the trapping frequencies for 1D, 2D and 3D optical lattices and dipole traps can be seen in table 6.2. In figure 6.1 potential isosurfaces can be seen for a 1D, 2D and 3D optical lattice. All the calculations are done using rubidium 87 and 915 nm laser beams.

---

[3] Also the intensity differs by a factor of 4—see e.g. eq. (C.3).

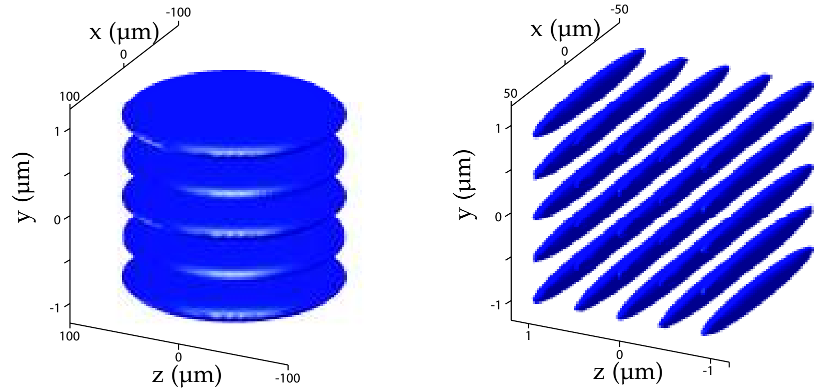

(a) 1D lattice (y-axis)—$V_0 = 66E_r$. Isosurfaces are placed $45E_r$ above the trap minimum.

(b) 2D lattice (y- and z-axis)—$V_0 = 66E_r$. The maximum depth is $2 \times 66E_r$. Isosurfaces are placed $40E_r$ above the trap minimum.

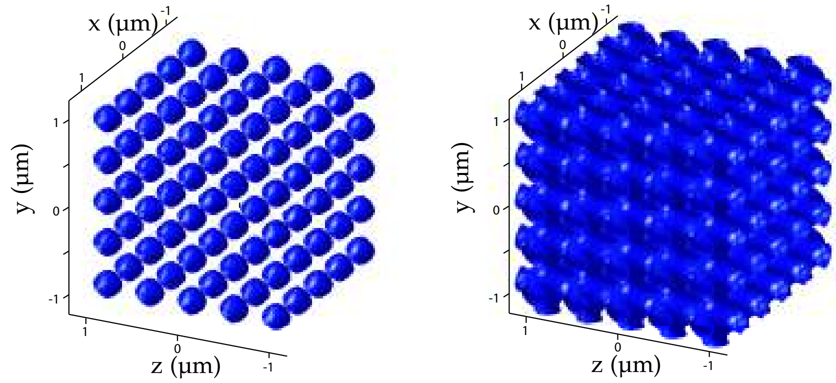

(c) 3D lattice—$V_0 = 66E_r$. The maximum depth is $3 \times 66E_r$. Isosurfaces are placed $50E_r$ above the trap minimum.

(d) 3D lattice—Same as (c) but here the isosurfaces are placed $100E_r$ above the trap minimum. Here the isosurfaces of the different sites "melt together".

Figure 6.1: Potential energy isosurfaces of a 1D, 2D and 3D optical lattice. The figures show that if the atoms are cooled enough they will form disks (a 2D gas) in a 1D lattice (a), cigars (a 1D gas) in a 2D lattice (b), and a simple cubic lattice (one can say a 0D gas) in a 3D lattice ((c) and (d)). For all beams, the laser power is 200 mW and the beam waist is 120 μm. The script used for calculating these isosurfaces use the completely general expressions found in appendix C—this means that we can e.g. have different beam waists, powers and wavelengths in all the different beams and we can actually also turn on gravity and the QUIC-trap.

Trap type	Trap depth	ω_x	ω_y	ω_z
1D dipole (z)	V_{dip}	$\sqrt{\frac{4V_{dip}}{mw_0^2}}$	$\sqrt{\frac{4V_{dip}}{mw_0^2}}$	$\sqrt{\frac{2V_{dip}}{mz_0^2}}$
2D dipole (y/z)	$2V_{dip}$	$\sqrt{\frac{8V_{dip}}{mw_0^2}}$	$\sqrt{\frac{4V_{dip}}{mw_0^2}}$	$\sqrt{\frac{4V_{dip}}{mw_0^2}}$
3D dipole	$3V_{dip}$	$\sqrt{\frac{8V_{dip}}{mw_0^2}}$	$\sqrt{\frac{8V_{dip}}{mw_0^2}}$	$\sqrt{\frac{8V_{dip}}{mw_0^2}}$
1D lattice (z)	V_0	$\sqrt{\frac{4V_0}{mw_0^2}}$	$\sqrt{\frac{4V_0}{mw_0^2}}$	$\sqrt{\frac{2V_0k^2}{m}}$
2D lattice (y/z)	$2V_0$	$\sqrt{\frac{8V_0}{mw_0^2}}$	$\sqrt{\frac{2V_0k^2}{m}}$	$\sqrt{\frac{2V_0k^2}{m}}$
3D lattice	$3V_0$	$\sqrt{\frac{2V_0k^2}{m}}$	$\sqrt{\frac{2V_0k^2}{m}}$	$\sqrt{\frac{2V_0k^2}{m}}$

Table 6.1: Trap depths and angular trapping frequencies. Here $V_{dip} \equiv \frac{1}{4}V_0$. (y/z) means that the beams are propagating in the y- and z-direction. The calculations leading to these expressions can be found in appendix C.

Trap type	Depth (in E_r)	Depth (in µK)	ν_x (Hz)	ν_y (Hz)	ν_z (Hz)
1D dipole (z)	16	2.2	38	38	0.066
2D dipole (y/z)	2×16	2×2.2	54	38	38
3D dipole	3×16	3×2.2	54	54	54
1D lattice (z)	66	8.7	76	76	44×10^3
2D lattice (y/z)	2×66	2×8.7	108	44×10^3	44×10^3
3D lattice	3×66	3×8.7	44×10^3	44×10^3	44×10^3

Table 6.2: Examples of trap depths and trapping frequencies. All beams have a beam waist of 120 µm and a power of 200 mW. The trap depth is given in recoil energies and in temperatures T equivalent to an energy of k_BT (where k_B is the Boltzmann constant). Notice that even though the total trap depth for a 3D optical lattice is e.g. 3×66 E_r we still say that we have a 3D optical lattice with $s = 66$.

6.3 Band structure

In many of the optical lattice experiments we have performed so far, band structure calculations have proved very useful in the analysis of the outcome of these measurements. In this chapter we will look at the theory and then use it in chapter 8.

As already mentioned atoms in an optical lattice can to some extend be compared to electrons in a solid. In solid state physics the nearly free electron model is used to describe an electron moving in a periodic potential. The model introduces the electronic band structure and was introduced by Felix Bloch in 1928. In an optical lattice the electrons are replaced by atoms but still the same formalism can be used. We consider a single atom (no interactions between the atoms are included) moving in a periodic potential (a one-dimensional optical lattice)[4]. We need to solve the time independent Schrödinger equation

$$-\frac{\hbar^2}{2m}\frac{d^2\psi(x)}{dx^2} + V(x)\psi(x) = E\psi(x). \tag{6.10}$$

---

[4] This can easily be generalized to the 3D case since the corresponding Hamiltonian is fully separable.

Here we will use the simplified optical lattice potential $V(x) = V_0 \cos^2(kx)$, which is infinite and does not have a harmonic confinement like real (Gaussian) laser beams have[5]. If we divide each term in this equation with the recoil energy $E_r = \hbar^2 k^2/(2m)$ we get

$$-\frac{1}{k^2}\frac{d^2\psi_{n,q}(x)}{dx^2} + s\cos^2(kx)\psi_{n,q}(x) = \epsilon_{n,q}\psi_{n,q}(x), \tag{6.11}$$

where we have used that

$$V(x) = V_0\cos^2(kx) = sE_r\cos^2(kx) \quad \text{and} \quad E = \epsilon_{n,q}E_r. \tag{6.12}$$

According to Bloch's theorem [76], the solutions (eigenfunctions) to a Schrödinger equation with a periodic potential is given by Bloch wave functions $\psi_{n,q}(x)$ which are the products of a plane wave $\exp(iqx/\hbar)$ and a function $u_{n,q}(x)$ with the periodicity of the optical lattice ($u_{n,q}(x) = u_{n,q}(x+a)$, where $a = \pi/k = \lambda/2$)

$$\psi_{n,q}(x) = e^{iqx/\hbar} \cdot u_{n,q}(x). \tag{6.13}$$

We will in the following refer to n as the band index[6] and q as the quasi-momentum. If this ansatz is put into eq. (6.11) we get a Schrödinger equation for $u_{n,q}(x)$

$$-\frac{1}{k^2}\left[\frac{-q^2}{\hbar^2} + \frac{2iq}{\hbar}\frac{d}{dx} + \frac{d^2}{dx^2}\right]u_{n,q}(x) + s\cos^2(kx)u_{n,q}(x) = \epsilon_{n,q}u_{n,q}(x). \tag{6.14}$$

An arbitrary periodic function (e.g. $u_{n,q}(x)$ and $s\cos^2(kx)$) can be written as a Fourier series

$$u_{n,q}(x) = \sum_{l=-\infty}^{\infty} c_{l,n,q}e^{i2klx} \quad \text{and} \quad s\cos^2(kx) = s\Big(\tfrac{1}{2} + \tfrac{1}{4}(e^{-i2kx} + e^{i2kx})\Big). \tag{6.15}$$

Here the periodicity is $2\pi/2k = \pi/k = \lambda/2$ as it should be. Note that the Bloch wavefunctions are then given by

$$\psi_{n,q}(x) = \sum_{l=-\infty}^{\infty} c_{l,n,q}|\phi_{2l\hbar k+q}\rangle, \tag{6.16}$$

where $|\phi_{2l\hbar k+q}\rangle = e^{i(2lk+q/\hbar)x}$.

Inserting (6.15) into (6.14) and comparing all terms with the same e^{i2klx} term we obtain equations where the coefficients $c_{l,n,q}$ and the eigenenergies $\epsilon_{n,q}$ are unknown. These equations can be written in matrix form

$$\sum_{l=-\infty}^{\infty} H_{l,l'} \cdot c_{l,n,q} = \epsilon_{n,q}c_{l,n,q}, \quad \text{where} \quad H_{l,l'} = \begin{cases} (2l + \frac{q}{\hbar k})^2 + s/2 & \text{if } l = l' \\ -s/4 & \text{if } |l-l'| = 1 \\ 0 & \text{else} \end{cases} \tag{6.17}$$

---

[5] The lattice spacing in this potential is $a = \frac{1}{2}\lambda = \pi/k$ and the reciprocal lattice vector G is given by $G = \frac{2\pi}{a} = \frac{4\pi}{\lambda} = 2k$.

[6] Throughout this thesis we will name the lowest band (the ground band) zero—in the literature it is sometimes named zero and sometimes one.

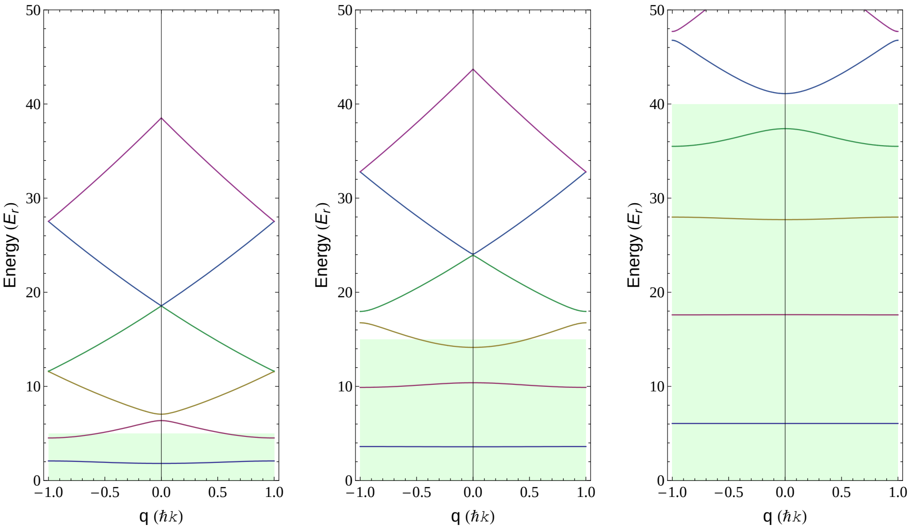

(a) $s = 5$.

(b) $s = 15$.

(c) $s = 40$.

Figure 6.2: Band structure of an optical lattice for different lattice depths. The energy of the Block bands ($\epsilon_{n,q}$) are plotted as a function of quasi-momentum q (in the range from $-\hbar k$ to $\hbar k$—the first Brillouin zone) for different band number n and different depths s. Only the first six bands (band 0–5) are plotted. Notice that for a deep lattice the first bands are almost flat. It is also seen that the distance between the bands increase with increasing depths s.

Solving this equation numerically (just using e.g. $l = -20, -19, -18, \ldots, 20$) we obtain the coefficients $c_{l,n,q}$ and the eigenenergies $\epsilon_{n,q}$. The coefficients can be used together with e.q. (6.16) to find the wavefunctions $\psi_{n,q}(x)$ and thereby e.g. determine the probability for finding the atom at a certain position in the optical lattice. I will not show these results here (they can be found in e.g. [77]) but I will show the eigenenergies. In figure 6.2 a band structure calculation (a plot of the eigenenergies) can be seen for different lattice depths.

6.4 Interference peaks

In optical lattice experiments we (also) normally measure the momentum distribution of the condensate instead of the spatial distribution (the optical resolution is not good enough to see individual lattice sites). This means that we turn off the optical lattice potential and wait for a certain TOF before we grab the images[7].

If a BEC is loaded into an optical lattice which is later turned off, the momentum distribution of the condensate will have discrete peaks (in multiples of the momentum corresponding to the reciprocal lattice vector $2\hbar k$, where $k = 2\pi/\lambda$)[8]. These peaks are a consequence of the periodic potential (and the spatial periodic eigenfunctions). One way to think about this is that coherent matter waves are emitted from each lattice site and the

---

[7] It should be mentioned that single site (actually single atom) resolution has been gained experimentally in an experiment where a scanning electron microscope is combined with a BEC/1D optical lattice setup [78].

[8] These momentum peaks have many names—e.g. interference or diffraction patterns/fringes/peaks.

interference between these waves generate discrete peaks. Another way to say the same is that it is just the reciprocal lattice we see and that we can get the pattern by taking the Fourier transform of the wavefunction in the optical lattice (which we found in the previous section). If this calculation is carried out we get delta functions separated with $2\hbar k$ because the lattice potential used in the previous section was infinite. However in the case where we have a certain distribution of atoms in only a finite number of lattice sites (like in a real experiment) the momentum peaks will have a certain width.

6.5 Non-adiabatic loading

In chapter 8 we will look (experimentally) at what happens when we load and turn off an optical lattice in different ways. In this section we will see what happens (theoretically) when an optical lattice is turned on and off suddenly.

A BEC can be described as a plane wave $\phi_q = e^{iqx/\hbar}$ with momentum q. If we load this condensate into an optical lattice abruptly to $t = 0$ the wavefunction can be written as a superposition of Bloch states $|n, q\rangle = \psi_{n,q}(x)$ (given by eq. (6.16)—we are basically following [79])

$$|\Psi(t=0)\rangle = \sum_{n=0}^{\infty} |\Psi_n(t=0)\rangle, \tag{6.18}$$

where

$$|\Psi_n(t=0)\rangle = |n,q\rangle\langle n,q|\phi_q\rangle. \tag{6.19}$$

Using eq. (6.16) we can find $\langle n, q|\phi_q\rangle$

$$\langle n,q|\phi_q\rangle = \sum_{l=-\infty}^{\infty} c^*_{l,n,q}\langle \phi_{2kl\hbar+q}|\phi_q\rangle = c^*_{0,n,q}. \tag{6.20}$$

Now we can write down the time evolution of the BEC-wavepacket

$$|\Psi_n(t)\rangle = |\Psi_n(t=0)\rangle e^{-i\epsilon_{n,q}\hbar k^2 t/2m} \tag{6.21}$$

$$= c^*_{0,n,q}|n,q\rangle e^{-i\epsilon_{n,q}\hbar k^2 t/2m} \tag{6.22}$$

If we keep the lattice on for a time Δt and then abruptly turn it off again we can project the BEC-wavepacket onto the planewave basis $|\phi_{2l\hbar k+q}\rangle$

$$|\Psi_n(\Delta t)\rangle = \sum_{l=-\infty}^{\infty} b_{l,n,q}|\phi_{2kl\hbar+q}\rangle, \tag{6.23}$$

where the coefficients $b_{l,n,q}$ are given by

$$b_{l,n,q} = c^*_{0,n,q} c_{l,n,q} e^{-i\epsilon_{n,q}\hbar k^2 \Delta t/2m}. \tag{6.24}$$

In figure 6.3 the population $|b_{l,n,q}|^2$ is plotted as a function of different bands n and different momentum orders l for fixed s and q. This figure shows how the different planewave states are occupied when the lattice is loaded suddenly. Notice that we are only occupying even bands when $q = 0$.

Figure 6.4 shows basically the same (this figure is plotted for $s = 40$). However in this figure we focus on the time development of a BEC trapped in a (stationary $q = 0$)

optical lattice. If we do this experiment in the laboratory we measure in the momentum basis, and we want to know what the populations in the different momentum peaks ($0\hbar k, \pm 2\hbar k, \dots$) are as a function of pulse time Δt. These populations are calculated by summing over the different bands (for fixed momentum order l) and then take the norm squared ($|\sum_{n=0}^{\infty} b_{l,n,q}|^2$). Notice that the population in the different bands n and in the different momentum orders l is time-independent (like seen in figure 6.3), however $|\sum_{n=0}^{\infty} b_{l,n,q}|^2$ is not. The $b_{l,n,q}$ coefficients oscillate in time (see eq. (6.24)) with a frequency that depends on the energies $\epsilon_{n,q}$ of the different bands (this can be seen by looking at the real and imaginary parts of $b_{l,n,q}$ which is shown in figure 6.4). The interference between these coherent oscillations of the $b_{l,n,q}$ coefficients give rise to complicated oscillations of e.g. the $0\hbar k$ momentum component/peak (figure 6.4a) and the $\pm 2\hbar k$ peak (figure 6.4b). Up to around $s = 12$ we will primarily populate two bands (band 0 and 2—see figure 6.5a). In this case we will just see coherent oscillations between these two bands. This corresponds to simple "Rabi oscillations" with an angular frequency $(\epsilon_{n=2,q=0} - \epsilon_{n=0,q=0})E_r/\hbar$ given by the energy difference between the two bands. When more than two bands are involved we get complicated oscillations like the ones seen in figure 6.4a and 6.4b. The coherent oscillations are only present if the lattice is loaded suddenly—if the lattice is loaded adiabatically we will only populate the ground band and there will be no time-development of the momentum peaks.

In the rest of this thesis we will refer to the case where the optical lattice is applied for a short period Δt as Kapitza-Dirac scattering [80].

6.6 Adiabatic loading

Often we are interested in loading a BEC adiabatically into the optical lattice potential. Here we will look (briefly) at what that means.

Adiabaticity with respect to band structure

In order to prevent atoms from being excited into higher bands during the loading we have to ramp the lattice on slowly. How slowly depends on the spacing between the bands. For $q = 0$ we can potentially load atoms into band 0 and 2 (and higher order even bands) and here the distance will always be at least 4 E_r no matter how week the lattice is (this can be seen by making a band structure calculation like the one shown in figure 6.2 using a small s). In e.g. [79] it is shown that the ramping must then be done in accordance with $ds/dt \ll 16E_r/\hbar$. With our experimental parameters this gives $ds/dt \ll 8 \times 10^5\,\mathrm{s}^{-1}$, which means that it should be done on a time scale much larger than a few tens of microseconds (assuming we ramp up to $s = 15$).

In figure 6.6 it can be seen how the $0\hbar k$, $\pm 2\hbar k$ and $\pm 4\hbar k$ peaks will be populated in the case where the lattice is loaded adiabatically (which means that all the atoms are loaded into band 0) and then turned off suddenly.

Adiabaticity with respect to many-body dynamics

When the lattice is turned on adiabatically it provides an adiabatic compression of the condensate (because of the enhanced confinement or higher trapping frequencies provided by the optical lattices). This means that the atoms have to redistribute in the lattice through tunneling to get the right Thomas-Fermi shape in the new potential. In [69] they

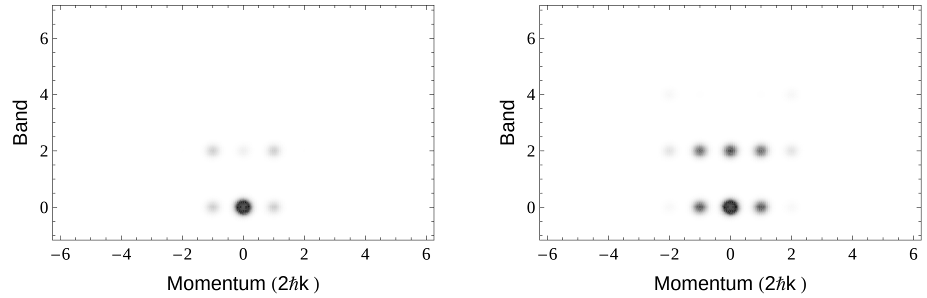

(a) Here $s = 5$, $q = 0$. In this case the atoms are loaded primary into the two lowest allowed bands (band 0 and 2).

(b) Here $s = 15$ and $q = 0$.

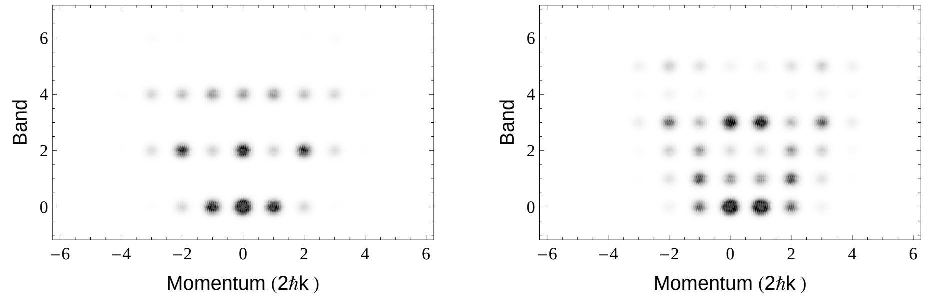

(c) Here $s = 40$ and $q = 0$.

(d) Here $s = 40$ and $q = 1$. Because $q = 1$ it is also allowed to load atoms into odd band numbers.

Figure 6.3: Population (or $|b_{l,n,q}|^2$) plotted for different momentum orders l and for different bands n (the dots are plotted with a Gaussian profile with a fixed width and the amplitude represents the population). When the lattice becomes deeper and deeper more bands are populated because we are turning on the lattice non-adiabatically. Notice that the population in the different bands (and in the different momentum orders) are time-independent. However if we turn off the lattice (also abruptly) we get momentum components (e.g. $-2\hbar k$, 0 and $2\hbar k$) but the populations in these peaks depend on the pulse time Δt because of coherent oscillations between the different bands—this is illustrated in figure 6.4. All the dots in e.g. band zero in (c) (where $s = 40$ and $q = 0$) can be represented by only one dot in the band structure shown in figure 6.2(c)—namely a dot in band $n = 0$ at $q = 0$. The normalization is $\sum_{l=-\infty}^{\infty} \sum_{n=0}^{\infty} |b_{l,n,q}|^2 = 1$.

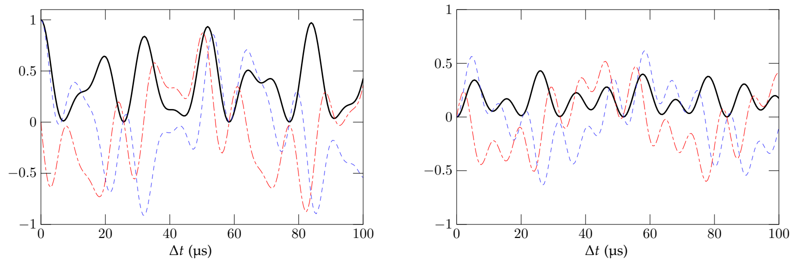

(a) Here the norm squared (black), the real part (blue) and the imaginary part (red) of $\sum_{n=0}^{\infty} b_{l=0,n,q=0}$ is plotted. The black curve shows the expected population in the $0\hbar k$ peak when the optical lattice is turned off after different times Δt.

(b) Here the norm squared (black), the real part (blue) and the imaginary part (red) of $\sum_{n=0}^{\infty} b_{l=\pm 1,n,q=0}$ is plotted. The black curve shows the expected population in the $\pm 2\hbar k$ peaks when the optical lattice is turned off after different times Δt.

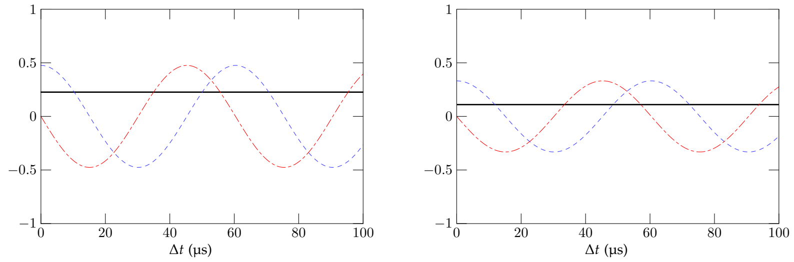

(c) Here the norm squared (black), the real part (blue) and the imaginary part (red) of $b_{l=0,n=0,q=0}$ is plotted.

(d) Here the norm squared (black), the real part (blue) and the imaginary part (red) of $b_{l=\pm 1,n=0,q=0}$ is plotted.

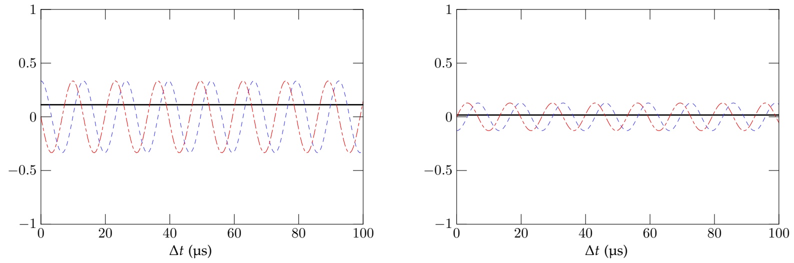

(e) Here the norm squared (black), the real part (blue) and the imaginary part (red) of $b_{l=0,n=2,q=0}$ is plotted.

(f) Here the norm squared (black), the real part (blue) and the imaginary part (red) of $b_{l=\pm 1,n=2,q=0}$ is plotted.

Figure 6.4: (a) shows the population in the $0\hbar k$ momentum peak as a function of lattice pulse time Δt. (c) and (e) illustrate where the complicated evolution shown in (a) comes from. (b), (d) and (f) are plotted for the $\pm 2\hbar k$ momentum peak ($l = \pm 1$). The depth of the lattice is $s = 40$. The black line in e.g. (c) corresponds to the dot shown in figure 6.3c for $l = 0$ and $n = 0$. Notice that only a few graphs are plotted to give the idea—(a) and (b) have also contributions from higher order bands (especially from $n = 4$—see figure 6.3c).

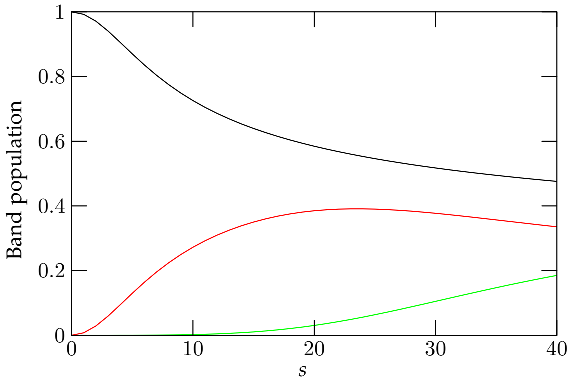

(a) The populations in band 0, 2 and 4 as a function of s for $q = 0$. Black: 0th band ($\sum_{l=-\infty}^{\infty} |b_{l,n=0,q=0}|^2$). Red: 2nd band ($\sum_{l=-\infty}^{\infty} |b_{l,n=2,q=0}|^2$). Green: 4th band ($\sum_{l=-\infty}^{\infty} |b_{l,n=4,q=0}|^2$). The graphs correspond to summing horizontally (for constant band number) in figure 6.3. The different populations sum to 1.

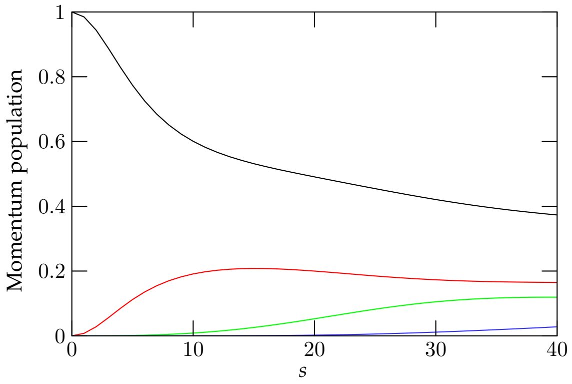

(b) Population in different momentum orders l as a function of the lattice depth s. Black: $l = 0$ ($\sum_{n=0}^{\infty} |b_{l=0,n,q=0}|^2$). Red: $l = \pm 1$ ($\sum_{n=0}^{\infty} |b_{l=\pm 1,n,q=0}|^2$). Green: $l = \pm 2$ ($\sum_{n=0}^{\infty} |b_{l=\pm 2,n,q=0}|^2$). Blue: $l = \pm 3$ ($\sum_{n=0}^{\infty} |b_{l=\pm 3,n,q=0}|^2$). These graphs correspond to summing vertically (for constant momentum order) in figure 6.3.

Figure 6.5: This figure shows the populations in the different bands n and in the different momentum orders l as a function of the optical lattice depth s in the case where the lattice is loaded abruptly.

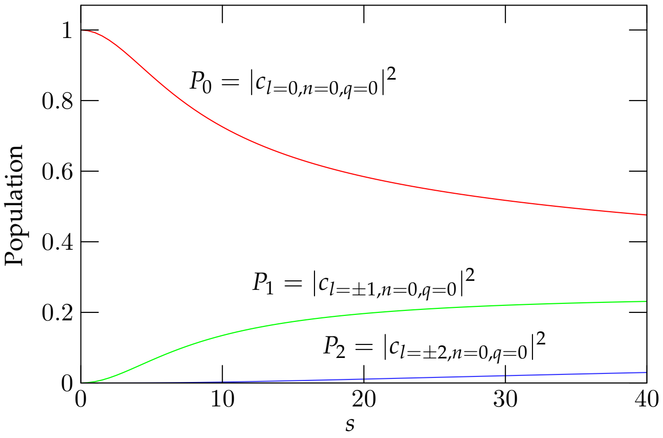

(a)

Figure 6.6: Populations of the momentum components in band zero as a function of s. This figure shows the population in the different momentum peaks (e.g. the $0\hbar k$ and the $\pm 2\hbar k$ peaks) if the lattice is loaded adiabatically and then turned off abruptly. When we turn off the lattice suddenly and wait for a certain TOF we are able to measure the population in the different momentum components. At $s = 0$ all the atoms will appear in the center peak $0\hbar k$, but as the lattice depth s is increased the population in e.g. the $\pm 2\hbar k$ peaks will start to increase. Notice that e.g. P_1 in this figure shows the population of only one of the two first order ($\pm 2\hbar k$) peaks (the coefficient c is normalized such that $\sum_{l=-\infty}^{\infty} |c_{l,n,q}|^2 = 1$).

find that the ramp time is limited by single-particle tunneling. They also find that with a ramp time of more than 100 ms the optical lattice is loaded adiabatically[9].

More details concerning adiabatic loading can be found in e.g. [81].

6.7 Hybrid optical trap

In the design process of the optical lattices we decided to make a setup where it is possible to change continuously from a dipole trap to an optical lattice—a hybrid dipole/lattice trap. We can do this because we can control the intensity/power of the retro-reflected beams independently of the ingoing beams. By ramping up the retro-reflected beams we go from a dipole trap to an optical lattice. In general this setup gives us a lot of possibilities for making different kinds of trapping configurations and change them dynamically.

A 3D optical lattice is normally loaded on top of a magnetic trap that prohibits the atoms from falling in the field of gravity (when the intensity of the optical lattice beams is low). With our setup we can e.g. study the superfluid to Mott insulator transition in a purely optical trap. This can be done by transferring a thermal cloud to the 3D optical dipole trap and then make a BEC by lowering the intensity of the beams and thereby allowing the most energetic atoms to disappear from the trap (just like in the case of RF-evaporative cooling—here we just make the evaporative cooling by lowering the laser intensity). In this trap we can now slowly ramp up the retro-reflected beams to a certain

---

[9] This value is found by investigating how the visibility of the interference pattern evolves for different lattice hold times (if the visibility decays slowly it is assumed that the ramping is done adiabatically).

hight and thereby get an optical lattice on top of the confining dipole trap. Since the atoms are trapped in a purely optical trap we can e.g. also do experiments with atoms in different m_F-states (which connot be trapped in the normal magnetic QUIC-trap).

The normal combined magnetic and optical trap has a harmonic contribution to the potential from the QUIC-trap but also from the optical trap due to the Gaussian profile of the laser beams (these two contributions to the trapping frequencies of the combined trap are comparable in size for the relevant parameters). This means that when we ramp up the laser beams the harmonic trapping frequencies changes. This results in an adiabatic compression (as already described in section 6.6) which in the end results in a temperature increase of the atoms [81].

We hope that we with the hybrid trap can ramp up the optical lattice and at the same time adjust the trapping frequencies to counteract the effects of the adiabatic compression. This way we might be able to load atoms into an optical lattice faster (or more adiabatic) than in the current experiments [82].

Currently we are looking into different loading schemes. Figure 6.7 and 6.8 show an example of one way of ramping up the optical lattice. In this scheme we simply maintain the harmonic trapping frequency during the loading of a one-dimensional optical lattice. We start with a dipole beam with a certain trapping frequency (and no magnetic trap). Then we ramp up the retro-reflection (and ramp down the ingoing beam) in order to get an optical lattice with the same "overall" harmonic trapping frequency as we started out with.

In chapter 7 I will go through how we have implemented this hybrid trap in the setup. However, unfortunately we are still only beginning to look into ways to improve the loading techniques, so I cannot say if the preliminary proposal presented above actually works. More work is needed on this issue.

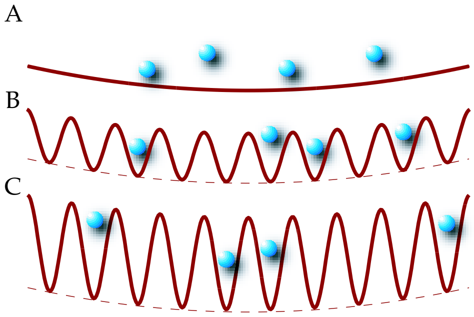

Figure 6.7: The transfer from a dipole trap to an optical lattice (in 1D). The optical lattice has a high (many kHz) trapping frequency because of the standing wave pattern but it also has a much smaller harmonic trapping frequency due to the fact that it is formed by Gaussian beams with a certain Raleigh length. During the loading of the optical lattice this harmonic trapping frequency will stay the same if the in-going power P_1 and retro-reflected power P_2 is ramped correctly (the "curvature" of the potential is kept constant like in the figure). This way the hybrid trap can provide both an optical lattice but also a confining harmonic (optical dipole) potential. Because of the harmonic potential we have no need of a magnetic trap to hold the atoms against gravity (not even at very small lattice depths).

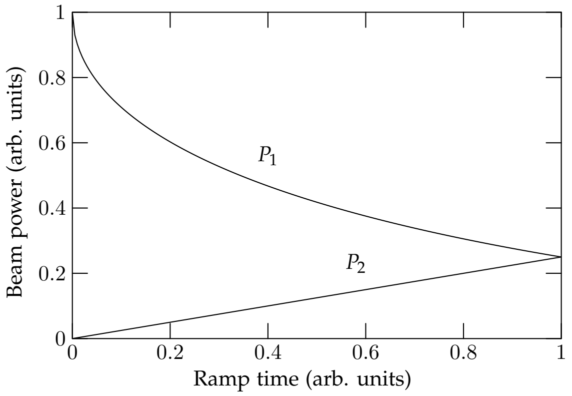

Figure 6.8: This graph shows how the power of the ingoing beam P_1 and retro-reflected beam P_2 should be adjusted in order to maintain the same "overall" harmonic trapping frequency. P_2 is in this example ramped linearly up (but we could use an arbitrary waveform) and P_1 is ramped down in accordance with eq. (C.24) (in appendix C it is calculated how the relationship between P_1 and P_2 should be). The two graphs meet at 1/4 of the initial power of P_1 as expected. After this point $P_2 > P_1$ and this is not possible in the setup (at least not for the x- and y-axis as we will see in chapter 7) and therefore not shown here. Initially P_1 should be chosen high so that the BEC can be trapped in a purely optical dipole trap (and not fall in the field of gravity). This means that if we turn up P_2 until $P_1 = P_2$ then we will be deep in the Mott insulator regime. Instead we can just stop somewhere on the graph (to a certain ramp time). Here the dipole trap will be strong enough to hold the atoms (if P_1 initially is chosen high enough) and moreover we can have an optical lattice that can be arbitrary small (such that the atoms are still in the superfluid region).

Seven

Optical lattice setup

7.1 Introduction

In this chapter it is described how a simple-cubic 3D optical lattice was implemented in our BEC setup. The optical lattice is based on a diode laser system operating at a wavelength of $\lambda = 915\,\mathrm{nm}$. Along the z-axis of the optical lattice we can detune the the two counter propagating beams and thereby get a moving lattice. A moving optical lattice makes it e.g. possible to load atoms with a non-zero quasi-momentum q.

I will start this chapter with a few comments on the design of the setup in section 7.2. After this I will describe the laser system for the optical lattices in section 7.3. In section 7.4 the part of the optical lattice setup that is placed around the science chamber on the vacuum table is described.

Christian Schori made the initial design of the experimental setup for the optical lattices and ordered much of the equipment. Later I took over and finished the design and built the setup.

7.2 Design considerations

We have chosen to build a 915 nm optical lattice mainly because of the available laser power (with diode lasers) at this wavelength. Initially it was discussed to use a certain "magic" wavelength for which the optical lattice would have the same trapping frequencies ω_i for both rubidium and lithium. Having the same trapping frequencies means that the gravitational sag ($\Delta y_{\mathrm{sag}} = -g/\omega_i^2$) for the two species will be the same. Because of this the atoms will be confined at the same "height" in the optical lattice even though the two species have different masses. It might sound like a negligible effect but the gravitational sag can easily differ by several hundred micrometers for two different species. In the case of rubidium 87 and lithium (6 or 7) the "magic" wavelength is around $\sim 807\,\mathrm{nm}$[1].

It was also discussed to make an optical lattice with the same depth (in terms of recoil energies) for the two species—this is fulfilled for another "magic" wavelength. For

---

[1] The "magic" wavelength can be calculated using one of the expressions for the trapping frequencies (e.g. eq. (C.6)). This gives the demand that $V_{Li}/m_{Li} = V_{Rb}/m_{Rb}$ in order to have the same trapping frequencies. V_{Li} and V_{Rb} can be calculated using eq. (6.4) where the lattice wavelength enters. The relevant wavelengths and linewidths for the different species can be found in [75]. A fun coincidence is that the "magic" wavelength for rubidium and lithium is very close to the "magic" wavelength for rubidium and potassium [83].

rubidium and lithium this wavelength is around 666 nm (i.e. blue detuned)[2].

These two examples illustrate that it is difficult to built an optical lattice that meets all possible demands when two species have to be trapped in the same lattice. The "magic" wavelength depends on the kind of experiment we want to perform. In the end we just chose the "non-magical" wavelength 915 nm. This wavelength is far detuned from the D1 and D2 lines in rubidium (and lithium) which means that the scattering rate is low. In the future we could e.g. make an extra horizontal dipole beam through the science chamber around 807 nm to set a confining potential in the direction of gravity and thereby compensate for the different gravitational sags if lithium is implemented [83].

Two different regimes have to be considered when the necessary power for the optical trap is calculated. For the optical lattices it is often enough to be able to go up to 25 recoil energies. For our experimental parameters this can e.g. be done with 75 mW power and a beam waist of 120 µm or with 50 mW and a beam waist of 100 µm. This is easily achieved with our setup since we have between 100 mW and 200 mW in each beam. However for the dipole trap it is another story. Here we need all the power we can get since we want to be able to load thermal clouds into the dipole trap. So, basically we have more than enough power for the optical lattices, but only barely enough power for the dipole traps. Instead of increasing the power one can also decrease the beam waists in order to get a higher intensity of the beams. The intensity is inversely proportional to the beam waist squared (see e.g. eq. (6.5)), which means that we can gain a lot by decreasing the beam waists. The good thing is that we only need to have one strong dipole beam to hold the atoms against gravity (or only one axis with a small beam waist). We use the horizontal z-axis (the Ioffe-axis) for this purpose because the trapping frequencies of the dipole beam are best matched with the trapping frequencies of the magnetic trap along this axis. With a strong dipole beam along the z-axis we can hold the atoms against gravity but since the axial trapping frequency is very low we need a second beam to confine the atoms in the axial direction (this beam does not have to be very strong). We normally use the most powerful beam and also the smallest beam waist along the z-axis. So far we have used a beam waist of around 65 µm along the z-axis, around 100 µm along the x-axis and around 120 µm along the y-axis (this is also used in all the experiments presented in chapter 8). With these waists we have enough power to make a strong dipole trap (and a very strong optical lattice). We have not observed any significant differences between the optical lattice axes when we e.g. look a the diffraction patterns produced along the three different axes—luckily the *radial* extension of the BEC is the smallest, which fits good with having the smallest beam waist along the axial direction. However, in the future it might be necessary to adjust the beam waists depending on what kind of experiments we are performing. If the beam waist gets to small (compared to the BEC) it is no longer a good approximation to say that the optical lattice is described by a $\cos^2$ potential because of the Gaussian shape of the laser beams.

An important feature for an optical lattice is to have active control of the beam power along the different axes. Otherwise the laser power will drift in time and produce a vibrating trapping potential (and cause heating and losses). Also we need to be able to ramp up the lattice beams in a controlled fashion, since we often want to do the loading of the atoms in the optical lattice adiabatically (slowly) using a certain ramp type. Moreover, even though we use orthogonal polarizations there is always a little interference between the different axes. To minimize this effect it is important to use slightly different frequencies along the different axes. This way the residual interference between the beams along the different axes will be time averaged to zero, and for this reason not influence on the atoms. To solve all these problems we use an AOM along each axis. With the AOMs we can shift the frequencies, have active power control, and

---

[2] In this case we require that $s = V_{0,Rb}/E_{r,Rb} = V_{0,Li}/E_{r,Li}$ which means that $V_{Rb}m_{RB} = V_{Li}m_{Li}$.

ramp up the different beams exactly the way we want to.

The power in each lattice beam is controlled by a PID-circuit [84][3]. A detector measures the beam power that goes through the science chamber and the power is adjusted by the AOMs (and also by EOMs—described in the next section). The rise time of the PID-circuit in response to a TTL signal is a few microseconds. Up to around 100 kHz we can modulate the intensity of the optical lattice through the PID-circuit. This is done with a bias-tee on the input to the PID-circuit where a dc- and an ac-signal can be combined. If higher frequencies are needed we can bypass the PID-circuit (when doing this it remembers the last output value to the AOM-driver), and then we can put an ac-signal on top of this offset. Modulating the intensity of the optical lattices can e.g. be used to transfer atoms between the different Bloch bands.

For the moving z-lattice we need to have extremely good control over the relative frequency between the two counter propagating beams. We use two DDSs (already described in chapter 3) to generate the frequencies for the AOMs. This is really an ideal solution since they can be locked to each other and the relative frequency and phase of the counter propagating beams can be set digital. The AOM-drivers for the x- and y-axis are instead based on a VCO and they just put out a preset frequency (although it is possible to control the frequency with an analog $0 - 5\,\mathrm{V}$ input voltage). All our AOM-drivers have a "fast-off" feature that makes it possible to turn off the signal to the AOMs in a fraction of a microsecond (they are also home-built).

Hybrid trap (or EOM setup)

Along the z-axis (where the moving optical lattice is implemented) we have the possibility to go from a dipole trap to an optical lattice since we can control the amount of power in each of the counter propagating beams completely independently. However along the x- and y-axis (where we use retro-reflectors) we do not have this possibility.

Instead we have implemented a setup (on the x- and y-axis) where an EOM and a PBS is used to control the amount of light that is retro-reflected back through the science chamber. The PBS is placed just after the science chamber and the EOM is placed between the PBS and the retro-reflector (see e.g. figure 7.3a). Without going into details the EOM changes the polarization of the light which means that we can control the amount of light that is reflected on (or transmitted through) the PBS. With this setup we can go from a dipole trap to an optical lattice just by changing the control voltage to the EOM[4].

The EOM also changes the phase of the retro-reflected beam when it is ramped on. This means that the lattice sites will move when the retro-reflected beam is ramped up or down. However as long as we do the ramping relatively slow we do not expect this to be a big problem.

We are not aware of any other group with this kind of setup—however in the group of I. Bloch they used a similar setup (with a quarter-wave plate instead of the PBS[5]) in a famous experiment where they realized spin-dependent optical lattice potentials, which they used to "split" the wave function of a single atom (prepared in superposition of two internal spin states) and transported the corresponding wave packets in opposite directions [85].

---

[3] The PID-circuits were designed by measuring the response of the system to a TTL pulse. This way it is possible to calculate the appropriate gains and time constants.

[4] The EOM is controlled by a home-built $0 - 4\,\mathrm{kV}$ power supply that can be controlled with a $0 - 5\,\mathrm{V}$ input.

[5] This means that the retro-reflected (linear polarized) beam will always have the same power.

7.3 Laser system

A sketch of the laser system for the optical lattices can be seen in figure 7.1. The setup is placed on the laser table and the light is transfered to the vacuum table through single mode, polarization maintaining optical fibers. The system is based on a master diode laser[6] that sets the wavelength of the entire optical lattice (all three axes). A second (identical) slave diode laser is injection locked to the master and they can each deliver around 100 mW of optical power. This is not enough for the optical trap so the two lasers seed three TAs[7] (one for each of the three axes through the science chamber). This gives an optical output power between 500 mW and 1 W for each axis. However after going through optical fibers, AOMs and a lot of other optics the available power at the science chamber is only between 100 mW and 200 mW depending on the axis. The optical fibers are used as spatial filters since the beam qualities of the TAs are really poor.

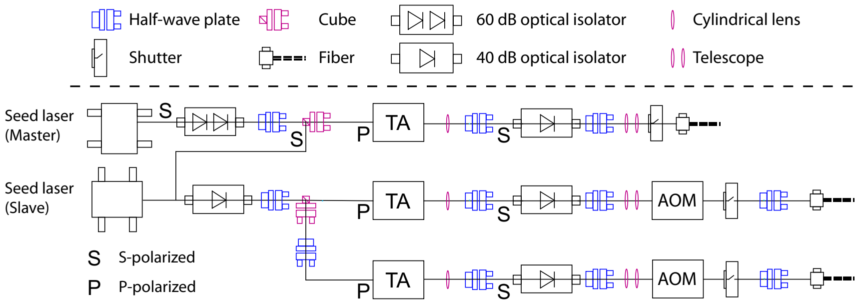

Figure 7.1: A sketch of the experimental setup on the laser table. The three TAs are seeded by two seed diode lasers—a master and a slave. On the x- and y-axis we have the AOMs before the optical fiber. On the z-axis the AOMs are placed after the fiber on the vacuum table. Since the TAs can be destroyed by feedback the cylindrical lenses and the half-wave plates after the TAs are slightly angled. To improve the AOM efficiency and the fiber in-coupling we have a telescope on all three axes.

The master laser is similar to our rubidium lasers (described in chapter 2). It is tuned to 915 nm which is the specified wavelength of the TAs. Both the master and the slave laser have a free-running wavelength of around 925 nm so we pull them $\sim$ 10 nm down with the grating on the master laser. Initially the idea was to have a 920 nm optical lattice (a compromise between the diode lasers and the TAs). However because we had problems with getting the specified amount of power from the TAs we changed the wavelength to 915 nm (more about this in section 7.3.2).

7.3.1 Locking scheme

Initially it was decided *not* to frequency lock the lasers because for most experiments it does not matter if the frequency of the optical lattices would drift a few GHz during a day—the corresponding drift in wavelength is negligible. However, when we started to make experiments with the system the lasers were often unstable and would typically only stay single mode for perhaps an hour in average. Because of this we decided to lock the lasers.

---

[6] The laser diode is an uncoated EYP-RWL-0940-00100-0750-SOT01-0000 manufactured by Eagleyard Photonics.

[7] Three EYP-TPA-0915-01500-3006-CMT03-0000 manufactured by Eagleyard Photonics.

The master lattice laser is locked to a Fabry-Perot interferometer (or cavity), which again is locked to the rubidium master laser. The reason for choosing to injection lock the slave laser was mostly because it demanded less optics/electronics than e.g. to lock to another cavity. The master lock is achieved by sending light from the rubidium master and the lattice master into the Fabry-Perot interferometer with orthogonal polarizations. On the other side of the interferometer the light is divided on a PBS and send to two detectors. Since the rubidium master is already locked using a FMS-lock we already have 10 MHz side-bands on the master. Because of this a Pound-Drever-Hall (PDH) locking signal for the cavity is formed (using the rubidium light) when we scan the cavity. A servo-amplifier gives an output to the piezo in the Fabry-Perot interferometer which ensures that the length of the cavity will stay constant. The current to the lattice master is also modulated with a small 10 MHz signal. When the master lattice laser is scanned another PDH locking signal is formed and the laser is locked to the Fabry-Perot interferometer. The generated side-bands on the optical lattice light could potentially be bad since it is also in the light that goes through the science chamber. However, we have tested the influence of this modulation by running identical runs with the modulation on and off. In this test we were not able to see a reduced lifetime or any difference in the interference peaks formed by the optical lattice.

An alternative locking scheme would be to use the 917 nm transition in cesium to lock to. It demands an extra 852 nm laser—more info on locking to this transition can be found in [86]. This way we would also know the exact frequency of the lattice beams.

The lattice laser system has in general been running fine for the last half a year. Often we just lock it in the morning and then it will stay single-mode and locked the entire day. Another Fabry-Perot interferometer is used to see if the slave laser is single-mode. Also we use a fiber optic spectrometer[8] to check whether or not the slave laser is in fact injection locked to the master (if it is *not* it will be free running at ~ 925 nm).

7.3.2 TA

A drawing of our TA setup can be seen in figure 7.2. The TA-chip is placed between two lenses. The lens in front of the TA is used to focus the beam down to the waveguide of the TA. The lens after is used to collimate the light from the TA, which is diverging a lot (especially in one direction because of the "tapered" shape of the TA). Outside the TA-housing a cylindrical lens is used to make the beam more Gaussian (see e.g. figure 7.1).

A peltier element, a thermistor and a PID-circuit is used to stabilize the temperature of the TA setup to 1 mK (like for the normal diode laser setup). We can run the TAs at a current up to 3 A which is much more than the few hundred mA consumed by our diode lasers. For this reason we have prepared the TAs for water cooling but most of the time we have been running them without since the laser table (and the surrounding air) is a large enough heat reservoir to keep the temperature of the TA-setup constant.

The TAs are extremely sensitive to seed and therefore especially the focusing lens in front of the TA must be positioned extremely accurate. Both lenses can be roughly aligned by looking at the spontaneous emitted light that comes out of the TA in both directions. The lenses are glued onto the base plate[9], and this can be very hard to do for the focusing lens without loosing the seed. The lens after the TA does not have to be placed with the same accuracy, however the beam shape from the TAs can be really bad and it can take days to find the right combination (and position) of the cylindrical lens and the following telescope to be able to get 50% of the light from the TAs into the optical fibers.

---

[8] Avantes AvaSpec-2048-USB2-FCPC.

[9] We have mostly used glue that hardens when it is exposed to UV-light—Ivoclar Vivadent, Heliobond.

Feedback to a TA can potentially destroy or damage it. Therefore (and for the stability of the optical lattice beams) we have Optical Isolators (OIs) after each TA and again after the fibers on the vacuum table. The spontaneous emitted light from the TAs will go back to the seed lasers so also here we have OIs. After the master laser we have a 60 dB OI (which is *not* overkill) since it also needs to shield the master from the injection locked slave laser.

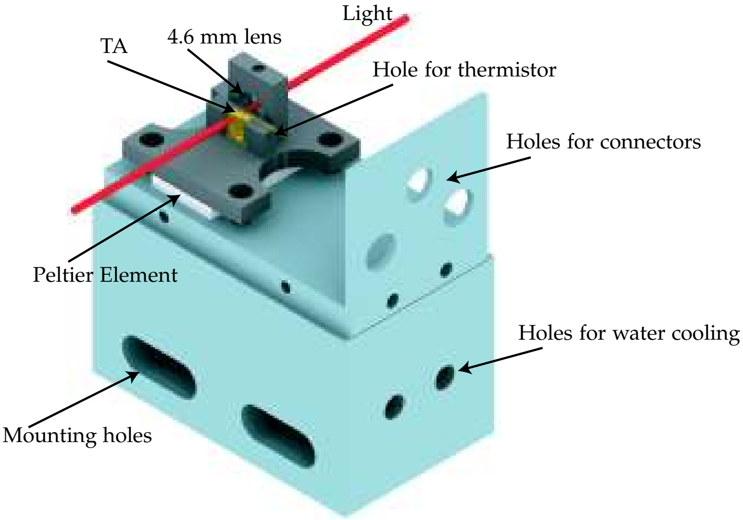

Figure 7.2: A drawing of our TA-design. Here the lens used for collimating the beam after the TA has been removed for clarity. Both lenses have a focal length of 4.6 mm. The blue parts are made in aluminum and the gray parts are made in a $CuNi_7Zn_{39}Pb_3Mn_2$ alloy (also used for the normal diode laser setup). This TA setup is identical to the one used for the rubidium laser system described in chapter 2.

The TAs have really given us a lot of problems and it has taken months to get the setup on the laser table working properly (mainly due to TA related issues). The really strange thing is that the rubidium TA is placed in an identical setup and we have never had any problems with it. It just worked right away and it has now been working fine for at least three years at the specified power. The lattice TAs are specified to provide 1.5 W of optical power but we have never managed to get this.

We have tested all kinds of things. Dependency on temperature, seed wavelength, different in-coupling lenses/telescopes and so on. These parameters all meant something to the output power but unless really crazy settings were used it did not change the output power a lot. The TAs are specified to saturate at around 50 mW of seed power. In the beginning we had only one seed laser so we were only able to seed each TA with around $20 - 25$ mW of power. In the end we decided to put in the second (already mentioned) slave seed laser which gave us $40 - 50$ mW of seed power for all three TAs. That helped a lot. Another mistake we had done in the design of the TA mount (see figure 7.2) was that with our initial mount we actually cut away some of the seed beam. A stupid mistake which we actually also did for the rubidium TA but here we have not changed it since it works fine even though we cut away 20% of the seed beam. These two changes gave us a factor of two and we are now between 500 mW and 1 W of output power for each TA (at 2.5 A). The lattice TAs can be run at 3 A (if this current gives less than 1.5 W of optical power) without significant degrading so we could run them a bit harder.

Our problems could be due to a manufacturing fault but it seems unlikely since all three TAs behave very similar. The most likely explanation we can come up with is that our setup is not stable enough. An interesting feature about our setup is namely that the in-coupling into the TAs depend critically on the current to the TAs. This means that if we have perfect seed at 0.5 A we have no seed at all at e.g. 2 A. According to the manufacturer this is not expected behavior.

We do not really understand why we have this problem since the setup is compact and seems very stable. One guess could be that the peltier element itself or the aluminum base block expands when the TA is turned on since more heat has to be transfered to the base block in order to keep the temperature of the TA-mount constant. This theory is supported by the fact that the seed beam also mainly needs to be steered in the vertical direction when the current is changed. On the other hand the seed is perfectly stable as long as we just choose the same current every day—this we find a bit strange if it is related to the temperature. However since the seed is stable at a constant TA-current one could argue that it does not really matter.

The real problem might be that the lens in front of the TA is mounted (for all three TAs) with a current of around 0.5 A. Since the seed depends on the current to the TAs the seed is (kind of) optimized at the wrong settings. The reason why we did this was mostly that we were afraid of destroying the TAs since it can be a problem to run a TA without seed[10].

The rubidium TA has the same characteristic behavior as the lattice TAs—here we also have to turn the TA up to around 60%-70% of the total wanted current before the output really starts to go up (which means that it is not really seeded before this point). In the case of the rubidium TA the lens was placed at a relatively higher current since this TA is only run at 1.2 A (this might explain why it works better).

Today I would definitely take the chance and insert the lens in front of the TA at 2.5 A and then see (if the TA survives this treatment) if this will give us the specified amount of optical power (1.5 W). However, I came up with this idea late in the process after we got the second seed laser, and at this time we had enough power to do experiments and the entire setup was working. It can easily take a few days to face the consequences of a removed focusing lens, so we decided not to take any chances (but I will test it soon since it is always nice to have more power).

7.4 Vacuum table setup

A sketch of the setup for the x- and y-axis on the vacuum table can be seen in figure 7.3a. After the fiber (and a collimation lens) the lattice light propagates through an OI and a few percent of the light is reflected (using a sampler) to a detector which monitors the power that goes into the science chamber. A telescope and a 500 mm lens (placed around 500 mm from the center of the science chamber) is used to produce the right beams waist (both regarding size and focus plane) inside the science chamber. The lattice light is combined with imaging light on a PBS and after the science chamber the beams are again divided on another PBS. The lattice light is sent through an EOM before it hits the retro-reflector that sends it back. Another detector measures the amount of the light that comes back from the retro-reflector and does *not* go back through the science chamber (since the EOM has changed the polarization of the light—this way we can of course also control

---

[10] It is impossible to place the lens in front of the TA without blocking the seed beam many times (since just touching the sub-micrometer positioning system can destroy the seed during the placement of the in-coupling lens).

the amount of light that actually *goes* back through the science chamber). The imaging light goes through the imaging system and hits a camera. All the optics (PBSs, lenses and mirrors) used in the imaging systems are 50 mm in diameter to avoid reducing the optical resolution of these systems.

We have chosen to reflect the lattice beam on both PBSs (both the one in front of the science chamber and the one after) because this gives a higher efficiency when the beams are reflected (S-polarized) on the PBSs. The price we have to pay for this is in terms of the purity of the polarization. However the beam is transmitted (P-polarized) through a really good PBS on the OI so it should already have a good polarization before it hits the PBS in front of the science chamber.

The z-axis is slightly different because here the beam from the optical fiber is divided into two before the light propagates through the science chamber in opposite directions (see figure 7.3b). Alternatively we could place the AOMs on the laser table and then use two optical fibers to send the light to the vacuum table. However we were afraid that with no active phase control we would run into troubles since the relative phase between light in two fibers can easily drift[11]. By dividing the beam on the vacuum table we are sure that we have control over the phase.

The AOMs are placed in such a way that the two AOMs on the vacuum table (for the z-axis) shift the same way, and the two AOMs on the laser table shift the other way. This way we can detune the z-plus beam compared to the z-minus beam and thereby make a moving lattice. Since the two other AOMs (x- and y-axis) shift the other way we only have to be careful to use different frequencies on these two. So far the x- and y-axis have been detuned around 10 MHz and it seems to work fine. We have tried to use the same frequency on the two axes and here the lifetime was drastically reduced when both beams were on—detuning is important.

Even though both the z-minus and the z-plus setup is finished we have not yet had time to test the moving lattice system. So far we have generated the z-lattice using a retro-reflector. However we plan to test the moving lattice feature soon.

In figure 7.3 a drawing of the setup on the vacuum table can be seen. More detailed drawings and comments on the setup is given in appendix D[12].

7.4.1 Beam propagation and imaging

The propagation of the lattice beams through all the optics on the vacuum table was also simulated before the setup was built. These simulations are based on the ABCD-matrix formalism for Gaussian beams [71]. In figure 7.5 and figure 7.6 the spot size of the y-lattice beam can be seen as it moves through the optical components on the vacuum table (the similar curves for the other axes are almost identical). It is my general impression that these simulations are in good agreement with what we actually have in the laboratory.

In the case of the y-lattice we produce a 120 µm beam waist inside the science chamber. If we want to change the beam waist by less than perhaps 20 µm it can be done by moving the lens on the fiber coupler to get a slightly different collimated/focused beam. Afterwards we correct for the change by moving one of the lenses in the telescope to focus the beam again (see figure 7.3). If we want to change the beam waist a lot we replace one (or both) of the lenses in the telescope. If the lens after the science chamber is placed at the right position (one focal length away from the center of the science chamber—in this case 140 mm) the beam will also have a 120 µm beam waist inside the science chamber

---

[11] We are aware of at least one other group that uses a similar setup where the AOMs are placed before the fibers (with no active phase control), so it is perhaps not crucial.

[12] I made detailed 3D drawings of the entire lattice setup before I started building. I am sure this saved me a lot of time because it would have been almost impossible to see if everything would fit in using 2D drawings. We use Autodesk Inventor for this kind of drawings.

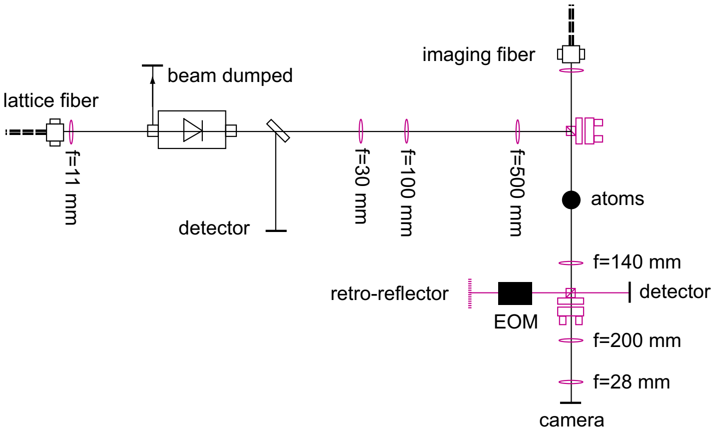

(a) x and y-lattice setup. Here the focal lengths of the lenses are specified for the y-axis—for the x-axis some of the focal lengths are slightly different.

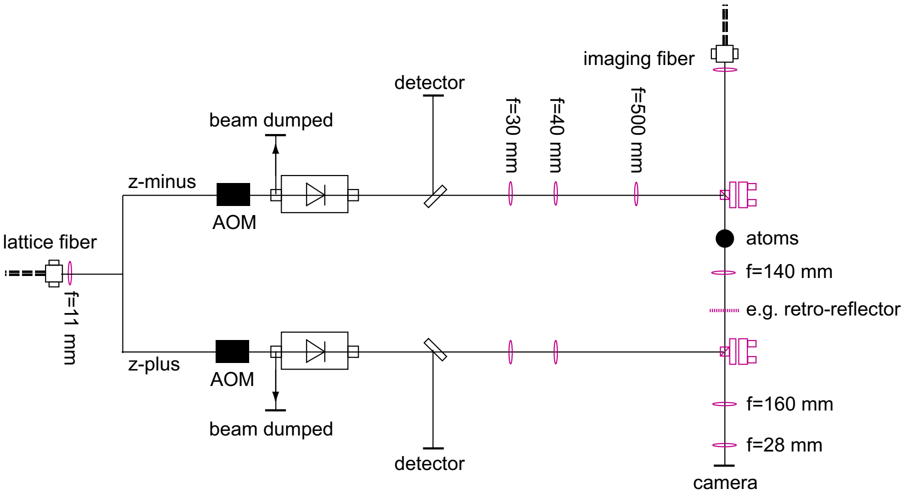

(b) z-lattice setup. We can choose to retro-reflect the z-minus beam if we do not want to use the moving lattice. In this case we use a dielectric mirror with a large transmission at 780 nm and large reflection at 915 nm as the retro-reflector.

Figure 7.3: A sketch of the setup on the vacuum table—not to scale. Only the most important parts are shown on these drawings. More information is given in the text.

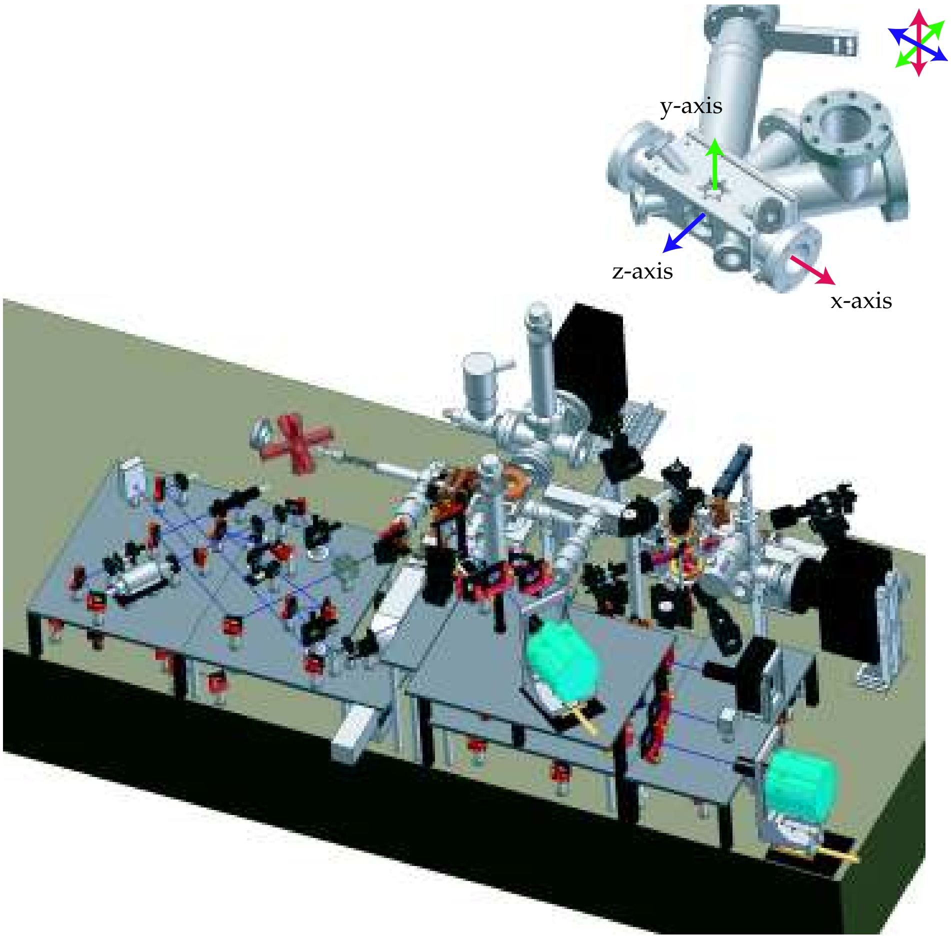

Figure 7.4: A drawing of the entire setup on the vacuum table. The optical lattice system is placed on three different levels around the science chamber. Detailed 3D drawings of each of the optical lattice axis can be seen in appendix D. We have focused on making the setup as stable as possible and therefore the breadboards for level two and three are placed on approximately 30 construction rails and around 50 angle brackets were used to stabilize the setup and attach the different breadboards to each other. Even though we have a lot of optics mounted around the science chamber the setup is designed in such a way that there will also be room for e.g. a second species imaging system, extra laser beams and so on. The polarization of the different lattice beams are shown on the inset—e.g. the x-axis lattice beam is polarized along the y-axis.

when it comes back (see figure 7.6). Since the spot size is almost constant between the lens and the retro-reflector (because the 140 mm lens is placed 140 mm from the beam waist inside the science chamber) the position of the retro-reflector is not very critical—however we have placed it around 140 mm from the lens anyhow. The x- and z-axis are basically made the same way.

We use the cameras to focus the lattice beams on the right planes and also to measure the beam waists, since the cameras are already focused (on the atoms) and the pixel sizes are well-calibrated. Because of chromatic aberration the focusing plane of the 780 nm light will be located at a slightly different position than the focusing plane of the 915 nm light. However, since the Raleigh length of the optical lattice is several millimeters we can live with a small focusing error. A way to test if the beam waist of an optical lattice beam is located at the right position is to trap atoms in a 1D dipole beam with imaging along one of the other axes. This way we can see if the "center of mass" moves away from the starting position if we hold the atoms in a purely 1D optical dipole trap. However, it only works for the horizontal beams and only if the beams are in fact aligned exactly horizontally (otherwise gravity will try to pull them in one direction). We have played a bit with this but it is not really necessary—just focusing the beams on the cameras works fine.

In figure 7.7 the y-imaging system can be seen (the x- and z-imaging systems are almost identical). The imaging setup is prepared for phase-contrast-imaging, which makes it possible to image a BEC non-destructively [87]. Phase-contrast-imaging can be achieved by inserting a phase plane (a glass substrate with a small dot at the center) in the imaging system as indicated in figure 7.7. This way we can selectively shift the phase of the unscattered light[13].

---

[13] The know-how to make this kind of phase plate is already at the University of Aarhus.

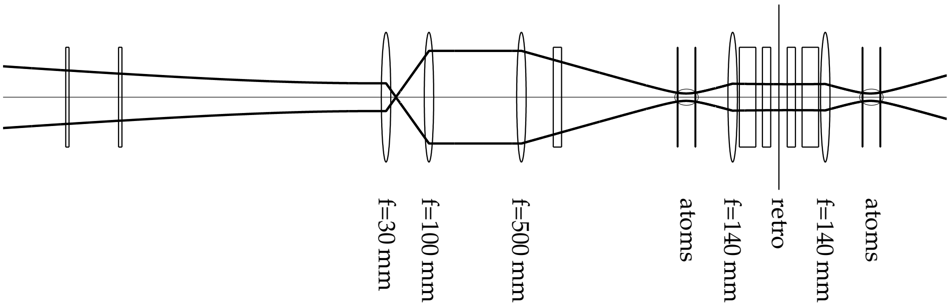

Figure 7.5: A schematic diagram of a Gaussian beam calculation (ABCD formalism) through the different y-lattice optics. The beam propagates from the left to the right. The starting point is a collimated/focused beam from the fiber mount. Initially the beam propagates through two PBSs (an OI). After this it goes through a (Keplerian) telescope and hits a 500 mm lens which is placed approximately 500 mm from the center of the science chamber. The telescope (the distance between the lenses is equal to the sum of their focal lengths) is used to magnify the beam before it hits the 500 mm lens which produces a beam waist inside the science chamber. In this case the magnification is 3.3 (because of the 30 mm/100 mm lenses). The lattice beam is then combined with the imaging beam on a PBS and before the beam hits the atoms it goes through a window on the way into the science chamber. On the way out it goes through another window, through a 140 mm lens placed approximately 140 mm from the center of the science chamber and 140 mm away from the retro-reflector. Between the 140 mm lens and the retro-reflector a PBS and an EOM is placed. Notice that on this drawing the beam is also traced back through the science chamber. The spot size can be seen more clearly in figure 7.6.

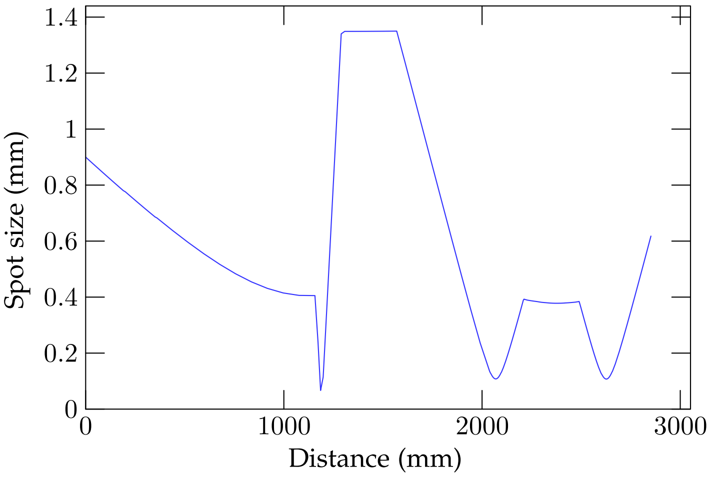

Figure 7.6: Here the spot size of the y-lattice beam is shown as a function of the propagation distance (the optics are the same as shown in figure 7.5). The beam from the fiber coupler is focused onto a point close to the telescope. The spot sizes of the x- and z-lattice beams have the same tendencies as shown here.

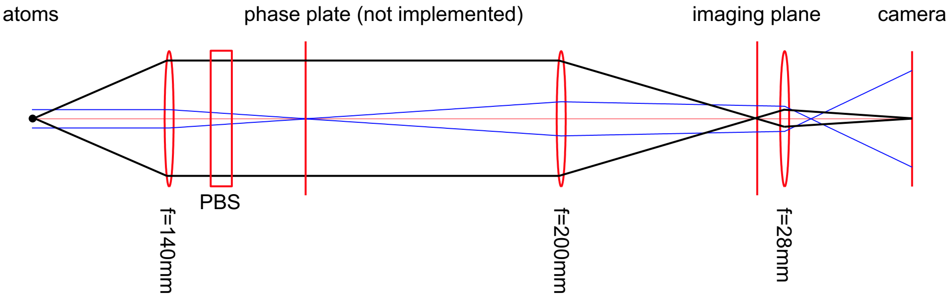

Figure 7.7: A sketch of the y-imaging system. The (unscattered) imaging beam (blue lines) propagates through the imaging system. The light has a large spot-size and it is collimated when it hits the atoms. The black lines indicate how the scattered light from the BEC propagates through the optics. The 140 mm lens is placed one focal length from the atoms. For this reason the distance between the 140 mm lens and the 200 mm lens can be varied. Around 200 mm from the 200 mm lens an imaging plane is formed (on the z-axis the mask for the DTA-camera is placed in this plane). A "macro" lens (Schneider Kreuznach CPN 2.8/28—focal length 28 mm) is placed between the imaging plane and the camera. A rough estimate of the magnification is 200 mm/140 mm = 1.4 for the achromatic lens pair times approximately 3 provided by the macro lens (this can be found in a datasheet for the lens given the placement). This gives 1.4 × 3 = 4.3, which is close to what we get if we use the measured pixel size of the imaging system (3.17 µm) and the pixel size of the camera (13 µm)—13 µm/3.17 µm = 4.1. The magnification of the x-imaging axis is also 4.1 and for the z-axis it is 4.8.

Eight

Optical lattices—methods and results

8.1 Introduction

The 3D optical lattice system has been working in all of 2008. However, it took some months to obtain the necessary experience with the system to do experiments. In this period we developed e.g. an alignment procedure for the optical lattice beams (it can be found in appendix E). In this chapter I will present some of the experiments we have performed with the dipole beams and the optical lattices. Most of these experiments have been performed within the last few months, and for this reason many of the results presented in this chapter are preliminary.

I will start with a few comments on different loading and turn-off techniques of the the optical trap in section 8.2. In section 8.3 I will present different calibration methods for calibrating the depth of an optical lattice. In the end of this chapter (section 8.4 and 8.5) I will present some other experiments we have performed with the setup.

Sung Jong and I have worked together on the results presented in this chapter.

8.2 Loading and turn-off techniques

Immediately after the RF-evaporative cooling sequence is done we “open up” the magnetic trap adiabatically by ramping the current in the QUIC-coils linearly from 300 A to 150 A and the current in the z-shim coils linearly from 0.35 A to 4.00 A in 500 ms. The BEC survives this treatment and in the reshaped trap the lifetime of the BEC is several seconds (since the number of three-body collisions is significantly reduced in this trap). The lower trapping frequencies of this trap (see e.g. table 5.3 in chapter 5) are also better matched with the trapping frequencies of the dipole trap. When the magnetic trap is reshaped, we load the optical trap on top of the magnetic trap.

The loading and the turn-off of the optical lattice can be done in different ways—here we will look at four examples:

a) Adiabatic loading, non-adiabatic turn-off (figure 8.1a): Here all the atoms are loaded into band zero, and interference peaks are seen when the BEC is released suddenly from the optical lattice (we are effectively measuring the momentum distribution of the atoms in the optical lattice).

b) Adiabatic loading, adiabatic turn-off (figure 8.1b): Here the optical lattice is turned off adiabatically to map the bands into individual momentum components. However, since the lattice was loaded adiabatically we only have atoms in band zero and for this reason we only see a $0\hbar k$ component.

c) Non-adiabatic loading, non-adiabatic turn-off (figure 8.1c): Here the lattice is pulsed on for a certain time Δt. This method is referred to as Kapitza-Dirac scattering and it was described in section 6.5.

d) Non-adiabatic loading, adiabatic turn-off (figure 8.1d): Here higher bands are populated because of the non-adiabatic loading. When the lattice is turned off adiabatically the quasi-momentum q is mapped onto the free particle momentum. This means that atoms in band n will be mapped onto the n-th free particle momentum state [77].

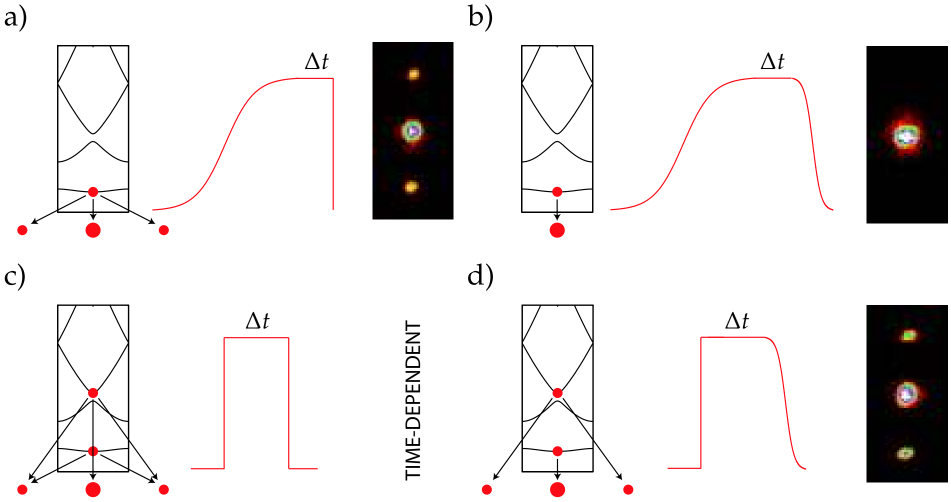

Figure 8.1: Different loading and turn-off techniques. In (a) and (b) the lattice is loaded adiabatically, so all the atoms go into band zero. In (c) and (d) the lattice is loaded non-adiabatically, so higher bands are populated (in this figure we assume that $q = 0$ and that the lattice is week so primary band zero and two are populated). In (a) the lattice is turned off non-adiabatically and the different momentum components are seen. In (b) the lattice is turned off adiabatically and all atoms are in band 0 so the BEC is just mapped onto the $0\hbar k$ momentum component. In (c) the optical lattice is turned off non-adiabatically and this corresponds to Kapitza-Dirac scattering—what we see depends on the pulse time Δt as described in chapter 6. In (d) the lattice is turned off adiabatically and here the atoms in band zero is mapped onto the $0\hbar k$ peak and the atoms in band two are mapped onto the $\pm 2\hbar k$ momentum peaks. Based on 140808 run 11, 24, 27.

8.3 Lattice depth calibration

In principle one can get the depth of the optical lattice just from the beam waist (located at the atoms) and the beam power (see e.g. eq. (6.6)). However, it can be difficult to measure these parameters exactly and an even bigger problem is that we do not know the power of the retro-reflected beam (all though it could be estimated by looking at the losses on the different optics), or how perfect the interference pattern really is—in principle we could have unwanted reflections that destroyed/reduced the standing wave pattern, bad polarizations, or a bad aligned optical lattice.

Fortunately the lattice depth can be accurately measured using different methods [40]. Some methods exist where a moving lattice is used to calibrate the optical lattice depth. However since we have a static optical lattice along the x- and y-axis we have only investigated methods that can be used for a static optical lattice. We have used four different calibration methods:

1. Parametric heating—atom loss
2. Parametric heating—population in band two
3. Kapitza-Dirac scattering
4. Expansion from an optical lattice (size of the diffraction peaks)

1. Parametric heating—atom loss

In this method we load the atoms adiabatically into a deep, one-dimensional optical lattice and turn off the magnetic confinement. While the atoms are only trapped in the optical lattice, we modulate the intensity (or the depth of the lattice) periodically with approximately twice the harmonic trapping frequency [88]. This results in parametric heating (and losses) since the atoms are exited from the ground band to the second excited band (only the 0th to 2nd band transition is allowed with intensity modulation). When we use this method we grab in-trap images.

The connection between the lattice depth V_0 and the axial trap frequency ω_a is calculated in appendix C. Here we will just use the result

$$V_0 = \frac{\omega_a^2 m}{2k^2}, \tag{C.17}$$

where m is the mass of an atom, $k = \frac{2\pi}{\lambda}$, and λ is the wavelength of the lattice laser.

From this we can find the depth of the optical lattice in terms of the recoil energies s

$$s \equiv \frac{V_0}{E_r} = \frac{\omega_a^2 m}{2k^2}\frac{2m}{\hbar^2 k^2} = \frac{\omega_a^2 m^2}{\hbar^2 k^4}. \tag{8.1}$$

This means that if we know the axial trapping frequency we also know the depth s of the optical lattice. When we calibrate the lattice using this method, we measure the atom loss as a function of the modulation frequency. For a harmonic oscillator potential the trapping frequency ω_a corresponds exactly to half the distance between "band" zero and "band" two. It is tempting also to use this result for a deep optical lattice, but it is a bad approximation. In figure 8.2 we compare the harmonic oscillator case with a band structure calculation. In the band structure calculation we take the energy difference between e.g. band 0 and 1, between band 0 and 2 and so on, at quasi-momentum $q = 0$ and $q = \hbar k$ (this is calculated using the theory described in section 6.3). No matter how deep the optical lattice is, it is a bad approximation to use the harmonic potential[1].

In figure 8.3 an example of this measurement can be seen. In this case we loose most of the atoms at a modulation frequency of around 62 kHz. The corresponding depth, using eq. 8.1 with $\omega_a = 31\,\text{kHz}$, is $s = 31.8$. Using the band-structure calculation we get $s = 42.7$ (for $q = 0\,\hbar\text{k}$)—the two different results can also be seen in figure 8.2. So, by

---

[1] We have seen articles and theses where a factor of 2 is used which gives rise to a wrong calibration.

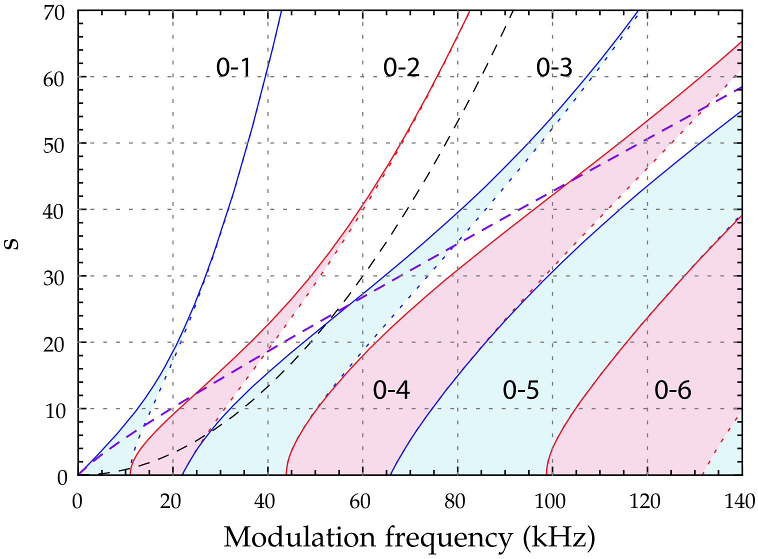

Figure 8.2: Lattice depth (in recoil energies) as a function of the modulation frequency. The red and blue areas are calculated by looking at the energy difference (in terms of the frequency) between band 0 and band 1–6 for different lattice depths s. It can e.g. be seen that for $s = 40$, the distance between band zero and band two is around 60 kHz. The left side of the different areas (or bands) are limited by a solid line indicating quasi-momentum zero ($q = 0$), and the right side of the areas are limited by a dotted line indicating $q = \hbar k$. The dashed, thick blue line indicates the "height" of the optical lattice—everything above the line is below the top of the lattice potential. E.g. for $s = 40$ there are four bands (including band zero) below the top of the potential, but band 4 (and higher bands) is outside (this can also be seen in figure 6.2c). The black dashed curve (part of a parabola) is for the case where the trapping frequency is equal to half the modulation frequency (harmonic oscillator case)—so it is eq. 8.1 with ω_a equal to half the angular modulation frequency.

using the harmonic potential approximation we get in this case a value that is 25% too low.

We can compare this result to the calculated value using the knowledge about the beam power and waist. In this case the beam waist was measured to 120 µm and the ingoing beam power was measured to 135 mW yielding a lattice depth of $s = 44.6$ (using C.15).

We have measured a lot of parametric heating profiles like the one shown in figure 8.3. With higher amplitudes we see a number of peaks, which we believe can be explained using figure 8.2. Also, we have studied the shape of the peaks. However we are not done with this analysis, so I will not go into details.

2. Parametric heating—population in band two

We can make a small twist to the parametric heating method just described (this is basically the same as what they do in [79]). If we load the lattice adiabatically and also turn it off adiabatically, we will normally only see one peak because the atoms are loaded into

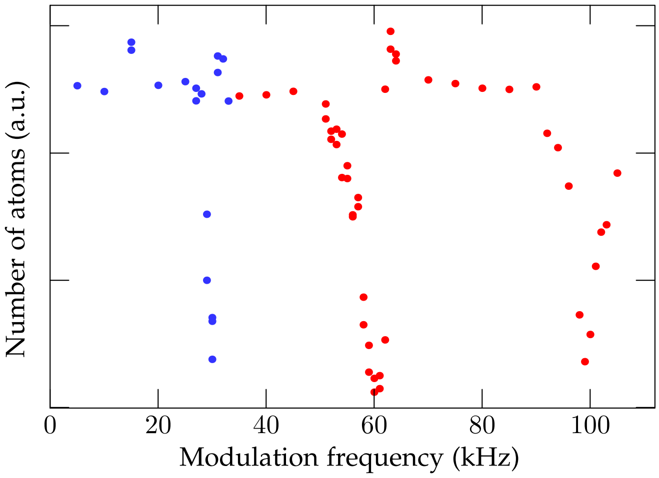

Figure 8.3: This graph shows a typical curve for a parametric heating measurement used to calibrate the optical lattice depth. Normally we only focus on the second peak/dip (the 0th to 2nd band transition). The red points are measured with a modulation amplitude of around 7% and the blue points around 20%. We believe that the blue peak is a "two photon" transition from band zero to band two (and not the forbidden 0th to 1st band transition) located exactly at half the real 0th to 2nd band transition frequency. The third peak is the 0th to 4th band transition. The data is from the y-lattice (beam waist 120 µs and a beam power of about 135 mW, which gives a theoretical value around $s = 44.6$). The 0th to 2nd band transition is here located around 62 kHz which corresponds to $s = 42.7$. Based on data from 110608 run 141-233.

band zero (as seen in figure 8.1b). However, if we now put on an intensity modulation while the lattice is on, we will transfer atoms into band two, if we hit the resonance frequency. In this case the atoms in band two will be mapped onto the $\pm 2\hbar k$ peaks because we turn off the lattice adiabatically (like in figure 8.1d).

This experiment is done in figure 8.4. Here it is clearly seen when we hit the resonance frequency. The advantage by using this method is that it is independent of the total number of atoms, since we always compare the number of atoms in the first order peaks to the center peak. Also, the width of the resonance in this case is more narrow than in the normal parametric heating experiment.

3. Kapitza-Dirac scattering

A really nice way of calibrating the optical lattice depth is to use Kapitza-Dirac scattering (this was described theoretically in section 6.5). This method involves pulsing on a one-dimensional optical lattice for a short time Δt and looking at the interference pattern (or scattering orders) afterwards. This method is also described in e.g. [79, 80].

Figure 8.5a shows typical images of the different diffraction orders when this method is used. Of course this pattern can also be generated in 2D or 3D (see figure 8.5b). For a deep lattice ($s = 40$) we can reproduce the pattern up to a few 100 µs—however there is a decay/dephasing of the interference peaks for longer pulse times. For small s the dephasing times are significantly longer.

Figure 8.6 shows examples of the method used for different lattice depths s. In all

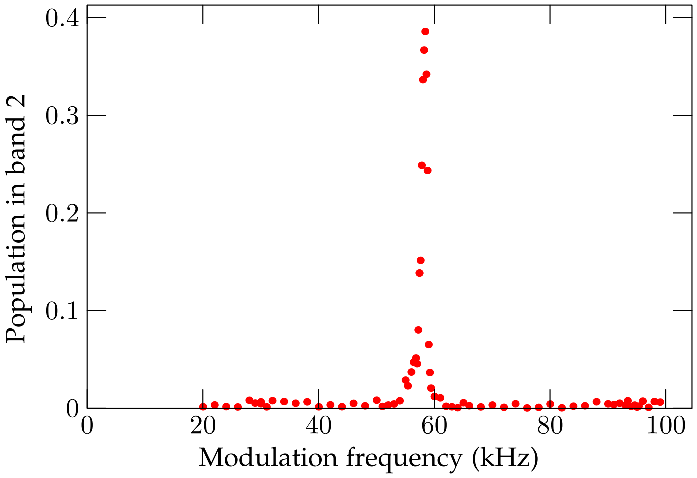

Figure 8.4: The population in band two as a function of the modulation frequency. Here we load the lattice adiabatically and also turn it off adiabatically (we use an s-curve for both ramps). While the lattice is on we use a 3% intensity modulation, which is on for 10 ms. The peak is located around 58 kHz and at this frequency about 40% of the atoms are transfered into band two. Notice how narrow the peak is compared to the “normal” parametric heating experiment (in this case the FWHM is only around 1 kHz). For this modulation depth we only see the 0th to 2nd band transition. We have not had time to try this method with a larger modulation, but we expect to see more peaks/transitions in this case. Based on 200808 run 166-277.

cases the different momentum components are plotted as a function of pulse time Δt. The different momentum components are fitted to the theoretical expression $|\sum_{n=0}^{\infty} b_{l,n,q=0}|^2$, given in section 6.5 (see e.g. also figure 6.4). The only unknown (if the data is normalized) is the lattice depth s. Figure 8.6 indicates how good this method is—there is a remarkable agreement between the measurements and the theory[2].

This method can be used for all lattice depths which can be seen in figure 8.7.

4. Expansion from an optical lattice (size of the diffraction peaks)

The fourth method we have used is based on measuring the lattice depth by looking at the “normal” diffraction peaks. The BEC is adiabatically loaded into a one-dimensional optical lattice and the traps (both the optical and the magnetic trap) is suddenly switched off (like shown in figure 8.3a). In [89] it is shown that the lattice depth $V_0 = sE_r$ can be calibrated using

$$s = \frac{16}{[\ln(P_\pm)]^2} P_\pm^{-1/4} \quad \text{for } s \gtrsim 5, \tag{8.2}$$

---

[2] In figure 8.6 we fit to the band theory introduced in section 6.5 but alternatively one can approximate the different orders with Bessel functions as done in [80]. This way they are able to describe the first microseconds of the evolution of the different orders. However since it is easy to do the real calculation one might as well fit to the real expressions, which also should give a smaller uncertainty since our fit does not only rely on one “period” but many.

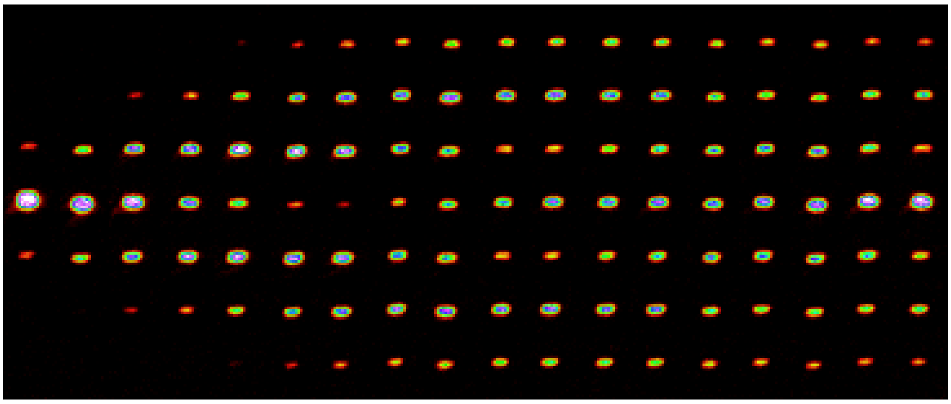

(a) One-dimensional Kapitza-Dirac scattering. Here we have used the y-lattice and x-axis imaging. The images show the different momentum orders for different pulse times Δt. To the left $\Delta t = 1\,\mu\text{s}$ and then increasing in steps of $1\,\mu\text{s}$ (to $\Delta t = 18\,\mu\text{s}$ to the right). We can use this experiment to calibrate the depth of the optical lattice by looking at the size of the different momentum peaks. In this case the depth is around $s = 40$. Based on data from 030608 run 185-202.

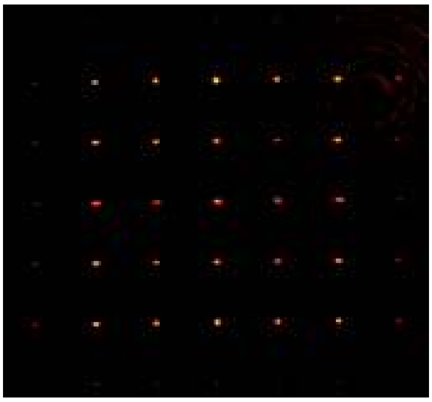

(b) Two-dimensional Kapitza-Dirac scattering. Here we have used the x- and z-lattice (and y-axis imaging). The pulse time is $\Delta t = 100\,\mu\text{s}$. So far we have *not* tried to use the two-dimensional experiment to calibrate the optical lattice. Based on 150408 run 91.

Figure 8.5: Kapitza-Dirac scattering in an optical lattice.

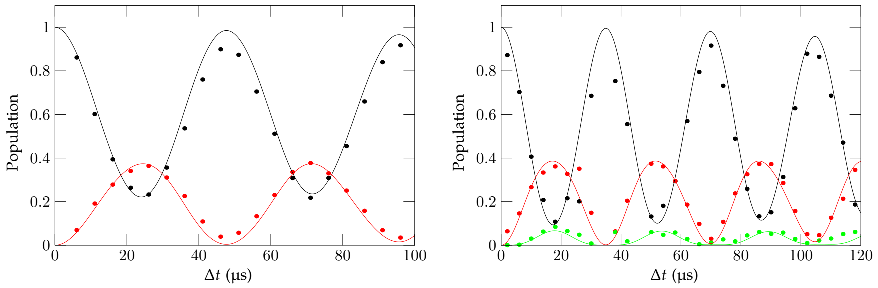

(a) $s = 9.7 \pm 0.1$. Based on 100608 191-209.

(b) $s = 14.8 \pm 0.1$. Based on 030908 run 120-149.

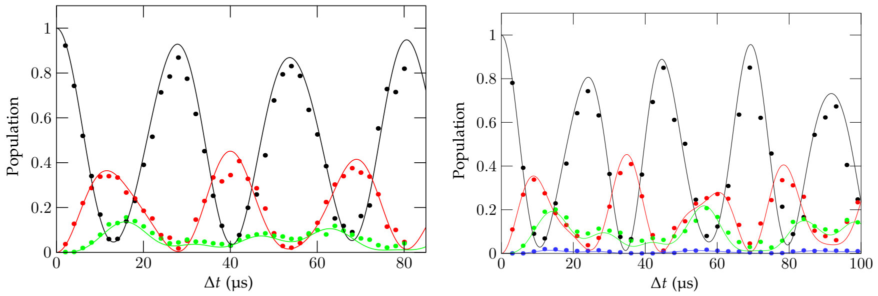

(c) $s = 20.4 \pm 0.1$. Based on 220808 run 150-192.

(d) $s = 25.4 \pm 0.1$. Based on 020908 - run 112-161.

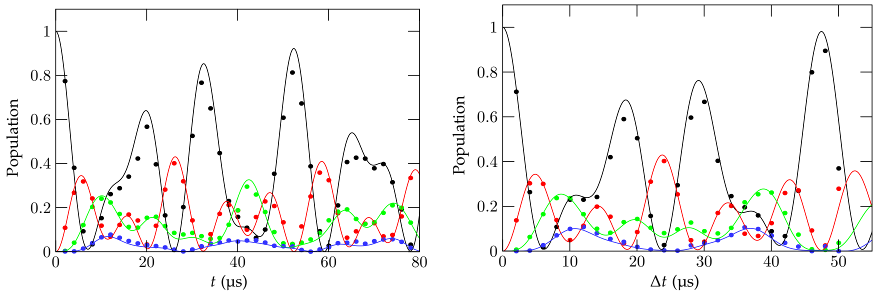

(e) $s = 39.0 \pm 0.1$. Based on 290808 run 74-118.

(f) $s = 44.9 \pm 0.1$. Based on 280808 run 8-45.

Figure 8.6: Typical curves for the momentum components as a function of the pulse time Δt in the case of Kapitza-Dirac scattering. Black: $0\hbar k$ peak, Red: average of $\pm 2\hbar k$ peaks, Green: average of $\pm 4\hbar k$ peaks, Blue: average of $\pm 6\hbar k$ peaks. Normalized and fitted to just one unknown, namely the lattice depth s. We have only seen groups using this method for a relatively week lattice, however as indicated here, it also works fine for deep lattices. Our PID-circuit is not fast enough to pulse the lattice on for e.g. just 1 μs. Therefore we "calibrate" the lattice a few seconds before the measurement is done (in the RF-evaporative cooling sequence), and then fix the calibrated output voltage for the AOM-driver in order to be able to pulse the lattice on fast (using the hold-TTL and fast-off feature on the power controller). This will result in fluctuations of the lattice power from run to run—however it is hard to see these fluctuations here so they are small.

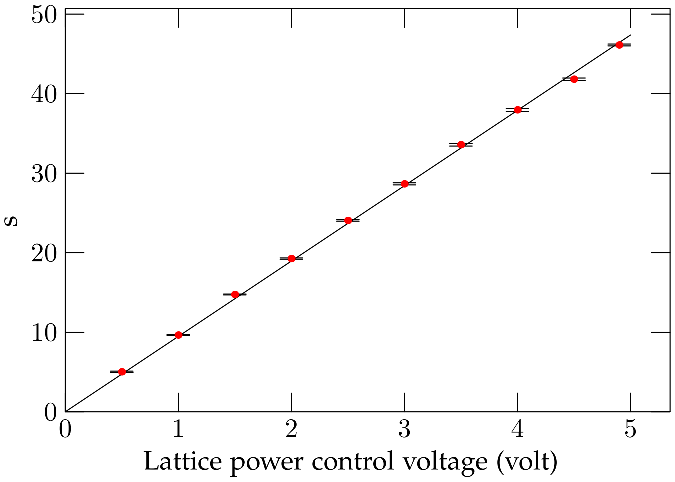

Figure 8.7: Optical lattice depth as a function of the control voltage to the PID-circuit controlling the power of the optical lattice beam. s is measured using the Kapitza-Dirac method. This graph shows that the Kapitza-Dirac method can be used for calibrating the optical lattice for all depths. The fit has been forced to go through the origin (there is a small offset on the power controller but it is only a few 10 mV, so it is neglected here). Based on 100608 run 038-261 (all these measurements were performed the same day to be sure that e.g. the detector was not moved and that the gain was not changed).

where $P_{\pm}$ is the number of atoms in one of the $\pm 2\hbar k$ peaks compared to the number of atoms in the $0\hbar k$ peak (since the 1st order peaks are symmetric, it should yield the same value no matter which of the first order peaks that is used—in practice we use the average value).

In figure 8.8 we do a (preliminary) test of this method in a well-calibrated optical lattice. Although this method gives fairly good results, it has not the same precision as the Kapitza-Dirac method or the parametric heating methods. For this reason we (so far) only use this method if we want to have a rough estimate of the depth of the optical lattice (actually CCS can automatically estimate the depth of the lattice based on this method).

Conclusion

In the daily work we use the Kapitza-Dirac scattering method because it is easy to use, works for all lattice depths and gives a great precision. In general we take 10–20 data points to get the uncertainty in s below 0.1. Alternatively the parametric heating method (where we measure the population in band two) also works very nicely, and here we do not have to take as many points (at least if we know approximately where the peak is located), however the uncertainty is slightly higher than in the Kapitza-Dirac method.

We have compared the Kapitza-Dirac method to the parametric heating method, and we have found very good agreement in all our measurements (it was from the beginning our goal to find at least two different methods that gave the same result). Also, when we compare to the calculation based on the measured beam waist and power, we get a good agreement along all three axes (within 5–10%). This indicates that the overlap between the beams and the quality of the interference pattern is good—but it is the measured depth we trust.

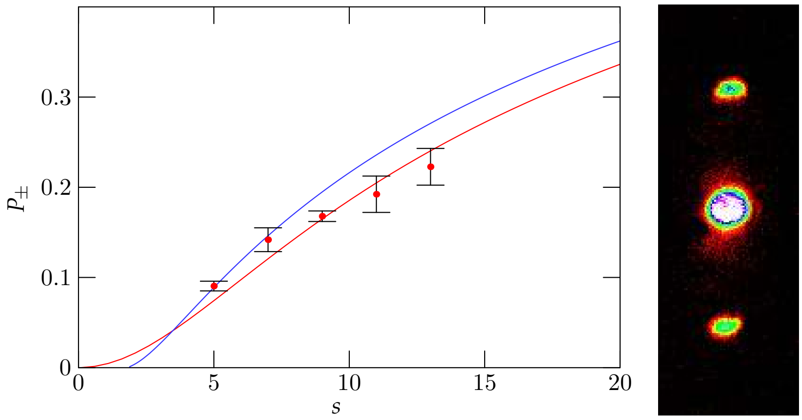

Figure 8.8: Here eq. 8.2 is plotted (red line) together with experimental points. In this case we have calibrated the optical lattice using the Kapitza-Dirac method before we did this test. The image shows an example of a picture grabbed for this data series (one of the pictures for the $s = 7$ measurement). There is a reasonable agreement between the data points and the theoretical line but the uncertainty in this method is large. When the lattice is increased to around $s = 15$, it becomes very hard to resolve the first order interference peaks from the center peak, so s has to be between $s \sim 5$ to $s \sim 15$—to get clearly separated peaks we normally use the method between $s = 5$ and $s = 10$. The blue curve is a band structure calculation—here we have just divided P_1 with P_0 in figure 6.6. Based on data from 100608 run 263-294 (the example image is run 271).

8.4 Lifetime and optical BEC

In figure 8.9 we have tested the stability of the optical lattice system. A BEC is formed and loaded into a one-dimensional optical lattice and the magnetic trap is turned off. After 40 s we still have atoms (not a BEC) left in the lattice, which indicates that the lasers are stable and that the technical noise is low.

In figure 8.10 a BEC is formed in the magnetic trap and loaded into a three-dimensional dipole trap. The BEC survives for more than 7 s in the dipole trap before it heats up and becomes thermal. Alternatively we can sweep the RF-frequency down to e.g. 50 kHz to 100 kHz above the trap bottom and then form the BEC in the dipole trap by doing the rest of the evaporative cooling here—an example is shown in figure 8.11.

In conclusion we can make an optical BEC which can survive for several seconds in the optical trap. We have not optimized these experiments in great details, so it might be possible to improve the lifetime further.

8.5 Other preliminary experiments

In the last part of this chapter we will look at some of the other experiments, we have performed with the optical lattice setup. We have done many of these experiments to improve our insight into the field (and not necessarily because we want to work on them in the future).

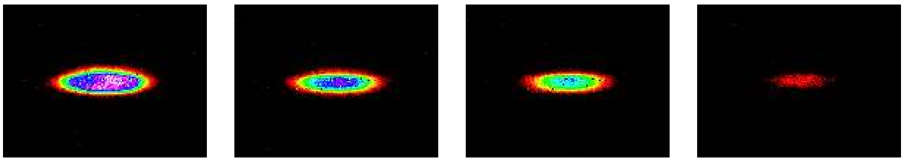

Figure 8.9: One-dimensional optical y-lattice—x-axis imaging. Atoms trapped for 1 s, 10 s, 20 s and 40 s (the magnetic trap was off in this experiment). The lifetime is more than 10 s. Based on data from 260208 run 160, 161, 162, 164.

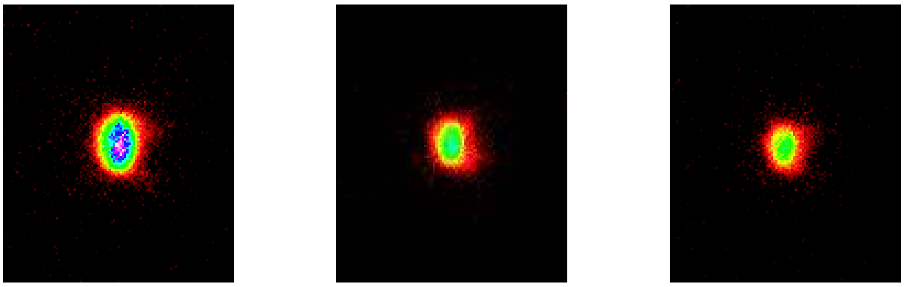

Figure 8.10: A BEC trapped in a purely 3D optical dipole trap for 2 s, 5 s and 7 s and afterwards allowed to expand for TOF=10 ms. In this case the BEC is created in the magnetic trap before the optical trap is turned on (and the magnetic trap off). Based on 231007 run 101.

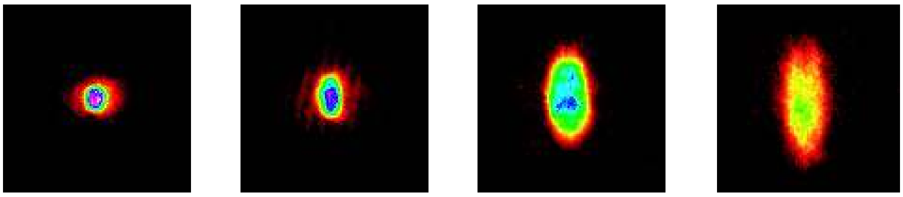

Figure 8.11: A BEC created in a 3D optical dipole trap. Here the RF-frequency is swept down to 815 kHz (around 40 kHz from the trap bottom, but it also worked with sweeping down to only 900 kHz). After this the BEC is formed in the optical trap—in this example we just turned on a dipole trap at a certain (static) depth which allowed the most energetic atoms to disappear. In this case the optical dipole trap has been on for 500 ms and TOF is 5 ms, 10 ms, 20 ms, 30 ms. Based on 241007 run 30, 29, 31, 35.

8.5.1 Microwave transfers

In the optical trap we can transfer atoms into m_F states which cannot be trapped in a magnetic trap. An example where we have transfered atoms into all the different m_F states (of $F = 2$) can be seen in figure 8.12. We have also done similar experiments with RF-transfers in a dipole trap [90].

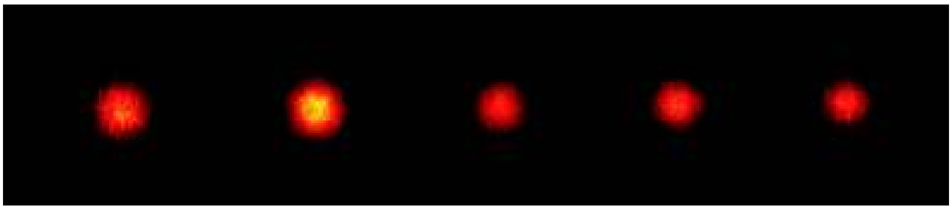

Figure 8.12: An example of microwave transfers in a one-dimensional optical lattice for atoms prepared in the $F = 2$, $m_F = +2$ state. This figure shows the different m_F components—$m_F = -2$ (left) to $m_F = +2$ (right). The different m_F states are separated using the Stern Gerlach method (TOF is 10 ms). We can transfer almost all of the atoms into each of the different m_F states using a number of microwave sweeps. If we wait too long between the sweeps, we start to populate other m_F states—we assign this behavior to spin exchange collisions. This experiment was done together with Sune. Based on 010408 run 138.

8.5.2 Superfluid to Mott-insulator transition

We have not studied the superfluid to Mott insulator transition in great detail so far, but seeing the transition was one of our large goals when we started to built the optical lattice setup, so we have (of course) also reproduced the experiment [29]. In figure 8.13 the superfluid to Mott insulator transition can be seen in a 3D optical lattice[3]. For low lattice depths clear interference peaks are observed. This indicates that the system is in the superfluid phase, where each atom is delocalized over the entire lattice (the number of atoms in each lattice site fluctuates), and phase coherence between the atomic wave-functions on the different lattice sites form the pattern. In the Mott insulator phase, an exact number of atoms are localized in each of the individual lattice sites and the phase coherence is destroyed.

It is our general impression that this experiment depends quite critically on the way the optical lattice is turned on, so this is definitely an experiment that we want to look into again in connection with the adiabatic loading experiments in the optical hybrid trap.

8.5.3 Population in band two

An example of population transfer from band zero to band two can be seen in figure 8.14. This experiment is similar to the technique we use, when we want to calibrate the lattice (i.e. measure the distance from band zero to band two) using the parametric heating method, where we transfer atoms into band two. However, in this case we sweep the modulation frequency through the resonance frequency to get a higher fraction of the

---

[3] There is an asymmetry between the peaks in these images which we did not see a few months ago, when we did this experiment the first time (the reason why these images are shown and not the old (and better quality) images is that at that time the lattice was not well-calibrated). We do not know why this asymmetry has appeared, but we believe that it is related to the z-lattice, since we do not see it when we grab images along the z-axis.

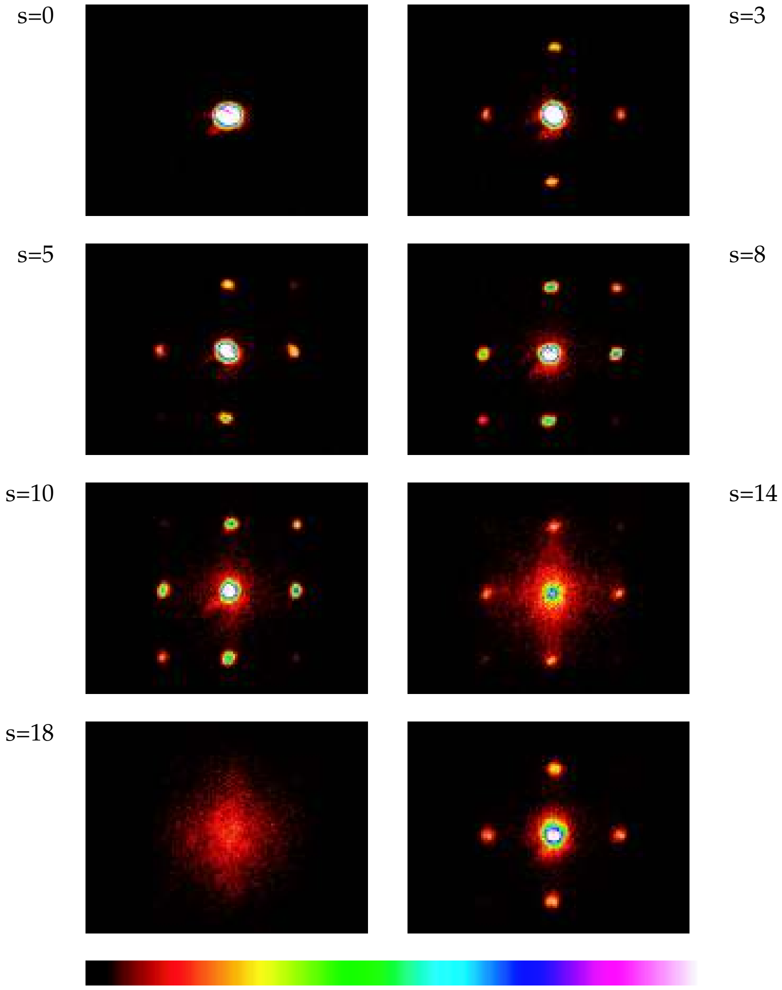

Figure 8.13: The superfluid to Mott insulator transition (x-axis imaging). The 3D optical lattice is adiabatically ramped up to different values of s and kept on for 20 ms before the lattice is suddenly turned off and the atomic cloud is allowed to expand for 12 ms TOF. It is difficult to define the exact point of the superfluid to Mott insulator transition, but up to around $s = 15$ clear interference peaks can still be seen. After this point the incoherent background is dominating. The last image shows the restoring of coherence when the system is brought into the Mott insulator phase and back again. In this image s is increased to 25 and then after a 20 ms hold time decreased to $s = 7$ in 10 ms. Based on 040908 run 14, 36, 12, 18, 13, 21, 23, 113.

atoms into band two. So far we have been able to transfer more than 65% of the atoms to band two, but this experiment has not been optimized at all, so we expect to be able to do more efficient transfers. Using several sweeps it should also be possible to transfer atoms to higher bands in a deep optical lattice in a controlled fashion.

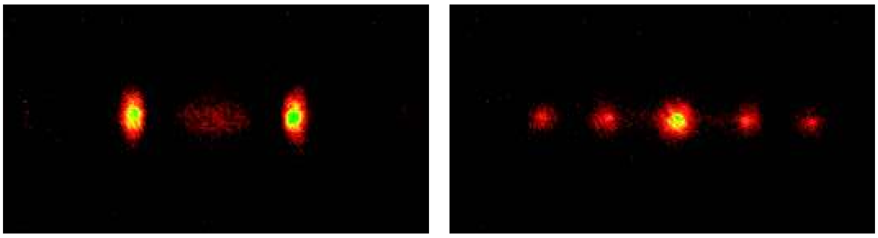

(a) Here the optical lattice is turned off adiabatically and the presence of $\pm 2\hbar k$ peaks indicates that we have transfered around 65% of the atoms to band two.

(b) Like (a) but here the optical lattice is turned off abruptly so the different momentum orders can be distinguished. These momentum orders come from both band zero and band two.

Figure 8.14: Here we transfer most of the atoms into band two. A one-dimensional optical lattice is loaded adiabatically to $s = 30.0$. After this we turn on a 49.0 kHz intensity modulation and sweep the depth of the lattice down to $s = 27.0$. Since the resonance frequency between band 0 and 2 in this case was 49.2 kHz (for $s = 30.0$) this corresponds to sweeping the frequency upwards through the resonance. In this case this procedure transfered about 65% of the atoms into band two. We could also have swept the modulation frequency instead of the depth but this was just easier. Based on 150908 run 250 and 251.

8.5.4 The first Brillouin zone

In figure 8.15 we have grabbed an image of the first Brillouin zone in a 2D optical lattice. In this case band zero is homogeneously populated and when the lattice is turned off adiabatically it is mapped onto the first Brillouin zone (the width of the squared momentum distribution is $2\hbar k$). We also want to do this experiment with a large population of atoms in higher bands in order to see higher order Brillouin zones in a 2D optical lattice.

8.5.5 Atomic "droplets"

Another type of experiments that we are currently looking into is what happens when a BEC is allowed to fall (i.e. accelerate) through a shallow one-dimensional, vertical optical lattice under influence of gravity [40, 89, 91, 92]. In the following we will describe the interaction time between the condensate and the optical lattice in terms of the Bloch period, which we will define as $t_B = 2\hbar k/mg$ ($= 1.024\,\text{ms}$ for our lattice), where g is the acceleration due to gravity, $k = 2\pi/\lambda$, λ is the wavelength of the optical lattice, and m is the mass of an atom. With this definition the (Bloch) velocity of an atom after one Bloch period of free fall will be $v_B = gt_B = 2\hbar k/m$.

Two different examples of this kind of experiment can been seen in figure 8.16 and figure 8.17. These figures can be understood in terms of "Bloch oscillations" and "Landau-Zener tunneling" [89, 92]. In these examples the condensate is accelerated in the field of gravity. At a specific time ($0.5t_B$) the condensate reaches the edge of the Brillouin zone.

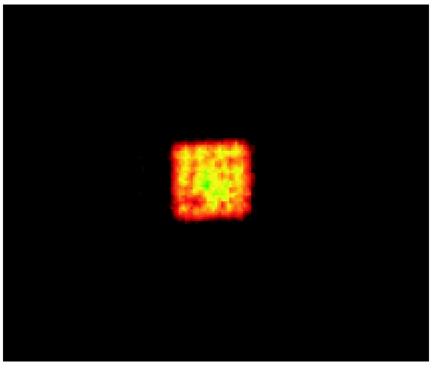

Figure 8.15: Here the first Brillouin zone in a 2D optical lattice is seen. A BEC is prepared in a 2D optical lattice and the phase coherence of the condensate is destroyed with a deep third optical lattice axis. The dephased condensate leads to a homogeneously populated first Brillouin zone. Here TOF was 7 ms. Based on 040908 run 8.

Depending on the depth of the optical lattice the condensate can either be Bragg reflected (and reappear at the opposite side of the Brillouin zone), or perform a Landau-Zener tunneling, where it "jumps" from band zero into band one (across the band-gab). In the latter case the condensate will not be reflected but instead continue up and out of the lattice. For a shallow lattice a large fraction of the condensate will undergo Landau-Zener tunneling and be lost from the optical lattice.

This kind of experiment can also be used to calibrate the depth of a (vertical) optical lattice as described in [89]. Here it is shown that the number of atoms trapped in a shallow, vertical optical lattice after a certain interaction time t_{int} is given by

$$N_{trapped} = N_0 (1 - e^{-a_c/g})^{t_{int}/t_B}, \tag{8.3}$$

where N_0 is the total number of atoms, and a_c is the critical acceleration given by $a_c = \frac{\lambda}{8\hbar^2}(\Delta E)^2$, where ΔE is the band gap between band zero and band one at the edge of the Brillouin zone. In eq. (8.3) it is assumed that t_{int} is an integer multiple of the Bloch period t_B. An example where this method has been used can be seen in figure 8.18. Here we have first calibrated a one-dimensional optical lattice using the Kapitza-Dirac method and afterwards checked it against the "droplets method". The two methods are in good agreement. However, again it should be stressed that this is only preliminary results. We will have to do a more profound comparison.

Some really interesting experiments can be made using this simple idea, where a condensate is accelerated through an optical lattice. Above we just played with different details and got quite different outcomes (although it is of course basically the same experiment). We want to work more on this and understand all the measurements presented above (and others) in full details. Also, similar experiments can be done with the moving optical lattice (but here we decide the acceleration—and the velocity).

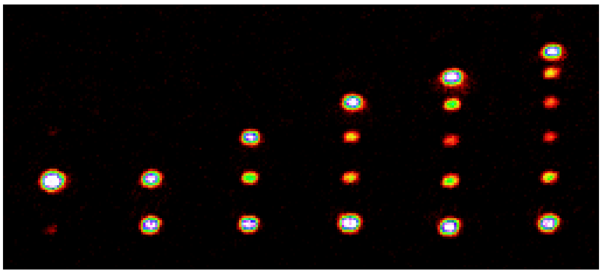

Figure 8.16: Atomic "droplets". These images are produced by letting a BEC fall in a one-dimensional, vertical optical lattice under influence of gravity. The optical lattice is loaded adiabatically and the magnetic field is then turned off. The acceleration/interaction time in the optical lattice is $0t_B$, $0.5t_B$, $1.5t_B$, $2.5t_B$, $3.5t_B$ and $4.5t_B$, respectively. The number of droplets can be controlled by changing the interaction time with the lattice (the time from the magnetic trap was turned off to the pictures were taken was 8 ms). The exact pattern of droplets can be interpreted in terms of Bloch oscillations and Landau-Zener tunneling. Although we understand the basic mechanisms we cannot explain all the details in this figure yet. Based on 280808 run 83, 82, 86, 88, 94, 96.

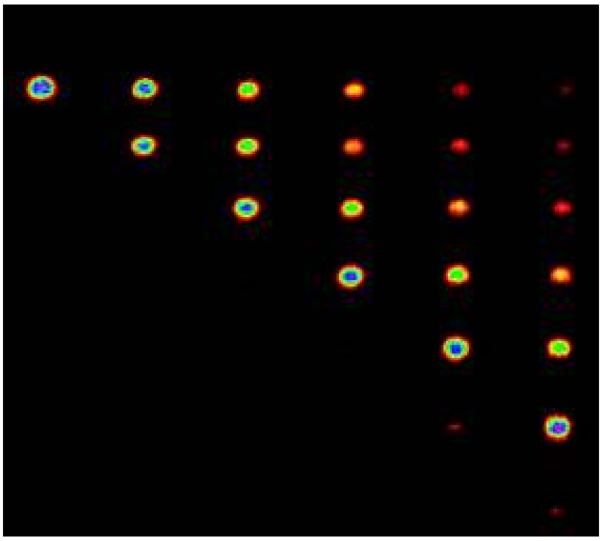

Figure 8.17: Atomic "droplets". Here the magnetic field is turned off and the atoms are allowed to fall for $0.5t_B = 512\,\mu\text{s}$ (at this time the atoms will see the lattice move with a velocity that corresponds to the momentum $\hbar k$—the edge of the Brillouin zone in the band structure picture) before we pulse the optical lattice on for a variable time ($0t_B$, $1t_B$, $2t_B$, $3t_B$, $4t_B$ and $5t_B$). After this we wait for 10 ms before we grab the image (so the time from the magnetic trap is turned off to the image is grabbed is varied). Here it is more clearly seen how we have Landau-Zener tunneling and Bloch oscillations, which result in a new free falling peak for each new image. The depth of the optical lattice was chosen in such a way that we loose approximately 50% of the atoms from the lattice (due to Landau-Zener tunneling) for each image (i.e. the "top" or "trapped" peak is always split into two equally sized peaks for each Bloch period). Based on 220908 run 225-230.

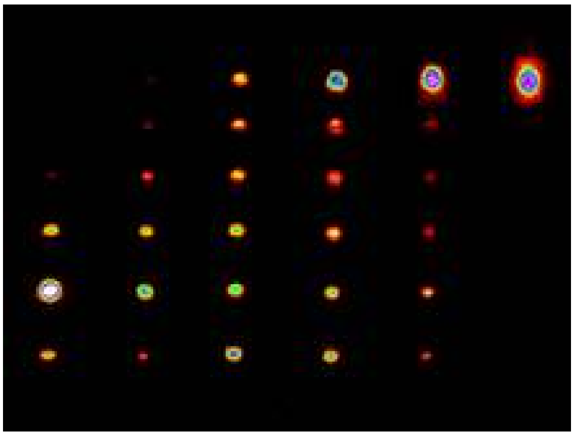

(a) The different images show the amount of atoms that are trapped (the “top row”) in the optical lattice for $s = 0.4$, 0.8, 1.2, 1.6, 2.0 and 2.6. In the first (left) image (where $s = 0.4$) no atoms are trapped in the optical lattice whereas in the last (right) image all of the atoms are confined in the lattice. Based on 230908 run 159, 161, 163, 165, 167, 170.

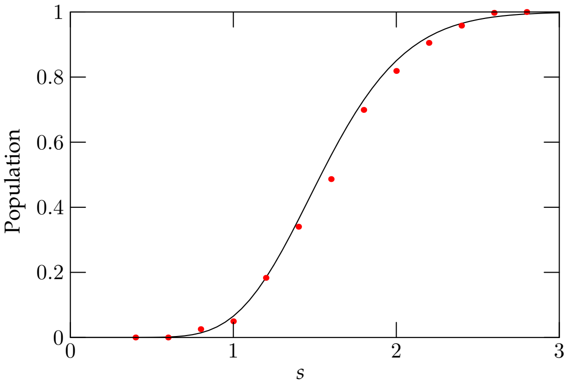

(b) Here the relative number of atoms trapped in the optical lattice is plotted as a function of s (the value of s is based on a “Kapitza-Dirac calibration”). Also a theoretical curve is plotted (no fitting is done). The data points follow the theoretical curve (eq. (8.3)) quite nicely indicating a good agreement between this method and the Kapitza-Dirac calibration method. Based on 240908 run 76-99.

Figure 8.18: These figures indicate how the lattice depth can be calibrated by the “droplets method”. A BEC is loaded adiabatically into a vertical, one-dimensional optical lattice of depth s. The magnetic trap is then turned off and the condensate is allowed to interact with the lattice for a certain (constant) time (here $t_{int} = 5 \times t_B = 5.121\,\text{ms}$). After this the lattice is turned off abruptly and a 10 ms TOF expansion follows (so the time from the magnetic trap is turned off to the image is grabbed is $5 \times t_B + 10\,\text{ms}$).

Nine

Conclusion and outlook

In this thesis it has been described how a BEC-apparatus and a three-dimensional optical lattices experiment were built and characterized. The current status of the laboratory is that everything is working and that we are doing experiments with the setup.

In the entire building process we have tried to think ahead and make flexible solutions that can be adjusted to meet many different demands. Some places we could have chosen faster or easier solutions, but we have always believed that if we want to build a quantum gas laboratory with all the state of the art features, we need to have a good basis—otherwise these experiments are just too complicated. The time, we have spent on making the experiment stable and easy to work with, will come back in the long term—this is also somehow demonstrated in this thesis—almost all the data presented here were made within a few months, although we were doing these experiments for the first time. This would not have been possible if we were not able to just turn on the experiment and start to take data.

In chapter 2 the basic BEC setup was described. The setup stands out by the fact that it is designed to do mixtures of two different species, where each species is prepared in its own MOT. So far only the rubidium 87 part has been implemented, but much of the current setup can be copied to the second species part (e.g. much of the control electronics and perhaps even the laser system).

In chapter 3 and 4 the computer systems (ECS and CCS) used to control the laboratory were introduced. These systems are capable of doing everything from controlling the current in the different coils to control cameras and move the optical lattice beams. We have focused on making user-friendly systems, that can be expanded and further developed. Right now we are working on a large expansion of ECS, which e.g. will make it possible for ECS to control much more equipment than currently.

In chapter 5 we looked at experiments performed with the basic BEC setup. Besides going through different characterization experiments, we also analyzed images of thermal clouds and BECs, and we saw that the system was able to make reproduceable BECs with 3×10^5 atoms.

The second part of this thesis dealt with BECs in optical lattices. In chapter 6 we calculated the trapping potential of different optical traps, and we also did different band structure calculations. Moreover, a proposal for decreasing the adiabatic loading time of a BEC into an optical lattice was presented in this chapter.

In chapter 7 the 3D optical lattice setup was described. In the setup we have implemented a moving lattice along one of the axes, and hybrid optical traps along the other two axes. This setup gives us many possibilities for changing the optical trapping potential dynamically—we can e.g. change continuously from a dipole trap to an optical lattice along all three axes.

In chapter 8 we looked at different experiments performed with the 3D optical lattice setup. So far we have mostly focused on different basic experiments because we wanted to be sure that our setup was working, and also we needed to learn many basic things. In the chapter it was e.g. described how we have studied different ways of calibrating the depth of an optical lattice, but also many other experiments were presented.

I am *very* glad that we got to do these optical lattice experiments before I had to end my PhD-studies. Although I also liked the building phase, it has been very fun to work on the optical lattice experiments and develop new ideas. Of course, I would have liked to do more experiments with the setup, but given the manpower we have had in the laboratory, it was just not possible.

9.1 Outlook

Currently the long-time planning of experiments in the laboratory is lacking due to the absence of a group leader. Instead Sung Jong and I have made some goals for the optical lattice experiments, which we are currently working on. Since the laboratory is now working, it has been important for us to focus all our energy on making results with the setup. During our work, we have gotten some ideas for different experiments, that we would like to do with the optical lattice setup. Here I will only mention two projects related to what has already been presented in this thesis.

We are currently working on a paper about calibrating the depth of an optical lattice. We want to compare the different calibration methods in details, and with the work we have already performed, we are close to having all the data. Although this subject is a *part* of quite a few articles within the field, we would have liked (when we started in this area) to have a central paper where many of the different methods are compared (using data) and described in details.

We also want to study and improve adiabatic loading into an optical lattice using the optical hybrid trap (described in chapter 6 and 7). We have not performed any experiments yet, but we are ready to do it, and I think we will quite fast be able to judge if this is something worth spending time on (which we of course hope).

A

Appendix A - The offset lock

The offset lock is used to offset lock the Cooler laser to the Master laser as described in section 2.3.3. In this appendix the relevant formulas and a schematics of the electronic circuit is shown. I will not go through our FMS-locking scheme since it is generally well known, but in the end of this appendix I will show the locking signals for the Master and the Repumper (which are both locked using a FMS-locking scheme)[1].

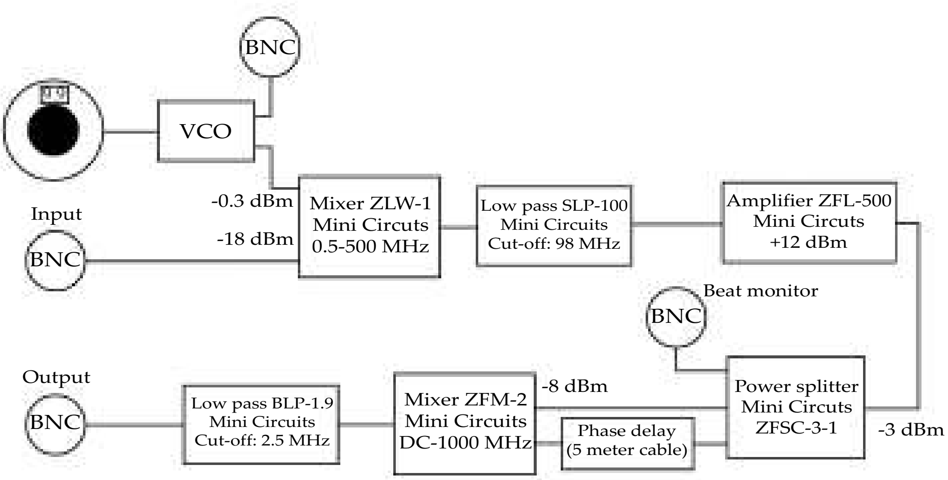

Figure A.1: A schematic overview of the offset locking system.

Light from the Master and the Cooler is mixed on a 300 MHz detector. The detector will see the beat frequency between the two lasers $\nu_{beat} = \nu_1 - \nu_2$. This means that the electronic signal from the detector will have the following form

$$V_s \cos(\omega_{beat} t + \varphi_{beat}), \tag{A.1}$$

---

[1] A detailed description of the FMS-locking method we use can be found in [50].

where V_s is the amplitude of the detector voltage, $\omega_{beat} = 2\pi\nu_{beat}$ and φ_{beat} is a phase delay. This signal is mixed with a signal from a VCO

$$V_{VCO}\cos(\omega_{VCO}t + \varphi_{VCO}). \tag{A.2}$$

Here V_{VCO} is the amplitude of the VCO-voltage, $\omega_{VCO} = 2\pi\nu_{VCO}$ and φ_{VCO} is a phase delay on the VCO signal.

These two signal are mixed (see figure A.1). The output from the mixer will have the following form

$$\begin{aligned} & V_s\cos(\omega_{beat}t + \varphi_{beat}) \cdot V_{VCO}\cos(\omega_{VCO}t + \varphi_{VCO}) \\ = \; & \frac{V_sV_{VCO}}{2}\Big[\cos\Big((\omega_{beat} - \omega_{VCO})t + (\varphi_{beat} - \varphi_{VCO})\Big) + \\ & \cos\Big((\omega_{beat} + \omega_{VCO})t + (\varphi_{beat} + \varphi_{VCO})\Big)\Big]. \end{aligned}$$

The second term in this expression is filtered out in a low-pass filter after the mixer ($\nu_{beat} + \nu_{VCO} \approx 494\,\mathrm{MHz}$). After the low-pass filter the signal is proportional to

$$\frac{V_sV_{VCO}}{2}\cos(\Delta\omega t + \Delta\varphi),$$

where $\Delta\omega \equiv \omega_{beat} - \omega_{VCO}$ and $\Delta\varphi \equiv \varphi_{beat} - \varphi_{VCO}$.

After the low-pass filter the signal is split into three parts in a power splitter. One of the outputs is used as a monitor, where $\nu_{beat} - \nu_{VCO}$ can be measured using e.g. a frequency counter (here we can check that the offset frequency is right). The two other outputs go through a short and a long cable, respectively, which puts different phase delays on the signals before they are mixed again in a second mixer.

The output from this mixer is proportional to

$$\begin{aligned} & \frac{V_sV_{VCO}}{2}\cos(\Delta\omega t + \Delta\varphi) \cdot \frac{V_sV_{VCO}}{2}\cos(\Delta\omega t + \Delta\varphi + \Phi) \\ = \; & \frac{V_s^2V_{VCO}^2}{8}\Big(\cos(\Phi) + \cos(2\Delta\omega t + 2\Delta\varphi + \Phi)\Big). \end{aligned}$$

The last term is removed with another low-pass filter ($2\Delta\nu \approx 20\,\mathrm{MHz}$). The first term (which is proportional to $\cos(\Phi)$) is the output from the offset lock and it goes into a servo-amplifier. The phase delay Φ depends on the time τ it takes for the signal to move though the cable

$$\Phi = (\omega_{beat} - \omega_{VCO})\tau = 2\pi(\nu_{beat} - \nu_{VCO})\tau. \tag{A.3}$$

τ depends on the velocity of the electronic signal through the cable—if we put this velocity equal to $\frac{2}{3}c$ we get 5 ns/m. In our case we chose two cables with a 5 m length difference, which gives a delay of around 25 ns (the number of locking points also depends on the cable length).

Different VCO frequencies can be chosen that gives the same beat frequency because there are many locking points for each VCO frequency. However, if we choose to lock to (one of) the first locking points (90 degrees away) the VCO frequency should be set to

$$2\pi(\nu_{beat} - \nu_{VCO})\tau = \frac{\pi}{2}$$

$$\Downarrow$$

$$\nu_{VCO} = \nu_{beat} - \frac{1}{4\tau}.$$

For $\tau = 25\,\text{ns}$ and a wanted beat frequency of 190 MHz, ν_{VCO} should be set to 180 MHz. The locking signal can be seen in figure A.2.

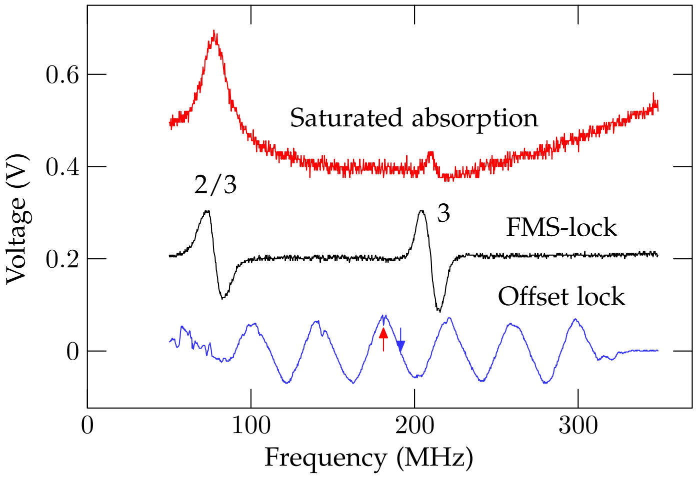

Figure A.2: The locking signal of the Cooler. The zero-point for the frequency axis is chosen to be the Master locking point ($F = 2$ to $F' = 1/3$ transition). The blue signal is the signal from the offset lock. The VCO frequency is chosen to be 180 MHz (marked by the red arrow—there is a small dip in the signal at the VCO frequency). This gives a locking point at 190 MHz which we lock to (marked by the blue arrow). Also shown is the saturated absorption signal and the FMS-locking signal.

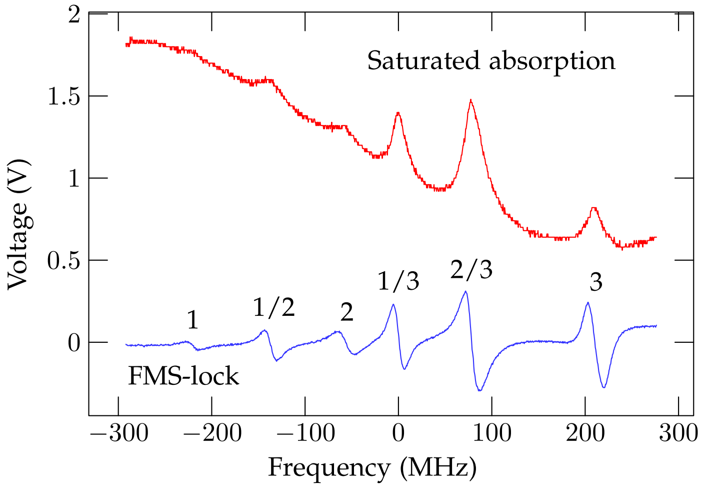

Figure A.3: The locking signal of the Master. The zero-point for the frequency axis is chosen to be the Master's locking point ($F = 2$ to $F' = 1/3$ transition).

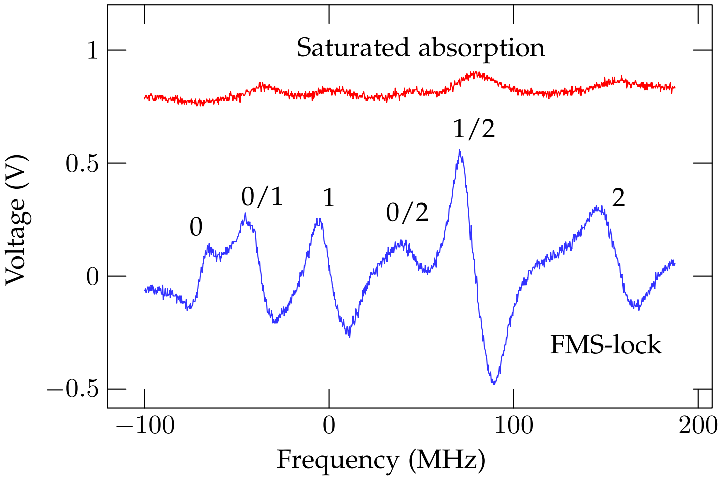

Figure A.4: The locking signal of the Repumper. The zero-point for the frequency axis is chosen to be the Repumper's locking point ($F = 1$ to $F' = 1$ transition).

B

Appendix B - Screenshots

In this appendix some ECS and CCS screenshots are shown.

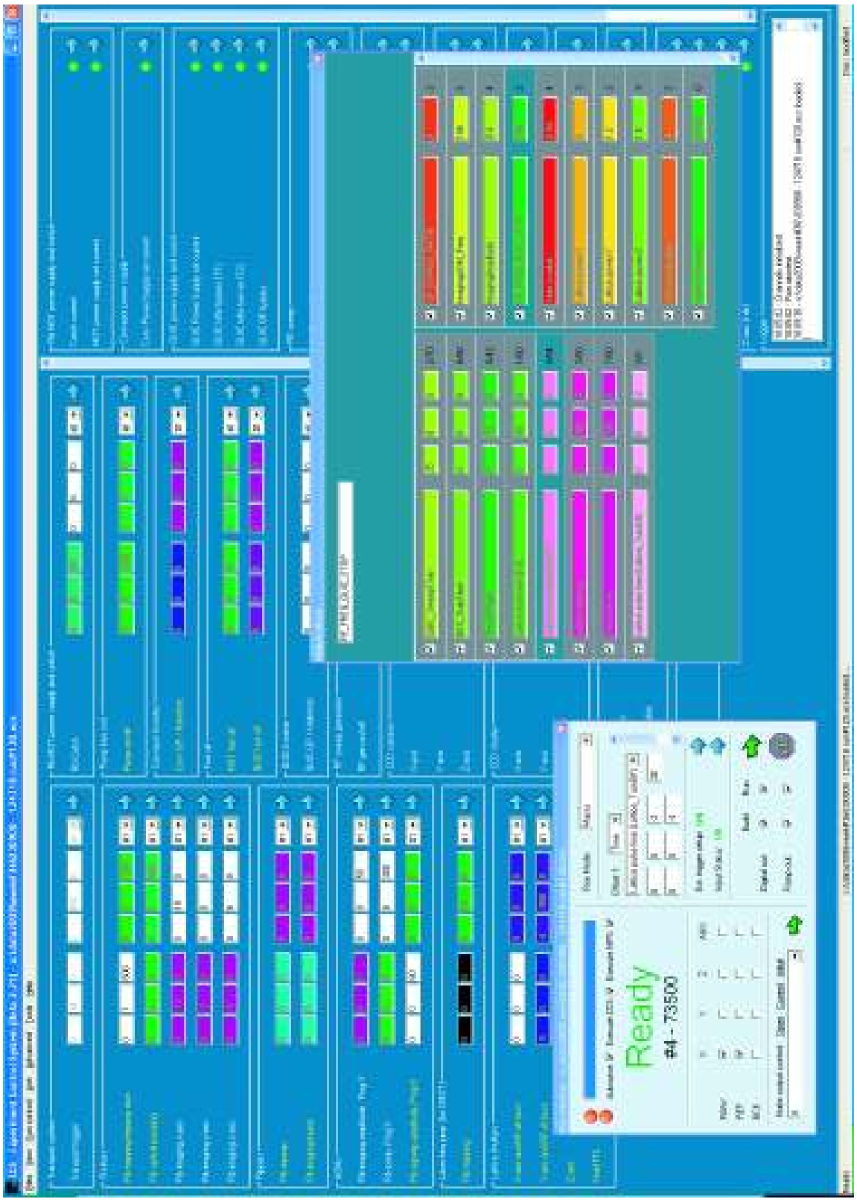

Figure B.1: The `Main Window` of ECS. This window is the primary interface between the user and ECS. The window is divided into two panels. The left panel gives access to the digital channels and the right panel gives access to the `Analog Window`, where the analog outputs are programmed. Also shown is the `Global Offset Window` and the `Run window`.

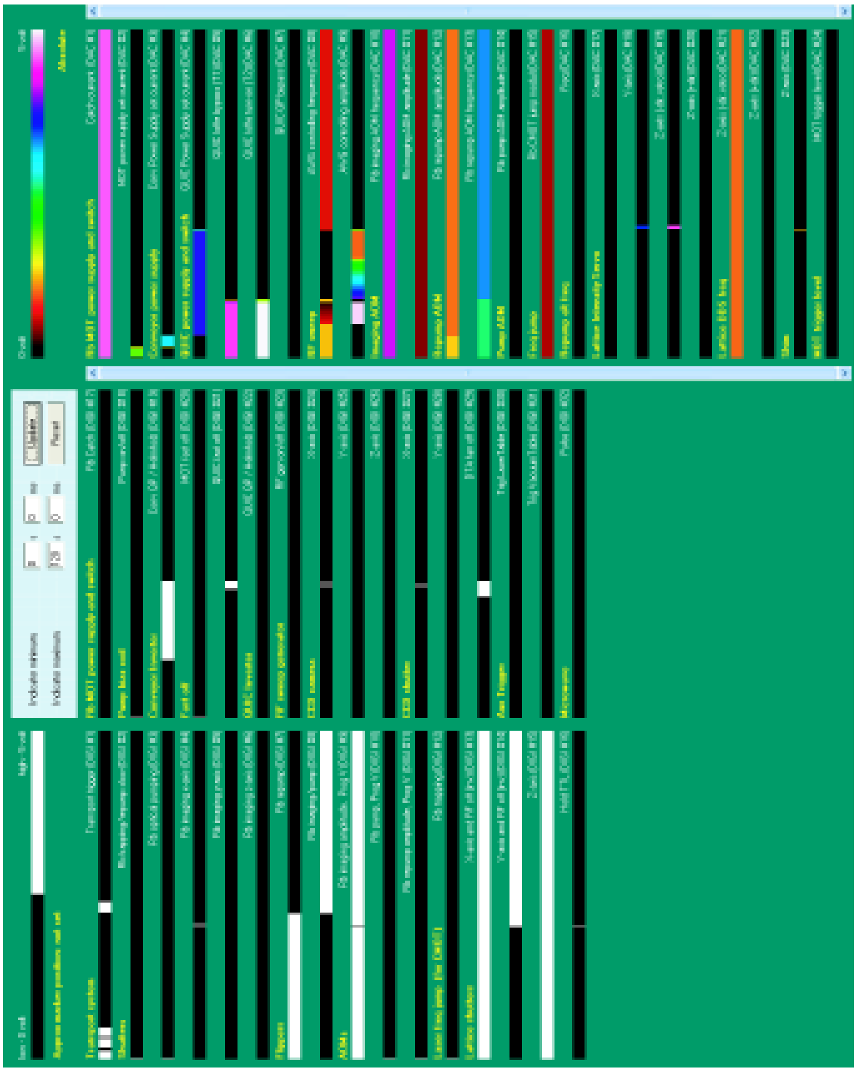

Figure B.2: The `Event Browser` of ECS. This browser is always shown on the second monitor. In this window all the different events on the different channels are shown. By left-clicking the mouse and dragging it left the user can zoom in on an event. Subsequently, a marker can be placed which makes it very easy to see if events on different channels are synchronized. For this reason much of the debugging concerning the output on the different channels are done in this window.

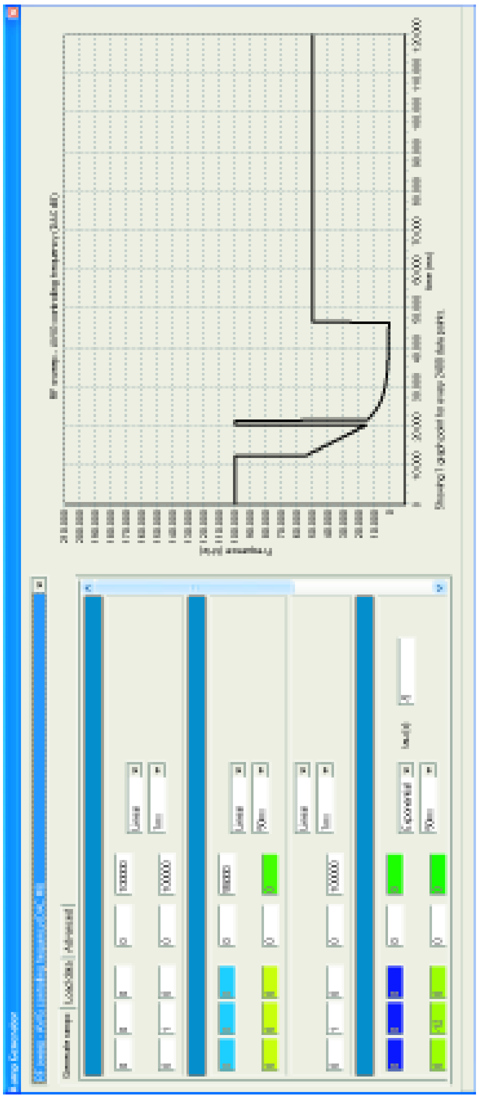

Figure B.3: The `Analog Window` of ECS. Here the time nodes defining the different ramps can be created. All of the times and values set in this window can be combined with global offsets. The colors indicate that an offset is used. Also here the user can zoom in on the graph to see the output. The graph shows exactly what is sent out—this means that by zooming enough the user can see every point in the ramp. The user can therefore also see how the lines get smoother when the time resolution of a time segment (ramp) is increased (e.g. from 1 ms to 50 µs). If something is changed in this window the graph and the `Event Window` is automatically updated. The different groups and elements are created by right-clicking and choosing the appropriate functions.

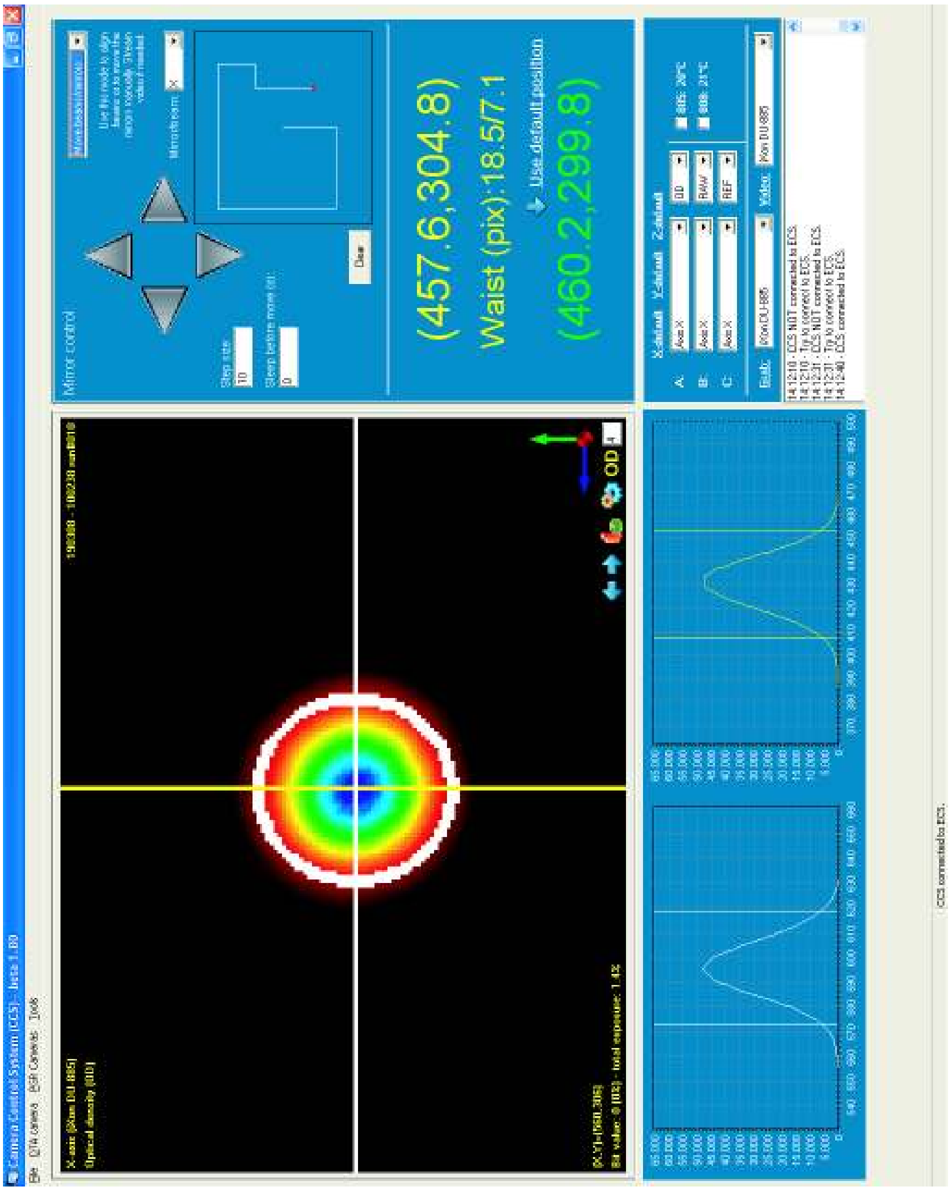

X-axis (iXon DU-885)
Optical density (OD)
190308 - 100238 run0010
(X,Y)=(560,306)
Bit value: 0 (0%) - total exposure: 1.4%

(457.6,304.8)
Waist (pix):18.5/7.1
(460.2,299.8)

Step size: 10
Sleep before move (s): 0
Mirror/beam: X

A: Axis X OD
B: Axis X RAW
C: Axis X REF
☐ 885: 20°C
☐ 888: 21°C
Grab: iXon DU-885
Video: iXon DU-885

```
14:12:10 - CCS NOT connected to ECS.
14:12:10 - Try to connect to ECS.
14:12:31 - CCS NOT connected to ECS.
14:12:31 - Try to connect to ECS.
14:12:40 - CCS connected to ECS.
```

CCS connected to ECS.

Figure B.4: The `Main Window` of CCS. Here we are streaming video images of one of the lattice beams. By clicking the gray arrows the beam will move across the image. A history of the movement of the mirror is shown, so it is easy to remember how to get back. CCS also automatically shows the beam profiles and calculates the beam waist.

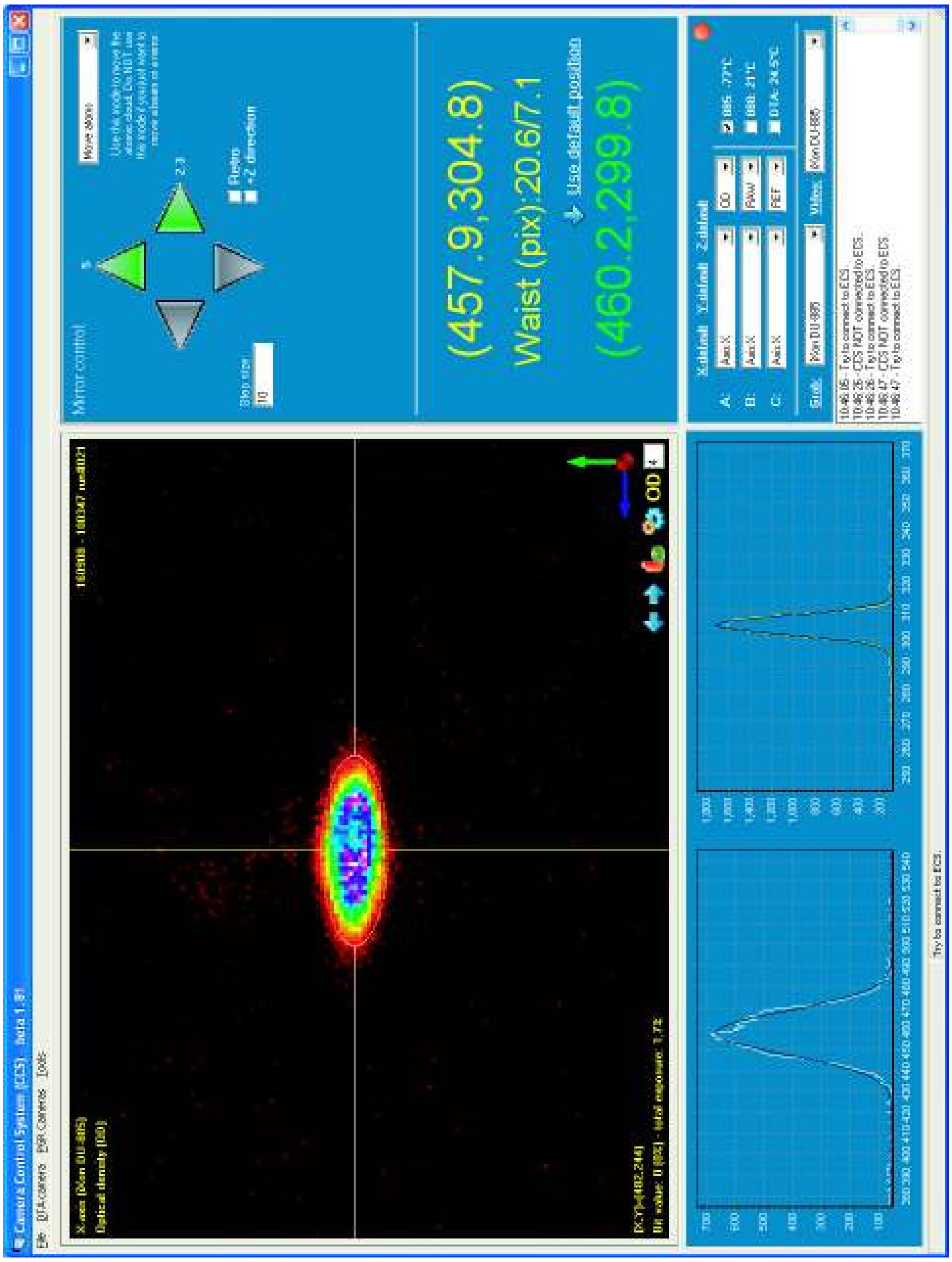

(457.9,304.8)

Waist (pix):20.6/7.1

(460.2,299.8)

```
10:46:05 - Try to connect to ECS.
10:46:26 - CCS NOT connected to ECS.
10:46:26 - Try to connect to ECS.
10:46:47 - CCS NOT connected to ECS.
10:46:47 - Try to connect to ECS.
```

Figure B.5: The `Main Window` of CCS. Here we are aligning the y- and z-lattice beams. In the image we see a BEC which is hold in a 2D dipole trap. CCS tells us how the atomic cloud should be moved in order to be placed exactly on top of the magnetic trap. In this case the lattice beams should be moved in such a way that the atomic cloud will move slightly up and right (indicated by the arrows). With this feature we can align the 3D optical lattice (including the retro-reflectors) very precisely in an hour. This alignment does *not* have to be done every day.

C

Appendix C - Optical traps

C.1 Introduction

In this appendix the optical dipole (and lattice) intensities, potentials and trapping frequencies are calculated for a 1D, 2D and 3D optical trap.

C.2 Definitions

We will use a Gaussian laser beam in the calculations [71]:

Laser wavelength: λ
Angular laser frequency: ω
Beam waist: w_0
Wave vector: $k = \frac{2\pi}{\lambda}$
Rayleigh range: $z_0 = \frac{\pi w_0^2}{\lambda}$
$\phi(z) = \tan^{-1}(z/z_0)$
Radius of curvature: $R(z) = z + \frac{z_0^2}{z}$
Spot size: $w(z) = w_0\sqrt{1 + z^2/z_0^2}$
$\mathcal{E}_0(x,y,z) = \frac{A \cdot e^{-i\phi(z)}}{\sqrt{1 + z^2/z_0^2}} e^{ik(x^2+y^2)/2R(z)} e^{-(x^2+y^2)/w^2(z)}$
$\mathcal{E}(x,y,z) = \mathcal{E}_0(x,y,z) \cdot e^{ikz}$
Electric field: $E(x,y,z,t) = \mathcal{E}(x,y,z) \cdot e^{-i\omega t}$

C.3 1D dipole trap

The intensity of a 1D dipole beam propagating in the z-direction is given by

$$\begin{aligned} I_{\text{1D-dip}}(x,y,z) &= \frac{c\cdot\epsilon_0}{2}|E(x,y,z,t)|^2 \\ &= \frac{c\cdot\epsilon_0}{2}|\mathcal{E}(x,y,z)|^2 \\ &= \frac{c\cdot\epsilon_0\cdot|A|^2}{2(1+z^2/z_0^2)}e^{-2(x^2+y^2)/w^2(z)}. \end{aligned}$$

The beam power is given by

$$P = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} I_{\text{1D-dip}}(x,y,z)dxdy = \frac{c\epsilon_0}{4}|A|^2(\pi w_0^2),$$

which gives

$$|A|^2 = \frac{4P}{c\epsilon_0\cdot\pi w_0^2}$$

so

$$I_{\text{1D-dip}}(x,y,z) = \frac{2P}{\pi w^2(z)}e^{-2(x^2+y^2)/w^2(z)}. \tag{C.1}$$

The potential is given by eq. (6.4)

$$\begin{aligned} V(x,y,z) &= \frac{\pi c^2\Gamma}{2\omega_0^3}\left(\frac{2+\mathcal{P}g_{\text{F}}m_{\text{F}}}{\Delta_{2,F}} + \frac{1-\mathcal{P}g_{\text{F}}m_{\text{F}}}{\Delta_{1,F}}\right)I(x,y,z) \\ &= BI(x,y,z), \end{aligned} \tag{C.2}$$

where

$$B \equiv \frac{\pi c^2\Gamma}{2\omega_0^3}\left(\frac{2+\mathcal{P}g_F m_F}{\Delta_{2,F}} + \frac{1-\mathcal{P}g_F m_F}{\Delta_{1,F}}\right).$$

So, the one-dimensional dipole potential becomes

$$V_{\text{1D-dip}}(x,y,z) = V_{dip}\frac{1}{(1+z^2/z_0^2)}e^{-2(x^2+y^2)/w^2(z)}, \tag{C.3}$$

where

$$V_{dip} = B\cdot\frac{2P}{\pi w_0^2}. \tag{C.4}$$

Trapping frequencies

If we Taylor expand eq. (C.3) we get ($z \ll z_0$, $x^2+y^2 \ll w_0^2$)

$$\begin{aligned}
V_{\text{1D-dip-approx}} &= V_{dip}\left[1 - z^2/z_0^2\right]\left[1 - \frac{2(x^2+y^2)}{w_0^2}\cdot\left(1 - z^2/z_0^2\right)\right] \\
&= V_{dip}\left[1 - 2\frac{x^2+y^2}{w_0^2} - \left(\frac{z}{z_0}\right)^2\right] \\
&= V_{dip}\left[1 - 2\frac{x^2+y^2}{w_0^2} - \left(\frac{z}{z_0}\right)^2\right].
\end{aligned}$$

From this follows the trapping frequencies ($r^2 = x^2 + y^2$)

$$\begin{aligned}
V_r &= V_{dip}\cdot 2\frac{r^2}{w_0^2} = \tfrac{1}{2}k_r r^2 \\
&\Downarrow \\
\omega_r &= \sqrt{\frac{k_r}{m}} = \sqrt{\frac{4V_{dip}}{mw_0^2}},
\end{aligned} \tag{C.5}$$

and likewise

$$\omega_z = \sqrt{\frac{2V_{dip}}{mz_0^2}}. \tag{C.6}$$

C.4 2D dipole trap

Here we will look at a 2D dipole trap (in the y- and z-direction). If the polarizations are orthogonal the intensity will simply be given by (assuming equal beam powers $P = P_1 = P_2$ and waists $w_0 = w_{01} = w_{02}$)

$$\begin{aligned} I_{\text{2D-dip}}(x,y,z) &= \frac{c\cdot\epsilon_0}{2}\Big(|E_1(x,y,z,t)|^2 + |E_2(x,y,z,t)|^2\Big) \\ &= \frac{2P}{\pi w^2(z)}e^{-2(x^2+y^2)/w^2(z)} + \frac{2P}{\pi w^2(y)}e^{-2(x^2+z^2)/w^2(y)} \end{aligned}$$

The 2D dipole potential will then be given by (since the beam waists and wavelengths are the same the Raleigh lengths are also—$z_0 = y_0$)

$$V_{\text{2D-dip}} = V_{dip}\left(\frac{1}{(1+z^2/z_0^2)}e^{-2(x^2+y^2)/w^2(z)} + \frac{1}{(1+y^2/z_0^2)}e^{-2(x^2+z^2)/w^2(y)}\right). \tag{C.7}$$

This can be approximated by

$$\begin{aligned} V_{\text{2D-dip-approx}} &= B\cdot\frac{2P}{\pi}\left(\frac{1}{w_0^2(1+z^2/z_0^2)}e^{-2(x^2+y^2)/w^2(z)}\right. \\ &\quad\left. + \frac{1}{w_0^2(1+y^2/z_0^2)}e^{-2(x^2+z^2)/w^2(y)}\right) \\ &= B\cdot\frac{2P}{\pi w_0^2}\left\{\left[1 - 2\frac{x^2+y^2}{w_0^2} - \Big(\frac{z}{z_0}\Big)^2\right]\right. \\ &\quad\left. + \left[1 - 2\frac{x^2+z^2}{w_0^2} - \Big(\frac{y}{z_0}\Big)^2\right]\right\} \\ &= V_{dip}\left[2 - 4\frac{x^2}{w_0^2} - \Big(\frac{2}{w_0^2} + \frac{1}{z_0^2}\Big)y^2 - \Big(\frac{2}{w_0^2} + \frac{1}{z_0^2}\Big)z^2\right], \end{aligned}$$

which gives the trapping frequencies

$$\begin{aligned} \omega_z = \omega_y &= \sqrt{\frac{2V_{dip}}{m}\cdot\Big(\frac{2}{w_0^2} + \frac{1}{z_0^2}\Big)} \\ &\approx \sqrt{\frac{4V_{dip}}{mw_0^2}}, \end{aligned} \tag{C.8}$$

and

$$\omega_x = \sqrt{\frac{8V_{dip}}{mw_0^2}}. \tag{C.9}$$

C.5 3D dipole trap

Like the 2D case but now in 3D

$$\begin{aligned}I_{\text{3D-dip}}(x,y,z) = {} & \frac{2P}{\pi w^2(z)} e^{-2(x^2+y^2)/w^2(z)} + \frac{2P}{\pi w^2(y)} e^{-2(x^2+z^2)/w^2(y)} \\ & + \frac{2P}{\pi w^2(x)} e^{-2(y^2+z^2)/w^2(x)}. \end{aligned} \tag{C.10}$$

Giving a potential (assuming same beam waist w_0 and Raleigh length z_0 on all axes)

$$\begin{aligned}V_{\text{3D-dip}} = V_{dip} \Bigg(& \frac{1}{(1+z^2/z_0^2)} e^{-2(x^2+y^2)/w^2(z)} \\ & + \frac{1}{(1+y^2/z_0^2)} e^{-2(x^2+z^2)/w^2(y)} + \frac{1}{(1+x^2/z_0^2)} e^{-2(y^2+z^2)/w^2(x)} \Bigg). \end{aligned} \tag{C.11}$$

This can be approximated by

$$\begin{aligned}V_{\text{3D-dip-approx}} = V_{dip} \Bigg[& \bigg[1 - 2\frac{x^2+y^2}{w_0^2} - \Big(\frac{z}{z_0}\Big)^2 \bigg] \\ & + \bigg[1 - 2\frac{x^2+z^2}{w_0^2} - \Big(\frac{y}{z_0}\Big)^2 \bigg] + \bigg[1 - 2\frac{y^2+z^2}{w_0^2} - \Big(\frac{x}{z_0}\Big)^2 \bigg] \Bigg] \\ = V_{dip} \bigg[& 3 - \Big(\frac{4}{w_0^2} + \frac{1}{z_0^2}\Big)x^2 - \Big(\frac{4}{w_0^2} + \frac{1}{z_0^2}\Big)y^2 - \Big(\frac{4}{w_0^2} + \frac{1}{z_0^2}\Big)z^2 \bigg], \end{aligned}$$

which gives the trapping frequencies

$$\begin{aligned}\omega_x = \omega_y = \omega_z &= \sqrt{\frac{2V_{dip}}{m} \cdot \Big(\frac{4}{w_0^2} + \frac{1}{z_0^2}\Big)} \\ &\approx \sqrt{\frac{8V_{dip}}{m w_0^2}}. \end{aligned} \tag{C.12}$$

C.6 1D optical lattice

Now, two beams propagating in z-plus and z-minus direction with the same linear polarization. The z-plus direction (with A_1, ω_1 and $k_1 = 2\pi/\lambda_1$)

$$\begin{aligned}
\mathcal{E}_{01}(x,y,z) &= \frac{A_1 \cdot e^{-i\phi_1(z)}}{\sqrt{1+z^2/z_{01}^2}} e^{ik_1(x^2+y^2)/2R_1(z)} e^{-(x^2+y^2)/w_1^2(z)} \\
\mathcal{E}_1(x,y,z) &= \mathcal{E}_{01}(x,y,z) \cdot e^{ik_1 z} \\
E(x,y,z,t) &= \mathcal{E}_1(x,y,z) \cdot e^{-i\omega_1 t}.
\end{aligned}$$

Likewise the z-minus direction (with A_2, ω_2 and $k_2 = 2\pi/\lambda_2$)

$$\begin{aligned}
\mathcal{E}_{02}(x,y,z) &= \frac{A_2 \cdot e^{-i\phi_2(z)}}{\sqrt{1+z^2/z_{02}^2}} e^{ik_2(x^2+y^2)/2R_2(z)} e^{-(x^2+y^2)/w_2^2(z)} \\
\mathcal{E}_2(x,y,z) &= \mathcal{E}_{02}(x,y,z) \cdot e^{-ik_2 z} \\
E(x,y,z,t) &= \mathcal{E}_2(x,y,z) \cdot e^{-i\omega_2 t}.
\end{aligned}$$

If $P = P_1 = P_2$ and $w_0 = w_{01} = w_{02}$ then $|A|^2 = |A_1|^2 = |A_2|^2$, and if also $k = k_1 = k_2$ then $z_0 = z_{01} = z_{02}$, $\phi(z) = \phi_1(z) = \phi_2(z)$, $R(z) = R_1(z) = R_2(z)$ and $w(z) = w_1(z) = w_2(z)$, so

$$\begin{aligned}
I_{\text{1D-lat}}(x,y,z) &= \frac{c \cdot \epsilon_0}{2} |E_1(x,y,z,t) + E_2(x,y,z,t)|^2 \\
&= \frac{c \cdot \epsilon_0}{2} \left(|\mathcal{E}_1|^2 + |\mathcal{E}_2|^2 + \mathcal{E}_1^* \mathcal{E}_2 + \mathcal{E}_2^* \mathcal{E}_1 \right) \\
&= \frac{c \cdot \epsilon_0}{2} \left[2 \frac{|A|^2}{1+z^2/z_0^2} e^{-2(x^2+y^2)/w^2(z)} \right. \\
&\left. + \frac{|A|^2}{1+z^2/z_0^2} \cdot e^{-2(x^2+y^2)/w^2(z)} \left(e^{2ikz} + e^{-2ikz} \right) \right] \\
&= \frac{c \cdot \epsilon_0}{2} \frac{4|A|^2}{1+z^2/z_0^2} e^{-2r^2/w^2(z)} \cos^2 kz \\
&= \frac{8P}{\pi w^2(z)} e^{-2r^2/w^2(z)} \cos^2 kz, \qquad \text{(C.13)}
\end{aligned}$$

where $r^2 = x^2 + y^2$. This result is 4 times the dipole intensity and with a $\cos^2(kz)$ modulation. The potential can again be calculated

$$\begin{aligned}
V_{\text{1D-lat}} &= B \cdot \frac{8P}{\pi} \frac{1}{w_0^2(1+z^2/z_0^2)} e^{-2r^2/w^2(z)} \cos^2 kz \\
&= V_0 \frac{1}{(1+z^2/z_0^2)} e^{-2r^2/w^2(z)} \cos^2 kz. \qquad \text{(C.14)}
\end{aligned}$$

The connection between V_0 and V_{dip} is

$$V_0 = \frac{B \cdot 8P}{\pi w_0^2} = 4V_{dip}. \tag{C.15}$$

Trapping frequencies

We can Taylor expand eq. (C.14)

$$\begin{aligned} V_{\text{1D-lat-approx}} &\approx V_0 \Big[1 - 2\Big(\frac{r}{w_0}\Big)^2 - \Big(\frac{z}{z_0}\Big)^2\Big]\Big[1 - (kz)^2\Big] \\ &\approx V_0 \Big[1 - (kz)^2 - 2\Big(\frac{r}{w_0}\Big)^2 - \Big(\frac{z}{z_0}\Big)^2\Big] \\ &\approx V_0 \Big[1 - 2\Big(\frac{r}{w_0}\Big)^2 - \Big(\frac{1}{z_0^2} + k^2\Big)z^2\Big] \end{aligned}$$

and thereby calculate the trapping frequencies. The axial trapping frequency is significantly changed compared to the dipole case because it is now the wavelength of the lattice that sets the frequency

$$\omega_z = \sqrt{\frac{2V_0}{m} \cdot \Big(\frac{1}{z_0^2} + k^2\Big)} \approx \sqrt{\frac{2V_0}{m} \cdot k^2} \tag{C.16}$$

The lattice depth V_0 can be found if the axial trapping frequency is measured

$$V_0 = \frac{\omega_z^2 m}{2k^2}. \tag{C.17}$$

Notice that V_0 depends slightly on the Raleigh length z_0—however as long as $w_0 \gg \lambda$ it is a *very* small effect. For beam waists in the range from 30 µm to 120 µm, z_0 is a few millimeters to a few centimeters, respectively.

The radial trapping frequency is the same as in the dipole case except that it is now V_0 that enters instead of V_{dip}

$$\omega_r = \sqrt{\frac{2V_0}{m} \cdot \Big(\frac{2}{w_0^2}\Big)} = \sqrt{\frac{4V_0}{m w_0^2}}. \tag{C.18}$$

If none of the parameters leading to eq. (C.13) are equal then the intensity I is given by (assuming that A_1 and A_2 are real)

$$
\begin{aligned}
I(x,y,z,t) &= \frac{c\cdot\epsilon_0}{2}|E_1(x,y,z,t)+E_2(x,y,z,t)|^2\\
&= \frac{c\cdot\epsilon_0}{2}\Big(|\mathcal{E}_1|^2+|\mathcal{E}_2|^2+\mathcal{E}_1^*\mathcal{E}_2e^{i(\omega_1 t-\omega_2 t)}+\mathcal{E}_2^*\mathcal{E}_1e^{i(\omega_2 t-\omega_1 t)}\Big)\\
&= \frac{c\cdot\epsilon_0}{2}\Bigg(\frac{A_1^2}{1+z^2/z_{01}^2}e^{-2(x^2+y^2)/w_1^2(z)}\\
&+\frac{A_2^2}{1+z^2/z_{02}^2}e^{-2(x^2+y^2)/w_2^2(z)}\\
&-\frac{2\cdot A_1A_2}{\sqrt{1+z^2/z_{01}^2}\sqrt{1+z^2/z_{02}^2}}e^{-(x^2+y^2)/w_1^2(z)}e^{-(x^2+y^2)/w_2^2(z)}\\
&+\frac{4\cdot A_1A_2}{\sqrt{1+z^2/z_{01}^2}\sqrt{1+z^2/z_{02}^2}}e^{-(x^2+y^2)/w_1^2(z)}e^{-(x^2+y^2)/w_2^2(z)}\\
&\cdot\cos^2\Big[\tfrac{1}{2}\Big(\Big[\phi_2(z)-\phi_1(z)\Big]+\Big[\frac{(x^2+y^2)}{2}\Big(\frac{k_1}{R_1(z)}-\frac{k_2}{R_2(z)}\Big)\Big]\\
&+\Big[(k_1+k_2)z)\Big]+\Big[(\omega_2-\omega_1)t\Big]\Big)\Big]\Bigg).
\end{aligned}
\tag{C.19}
$$

This is the expression we use in our simulations because then we can vary the parameters of the different beams independently. If $P_1 = P_2$ and $w_{01} = w_{02}$ then eq. (C.19) reduces to

$$
\begin{aligned}
I(x,y,z,t) = \frac{8P}{\pi w^2(z)}&e^{-2(x^2+y^2)/w^2(z)}\\
&\cdot\cos^2\Big[\tfrac{1}{2}\Big(\Big[\frac{(x^2+y^2)}{2}\Big(\frac{k_1}{R(z)}-\frac{k_2}{R(z)}\Big)\Big]\\
&+\Big[(k_1+k_2)z)\Big]+\Big[(\omega_2-\omega_1)t\Big]\Big)\Big].
\end{aligned}
\tag{C.20}
$$

Notice that the time dependent term $(\omega_2-\omega_1)t$ can describe a moving optical lattice.

From C.19 we can also derive how to change from a dipole trap to an optical lattice without changing the "overall" trapping frequency (see chapter 6 for more information). In this example we will allow for different powers P_1 and P_2 (but the rest of the beam parameters are the same for the two beams). We can now calculate the axial trapping frequency

$$
\omega_z^2 = \frac{4B}{\pi w_0^2 z_0^2 m}\Big(P_1+P_2-2\sqrt{P_1P_2}+4\sqrt{P_1P_2}k^2z_0^2+4\sqrt{P_1P_2}\Big). \tag{C.21}
$$

The k term is the term we normally keep because it gives the trapping frequency of the optical lattice (many kHz). However in this case we are interested in the lower "overall" trapping frequency due to the fact that we use Gaussian beams which a certain Raleigh length. This means that here we throw away the k term. If we define a constant

$$c_1 \equiv \frac{\omega_z^2 \pi w_0^2 z_0^2 m}{4B}, \tag{C.22}$$

we get that

$$P_1 = P_2 + c_1 - 2\sqrt{P_2 c_1} \tag{C.23}$$
$$P_2 = P_1 + c_1 - 2\sqrt{P_1 c_1}. \tag{C.24}$$

This means that we can choose a specific ramp for e.g. P_1 and then calculate how P_2 should be ramped to give the same "overall" trapping frequency. We can check these formulas in the limits. If we set $P_2 = 0$ (dipole trap) we get C.6. In the other limit we put $P_1 = P_2$ (optical lattice). Here we get the z_0 term of C.16.

C.7 2D optical lattice

Now the same for a 2D optical lattice in the y- and z-direction (here we will only write up the simple formulas where all the parameters of the counter propagating beams are identical)

$$\begin{aligned} I_{\text{2D-lat}}(x,y,z) &= \frac{8P}{\pi w^2(z)} e^{-2(x^2+y^2)/w^2(z)} \cos^2 kz \\ &+ \frac{8P}{\pi w^2(y)} e^{-2(x^2+z^2)/w^2(y)} \cos^2 ky, \end{aligned} \tag{C.25}$$

and

$$\begin{aligned} V_{\text{2D-lat-approx}} = V_0 \Big[2 &- \Big(k^2 + \frac{1}{z_0^2} + \frac{2}{w_0^2} \Big) z^2 \\ &- \Big(k^2 + \frac{1}{z_0^2} + \frac{2}{w_0^2} \Big) y^2 - \frac{4}{w_0^2} x^2 \Big]. \end{aligned}$$

Again we can find the trapping frequencies

$$\omega_y = \omega_z \approx \sqrt{\frac{2V_0 k^2}{m}} \tag{C.26}$$

and

$$\omega_x = \sqrt{\frac{8V_0}{m w_0^2}}. \tag{C.27}$$

C.8 3D optical lattice

Like the 2D case but now in 3D

$$\begin{aligned} I_{\text{3D-lat}}(x,y,z) &= \frac{8P}{\pi w^2(z)} e^{-2(x^2+y^2)/w^2(z)} \cos^2 kz \\ &+ \frac{8P}{\pi w^2(y)} e^{-2(x^2+z^2)/w^2(y)} \cos^2 ky \\ &+ \frac{8P}{\pi w^2(x)} e^{-2(y^2+z^2)/w^2(x)} \cos^2 kx, \end{aligned} \tag{C.28}$$

and

$$V_{\text{3D-lat-approx}} = V_0 \Big[3 - \Big(k^2 + \frac{1}{z_0^2} + \frac{4}{w_0^2} \Big) (x^2 + y^2 + z^2) \Big],$$

so

$$\omega_x = \omega_y = \omega_z = \sqrt{\frac{2V_0 k^2}{m}}. \tag{C.29}$$

D

Appendix D - Drawings of the optical lattice setup

Detailed drawings of the 3D optical lattice setup can be seen in this appendix. All the optics for the lattices are placed on three different levels around the science chamber. First level (the vacuum table) can be seen in figure D.1 and D.2. Here the z-lattice beam is divided into two beams, and a part of the y-lattice/imaging setup is also placed here. Level two is placed 75 mm (our normal beam height) below the center of the science chamber and here the optics for the x- and z-lattice/imaging axes are placed. The third level holds the other part of the y-lattice/imaging setup (the part after the science chamber). Dichromatic mirrors and wave-plates are used everywhere in the setup where we have both 780 nm and 915 nm light.

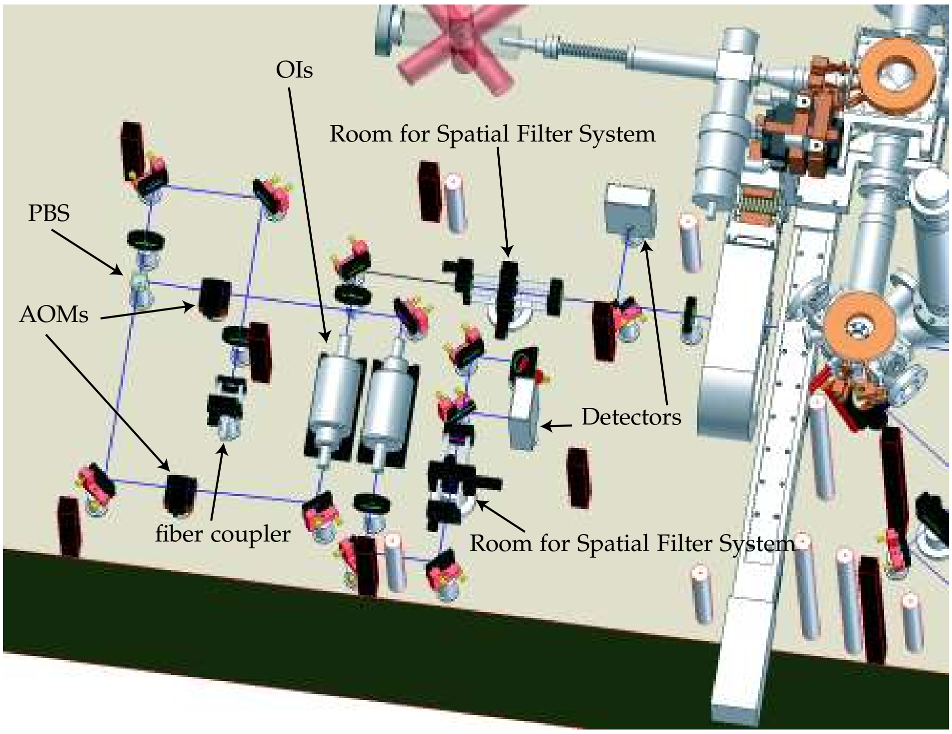

Figure D.1: A drawing of the part of the z-lattice setup where the beam is divided into two (a z-plus and a z-minus beam). When the light comes out of the fiber coupler it is split into two beams on a PBS. Before the PBS we have a half-wave plate, where we can adjust the ratio between the two beams. Each beam goes through an AOM and then through an OI. The OIs points in opposite directions to minimize stray magnetic fields. Detectors monitor the beam power and it is kept constant using the AOMs and some control electronics. The z-minus beam is sent up through a breadboard on this side of the science chamber. The z-plus beam is sent below the science chamber and then reflected up through another breadboard on the other side of the science chamber. Even though the beam passes below the transport (MPS) system the beam is never blocked no matter where the conveyor coils are placed. Room for spacial filters is reserved if the beam profiles from the AOMs need to be cleaned up. So far this has not been necessary.

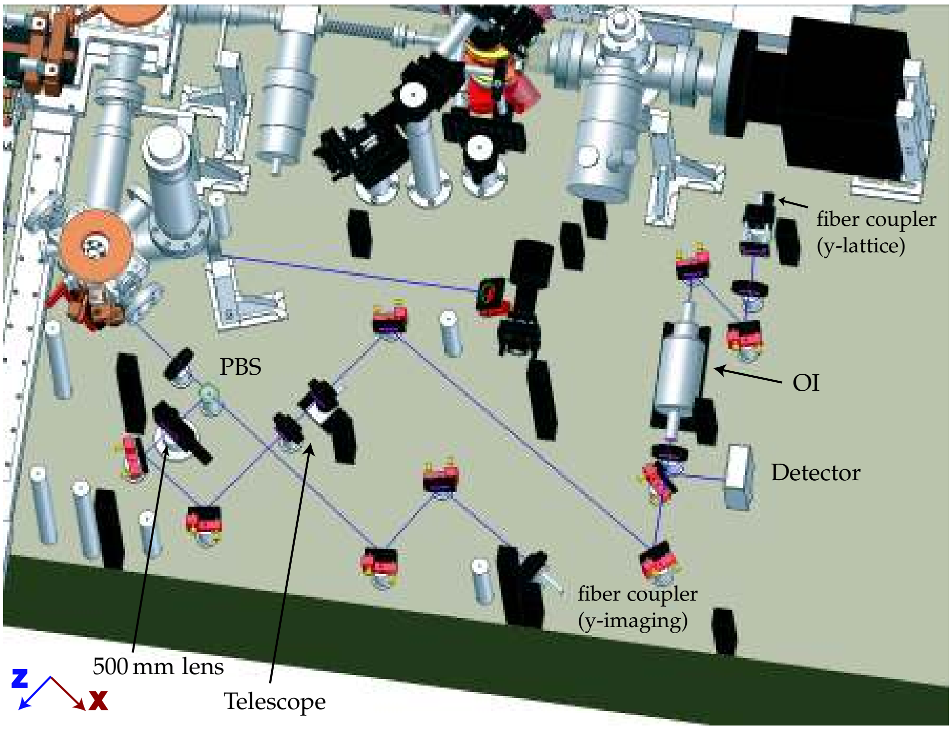

Figure D.2: A drawing of the y-lattice and y-imaging setup located before the science chamber. The lattice light passes through an OI before it goes through a telescope and hits a 500 mm lens, which produces a beam waist inside the science chamber. After the 500 mm lens the s-polarized (vertical polarized) lattice beam is reflected on a PBS, where it is also combined with the p-polarized imaging beam. A dichromatic half-wave plate is placed after the PBS and it changes the polarization of both beams. This means that the polarization of the y-lattice beam is parallel to the z-axis when it goes up through the science chamber. The z-plus lattice beam (from figure D.1) can also be seen here—it goes below the vacuum chamber and hits a mirror that reflects it upwards.

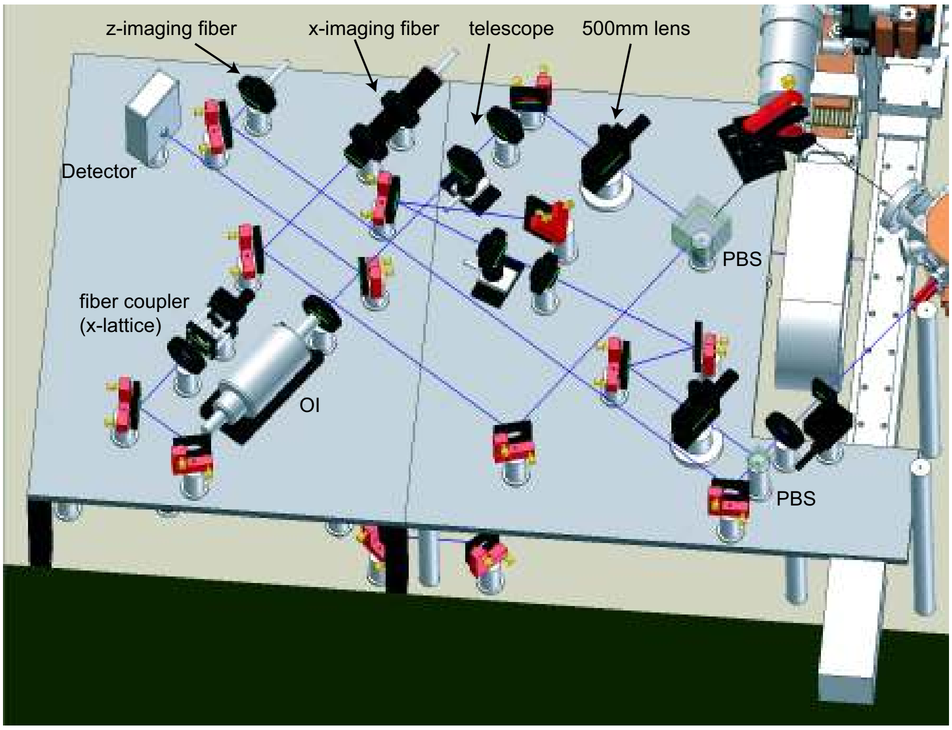

Figure D.3: The x- and z-lattice/imaging setup before the science chamber (on level 2). Again the beams are combined on a PBS before they are sent through the science chamber. Both lattice beams are s-polarized when they hit the PBS, which means that the x-lattice beam is polarized along the y-direction. The z-lattice beam goes through a dichromatic half-wave plate after the PBS, which makes the lattice light polarized along the x-axis.

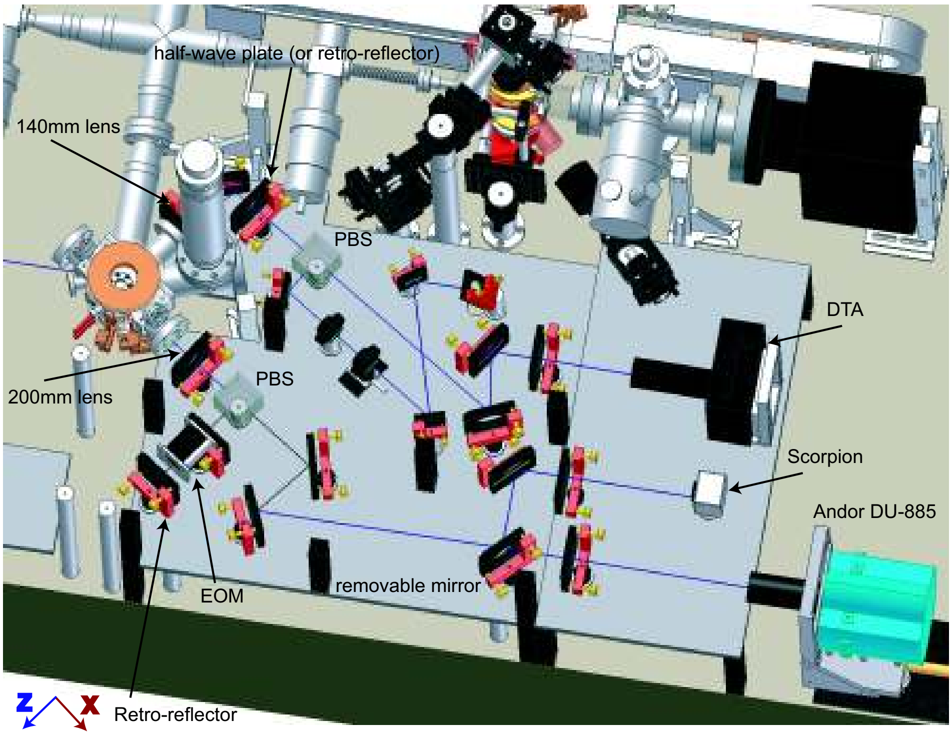

Figure D.4: This figure shows the x- and z-axis lattice/imaging beams after the science chamber (on level 2). The x-axis lattice/imaging beam comes out through the science chamber and hits a 200 mm lens and a PBS. The x-lattice beam is reflected and goes through an EOM before it hits the retro-reflector. The x-imaging beam goes through the PBS. Depending on whether or not the "removable mirror" is in the setup, the imaging beam will either hit the Scorpion or the Andor camera. The Andor camera is focused on the center of the QUIC trap, whereas the Scorpion camera is looking at the QP center (and has a smaller magnification). The z-minus lattice/imaging beam comes out of the science chamber and goes through a 140 mm lens, through a half-wave plate and hits a PBS, where the lattice beam is reflected and the imaging beam is transmitted. The z-minus beam goes back through a telescope and down through the breadboard. The z-plus lattice beam goes in the opposite direction. If we want to use the retro-reflector (instead of having both a z-plus and a z-minus beam), we can place it 140 mm from the lens (before the half-wave plate). The retro-reflector reflects the 915 nm lattice beam but transmits the 780 nm imaging beam. In front of all the cameras we have an achromatic lens, and on the Andor and DTA cameras we have an additional macro lens.

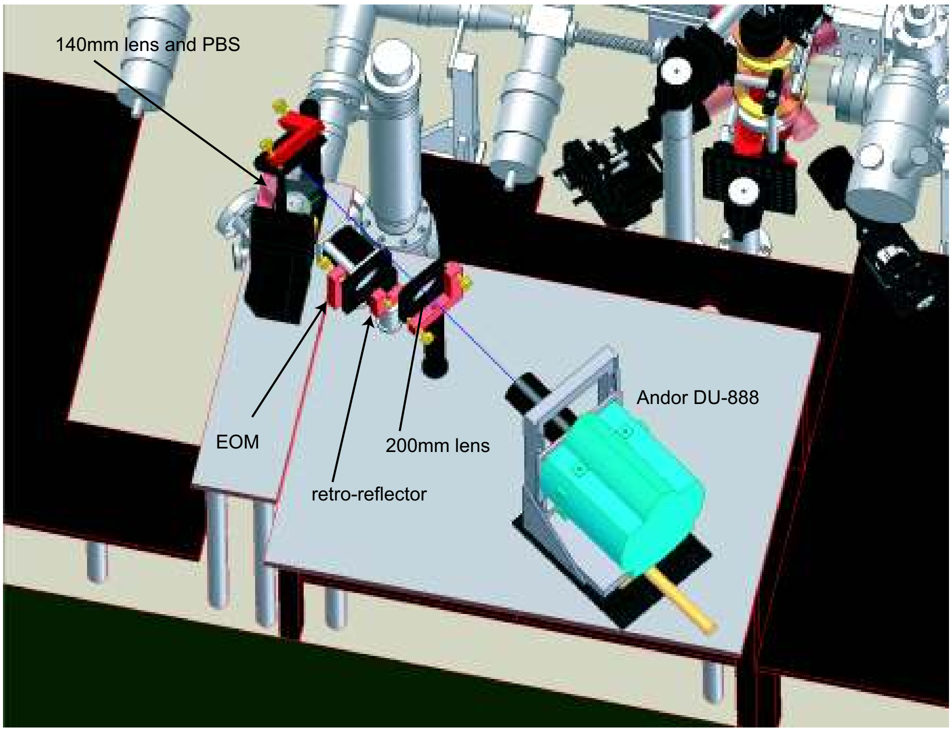

Figure D.5: Here the y-lattice/imaging beam comes up through the science chamber and hits a 140 mm lens. This lens is automatically moved when we want to grab images after a certain TOF on the y-axis (slightly unexpected it is actually possible to move this lens and still have an optical lattice along the y-axis). The light then hits a PBS where the imaging light is transmitted and the lattice light is reflected. The lattice beam goes through an EOM and hits the retro-reflector located approximately one focal length away from the 140 mm lens. The imaging beam hits a mirror and then goes through a 200 mm lens before it hits the macro lens mounted on the camera. Notice that the beam height of the lattice beam (80 mm) is different from the beam height of the imaging beam (180 mm). The microwave antenna (not visible on this figure) is mounted below this breadboard. It sends microwave radiation through the science chamber in the negative x-direction.

E

Appendix E - Aligning the 3D optical lattice

In this appendix it is described how we align the 3D optical lattice beams. The goal is to align three ingoing and three retro-reflected beams with a precision of a few micrometers exactly on top of the magnetic trap minimum.

For the alignment we use motorized mirror mounts with a resolution of 0.7 µrad as described in chapter 4. We chose the motorized mirror mounts because we liked the idea of the process being absolutely reproducible. Also (as we will see in the following), we almost entirely base the alignment procedures on images of the atomic clouds. For this reason it would be very difficult to do manually—at least it would require mirror mounts with a scale to be able to see how much the mirrors were moved. We align the ingoing beams with a precision of approximately 5 steps ($5 \times 0.7\,\mu$rad), and the retro-reflected beam with a precision of about 20 steps. Since the mirrors in front of the chamber is placed around 1 m away (and the retro-reflectors around 30 cm away) from the center of the science chamber, this corresponds to about 4 µm at the position of the atoms.

CCS knows the "reference" position of the center of the magnetic trap (the standard deviation of this position is typically 0.5 to 2 pixels corresponding to a few micrometers). For this reason CCS can help us moving the right mirrors as described in chapter 4. The position of the center of the magnetic trap depends on the value of I_{z-shim} and I_{QUIC} (as described in chapter 5) so we cannot e.g. change the trapping frequencies of the magnetic trap without having to realign the lattice.

Rough alignment

1. We start by streaming video images of the ingoing lattice beam and move it so that it is placed exactly at the reference position (since the cameras are extremely sensitive we flip in an $OD = 2$ to $OD = 3$ filter before the optical fiber to reduce the power of the beams). We do this initial alignment on all three axes, however since the wavelength of the lattice beams is not the same as the wavelength of the imaging beams, it is not a very precise alignment (because of chromatic aberration).

2. We then turn on the dipole trap (retro-reflectors blocked—e.g. with the EOMs) and look at the position of the atomic cloud. Even though the magnetic field is still on the atoms will typically be shifted because the dipole beams are not well-aligned yet, and this will pull the atoms away from the reference position. With e.g. the x-imaging we can align the y- and z-dipole beams in one direction. We then do the same with y- and z-imaging. Often we need to move the mirrors up to 50 steps in this procedure.

3. The retro-reflectors are then roughly aligned. This is done by watching the retro-reflected beam on an iris as far away from the retro-reflector as possible. This has to be done precisely since a 1 mm misalignment on the iris can mean a few hundreds µm at the position of the atomic cloud. To check that this step is done accurately, we either see if we can hold the atoms in a strong 1D optical lattice beam (with the magnetic trap off), or we look for interference peaks (with a week lattice and TOF). If we do not see atoms or interference peaks we align the retro-reflected beam again (normally we have to try one or two times before we have an optical lattice). The retro-reflector will often be many hundred steps away from the right position, when we use this method, but it does not matter since we also go through a fine alignment procedure.

We only use this procedure if we have changed something in the setup—otherwise it is not necessary.

Fine alignment

In this procedure we go through nine steps. In step 1–3 we align the dipole beams (6 degrees of freedom), and in step 4–9 we align the retro-reflectors (also 6 degrees of freedom):

1. X-axis imaging: We turn on the y- and z-dipole beams (retro-reflectors blocked) and turn off the magnetic trap and hold the atoms for 1 s using only the dipole cross. We then place the atomic cloud at the reference position by moving the mirrors.

2. Y-axis imaging: Like step 1, but here we align the x- and z-dipole beams—also 1 s hold time.

3. Z-axis imaging: Here we align the x- and y-dipole beams. However, since the z-beam is off, we cannot hold the atoms, so this step is done with the magnetic trap on.

4. Z-axis imaging: We trap the atoms in a 1D optical lattice along the y-axis and move the retro-reflector until the atomic cloud is located at the reference position (1 s hold time).

5. X-axis imaging: Like step 4, but here we align the retro-reflector in the other direction.

6. X-axis imaging: Like step 4, but for the z-lattice.

7. Y-axis imaging: Like step 5, but for the z-lattice.

8. Y-axis imaging: Here we use the z-dipole beam (retro-reflector blocked) together with the x-lattice beam. We align the retro-reflector for the x-lattice beam—1 s hold time.

9. Z-axis imaging: In this step we align the retro-reflector for the x-lattice beam in the other direction, but in this case with the magnetic trap on (since we cannot use the z-dipole beam, when we are using the z-axis camera to grab in-trap images).

The magnetic trap is off except in step 3 and 9. This is done in order to make it easier to see the effects of moving the beams.

We have tried different approaches for aligning the optical lattices, but we ended up with the one described above. A common way to align the retro-reflectors are to couple the beams back through the optical fibers. We can (in principle) easily see the retro-reflected beams that go back through the fibers because they are coupled out on a PBS in the OIs after the TAs on the laser table (see e.g. figure 7.1). However, we have a 40 dB OI on the vacuum table and this means that the intensity of the retro-reflected beam is really small when we measure it on the laser table (compared to the reflections we get from e.g. the lenses on the laser table). This makes it hard to get a good signal, and even though we have a signal we cannot align the retro-reflectors nearly as good using this method as when we are looking at the atoms.

Initially we thought that we had to align the optical lattice beams every day (e.g. due to temperature variations). For this reason we used a lot of time on refining our methods for doing the alignment good and fast. Currently, we can align the 3D optical lattice beams in less than an hour. However after having used the optical lattices for a few months, we can conclude that the setup is very stable. Without having been in the laboratory for a months, we were e.g. able to see the superfluid to Mott insulator transition without doing any alignment. With this in mind we have concluded, that it is fine just to do the alignment procedure once a week, and then instead just measure the depth of the optical lattices from day to day.

Bibliography

[1] W. Ketterle, D. S. Durfee and D. M. Stamper-Kurn, *Making, probing and understanding Bose-Einstein condensates*, arxiv:cond-mat/9904034 (1999).
URL link

[2] V. Bagnato, D. E. Pritchard and D. Kleppner, *Bose-Einstein Condensation In An External Potential*, Physical Review A **35** (1987), 10, pp. 4354, doi:10.1103/PhysRevA.35.4354.
URL link

[3] F. Mandl, *Statistical Physics* (Wiley; 2 edition, 1988).

[4] C. Pethick and H. Smith, *Bose-Einstein Condensation in Dilute Gases* (University Press, Cambridge, 2002).

[5] T. Weber, J. Herbig, M. Mark, H. C. Nagerl and R. Grimm, *Bose-Einstein condensation of cesium*, Science **299** (2003), 5604, pp. 232, doi:10.1126/science.1079699.
URL link

[6] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman and E. A. Cornell, *Observation Of Bose-Einstein Condensation In A Dilute Atomic Vapor*, Science **269** (1995), 5221, pp. 198, doi:10.1126/science.269.5221.198.
URL link

[7] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. Vandruten, D. S. Durfee, D. M. Kurn and W. Ketterle, *Bose-Einstein Condensation In A Gas Of Sodium Atoms*, Physical Review Letters **75** (1995), 22, pp. 3969, doi:10.1103/PhysRevLett.75.3969.
URL link

[8] E. A. Cornell and C. E. Wieman, *Nobel Lecture: Bose-Einstein condensation in a dilute gas, the first 70 years and some recent experiments*, Reviews of Modern Physics **74** (2002), 3, pp. 875, doi:10.1103/RevModPhys.74.875.
URL link

[9] W. Ketterle, *Nobel lecture: When atoms behave as waves: Bose-Einstein condensation and the atom laser*, Reviews of Modern Physics **74** (2002), 4, pp. 1131, doi:10.1103/RevModPhys.74.1131.
URL link

[10] T. Fukuhara, S. Sugawa and Y. Takahashi, *Bose-Einstein condensation of an ytterbium isotope*, Physical Review A **76** (2007), 5, p. 051604, doi:10.1103/PhysRevA.76.051604.
URL link

[11] G. Roati, M. Zaccanti, C. D'Errico, J. Catani, M. Modugno, A. Simoni, M. Inguscio and G. Modugno, *K-39 Bose-Einstein condensate with tunable interactions*, Physical Review Letters **99** (2007), 1, p. 010403, doi:10.1103/PhysRevLett.99.010403.
URL link

[12] B. DeMarco and D. S. Jin, *Onset of Fermi degeneracy in a trapped atomic gas*, Science **285** (1999), 5434, pp. 1703, doi:10.1126/science.285.5434.1703.
URL link

[13] A. G. Truscott, K. E. Strecker, W. I. McAlexander, G. B. Partridge and R. G. Hulet, *Observation of Fermi pressure in a gas of trapped atoms*, Science **291** (2001), 5513, pp. 2570, doi:10.1126/science.1059318.
URL link

[14] M. R. Andrews, C. G. Townsend, H. J. Miesner, D. S. Durfee, D. M. Kurn and W. Ketterle, *Observation of interference between two Bose condensates*, Science **275** (1997), 5300, pp. 637, doi:10.1126/science.275.5300.637.
URL link

[15] M. R. Matthews, B. P. Anderson, P. C. Haljan, D. S. Hall, C. E. Wieman and E. A. Cornell, *Vortices in a Bose-Einstein condensate*, Physical Review Letters **83** (1999), 13, pp. 2498, doi:10.1103/PhysRevLett.83.2498.
URL link

[16] J. R. Abo-Shaeer, C. Raman, J. M. Vogels and W. Ketterle, *Observation of vortex lattices in Bose-Einstein condensates*, Science **292** (2001), 5516, pp. 476, doi:10.1126/science.1060182.
URL link

[17] S. Inouye, M. R. Andrews, J. Stenger, H. J. Miesner, D. M. Stamper-Kurn and W. Ketterle, *Observation of Feshbach resonances in a Bose-Einstein condensate*, Nature **392** (1998), 6672, pp. 151, doi:10.1038/32354.
URL link

[18] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, S. Riedl, C. Chin, J. H. Denschlag and R. Grimm, *Bose-Einstein condensation of molecules*, Science **302** (2003), 5653, pp. 2101, doi:10.1126/science.1093280.
URL link

[19] M. Greiner, C. A. Regal and D. S. Jin, *Emergence of a molecular Bose-Einstein condensate from a Fermi gas*, Nature **426** (2003), 6966, pp. 537, doi:10.1038/nature02199.
URL link

[20] M. W. Zwierlein, C. A. Stan, C. H. Schunck, S. M. F. Raupach, S. Gupta, Z. Hadzibabic and W. Ketterle, *Observation of Bose-Einstein condensation of molecules*, Physical Review Letters **91** (2003), 25, doi:10.1103/PhysRevLett.91.250401.
URL link

[21] L. V. Hau, S. E. Harris, Z. Dutton and C. H. Behroozi, *Light speed reduction to 17 metres per second in an ultracold atomic gas*, Nature **397** (1999), 6720, pp. 594, doi:10.1038/17561.
URL link

[22] N. S. Ginsberg, S. R. Garner and L. V. Hau, *Coherent control of optical information with matter wave dynamics*, Nature **445** (2007), 7128, pp. 623, doi:10.1038/nature05493.
URL link

[23] T. Konemann, W. Brinkmann, E. Goklu, C. Lammerzahl, H. Dittus, T. Van Zoest, E. M. Rasel, W. Ertmer, W. Lewoczko-Adamczyk, M. Schiemangk, A. Peters, A. Vogel, G. Johannsen, S. Wildfang, K. Bongs, K. Sengstock, E. Kajari, G. Nandi, R. Walser and W. P. Schleich, *A freely falling magneto-optical trap drop tower experiment*, Applied Physics B-Lasers And Optics **89** (2007), 4, pp. 431, doi:10.1007/s00340-007-2863-8.
URL link

[24] M. Lewenstein, A. Sanpera, V. Ahufinger, B. Damski, A. Sen(De) and U. Sen, *Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond*, Advances In Physics **56** (2007), 2, pp. 243.
URL link

[25] L. Fallani, C. Fort, J. E. Lye and M. Inguscio, *Bose-Einstein condensate in an optical lattice with tunable spacing: transport and static properties*, Optics Express **13** (2005), 11, pp. 4303, doi:10.1364/OPEX.13.004303.
URL link

[26] S. Peil, J. V. Porto, B. L. Tolra, J. M. Obrecht, B. E. King, M. Subbotin, S. L. Rolston and W. D. Phillips, *Patterned loading of a Bose-Einstein condensate into an optical lattice*, Physical Review A **67** (2003), 5, p. 051603, doi:10.1103/PhysRevA.67.051603.
URL link

[27] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner and P. Zoller, *Cold bosonic atoms in optical lattices*, Physical Review Letters **81** (1998), 15, pp. 3108, doi:10.1103/PhysRevLett.81.3108.
URL link

[28] M. P. A. Fisher, P. B. Weichman, G. Grinstein and D. S. Fisher, *Boson Localization and the Superfluid-Insulator Transition*, Physical Review B **40** (1989), 1, pp. 546, doi:10.1103/PhysRevB.40.546.
URL link

[29] M. Greiner, O. Mandel, T. Esslinger, T. W. Hansch and I. Bloch, *Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms*, Nature **415** (2002), 6867, pp. 39, doi:10.1038/415039a.
URL link

[30] T. Stoferle, H. Moritz, C. Schori, M. Kohl and T. Esslinger, *Transition from a strongly interacting 1D superfluid to a Mott insulator*, Physical Review Letters **92** (2004), 13, doi:10.1103/PhysRevLett.92.130403.
URL link

[31] B. Paredes, A. Widera, V. Murg, O. Mandel, S. Folling, I. Cirac, G. V. Shlyapnikov, T. W. Hansch and I. Bloch, *Tonks-Girardeau gas of ultracold atoms in an optical lattice*, Nature **429** (2004), 6989, pp. 277, doi:10.1038/nature02530.
URL link

[32] T. Kinoshita, T. Wenger and D. S. Weiss, *Observation of a one-dimensional Tonks-Girardeau gas*, Science **305** (2004), 5687, pp. 1125, doi:10.1126/science.1100700.
URL link

[33] T. Kinoshita, T. Wenger and D. S. Weiss, *A quantum Newton's cradle*, Nature **440** (2006), 7086, pp. 900, doi:10.1038/nature04693.
URL link

[34] H. Ott, E. de Mirandes, F. Ferlaino, G. Roati, G. Modugno and M. Inguscio, *Collisionally induced transport in periodic potentials*, Physical Review Letters **92** (2004), 16, p. 160601, doi:10.1103/PhysRevLett.92.160601.
URL link

[35] S. Ospelkaus, C. Ospelkaus, O. Wille, M. Succo, P. Ernst, K. Sengstock and K. Bongs, *Localization of bosonic atoms by fermionic impurities in a three-dimensional optical lattice*, Physical Review Letters **96** (2006), 18, p. 180403, doi:10.1103/PhysRevLett.96.180403.
URL link

[36] M. Takamoto, F. L. Hong, R. Higashi and H. Katori, *An optical lattice clock*, Nature **435** (2005), 7040, pp. 321, doi:10.1038/nature03541.
URL link

[37] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hansch and I. Bloch, *Controlled collisions for multi-particle entanglement of optically trapped atoms*, Nature **425** (2003), 6961, pp. 937, doi:10.1038/nature02008.
URL link

[38] D. Schrader, I. Dotsenko, M. Khudaverdyan, Y. Miroshnychenko, A. Rauschenbeutel and D. Meschede, *Neutral atom quantum register*, Physical Review Letters **93** (2004), 15, p. 150501, doi:10.1103/PhysRevLett.93.150501.
URL link

[39] I. Bloch and M. Greiner, *Exploring quantum matter with ultracold atoms in optical lattices*, Advances In Atomic Molecular And Optical Physics, Vol 52 **52** (2005), pp. 1, doi:10.1088/0953-4075/38/9/013.
URL link

[40] O. Morsch and M. Oberthaler, *Dynamics of Bose-Einstein condensates in optical lattices*, Reviews of Modern Physics **78** (2006), 1, pp. 179, doi:10.1103/RevModPhys.78.179.
URL link

[41] J. F. Bertelsen, *Ultracold Atomic Gases*, Ph.D. thesis, University of Aarhus (2007).
URL link

[42] H. J. Lewandowski, D. M. Harber, D. L. Whitaker and E. A. Cornell, *Simplified system for creating a Bose-Einstein condensate*, Journal of Low Temperature Physics **132** (2003), 5-6, pp. 309, doi:10.1023/A:1024800600621.
URL link

[43] W. Ketterle and N. van Druten, *Evaporative cooling of trapped atoms*, Adv. At. Mol. Opt. Phys. **37** (1996), p. 181.
URL link

[44] C. V. Sukumar and D. M. Brink, *Spin-flip transitions in a magnetic trap*, Physical Review A **56** (1997), 3, pp. 2451, doi:10.1103/PhysRevA.56.2451.
URL link

[45] M. Yamashita and T. Mukai, *Stabilization of the number of Bose-Einstein-condensed atoms in evaporative cooling via three-body recombination loss*, Physical Review A **68** (2003), 6, p. 063601, doi:10.1103/PhysRevA.68.063601.
URL link

[46] K. B. Macadam, A. Steinbach and C. Wieman, *A Narrow-Band Tunable Diode-Laser System With Grating Feedback, And A Saturated Absorption Spectrometer For Cs And Rb*, American Journal Of Physics **60** (1992), 12, pp. 1098, doi:10.1119/1.16955.
URL link

[47] C. E. Wieman and L. Hollberg, *Using diode lasers for atomic physics*, Review of Scientific Instruments **62** (1991), p. 1.
URL link

[48] A. S. Arnold, J. S. Wilson and M. G. Boshier, *A simple extended-cavity diode laser*, Review Of Scientific Instruments **69** (1998), 3, pp. 1236, doi:10.1063/1.1148756.
URL link

[49] C. J. Hawthorn, K. P. Weber and R. E. Scholten, *Littrow configuration tunable external cavity diode laser with fixed direction output beam*, Review Of Scientific Instruments **72** (2001), 12, pp. 4477, doi:10.1063/1.1419217.
URL link

[50] H. K. Andersen, *Opbygning af et rubidium MOT lasersystem (only available in danish)*, Bachelor rapport (2003).

[51] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley and H. Ward, *Laser Phase And Frequency Stabilization Using An Optical-Resonator*, Applied Physics B: Lasers and Optics **31** (1983), 2, pp. 97, doi:10.1007/BF00702605.
URL link

[52] E. D. Black, *An introduction to Pound-Drever-Hall laser frequency stabilization*, American Journal Of Physics **69** (2001), 1, pp. 79, doi:10.1119/1.1286663.
URL link

[53] U. Schunemann, H. Engler, R. Grimm, M. Weidemuller and M. Zielonkowski, *Simple scheme for tunable frequency offset locking of two lasers*, Review Of Scientific Instruments **70** (1999), 1, pp. 242, doi:10.1063/1.1149573.
URL link

[54] E. L. Raab, M. Prentiss, A. Cable, S. Chu and D. E. Pritchard, *Trapping Of Neutral Sodium Atoms With Radiation Pressure*, Physical Review Letters **59** (1987), 23, pp. 2631, doi:10.1103/PhysRevLett.59.2631.
URL link

[55] C. Monroe, W. Swann, H. Robinson and C. Wieman, *Very Cold Trapped Atoms In A Vapor Cell*, Physical Review Letters **65** (1990), 13, pp. 1571, doi:10.1103/PhysRevLett.65.1571.
URL link

[56] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping* (1999).

[57] C. Wieman, G. Flowers and S. Gilbert, *Inexpensive Laser Cooling And Trapping Experiment For Undergraduate Laboratories*, American Journal Of Physics **63** (1995), 4, pp. 317, doi:10.1119/1.18072.
URL link

[58] W. H. Wing, *On Neutral Particle Trapping in Quasistatic Electromagnetic-Fields*, Progress in Quantum Electronics **8** (1984), 3-4, pp. 181, doi:10.1016/0079-6727(84)90012-0.
URL link

[59] M. D. Barrett, J. A. Sauer and M. S. Chapman, *All-optical formation of an atomic Bose-Einstein condensate*, Physical Review Letters **8701** (2001), 1, doi:10.1103/PhysRevLett.87.010404.
URL link

[60] T. Kinoshita, T. Wenger and D. S. Weiss, *All-optical Bose-Einstein condensation using a compressible crossed dipole trap*, Physical Review A **71** (2005), 1, p. 011602, doi:10.1103/PhysRevA.71.011602.
URL link

[61] C. J. Myatt, N. R. Newbury, R. W. Ghrist, S. Loutzenhiser and C. E. Wieman, *Multiply loaded magneto-optical trap*, Optics letters **21** (1996), 4, p. 290.
URL link

[62] M. Greiner, I. Bloch, T. W. Hansch and T. Esslinger, *Magnetic transport of trapped cold atoms over a large distance*, Physical Review A **6303** (2001), 3, doi:10.1103/PhysRevA.63.031401.
URL link

[63] D. J. Griffiths, *Introduction to electrodynamics* (Benjamin Cummings; 3rd edition, 1999).

[64] N. R. Thomas, A. C. Wilson and C. J. Foot, *Double-well magnetic trap for Bose-Einstein condensates*, Physical Review A **65** (2002), 6, p. 063406, doi:10.1103/PhysRevA.65.063406.
URL link

[65] J. Goldwin, S. B. Papp, B. DeMarco and D. S. Jin, *Two-species magneto-optical trap with K-40 and Rb-87*, Physical Review A **65** (2002), 2, p. 021402, doi:10.1103/PhysRevA.65.021402.
URL link

[66] S. Hensler, A. Griesmaier, J. Werner, A. Gorlitz and T. Pfau, *A two species trap for chromium and rubidium atoms*, Journal Of Modern Optics **51** (2004), 12, pp. 1807, doi:10.1080/09500340408232492.
URL link

[67] M. Taglieber, A. C. Voigt, F. Henkel, S. Fray, T. W. Hansch and K. Dieckmann, *Simultaneous magneto-optical trapping of three atomic species*, Physical Review A **73** (2006), 1, p. 011402, doi:10.1103/PhysRevA.73.011402.
URL link

[68] D. T. Ashley, J. P. DeVoe, K. Perttunen and C. Pratt, *On Best Rational Approximations Using Large Integers*.

[69] T. Gericke, F. Gerbier, A. Widera, S. Folling, O. Mandel and I. Bloch, *Adiabatic loading of a Bose-Einstein condensate in a 3D optical lattice*, Journal Of Modern Optics **54** (2007), 5, pp. 735, doi:10.1080/09500340600777730.
URL link

[70] S. Mai, *Part A report—will be available at http://www.phys.au.dk/quantumgas sometime in the future.*

[71] P. W. Milonni and J. H. Eberly, *Lasers* (1988).

[72] D. A. Steck, *Rubidium 87 D Line Data*, revision 1.6, 14 October 2003.
URL link

[73] Y. Castin and R. Dum, *Bose-Einstein condensates in time dependent traps*, Physical Review Letters **77** (1996), 27, pp. 5315, doi:10.1103/PhysRevLett.77.5315.
URL link

[74] U. Ernst, J. Schuster, F. Schreck, A. Marte, A. Kuhn and G. Rempe, *Free expansion of a Bose-Einstein condensate from an Ioffe-Pritchard magnetic trap*, Applied Physics B-Lasers And Optics **67** (1998), 6, pp. 719, doi:10.1007/s003400050571.
URL link

[75] R. Grimm, M. Weidemuller and Y. B. Ovchinnikov, *Optical dipole traps for neutral atoms*, Advances In Atomic Molecular, And Optical Physics, Vol. 42 **42** (2000), pp. 95.
URL link

[76] C. Kittel, *Introduction to Solid State Physics* (Wiley; 7th edition, 1995).

[77] M. Greiner, *Ultracold quantum gases in three-dimensional optical lattice potentials*, Ph.D. thesis (2003).
URL link

[78] T. Gericke, P. Würtz, D. Reitz, T. Langen and H. Ott, *High Resolution Imaging of Single Atoms in a Quantum Gas*, arxiv.org/abs/0804.4788 (2008).
URL link

[79] J. H. Denschlag, J. E. Simsarian, H. Haffner, C. McKenzie, A. Browaeys, D. Cho, K. Helmerson, S. L. Rolston and W. D. Phillips, *A Bose-Einstein condensate in an optical lattice*, Journal Of Physics B-Atomic Molecular And Optical Physics **35** (2002), 14, pp. 3095, doi:10.1088/0953-4075/35/14/307.
URL link

[80] R. Sapiro, R. Zhang and G. Raithel, *1-D Mott insulator transition of a Bose-Einstein condensate*, Unpublished .
URL link

[81] T. L. Ho and Q. Zhou, *Intrinsic heating and cooling in adiabatic processes for bosons in optical lattices*, Physical Review Letters **99** (2007), 12, p. 120404, doi:10.1103/PhysRevLett.99.120404.
URL link

[82] J. Plata, *Loading of a Bose-Einstein condensate into an optical lattice: The excitation of collective modes*, Physical Review A **69** (2004), 3, p. 033604, doi:10.1103/PhysRevA.69.033604.
URL link

[83] S. Ospelkaus-Schwarzer, *Quantum Degenerate Fermi-Bose Mixtures of 40K and 87Rb in 3D Optical Lattices*, Ph.D. thesis (2006).
URL link

[84] P. Horowitz and H. Winfield, *The Art of Electronics* (Cambridge University Press, 1989).

[85] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hansch and I. Bloch, *Coherent transport of neutral atoms in spin-dependent optical lattice potentials*, Physical Review Letters **91** (2003), 1, doi:10.1103/PhysRevLett.91.010407.
URL link

[86] V. Horvatic, T. L. Corrella, N. Omenetto, C. Vadla and J. D. Winefordner, *The effects of saturation and velocity selective population in two-step 6S(1/2)-> 6P(3/2)-> 6D(5/2) laser excitation in cesium*, Spectrochimica Acta Part B-Atomic Spectroscopy **61** (2006), 12, pp. 1260, doi:10.1016/j.sab.2006.10.010.
URL link

[87] M. R. Andrews, M. O. Mewes, N. J. vanDruten, D. S. Durfee, D. M. Kurn and W. Ketterle, *Direct, nondestructive observation of a bose condensate*, Science **273** (1996), 5271, pp. 84, doi:10.1126/science.273.5271.84.
URL link

[88] R. Jauregui, N. Poli, G. Roati and G. Modugno, *Anharmonic parametric excitation in optical lattices*, Physical Review A **6403** (2001), 3, p. 033403, doi:10.1103/PhysRevA.64.033403.
URL link

[89] M. Cristiani, O. Morsch, J. H. Muller, D. Ciampini and E. Arimondo, *Experimental properties of Bose-Einstein condensates in one-dimensional optical lattices: Bloch oscillations, Landau-Zener tunneling, and mean-field effects*, Physical Review A **65** (2002), 6, p. 063612, doi:10.1103/PhysRevA.65.063612.
URL link

[90] S. Mai, *Magnetic Sisyphus Cooling*, Master's thesis, University of Aarhus (2007).
URL link

[91] B. P. Anderson and M. A. Kasevich, *Macroscopic quantum interference from atomic tunnel arrays*, Science **282** (1998), 5394, pp. 1686, doi:10.1126/science.282.5394.1686.
URL link

[92] G. Modugno, E. de Mirandes, F. Ferlaino, H. Ott, G. Roati and M. Inguscio, *Atom interferometry in a vertical optical lattice*, Fortschritte Der Physik-Progress Of Physics **52** (2004), 11-12, pp. 1173, doi:10.1002/prop.200410187.
URL link