Spectrally Probing Ultracold Bosonic Clouds

Rasmus Malthe Fiil Andersen
Department of Physics and Astronomy
Aarhus University, Denmark

Advisor: Jan Joachim Arlt

PhD thesis
October 22, 2025

Front cover illustration: The diffracted electric field intensity right after passing through a bimodal cloud at intermediate light detuning.

Abstract

Quantum mechanics has played a pivotal role in the evolution of both the fundamental understanding of physics and technological advancements in the recent century. As science and technology progress, systems with yet greater complexity are tackled. Ultracold dilute gases and Bose–Einstein condensates (BECs), first realized in 1995, provide a highly controlled platform for investigating many-body quantum phenomena. However, increasing demands on experimental precision place stringent requirements on probing techniques. Conventional absorption imaging after time-of-flight expansion avoids the complexity of atom–light interactions in dense media but is typically destructive and subject to systematic errors arising from cloud expansion modeling, imaging moving clouds, and the pressure from the imaging light.

This thesis addresses both free-expansion and in-trap probing of ultracold atoms. First, the expansion of partially condensed clouds is modeled with high accuracy, incorporating interatomic interactions, and the systematic errors arising from neglecting these effects are quantified. Second, a general model of in-trap probing is developed, accounting for spatially dependent absorption and phase shifts for arbitrary probe detunings. The model reveals complex spectral features in dense clouds and is shown to yield quantitative agreement with established methods, while enabling less destructive measurements and increased sensitivity.

To realize these advances experimentally, a new detection setup with high flexibility and sensitivity was constructed and benchmarked against conventional approaches. The system achieves high measurement accuracy and enables extended measurement protocols. Together, these developments provide a foundation for probing complex quantum behavior in ultracold gases.

Resume

Kvantefysikken har spillet en afgørende rolle i udviklingen af både den fundamentale forståelse af fysik og de teknologiske fremskridt i det seneste århundrede. Efterhånden som videnskab og teknologi udvikler sig, undersøges systemer med stadig større kompleksitet. Ultrakolde gasser og Bose–Einstein-kondensater (BEC'er), som første gang blev realiseret i 1995, udgør en højt kontrollerbar platform til at undersøge mange-partikel kvantefænomener. En stigende eksperimentel præcision stiller imidlertid skærpede krav til måleteknikker. Konventionel absorptionsafbildning efter time-of-flight ekspansion undgår kompleksiteten i atom–lys-interaktioner under høj tæthed, men er typisk destruktiv og udsat for systematiske fejl forbundet med modelleringen af ekspansionen, afbildning af skyer i bevægelse of trykket forårsaget af afbildningslyset.

Denne afhandling behandler både målinger efter ekspansion og in-trap probing af ultrakolde atomer. For det første modelleres ekspansionen af delvist kondenserede skyer med høj nøjagtighed, idet interatomare vekselvirkninger inddrages, og de systematiske fejl ved at ignorere disse effekter kvantificeres. For det andet udvikles en generel model for in-trap probing, som tager højde for den rumligt afhængige absorption og faseskift ved vilkårlige probefrekvenser. Modellen afslører komplekse spektrale strukturer i tætte skyer og viser sig at være i kvantitativ overensstemmelse med etablerede metoder, og muliggør mindre destruktive målinger samt øget følsomhed.

For at realisere disse fremskridt eksperimentelt er et nyt detektionssystem med høj præcision, fleksibilitet og følsomhed blevet konstrueret og sammenlignet med konventionelle tilgange. Systemet opnår forbedret målenøjagtighed og muliggør mere avancerede måleprotokoller. Samlet set danner disse udviklinger grundlag for at undersøge komplekse kvantemekaniske fænomener i ultrakolde gasser.

Preface

My PhD journey began in September 2022, when I joined Jan Arlt's group on ultracold bosonic gases in the Lattice Lab experiment. I entered with a primarily theoretical background, having worked on QFT simulations in my bachelor's project [1] and on quantum computing error-correction algorithms during my master's project. At that time, my interest in physics was starting to wane, and I was facing a minor life crisis about my future direction. I realized, however, that I wished to engage with something more technical, and I took the gamble of pursuing an experimental PhD. Three years later, I can now say the gamble paid off. Throughout this project, I have acquired a wealth of new skills, both in physics and more generally, I have met many inspiring people, and regained my passion for the field.

When I joined the group, two colleagues were already working on the experiment. Laurits Stokholm was finishing his master's degree, and Toke Vibel was involved in the technical analysis of previous results from the group. In the early stages of my PhD, I primarily worked with Toke, and I am grateful for his excellent introduction to both the experiment and to life as a PhD student. In the demanding world of academia, his insistence on respect for his time and boundaries has left a lasting impression on me.

Laurits remained in the group, first as a technical assistant and later as a PhD student, and as Toke shifted his focus to writing his thesis, Laurits became my main collaborator on the experiment. The experiment, which is now about 20 years old, gradually revealed a pressing need for extensive maintenance. Consequently, both Laurits and I spent significant portions of our PhDs troubleshooting and completely rebuilding worn-out sections of the setup. While this inevitably resulted in fewer immediate scientific results than we might have hoped for, it also led to a renewed experiment that is now well-equipped to run for many years to come. In hindsight, I believe this challenge was ultimately a gift, forcing Laurits and me to gain a much deeper understanding of the technical details of such setups. I especially want to thank Laurits for maintaining high spirits through numerous setbacks and limited successes. The experiment would not have reached its current maturity without his persistence and contributions.

In the midst of experimental turbulence, a new PhD student, Ilja Zebergs, joined the Lattice Lab. Despite the challenges, he quickly familiarized himself with the experiment and began contributing meaningfully. With both technical skills and theoretical knowledge, Ilja has helped drive the project forward under significant time pressure, and I am grateful for his dedication and positive spirit.

I also wish to thank Adam Chatterley, whose technical expertise has been invaluable. He provided crucial guidance during some of the most difficult phases of the rebuild process, and I attribute part of the success of my PhD to his support and technical proficiency.

A special thanks is owed to my supervisor, Jan Arlt, for his excellent guidance in the academic world. He approaches mistakes with understanding and constructive advice, never with judgment. I am grateful for the freedom he granted me, allowing me to play an active role in shaping the decision-making and direction of the experiment.

On a broader note, I want to thank all members of the group — the PhDs of the Lattice Lab, Jan Arlt, Adam Chatterley, as well as the PhDs from the other experiment in the group, Andreas Morgen, Søren Balling, and Morten Strøe. Together, they have created a safe and welcoming working environment where all questions can be asked freely, and where coming to work feels less like an obligation and more like a privilege.

Finally, I am grateful to my family, friends, and partner, Vy Vũ, for their support and for providing a space where I could disconnect from work. While they did not contribute directly to the research, their presence in my life has been essential to the completion of this degree.

Contents

Bibliography

Chapter 1

Introduction

When Satyendra Nath Bose and Albert Einstein published a theoretical paper predicting a new, purely quantum state of matter at ultracold temperatures [2], now called Bose-Einstein condensation (BEC), based on Bose's initial predictions about photon statistics [3], it was seen mostly as a theoretical curiosity of statistical quantum mechanics. Fast forward to modern physics, BEC has proven to be an important experimental platform to investigate complex many-particle quantum systems, with both high precision and high controllability. Additionally, as the field of ultracold atoms has matured, technological applications are coming within reach [4], such as quantum metrology [5–8], quantum simulators [9, 10], and quantum computing platforms [11].

As applications of ultracold atoms develop, the required precision and diversity of probing methods follow. To achieve the best possible results, it is imperative to have access to a variety of different high precision detection tools, such that the optimal methodology can be chosen in each scenario. Developing and optimizing such techniques is an on-going scientific process, with detailed investigations into the widely adopted absorption imaging method [12], and the appearance of novel detection schemes with multiple cameras working together to reconstruct the 3-dimensional propagation of light fields [13].

This dissertation revisits the free fall expansion model of partially condensed ultracold clouds, which is used to extract information such as the number of atoms and the temperature. By including the repulsive forces between atoms inside the trap, an important source of inaccuracy can be removed. Since the measurement of parameters in cold clouds is important for a significant amount of experiments this has the potential to increase accuracy in a wide range of scenarios.

The main outcome of this thesis is the development of an entirely new method of detecting ultracold clouds non-destructively, with potentially significant increase in bandwidth compared to any methods implementing cameras, and for a broad range of optical densities. Instead of measuring the spatial distribution with cameras, a time-resolved spectrum is recorded with a sensitive single photon detector. It was found that intricate spectroscopic details emerge and carries the information of the cold clouds.

To my knowledge, the only implementation of spectra to measure the fundamental properties of ultracold clouds was in hydrogen BECs [14], by detecting the shift in energy caused by interactions in the BEC on a narrow linewidth transition. Our new approach is fundamentally different, and extracts the cloud parameters from spectral structures arising solely as a consequence of spatial variation in optical densities.

While the newly developed spectral detection method works well at extracting cloud parameters, it also enables measurements of various interesting physical phenomena in mesoscopic atomic ensembles. The refined spectral model leads to sensitivity to any deviations from the predicted spectra, which can be exploited as a tool to measure the underlying effects. One such example could be the light induced dipole-dipole interaction between neighboring atoms causing a broadened natural linewidth. This effect occurs at high densities of atoms, and a measurement is therefore likely to be accompanied by high optical densities. Analysis of such spectra requires the incorporation of the effects described in this model to yield meaningful results.

The experimental setup developed to detect the spectra has the additional benefit of allowing for time-resolved measurements. This can be used in a number of novel experiments, and will benefit from the detailed description of spectra. Measurements of optical Rabi oscillations can be a tool to observe several many-particle effects, such as sub- and superradiance, or even as a complementary measurement of the light induced dipole-dipole broadening. Experimental parameters can be extracted substantially quicker with time-resolved measurements, and real-time detections of evolving parameters can be performed.

Production, manipulation, and probing of BECs is a complicated process. A part of this thesis deals with the maintenance of equipment and addition of new setups in the experiment. After many years of reliable BEC production, the experiment was in need of an overhaul. This has been completed, and the experiment is once again performing well. The cycle time of the experiment has been significantly improved, which is important for faster scientific progress.

Structure of the thesis

Chapter 2, the first chapter after the introduction, covers the standard knowledge on ultracold gases in harmonic traps.

Chapter 3 investigates the expansion model for partially condensed clouds in detail by estimating the consequences of repulsive interatomic forces.

Chapter 4 introduces the theory of light propagation in ultracold clouds, and develops a detailed spectroscopic model.

Chapter 5 describes the experimental setup used to produces BEC, and includes detailed descriptions of all the new additions to the experiment during this project.

Chapter 6 is split in two parts. The first part includes the main results from deploying the newly developed spectroscopic measurement technique to extract the cloud parameters non-destructively. The second part illustrates how the time-resolved measurements can potentially be utilized.
Chapter 7, the final chapter, concludes the thesis and provide an outlook.

Chapter 2

Ultracold atoms

This chapter describes the physics of ultracold bosonic clouds and the interesting phenomenon of Bose-Einstein condensation (BEC). It was first predicted theoretically by Satyendra Nath Bose and Albert Einstein in the 1920s [2], but due to the extremely cold temperatures required to produce it, about 70 years would pass before an experimental observation of BEC was made. The observations took place in three groups separated in time by only a few months [15–17].

The chapter is split into two sections. First, Section 2.1 covers the in-trap shapes of ultracold clouds, followed by Section 2.2 which deals with the expansion of a released, free-falling cloud. The chapter is based mainly on [18, 19].

2.1 Cloud distributions

Statistically, there exist two kinds of particles: distinguishable particles and indistinguishable particles. Distinguishable particles require each permutation to be considered separately, while indistinguishable particles only care about the total number of particles occupying each

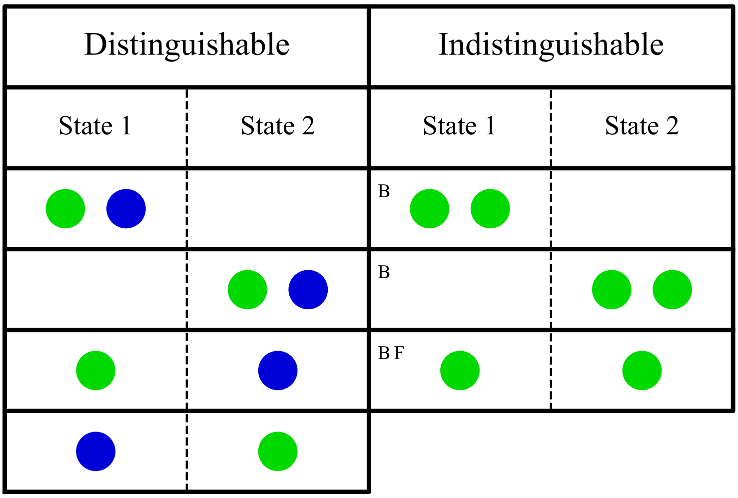

Figure 2.1: Particle statistics in a system with two particles and two available states. There are two main columns: distinguishable and indistinguishable particles. The system with distinguishable particles has 4 different possible states, while the indistinguishable particle system has either 3 possible states for bosons (indicated with a ”B” letter in the top left) or only 1 possible state for fermions (indicated with a ”F” letter in the top left).

state. Indistinguishable particles can be split into two sub-categories, named bosons and fermions. Bosons can have several particles in the same state, while fermions can have at most one particle in each state.

Figure 2.1 illustrate the different types of particles in a system with 2 particles and 2 possible states to occupy. The distinguishable particles have 4 possible states, the indistinguishable bosonic particles have three different possible states, and the indistinguishable fermionic particles have only one possible state. Notice how the bosonic particles have a higher chance of existing in the same state ($\frac{2}{3}$ in this example) compared to all other types, simply due to statistics. This is, ultimately, the mechanics behind BEC. The rest of the chapter deals with trapped indistinguishable boson gases only.

2.1.1 Bose-Einstein statistics

For non-interacting bosonic particles in a grand-canonical ensemble, it can be shown [19] that a system with temperature T and chemical potential μ follows the Bose-Einstein distribution

$$N_i = \frac{1}{\mathrm{e}^{(E_i-\mu)/k_B T} - 1}, \tag{2.1}$$

where N_i is the average number of particles in state i with single particle energy E_i. When the exponent $(E_i - \mu)/k_B T$ becomes large, the distribution approaches the Boltzmann distribution. In this limit, N_i is small, which happens either if the total number of particles in the system is small or if the particles are spread over many states at high temperature.

The particles in an ultracold experiment are neither non-interacting nor is the ensemble grand-canonical. It is rather micro-canonical since there is limited exchange of particles or energy with the surroundings. However, the non-interacting, grand-canonical case is far simpler to deal with. By requiring the total energy E and the total number of particles N to be held constant in eq. 2.1 it is possible to get results that agree with the micro-canonical description

$$E = \sum_i E_i N_i \quad \text{and} \quad N = \sum_i N_i. \tag{2.2}$$

These equations determine the temperature and chemical potential of the system. This approach only works to some extent, however. If the same method is used to calculate the fluctuation in particle numbers, the result diverges to infinity at low temperatures [20–22].

When the exponential in the denominator approaches one, the distribution starts to differ significantly from the Boltzmann distribution. The consequence is that a macroscopic number of particles fall into the ground state of the system. It therefore makes sense to distinguish between the particles in the ground state, which will be called "condensed particles" and all other particles, henceforth called "thermal particles". Section 2.1.2 explores the thermal particles and Section 2.1.3 explores the condensed particles.

2.1.2 Thermal clouds

In all trapping potentials $V(\mathbf{r})$, the energy levels are discrete, but the thermal particles are distributed over enough energy levels to approximate the distribution as continuous. This is equivalent to approximating the distribution as non-quantum, giving the energies $E = \frac{\mathbf{p}^2}{2m^2} + V(\mathbf{r})$. The distribution in eq. 2.1 is multiplied by a density of states $g(\mathbf{r}, \mathbf{p}) = \frac{1}{(2\pi\hbar)^3}$ to get a probability density distribution ρ_{th} of thermal particles

$$\rho_{\text{th}}(\mathbf{r}, \mathbf{p}) = \frac{1}{(2\pi\hbar)^3} \frac{1}{\mathrm{e}^{\left(\frac{\mathbf{p}^2}{2m^2} + V(\mathbf{r}) - \mu\right)/k_B T} - 1}. \tag{2.3}$$

The spatial distribution can be found by integrating over all momenta and gives

$$\rho_{\text{th}}(\mathbf{r}) = \int \rho_{\text{th}}(\mathbf{r}, \mathbf{p}) \mathrm{d}\mathbf{p} \tag{2.4}$$

$$= \left(\frac{m k_B T}{2\pi\hbar^2}\right)^{3/2} g_{3/2}\left(\mathrm{e}^{-(V(\mathbf{r}) - \mu))/k_B T}\right), \tag{2.5}$$

where the polylogarithm function $g_\gamma(x)$ is

$$g_\gamma(x) = \sum_{i=1}^{\infty} \frac{x^i}{i^\gamma}. \tag{2.6}$$

In a harmonic trap $V(\mathbf{r}) = \frac{1}{2}m\sum\omega_i^2 r_i^2$, and the distribution is called the Bose-enhanced distribution. It resembles a Gaussian distribution except it is more concentrated in the center. When the exponent in the ρ_{th} distribution is large and the distribution approaches the Boltzmann distribution, the exponent in eq. 2.4 is small, at the polylogarithm function can be taken to first order

$$\rho_{\text{th}}(\mathbf{r}) \approx \left(\frac{mk_BT}{2\pi\hbar^2}\right)^{3/2} \mathrm{e}^{-(V(\mathbf{r})-\mu))/k_BT}, \tag{2.7}$$

which is now a true Gaussian distribution. The takeaway is that at high temperatures or low particle numbers, the thermal particles are spatially distributed like a Gaussian distribution, and when the cloud gets cold enough, the particles start to fall towards the center to a greater extent than would be expected by distinguishable (or classical) particles. This effect is sometimes described as an attractive force between bosonic particles, which is a useful way to picture the complicated statistical effects. There is no attractive force[1], however, and the gathering of particles in the center is purely a consequence of bosonic statistics.

2.1.3 Bose-Einstein condensates

This section focuses on the particles in the ground state of the system, and it is now necessary to treat the particles quantum mechanically. The interactions between particles can no longer be ignored since a macroscopic number of particles gathered in the lowest energy state will lead to a substantially higher particle density than that of the thermal particles. In a mean-field approximation, the point-like, S-wave scattering between particles leads to an effective additional potential $U\rho_0(\mathbf{r})$, $U = \frac{4\pi\hbar^2 a}{m}$ where a is the scattering length, $\rho_0(\mathbf{r}) = |\psi(\mathbf{r})|^2$

---

[1] The description so far still assumes no interactions between particles.

is the ground state particle density, and ψ is the wave function of the ground state but normalized to the number of ground state particles N_0 rather than 1. The Schrödinger equation for the ground state becomes [18, 23–25]

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + U|\psi(\mathbf{r})|^2\right)\psi(\mathbf{r}) = \mu\psi(\mathbf{r}). \tag{2.8}$$

This equation is called the Gross-Pitaevskii (GP) equation. The non-linear interaction causes the distribution of condensed particles to differ from that of the potential $V(\mathbf{r})$ alone.

2.1.3.1 Thomas-Fermi approximation

The three terms on the left side of eq. 2.8 is, from left to right, the kinetic energy of the system, the external potential energy of the system, and the effective interaction energy of the system. Assume now the interaction is repulsive. At low enough temperatures and enough particles for the macroscopically large ground state to form, the kinetic term becomes relatively small compared to the potential terms. To reach an approximate solution to the non-linear GP differential equation, which is otherwise complicated to solve, the kinetic term can be neglected, giving

$$\left(V(\mathbf{r}) + U|\psi(\mathbf{r})|^2\right)\psi(\mathbf{r}) = \mu\psi(\mathbf{r}). \tag{2.9}$$

Since only the kinetic term contains derivatives, the solution is now given simply by solving the algebraic expression

$$\rho_0(\mathbf{r}) = |\psi(\mathbf{r})|^2 = \left\{ \begin{array}{cc} \frac{1}{U}(\mu - V(\mathbf{r})), & \mu - V(\mathbf{r}) \geq 0 \\ 0, & \mu - V(\mathbf{r}) < 0 \end{array} \right. \tag{2.10}$$

The repulsive interaction has forced the distribution away from the ground state of the trap and into the reverse shape of the potential.

The approximation is called the Thomas-Fermi (TF) approximation. In a harmonic trap, the cloud is a reversed parabola and can be written in a more convenient form

$$\rho_0(\mathbf{r}) = \frac{15}{8\pi}\frac{N_0}{R_1R_2R_3}\max\left(0, 1 - \frac{x_1^2}{R_1^2} - \frac{x_2^2}{R_2^2} - \frac{x_3^2}{R_3^2}\right). \tag{2.11}$$

The radii of the cloud R_i, $i = 1, 2, 3$ are when $\mu - V(\mathbf{r}) = 0$ and for a harmonic trap are

$$R_i = \sqrt{\frac{2\mu}{m}\frac{1}{\omega_i}}, \tag{2.12}$$

and are called the Thomas-Fermi radii. Integrating over the cloud gives the relation between the chemical potential and the condensed particle number

$$N_0 = \frac{1}{15}\frac{a_{\text{HO}}}{a}\left(\frac{\mu}{\hbar\overline{\omega}}\right)^{5/2}, \tag{2.13}$$

where $a_{\text{HO}} = \sqrt{\hbar/m\overline{\omega}}$ is the harmonic oscillator length. The TF approximation is remarkably robust considering how extreme the simplification is.

2.1.3.2 Critical temperature

To find the temperature where the particles start to gather in the ground state, it is useful to rewrite eq. 2.3 into a continuous function of energy rather than phase-space coordinates

$$\rho_{\text{th}}(E) = g(E)\frac{1}{\mathrm{e}^{(E-\mu)/k_BT} - 1}. \tag{2.14}$$

The density of states in a harmonic trap is now $g(E) = \frac{E^2}{2(\hbar\overline{\omega})^3}$ where $\overline{\omega} = (\omega_1\omega_2\omega_3)^{1/3}$ is the geometric mean of the trap frequencies. According to the TF approximation (eq. 2.10), the particles start to fall to the ground state when μ crosses from negative to positive. The critical temperature T_c must therefore be the temperature at $\mu = 0$ that fulfills the condition that the number of all particles N is still thermal

$$N = N_{\mathrm{th}} = \int_0^\infty \frac{E^2}{2(\hbar\overline{\omega})^3} \frac{1}{\mathrm{e}^{E/k_B T_c} - 1} \mathrm{d}E \tag{2.15}$$

$$= \zeta(3) \left(\frac{k_B T_c}{\hbar\overline{\omega}} \right)^3 , \tag{2.16}$$

where ζ is the Riemann zeta function. The critical temperature is therefore

$$T_c = \frac{\hbar\overline{\omega}}{k_B \zeta(3)^{1/3}} N^{1/3}. \tag{2.17}$$

The formulas used here assumed no particle-particle interactions and that the lowest energy level is truly zero. In a harmonic oscillator, the lowest energy level is instead $\frac{3}{2}\hbar\overline{\omega}$ and μ should be set to this value instead of 0 to find T_c. Correcting for this gives a shift ΔT_c in the critical temperature [26]

$$\frac{\Delta T_c}{T_c} = -0.73 \frac{\overline{\omega}}{\langle \omega \rangle} N^{-1/3}, \tag{2.18}$$

where $\langle \omega \rangle$ is the mean of the trap frequencies. The correction due to interactions is

$$\frac{\Delta T_c}{T_c} = -1.33 \frac{a}{a_{\mathrm{HO}}} N^{1/6}. \tag{2.19}$$

Further discussion on critical temperature can be found in [27, 28].

2.1.4 Semi-ideal model and bimodal clouds

The two previous sections dealt with thermal and condensed clouds separately. However, the interactions shaping the condensed cloud into the shape of the reversed potential also happen between thermal and condensed particles. In [29], this is taken care of by adding the mean field potential from the condensed cloud $U\rho_0(\mathbf{r})$ to the external potential in eq. 2.4 which becomes

$$\rho_{\text{th}}(\mathbf{r}) = \left(\frac{mk_BT}{2\pi\hbar^2}\right)^{3/2} g_{3/2}\left(\mathrm{e}^{-|V(\mathbf{r})-\mu)|/k_BT}\right). \tag{2.20}$$

The result is a dip in the middle of the thermal distribution and an increased thermal density around the condensed cloud. Since the thermal cloud density is much lower than the condensed cloud, the effect of the thermal particles on the condensate is ignored, leaving the TF-distribution unchanged. This model is called the semi-ideal model.

The only free parameters in the model are the temperature and the chemical potential. The chemical potential directly gives the condensed number of particles through the TF-approximation, while the thermal particle number comes from the integral of eq. 2.20 and is given by both the temperature and the chemical potential. Since the number of thermal particles depends on the chemical potential, it is not possible to express the condensed fraction N_0/N in terms of temperature alone. Approximate expressions for this are given in [29], with the most precise being an equation that has to be solved numerically

$$\frac{N_0}{N} = 1 - \left(\frac{T}{T_c}\right)^3 - \eta\frac{\zeta(2)}{\zeta(3)}\left(\frac{T}{T_c}\right)^2\left(\frac{N_0}{N}\right)^{\frac{2}{5}}, \tag{2.21}$$

where

$$\eta = \frac{1}{2}\zeta(3)^{\frac{1}{3}}\left(15N^{\frac{1}{6}}\frac{a}{a_{\text{HO}}}\right)^{\frac{2}{5}}. \tag{2.22}$$

2.2 Cloud expansion - analytical expressions

One of the most used techniques to acquire data about cold clouds is to rapidly turn off the trap and take images after a given duration of free fall, called time-of-flight (TOF). The distributions after TOF are related to a combination of the spatial and momentum distributions in-trap and can be used to measure temperature more easily than imaging the very small in-trap clouds. To do this, it is necessary to have expressions for the expansion of the clouds. Section 2.2.1 deals with the free fall expansion of pure thermal clouds and Section 2.2.2 with the expansion of pure condensates. There are no analytical expressions for the expansion of a bimodal cloud. One way to approximate the expansion is to use the expressions for pure thermal clouds and condensates separately, which is the standard method. Chapter 3 delves deeper into this question via numerical simulations.

2.2.1 Thermal clouds

Assume the particles move ballistically after the release of the trap. The momenta in the in-trap phase-space distribution of pure thermal clouds (eq. 2.3) can then be used to find the distribution of particles after a duration t of TOF, $\rho(\mathbf{r}, \mathbf{p}, t)$, by integrating over the momenta and using $\mathbf{r}_{\text{in-trap}} = \mathbf{r} - \mathbf{p}/m$

$$\rho(\mathbf{r}, \mathbf{p}, t) = \int \rho_{\text{th}}(\mathbf{r} - \mathbf{p}/m, \mathbf{p}, t = 0)\mathrm{d}\mathbf{p}. \tag{2.23}$$

In a harmonic oscillator, there is an analytical expression for the integral [30]

$$\rho(\mathbf{r},\mathbf{p},t) = \frac{N_{\mathrm{th}}}{(2\pi)^3 g_3(\tilde{z}) \prod w_i} g_{3/2}\left(\tilde{z}\mathrm{e}^{-\sum r_i^2/(2w_i^2)}\right) \tag{2.24}$$

$$\tilde{z} = \mathrm{e}^{\mu/k_B T}, \quad w_i = \sqrt{\frac{k_B T}{m}}\sqrt{1+\omega_i^2 t^2}. \tag{2.25}$$

Interestingly, the cloud shape stays the same and only the widths w_i change with time. This can be approximated to first order in the polylogarithm function for temperatures that are not close to or below the critical temperature for condensation and gives a Gaussian distribution with the same widths as in the equation above.

2.2.2 Bose-Einstein condensates

Just like pure thermal clouds, an expanding condensate also keeps its shape while the radii change [31]. Write the radii after TOF as $R_i(t) = \lambda_i R_{i,0}$, where $R_{i,0}$ are the in-trap TF-radii from eq. 2.12. The radii during TOF then obey a differential equation

$$\ddot{\lambda}_i = \frac{\omega_i^2}{\lambda_i \lambda_1 \lambda_2 \lambda_3}. \tag{2.26}$$

A condensate does not expand ballistically but is instead governed by quantum mechanics. A smaller in-trap spatial distribution means the momentum distribution must be wider and therefore contain higher momentum components, resulting in a faster expansion and a larger cloud after TOF. This is true for each dimension separately. The result is that a cloud which is initially smaller in one direction than another will expand into the opposite scenario, such that the small axis turns into the large axis after TOF.

These equations still hold in the semi-ideal model since that model assumed no effect on the condensate from the thermal cloud.

Chapter 3

Expanding cloud simulations

In many ultracold atom experiments, information is gained from absorption imaging. However, the in-trap optical densities are too high for any meaningful amount of light to penetrate the cloud. Consequently, the standard procedure is to release the cloud in free fall for a short period, resulting in the expansion of the cloud and thus a reduction in optical density.

To extract information from such images, it is necessary to model the expansion of ultracold atom clouds. This is a relatively simple problem to solve analytically for ideal gas clouds (released from a harmonic trap potential) [12, 32], which is a good approximation in many cold atom experiments since the densities of the clouds are low enough for interactions to be negligible. However, when a gas approaches degeneracy and obtains a BEC component, a macroscopic number of atoms fall into the same ground state, resulting in a sharp increase in density. The consequence is that both the ideal gas assumption and classical physics break down.

In a mixed state of BEC and thermal atoms, the thermal component can be approximated as an ideal gas, and any interactions between the two components is typically be ignored, thus decoupling the BEC and thermal atoms.

However, the high-density BEC atoms, which break the ideal gas assumption, interact with all atoms, including the thermal atoms. This leads to a more complicated system where an analytical solution to the free expansion has not yet been found. A numerical model is therefore used to estimate what effect the presence of BEC atoms will have on the thermal cloud and specifically what the consequential errors on the estimation of the two key parameters, temperature, T, and thermal atom number, N_{th}, is when this effect is ignored.

The chapter is structured as follows. Section 3.1 discusses the simulation of the expansion of partially condensed clouds. The errors which occur in a typical image analysis are discussed in Section 3.2 based on the model. Section 3.3 show that the model can be used to fit experimental data and extract more accurate temperatures and atom numbers.

3.1 Expanding partially condensed Bose gases

Due to the high density, the BEC atoms will affect the distribution of thermal atoms. This can be modeled by a mean-field approximation, which leads to an extra term in the potential energy. Consequently, the in-trap densities of BEC and thermal atoms become coupled [26, 33]

$$\rho_0 = \frac{\mu - V - 2U_0\rho_{\mathrm{th}}}{U_0}\Theta(\mu - V - 2U_0\rho_{\mathrm{th}}), \tag{3.1}$$

$$\rho_{\mathrm{th}} = \frac{N_{\mathrm{th}}}{(2\pi)^{\frac{3}{2}}\sigma_1\sigma_2\sigma_3}\frac{1}{g_3(\tilde{z})}g_{\frac{3}{2}}\left(\tilde{z}\mathrm{e}^{-(V+2U_0(\rho_{\mathrm{th}}+\rho_0))/k_BT}\right). \tag{3.2}$$

Here, V is now the external trap potential and $U_0 = 4\pi\hbar^2 a/m$ is the interaction energy. The Heaviside step function Θ is used in the condensed distribution. Since the density of thermal atoms is much smaller than that of condensed atoms, the effects on the BEC are neglected, as well as the $2U_0\rho_{\text{th}}$ term in the exponential in the thermal distribution.

The expansion of a released thermal cloud given this initial distribution is found by ignoring the interactions after the trap is turned off, resulting in ballistic motion. In reality, the interactions between the two components will quickly decrease as they expand rather than suddenly disappear.

The time-dependent thermal cloud profile is again given by integrating over the momentum space

$$\begin{aligned}\rho_{\text{th}}(\mathbf{r},t) &= \int \frac{\mathrm{d}^3\mathbf{p}}{(2\pi\hbar)^3}\frac{1}{\mathrm{e}^{(p^2/2m+V_{\text{eff}}(\mathbf{r},\mathbf{p},t))/k_BT}-1},\\ V_{\text{eff}}(\mathbf{r},\mathbf{p},t) &= V(\mathbf{r}-\mathbf{p}/m\cdot t)+2U_0\rho_0(\mathbf{r}-\mathbf{p}/m\cdot t)-\mu\\ &= |\mu - V(\mathbf{r}-\mathbf{p}/m\cdot t)|, \end{aligned} \tag{3.3}$$

but this time, the integral cannot be solved analytically. Here, t is the TOF.

To estimate the errors from ignoring the repulsion from the BEC, the thermal distribution after TOF is calculated with eq. 3.3 and then fitted with the non-interacting model from the previous chapter.

To calculate the distribution after TOF, the physical parameters $N = N_0 + N_{\text{th}}$, T, and t are chosen in a trap with frequencies ω_i. Then, an $n \times n \times n$ grid in space on which to calculate the cloud profile after TOF is defined, called the spatial grid. The range of that grid is defined to be 4 times the width of the thermal cloud calculated with the non-interacting model. Another $m \times m \times m$ grid in space around the in-trap cloud at $t = 0$ is defined, called the momentum grid, with the range again given by 4 times the in-trap non-interacting thermal width.

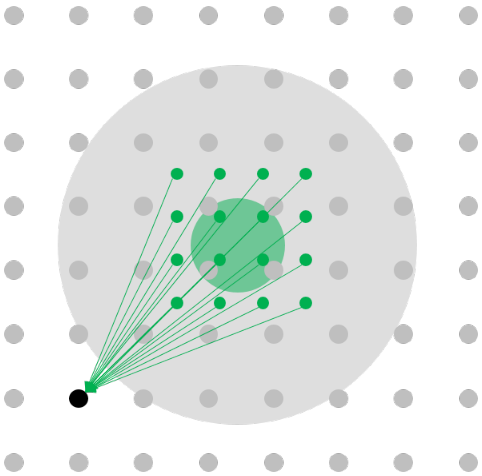

Figure 3.1: Illustration of how to calculate the thermal cloud profile after TOF. The initial cloud, as well as the momentum grid, is illustrated in green (small). The final cloud, as well as the spatial grid, is shown in grey (large). A black point in the spatial grid is highlighted. To calculate the density of atoms at this point, the probability of starting at a specific momentum grid point and with the correct momentum to end up at the black point is calculated for each momentum grid point and then added. This is done for all spatial grid points to get the distribution. The grids and clouds shown here are not representative of those used in the real calculations, but are used only as an illustration of the method. The clouds are not assumed to be spherical, and the calculations are performed in 3 dimensions.

For each point in the spatial grid, the required momentum to have originated from each point in the momentum grid is calculated. Eq. 3.3 is then used to calculate the density at that spatial point, by summing over the momentum grid. This step is illustrated in Figure 3.1. Integration along x_3 is done to obtain column densities.

The grid finesses n and m are found by repeating the calculations with varying values on the most demanding clouds investigated later, with the parameters $N = 10^6$, $T = 0.3T_c$, $t = 3\,\mathrm{ms}$, $\omega_x = 2\pi \cdot 10\,\mathrm{Hz}$,

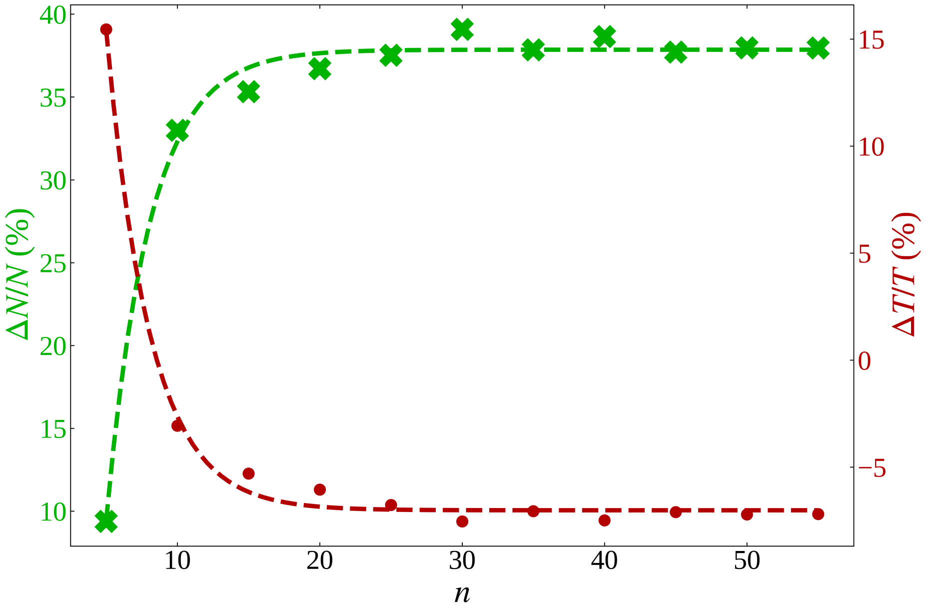

Figure 3.2: Scanning the value of n with $m = 60$ held constant and for the cloud parameters specified in text. Exponential fits are plotted on top. Green ✖-markers show the error in atom number, and red ●-markers show the temperature error.

and $\omega_{y,z} = 2\pi \cdot 100\,\mathrm{Hz}$. The scan for the spatial finesse n is shown in Figure 3.2 with $m = 60$ held constant. The results show clear convergence in both the atom number and temperature estimation errors, and the chosen value was $n = 40$.

Figure 3.3 displays the result of the momentum finesse scan with $n = 40$ held constant. There is again clear convergence in the results, and the chosen value was $m = 60$. Both of these choices can be relaxed if speed is to be optimized.

The complexity of the calculations scales quickly as $\mathcal{O}(n^3 \cdot m^3)$. Therefore, calculations are performed in parallel on a GeForce GTX 1050 Mobile graphics card.

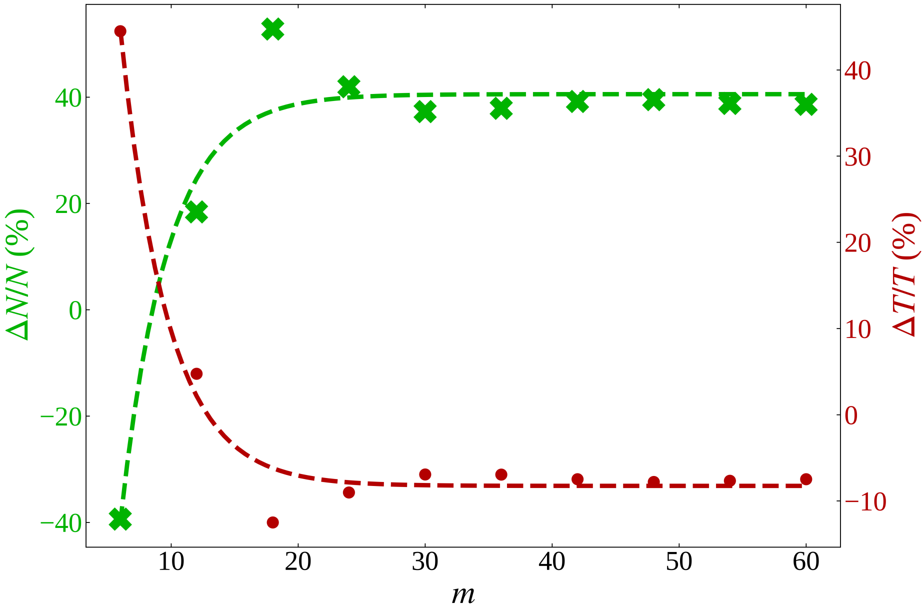

Figure 3.3: Scanning the value of m with $n = 40$ held constant and for the cloud parameters specified in text. Exponential fits are plotted on top. Green ✖-markers show the error in atom number, and red ●-markers show the temperature error.

It is possible to avoid calculating every spatial point by exploiting the symmetry of the system. If the trap geometry is assumed to be at least cylindrically symmetric, meaning that at least two out of the three trap frequencies are equal, it is possible to do the calculation of the cloud profile in 2 dimensions and then exploit the rotational symmetry around the axis with a different trap frequency to get the 3-dimensional distribution. This assumption is used here, and typical cigar-shaped clouds from experiments fall within this assumption.

There are three categories of cloud shapes within this assumption: a cigar-shaped cloud if the two equal frequencies are larger than the last frequency, a spherical cloud if all frequencies are the same, and a

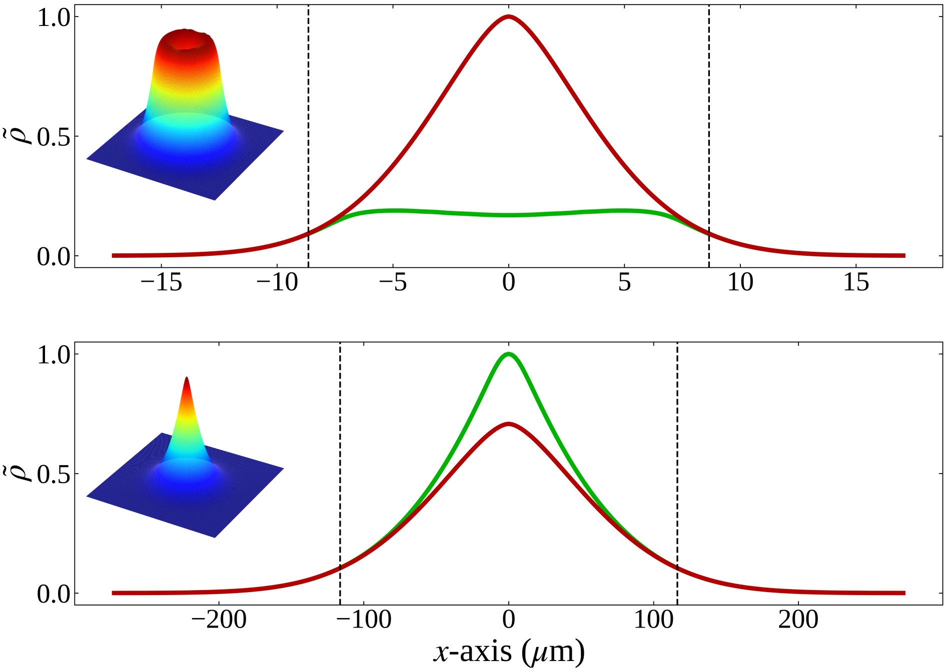

Figure 3.4: Thermal column densities in a symmetrical trap configuration with $N = 200\,\mathrm{k}$, $T/T_c = 0.2$, $\omega_i = 2\pi \cdot 100\,\mathrm{Hz}$, and $N_0/N = 0.97$, for two different TOFs, normalized to the highest column density. The top has $t = 1\,\mathrm{ms}$ and the bottom has $t = 30\,\mathrm{ms}$. A slice through the center of the simulated column densities is shown in green, and a slice through the center of a Bose-enhanced fit to the simulated data, fitted outside the BEC region illustrated by vertical dashed lines, is included in red. Inserts in both figures show the simulated column densities. The BEC is not plotted, and the grids used here are not the same as in the later simulations.

pancake-shaped cloud if the two equal frequencies are lower than the last frequency. It is only necessary to calculate the distribution in the positive quadrant ($x_i \geq 0$, $i = 1, 2, 3$).

The results of such a simulation are shown in Figure 3.4 for a very cold, partially condensed cloud with total atom numbers $N = 200\,\mathrm{k}$, $T/T_c = 0.2$, and $\omega_i = 2\pi \cdot 100\,\mathrm{Hz}$ which results in a condensed fraction of 97% according to the semi-ideal model [29]. The top has $t = 1\,\mathrm{ms}$ and the bottom $t = 30\,\mathrm{ms}$. The figure shows cuts through the center of both simulated and Bose-enhanced fitted thermal column densities normalized to the highest value. The Bose-enhanced fit is performed in a region outside the BEC region illustrated by vertical dashed lines and with the method described in Section 5.6.1.1 which fits the thermal and condensed atom numbers, as well as the temperature in both the horizontal and vertical directions. The temperatures of both directions are averaged to get an estimate of the cloud temperature. The fit is performed twice; first to get the ROIs and then to extract the parameters. The thermal fit is performed outside 1.2 times the BEC radii. The BEC is not plotted, and the x-axis has been rescaled between the two plots.

The repulsive effect of BEC is most clearly seen in the short TOF figure, where a dip in the thermal column density in the center occurs, and while the Bose-enhanced fit to the wings is good, it disagrees greatly in the center and overestimates the temperature by a factor of 7.2% and the atom number by 64.5%. The temperature is fitted to be higher because of the broader wings from the repulsive BEC, and the atom number is fitted much higher because the fit cannot "see" the dip in density in the center. After 30 ms TOF, some of the thermal cloud has fallen into the hole in the middle, and the cloud now has a slightly higher peak column density than the Bose-enhanced distribution. The error in the fit is now -1.3% for temperature and -16.9% for atom number. The sign of the estimation has reversed because the simulated distribution is now more peaked than the Bose-enhanced distribution.

3.2 Analysis of Atom Number and Temperature Errors

In the following, the errors in thermal atom number and temperature determination that occur when the BEC repulsion is ignored are discussed. These errors are obtained by comparing the input parameters of the simulation with the results of a fit similar to the one shown in Figure 3.4. The errors are defined as $\Delta N_{th} = N_{th} - N_{th,s}$ and $\Delta T = T_{th} - T_{th,s}$, where the subscript s indicates the input values of the simulation.

Figure 3.5 shows the error in the thermal atom number determination for various cloud and trap parameters. The relative error $\Delta N_{th}/N$ is shown for for three different trap geometries (cigar shaped trap with $\omega_x = 10\,\mathrm{Hz}$, $\omega_{y,z} = 100\,\mathrm{Hz}$; spherical trap with $\omega_{x,y,z} = 100\,\mathrm{Hz}$; pancake shaped trap $\omega_x = 1000\,\mathrm{Hz}$, $\omega_{y,z} = 100\,\mathrm{Hz}$) and two different TOFs ($t = 3\,\mathrm{ms}$; $t = 30\,\mathrm{ms}$) as function of the total atom number and the reduced temperature. Positive values of the error indicate an overestimation by the Bose-enhanced fit function. The errors are zero at T_c as expected when no BEC is present.

The two columns in the figure illustrate the difference between short (left) and long (right) TOF. A longer TOF generally reduces the magnitude of the errors; however, it does not approach zero. At $t = 30\,\mathrm{ms}$, the errors are at most ±10%. The sign of the errors reverses from short to long TOF because the thermal cloud goes from flattened to peaked compared to a Bose-enhanced fit, as was also seen in Figure 3.4.

The errors are weakly dependent on atom numbers. The effect comes from the fact that the thermal distribution widths are independent of the number of atoms, while the BEC will grow in size with more atoms. Thus, the errors of a large cloud are equal to the errors of a smaller cloud at a colder temperature. Since the size of a BEC scales proportionally to $N_0^{1/5}$, this is a small effect.

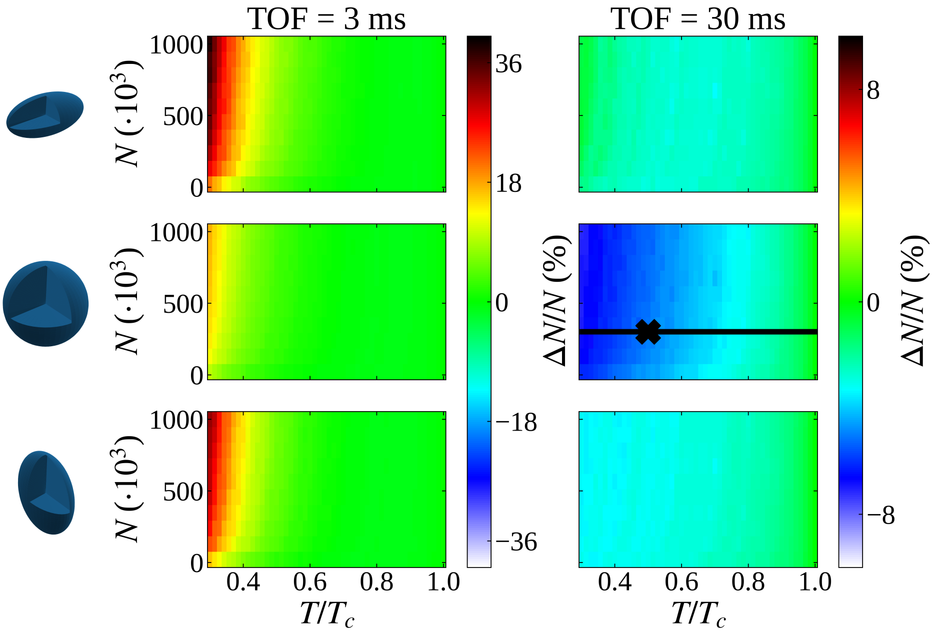

Figure 3.5: Errors in thermal atom number estimation as a function of reduced temperature and total atom numbers for different trap configurations and TOFs. The left column has a TOF of 3 ms and the right column has a TOF of 30 ms. The top row is cigar shaped with $\omega_x = 2\pi \cdot 10\,\mathrm{Hz}$ and $\omega_{y,z} = 2\pi \cdot 100\,\mathrm{Hz}$, the middle row is spherical with $\omega_{x,y,z} = 2\pi \cdot 100\,\mathrm{Hz}$, and the bottom row is pancake shaped with $\omega_x = 2\pi \cdot 1000\,\mathrm{Hz}$, $\omega_{y,z} = 2\pi \cdot 100\,\mathrm{Hz}$.

The temperature is more directly related to the size of the BEC than the number of atoms and will therefore have a much larger effect. Additionally, falling temperatures will decrease the thermal cloud width, making the overlap between the two components even larger.

At long TOF, the errors in a spherical cloud are worse than in both asymmetrical clouds. One explanation could be the following. The volume of a BEC at a given temperature and atom number is relatively fixed at different trap geometries. Therefore, the spatial

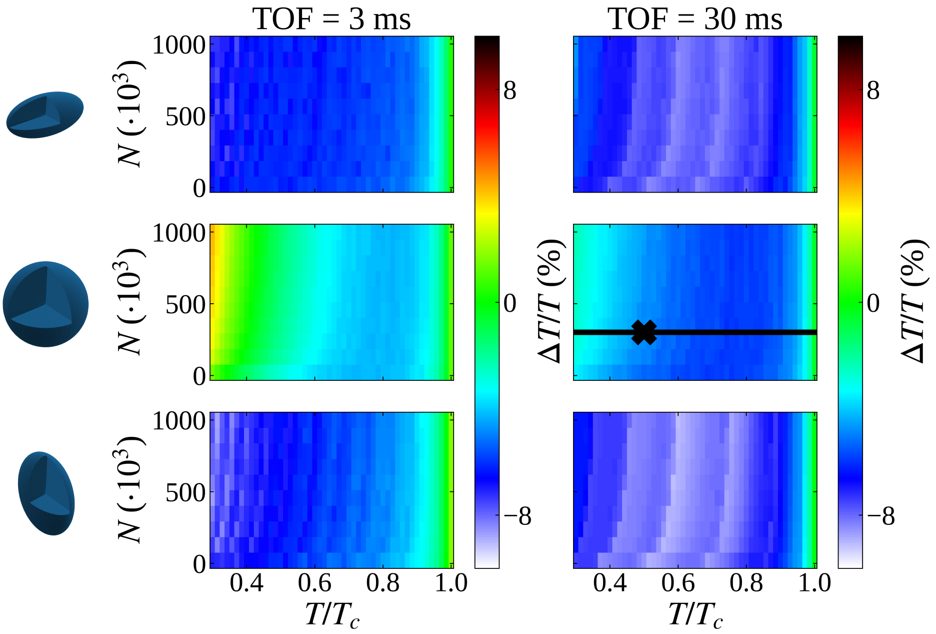

Figure 3.6: Errors in temperature estimation as a function of reduced temperature and total atom numbers for different trap configurations and TOFs. The parameters are the same as in Figure 3.5, specifically: the left column has a TOF of 3 ms and the right column has a TOF of 30 ms. The top row is cigar shaped with $\omega_x = 2\pi \cdot 10\,\text{Hz}$ and $\omega_{y,z} = 2\pi \cdot 100\,\text{Hz}$, the middle row is spherical with $\omega_{x,y,z} = 2\pi \cdot 100\,\text{Hz}$, and the bottom row is pancake shaped with $\omega_x = 2\pi \cdot 1000\,\text{Hz}$, $\omega_{y,z} = 2\pi \cdot 100\,\text{Hz}$

extent in the weak trap directions in asymmetrical clouds must be larger than the spherical cloud radius. A larger radius will expand more slowly, and the thermal fit can be performed in a region closer to the center, allowing it to catch the volume of the cloud more precisely. The opposite happens at short TOF since the asymmetrical BECs are now larger than the spherical ones.

Figure 3.6 shows the errors in the temperature estimation for the

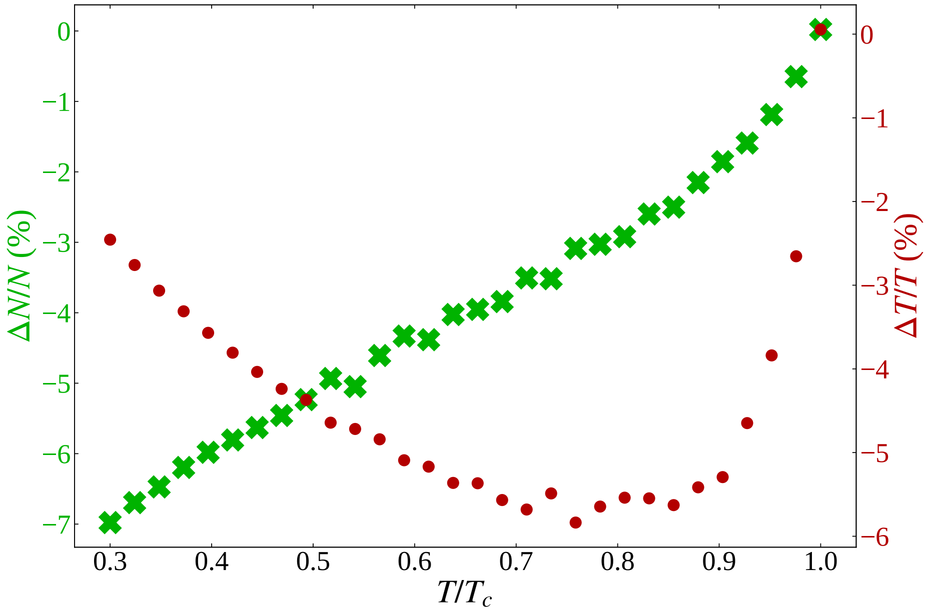

Figure 3.7: Errors in thermal atom numbers and temperature (average of both directions) for different temperatures. The simulations are done in a spherical trap at $N = 3 \cdot 10^5$ and $t = 30\,\text{ms}$. The green ✖-markers show the error in atom numbers, and the red ●-markers show the error in temperature.

same parameters as Figure 3.5. For all parameters simulated here, the temperature errors lie within ±10%.

The temperature errors are weakly dependent on atom number in the same way as atom number errors. The temperature dependence is also generally similar. An exception is the spherical short TOF example, which shows a reversal in the sign of the error at low temperatures. The reason is the same as the reversal in sign explained with Figure 3.4. $t = 3\,\text{ms}$ is enough to reverse the sign of the error for moderately low temperatures, but for the coldest clouds, more TOF is needed before the reversal happens. Interestingly, this means there

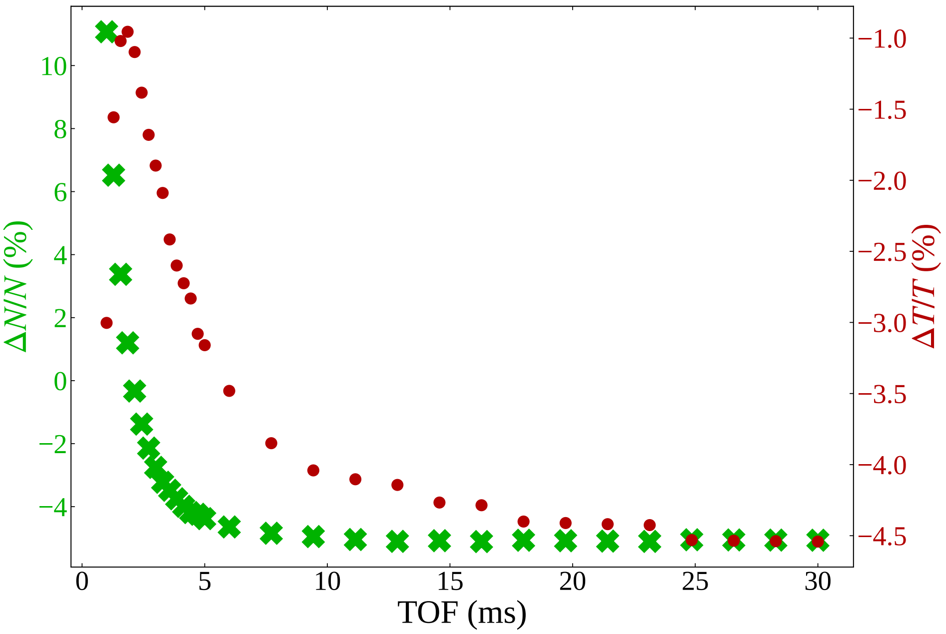

Figure 3.8: Errors in thermal atom numbers and temperature for different TOFs from a spherical trap. The simulations are done at $N = 3 \cdot 10^5$ and $T/T_c = 0.5$. The green ✖-markers show the error in atom numbers and the red ●-markers show the error in temperature.

are regions with no errors. The same effect can also be seen in the spherical long TOF plot, where $t = 30\,\mathrm{ms}$ is not enough to fully reverse the error for the very coldest clouds.

The behavior at different trap geometries is the opposite of that seen for atom number errors. This is because a larger fit region allows the fit function to access more of the central region where the cloud is deformed. While this gives a better estimate of the area under the curve and thus lower atom number errors, it will require the Bose-enhanced function to adapt the temperature more to fit the sharpness of the peak. Spherical clouds are thus the best for estimating cloud

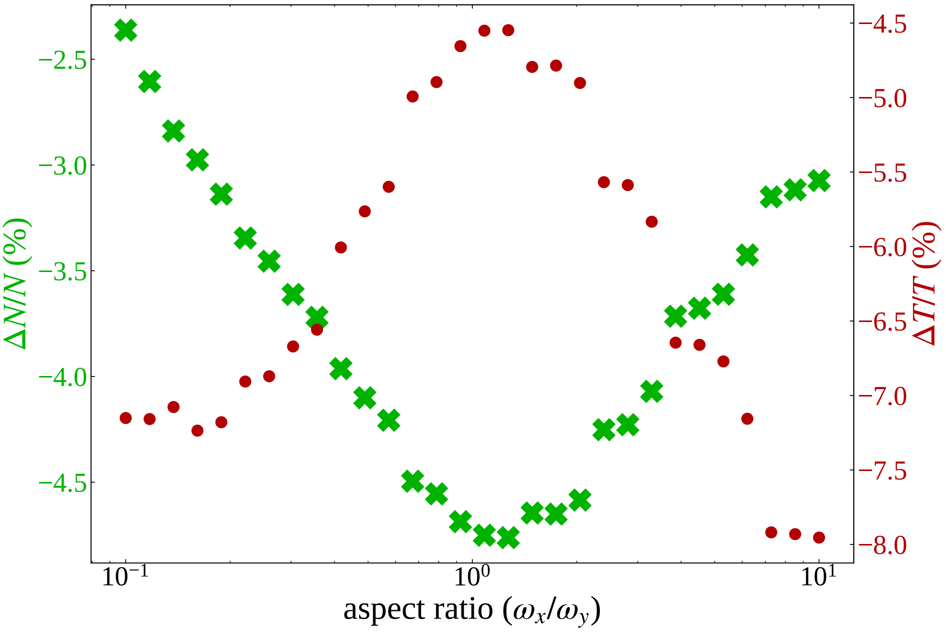

Figure 3.9: Errors in thermal atom numbers and temperature (average of both directions) for different aspect ratios in the trap, but with a constant geometrical mean of the aspect ratios. The simulations are done at $N = 3 \cdot 10^5$, $T/T_c = 0.5$ and $t = 30\,\mathrm{ms}$. The green ✖-markers show the error in atom numbers and the red ●-markers show the error in temperature.

temperatures.

The errors in atom number and temperature were explored further for a selected range of parameters. In the long TOF, spherical shape panels of Figures 3.5 and 3.6 a trajectory at a constant number of $3 \cdot 10^5$ atoms (solid line) is shown. Fig. 3.7 shows the errors in atom number and temperature as a function of temperature along this line. The thermal atom number underestimate grows for low temperatures, reaching $\approx$ 7%. Similarly, the temperature is underestimated, but reaches a maximal underestimation by $\approx$ 6% for $T/T_c = 0.8$. This corresponds to the general behavior of the temperature error.

The dependence of the atom number and temperature errors on the TOF is explored in further detail in Figure 3.8 for the parameters indicated by a cross in the long TOF, spherical shape panels of Figures 3.5 and 3.6. The atom number errors converge faster than the temperature errors. This is in agreement with the behavior seen in the long TOF, spherical shape panel of Figure 3.6, where the temperature error is still reversing for low temperatures while in the same panel of Figure 3.5, the atom number error sign reversal has fully occurred at $t = 30\,\mathrm{ms}$. The errors do not converge to zero for long TOF because the BEC alters the energy levels populated by the thermal atoms, and thus the distribution of momentum. This affects the expansion of the cloud for all TOFs and, consequentially, the estimation of parameters.

Finally, the effect of the trap geometry is investigated from a cigar to a pancake-shaped cloud with 30 ms TOF at otherwise the same parameters as in Figure 3.8. This once again shows how the spherical cloud is best suited for atom number estimation but not for temperature estimation, as was also discussed above.

3.3 Bimodal fits of partially condensed clouds

The simulation of the expansion in TOF can be used within a fitting routine to obtain the thermal atom number and temperature more accurately from experimental data. The fit follows the same steps as explained in Section 5.6.1.2, which fits the thermal and the BEC component separately, once to extract ROIs and a second time to extract cloud parameters. The thermal cloud is fitted outside 1.2 times the BEC radii, and the thermal distribution comes from the simulation instead of the Bose-enhanced function.

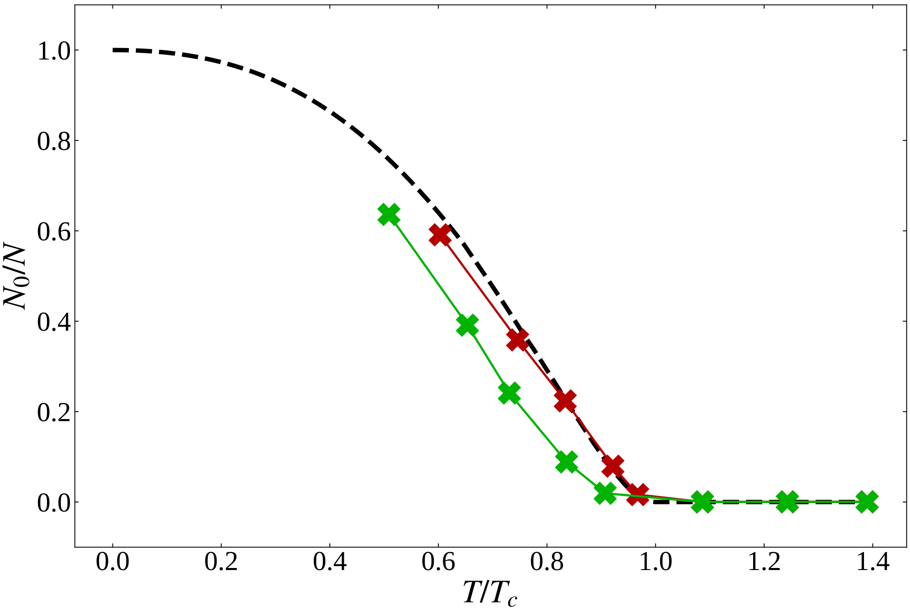

Figure 3.10: BEC fraction as a function of reduced temperature. The results of an analytic fit with the Bose-enhanced distribution are shown as green data points, and the fit with the simulated distribution is shown in red. The black dashed line shows the BEC fraction predicted by the semi-ideal model [29]. Each point is an average of 6 to 8 fits.

The experimental apparatus to produce partially condensed BECs for the dataset presented here was previously described in detail [34] and differs slightly from the current implementation since it used a QUIC trap instead of a crossed dipole trap. Briefly, ^{87}Rb atoms are initially cooled by radio-frequency (RF) evaporation in a Ioffe-Pritchard type magnetic trap. To stabilize the production of ultracold clouds, the cooling process was interrupted at a temperature of $\sim 14\ \mu$K, and the atomic clouds were probed using minimally destructive Faraday imaging [35, 36]. Based on the outcome, the atom number was corrected by removing excess atoms. Subsequently, the magnetic trap was decompressed to obtain radial and axial trapping frequencies of $\omega_\rho = 2\pi \times 93.4\,\mathrm{Hz}$ and $\omega_z = 2\pi \times 17.7\,\mathrm{Hz}$. Finally, BECs were produced by applying an RF sweep that led to the desired average BEC occupation.

The partially condensed clouds were probed using absorption imaging after a 27.5 ms TOF. To accurately detect the density distribution of dense clouds, the imaging system was calibrated following [37, 38] and an intensity $I = 2.3\ I_{\mathrm{sat}}$ was used.

The fit routine for both the simulated clouds and the Bose-enhanced fit minimizes the sum of the squared differences for each pixel and thus obtains the parameters N_0,N_{th}, and T. A fit with the simulated distribution after TOF takes about 10 minutes to complete, and thus the fits were performed on a small dataset with 6 to 8 images at 5 different final RF frequencies. The Bose-enhanced fits were also performed at 3 RF frequencies above condensation. Figure 3.10 shows the results of the simulated (red) and Bose-enhanced (green) fits in a BEC fraction vs. reduced temperature plot and are compared to predictions from the semi-ideal model [29]. The negative temperature error from using the Bose-enhanced fit results in those points lying below the expected curve, while the simulated fit shows better agreement with the semi-ideal model. The simulation ignores the interactions during TOF, which might explain some of the remaining discrepancies.

Chapter 4

Light in high-density media

This chapter describes the phenomenon of light propagating through "dense"[1] atomic clouds. The chapter begins with an introduction to light propagation in a medium in Section 4.1. Section 4.2 deals with the resulting spectra after the light has passed through a full atomic cloud, which is described for both bright field and dark field setups. Finally, Section 4.3 uses the results from the previous sections to discuss the absorption of light in a three-level atom in the presence of a second light field.

4.1 The complex refractive index

This section covers the basic theory of light propagation through a medium and is based on standard textbook knowledge, e.g. [39]. Let the light field have angular frequency ω and wavenumber $k = \omega/c$ with c being the speed of light in vacuum, and let the propagation direction be along $+z$. Assume a continuous incoming wave with no temporal

---

[1] The exact meaning of dense is specified in Section 4.1, and the atoms are, in fact, very dilute.

variation in amplitude. The following calculations are applicable in a sufficiently small region along z such that any diffraction of the wave does not need to be taken into account, analogous to the thin-lens approximation.

This is typically the case for light waves propagating through cold clouds, which are somewhere around $1\,\mu$m to $100\,\mu$m in size. A back-of-the-envelope calculation shows the length scale on which it breaks down by setting the Rayleigh length equal to the waist, giving $\lambda/\pi = 250\,\mathrm{nm}$. While the clouds are not that small, some of the phase modulation patterns of moderately detuned light going through the cloud can be.

The complex E-field is described by

$$E(t,x,y,z) = A(x,y,z)\mathrm{e}^{\mathrm{i}\phi(t,x,y,z)}, \quad \phi = zk - \omega t, \tag{4.1}$$

where A is the complex field amplitude, ϕ is the phase, and the actual E-field is the real part of the complex E-field. The time-averaged intensity of the light wave is $I = \frac{1}{4}c\epsilon_0|A|^2$.

In a medium, the propagation speed of the light wave can be altered to any value[2] v. Now introduce the refractive index $n = c/v$. The equation for the phase of the E-field becomes

$$\mathrm{d}\phi = nk\mathrm{d}z - \omega\mathrm{d}t. \tag{4.2}$$

The amplitude can decrease while the light moves through a medium. This can be described by rewriting the complex amplitude as $\mathrm{d}A = -\kappa A\mathrm{d}z$ where κ describes the amount of absorption caused by the medium. The z-dependence of A can now be moved to the phase in eq. 4.2 by making the refractive index complex

$$\tilde{n} = n + \mathrm{i}\kappa, \quad \mathrm{d}\phi = \tilde{n}k\mathrm{d}z - \omega\mathrm{d}t. \tag{4.3}$$

---

[2] It can even become faster than c; however, this superluminal speed cannot be used to transfer information and is therefore not contradicting relativity.

The propagation of light through a small atomic cloud can be achieved by choosing an appropriate spatial dependence for $\tilde{n}(x, y, z)$. In this model, the only parameter to be calculated is the total complex phase change $\Delta\phi$ throughout the cloud relative to the phase change of vacuum with $\tilde{n} = 1$

$$\Delta\phi(x, y) = \int_{-\infty}^{\infty} (\tilde{n}(x, y, z) - 1) k \mathrm{d}z. \tag{4.4}$$

The value of $\tilde{n}$ in a homogeneous medium of two-level atoms with resonance ω_0 and FWHM linewidth Γ is given by

$$\tilde{n} = \sqrt{1 + \chi}, \tag{4.5}$$

where the polarizability χ is [39]

$$\begin{aligned} \chi &= \frac{\sigma_0 \rho}{k} \frac{-1}{2\tilde{\delta} + \mathrm{i}} \frac{1}{1 + \tilde{I}} & (4.6) \\ &= \frac{\sigma_0 \rho}{k} \frac{-2\tilde{\delta}}{4\tilde{\delta}^2 + 1} \frac{1}{1 + \tilde{I}} + \mathrm{i} \frac{\sigma_0 \rho}{k} \frac{1}{4\tilde{\delta}^2 + 1} \frac{1}{1 + \tilde{I}}, & (4.7) \end{aligned}$$

with $\tilde{\delta} = (\omega - \omega_0)/\Gamma$ being the normalized detuning, $\tilde{I} = I/I_{\mathrm{sat}}$ the normalized intensity to the detuning dependent saturation intensity $I_{\mathrm{sat}} = \hbar\omega\Gamma/(2\sigma_0)\left(1 + 4\tilde{\delta}^2\right)$, ρ the number of atoms per volume, and $\sigma_0 = 6\pi/k^2$ the resonant cross section[3]. It is assumed that $|\omega - \omega_0| \ll \omega$ for the RWA.

Because an atomic cloud is not a continuous medium, the electric field experienced by each atom will be modified. The Lorentz-Lorenz model [40, 41] approximates the situation by placing each atom in a sphere of vacuum surrounded by a polarized medium, which leads to an enhancement of the electric field, and the resulting refractive index

$$\tilde{n} = \sqrt{1 + \frac{\chi}{1 - \frac{1}{3}\chi}}. \tag{4.8}$$

---

[3] In a multi-level system the expression for the cross section gets more complicated. See Section 5.1.1 for details in ^{87}Rb.

This approximation ignores the bosonic statistics of the atoms, which can lead to a myriad of other effects (see [42]) due to the bunching nature of bosons. It can be accounted for by extra correction terms to the index of refraction [43], but is not included in our models.

It is often the case in cold atom experiments that the product $\sigma_0\rho/k$ is much smaller than one, such that eq. 4.8 can be expanded to first order

$$\tilde{n} \approx 1 + \frac{1}{2}\chi \tag{4.9}$$

The density will be called ”high” whenever this approximation does not hold for all detunings, although that happens at a very low density compared to most everyday objects or even air at atmospheric pressure.

At low density, the phase shift caused by the real part of $\tilde{n}$ is zero on resonance. But at high density, when the square root in the refractive index cannot be approximated to first order, the real and imaginary parts of χ start to mix, and there is a non-zero phase shift even on resonance.

4.1.1 Absorption

It is necessary to calculate the time average over one period to go from complex E-fields to intensities

$$I = \frac{c\epsilon_0}{2}\int_0^{2\pi/\omega} (\mathrm{Re}(E))^2 \, \mathrm{d}t \tag{4.10}$$

Assume the light has acquired a complex phase change $\Delta\phi$ after passing through a cloud. Inserting this into the expression for the E-field in eq. 4.1 gives the intensity I_{out} after having passed through that medium

$$I_{\text{out}} = \frac{c\epsilon_0}{2}|A|^2 \int_0^{2\pi/\omega} \left(\operatorname{Re}\left(\mathrm{e}^{\mathrm{i}(\Delta\phi-\omega t)}\right)\right)^2 \mathrm{d}t \tag{4.11}$$

$$= 2I_{\text{in}} \int_0^{2\pi/\omega} \mathrm{e}^{-2\operatorname{Im}(\Delta\phi)} \cos^2\left(\operatorname{Re}(\Delta\phi) - \omega t\right) \mathrm{d}t, \tag{4.12}$$

The value $\frac{c\epsilon_0}{2}|A|^2$ is twice the incoming intensity I_{in}. $\operatorname{Im}(\Delta\phi)$ is only dependent on space, not time, and can thus be taken outside the integral. $\operatorname{Re}(\Delta\phi)$ is time-independent, and the integral over a full period of $\cos^2$ will therefore give $\frac{1}{2}$

$$I_{\text{out}} = I_{\text{in}}\mathrm{e}^{-2\operatorname{Im}(\Delta\phi)}. \tag{4.13}$$

This is the general expression for the absorption of light through a thin medium, but $\Delta\phi$ has to be calculated numerically for arbitrary clouds.

4.1.1.1 Normalized transmission

In an experiment, it is necessary to block the background signal around the atom cloud with a small pinhole of area[4] A to get a better signal-to-noise ratio. The power making it through this pinhole without the presence of atoms is $P_{\text{in}} = I_{\text{in}}A$.

With atoms present, assume the intensity distribution of light just after the cloud is $I(x, y)$. The power P is then given by

$$P = \int_{\text{pinhole}} I \mathrm{d}A \tag{4.14}$$

---

[4] Since it is not possible to place a pinhole just after the atoms, it is placed in the intermediate image, and the area must be scaled according to any magnification taking place in the optical setup.

From this, a very useful quantity called the normalized transmission can be defined

$$\mathcal{T} = \frac{P}{P_{\text{in}}} \tag{4.15}$$

Many results in the thesis is presented in the form of the normalized transmission.

4.1.1.2 Low density resonant absorption

On resonance and in the low density approximation from eq. 4.9, there is a simple relation between the absorbed light and the number of atoms. These conditions are typically applicable to imaging of a cloud after TOF expansion, but the density of a BEC in-trap is too high for the low-density approximation to hold.

Let $I(z)$ denote the intensity along the z-direction and let $\tilde{\delta} = 0$. Then

$$I(z) = I_{\text{in}} \mathrm{e}^{-2\,\mathrm{Im}(\Delta\phi(z))}, \tag{4.16}$$

with

$$\begin{aligned} 2\,\mathrm{Im}(\Delta\phi(z)) &= 2k \int_{-\infty}^{z} \mathrm{Im}(\tilde{n}) \mathrm{d}z & (4.17)\\ &\approx k \int_{-\infty}^{z} \mathrm{Im}(\chi) \mathrm{d}z & (4.18)\\ &= \sigma_0 \int_{-\infty}^{z} \rho \frac{1}{1+\tilde{I}} \mathrm{d}z. & (4.19) \end{aligned}$$

Differentiating both sides of eq. 4.16 with respect to z gives

$$\begin{aligned} \frac{\mathrm{d}I}{\mathrm{d}z} &= I_{\text{in}} \mathrm{e}^{-2\,\mathrm{Im}(\Delta\phi)} \cdot \frac{\mathrm{d}}{\mathrm{d}z} \left(-2\,\mathrm{Im}(\Delta\phi)\right) & (4.20)\\ &= -I\sigma_0\rho \frac{1}{1+\tilde{I}} & (4.21) \end{aligned}$$

This differential equation can be solved by gathering all terms dependent on I on the left side and integrating

$$-\int_{I_{\mathrm{in}}}^{I_{\mathrm{out}}} \frac{1+\frac{I}{I_{\mathrm{sat}}}}{I} \mathrm{d}I = \sigma_0 \int_{-\infty}^{\infty} \rho \mathrm{d}z. \tag{4.22}$$

The integrals give

$$\ln\left(\frac{I_{\mathrm{in}}}{I_{\mathrm{out}}}\right) + \frac{I_{\mathrm{in}} - I_{\mathrm{out}}}{I_{\mathrm{sat}}} = \sigma_0 \tilde{\rho}, \tag{4.23}$$

where the integral of the atom density along z is called the column density of the atoms and is denoted with $\tilde{\rho}$. The product $\sigma_0\tilde{\rho}$ is called the optical density or *od*. This relation between absorbed light and *od* is called the Beer-Lambert law. In the limit of $I \ll I_{\mathrm{sat}}$, the law reduces to a simple exponential decay of intensity.

For reference, the column densities of a Bose-enhanced (BE) distribution, a Gaussian (G) distribution (for higher temperatures), and a TF-distribution (see Appendix A in [30]) are

$$\tilde{\rho}_{\mathrm{BE}} = \frac{N_{\mathrm{th}}}{2\pi w_x w_y g_3(\tilde{z})} g_2\left(\tilde{z}\mathrm{e}^{-x^2/(2w_x^2)-y^2/(2w_y^2)}\right) \tag{4.24}$$

$$\tilde{\rho}_{\mathrm{G}} = \frac{N_{\mathrm{th}}}{2\pi w_x w_y} \mathrm{e}^{-x^2/(2w_x^2)-y^2/(2w_y^2)} \tag{4.25}$$

$$\tilde{\rho}_{\mathrm{TF}} = \frac{5}{2\pi}\frac{N_0}{R_x R_y} \max\left(0, \left(1-\frac{x^2}{R_x^2}-\frac{y^2}{R_y^2}\right)^{3/2}\right), \tag{4.26}$$

where the imaging direction is taken to be z.

The BE column density above is normalized to a specific atom number, while the only true parameters are the temperature and the chemical potential. If those are to be fitted instead, the following BE column density can be used

$$\tilde{\rho}_{\mathrm{BE}} = \frac{\sqrt{2\pi} w_z(0)}{\lambda_T^3 \beta_x \beta_y} g_2\left(\tilde{z}\mathrm{e}^{-x^2/(2w_x^2)-y^2/(2w_y^2)}\right), \tag{4.27}$$

where $w_z(0) = \sqrt{\frac{k_B T}{m\omega_z^2}}$ is the in-trap width in the imaging direction z, $\beta_i = \sqrt{1 + \omega_i^2 t^2}$ are the expansion factors of the cloud after TOF in each direction, $\lambda_T = \sqrt{\frac{2\pi\hbar^2}{mk_B T}}$ is the thermal de Broglie wavelength, and $w_{x,y} = \sqrt{\frac{k_B T}{m\omega_{x,y}^2}}\beta_{x,y}$ are the cloud widths are TOF in the x, y-directions.

4.2 High density spectra

This section investigates the spectra that result from high-density clouds. From now on, the assumption will always be that $I \ll I_{\text{sat}}$ since this analysis will be applied to single photon measurements. The analysis is split into two parts, Section 4.2.1 for bright field and Section 4.2.2 for dark field. Bright field and dark field are two complementary ways of extracting information from the light field that has passed through a sample. Where bright field measures the transmitted light field directly, dark field applies some sort of special preparation of the incoming light field or post-processing of the transmitted light field to remove the part of the light field that has not interacted with the sample.

There are countless ways to apply both the bright field and the dark field techniques, and each method requires its own analysis and has its own advantages and disadvantages. On top of standard bright field imaging, novel bright field techniques include, but are not limited to, frequency modulation imaging [44], dual-port Faraday imaging [45], bright field partial-transfer absorption imaging [46], phase contrast imaging with a phase spot [41], off-resonant defocus-contrast imaging [47], and inverse holography [48]. Similarly, there are countless different dark field techniques, such as dark field scalar imaging, which have been used on resonance [49, 50] as well as far-off resonance to achieve phase contrast imaging [51–54], spatial heterodyne dark field imaging [55], and dark field Faraday imaging [35, 36, 45]. In principle, many of the bright field techniques could be used in a dark field setup, and similarly, the dark field methods could be used in bright field. All the dark field methods listed here use some kind of post-processing of the light, but it is also possible to shape the incoming light beam such that the resulting image becomes dark ground. I am not aware of any such techniques used in the field of cold atoms, but it is possibly the most common technique used in other fields such as nano-science, chemistry, and microbiology (see, for example [56]).

All the imaging techniques referenced here use either on-resonant or far off-resonant probe light. The following sections will bridge the gap between the two methods by focusing on the entire spectrum.

4.2.1 Bright field

In bright field, the transmitted light is imaged directly onto a camera or a detector, and the signal is given by eq. 4.13 with the spatial distribution of atoms described in Chapter 2. The normalized transmission is

$$\mathcal{T} = \frac{1}{A}\int_{\text{pinhole}} \mathrm{e}^{-2\,\mathrm{Im}(\Delta\phi)}\mathrm{d}A. \tag{4.28}$$

Thermal clouds have two degrees of freedom: the number of atoms and the temperature. Figure 4.1 illustrates in a range of scenarios how the transmission of the light wave looks. An effective in-trap pinhole radius of $75\,\mu$m is used and trap frequencies in all directions are $2\pi \cdot 100\,$Hz. The top of the figure shows the intensity distribution from eq. 4.13, the integrand of eq. 4.28, through a pinhole with effective radius just in front of the cloud equal to $75\,\mu$m for varying atom numbers and temperatures. The bottom part of the figure shows the normalized transmission from eq. 4.28 for the same atom numbers and temperatures plotted above.

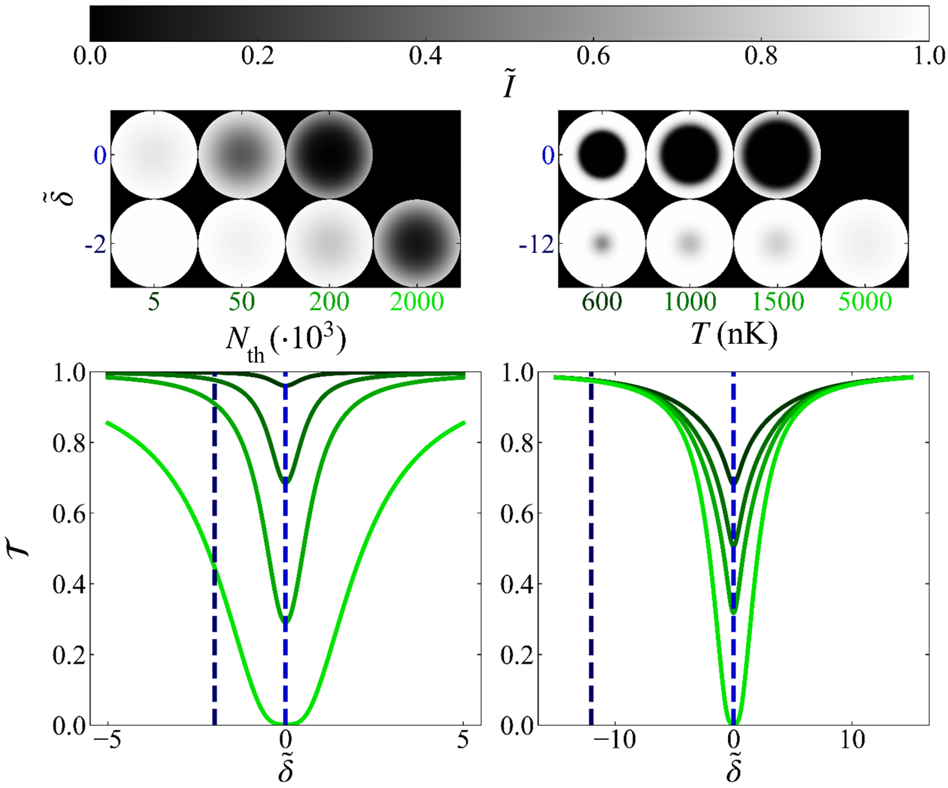

Figure 4.1: Examples of the bright field signal through different thermal clouds. Top left: the light field intensity after passing a cloud with temperature $T = 5000\,\mathrm{nK}$ for different numbers of thermal atoms and with $\tilde{\delta} = 0$ and $\tilde{\delta} = -2$. Top right: the light field intensity after passing a cloud with $N = 2000\,\mathrm{k}$ for different temperatures and with $\tilde{\delta} = 0$ and $\tilde{\delta} = -12$. Bottom left: the normalized intensity spectra corresponding to the plots above. The atomic number increases from black to green, corresponding to the same colors in the x-axis labels above. Bottom right: the normalized intensity spectra corresponding to the plots above. The temperature increases from black to green, corresponding to the same colors in the x-axis labels above.

On the top left, it can be seen that the center of the cloud quickly becomes totally opaque at resonance. Any more atoms added to this central region will not affect the resonant part of the spectrum. In the extreme case, the opaque region of the cloud is larger than the pinhole itself and blocks all light, which results in a flattened spectral signal around resonance.

At detunings close to resonance, the normalized transmission decreases with the number of atoms in the cloud in a complicated nonlinear fashion because of this saturation. At far-detuning, where not even the cloud center is able to block all the incoming light, the absorption approaches a linear response to atom numbers. That is, the exponent in eq. 4.28 becomes much less than 1 everywhere. The in-trap thermal cloud density is low enough to use the approximations from Section 4.1.1.2, that $\mathrm{Im}(\tilde{n}) = \frac{1}{2}\,\mathrm{Im}(\chi)$. For far detuning, the intensity through the thermal cloud is

$$
\begin{aligned}
I_{\text{out}} &\approx I_{\text{in}} \left(1 - \sigma_0 \int_{-\infty}^{\infty} \rho \frac{1}{4\tilde{\delta}^2 + 1} \mathrm{d}z \right) && (4.29) \\
&= I_{\text{in}} \left(1 - \sigma_0 \tilde{\rho} \frac{1}{4\tilde{\delta}^2 + 1} \right), && (4.30)
\end{aligned}
$$

where $\tilde{I} = 0$ has been used as an approximation for low intensity light fields. The normalized transmission of a cloud with N atoms, and assuming the entire cloud is within the pinhole, is then

$$
\begin{aligned}
\mathcal{T} &= \frac{1}{P_{\text{in}}} \int_{\text{pinhole}} I_{\text{in}} \left(1 - \sigma_0 \tilde{\rho} \frac{1}{4\tilde{\delta}^2 + 1} \right) \mathrm{d}A && (4.31) \\
&= \frac{1}{P_{\text{in}}} \left(P_{\text{in}} - I_{\text{in}} \sigma_0 N \frac{1}{4\tilde{\delta}^2 + 1} \right) && (4.32) \\
&= 1 - \frac{I_{\text{in}} \sigma_0 N}{P_{\text{in}}} \frac{1}{4\tilde{\delta}^2 + 1} && (4.33) \\
&= 1 - \frac{I_{\text{in}} \sigma(\tilde{\delta}) N}{P_{\text{in}}}, \qquad \sigma(\tilde{\delta}) = \sigma_0 \frac{1}{4\tilde{\delta}^2 + 1}. && (4.34)
\end{aligned}
$$

The result can be understood as follows. In this regime, the probability of a photon seeing an atom in the cloud is small, so the probability of a photon getting absorbed is given by the probability of getting absorbed by one atom ($I_{\mathrm{in}}\sigma(\tilde{\delta})/P_{\mathrm{in}}$) times the total number of atoms in the cloud. The consequence is that for any cloud entirely contained within the pinhole and with the same number of atoms but different temperatures or trap frequencies, the spectra will always converge to the same values at far enough detuning.

This can be seen in the right column in Figure 4.1. At $\tilde{\delta} = -12$, the three lowest spectra have already converged, but the highest spectrum is still slightly different. Looking at the corresponding intensity at $\tilde{\delta} = -12$ and $T = 600\,\mathrm{nK}$, the center of this cloud is still somewhat saturated, which means the approximations break down for this cloud specifically. For the three other clouds, the intensity distributions vary greatly, but the spectra are the same.

Only in the regime where all of the approximations above hold does the spectrum take the shape of a Lorentzian function for all detunings. This is typically the case after TOF, but rarely the case in-trap.

From the considerations above, it can be seen that the spectra around resonance are sensitive to temperature because of the consequential variation in the size of the cloud interacting with the incoming light. In the bottom right of Figure 4.1 the spectra show a sharp response to temperature changes close to resonance and a minimal response far off-resonant. The spectra are also sensitive to the number of atoms for obvious reasons, but at far detuning, the spectra become sensitive only to the number of atoms. Because of this variation in relative sensitivity to atom number and temperature, the full spectrum can be fitted to obtain both values independently.

Bimodal clouds also have two degrees of freedom, the chemical potential directly related to the number of BEC atoms N_0 through the TF-approximation and the temperature, which together with the chemical potential gives the number of thermal atoms. Figure 4.2

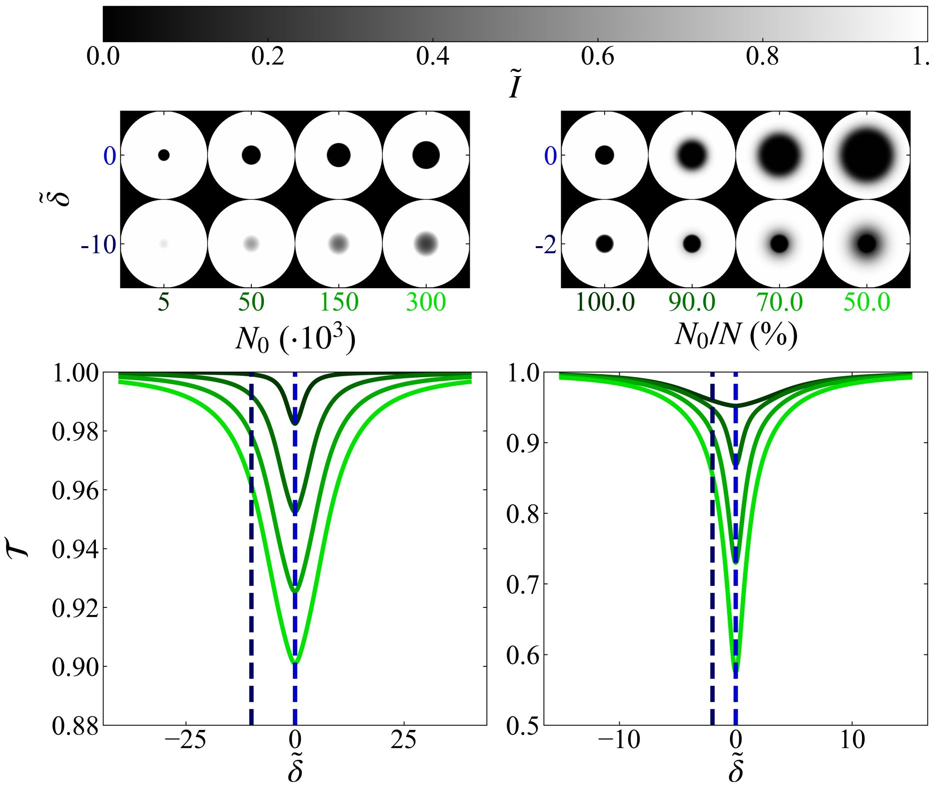

Figure 4.2: Examples of the bright field signal through different bimodal or purely condensed clouds. Top left: the light field intensity after passing a pure condensate cloud for different numbers of atoms and with $\tilde{\delta} = 0$ and $\tilde{\delta} = -10$. Top right: the light field intensity after passing a bimodal cloud with $N_0 = 50\,\mathrm{k}$ for different condensate fractions and with $\tilde{\delta} = 0$ and $\tilde{\delta} = -2$, indicated with vertical dashed lines below in different shades of blue. Bottom left: the normalized intensity spectra corresponding to the plots above. The condensed atom number increases from black to green, corresponding to the same colors in the x-axis labels above. Bottom right: the normalized intensity spectra corresponding to the plots above. The condensed fraction decreases from black to green, corresponding to the same colors in the x-axis labels above.

shows a similar result as Figure 4.1 does for pure thermal clouds, but now for bimodal clouds. The effective in-trap pinhole radius is $25\,\mu$m and the trap frequencies in all directions are $2\pi \cdot 100\,$Hz.

The left side of the figure shows different pure BEC clouds. Starting at the top, it can be seen that the number of atoms now alters the size of the cloud. This does not happen for thermal distributions. On resonance, the *od* of the BEC is enormous (several hundreds) and essentially the entire cloud is saturated, no matter how few BEC atoms there are in the cloud. This also leads to a much broader spectrum, and the clouds are still clearly visible (and even saturated for the two larger clouds) at $\tilde{\delta} = -10$.

On the right side, the number of BEC atoms is set to $N_0 = 50\,$k and the condensed fraction (or the temperature) is varied from 1 to 0.5. At $\tilde{\delta} = -2$, the BEC part is essentially unchanged, but the thermal atom absorption has decreased significantly. At even higher detunings, only the BEC would be visible.

The result of the vastly different spatial distributions of a BEC and a thermal cloud is a bimodal spectrum. Contrary to the spatial bimodal distribution with a narrow but high BEC peak in the center surrounded by a broad and low thermal tail, the spectral bimodal shape is turned around. It is now the BEC which is broad because of the high *od*, but shallow because of the small extent spatially, meaning a smaller fraction of the incoming light hits the BEC and gets absorbed. The thermal cloud is large, and at resonance absorbed much more light than the BEC; however, the low *od* means the large dip quickly disappears at relatively small detunings.

The bimodal shape of the cloud means that it is possible to fit all relevant cloud parameters in-trap from a spectrum instead of an image. Additionally, it is much easier to compensate for low resolution in the optical setup, because only the total power of the signal has to be adjusted.

Perhaps one of the largest benefits of this method is the enhanced sensitivity to thermal atom numbers at very high BEC fractions. Standard TOF imaging techniques typically start to struggle with separating the two components at a BEC fraction around 70%. But looking at the bottom right panel in Figure 4.2, it can be seen that the spectrum is highly sensitive to the presence of thermal atoms well into the 90's. The spectral method also does not require the release of atoms, and with a sensitive single-photon detector, the measurements can therefore be performed non-destructively.

Ironically, the main drawback must be the low sensitivity to the number of BEC atoms, since the amount of absorption is so small. Luckily, this can be fixed by using dark field techniques in combination with the spectrum acquisition, or by combining spectral and TOF imaging data.

4.2.2 Dark field

This section describes the light intensities and spectra seen in a scalar dark field setup similar to those used in [49–53]. The results will differ from other dark field setups. Interestingly, however, none of the methods seem to be vastly superior or inferior to the others [35] except perhaps in very specific cases. The formula for dark field signals is derived in [57]. It is summarized again here.

Assume the electric field of the incoming light is linearly polarized and that the atoms form a 2-level system. Later, it is shown that the same equations hold up for circularly polarized light. The probe field would have the wavefront electric field just after the position of the cloud $E_{\text{in}}(x, y, t)$ if no interactions had taken place. Denote the actual electric field after the cloud as $E(x, y, t)$, which can be described as the untouched field E_{in} plus an extra diffraction field emitted by the atoms[5] E_{diff}

$$E = E_{\text{in}} + E_{\text{diff}}. \tag{4.35}$$

A dark field setup can sort away only E_{in} leaving behind $E_{\text{diff}} = E - E_{\text{in}}$ going to the detector or the camera (see Figure 4.3 for a simple example of how it can be achieved), leaving behind the time-averaged diffracted intensity

$$I_{\text{diff}} = I_{\text{in}} \left| \mathrm{e}^{\mathrm{i}\Delta\phi} - 1 \right|^2, \tag{4.36}$$

where I_{in} is the time-averaged intensity of the incoming field E_{in}. Splitting $\Delta\phi$ into the real and imaginary components gives

$$\begin{aligned} I_{\text{diff}} &= I_{\text{in}} \left| \mathrm{e}^{-\operatorname{Im}(\Delta\phi)} \Big(\cos\Big(\operatorname{Re}(\Delta\phi)\Big) + \mathrm{i}\sin\Big(\operatorname{Re}(\Delta\phi)\Big)\Big) - 1 \right|^2 & (4.37) \\ &= I_{\text{in}} \left(\mathrm{e}^{-2\operatorname{Im}(\Delta\phi)} + 1 - 2\cos\Big(\operatorname{Re}(\Delta\phi)\Big)\right). & (4.38) \end{aligned}$$

Here, $\exp\Big(-2\operatorname{Im}(\Delta\phi)\Big)$ is the transmission of the intensity (see eq. 4.13) and is denoted T. Then

$$I_{\text{diff}} = I_{\text{in}} \Big(T + 1 - 2\sqrt{T} \cos\Big(\operatorname{Re}(\Delta\phi)\Big)\Big), \tag{4.39}$$

and at resonance where there is no phase shift ($\operatorname{Re}(\Delta\phi) = 0$)

$$I_{\text{diff}} = I_{\text{in}} \left(1 - \sqrt{T}\right)^2. \tag{4.40}$$

The power P is the integral over the pinhole with area A

$$P_{\text{diff}} = \int_{\text{pinhole}} I_{\text{diff}} \mathrm{d}A, \tag{4.41}$$

---

[5] This is more than just a smart trick. Maxwell's equations contain only source terms and no "absorption" terms. The only way to absorb light is to emit new light out of phase, which is exactly what happens when the atomic dipoles are set into oscillation by the incoming light.

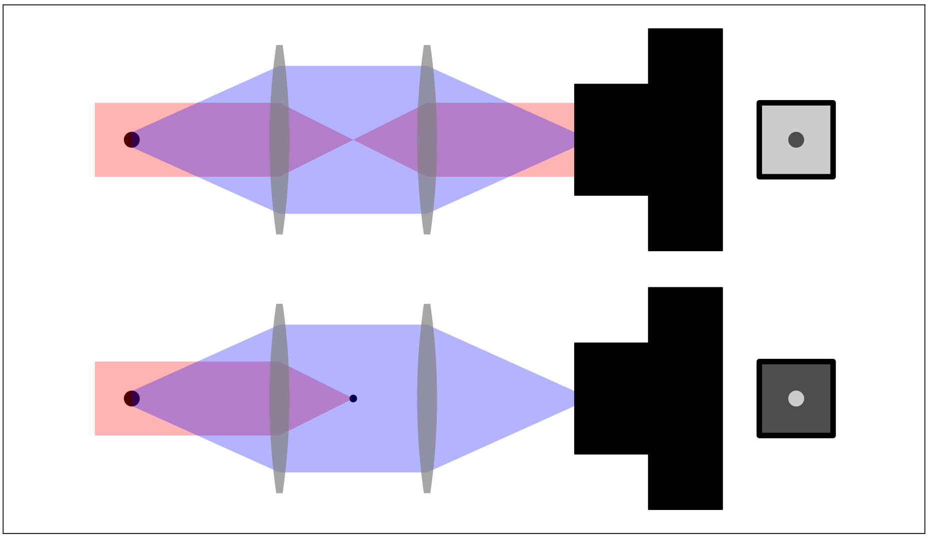

Figure 4.3: A bright field setup compared to a simple scalar dark field setup. A large waist probe field (red) is sent onto a small target. The diffracted electric field E_{diff} (blue) expands much faster than the probe field because of the smaller waist. A focal length away, a lens collimates the diffracted field and focuses the probe field. Top: A second lens focuses the diffracted field onto a camera (or a detector) through a pinhole (not shown). The diffracted field interacts with the probe field and results in a dark spot. Bottom: a focal length away from the first lens, the now small probe field hits a small obstruction target while the large diffracted field passes by. The diffracted field is focused onto a camera (or a detector) through a pinhole (not shown) by a second lens, resulting in a bright spot.

and the normalized transmission (eq. 4.15), where $P_{\text{in}} = I_{\text{in}}A$ is incoming power, is

$$\mathcal{T}_{\text{diff}} = \frac{1}{A}\int_{\text{pinhole}} T + 1 - 2\sqrt{T}\cos\Big(\text{Re}(\Delta\phi)\Big)\text{d}A, \tag{4.42}$$

and it represents the transmission through both the cloud and past the dark field target.

Now consider circularly polarized light. Eq. 4.36 gives the time-averaged diffracted intensity in terms of the incoming time-averaged intensity and the norm of the complex electric field difference. The relation between field amplitude and time-averaged intensity is different for circularly polarized light and leads to a factor of two for circularly polarized light compared to linear polarization. However, this factor of two is absorbed into I_{in} and the equation remains the same. This is not surprising since circularly polarized light is a superposition of two perpendicular light fields, and all the results are stated relative to the incoming intensity.

The equations do not include the Faraday rotation, however, and this can affect linearly and circularly polarized light differently. Faraday rotation happens if the light field couples to more than one excited state, such as linearly polarized light in a magnetic field along the propagation direction or circularly polarized light in a magnetic field perpendicular to the propagation direction.

It is worth considering the result in eq. 4.42 in a few simple cases. When there are no atoms, and bright field would give a full signal, dark field gives zero signal since $T = 1$ and $\mathrm{Re}(\Delta\phi) = 0$. The opposite case is when the atoms block all signals and $T = 0$, resulting in $\mathcal{T}_{\mathrm{diff}} = 1$. From this, it can be seen that dark field and bright field are reversed, although the relation between the two is not linear. An immediate benefit of this is that for small $\mathcal{T}_{\mathrm{diff}}$, the signal-to-noise ratio is significantly improved compared to bright field. In practice, $\mathcal{T}_{\mathrm{diff}}$ will always be small, unless the pinhole is approximately the size of the cloud or smaller, and this is rarely a good idea since the edge of the pinhole can induce complicated focus effects, and since such a setup is very dependent on the exact position of the center of the cloud relative to the pinhole.

The more surprising scenario is when $T = 1$ but $\mathrm{Re}(\Delta\phi) \neq 0$. In this case, bright field also gives full power and therefore no sign of atoms because it cannot detect phase shifts. Dark field, however, can detect phase shifts, even though there is no interference occurring on the detector or camera. This is because the phase shift is spatially dependent and works akin to a lens. If the phase shift was global, an interference setup would indeed be necessary. The normalized transmission becomes

$$\mathcal{T}_{\mathrm{diff},T=1} = \frac{2}{A}\int_{\mathrm{pinhole}} 1 - \cos\Big(\mathrm{Re}(\Delta\phi)\Big)\mathrm{d}A, \tag{4.43}$$

which varies between 0 signal when there is no phase shift, and

$$\mathcal{T}_{\mathrm{diff},T=1} = 4, \tag{4.44}$$

when the phase shift is exactly π everywhere. In this case, $E_{\mathrm{diff}} = -2E_{\mathrm{in}}$ and the signal is four times the incoming power.

How is it possible that four times the incoming power is hitting the detector? The explanation lies in how the E_{in} field is removed from the setup. Consider again Figure 4.3. The setup works by exploiting the different waists of the two fields E_{diff} and E_{in} to block only the latter. It is therefore necessary that E_{in} is much larger than the target. The "missing" energy comes from the part of E_{in} passing around the cloud and diffracting around the obstruction target after the first lens. An equivalent explanation is that just before the obstruction target the diffracted field interacts destructively with the focused original field. The change in electric field amplitude is small, but since the focused electric field has a huge amplitude and intensity scales as field amplitude squared, the change in intensity is large. This missing absorption is what goes to the detector.

An important case to study is the far-detuned region where $|\chi| \ll 1$ and $|\Delta\phi| \ll 1$. Then

$$\chi \approx \frac{\sigma_0\rho}{k}\left(-\frac{1}{2\tilde{\delta}} + \mathrm{i}\frac{1}{4\tilde{\delta}^2}\right) \tag{4.45}$$

$$\tilde{n} \approx 1 + \frac{1}{2}\chi. \tag{4.46}$$

The complex phase change in this region is

$$\Delta\phi \approx od\left(-\frac{1}{4\tilde{\delta}} + \mathrm{i}\frac{1}{8\tilde{\delta}^2}\right), \tag{4.47}$$

where $od = \sigma_0\tilde{\rho}$ is the resonant optical density. The transmission is then

$$T = \mathrm{e}^{-2\,\mathrm{Im}(\Delta\phi)} \approx 1 - od\frac{1}{4\tilde{\delta}^2} \tag{4.48}$$

$$\sqrt{T} \approx 1 - od\frac{1}{8\tilde{\delta}^2}. \tag{4.49}$$

The last term in the expression for $\mathcal{T}_{\text{diff}}$ (eq. 4.42) is $\cos\Big(\mathrm{Re}(\Delta\phi)\Big)$ which can be approximated to

$$\cos\Big(\mathrm{Re}(\Delta\phi)\Big) \approx 1 - \frac{1}{2}\mathrm{Re}(\Delta\phi)^2 \approx 1 - od^2\frac{1}{32\tilde{\delta}^2}. \tag{4.50}$$

Inserting these into the expression for $\mathcal{T}_{\text{diff}}$ gives

$$\mathcal{T}_{\text{diff}} \approx \frac{1}{A}\frac{1}{32\tilde{\delta}^2}\int_{\text{pinhole}} od^2\mathrm{d}A. \tag{4.51}$$

The contributions for T have canceled, and only the effect of the phase shift is left. Interestingly, the signal is proportional to od^2, which means higher optical densities can lead to a much more far-detuned signal.

The absorption, which indicates the amount of destruction, is $1-T$, which scales as $\tilde{\delta}^{-2}$, similar to the scaling of the signal. Therefore, the signal-to-destruction ratio approaches a constant in the far-detuned limit. Define the normalized destruction $\mathcal{D}$ similar to $\mathcal{T}$ as

$$\mathcal{D} = \frac{1}{A}\int_{\text{pinhole}} 1 - T\mathrm{d}A. \tag{4.52}$$

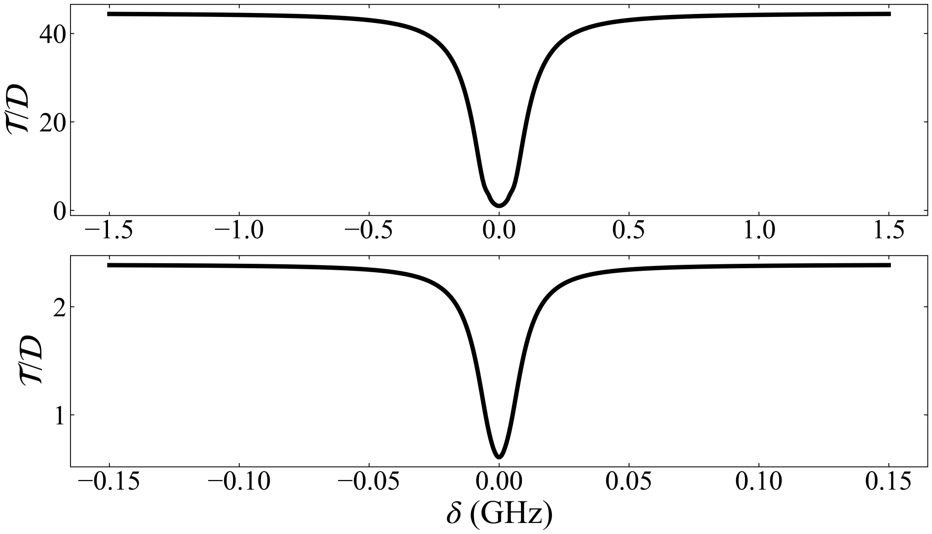

Figure 4.4: The signal to destruction ratio for a pure condensed cloud (top) and a pure thermal cloud (bottom) with $T = 500\,\mathrm{nK}$, both with $N = 100\,\mathrm{k}$, $\omega_i = 100\,\mathrm{Hz}$ and the ^{87}Rb linewidth of $\Gamma = 6.065\,\mathrm{MHz}$.

Figure 4.4 shows the ratio $\mathcal{T}/\mathcal{D}$ for a pure condensed cloud in the top panel and a pure thermal cloud at $T = 500\,\mathrm{nK}$ at the bottom panel both with $N = 100\,\mathrm{k}$ atoms as a function of detuning with $\Gamma = 6.065\,\mathrm{MHz}$ and $\omega_i = 100\,\mathrm{Hz}$. Even better ratios can be achieved with cigar-shaped clouds and more compressed clouds to increase *od*. The condensed cloud reaches much higher signal-to-detection ratios due to the high *od*, and the far-detuned region requires much more detuning than a thermal cloud. The figure also shows that the constant signal to destruction approached in the far-detuned limit is higher than any values closer to resonance.

Eq. 4.51 shows that the parameter to optimize the signal in the far-detuned regime is *od*. The *od* can be changed by minimizing the trap frequency along the light propagation direction to shape the cloud

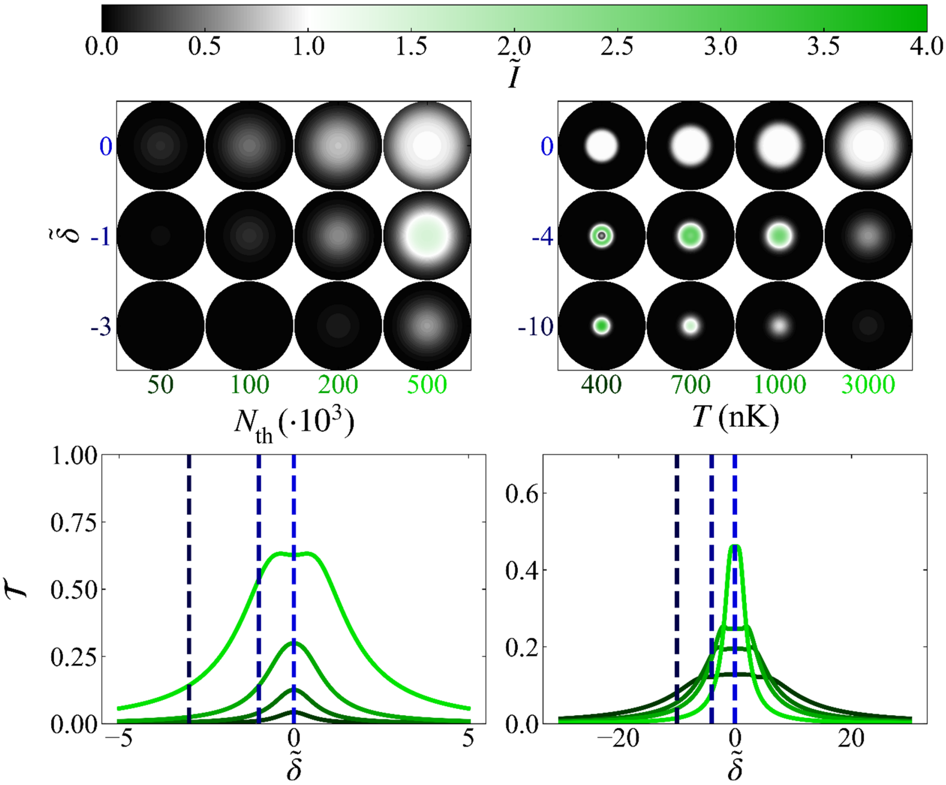

Figure 4.5: Examples of the dark field signal through different thermal clouds. Top left: the light field intensity after passing a cloud with temperature $T = 5000\,\text{nK}$ for different numbers of thermal atoms and with $\tilde{\delta} = 0$, $\tilde{\delta} = -1$, and $\tilde{\delta} = -3$. Top right: the light field intensity after passing a cloud with $N = 500\,\text{k}$ for different temperatures and with $\tilde{\delta} = 0$, $\tilde{\delta} = -4$, and $\tilde{\delta} = -10$. Bottom left: the normalized intensity spectra corresponding to the plots above. The atomic number increases from black to green, corresponding to the same colors in the x-axis labels above. Bottom right: the normalized intensity spectra corresponding to the plots above. The temperature increases from black to green, corresponding to the same colors in the x-axis labels above.

like a cigar, making the signal-to-destruction essentially limitless. In practice, the limit will come from how much the trap frequencies can be pushed and still keep a condensed cloud, as well as the imaging optics, which now have to catch the light coming from an extremely small cross-section. At some point, the thin-lens approximation will stop working as well. For bright field, the signal-to-destruction ratio will always be 1.

The focus now turns to the complicated results closer to resonance. When the phase shift becomes greater than π in some areas of the cloud, the signal will decrease again. Going from outside the cloud and towards the middle, the signal magnitude can, in principle, change from low to high to low multiple times, resulting in ring structures in the dark field intensity. Figure 4.5 shows the dark field intensities and the spectra for pure thermal clouds at different atom numbers, temperatures, and detunings, as well as the full signal spectrum. The effective in-trap pinhole radius is $75\,\mu$m and the trap frequencies in all directions are $2\pi \cdot 100\,\mathrm{Hz}$. Notice the colorbar in the top now maxes out at 4 instead of 1 as it did for bright field. Intensities above 1 is indicated with green colors, while anything in between 0 and 1 is in gray scale.

Start by looking at $\tilde{\delta} = 0$. On the top left, as more atoms are introduced, more of the light is blocked, and the middle of the clouds starts to reach saturation at $\tilde{I} = 1$ when all light is absorbed. On the top right, the cloud increases in size when the temperature increases, resulting in more signal on resonance.

When the detuning increases to $\tilde{\delta} = -1$ on the top left, the lowest atom number clouds start to dim. However, the largest actually increases in signal compared to the signal at resonance. It is also possible to spot a ring in the intensity profile. For this cloud specifically, $\tilde{\delta} = -1$ is close to the optimal detuning for signal (but not signal to destruction ratio), because here the maximum effect of the phase shift, where $\mathrm{Re}(\Delta\phi) = \pi$ happens in a radius that also maximizes the

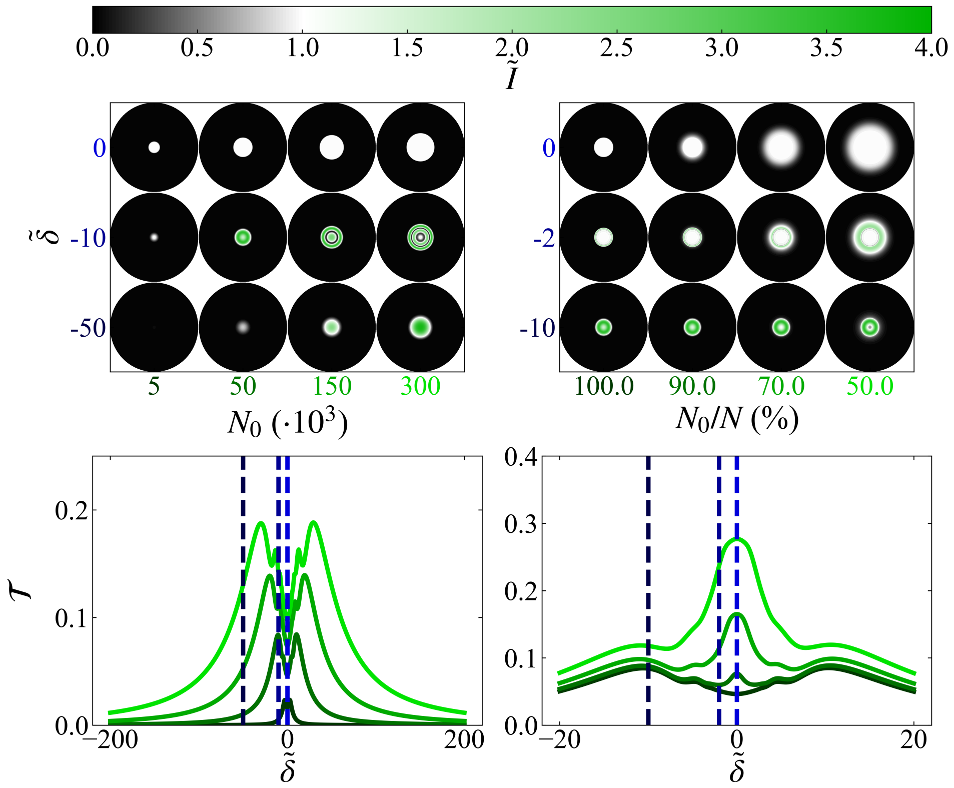

Figure 4.6: Examples of the dark field signal through different bimodal or purely condensed clouds. Top left: the light field intensity after passing a pure condensate cloud for different numbers of atoms and with $\tilde{\delta} = 0$, $\tilde{\delta} = -10$, and $\tilde{\delta} = -50$. Top right: the light field intensity after passing a bimodal cloud with $N_0 = 50\,\mathrm{k}$ for different condensate fractions and with $\tilde{\delta} = 0$, $\tilde{\delta} = -2$, and $\tilde{\delta} = -10$. Bottom left: the normalized intensity spectra corresponding to the plots above. The condensed atom number increases from black to green, corresponding to the same colors in the x-axis labels above. Bottom right: the normalized intensity spectra corresponding to the plots above. The condensed fraction decreases from black to green, corresponding to the same colors in the x-axis labels above.

cross-sectional area.

A detuning if -1 is not large enough to remove a significant amount of absorption, and the maximum intensity never reaches the highest possible of 4. The cloud *od*, even for the largest cloud, is not large enough to cause more than one full rotation in phase, and the result is a single shoulder on each side of the resonance. At even larger detunings, all clouds start to lose signal.

Some of the clouds on the top right, showing the temperature dependence, have higher *od* than the clouds on the top left. The result is that more than one full rotation in phase shift is taking place and can be seen at $\tilde{\delta} = -4$. The result is more than one shoulder appearing in the spectra. Whenever a new ring of high intensity enters the cloud radius, the power increases. The effect is strongest for the lowest temperature clouds as they have the highest *od*. Since the detuning is larger here than on the top left, the absorption is starting to wane, resulting in stronger intensities caused by the phase shift. It also makes the spectra broader in general. Notice how the high temperature cloud spectrum peaks at a higher value than the other clouds around resonance because of the large cross section, but how the lower temperature clouds have higher intensities at large detunings because of their higher *od*.

Figure 4.6 is similar to Figure 4.5 but now shows pure condensates and bimodal clouds. The effective in-trap pinhole radius is $25\,\mu$m and the trap frequencies in all directions are $2\pi \cdot 100\,$Hz. The important effect of a BEC is its extremely high *od* resulting in even broader spectra with signal remaining beyond even $\tilde{\delta} = 100$ and much more detailed ring structures and consequently more shoulders in the spectra. Similar to bright field, the thermal component of the bimodal clouds adds a narrow but tall peak in the spectra around resonance. Contrary to bright field, however, the difference in spectral features between thermal and BEC is much greater in dark field as a consequence of the phase shift. Because dark field has no background signal, it is

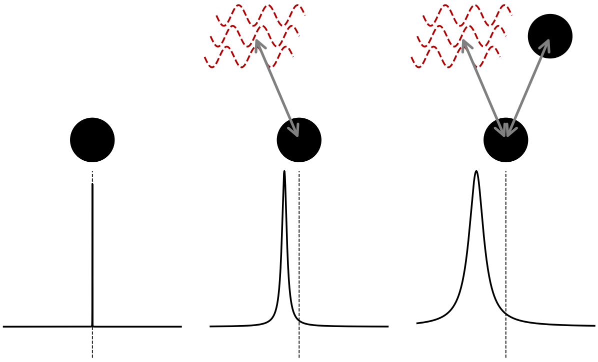

Figure 4.7: Illustration of the cause of linewidth. From left to right: more decay channels are added causing shifts in energy and increase in linewidth. Nearby atoms can function as decay channels and cause broadening.

much easier to get high-quality measurements of the BEC component in dark field compared to bright field.

4.2.3 Dipole-dipole broadening

Any orbital of an atom is an eigenstate of the Hamiltonian and should therefore have an infinite lifetime. This system is illustrated on the left in Figure 4.7 where the black dot is an atom, and with the spectrum of a single transition below. However, the atom exists in a bath of quantum field fluctuations, adding extra terms to the Hamiltonian. Consequently, the orbitals are no longer eigenstates of the full Hamiltonian of the system and will therefore have a finite lifetime, or equivalently, an uncertainty in the energy of the state. The mean energy will also be shifted due to the interactions, called the Lamb shift. This is illustrated in the middle of Figure 4.7, where the gray arrow symbolizes interactions with the red waves, which are the quantum electromagnetic field. The interaction happens via the exchange of photons.

Adding any other channels for interaction will have a similar effect to the presence of the quantum electromagnetic field background. In a cloud with more than one atom, one such channel could be the direct exchange of photons between atoms. The right side of Figure 4.7 illustrate these two channels, and how the spectrum will be broadened further and shifted in energy as a consequence. This will be called light-induced dipole-dipole broadening or simply dipole-dipole broadening. Unlike the broadening and shift due to the background, this effect will depend on the density of atoms and on the intensity of light sent into the cloud.

The main obstacle to overcome when trying to measure this effect is to achieve a high enough atom density. It is also necessary to perform the measurements in-trap, and the signal will therefore most likely be affected by the high OD of the cloud, which will disturb the apparent linewidth in a spectrum. Previous measurements were performed in very tight dipole traps, which could achieve a high atom number density with only hundreds of atoms and therefore keep the *od* low [58–60]. Other investigations have considered the impact of the light-induced forces on the distribution of condensed atoms and found non-trivial effects [61].

One of the goals of this PhD has been to measure the broadening directly via spectroscopy of ^{87}Rb-clouds and compare results with simulations done by Jan Kumlin, who is currently working as a postdoc at Technische Universität Wien. It has proven to be more complicated than initially estimated, and the work so far has instead managed to lay the groundwork for performing such measurements by understanding the many spectroscopic structures arising from high-density cloud spectra in detail.

4.3 Electromagnetically induced transparency

Adding a third level and a second light field to the system can lead to interesting effects that can be exploited experimentally to increase control and precision. One of the consequences is called electromagnetically induced transparency (EIT) and is the focus of this section.

Consider a system with three energy levels $|g_1\rangle$, $|g_2\rangle$, and $|e\rangle$, where the first two levels are much lower in energy than the third level and act like separate ground levels. The probe field, which is still the measured field after having passed through the cloud, goes from $|g_1\rangle$ to $|e\rangle$. A second field, called the coupling field, is now added from $|g_2\rangle$ to $|e\rangle$. It is assumed that neither of the fields interacts with the wrong ground state level, but $|g_2\rangle$ can still experience decoherence relative to $|g_1\rangle$.

The left side of Figure 4.8 shows the energy levels and light fields and is called a lambda scheme. The probe field intensity is assumed to be very low and might not have a well-defined Rabi frequency. Its detuning from the two-level transition resonance is δ_p. The coupling field has Rabi frequency Ω_c and detuning from its two-level resonance δ_c. The two-photon detuning $\delta = \delta_p - \delta_p$ gives the detuning of the transition $|g_1\rangle \to |g_2\rangle$ via a two-photon process.

Since the coupling field is much stronger than the probe field, a simple model of the system can be found by ignoring the probe field effects and considering only the effects of the coupling field. The system of two levels $\{|g_1\rangle, |e\rangle\}$ plus a coupling between them can be described by an interaction Hamiltonian H_I, which for Rabi oscillations is (see any book on atomic physics, e.g. [62])

$$H_I = \begin{pmatrix} 0 & \frac{\Omega_c}{2} \\ \frac{\Omega_c}{2} & \delta_c \end{pmatrix}. \tag{4.53}$$

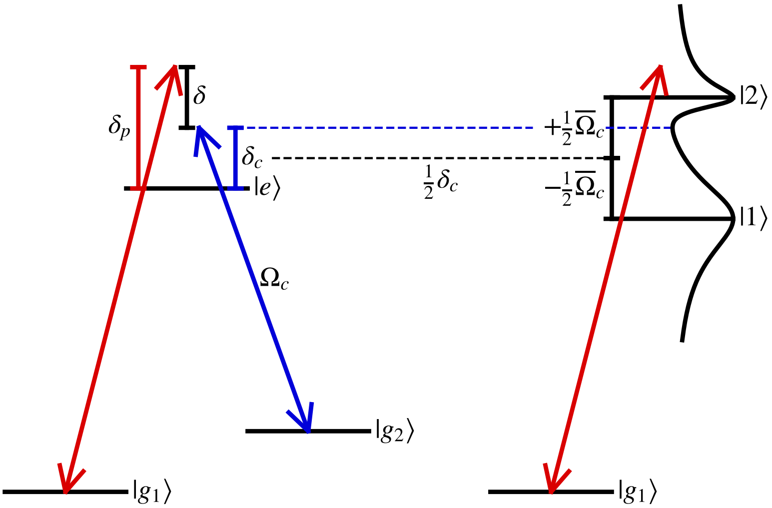

Figure 4.8: The three-level lambda scheme used for EIT (left) and the effective energy levels seen by the probe (right). The coupling laser is detuned by δ_c and leads to a set of dressed states separated by the generalized Rabi frequency $\overline{\Omega}_c$ and centered on half the coupling laser detuning (black, dashed line). The probe laser is detuned by δ_p and the two-photon detuning is $\delta = \delta_p - \delta_c$. When the probe laser frequency is scanned across the dressed states, it will encounter the polarizability from eq. 4.55 plotted across the dressed states on the right as a solid black curve. Transparency occurs at the coupling detuning (dashed blue line) when $\delta = 0$.

The eigenvalues of this Hamiltonian are

$$E_{\pm} = \frac{1}{2}\delta_c \pm \frac{1}{2}\overline{\Omega}_c, \quad \overline{\Omega}_c = \sqrt{\Omega_c^2 + \delta_c^2}, \tag{4.54}$$

where $\overline{\Omega}_c$ is the generalized Rabi frequency. The excited level $|e\rangle$ and the ground level $|g_2\rangle$ together with the coupling field have been transformed into two photonically dressed states $|1\rangle$ and $|2\rangle$ with energies relative to $|e\rangle$ given by the eigenenergies in eq. 4.54. The low-power probe field can couple to both $|1\rangle$ and $|2\rangle$ without disturbing the structure. Importantly, the same photon polarization can couple to both states, which means the interactions can interfere[6]. The offset in eigenenergies $\frac{1}{2}\delta_c$ is often not included in dressed state energies since there is no outside observation of the energy offset in a two-level system.

The right side of Figure 4.8 depicts the dressed state picture of the lambda scheme. The full derivation of the polarizability experienced by the probe field can be found in [63–66] and is

$$\chi = \frac{\mathrm{i}\tilde{\Gamma}_2 + 2\tilde{\delta}}{\tilde{\Omega}_c^2 + (\tilde{\Gamma}_2 - \mathrm{i}2\tilde{\delta})(1 - \mathrm{i}2(\tilde{\delta}_c + \tilde{\delta}))}, \tag{4.55}$$

where $\tilde{\Gamma}_2 = \Gamma_2/\Gamma$ is the normalized dephasing rate of state $|g_2\rangle$, $\tilde{\delta} = \delta/\Gamma$ is the normalized two-photon detuning, and $\tilde{\delta}_c = \delta_c/\Gamma$ is the normalized coupling detuning, and where Γ is the decay rate of the excited state $|e\rangle$. Like previously, all linewidths refer to the FWHM linewidth. The real part of χ gives the absorption and is also plotted on the right side of Figure 4.8 on top of the dressed states. All equations expressed in terms of χ from Section 4.1 and Section 4.2 can still be used with this χ.

The system returns to a two-level system when the coupling field is not present ($\delta_c = \Omega_c = \Gamma_2 = 0$ and $\delta = \delta_p$), which can easily be seen by plugging these values into eq. 4.55. It also returns to a two-level system in the limit of the coupling field being infinitely detuned ($\delta_c \to \infty$ and $\delta_p \ll \delta_c$), but now in terms of the probe detuning δ_p ($\neq \delta$) specifically

$$\chi_{\delta_c \to \infty} = \frac{-1}{2\tilde{\delta}_p + \mathrm{i}}. \tag{4.56}$$

---

[6] This is in contrast with two closely separated magnetic m_F states, which require different polarizations for interaction and can therefore not interfere.

What happens here can also be seen on the right side in Figure 4.8. When the coupling field gets detuned, both the offsets of the dressed states move in the same direction by $\frac{1}{2}\delta_c$. In the far detuned limit, $\overline{\Omega}_c = \delta_c$, and one of the dressed levels lies at the same energy as $|e\rangle$, while the energy of the other dressed level diverges. Intuitively, a far detuned coupling field no longer interacts with the system[7].

When $\delta = 0$, the probe field is on the two-photon resonance and the polarizability becomes

$$\chi_{\delta=0} = \frac{\mathrm{i}\tilde{\Gamma}_2}{\tilde{\Omega}_c^2 + \tilde{\Gamma}_2(1 - \mathrm{i}\tilde{\delta}_c)}. \tag{4.57}$$

The consequence is most clear in the ideal case of zero dephasing $\Gamma_2 = 0$, where $\chi = 0$, and thus, there is no interaction with the cloud (absorption or phase shift) at the two-photon resonance. For $\delta_c = 0$, this is the case *on* the single photon resonance where absorption is usually strongest. This induced transmission is what gives the name to EIT. What happens is that the absorption from states $|1\rangle$ and $|2\rangle$ cancel exactly on the two-photon resonance. Alternatively, it can be explained by the sum of two amplitudes. One is the probe photon being absorbed, and the other is the probe being absorbed followed by one Rabi cycle to $|g_2\rangle$ and back to $|e\rangle$, which gives a minus sign to the wave function. These two amplitudes thus perfectly cancel in the ideal system.

The effects of EIT can become extremely narrow compared to the natural linewidth. This can lead to a very steep slope in the real part of the polarizability, and consequently, an extremely slow group velocity for a light pulse propagating through the medium. Additionally, the absorption is zero on the EIT resonance (if there is no dephasing of the second ground state), and this combination allows the light pulse to

---

[7] Unless the probe becomes equally detuned, in which case the system can experience narrow stimulated Raman transitions [67].

propagate undisturbed. An observation of this unusual phenomenon was made in BEC by Lene Hau in 1999, where she and her group managed to slow light[8] down to 17 m/s [68]. Other interesting effects derived from EIT can be found in [64, 69–72].

---

[8] While propagating in the cloud, it is more correct to describe it as a superposition of photons and excited internal states of atoms, called polaritons.

Chapter 5

The experiment

After a long period of taking data prior to my PhD project, the experiment needed some restoration. Additionally, the research was taken in a new direction at the start of my PhD. Consequently, during my PhD, the experiment underwent many changes, and rather than describing all the changes made, I will instead present the experiment in its current form.

The chapter is split into four sections, starting with an overview of the experimental setup in Section 5.1, followed by a summary of the laser systems in Section 5.2, a short description of the vacuum chamber in Section 5.3, the magnetic traps, and the magneto-optical trap in Section 5.4, the control setup in Section 5.5, and finally an overview of the different detection setups in Section 5.6.

5.1 Overview

In 2003, the experiment was initially built by PhD students Jesper Fevre Bertelsen and Henrik Kjær Andersen [73, 74] under the supervision of Michael Budde. It has since undergone several changes during

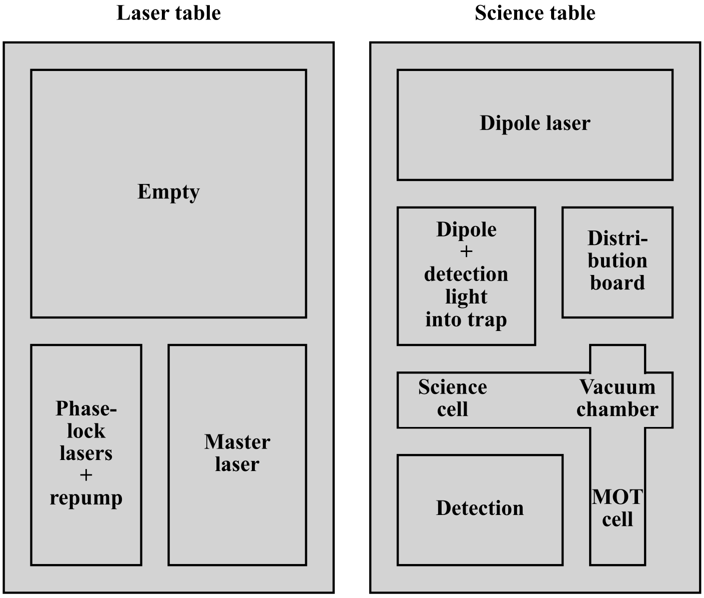

Figure 5.1: The layout of the two optical tables in the experiment. Not to scale. The left table, called the laser table, is used for the preparation and switching of the different light sources, except the dipole laser. There are three distinct areas on the laser table: an empty half not currently in use, a phase-lock setup in the bottom left, and the master laser setup in the bottom right. The right table, called the science table, contains the atoms and is used for the preparation and measurement of cold clouds. There are five distinct areas on the science table: the dipole laser setup in the top, an optical setup to control the light sources entering the trap on the middle left, a distribution board of cooling light to the six arms of the MOT on the middle right, an L-shaped vacuum chamber, and the detection setup on the bottom left.

many PhD projects [30, 75–78] and is a relatively old experiment.

The experiment uses a set of lasers and electronics to carefully manipulate ^{87}Rb atoms in a vacuum chamber to form an ultracold cloud. The cloud can then be finely manipulated and measured to perform different experiments.

Figure 5.1 displays the overall layout of the experiment. There are two optical tables, with one of them containing the vacuum chamber (right), called the science table, and the other one having all the resonant laser sources (left), called the laser table. Laser light is transported between the tables with optical fibers. It is wise to separate the resonant light of the lasers from the site of the atoms, because even a tiny amount of light entering the vacuum chamber over the course of an experiment can heat the cloud and prevent the formation of BEC. The only laser source on the science table is the 20 W dipole laser at a wavelength of 1064 nm, and stray light from this laser will not heat the atoms.

Figure 5.1 shows three areas on the laser table and five areas of the science table. Starting with the laser table, half of the optical table has been freed up during my PhD. This can be used in future projects, such as implementing new atomic species if desired. In the bottom left of the table is the phase-lock setup, which locks the beat frequency between two separate lasers and is described further in Section 5.2.2. One of the lasers has a frequency which can be used for repumping (see Figure 5.2) and currently plays a dual purpose of repumper and probe laser. In the bottom right of the table is the master laser setup. This setup is responsible for most of the manipulations of the atoms throughout the experimental sequence and is described in Section 5.2.1.

Consider now the science table. On the top is the dipole laser setup, which is responsible for handling the high-power laser light used to trap atoms (see Section 5.2.3). Below the dipole laser setup on the right is the distribution board, which splits the cooling light used in the magneto-optical trap (MOT, Section 5.4.2) into six beams with power balance controls. To the left of the distribution board is an optical setup that directs and shapes the light coming out of the optical fibers into the atom cloud. It is raised to match the height of the vacuum chamber, and below is an empty area which can be used for other projects. The vacuum chamber itself is L-shaped and is briefly described in Section 5.3. Most importantly, it is separated into sections called the MOT cell with ”high” pressure, where the atoms are initially caught, and the science cell with ultra-high vacuum, where the rest of the experiment takes place. Lastly, on the bottom left is the imaging optics leading to two cameras and a single photon detector, described in detail in Section 5.6. This part is also raised, and more free space exists below. Above the detection setup is a small optical setup, which has one of the cooling beams for the MOT as well as the repump and depump lasers leading into the science cell from above.

The overall sequence of the experiment follows a typical path [79, 80]. It begins with gathering atoms in the MOT, then transports the atoms with magnetic quadrupole coils on a transport stage from the MOT chamber to the science chamber, where forced RF evaporative cooling takes place until the cloud is cold enough to transfer to a hybrid trap. Here, further evaporative cooling takes place by lowering the dipole beam power until the cloud can be transferred to a pure crossed dipole trap, where the final evaporation steps take place until BEC is produced.

Typically, the experiment can produce $\approx 200\,\mathrm{k}$ BEC atoms. By carefully adjusting all parameters in the sequence and optimizing the MOT setup (see Section 5.4 for more details), the duration of the entire experiment (plus the loading of the MOT) was reduced from $\approx 90\,\mathrm{s}$ to $20 - 30\,\mathrm{s}$, depending on the dipole beam waist. A smaller waist increases the collision rate between atoms and results in faster evaporation.

5.1.1 Rubidium 87 properties

Rubidium is an alkali with one electron in the outer shell and thus, the energy scheme follows a similar structure to hydrogen. The ground state manifold of the ^{87}Rb is $5^2S_{1/2}$ and the first two excited manifolds are $5^2P_{1/2}$ and $5^2P_{3/2}$ of which only the D_2 transition to $5^2P_{3/2}$ is being used in the experiment, and the D_1 transition to $5^2P_{1/2}$ will therefore not be mentioned in any detail later. The nucleus has an angular momentum of $I = 3/2$, which results in two hyperfine energy levels in the $5^2S_{1/2}$ orbital with $F = 1, 2$, and in four hyperfine levels in the $5^2P_{3/2}$ orbital with $F = 0, 1, 2, 3$. These hyperfine levels are depicted in Figure 5.2 with the different laser frequencies used in the experiment.

The hyperfine levels are each split into $2F + 1$ magnetic sub-levels. These are shown in Figure 5.3 with the laser transitions, which need specific polarities. Further details on the energy levels and properties of ^{87}Rb can be found in [81].

Rubidium is very suitable for the production of BEC due to several key properties [57] and was the first atomic species used to achieve it [15]. The transitions of ^{87}Rb lie in a region where commercial diode lasers are available, which significantly improves the access to relatively cheap and highly stable lasers for the experiment. The mass of ^{87}Rb is such that the speed of atoms at room temperature is not too high for a MOT to catch a sufficient number of atoms from a low-pressure gas without the need for special equipment, such as Zeeman slowers, which are often used for lighter atoms, sodium for instance [82]. Efficient cooling without too much heating or atom loss is enabled by the ratio between the elastic and inelastic scattering rates of ^{87}Rb [83]. Importantly, the magnetic moment of ^{87}Rb is enough to allow for efficient magnetic trapping of the atoms. Unfortunately, the Feshbach resonances of ^{87}Rb are limited [84], which can restrict the application of ^{87}Rb in some experiments.

The level structure of ^{87}Rb in Figure 5.2 and 5.3 cannot always be

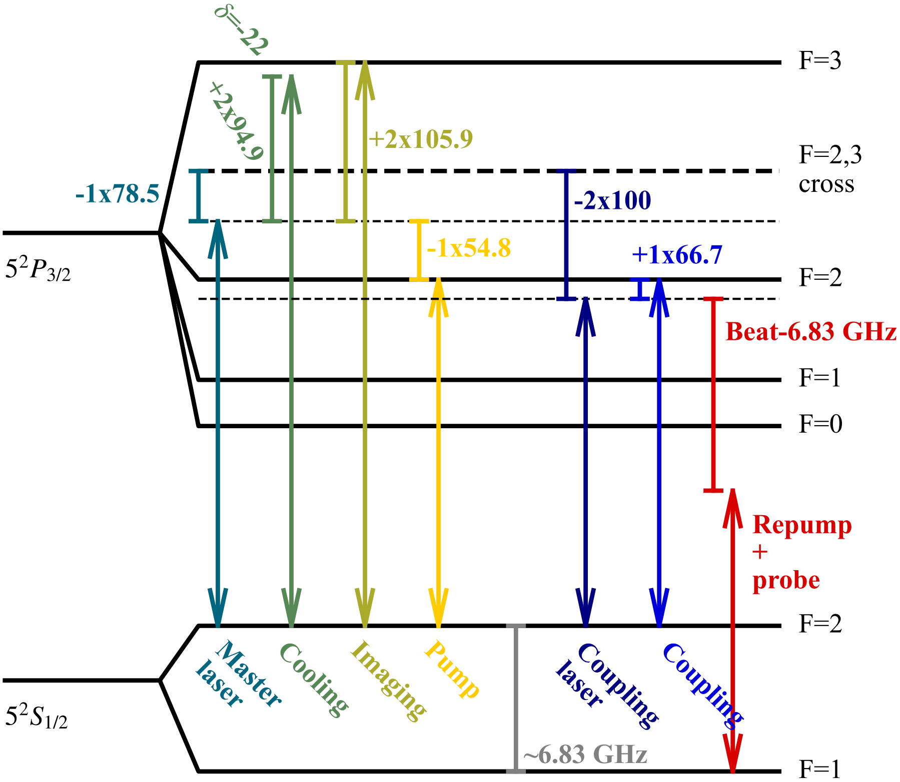

Figure 5.2: Transitions used in the experiment. From left to right: the master laser output frequency locked 78.5 MHz below the $F = 2 \leftrightarrow F' = 2, 3$ crossover, the cooling light at -22 MHz from $F = 2 \leftrightarrow F' = 3$ and can be varied, the imaging light resonant with $F = 2 \leftrightarrow F' = 3$ and can be varied, the pump light resonant with $F = 2 \leftrightarrow F' = 2$ and cannot be varied, the coupling laser output frequency locked 66.7 MHz below $F = 2 \leftrightarrow F' = 2$ and can be varied, the coupling light frequency resonant with $F = 2 \leftrightarrow F' = 2$ and can be varied, and the repump/probe light beat locked to the coupling laser. The vertical bar for the repump/probe light indicates the beat frequency minus the difference between $F = 1$ and $F = 2$ (gray). This color scheme is repeated in later figures.

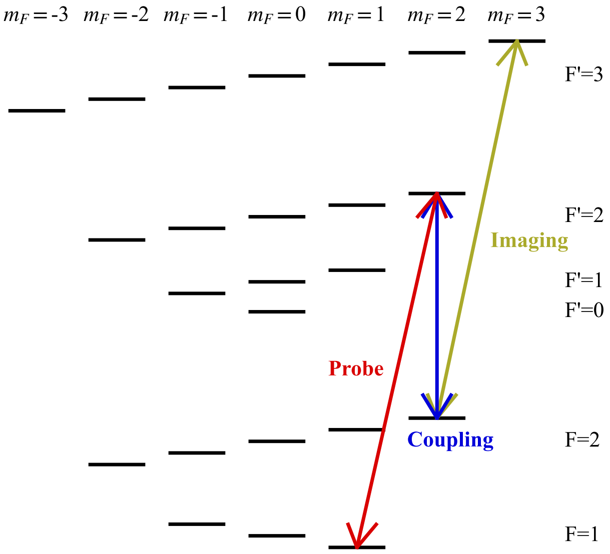

Figure 5.3: The specific m_F-states and polarizations used for the imaging, probe, and coupling light. The imaging light uses the cycling transition with σ^+ polarization. The probe light also uses σ^+ light, while the coupling light uses π-polarization. Because the probe and coupling beams are orthogonal, one of them must use circular polarization and the other linear polarization. This configuration is the least susceptible to pumping outside the lambda scheme levels in EIT, and avoids pumping to other m_F-states in the ground state $F = 1$ manifold via the probe light.

approximated as a two-level system. One effect is the change in the resonant cross section σ_0, which is $6\pi/k^2$ in an ideal two-level system. More generally, the resonant cross section of a transition from the

(J, F) manifold to the (J', F') manifold is given by [81]

$$\begin{aligned}\sigma_0 =&(2F+1)(2F'+1)(2J'+1)\\ &\cdot \begin{pmatrix} F & 1 & F' \\ m_F & q & -m'_F \end{pmatrix}^2 \begin{Bmatrix} J & J' & 1 \\ F' & F & I \end{Bmatrix}^2 \frac{6\pi}{k^2},\end{aligned} \tag{5.1}$$

with Wigner 3-j and Wigner 6-j symbols, and where $q = -1, 0, 1$ is the projection of the light polarization.

In our experiment, the cross section is used for the imaging light and for the probe light shown in Figure 5.3, both with $q = 1$. The imaging light from $J = \frac{1}{2}, F = 2, m_F = 2$ to $J' = \frac{3}{2}, F' = 3, m'_F = 3$ lies on the cycling transition and resembles an ideal two-level system because there are no dipole-allowed decay channels outside the two levels. In this case, the 3-j symbol squared is $\frac{1}{7}$, the 6-j symbol squared is $\frac{1}{20}$, and the cross section is $6\pi/k^2$, just like the ideal two-level system. The probe light from $J = \frac{1}{2}, F = 1, m_F = 1$ to $J' = \frac{3}{2}, F' = 2, m'_F = 2$ has a 3-j symbol squared equal to $\frac{1}{5}$ and a 6-j symbol squared equal to $\frac{1}{24}$ giving the cross section $\frac{1}{2}6\pi/k^2$ which is half the two-level cross section.

5.2 Laser systems

This section covers the three laser systems in the experiment: the master laser in Section 5.2.1, the phase-lock lasers in Section 5.2.2, and the dipole laser in Section 5.2.3. The setups are optimized to consist of a low number of optical components and short free-space propagation paths to increase stability. They are also optimized to be easily accessible on the laser table. The exception is the propagation of the repump laser from the phase-lock setup to the master laser setup, which is a temporary configuration. A separate repump laser setup is planned to be built in the future. All laser setup drawings are only roughly to scale.

5.2.1 Master laser setup

The master laser system is responsible for most of the manipulation and imaging of the atoms. The laser is a Toptica DLC Pro external cavity diode laser with an in-built optical isolator. Figure 5.2 shows the output frequency of the laser on the left as well as the three different frequencies derived from it via acousto-optic modulators (AOMs). The frequency of the laser light is shifted up by an AOM with frequency 78.5 MHz and is then locked to the Doppler free spectrum of ^{87}Rb in a cell, specifically to the $F = 2 \rightarrow F' = 2, 3$ crossover peak resulting in a detuning from the $F = 2 \rightarrow F' = 2$ transition of $+54.8$ MHz and from the $F = 2 \rightarrow F' = 3$ transition of -211.8 MHz. From left to right in Figure 5.2, the first frequency after the master laser's direct output, corresponds to the light used to trap and cool atoms in the MOT, the second frequency is the imaging light, and the last frequency is used for pumping atoms from $F = 2$ to $F = 1$ in the $S_{1/2}$ manifold.

The AOMs provide the frequency shifts, and by quickly turning on or off the RF-power, the light in that respective arm can be quickly switched on or off. The duration of the switch time is given by the time it takes the sonic wave in the AOM crystal to pass the width of the laser beam mode. A focused laser through the AOM can thus be switched faster, although the switching time is almost always fast enough even without focusing. The AOMs are not able to switch off the light completely, and since only a few resonant photons can heat to the ultracold clouds, there is always a shutter before each light beam enters the trap. The shutters are primarily placed on the laser table to avoid vibrations, and if they are used on the science table, they hang from the ceiling.

The cooling arm needs more power than the other arms, which is achieved by a tapered amplifier (TA) after the AOM. The TA emits collimated light in both directions, and it is sensitive to the reflections

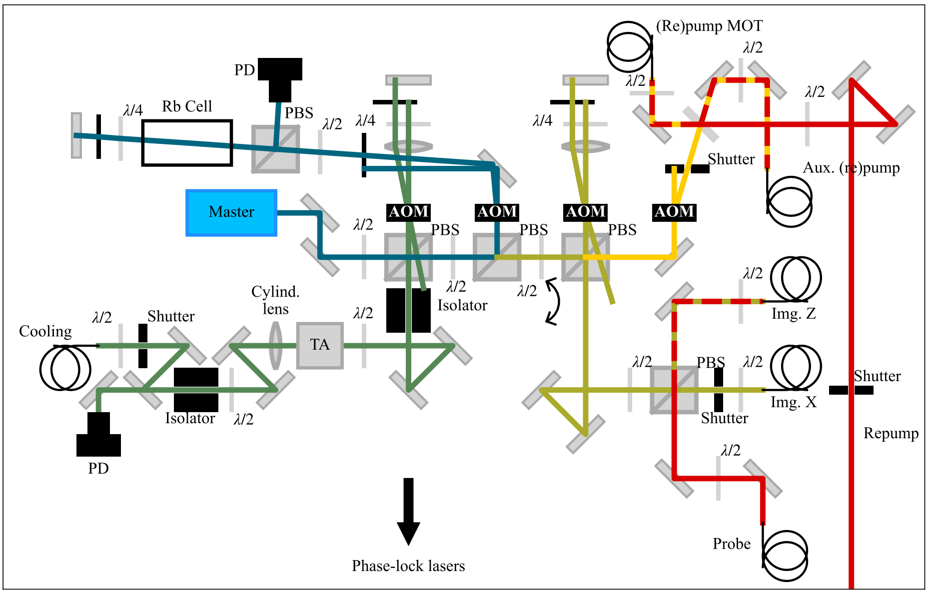

Figure 5.4: The master laser setup. A central row of polarizing beam splitters splits the laser beam into 4 different paths, each in a separate color matching those used in Figure 5.2. Cooling and imaging light go through double-pass AOMs, and spectroscopy and pump light go through single-pass AOMs. Only pump light uses the −1'st order, and cooling, spectroscopy, and imaging light use the +1'st order. Probe and repump light are fed into the setup from the phase lock lasers shown in Figure 5.6.

coming back into the TA from the high-power side. It is therefore necessary to place optical isolators both before and after the TA.

The setup of optical elements to achieve these goals is shown in Figure 5.4. The laser light passes through a row of polarizing beam-splitting cubes in the middle, which splits the beam into four beams moving upwards (on the figure), and the power distribution can be controlled with $\lambda/2$ wave-plates.

The first upward-moving beam from the left is the cooling light.

It enters an AOM which increases the light frequency dynamically between 70 MHz and 120 MHz. The frequency when loading the MOT is 94.9 MHz. It then passes a lens after one focal length to make sure the beam direction is independent of the AOM frequency. A $\lambda/4$ wave-plate turns the linearly polarized light into circularly polarized light, and an iris filters the correct diffraction order from the AOM. The beam is reflected back one focal length away from the lens, which flips the circular polarization, and after passing the $\lambda/4$ wave-plate a second time, the light becomes linearly polarized again, but in the orthogonal direction. The beam passes the lens again and becomes collimated. After passing through the AOM a second time, the frequency is shifted up again, and the light path in the correct diffraction order lies on top of the incoming beam. Since the polarization is rotated, it now passes through the cube and into the first optical isolator before the TA, which has a large aperture and does not need to be finely coupled with mirrors.

Two mirrors and a $\lambda/2$ wave-plate after the isolator are used to couple the light into the TA. The TA output mode is diverging in one direction, which is collimated with a cylindrical lens. After the second isolator, the beam is coupled into an optical fiber leading to the distribution board. The small leakage behind one of the mirrors is measured to monitor the TA output power. It is important that the TA outputs enough power when turned on; otherwise, the chip can possibly be damaged. It needs about 15 mW of input power to achieve approximately 700 mW of output. Only about 300 mW of that power is transferred through the fiber because of the ”bad” output mode of the TA, which is difficult to mode-match accurately to the optical fiber.

The second upwards-moving beam in Figure 5.4 is used for Doppler-free spectroscopy and for locking the laser [85]. After passing through an AOM which shifts the frequency up by 78.5 MHz, the light passes through a ^{87}Rb cell twice, and the power is measured on a photodiode. The frequency of the AOM is modulated by the laser control unit, which also demodulates the measured power to generate an error signal used for locking the laser. A double-pass $\lambda/4$ wave-plate ensures the correct reflection of the light by the polarizing beam splitter.

By beating two such locked lasers against each other, and assuming the frequency width of both lasers to be the same, the 1σ width of the locked laser light was measured to be 162 kHz. The beat signal can be seen in Figure 5.5.

Before the third upwards-moving light beam in Figure 5.4 there is a flippable $\lambda/2$ wave-plate, which shifts between most light passing through the last beam splitter and most light being reflected. This ensures the most power available for optical pumping in the MOT or the imaging when they are needed. The reflected light on the last cube leads to the imaging arm and starts by passing through a double pass AOM setup similar to that of the cooling arm, except the AOM frequency is now 105.9 MHz plus small corrections due to magnetic fields. It then reaches the two mirrors below the cube, which are used to couple into the imaging fiber for the X-direction. Degrees of freedom on the cube after the two mirrors and the mirror above that cube are used to couple into the imaging fiber for the Z-direction. A $\lambda/2$ wave-plate is used to distribute power between the two arms. The main imaging direction is X, and most of the light is set to go to that arm, unless more power is needed in Z for special purposes. From the phase-lock setup, a low-power probe beam comes from an optical fiber and is also coupled into the Z direction.

The final part of the setup in Figure 5.4 comes after the three cubes and produces the pump light. The light passes through an AOM once, and the frequency is shifted down by 54.8 MHz. Pump light is used both in the MOT cell and in the science cell (called aux. pump). The beam is therefore split in two by a 50/50 sampler and coupled into two fibers just after the AOM and the shutter. During optical pumping of atoms, the correct polarization of repump light is needed, and the

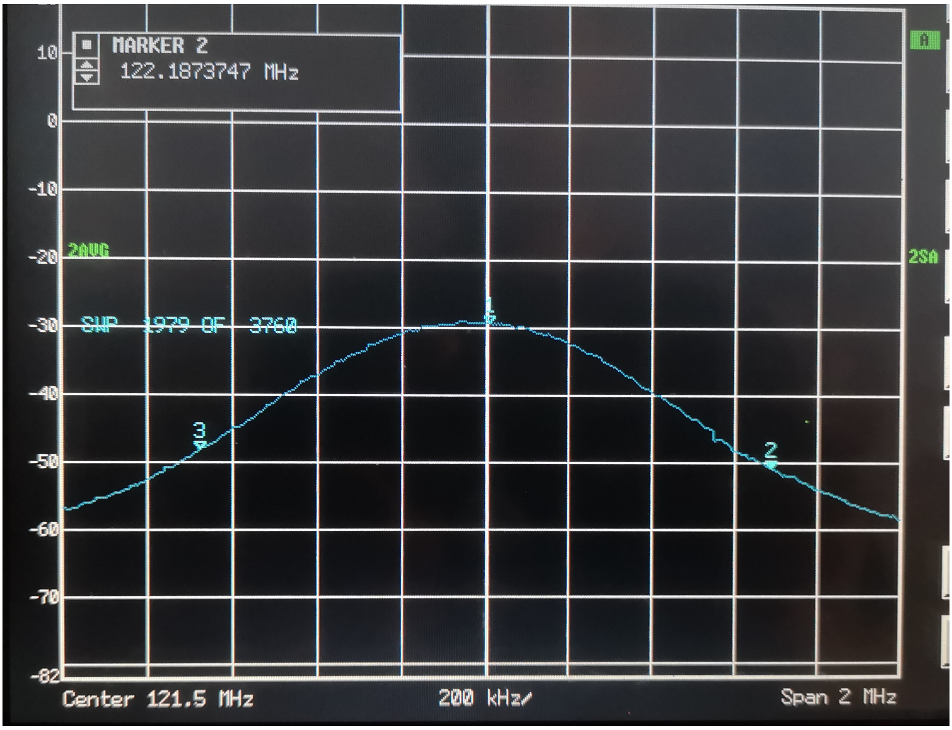

Figure 5.5: Beat signal between two similarly locked lasers to measure the frequency width. The signal was averaged 2000 times. Oscilloscope settings were $RBW = 1\,\mathrm{kHz}$, $VBW = 1\,\mathrm{kHz}$, $SWT = 10\,\mathrm{ms}$, and a span of $2\,\mathrm{MHz}$. A $19.5\,\mathrm{dB}$ drop corresponds to a 3σ distance from the peak indicated by marker 1 if a Gaussian shape is assumed, and this is indicated by marker 2. Dividing the distance by $3\sqrt{2}$ gives a 1σ frequency width of $162\,\mathrm{kHz}$ for both lasers, assuming they are equal.

repump light for loading the MOT cannot be used. Repump light is also needed in the science cell. Some of the repump light is therefore directed to the master setup from the phase lock setup and is coupled into the same optical fibers as the pump light. Since the polarizations have to be the same, it is not possible to replace the 50/50 sampler

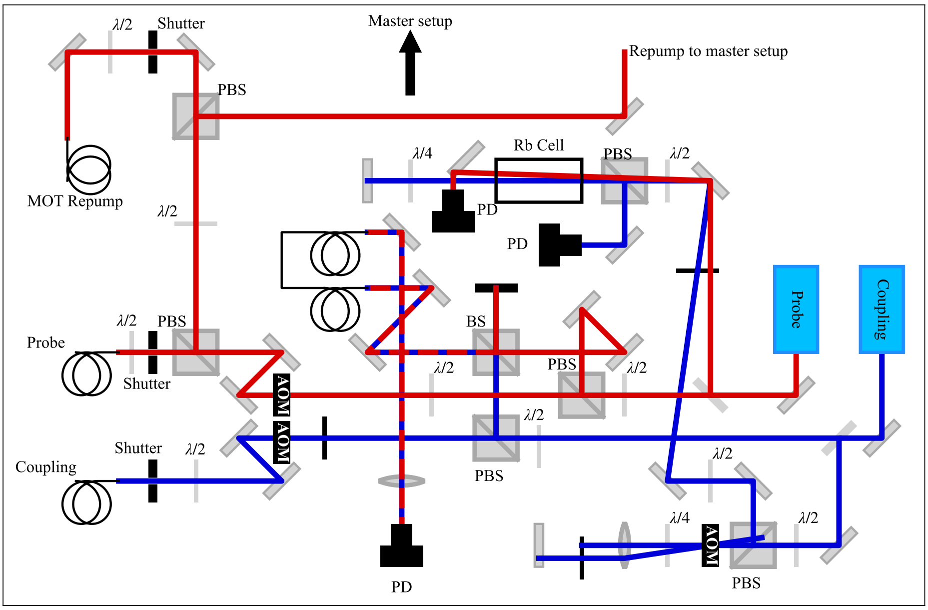

Figure 5.6: Optical setup of the phase-locked lasers. The coupling laser is locked by Doppler-free spectroscopy after a double pass AOM, and the probe laser is beat locked to the coupling laser to provide the probe and the repump light. The beat signal is measured with a ZX85-12G-S+ $0.2 - 12\,000$ MHz fast photo diode from Mini-Circuits. The coupling and probe/repump light is shifted by single-pass AOMs, which are locked together in frequency. The +1st order is used in all AOMs.

with a polarizing beam splitter.

5.2.2 Phase lock lasers and repump setup

The phase lock laser setup plays a very versatile role in the experiment. First, it provides the repump light needed in different parts of the sequence. Secondly, it provides a well-controlled probe laser which can be scanned over 12 GHz. Finally, it can produce a pair of coherent lasers beams used for electromagnetically induced transparency.

The setup consists of two Toptica DLC Pro external cavity diode lasers with internal optical isolators similar to the laser in the master laser setup. One of the lasers, called the coupling laser, is locked via Doppler-free spectroscopy to the $F = 2 \to F' = 2, 3$ crossover peak, and the other laser, called the probe laser, is beat locked to the coupling laser. Because the probe laser frequency is controlled by the beat lock with feedback directly into the piezo voltage and diode current, and not by an AOM, the frequency span is much larger and only limited by mode jumps or electronics.

The setup is shown in Figure 5.6. The coupling laser beam is depicted in blue and the probe laser beam in red. Consider first the coupling laser setup. Just after the laser output, a 10% sampler sends the light into a double-pass AOM setup. The AOM shifts the frequency up by 100 MHz on each pass and can be modulated to obtain a lock signal. A $\lambda/4$ wave-plate rotates the polarization to be reflected by the polarizing beam splitting cube, leading to a double pass in a ^{87}Rb cell to achieve Doppler-free spectroscopy. The result is the main laser beam being locked at 66.67 MHz below the $F = 2 \to F' = 2$ transition. The main laser beam proceeds to a polarizing beam splitter, which splits the light into an arm used for beat locking and another arm that leads to the science table. Before being coupled into the fiber to the science table, the frequency is shifted up by a single pass AOM by 66.67 MHz. The single-pass AOM frequency is fixed, and fine adjustments on the laser frequency are done on the double-pass AOM.

In the beat lock path, the light goes through a 50/50 beam splitter, and the reflected half is coupled into a fiber for cleaning the spatial mode, which then outputs into a fast photodiode with 12 GHz bandwidth. The 50/50 beam splitter is necessary for the polarizations of the coupling and the probe to be identical to create a beat signal. A minimum of 5 mW of power from the coupling laser should hit the diode.

The probe laser setup also has an early 10% sampler, which leads to a single pass of the ^{87}Rb cell. It is used only for monitoring the laser, and can be cut off with an iris to avoid interference with the coupling lock[1]. The main laser beam then proceeds to a single-pass AOM with frequency locked to the corresponding single-pass AOM for the coupling beam. Here, the light is split into two paths: a low-power probe beam going to the master laser setup via a fiber, and a repump beam. The repump beam is split into two parts, with most of the power going through the beam splitting cube and into a fiber leading directly into the MOT. This part is used for loading the MOT only. The rest of the light goes into the master laser setup and is used in the science cell for manipulating atoms and during the pumping phase of the MOT.

The repump part of the setup will be replaced by a separate laser with a Doppler-free spectroscopy lock in order to avoid the limits it puts on the probe laser. Because the probe laser has to be used for repump, it is currently not possible to scan too far away from the repump frequency if the repump has to be used soon after[2], because the lock might not be fast enough. In addition, it is not possible to switch the probe beam to go from $F = 2$ if that should be desired[3]. Finally, future experiments could require more power in the probe beam, which is not available if the repump light is derived from the same laser.

Part of the probe laser light is extracted in a polarizing beam splitting cube and overlapped with the coupling light in the 50/50 beam

---

[1] If it is not removed, room temperature EIT is visible in both lasers' spectroscopy.

[2] Such as for imaging right after a probe sweep.

[3] There could be benefits to the probe light being on the $F = 2, m_F = 2 \to F' = 3, m'_F = 3$ cycling transition, for instance.

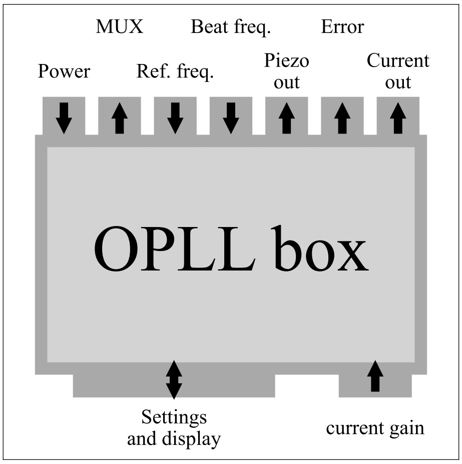

Figure 5.7: The OPLL box seen from above. Settings can be controlled from the front (down on the figure), and the inputs and outputs are on the back. ”MUX” and ”Error” are used only for monitoring.

splitter in order to get a beat frequency. A minimum of 5 mW of power from the probe laser should hit the diode after the cleaning fiber.

On the science table, after the fibers, both the probe light and the coupling light powers are monitored with photo diodes, and the power is stabilized. The probe light has to be sampled before potential OD filters.

5.2.2.1 Optical phase-lock loop

The output of the fast photo diode leads into two 15 dBm amplifiers from Mini-Circuits (ZX60.8008E-S+), a four times frequency divider from Analog Devices (HMC362S8G) designed for a maximum frequency input of 12 GHz, and a directional coupler for monitoring the beat signal. The coupler is rated for $2 - 8$ GHz (or $8 - 32$ GHz beat signal), which is not ideal, but it can be removed when monitoring is not needed. After this setup, a short SMA cable designed for MW frequencies leads to the lock control box called the OPLL (optical phase-lock loop) box, described here [86]. The box has three inputs (power, a reference frequency, and the beat frequency) and four outputs (a MUX signal for troubleshooting, a voltage signal for piezo control, an error signal, and a voltage signal for diode current control). To control its functions, the box has a potentiometer to adjust the gain of the current control output, a display, and a set of settings to choose on the display (see Figure 5.7). The input reference signal should have an amplitude between $2\,\mathrm{V_{pp}}$ and $3\,\mathrm{V_{pp}}$ and the input beat signal should be between $-10\,\mathrm{dBm}$ and $10\,\mathrm{dBm}$, preferably in the middle. An ATMEGA168-PU chip takes care of the display and the settings from the front, while an ADF4107 chip measures the difference between input and reference frequencies and outputs the piezo signal, error signal, and current signal. The ADF4107 can handle a beat signal frequency between 1 GHz and 7 GHz and a reference signal frequency between 20 MHz and 250 MHz. A table of recommended settings is shown in Table 5.1 with explanations of the most important ones.

The ADF4107 divides the beat frequency input f_{beat} (which is already divided by four by the frequency divider) by a factor N and the reference frequency input f_{ref} by a factor R that can both be changed in the settings. A better lock is achieved with a low N/R ratio. The error signal output is used for monitoring the lock and is between 0 V and 10 V. If the laser is locked, the signal is flat at 5 V. Any noticeable variation in the error signal means the lock is not working properly and is typically caused by too weak frequency input signals or current control gains.

The piezo signals are capped by the ADF4107 chip. When doing broad sweeps of the probe laser frequency in an experiment, it is important to make sure the piezo voltage can complete the entire sweep. This can be seen by monitoring the error signal during a sweep. If the experiment is running for a long time, the laser itself can drift, which drags the necessary piezo signal along with it. After some time, the piezo signal might not be able to do the full sweep anymore.

The beat signal measured through the directional coupler can be seen in Figure 5.8, where Figure 5.8a shows a wide beat spectrum and Figure 5.8b shows a zoomed-in version. The spectrum analyzer could not provide a resolution below 1 Hz, so it was not possible to measure the width of the lock signal, except that it is sub-Hertz level. To get the optimal lock, the potentiometer for the current control gain should be set to 9.25.

The upper limit for the rate of sweeping the laser frequency with the lock (by sweeping the reference frequency) is about 14 MHz/ms, which puts a limit on how fast spectroscopic data can be taken.

A back-of-the-envelope estimate of the broadening of the laser due to sweeping can be found by setting the broadening of measuring within a small time interval Δt equal to the swept frequency within that window

$$\sigma_f \approx \frac{1}{\Delta t} = \frac{\mathrm{d}f}{\mathrm{d}t}\Delta t, \tag{5.2}$$

which gives

$$\sigma_f \approx \sqrt{\frac{\mathrm{d}f}{\mathrm{d}t}}. \tag{5.3}$$

The broadening of a 14 MHz/ms sweep is approximately 118 kHz and is comparable to the Doppler free lock, while the broadening of a 100 kHz/ms sweep is approximately 10 kHz. It rarely makes sense to measure more precisely than 10 kHz since many effects in the experiment (such as laser frequency width, background magnetic field drifts, or decoherence effects) are of the same order, but it is important to keep an eye on the sweep rate if ultra-high precision is needed.

Setting	**Recommended setting**	**Explanation**
PFD Pol.	Negative	Whether the lock tends towards the middle or the extreme of the error signal. Try switching the polarity if the lock does not work.
N	96	Beat frequency division before comparison with reference frequency.
R	1	Reference frequency division before comparing with the beat frequency. Combined with N above, the criterion is $f_{\mathrm{beat}}/N = f_{\mathrm{ref}}/R$.
A. Backlash	1.3 ns	
CP Gain	6 pips	
Lockdet.	fast	
MUX	Disabled	Different settings are used to monitor outputs from the MUX channel in the back.
Ch.Pump	Enabled	
ErrorZero	512	Half of $2^{10} = 1024$ and defines the lock point of the error signal.

Table 5.1: OPLL box recommended settings, and explanations for the most important ones.

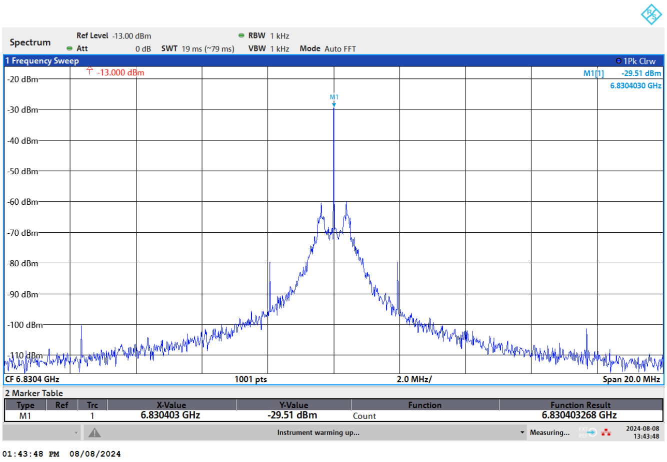

Spectrum
Ref Level -13.00 dBm RBW 1 kHz
Att 0 dB SWT 19 ms (~79 ms) VBW 1 kHz Mode Auto FFT

1 Frequency Sweep 1Pk Clrw

M1[1] -29.51 dBm
6.8304030 GHz

CF 6.8304 GHz 1001 pts 2.0 MHz/ Span 20.0 MHz

2 Marker Table

Type	Ref	Trc	X-Value	Y-Value	Function	Function Result
M1		1	6.830403 GHz	-29.51 dBm	Count	6.830403268 GHz

Instrument warming up... Measuring... 2024-08-08 13:43:48

01:43:48 PM 08/08/2024

(a) The beat lock spectrum. The span is 20 MHz.

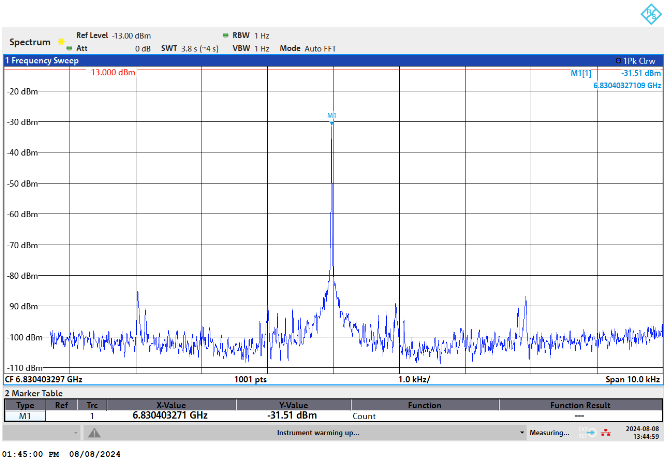

Spectrum
Ref Level -13.00 dBm RBW 1 Hz
Att 0 dB SWT 3.8 s (~4 s) VBW 1 Hz Mode Auto FFT

1 Frequency Sweep 1Pk Clrw

M1[1] -31.51 dBm
6.83040327109 GHz

CF 6.830403297 GHz 1001 pts 1.0 kHz/ Span 10.0 kHz

2 Marker Table

Type	Ref	Trc	X-Value	Y-Value	Function	Function Result
M1		1	6.830403271 GHz	-31.51 dBm	Count	---

Instrument warming up... Measuring... 2024-08-08 13:44:59

01:45:00 PM 08/08/2024

(b) A zoomed-in beat lock spectrum. The span is 10 kHz.

Figure 5.8: The beat lock spectrum locked at 6.83 GHz, which is roughly the $F = 1 \rightarrow F' = 2$ resonance for the probe laser. (a) shows a broad spectrum, and (b) shows a zoomed-in version. Spectrum analyzer settings on the bottom were $RBW = 1\,\text{Hz}$, $VBW = 1\,\text{Hz}$, and $SWT = 3.8\,\text{s}$.

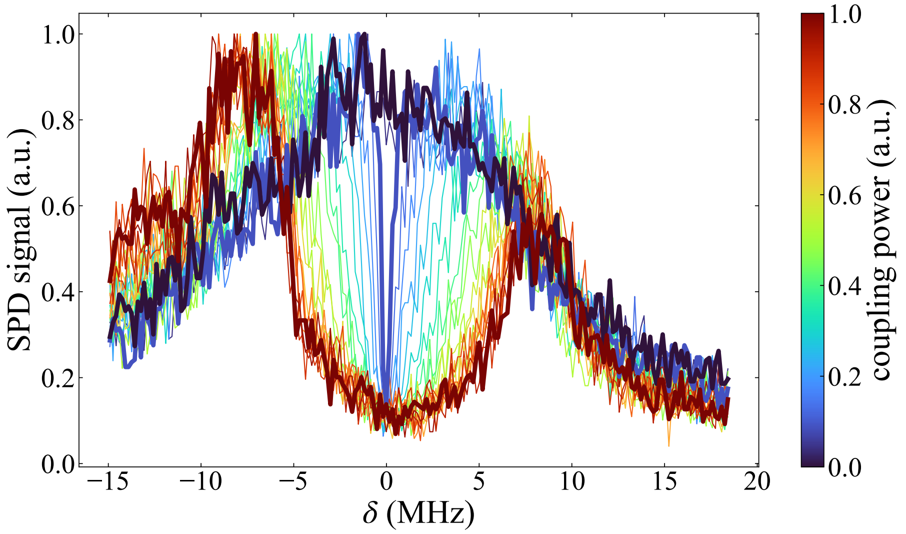

Figure 5.9: The dark field signal through a very small cloud with varying coupling field power as a function of two-photon detuning δ. Three of the curves have been emphasized with greater linewidth: zero coupling power (dark blue), very low coupling power (blue), and maximum coupling power (dark red). Notice how the dip in the middle widens with more power until the spectrum starts to approach two non-interfering Lorentzians (the dark red spectrum), which is the Autler-Townes regime.

5.2.2.2 EIT

The phase-locked laser system has been used to measure EIT in ultracold atoms with great controllability and precision. An example of such a measurement is displayed in Figure 5.9, which shows the dark field signal on a SPD (see Section 5.6.2 for the detection setup) while scanning the two-photon detuning δ across the EIT resonance for different coupling field strengths. The cloud contained very few atoms, for the spectroscopy to closely match the shape of χ.

As the coupling field strength is increased, the spectroscopy transitions from a Lorentzian to a Lorentzian with a narrow dip in the middle called the EIT regime, and finally to two separated Lorentzians called the Autler-Townes regime. Three of the curves are highlighted in the plot to illustrate the spectra in these three regions.

The original idea for building a phase-locked pair of lasers and measuring EIT was to use the narrow spectroscopic features for improved precision in measuring light-induced dipole-dipole broadening, and potentially the energy shift that accompanies the effect. There is, however, no immediate gain from using EIT in this case since any change in the width of the EIT transparency window scales relative to the width of the window itself. The relative precision needed to measure broadening effects is therefore unchanged, while the absolute precision necessary is higher than the one be needed to measure absorption linewidths. On top of that, the setup is more complicated, and describing the spectra involves more parameters. Additionally, due to the dephasing of the second ground state in the lambda-scheme, EIT never reaches complete transparency. Due to the high optical density of in-trap BEC, the measured transparency in bimodal clouds is mostly due to the thermal component, and the transparency spectrum essentially shields the contribution from BEC, which is the region of interest.

The EIT resonance can also not be used for improving energy shift measurements on the D2-line since it is a two-photon process, and only energy shifts between the two ground states will have a direct effect. It might be possible to use EIT for measuring the energy shift through the varying shape in the spectrum caused by coupling detunings, but this will most likely not be more precise than measuring the absorption spectrum directly, since the important parameter here is δ_c/Γ.

While the use of EIT in the experiment is currently on hold, it still holds the potential for interesting measurements because of the high precision, controllability, and the non-trivial effects deriving from EIT. Two such effects are the phenomenon of slow light and the existence of polaron-polaritons [87]. EIT can also be used to decrease the effect of the natural lifetime of the excited state on optical Rabi frequency measurements by measuring Raman Rabi frequencies [88–90]. For now, the phase-lock setup is used as a very versatile setup to control the probe laser frequency.

5.2.3 Dipole trap setup

This part of the experimental setup provides the optical confinement of atoms used both for hybrid trapping and crossed optical dipole trapping (cODT). It is a standard technique in the field with extensive literature, such as [62, 79, 91–94], and thus only a short overview of dipole trapping is given here. The setup was originally built by Laurits N. Stokholm [95]. It has since been modified, and is thus presented in this thesis in an up-to-date version and with a focus on the current experimental goals.

5.2.3.1 Single beam dipole traps

As atoms enter a light field, they are adiabatically transferred to one of the dressed states. If the field has a spatially varying intensity, the result is a similarly spatially varying energy of the atoms acting as an effective potential energy. If the light field frequency is far-detuned from a resonance by δ, the effective potential U is

$$U(\mathbf{r}) = \frac{\hbar\Gamma^2}{8\delta}\frac{I(\mathbf{r})}{I_{\text{sat}}}. \tag{5.4}$$

For a multi-level atom, the saturation intensity I_{sat} takes on an effective value which is a combination of interactions with all states. In case of ^{87}Rb and a far-detuned light field from the D2-line (but still not interacting with any other transitions) with π-polarization, the effective saturation intensity is $I_{\text{sat}} = 25.03\,\text{W/m}^2$ [81]. The absorption rate R_{abs} of photons for a single atom is

$$R_{\text{abs}} = \frac{\Gamma^3}{8\delta^2}\frac{I}{I_{\text{sat}}}. \tag{5.5}$$

The rate decreases with $1/\delta^2$ while the potential decreases with $1/\delta$, and it is thus possible to minimize the absorbed photons while keeping the trap potential useful by detuning the light field and turning up the power accordingly. To achieve this, the experiment uses a 1064 nm wavelength laser from Azurlight with a maximum power output of 20 W.

The intensity profile of a single Gaussian mode beam is given by

$$I(\mathbf{r}) = I_{\text{max}}(z)\mathrm{e}^{-\frac{2x^2}{w_x(z)^2}-\frac{2y^2}{w_y(z)^2}} \tag{5.6}$$

$$w_{x,y}(z) = w_{0,x,y}\sqrt{1+\frac{z^2}{R_{x,y}^2}} \tag{5.7}$$

$$R_{x,y} = \frac{\pi w_{0,x,y}^2}{\lambda}, \tag{5.8}$$

where z is the distance from a focus of the beam, $I_{\text{max}}(z)$ is the maximum intensity in the xy-plane at z, $w_{0,x,y}$ is the waist at the focus in the x- or y-direction, and $R_{x,y}$ is the Rayleigh length in the x- or y-direction. The relation between I_{max} and the power P is

$$P = \frac{\pi}{2}I_{\text{max}}w_x w_y. \tag{5.9}$$

The trap depth, ignoring the effects of gravity, is given by

$$U_{\text{max}} = \frac{\hbar\Gamma^2}{8\delta}\frac{I_{\text{max}}(0)}{I_{\text{sat}}} \propto \frac{P}{w_{0,x}w_{0,y}}. \tag{5.10}$$

The trap frequencies without gravity are given by

$$\omega_i^2 = \frac{1}{m}\frac{\partial^2}{\partial x_i^2}U\bigg|_{\mathbf{r}=0}, \quad x_i \in \{x, y, z\}. \tag{5.11}$$

The trap frequency along z can usually be ignored, except for very tightly focused beams. In the x- and y-direction

$$\omega_{x,y} = \sqrt{\frac{I_{\max}(0)}{m}}\frac{2}{w_{0,x,y}} \tag{5.12}$$

$$\propto \sqrt{\frac{P}{w_{0,x}w_{0,y}}}\frac{1}{w_{0,x,y}} \tag{5.13}$$

$$= \sqrt{U_{\max}}\frac{1}{w_{0,x,y}}. \tag{5.14}$$

The square root contains the increased trap frequency caused directly by more power in the beam and indirectly by higher intensity due to a more tightly focused beam. The second part contains the effect of larger derivatives due to smaller waists. It can be seen that the trap frequencies scale much more strongly with waists than with power.

The trap depth and the trap frequencies scale differently with power and beam waists. It is therefore possible to obtain BEC at different trap frequencies by varying the beam waist. A larger beam waist means BEC will first occur at a lower trap frequency.

5.2.3.2 Crossed dipole traps

A crossed dipole trap (cODT) is formed when two dipole beams cross at some angle. It results in a more confined trap because each beam can trap along the other beam's weak axis. In this experiment, the two dipole beams cross orthogonally at their foci. The beam propagating along direction z is denoted by superscripts z, and likewise along x. The vertical direction is along y.

The potential is a sum of two single-beam potentials. The trap depth and trap frequencies are

$$U_{\mathrm{max}} \propto \frac{P^x}{w^x_{0,y} w^x_{0,z}} + \frac{P^z}{w^z_{0,x} w^z_{0,y}} \tag{5.15}$$

$$\omega_x \propto \sqrt{\frac{P^z}{w^z_{0,x} w^z_{0,y}} \frac{1}{w^z_{0,x}}} \tag{5.16}$$

$$\omega_y \propto \sqrt{\frac{P^x}{\left(w^x_{0,y}\right)^3 w^x_{0,z}} + \frac{P^z}{w^z_{0,x} \left(w^z_{0,x}\right)^3}} \tag{5.17}$$

$$\omega_z \propto \sqrt{\frac{P^x}{w^x_{0,y} w^x_{0,z}} \frac{1}{w^x_{0,z}}}. \tag{5.18}$$

In case both powers are equal $P = P^{x,z}$ and all waists are equal $w_0 = w^{x,z}_{0,x,y,z}$ then

$$U_{\mathrm{max}} \propto \frac{2P}{w_0^2} \tag{5.19}$$

$$\omega_{x,z} \propto \frac{\sqrt{P}}{w_0^2} \tag{5.20}$$

$$\omega_y \propto \frac{\sqrt{2P}}{w_0^2}. \tag{5.21}$$

The y-direction provides the strongest confinement because it is trapped by both beams.

5.2.3.3 Gravitational sag

The gravitational potential $U_{\mathrm{grav}} = mgy$ should be added to the trapping potentials to describe the real potential experienced by the atoms. One consequence is a shallower trap depth because the atoms can now be pulled out of the trap by gravity, and another is a reduction in the trap frequency in the vertical direction y. Furthermore, the trap center is pulled downwards to the point where $\partial_y U_{\text{dipole}} = -mg$. This position is dependent on the dipole beam powers, and as a result, the cloud starts to move down (or ”sag”) when the powers are reduced. It also complicates the calculations of the trap frequencies, which have to be estimated around this position and not $\mathbf{r} = 0$.

The trap stops working entirely below a critical power $P_{\text{crit}} > 0$ where $\max\Big(-\partial_y U_{\text{dipole}}\Big) \leq mg$, and the trapping forces of the beams are no longer able to hold the atoms against gravity. When the power is slightly larger than P_{crit}, the potential becomes significantly anharmonic.

5.2.3.4 Hybrid traps

A hybrid trap consists of a single dipole beam and a trapping quadrupole (QP) magnetic field, described in Section 5.4.1. The beam is centered below the QP center to avoid Majorana spin flips.

Define the vertical offset position of the beam as $y = 0$. The QP potential U_{QP} from eq. 5.26 is then

$$U_{\text{QP}} = \hbar g_F m_F I B' \sqrt{x^2 + 4(y - d_y)^2 + z^2}. \tag{5.22}$$

In a correctly designed hybrid trap, the atoms will not get near the QP center and $x, y, z \ll d_y$. Then, the potential is approximately

$$U_{\text{QP}} \approx 2\hbar g_F m_F I B' d_y \left(1 + \frac{x^2}{8d_y^2} + \frac{y^2}{2d_y^2} + \frac{z^2}{8d_y^2} - \frac{y}{d_y}\right). \tag{5.23}$$

This is a harmonic potential plus a constant force upwards, which produces a new harmonic potential around the minimum. The trap frequencies are typically much smaller than those of a single dipole beam in the transverse directions, so the most important addition to the system is a trapping potential along the beam. Additionally, the

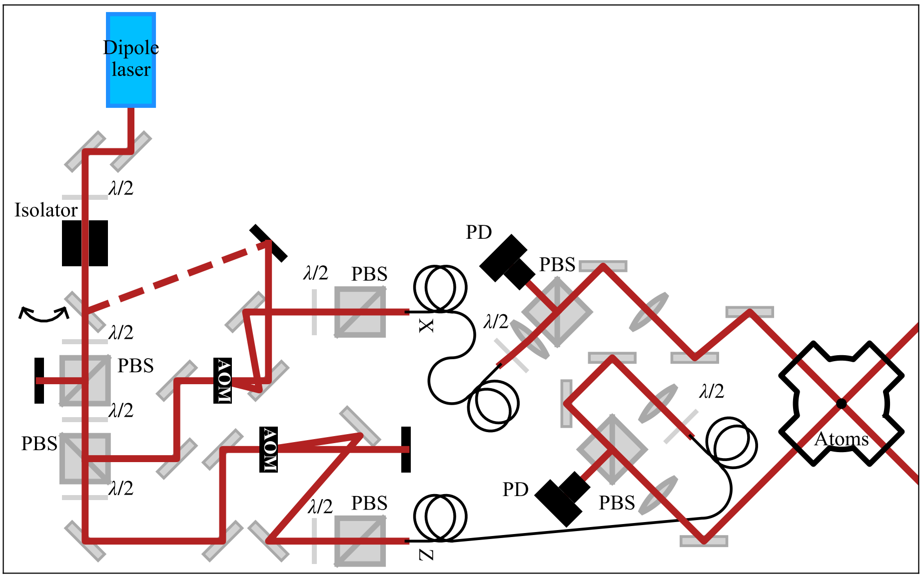

Figure 5.10: The dipole laser setup. It consists of two parts, the left half and the right half, separated by high-power fibers. This is to disconnect any work done on each part. The left part controls the power and frequency of the dipole beams. The X-beam (right) uses the -1'st order of the AOM, and the Z-beam (left) uses the $+1$'st order of the AOM. The right part directs and focuses the dipole beams to achieve a cODT. It also measures the dipole powers in each arm after the trap (not shown) for power stabilization.

upwards pointing force can be used to offset the effect of gravitational sag. During evaporation in a hybrid trap, it is important that the vertical QP force is at most slightly less than gravity to allow the evaporated atoms to escape the QP trap.

5.2.3.5 Dipole setup

The dipole setup is responsible for producing the two laser beams used for dipole trapping and is illustrated in Figure 5.10. The left half of the setup in the figure prepares the laser light and couples it into two high-power HC-1060 fibers from NKT Photonics. Depending on the trap configuration used, the laser outputs somewhere between 10 W and 20 W light with wavelength 1064 nm.

After the optical isolator, a flipper mirror is used to turn on or off the laser light completely. This is to avoid using shutters that would either absorb or reflect the high-power beam. After the flipper mirror, a PBS is used to turn down the power going to the rest of the setup without turning down the output power of the laser itself, which can alter the beam mode slightly. This allows the setup to be precisely optimized without the risk of damaging the equipment or personnel.

A second PBS is used to split the beam into two paths, which will eventually lead to the X- or Z-direction in the science chamber. Each beam passes through an AOM. These are used for power control and fast switching. The X-direction beam is shifted down in frequency, and the Z-direction beam is shifted up to avoid low-frequency beating between the beams when they cross.

The beams are led to the experimental setup (right side) by two optical fibers with mode field diameter $6.7 \pm 1\,\mu$m. The optical fibers separate the left half of the setup from the right half of the setup, such that optimizations can be performed without misaligning the cODT. Both paths after the fibers and up to the science chamber contain a cleaning PBS (preceded by a $\lambda/2$ wave plate), two lenses, a power monitoring photo diode, and mirrors to direct the beams into the trap. The PBS in the X-direction is rotated 90°around the beam axis, which ensures the two directions have orthogonal polarizations to minimize beating between the two arms further. The first lenses after the fibers are exchangeable and are used to determine the final beam waist around the atoms.

The setup currently have an $f = 50\,\mathrm{mm}$ outcoupling lens in the X-direction and an $f = 60\,\mathrm{mm}$ outcoupling lens in the Z-direction, both with a diameter of 1". The final lenses are always $f = 500\,\mathrm{mm}$ lenses with a 2" diameter and are placed on translational stages to adjust the beam focus positions. The vacuum chamber and the magnetic transport system limit the focus lengths of the last lenses to this value, but shorter focus lengths could make the setup more stable and achieve smaller waists.

Given these parameters, the waists of the beams in the trap are expected to be between 57 nm and 77 nm for the X-direction and 47 nm and 64 nm or the Z-direction. The uncertainty comes from the mode field diameter information provided by NKT. After measuring a collection of trap frequencies at different beam powers (shown in Figure 6.19 later in Section 6.2.1), the real waists were estimated to be

$$\begin{aligned} w_{0,y}^{x} &= 82.3\,\mu\text{m} \\ w_{0,z}^{x} &= 79.6\,\mu\text{m} \\ w_{0,x}^{z} &= 68.5\,\mu\text{m} \\ w_{0,y}^{z} &= 74.0\,\mu\text{m}. \end{aligned}$$

It is possible to increase the outcoupling lenses up to $f = 200\,\text{mm}$ with a 2" diameter, which can achieve a final waist of approximately $15\,\mu$m. BEC has been produced in such a setup, but the overlap between the two beams at this waist is unstable and has to be optimized daily. Such a small waist should only be used if necessary (see, for example, Section 5.2.3.6 about increasing cloud density).

After the science chamber, the dipole beams enter the detection setup (described in Section 5.6) with mirrors designed for 780 nm light, and much of the dipole power will therefore leak through the mirrors. A $\lambda/2$ wave plate, a PBS, a lens, and a photo diode are placed behind one of these mirrors in each arm to power stabilize the two beams (not shown in the figure). In order to achieve the lowest noise possible on the power stabilization, the photo diodes are always in the minimum gain, and hence they need the high power behind the 780 nm mirrors. The wave plate and PBS are used to adjust the power leading to the photo diodes such that the full dynamical range is in use. The lenses are used to focus the light onto the active area of the photo diodes and improve the stability of the setup when small adjustments to the beams are made. Furthermore, the electrical circuit of the power stabilization setup is separated from the power monitor setup (the photo diodes shown in the figure) to decrease the noise even more.

As a result, the power in the beams can be varied between 20 mW and 7 W without experiencing significant noise. A small waist configuration requires a better power stabilization setup than larger waist configurations because the power has to be decreased further to reach a sufficiently low trap depth. The current setup is able to function for waists down to 15 nm.

5.2.3.6 Producing high-density clouds

One of the goals in the experiment is to test the effects of the light-induced dipole-dipole interaction occurring at relatively high atom densities, and it is therefore necessary to be able to maximize this. The BEC has a much higher peak density than the thermal component of the cloud, so it is mostly interesting to look at the peak density of the BEC. The Thomas-Fermi model from Section 2.1.3.1 shows that the peak density ρ_{peak} is

$$\rho_{\mathrm{peak}} \propto N_0^{2/5}\overline{\omega}^{6/5} \propto N_0^{2/5} P^{3/5} \frac{1}{w_0^{12/5}}, \tag{5.24}$$

where the second proportionality was found by using the equations for the trap frequency in an ideal cODT without gravity (eq. 5.20-5.21). This model is too simple to predict the exact peak density of clouds, mainly due to the estimate of the trap frequencies. It is, however, a very important result in understanding which parameters to tweak to obtain the highest possible peak density.

There are three possible parameters to optimize: condensed atom number, dipole beam power, and dipole beam waist. The atom numbers in a BEC add very little to the peak density, and new atoms are instead added to the outer part of the cloud, increasing the radius. The power adds only a bit more than the number of atoms, and it has to be ramped down to a certain value during evaporation to obtain BEC. It can, however, be increased again adiabatically, which will turn some of the condensed atoms into thermal atoms, but result in an overall increase in density. Finally, the beam waist, which has a very strong impact on the peak density, scales the density with a power greater than 2. If the goal of the experiment is to produce high-density clouds, the beam waist should be minimized. In reality, the scaling will be slightly worse because the number of condensed atoms that can be produced in a tight trap, and the necessary power, are lower.

A typical experiment have a beam waist on the order of 100 μm to catch more atoms in the cODT [79], but some experiments work with waists to an order of magnitude of 1 μm [58–60]. On this scale, the Rayleigh length is so short that a single beam can trap in all three directions. Such a small waist requires a special high-NA setup.

The limit of our current experimental setup is 15 μm waists due to the transport coils, which prevents a close positioning of the final focusing lens to the atoms. Because of the 500 mm distance from the focusing lens to the focus, the drift of the beam is relatively large and starts to matter at this waist, limiting the stability of the experiment. My advice would therefore be to rethink the current dipole beam setup if experiments with small beam waists are to be performed again in the future.

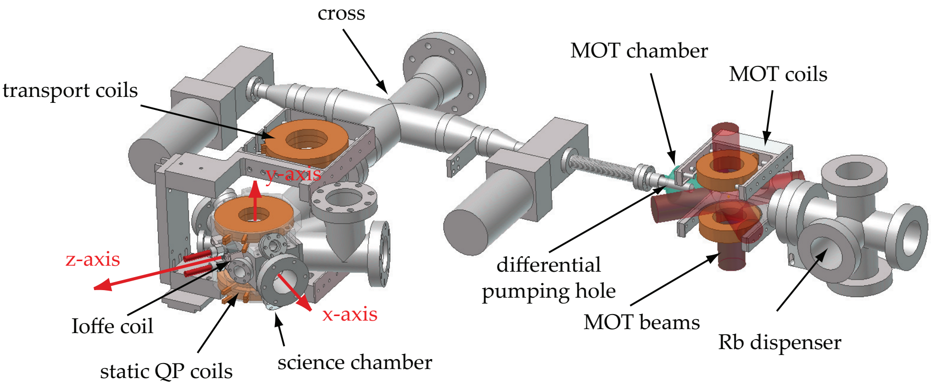

Figure 5.11: The original design of the L-shaped vacuum chamber [74], which is largely unmodified to this date except for the removal of the Ioffe coil. The design is made such that a second atomic species can be added in the cross, although this has not yet been achieved.

5.3 Vacuum chamber

The spine of the setup is the vacuum chamber shown in Figure 5.11 in which all experiments take place. It dates back to the beginning of the experiment and has not been modified significantly since then. It will therefore not be described in notable detail here, and I will instead refer to the early PhD projects [73, 74] in the group.

The vacuum chamber is split into two parts: a MOT chamber and a science chamber. They differ in pressure and are separated by a long, thin tube called a differential tube. Since the pressure in both chambers is so low, it is possible to maintain a pressure gradient between them without the need for valves. With the current equipment, the pressure in the science chamber is not detectable, and the pressure in the MOT chamber is 10^{-10}Torr to 10^{-8}Torr depending on the dispenser settings. The MOT plus the transport through the MOT chamber to the science chamber works most optimally around

$1-3\,\mathrm{nTorr}$.

The science chamber is surrounded by three sets of Helmholtz coils (called shim-X, Y, Z coils) to control the magnetic field background in the chamber. There is also a set of Helmholtz coils around the MOT chamber (called the MOT pump coils and not shown on the figure) to provide a magnetic field background during optical pumping. Crucially, the MOT does not feature any shim coils for magnetic field cancellation, and this has proved to be a major disadvantage in improving the MOT because it prevents efficient optical molasses from being implemented. A working optical molasses could make the cloud cold enough for a direct transfer to the cODT after transportation to the science chamber, although an intermediate step of RF-cooling could still be beneficial, and could in principle cut up to 8 s or about 30% of the entire sequence duration as well as retaining more atoms since some of the evaporation steps could be skipped. Efficient sub-Doppler cooling requires a low background magnetic field, however, and is thus not obtainable in the current setup. This is discussed more in Section 5.4.2.2.

The setup also features three sets of anti-Helmholtz coils. The coils used for the MOT (called the MOT coils), which can also be used to generate a high field for a QP trap, are placed on a transport stage that can move the coils from the MOT to the cross. A second set of coils (called the transport coils) can move from the cross to the science chamber. Here, the final set of coils (called the QP coils) is positioned and cannot be moved. The atoms are transferred from the transport coils to the QP coils for stability. This was initially imperative, as the original setup used a QUIC trap configuration (first designed by T. Esslinger et al. [96]) to avoid Majorana loss. Since the experiment now uses a cODT, which is relatively insensitive to the position of the QP coils, it might be possible to skip the transfer to the final set of coils. This could save the sequence about one half of a second or roughly 2%, and we have so far not deemed it worthy to pursue.

The MOT transport coils must align well with the differential pumping tube to minimize atom loss during the transport. The number of lost atoms after transport can be measured by moving the MOT coils back to the MOT again, turning on the cooling light, and comparing the fluorescence signal immediately after the light is turned on with the signal just before the MOT was terminated and the atoms transferred to the QP potential for transport. The recapture was measured in this way at different lengths of transport, and the result can be seen in Figure 5.12. The differential tube begins around 20k motor steps and is indicated with a vertical line. Note that the exact values depend on how much the MOT is loaded and the pressure in the MOT chamber. This data is taken to represent the most standard setup in the experiment. While the atoms are still in the ”high” pressure MOT chamber, the lifetime is relatively short. This is fitted with an exponential curve and indicates that 91% of the atoms are caught initially, and 76% of the atoms make it to the differential tube. The recapture decreases significantly just at the beginning of the tube, leaving 30% of the initial number of atoms in the trap (a reduction of 60% from just before the tube), making the tube the biggest loss of atoms during transportation. The atoms making it through the tube will also make it to the science chamber with ultra-high vacuum and have a very long lifetime.

Because a QUIC trap moves the potential minimum from the QP center, the vacuum chamber is designed with the QP coils placed off-center of the vacuum windows. The Ioffe coil was placed around one of the Z-direction windows, so the off-center position of the cloud in the current traps is worse for the X-direction. This results in a smaller NA and potentially worse imaging.

The off-center problem can be fixed in two ways. One is to simply detach the QP coils and position them correctly. Another is to move the QP center with shim coils just before transferring to the hybrid trap, which can then be centered on the windows. The second option

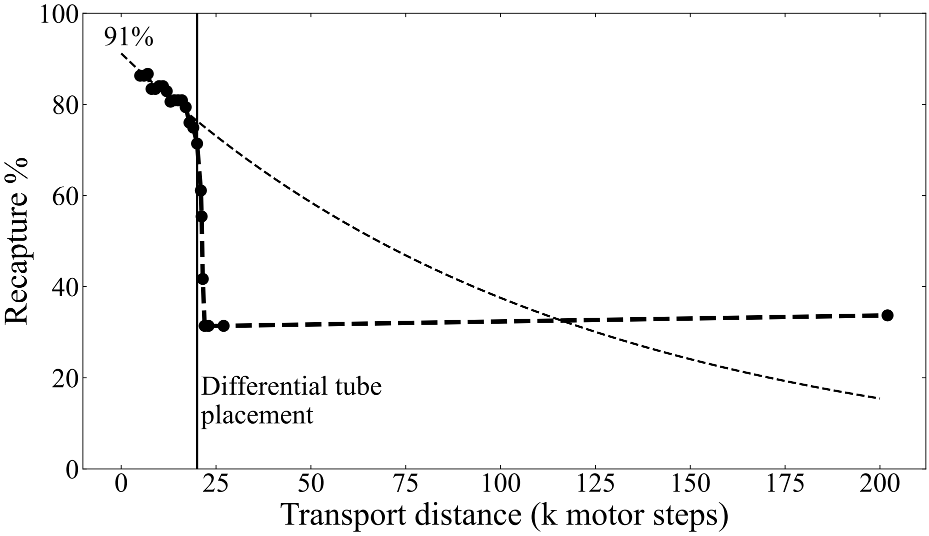

Figure 5.12: Recapture rate in the MOT as a function of transport distance away from the MOT. There is a short lifetime in the MOT chamber and a very long lifetime after the differential tube. The tube itself cuts about 60% of the atoms. This can be improved by achieving efficient molasses before transport.

is less severe, but requires accurate control of the shim current as the QP current is ramped. It has not been deemed necessary to correct for this yet.

5.4 Magnetic traps

Using magnetic fields to manipulate atoms is an essential tool in the experiment. The experiment has used three different kinds of magnetic traps: the quadrupole trap, the quadrupole Ioffe configuration (QUIC) trap, which is not used anymore, and the magneto-optical trap. The quadrupole trap and magneto-optical trap are described below.

5.4.1 Quadrupole traps

A magnetic quadrupole (QP) trap works because of the Zeeman shift in a magnetic field $\mathbf{B}$ given by $\Delta E = \hbar g_F m_F |\mathbf{B}|$, where g_F is the Landé g-factor in units of angular frequency per magnetic field strength and m_F is the projection of the magnetic dipoles along the local magnetic field direction. For a QP field produced by anti-Helmholtz coils stacked vertically (in the y-direction) with a current I, the B field is

$$\mathbf{B} = IB' \begin{pmatrix} x \\ 2y \\ z \end{pmatrix}, \tag{5.25}$$

and the potential U is

$$U = \hbar g_F m_F I B' \sqrt{x^2 + 4y^2 + z^2}, \tag{5.26}$$

where B' is a setup-specific constant which can be found by measuring the current for which atoms are exactly levitating against gravity. The projection m_F will adiabatically follow the magnetic field while the atoms are moving around, as long as they do not pass too close by the center, where the magnetic field direction flips quickly. The flipping of m_F states of those atoms is called Majorana spin flips, and prevents the use of QP magnetic traps for ultracold clouds on the order of micro Kelvins in temperature. QP traps are, however, fine for warmer clouds and are some of the simplest and most stable traps used for cold gases.

The QP center where $\mathbf{B} = 0$ can be moved around by adding a constant background B-field $\mathbf{B}_{\text{back}}$ with shim-coils. Denote the components of $\mathbf{B}_{\text{back}}$ as B_x, B_y, and B_z. Then the potential becomes

$$U = \hbar g_F m_F \sqrt{(IB'x + B_x)^2 + (2IB'y + B_y)^2 + (IB'z + B_z)^2}. \tag{5.27}$$

This can be used for moving around the cloud in a QP or hybrid trap. It can also be used for placing the QP center at a specific position relative to a cODT and then pulling the cloud in that direction, e.g., for measuring trap frequencies in one direction at a time.

The coils used to generate QP fields are made of relatively few windings compared to the high fields needed to have low inductance and fast response. This means the required current is high and generates a lot of heat. The coils are therefore made of copper tubes with cooling water running on the inside. However, the maximum current (400 A) should still not be on for long since some of the electrical circuit is not cooled.

The MOT coils and the QP coils have IGBT switches for fast turn-off. Because of the inductance from the coils, fast switching can generate large voltages. The setup, therefore, contains flyback diodes for protection.

5.4.2 Magneto-optical trap

The magneto-optical trap (MOT) is what initially catches and cools the atoms. A MOT can catch enough ^{87}Rb in a background vapor cell to produce BEC without having to slow down the atoms first. Figure 5.13 shows the key components in a MOT. It consists of a set of anti-Helmholtz coils to produce a weak QP field, which is not strong enough to capture atoms on its own, and six laser beams in counter-propagating pairs. All laser beams have σ^--polarization with respect to the magnetic field on the incoming side of that beam, which means the circular polarization for the vertical beams going through the coils is opposite from the horizontal ones. It also means the polarization of each beam on the outgoing side has turned to σ^+ because the magnetic field has inverted [97].

The incoming light fields are red detuned (specifically by 22 MHz from the $|F=2\rangle \rightarrow |F=3\rangle$ transition in this experiment), which means σ^- light has a higher probability of absorption than σ^+ light. The result is that the incoming beam will always push the atoms more towards the center than the outgoing beam will push them away, resulting in a trapping force. The beams also give rise to a cooling mechanism, described in more detail in Section 5.4.2.2.

For the MOT to work optimally, there has to be enough background vapor pressure for atoms to enter the trapping volume. However, increased pressure results in more background collisions, both in the MOT and during the transport from the MOT chamber to the science chamber. A MOT chamber pressure of $1 - 3\,\mathrm{nTorr}$ works well, while a pressure above $4\,\mathrm{nTorr}$ starts to be problematic. The background pressure is provided by a continuously running dispenser.

Some of the atoms fall to the $|F = 1\rangle$ ground state after decay and have to be repumped back into the MOT cycle by a separate laser. The direction and polarization of this laser are not important and are simply combined with one of the cooling beams in a polarization cleaning cube.

The number of atoms in the MOT is monitored by the fluorescence of the cloud by a photodiode. When the voltages reach a certain value, the experiment is triggered. This assures a stable number of atoms in each experiment.

Previously, the MOT was able to load up to a maximum voltage of $1\,\mathrm{V}$ while a trigger was typically set to $0.8\,\mathrm{V}$. It took about $30 - 60\,\mathrm{s}$ to load the MOT to this level consistently. With the addition of the distribution board, which gives both more total power and more stability than the previous setup, a complete refurbishment of the optical setup around the MOT and the cooling arm of the master laser setup, a change of dispenser, a new repump laser setup using a Toptica laser instead of home build one, and careful consideration of the optimal MOT chamber pressure the MOT is now able to load to a maximum of $3\,\mathrm{V}$ and the trigger is typically set to $1.5\,\mathrm{V}$ which can be consistently loaded in less than $10\,\mathrm{s}$. The experimental sequence has been modified to allow the MOT to start loading before the previous experiment is complete, which results in an essentially instant loading time.

The MOT setup is equipped with two cameras placed orthogonally to each other. These can be used for both monitoring and optimizing the MOT. For instance, aiming of the MOT beams can be achieved by placing an iris on the beam telescopes to work with a small MOT beam and then optimizing the MOT cloud size with feedback from the cameras. They can also be used to achieve beam power balance and, in principle, for optimizing optical molasses. If the MOT works sub-optimally, the cameras can be a precious tool for revealing the underlying cause.

Normally, the power in all six laser beams should be roughly equal. However, the center of the MOT is different than the center of the QP field at high current because of the presence of a magnetic background field, which is not being nulled in our experiment. If the MOT is loaded at the low current offset QP center and then transferred to the high current QP center for transport, the cloud will gain significant energy. It is therefore best to unbalance the beam powers such that the MOT cloud loads on top of the high current QP center, even if this results in a worse MOT. This also prevents the implementation of an optical molasses, which requires balanced beam powers; however, an efficient optical molasses also requires a nulled background magnetic field and would not be possible in the current configuration anyway.

5.4.2.1 CMOT

When atoms decay in the MOT, they emit photons which can be reabsorbed by the surrounding atoms. The consequence is a constant pressure in the cloud produced by the light, called light pressure. This makes the size of the cloud in the MOT much larger than the equivalent size of a cloud with the same temperature in a typical QP trap. If the atoms are directly transferred to a QP trap from the MOT, they will thus start with a very high potential energy, resulting in

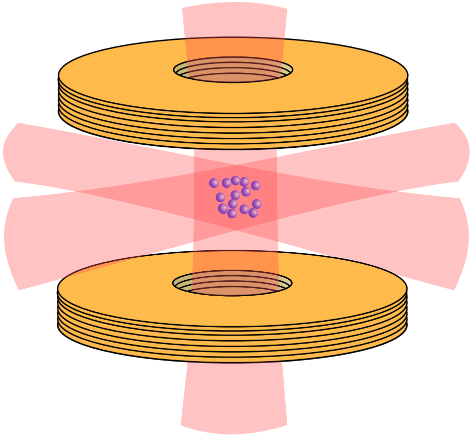

Figure 5.13: Illustration of the core setup of a MOT, borrowed from [78]. There are six beams, counter-propagating in pairs. All beams are σ^+-polarized with respect to the magnetic field on their incoming side.

significant warm-up. To avoid this, a short period of compression in the MOT is added before transferring the atoms to the QP potential, called a compressed MOT (CMOT) stage.

Compression can be achieved in three ways. The light pressure can be reduced by lowering the repump power and letting the atoms fall into a dark state ($|F = 1\rangle$). It can also be reduced by decreasing the cooling light scattering rate, which can be obtained by detuning more or less power. Finally, the compression can be accomplished by ramping up the QP coil current while the MOT is still on. The trap configuration in a CMOT is often not able to trap the atoms. The phase must therefore be short enough to keep the atoms in the trap region, and long enough to achieve compression.

Some implementations of a CMOT will result in heating during the compression. However, contrary to a QP trap, the CMOT phase can be followed by a molasses phase (see Section 5.4.2.2 below) for re-cooling, with the combined result of a compressed cold cloud. Even without a molasses phase, the cloud after a CMOT is typically colder than the heating from a QP catch without a CMOT, because some amount of cooling light is on during the CMOT.

Two different CMOT implementations have been used in this experiment. The original CMOT lasted 35 ms during which the QP current was ramped from 14 A to 18 A and the repump power was set from maximum power (about 7 mW) to about 10 μW during all of the CMOT. Neither the cooling power nor the frequency was changed. This type of CMOT is most optimally paired with a molasses phase, which was, however, not implemented in our experiment.

Since a molasses phase cannot be implemented without magnetic field nulling around the MOT, a different CMOT phase was found to be more optimal. This is the CMOT still in use today. It works by combining some of the mechanics from a molasses phase into the CMOT phase and is inspired by this thesis [98]. Perhaps counter-intuitively, in the new CMOT phase, the QP current was ramped down from 14 A to 4 A, which starts to mimic the zero magnetic field conditions during molasses. The repump power was kept at maximum, and instead the cooling light detuning was ramped continuously from the MOT detuning of −22 MHz to −55 MHz. The cooling light power was kept constant in the program, but the efficiency of the double-pass AOM setup over such a frequency span decreases, and the cooling light power was thus also decreased. Keeping the power to the maximum possible value was found to be best. This CMOT sequence obtains a compressed cloud by reduced light pressure and even better cooling than the previously implemented sequence.

In a particular experiment, just before making the switch from the first to the second CMOT sequence, the setup produced 230 k atoms with 100 k being condensed. Without changes to the rest of the experimental sequence, the switch from the first to the second CMOT sequence changed the result to a cloud with 450 k atoms and 230 k condensed atoms. The increased gain in performance is expected to come from the better cooling effects in the second CMOT sequence. The first CMOT sequence paired with a molasses phase could potentially work even better.

5.4.2.2 Optical molasses and sub-Doppler cooling

An optical molasses works similarly to a MOT, except there are no QP magnetic fields present. Atoms in a molasses experience a viscous force from the laser beams because of red- and blue-shifting due to their motion, with the consequence that atoms are more likely to absorb photons with the opposite direction of momentum than their own than photons with the same direction of momentum. Without the trapping effects of a QP field, however, the atoms will always propagate out of the laser beams if given enough time. A molasses phase, similar to a CMOT phase, should thus be kept at a relatively low duration in the order of milliseconds. As it turns out, the true mechanics of optical molasses are quite complicated, and more details can be found here [62, 99, 100].

In a simple model, the lowest temperature reached by optical molasses is called the Doppler limit, and is caused by the random effects of absorbing and emitting photons, leading to a random walk motion of a certain energy. But optical molasses turns out to be one of a few rare cases where the result in practice is much better than the predicted outcome, and the real lower temperature limit lies much below the Doppler limit. This is caused by many different effects, all collected under the term sub-Doppler cooling, and all of which rely on some form of light field polarization gradient throughout the molasses caused by the counter-propagating beams. Sub-Doppler cooling is also discussed in the citations above.

If sub-Doppler cooling can be achieved in the experiment, the atoms can be cooled far enough that they can, in principle, be transferred to a cODT without the need for pre-cooling in a QP trap. This can lead to a significant gain in the speed of the experiment since cooling in the QP trap is the slowest part of the sequence. It would also result in more atoms because of the skipped evaporation steps. Finally, it would result in a much better transfer of atoms through the differential tube because of the smaller cloud size during transportation due to the lower temperature.

Unfortunately, sub-Doppler cooling is only possible in a very low background magnetic field. The difference in Zeeman shift of m_F-states must be much smaller than the difference in dressed state energy shift of those states for sub-Doppler cooling to work, which requires a background magnetic field in the order of a few mGauss and thus necessitates magnetic field nulling coils. This is not yet implemented in our setup and prevents efficient sub-Doppler cooling from being achievable. A nulled magnetic field would have the additional benefit of the QP field center at low current being aligned with the center at high current.

5.5 Experiment control setup

The main control of the experiment is through a Delphi-based program called ECS (experimental control system), built mainly by Henrik Kjær Andersen when the experiment was initially constructed. It is described in detail in his PhD thesis [73], and has remained relatively unchanged since then. It currently provides 48 digital output channels with 1 μs resolution, 32 analog output channels with 50 μs resolution, 16 digital inputs channels where the trigger level can be adjusted on 4 of them, and a 10 kHz to 70 MHz RF frequency generator. The goal in the near future is to replace ECS with a new control system provided by ARTIQ, and ECS is thus not be described in much detail here.

The two computers used to control the setup are described in Section 5.5.1, the control of a FlexDDS synthesizer, used to generate MW frequencies and phase-lock reference frequencies, is described in Section 5.5.2, and the use of an oscilloscope from PicoScope to acquire SPD data is described in Section 5.6.2.2.

5.5.1 Computers

The ECS program is run on an old computer termed the ”ECS computer”. It is paired up with a second computer labeled the ”main computer”, which handles communication with data servers and a set of instruments in the experiment, namely two Andor EMCCD cameras for imaging of the clouds, two cameras to monitor the MOT, the lock control box for the master laser and optionally the coupling laser, the FlexDDS synthesizer for MW and beat frequency reference signals, the PicoScope for SPD data acquisition, and up to 7 motorized mirrors from New Focus where currently only one is in use for the cODT. It is also used for any immediate data analysis performed in the lab. The specs of the main computer are purposefully excessive since it is meant to run the whole experiment when ARTIQ is implemented, and should hold for many years. Whenever possible, Python is used as the programming language in the experiment.

There are two cloud storage systems in use, called ERDA, provided locally by Aarhus University, and ”the O-drive” provided by Microsoft. ERDA is a new and cheaper cloud storage system and was meant to replace the O-drive completely, but it has proven unstable so far and is used mostly for backup of large data files, while the O-drive is still implemented for important files used on a day-to-day basis.

The script ”main.py” on the main computer can detect when a new run file is created by the ECS computer via Ethernet connection. Because the ECS computer runs on Windows XP and the main computer runs on Windows 11, it is necessary to disable all security protocols of the main computer in order to allow access to the ECS computer’s shared folder. The security settings have to be updated each time the computer is restarted, which is accomplished with a batch script.

When a new ECS file is detected, it is copied to the main computer, and any necessary values are read from the run file and passed on to the different instruments. The main computer prepares all time-sensitive instruments in a state that requires triggers from ECS to continue in order to keep real-time synchronization with the experimental sequence. After a run is completed, run files and data are backed up on ERDA. If main.py is terminated, the last set of data is also backed up.

5.5.2 FlexDDS

MW frequencies for atom manipulation and reference frequencies for the OPLL are generated by a FlexDDS-NG-DUAL from Wieserlabs Electronics, which is controlled via USB connection by the main computer. The pySerial library is used to maintain the connection through Python. It is a new addition to the experiment, and the way it is used in our experiment is described in detail in this section.

A schematic of the FlexDDS is shown in Figure 5.14. It has a 10 MHz reference clock input, which is provided to the entire laboratory by an atomic clock. It also has an FPGA and two DDS cores. The FPGA has two DDS command processors (dcp's) which control the DDS's. Commands given to the FlexDDS will be processed and stored here and have to be actively pushed to the DDS Cores before anything updates on the output. The DDS cores are AD9910, and the manual can be found here [102].

To connect to the FlexDDS, use the lines of code below:

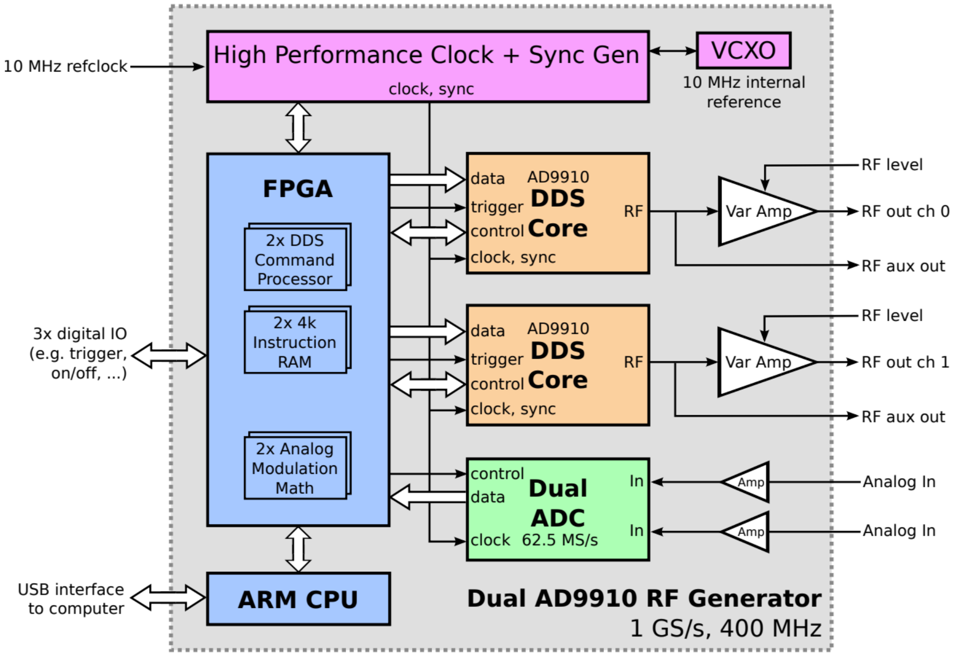

Figure 5.14: The FlexDDS interior structure from the manual [101]. An FPGA handles the communication with the main computer over USB and the control of the two DDS Cores. A 10 MHz reference input from an atomic clock is provided.

```
import serial
import time

ser = serial.Serial()
ser.port = "COM3"
ser.baudrate = 9600
ser.timeout = 1
time.sleep(1)
ser.open()
```

The object `ser` will be carried around to all other FlexDDS functions in Python. If a connection is already established, it is important to close that first. One way to do this is in the Spyder editor by going to ”Tools”→”Run”→”General settings” and choosing only the ”Run in console’s namespace instead of an empty one”-option and then running these lines of code before line 3 above:

```
try:
  ser.close()
except:
  pass
```

Writing to the FlexDDS is done by `ser.write(command)`, where `command` is a binary encoded string. A program is a combination of many commands. Each line in the `command` should end with `b"\r\n"` for immediate execution, where `b` makes the string encoding binary. The command is then stored on the FPGA. Updates to the actual output of the FPGA happen when this command is sent:

```
"dcp [CHAN] update:u[OPTIONS]\r\n"
```

Here `[CHAN]` refers to the channel and can be either 0, 1, or empty, in which case both channels are updated, and `[OPTIONS]` can be used to specify further options to the update. The FPGA runs through all the commands when given the line `"dcp start"`, which have to be at the end of any program.

One of the most important ways to control the FlexDDS is by updating the control function register (CFR) pins, where CFR2 (Table 19 in [102]) is often the one used. It is done by the following command:

```
"dcp [CHAN] spi:[REG]=[VAL]\r\n"
```

This updates the register `[REG]` for channel `[CHAN]` to the values `[VAL]`, which depends on the register being updated.

The FlexDDS can control frequency, phase, and amplitude of the two output channels. Since phase control is not being used in the experiment, it will not be described here. In order to control amplitude, it has to be activated first by pin 24 in CFR2, but it is also not used. A frequency f is given in terms of a frequency tuning word (FTW) which is defined by

$$\mathrm{FTW} = \mathrm{round}\left(2^{32}\frac{f}{f_{\mathrm{sysclk}}}\right), \quad f_{\mathrm{sysclk}} = 1\,\mathrm{GHz}. \tag{5.28}$$

It is sent to the FlexDDS as an 8-digit hexadecimal. A constant 100 MHz frequency on channel 0 is set by the single tone profile (stp0) register and the CFR2, like this:

```
amplitude = "3fff"
phase = "0000"
frequency = "1999999a"

command =\
f"""
dcp 0 spi:stp0=0x{amplitude}{phase}{frequency}\r
dcp 0 spi:CFR2=0x01000080\r
dcp 0 update:u\r
dcp start\r\n
"""

ser.write(command.encode("utf-8"))
```

The values for `amplitude` and `phase` are not important. The hexadecimal `0x01000080` changes the pins of CFR2 according to Table 19 in [102] to constant frequency mode.

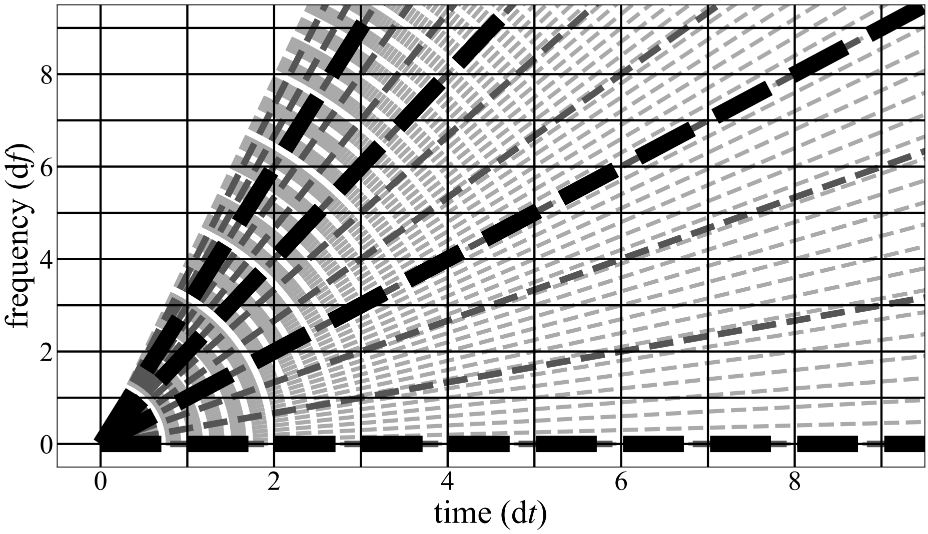

Figure 5.15: The four slowest possible sweeps for temporal steps sizes 1dt (thick, black), 3dt (medium, dark gray), and 20dt (small, light gray). The grid shows the size of one minimum step. A larger temporal step size gives a better resolution in sweep speed but rougher step sizes.

5.5.2.1 Frequency sweeps

A frequency sweep on the FlexDDS is given in terms of the step sizes of the sweep, in time and frequency space, and the limiting frequencies f_{min} and f_{max}. The step sizes are given as integers, N_t and N_f for time and frequency, respectively, in units of the smallest step sizes possible for time and frequency, called dt and df. The temporal step size dt is 4 ns and the frequency step size df is listed as 0.23 Hz. However, the frequency step size can vary slightly from device to device. By measuring the output signal from the FlexDDS on an oscilloscope during a sweep and comparing the duration of the sweep with the expected duration, the step size was calculated to be 0.2331 Hz on our

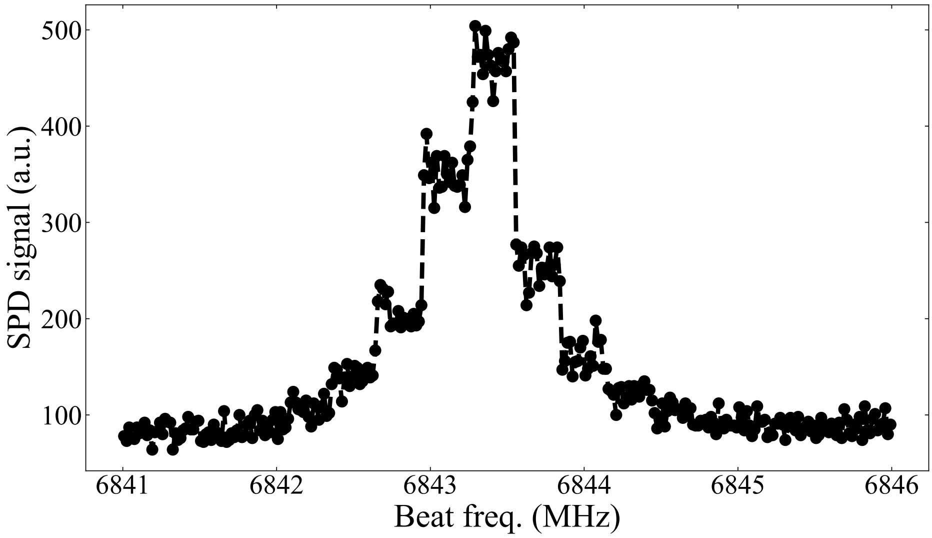

Figure 5.16: A time step size of $10^4\mathrm{d}t$ while scanning over an EIT peak in dark field. The step sizes of the reference frequency from the FlexDDS can be seen in the spectral data. The beat lock works too fast to be seen.

device.

In most sweeps, it is not possible to use the smallest possible step size to get the finest resolution. The reason why is illustrated in Figure 5.15 where the four slowest sweep speeds have been plotted for temporal step sizes $N_t = 1, 3, 20$. A fine time resolution (small temporal step size) results in a poor sweep rate resolution. Instead of choosing the best time resolution possible, it is necessary to strike a balance between the sweep rate resolution and the step size during the sweep.

Figure 5.16 shows an example of an EIT peak measurement with the SPD (the setup of which is described in Section 5.6.2) in dark field, where the temporal resolution of the FlexDDS was set too low. The step size of the frequency sweep is clearly visible in the signal, and the current resolution of the FlexDDS is set to be ten times finer.

There are algorithms to accomplish the best possible balance; however, in this experiment, the main interest is to measure spectroscopies. It is therefore not necessary to control the sweep rate perfectly, as long as the exact rate is well-known. A constant temporal step size of $10^3 \mathrm{d}t$ was therefore chosen, which means the frequency will be updated each $4\,\mu$s. The frequency step size $N_f \mathrm{d}f$ is then calculated as

$$N_f = \text{floor}\left(\frac{(f_{\text{max}} - f_{\text{min}})N_t \mathrm{d}t}{\mathrm{d}f T}\right), \tag{5.29}$$

where T is the duration of the sweep. The duration and frequency span will not be exactly the entered value in ECS, but the program will return the exact values from the FlexDDS, which can be used to analyze the spectra.

A sweeping frequency in channel 0 starting after a trigger in BNC input channel A from 100 MHz to 200 MHz over 100 ms is set by the following code:

```
3 frequency_low = "1999999a"
4 frequency_high = "33333333"
5 sweep_direction = "+"
6 time_step = "03e8"
7 frequency_step = "00004308"

8 command =\
9 f"""
10 dcp 0 spi:DRL=0x{frequency_high}{frequency_low}\r
11 dcp 0 spi:DRSS=0x{frequency_step}{frequency_step}\r
12 dcp 0 spi:DRR=0x{time_step}{time_step}\r
13 dcp 0 spi:CFR2=0x80080\r
14 dcp 0 wait::BNC_IN_A_RISING\r
15 dcp 0 update:u{sweep_direction}d\r
16 dcp start\r\n
"""

ser.write(command.encode("utf-8"))
```

Here, DRL is the digital ramp limit register, DRSS is the digital ramp step size register, and DRR is the digital ramp rate register. The step size and ramp rate registers have two inputs for upwards and downwards sweep settings. CFR2 is now set to single ramp mode. A ramp accumulator keeps track of where the frequency is in the sweep, and in single ramp mode, when the next frequency step would have jumped outside the frequency bounds, the sweep stops, and the frequency stays at the final value. If a triangular sweep is desired, line 13 can be switched with the following line:

```
13 dcp 0 spi:CFR2=0x000e0080\r
```

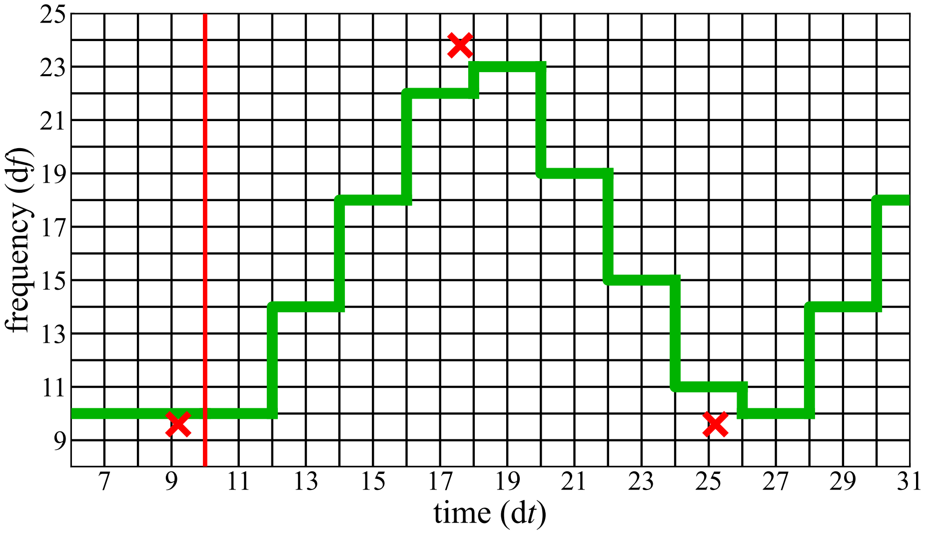

Figure 5.17: An example of a triangle sweep output from the FlexDDS when the time step size is 2dt. The grid shows the size of one minimum step. A single sweep follows the same pattern. The sweep tries to go through the red crosses. It begins at the vertical red line where, during the first step in time, the frequency remains constant. Each sweep always starts just above (for upwards sweeping) or just below (for downwards sweeping) the target frequency, no matter where it ended.

This resets the ramp accumulator when the sweep is completed and switches the direction. That kind of sweep will continue to go back and forth between the limiting values until terminated.

Figure 5.17 illustrates what the FlexDDS output frequency looks like over time through sweeps. The figure illustrates a triangle sweep, but a single sweep would simply follow the first upwards sweep on the plot. The sweep is trying to go through the red crosses with a temporal step size of 2dt. It will always start with one step at the initial frequency and thus takes four steps to accomplish this specific sweep, although it takes only three steps up or down in frequency. The upwards sweep starts at $f = 10\mathrm{d}f$ and ends at $f = 22\mathrm{d}f$; however, the downwards sweep starts at $f = 23\mathrm{d}f$ and ends at $f = 11\mathrm{d}f$ and is not the exact mirror of the upwards sweep. The actual step sizes are typically much smaller compared to the full sweep than illustrated here, and this effect typically does not matter.

5.5.2.2 Programs

The FlexDDS controls the MW frequency in the science chamber via channel 0. The frequency is mixed with a 7 dBm, 1.95 GHz signal from a Rohde & Schwarz signal generator, which also receives the 10 MHz reference frequency from the atomic clock, followed by a 2.27 GHz high pass filter, then amplified twice by $2 \times 9\,\mathrm{dB}$ ZX60-8008E-S+ amplifiers from Mini-Circuits and frequency multiplied by three by a ZX90-3-812-S+ from Mini-circuits. Two 45 dB ZFSWA2-63DR+ switches from Mini-circuits are used to turn on or off the MW signal after the multiplier, and this is followed by two amplifiers, a 34 dB ZVE-3W-83+ from Mini-Circuits and an $30 - 50\,\mathrm{dB}$ AM53-6.6-7-35-42 from Microwave Amplifiers LTD. The result is 40 dBm MW signal with frequency $3 \cdot (1.95\,\mathrm{GHz} + f_{\mathrm{FlexDDS}})$.

The most common use of the MW setup is to transfer atoms between $|F = 1\rangle$ and $|F = 2\rangle$. A very stable way to achieve this is by sweeping the frequency across the 6.835 GHz resonance[4] and adiabatically changing the dressed state from one of the F-states to the other. This is better than transferring with a constant frequency on resonance for a fixed amount of time, because it does not require precise knowledge about the resonance frequency, and the duration of the sweep is insignificant as long as the sweep is adiabatic.

The atoms will start in $|F = 2, m_F = 2\rangle$. To transfer the atoms to $|F = 1, m_F = 1\rangle$ (see Figure 5.3), the sweep should cross the res-

---

[4] Plus any changes caused by magnetic fields.

onance from below. This is because, if the sweep crosses the resonance from above, it will soon after cross the $|F = 1, m_F = 1\rangle \rightarrow |F = 2, m_f = 1\rangle$ resonance, and there is a risk of transferring some of the atoms back into $|F = 2\rangle$. If the atoms are to be transferred back to $|F = 2, m_F = 2\rangle$ from $|F = 1, m_F = 1\rangle$, the sweep should cross the resonance from above for the same reason. See Section 6.2.3 for a direct measurement of this.

A typical sequence prepares the atoms in the $|F = 2, m_F = 2\rangle$ state until reaching a cODT. The atoms can then be transferred to the magnetically untrappable $|F = 1, m_F = 1\rangle$ state with a MW sweep. This is beneficial because it removes the available internal energy, which can allow more inelastic collisions between atoms and thus a shorter lifetime. See Figure 5.18 for a quick measurement of this effect. It is also on atoms in this state that the probe laser is measuring. Depending on the goal of the experiment, the atoms have to be transferred back to $|F = 2\rangle$ for imaging. This can be done with a reverse MW sweep. Before the sweep begins, the MW frequency has been set to a constant value equal to the start of the sweep to avoid unpredictable, sudden jumps in the frequency.

The FlexDDS also controls the beat frequency reference signal described in Section 5.2.2.1. Because of the four times frequency divider after the fast diode, measuring the beat frequency between the probe and the coupling laser, and the 96 times frequency divider in the OPLL box, the reference frequency of the FlexDDS has to be the desired beat frequency divided by 384. If the coupling laser frequency is resonant and there is no magnetic field, the reference frequency should be 17.779 95 MHz for a resonant probe laser. During measurements, the reference frequency can be held or swept depending on the experiment. This will typically happen between two MW sweeps, such that the atoms are in $|F = 1\rangle$. Other experiments, such as those measuring the state preparation with MW directly, will work differently (see Section 6.2.3).

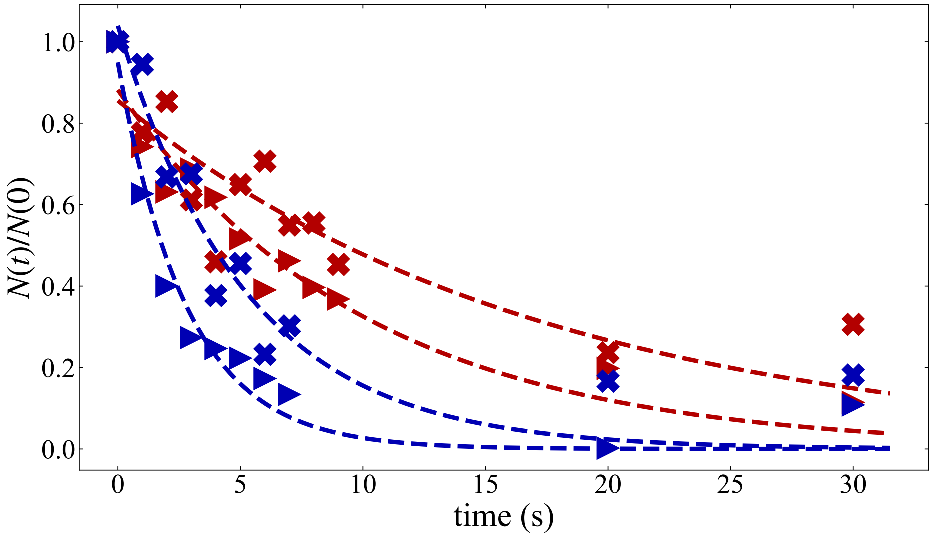

Figure 5.18: Lifetime of thermal (crosses) and condensed atoms (triangles) in $|F = 1\rangle$ (red) versus $|F = 2\rangle$ (blue). $|F = 1\rangle$ were transferred with a MW sweep and held for some time before being transferred back for imaging. $|F = 2\rangle$ were never transferred. $|F = 1\rangle$ thermal atoms have a lifetime of 17.2 s, $|F = 1\rangle$ condensed atoms have a lifetime of 10.0 s, $|F = 2\rangle$ thermal atoms have a lifetime of 5.3 s, and $|F = 2\rangle$ condensed atoms have a lifetime of 2.8 s.

The program of a typical experiment could be structured like this: 1) Initiate the MW and probe frequencies at constant values on channels 0 and 1, respectively, then wait with the update until an initiation trigger from ECS arrives to make sure the values have been read by the main computer. 2) Prepare the MW frequency for an upwards sweep and wait until the next trigger from ECS arrives, when the atoms are in the pure cODT trap and shim coil currents have rotated to a set value. 3) Prepare the probe for a sweep and wait for the next ECS trigger, which arrives when the cloud and magnetic fields are ready for measurement. 4) Sweep the remaining atoms back to $|F = 2\rangle$ with the same MW sweep but in the opposite direction at the next ECS trigger when shim coil currents have rotated back to the set value from the first MW sweep. 5) After imaging, the final ECS trigger sets the probe laser frequency to repump to allow for the MOT to start loading.

5.6 Detection

The experiment uses two different types of detection methods to probe the atoms for information. The first method is the well-established time-of-flight (TOF) imaging technique using cameras, and the second method is the newly introduced single-photon detector. The images of atoms were analyzed in great detail by the previous PhD student in the group, Toke Vibel [12, 38, 78], and I refer to that work for in-depth information on imaging.

The imaging optics is shown in Figure 5.19. The imaging light frequency and polarity are shown on an energy level diagram in Figure 5.2 and Figure 5.3, respectively, and goes from $|F = 2, m_F = 2\rangle$ to $|F' = 3, m'_F = 3\rangle$ with σ^+ polarization. There are two image directions. The X-direction imaging light comes from the top left and leads to the bottom camera. The Z-direction light comes from the bottom left and leads to the top camera. Probe light follows the same path as the imaging light in the Z direction, except it is much more attenuated and is reflected in a PBS towards the SPD, in between the two cameras, by flipping the $\lambda/2$ wave-plate just before the final PBS. Both cameras are the Andor iXon Ultra 897 EMCCD model, and the SPD is an SPDMH2 from Thorlabs.

5.6.1 Atom imaging

The X-direction is the main direction used for imaging the atoms. The imaging optics in this arm consists of four lenses, two mirrors,

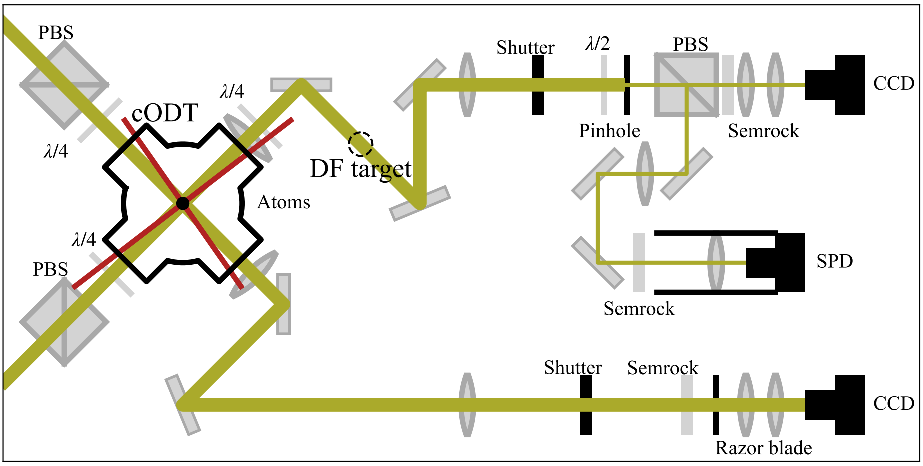

Figure 5.19: The imaging optics used for both the two cameras and the SPD. The bottom is for imaging the X-direction, and the top is for imaging the Z-direction and for the SPD. A flippable $\lambda/2$ wave-plate and a PBS can switch between the camera and SPD. A DF target can be inserted in the probe light focus in the Z-direction for dark field data acquisition. Both arms are preceded by polarization cleaning cubes. The dipole beams follow the same paths but at a small angle.

a large aperture shutter, a 780 nm optical filter (Semrock 780/12 nm BrightLine® single-band bandpass filter), and a razor blade. The Semrock filter removes any remnant of the dipole beam, which has already been reduced in power because the 780 nm mirrors leak most of the 1064 nm dipole light.

The placement of the four lenses is shown in Figure 5.20, taken from [78]. Ideally, the lenses should be placed in a $4f$ setup, where the first lens is one focal length away from the atoms, the distance between each consecutive pair of lenses is the sum of their focal lengths, and the distance from the final lens to the camera is one focal length. This would ensure that either the imaging light or the diffracted light from the atoms is collimated during the entire setup. Due to limitations in

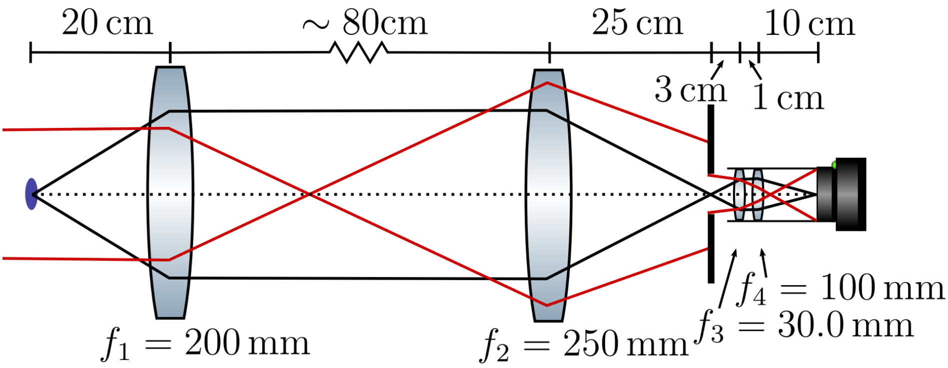

Figure 5.20: The optics in the X-imaging arm. Figure taken from [78]. Two lenses create an image of the atoms at a plane, which is being cut by a razor blade to avoid exposing half the region of the CCD camera, to allow for two images taken quickly in succession. Two more lenses focus the atom image onto the camera chip. From left to right, the focal lengths are 200 mm, 250 mm, 30 mm and 100 mm, which would give a magnification of 4.17, if all lenses were placed perfectly.

space, however, the distance between lens one and two is longer than this, and the distance between lens three and four is shorter. If the camera is placed correctly in focus, this does not disturb the image, but if the camera is slightly out of focus, the intensity of the imaging light relative to the diffracted light is either enhanced or decreased, resulting in defects in the *od* on top of the out-of-focus defects that always happen.

The calculated magnification from the setup is 4.170, which means the 16 μm wide pixels of the camera would give an image with 3.837 μm/px. However, if the lenses are not individually placed perfectly, even if the total setup is in focus, the magnification will differ slightly. It is therefore necessary to measure the magnification experimentally, which can be done by releasing the cloud, varying the TOF, and comparing the free fall with the expected acceleration of 9.82 $\mathrm{m/s^2}$ from gravity.

Figure 5.21 shows such a measurement, where both the x and

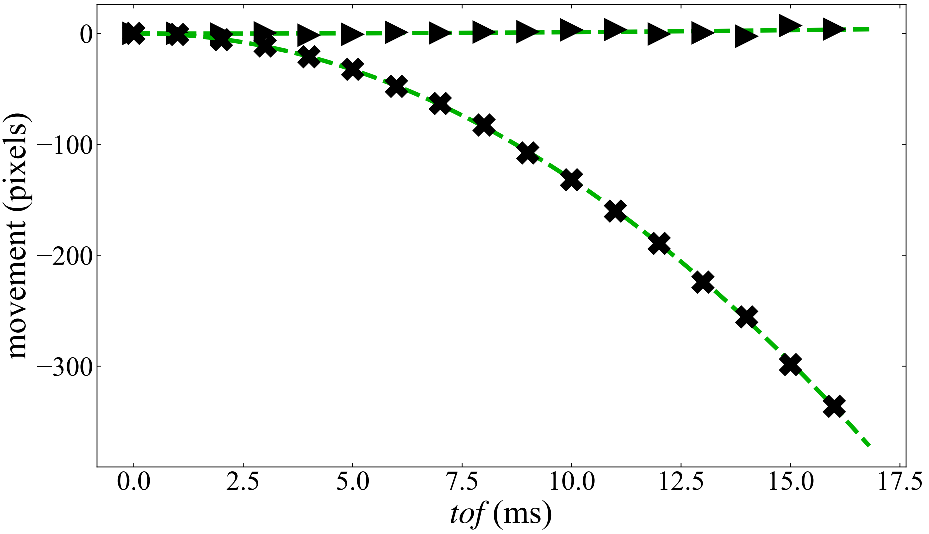

Figure 5.21: A measurement of the center position of a BEC after release as a function of TOF. The fitted acceleration is held against the expected value of $9.82\,\mathrm{m/s^2}$ and gives an in-trap pixel size of the X-camera of $3.723\,\mu\mathrm{m/px}$.

y components of the movement are tracked, in case the camera is not perfectly aligned vertically. The acceleration is then given by $a = \sqrt{a_x^2 + a_y^2}$, where $a_{x,y}$ is the acceleration in each direction. The resulting calibration was $3.723\,\mathrm{mm/px}$, which gives a magnification of 4.298.

Information about a cloud is inferred from the spatial distribution of optical density calculated with eq. 4.23 from Section 4.1.1.2. More details on absorption imaging can be found here [12, 37, 38, 57, 103–109]. It is necessary to measure the intensity of light with (I_{out} in eq. 4.23) and without[5] (I_{in} in eq. 4.23) atoms to calculate the *od*. This is accomplished by taking four images with light intensities $I_{1,2,3,4}$.

---

[5] The image without atoms can be omitted by applying deep learning to reconstruct the image without atoms and consequently reduce noise due to temporal

Image one is taken with the atoms in $|F = 2, m_F = 2\rangle$. In between image 1 and two, the atoms are pumped down to $|F = 1\rangle$ and the second image is thus taken ”without” atoms. Images 3 and 4 are taken without imaging light on for background subtraction. This gives

$$I_{\text{out}} = I_1 - I_3 \tag{5.30}$$

$$I_{\text{in}} = I_2 - I_4. \tag{5.31}$$

A correction to eq. 4.23 because of imperfections in light polarization and because of the multi-level structure of the atoms can be effectively modeled by scaling the saturation intensity by a factor α [37], which adds α as a scaling factor to the first term in eq. 4.23

$$\alpha \ln\left(\frac{I_{\text{in}}}{I_{\text{out}}}\right) + \frac{I_{\text{in}} - I_{\text{out}}}{I_{\text{sat}}} = \sigma_0 \tilde{\rho}, \tag{5.32}$$

In [78], following the method in [37], a value of $\alpha = 1.69$ was found and is used for analysis in this thesis.

In order to reduce fringes in the image, the time between image 1 and 2 and image 3 and 4 must be as short as possible, which is accomplished by blocking half the active area of the camera chip with a razor blade placed in the intermediate focus after the second lens. Two successive images can be taken by quickly exchanging the pixels of the bright and dark regions of the chip, with 340 μs between the images, which is the limit imposed by the camera. Uploading the image to the main computer and preparing for the next set of images takes about half a second, after which the background images are taken in the same way as images 1 and 2. To reduce the background light in the images, the final two lenses are situated in a tube, and to reduce the dark counts, the chip is cooled to −40 °C.

A typical TOF lasts 38 ms giving the cloud a speed of 0.37 μm/μs. Each image is exposed for 30 μs meaning the cloud will move 11 μm

---

variations in the imaging mode [110].

during the image which should be compared to a typical cloud size of $\gtrapprox 100\,\mu$m after expansion, and the result will be a slightly warmer temperature estimated from the vertical direction than from the horizontal direction.

The above analysis of images gives a grid of column densities which can then be fitted with equations from Chapter 2 to extract atom numbers and temperatures (see eq. 4.24-4.26). If the temperature is higher than T_c, the distribution quickly approaches the Gaussian distribution. Close to, but still above, T_c, the Bose-enhanced distribution must be used. When to switch between the two can be inferred by looking at the fits where a Gaussian distribution will not be able to follow the sharp rise in density at the center of clouds that are too cold for that approximation. Below T_c, either the Bose-enhanced distribution or the simulated expansion from Section 3 must be used for the thermal component, and the TF-distribution must be used for the condensed part of the cloud.

For a purely thermal cloud, the fit is performed once to extract approximate widths in both directions. The regions of interest (ROIs) are then updated, such that the cloud is fitted within a radius of $3w_{x,y}$. Outside a radius of $4w_{x,y}$, the average intensity of the image is calculated for both image 1 and 2. The intensity in image 2 is then scaled such that they match in average intensity in this region. This is necessary because the imaging light intensity might vary slightly between the two images. The column densities in the image are then calculated again, and the fit is performed once more to extract the values.

Fitting bimodal clouds is more complicated. Two methods can be applied depending on the desired outcome. The more stable and theoretically correct method of fitting uses eq. 4.27 for the thermal column density, which is a function of the temperature and the chemical potential, and eq. 4.26 for the BEC column density with the number of BEC atoms given by the chemical potential from eq. 2.13. The exact fitting procedure for this is described in Section 5.6.1.1 below. Alternatively, eq. 4.24 is used for the thermal column density, and the atom number and temperature are fitted. The number of condensed atoms is also fitted as an additional parameter. The fitting procedure is a bit more involved and is described in Section 5.6.1.2 below.

The first method is best to use if the goal is simply to extract the cloud parameters most precisely from the images. It is, however, more limited than the second method, which is necessary in several circumstances, such as when measuring any deviation from the semi-ideal model (e.g., atom number fluctuation measurements [20–22, 111, 112]) or when performing various tests in the experiment.

5.6.1.1 Bimodal distribution fit: method 1

The fitting parameters with this method are the position of the cloud, μ, T, and the two condensed radii $R_{x,y}$, and the thermal and condensed parts of the cloud are fitted together simultaneously. A preliminary fit is used to determine the thermal cloud widths to define the fitting ROI of a second fit. The two intensities in images 1 and 2 are scaled to match outside $4w_{x,y}$ of the cloud center, and the second fit is performed within $3w_{x,y}$. The number of thermal atoms comes from integrating the thermal part of the distribution, and the number of condensed atoms is given by the TF approximation.

5.6.1.2 Bimodal distribution fit: method 2

With this method [113], the fitting parameters are the position of the thermal cloud, the position of the condensed cloud, N_0, N_{th}, the widths of the thermal cloud $w_{x,y}$, and the two condensed radii $R_{x,y}$. The chemical potential is assumed to be 0. The widths of the thermal cloud are fitted only to the tail and subtracted from the measured column densities. The remaining column densities are then fitted to a TF-distribution. The first fit is used as a preliminary estimate of the size of the two components. Then, the two intensities in images 1 and 2 are scaled to match outside four times the thermal radii. A second thermal fit is then performed within three thermal radii and outside 1.2 TF-radii. The BEC is fitted within 1.2 TF-radii after the thermal part has been subtracted. The temperature can then be extracted from the thermal widths, giving a separate temperature in each direction. This method is especially sensitive to the chosen ROIs.

5.6.2 Single photon detection

Imaging ultracold atoms after TOF is a standard and well-tested technique, but it has the disadvantage of being destructive. For each data point, a new cloud will therefore have to be produced. Although there exist methods to image clouds with far-detuned light that are less destructive, these methods limit the frequencies used in the experiment and do not apply to measuring spectra. The frame rate of most cameras is also relatively low.

To be non-destructive, the measurement has to be taken in-trap, which requires high numerical aperture setups and careful imaging optics. An SPD works like a 1-pixel camera, which means there is a better signal-to-noise ratio at a given signal power. The ”frame rate” is limited only by the detector’s downtime; there is no need for high spatial resolution[6], and the better signal-to-noise ratio allows low enough power to non-destructively measure a cloud with resonant light.

There are four important advantages to implementing an SPD in the setup. Firstly, more data can be taken faster, which gives better statistics. Secondly, datasets can be taken on single clouds, which

---

[6] A low spatial resolution caused by numerical aperture will lead to loss of signal, instead.

eliminates the preparation noise from cloud to cloud, which is a big advantage. Previous achievements on atom number fluctuations measurements were mostly a matter of making this noise source less severe [12, 20, 36, 111, 112, 114]. Thirdly, taking in-trap low-intensity measurements avoids some of the complications in standard TOF imaging, such as imaging a moving cloud, modeling the expansion of the cloud, and the distorting effects of the light pressure from the imaging light. Finally, daily testing in the laboratory is accelerated significantly. Optimizing and measuring parameters, understanding new physics, and testing hypotheses can take minutes instead of hours.

Consider a data set like the one displayed in Figure 5.9, showing how an EIT spectrum reacts to 24 different coupling powers. This could be taken with cameras as well, and would require 4800 runs to obtain the same frequency resolution. With a runtime of 30 s per cloud, this results in 40 hours saved by using the SPD instead.

In the bottom right of the master laser setup in Figure 5.4, the probe laser is shown to be coupled into the imaging fibers. The power is directed to the Z fiber only, but could in principle also go through the X fiber in future experiments. While the imaging light is reflected off the final PBS, the probe light is transmitted through it, and this means the probe light will exit the fiber with orthogonal polarization to the imaging light. Because of this, there is a $\lambda/2$ flipper just after the fiber to select for either imaging or probe light. After this, a 10% sample of the power is measured by a diode, and the output is power stabilized. A shutter (which was not placed in the master laser setup for the Z direction, in contrast to the X direction) comes right after the power stabilization and allows the power to be continuously stabilized even if the signal to the atoms is turned off. Another flipper with an exchangeable amount of OD filters is placed just before the light enters the vacuum chamber. From here, the light follows the path from the bottom left of Figure 5.19. A PBS on the top right reflects the probe light to the SPD. A pinhole is placed just before the PBS in the intermediary image plane to sort away most of the background signal. The camera is used to place the pinhole correctly and in focus, and to judge the size of the pinhole needed.

Similar to the X imaging direction, there are four lenses in the Z direction. The first two lenses from the vacuum chamber have focal lengths 140 mm and 200 mm, respectively. The lenses in front of the camera are exchangeable, because the camera is used for different purposes requiring different magnifications, such as placing pinholes or imaging early clouds after arriving from the MOT. The two lenses in the SPD arm have focal lengths 200 mm and 40 mm. This achieves a demagnification, which is important to ensure the entire cloud is imaged on the $100\,\mu\text{m} \times 100\,\mu\text{m}$ active area of the SPD detector. It also leads to greater stability and allows the cloud to move around slightly before the image exits the active area. The final lens is only 1 ” while all other optical elements in the SPD arm are 2 ”. This lens will be upgraded to a 2" condenser lens in the near future. The first three mirrors and the second lens will be upgraded to 3 ” as well.

Just after the vacuum chamber, the dark field target can be inserted in the focus of the probe beam. The target has to be small enough to block a minimal amount of diffraction light, but large enough to block out a significant amount of the probe light. A back on the envelope calculation with a cloud width of $10\,\mu\text{m}$ and a probe beam mode width at the atoms of 5 mm shows that the amount of light allowed to pass to get the background signal down to the order of the actual signal is $\approx 10^{-6}$ or an ”OD” of 6. This is not an easy task because of the large non-Gaussian tails of power around any beam due to imperfections. Therefore, even in dark field mode, the pinhole is inserted to sort away the remaining background signal.

The optimal dark field target would be a small dot on a 3 ” window, but this has not been possible to obtain. Instead, a $500\,\mu\text{m}$ wide Allen key is used. The main loss of signal is not due to the Allen key, but instead to a limited numerical aperture, which is discussed further in

Section 5.6.2.1.

It is possible to combine the dark field setup with more intricate replacements of the current pinhole, such as various masks or extremely small pinholes, to select specific regions of interest in the cloud and enhance sensitivity. This will require precision in the placement of the pinhole relative to the intermediate cloud image, and any finite resolution effects would have to be taken into account.

In dark field, the probe beam must be larger than the cloud since the signal that ends up on the detector consists of the diffracted light going around the atoms. This can be stated in a different, but equivalent way. If the probe beam mode is comparable in size to the atomic cloud, it is not possible to sort the two modes efficiently in the Fourier plane. Consequently, a dark field setup must be constructed to be compatible with the desired target size. The current setup will not work for a cloud after TOF, which would require a larger probe beam and a smaller dark field target.

The SPD signal is a 15 ns long square pulse, with a minimum downtime of 35 ns after the pulse, which means the maximum detection rate of photons is 20 MHz. However, if this rate is detected, it means that a new photon must always be present just after the downtime is over, and this requires an incoming rate of photons much larger than that. In order to avoid overlap of photons, it is best to limit the detection rate to a few MHz and preferably less than a MHz. Section 5.6.2.2 discusses how the data is captured and analyzed in the current setup.

5.6.2.1 Numerical aperture effects

The first numerical aperture which can cut off the diffracted light is the vacuum chamber window and the first lens placed just behind it. The distance from the atoms is 140 mm and the diameter of the elements is 2 ”. It is very important that the imaging and probe light exits the window and the first lens is centered, because the diffracted light follows the same path. No further improvements can be made without rebuilding the vacuum chamber.

The diffracted light can be highly non-Gaussian in shape because of the phase shift through the cloud. It is therefore not necessarily collimated after the first lens. Throughout the rest of the setup, the light passes through more 2 ” diameter optics. Some of the mirrors are turned up to 45° which reduces the width by a factor $\sqrt{2}$.

Because the field mode depends both on the cloud and on the detuning of the probe light, the loss of signal will change between spectra and throughout the same spectrum, which can alter the physical parameters extracted. Simulations with Fourier optics by Ilja Zebergs from our group [115] suggest the biggest effect comes from the rest of the optical setup and not the vacuum window and first lens. This can be seen in Figure 5.22 which shows the percentage of diffraction light within a given diameter d from the center of the beam, denoted T, at the vacuum window on the left and at the dark field target, one focus length after the first lens, on the right. The second position is chosen because it is simplest to simulate the electric fields in the Fourier plane. The parameters used for the simulation mimic the values in a typical experiment and are $T = 152\,\mathrm{nK}$, $N_0 = 200\,\mathrm{k}$, $N_{\mathrm{th}} = 100\,\mathrm{k}$, and $f_{x,y,z} = 120, 120, 60\,\mathrm{Hz}$.

The transmission is 0 at $d = 0$, as expected, and then quickly rises as the radius is increased. Even at 1 ” diameter, essentially all the signal makes it through the vacuum window. The part that makes it through a 2 ” diameter at the dark field target must be close to collimated, and it is assumed that the values here resemble the efficiencies of the entire setup. The loss here is worse than at the vacuum windows, indicating that the diffracted light has diverged further.

Detunings that are in a range close to resonance produce high-resolution rings in the diffraction field (see Figure 4.6), and this explains why there is a dip in T at intermediary detunings. It is the plan to upgrade all optical elements up to the shutter, where the light is

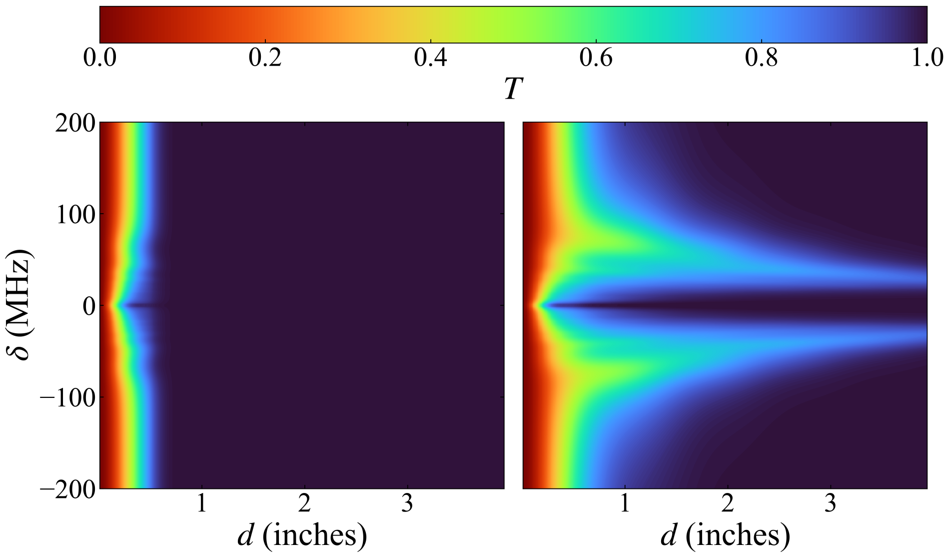

Figure 5.22: A simulation of the dark field signal fraction making it through a circular numerical aperture placed at the vacuum windows (left) and at the probe field focus (right) for various aperture diameters d. Even if the cloud is in focus, the dark field will not be collimated, and more signal is lost during the optical setup, indicated by the bigger losses on the right. Especially at medium detunings in the range $10\,\mathrm{MHz} \lesssim |\delta| \lesssim 150\,\mathrm{MHz}$, which coincides with the region with the most rings in the phase-shift of the probe light over the cross section of the cloud. The losses are frequency dependent and can mess with the measured spectra.

being focused again, from 2 ” to 3 ” diameter elements.

Because of the dependence on detuning, the Fourier optics simulations are included in the analysis of the spectra by simulating the loss at a calibrated radius and scaling the spectra accordingly (more on that in Section 6.1).

5.6.2.2 PicoScope

The square pulses from the SPD are picked up by a PicoScope 3000A, which measures the voltage from the SPD and sends it to the main computer as a 32-bit integer, which then extracts the arrival time of each photon from the rising edges in the signal. After the time stamp extraction, the voltage trace is deleted. To detect the 15 ns pulses, the PicoScope records the voltage every 4 ns and the maximum recording duration is 512 ms, limited by PicoScope's memory. Extracting the time stamps on the device and then transferring to the main computer (a time-to-digital converter, TDC) would be far more efficient, but it has not been possible to achieve this yet. The problem is the niche of detecting counts with nanosecond resolution over, preferably, many seconds.

A different option would be to detect several counts in pre-determined time bins; however, this results in less freedom in re-binning the data during analysis. It might still be the best option for very long traces (more than half a second), while PicoScope could be used for higher resolution and shorter measurements.

The communication between PicoScope and the main computer happens over USB 3.0. USB 2.0 is too slow to transfer the voltage traces within the cycle time of the experiment. The setup of PicoScope also happens over the serial connection, and similar to the FlexDDS, it is all handled by the main.py script. PicoScope is set to run in block mode, which means a single trace of a fixed duration is taken when triggered by ECS. The duration is read by the main computer from the ECS run files.

5.6.2.3 Analyzing SPD data

There are two main ways to analyze a list of time stamps[7]: binning and maximum likelihood estimation (MLE). Binning is the simplest and most efficient, but it is also prone to errors, while MLE analyzes directly on the time stamp data without modifying the data, but it is not as easily visualized. MLE is typically only necessary with low signals.

Binning works by dividing the entire duration of acquisition into intervals and counting how many events happened in each interval. This can then be plotted to visualize data and fits to data. The problem with binning is that too few intervals will wash out fine structures in the data, while too many intervals will struggle with the discrete nature of counting. The Poisson noise in each bin is given by $\sigma = \sqrt{N}$ where N is the number of counts in the bin, and the goal is typically to make σ less than the amplitude of differences to be measured in the data. This requires a minimal number of counts in the data set and is not necessarily achievable.

MLE analysis works by fitting a probability distribution $P_B(x)$ to the data without binning, where x is a value in the domain of the time stamps, such as time or frequency, and $B = \{b_i\}$ is a set of parameters $b_1, b_2, ...$ such as linewidth or resonance. The distribution must be normalized

$$\int_D P_B \mathrm{d}x = 1, \tag{5.33}$$

where D is the domain on which the data was taken. A measure of probability[8] $\mathbb{P}_S$ of observing the data set $S = \{x_i\}$ can be calculated

---

[7] The word ”time stamp” here refers to a list of measurements in general. In this experiment, it will often be a list of frequencies of the probe laser at which a photon arrived.

[8] Technically, the probability of observing exactly any given dataset is zero because the domain is a continuous space.

as a function of B

$$\mathbb{P}_S(B) = \prod_i P_B(x_i). \tag{5.34}$$

This number becomes extremely small quickly, and it is more practical to use the logarithm of $\mathbb{P}_S$ as a measure

$$\log\left(\mathbb{P}_S(B)\right) = \sum_i \log\left(P_B(x_i)\right) \tag{5.35}$$

The fit is now composed of choosing a set of parameters B such that the value $\log\left(\mathbb{P}_S(B)\right)$ is maximized. This way of analyzing the data does not involve the decision about bin widths, and it is stable with much fewer counts than what is needed for binned data to behave nicely. A downside is that the function $P_B(x_i)$ has to be evaluated at each time stamp x_i, contrary to having to evaluate the function at each bin. This can be problematic for a large data set, which can contain millions to billions of counts. Luckily, MLE is only necessary for small data sets, while binning works perfectly for very large data sets. If the model to calculate P_B is heavy, it can be sped up by calculating the probability on a few values of x spaced over the entire interval and then interpolating between them.

MLE analysis can also be applied to images to reach a higher precision and sensitivity on extracted cloud parameters [116].

Chapter 6

Non-destructive measurements

Measuring spectroscopies with destructive imaging techniques is time-consuming because each frequency requires the production of a new cloud. It also introduces cloud preparation noise from shot to shot. It was therefore decided to build a complementary detection setup that could achieve many measurements on the same cloud. For such a setup to function, it must possess sufficient time resolution to perform measurements before significant decay of the cloud, and it must be non-destructive. This was achieved with the single photon detector setup (described in Section 5.6.2), and the results are presented in this chapter.

The chapter is separated into two categories of non-destructive measurements. The first is the measurement of spectroscopies presented in Section 6.1, and the second is a collection of other measurements made possible by non-destructive time-resolved data acquisition presented in Section 6.2.

6.1 Spectroscopic measurements

The measurement of spectra forms the core outcome of this thesis, providing not only intrinsically valuable results but also insight into the broader capabilities and limitations of the experimental setup. The results are discussed for both bright field and dark field, and illustrates the strengths and weaknesses of each method.

6.1.1 F'-state interference

Spectra for atom number and temperature estimation are taken mainly on the blue detuned side of the peak. This is to avoid interference with the $|F' = 1\rangle$ exited state, which can happen if there is a small population in the $|F = 1, m_F = 0\rangle$ state or if the polarization has a small σ^--component.

If the reason is an impurity in the atom populations, the peaks will reach comparable heights near the resonances where nearly all light is absorbed, while an impurity in the light polarization will have a relative peak height equal to the fraction of power in each polarization. Moreover, an impurity in the polarization would also be able to detect the $|F' = 0\rangle$-state, which would otherwise require contamination to the $|F = 1, m_F = -1\rangle$-state.

A broad dark field spectrum with about 300 k thermal atoms and a small condensed component is shown at the top of Figure 6.1, with the resonance of the two F'-states indicated. Since both peaks reach a similar height, and since the $|F' = 0\rangle$-state is not visible, the main reason behind the detectable $|F' = 1\rangle$-state is most likely a small population in the $|F = 1, m_F = 0\rangle$-state.

In between the two resonances is a clear interference, which happens because the phase shift from each transition has the opposite sign. The canceling of the signal is greatest closest to the $|F' = 1\rangle$ resonance

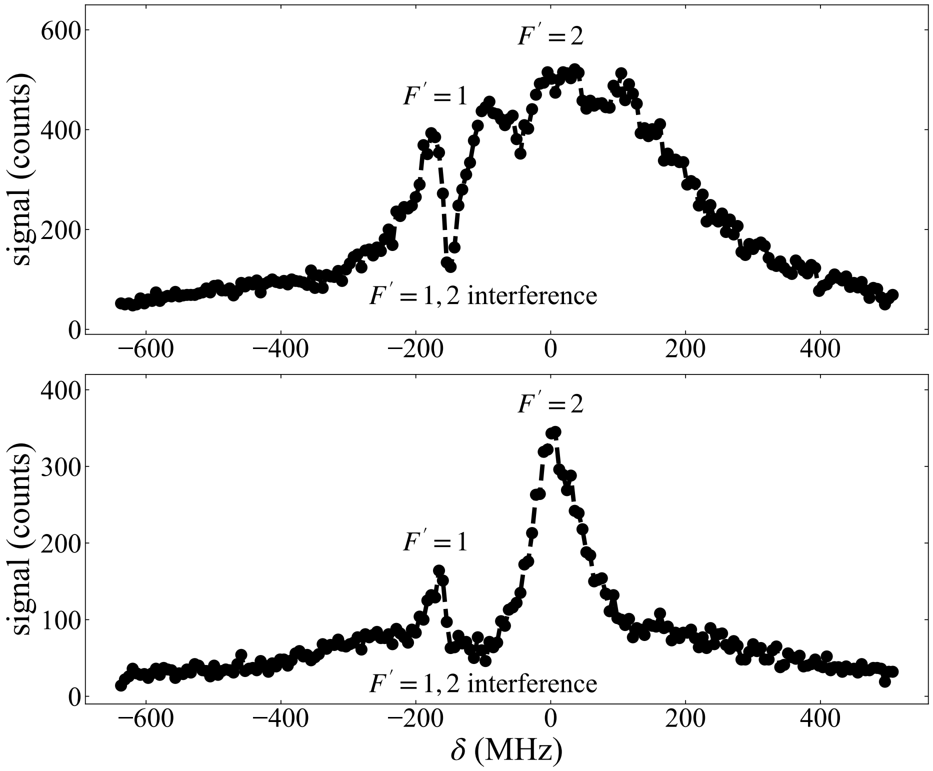

Figure 6.1: Broad dark field signals with some impurity from the $|F = 0\rangle \rightarrow |F' = 1\rangle$ transition for an almost pure thermal cloud (top) and an almost pure condensed cloud (bottom). An interference between the two transitions cancels the dark field signal in between and results in skewed peaks.

because the population in the $|F = 1, m_F = 0\rangle$-state is much smaller than the population in the correct $|F = 1, m_F = 1\rangle$-state.

The same scan with about 200 k condensed atoms and a small thermal component is shown at the bottom of Figure 6.1. The broad signal from the BEC and the large and narrow signal from thermal atoms are clearly distinguishable. The interference between the two transitions

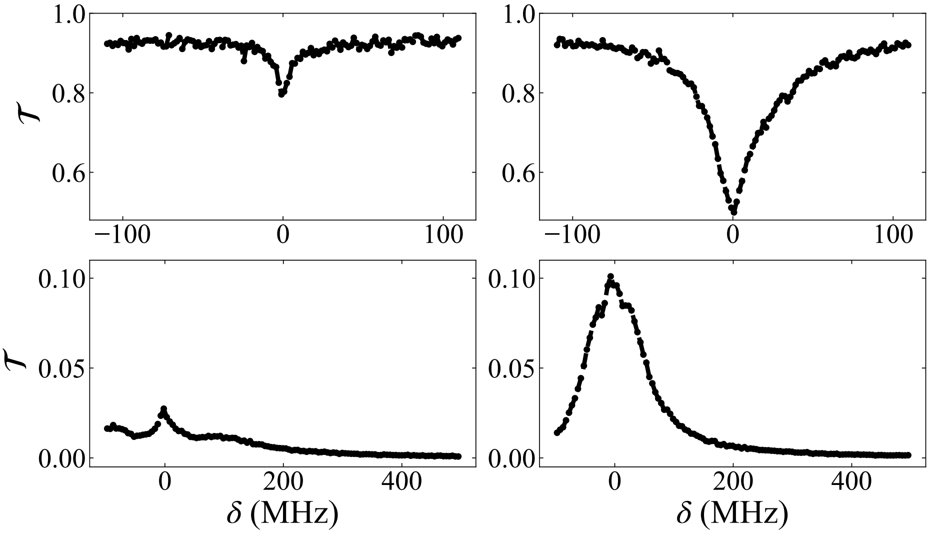

Figure 6.2: Four signals illustrating typical traces when measuring cloud parameters spectrally. Top left: $N_0 = 185\,\mathrm{k}$, $N_{\mathrm{th}} = 18\,\mathrm{k}$, and $T = 62\,\mathrm{nK}$. Top right: $N_0 = 146\,\mathrm{k}$, $N_{\mathrm{th}} = 290\,\mathrm{k}$, and $T = 239\,\mathrm{nK}$. Bottom left: $N_0 = 75\,\mathrm{k}$, $N_{\mathrm{th}} = 36\,\mathrm{k}$, and $T = 75\,\mathrm{nK}$. Bottom right: $N_0 = 0\,\mathrm{k}$, $N_{\mathrm{th}} = 244\,\mathrm{k}$, and $T = 395\,\mathrm{nK}$.

is still visible. More clearly seen here is the consequence of the blue detuned side of the $|F = 1\rangle \rightarrow |F' = 1\rangle$ being steeper than the red detuned side, and similarly but opposite for the $|F = 1\rangle \rightarrow |F' = 2\rangle$ peak. There are no other allowed transitions at $\delta > 0$ since interaction with the $|F' = 3\rangle$ state has $\Delta F = 2$. The steeper slope of the red detuned side of the main peak will always be visible in the datasets to some extent.

6.1.2 Cloud parameter measurements

The three cloud parameters N_0, N_{th}, and T can be extracted non-destructively from spectra by fitting the normalized transmission $\mathcal{T}$.

Moreover, the spectrally fitted parameters are not susceptible to error caused by the complicated cloud expansion dynamics discussed in Chapter 3.

Figure 6.2 shows four examples of data acquired according to the methods described in Section 6.1.3 for bright field and Section 6.1.4 for dark field. The top row is bright field results, and the bottom row is dark field results, while the left side is cold BECs and the right side is warmer clouds with many thermal atoms. The precise numbers are specified in the caption. Bright field results are an average of 20 runs, and dark field results are taken with a single run each. The pinhole was 75 μm in diameter for bright field and 150 μm for dark field.

The dip in the signals in bright field and the peak in the signals in dark field are clearly larger for hot clouds with more atoms, as expected. The bimodal structure of the spectra is most clearly visible in the cold dark field spectrum in the bottom left, with the narrow central peak around $\delta = 0$ being caused by the thermal component of the cloud and the much broader structure being caused by the BEC. The bimodal structure is less clear in the cold bright field spectrum in the top left, partly because of a limited NA and partly because of a worse signal-to-noise ratio.

The NA also causes the bright field signal to not approach $\mathcal{T} = 1$ until at much larger detunings than shown here. The "NA dips" in the bright field signal are comparable in width to the dark field signals. This is because the signal lost to the NA is, in fact, the diffracted electric field measured directly by the dark field setup. It is therefore necessary to take into account the NA aperture in the fitting model, which is discussed in more detail in the calibration sections for bright field and dark field below.

The NA also affects the dark field spectra, although it is less apparent. It causes the "dip" between the thermal peak and the shoulder at $\delta \approx 100\,\mathrm{MHz}$ in the bottom left plot to deepen, and the thermal peak becomes narrower. The effect on the extracted cloud parameters is less severe than for bright field, but it is still necessary to compensate for the NA in dark field.

For the extracted values to be accurate, a set of calibration parameters can be optimized across a subset of all datasets, called the calibration set, by comparing the results to those from TOF measurements, analyzed by the method described in Section 5.6.1.

A list of calibration parameters that have been considered is defined in Table 6.1. Each parameter can have different effects on the spectra depending on both cloud parameters and detuning. Some of the effects of different parameters will overlap, and choosing too many parameters can lead to over-fitting. It is not always immediately clear what effect each parameter has on the spectra. The best optimization parameters differs between BF and DF experiments. The choice of calibration parameters was done by repeating the analysis of the spectra with different calibrations and evaluating the impact and interplay of the parameters. The process is discussed further for BF and DF separately in their respective sections below (Section 6.1.3.1 and Section 6.1.4.1).

A calibration of $\Gamma(\rho)$ would be a direct measurement of the induced dipole-dipole interaction between atoms discussed in Section 4.2.3 and would by itself constitute an interesting result. It was found that calibration of the linewidth's dependence on density was susceptible to over-fitting, and even though the densities of the clouds are expected to lead to some broadening, it was decided not to calibrate it.

Future datasets could be taken with the sole purpose of calibrating $\Gamma(\rho)$, with a frequency sweep completely avoiding the thermal component of the spectrum near resonance and focusing specifically on the further detuned, BEC part of the spectrum, which has the highest spatial density. This would allow for much more data to be taken since the frequencies of the probe are less destructive in this region, and therefore ultimately a better calibration of $\Gamma(\rho)$.

The calculations of normalized transmission used as the fitting function to the spectra is done on a cylindrical grid with radial coordinate

Parameter	**Explanation**
Cross-section scaling factor α	An alpha calibration similar to [37].
Cross-section slope factor α'	The slope of a linear, spatially dependent alpha calibration similar to [38].
Power scaling factor β	A factor multiplied on the signal. It is not the efficiency of the setup, which has already been taken care of by using normalized transmission.
Aperture radius R_{NA}	The radius of the aperture used to estimate numerical aperture effects as discussed in Section 5.6.2.1. The aperture would be placed at the DF target.
Radius of in-trap image of the pinhole R^*_{ph}	Essentially the magnification M of the imaging setup since $R^*_{\mathrm{ph}} = R_{\mathrm{ph}}/M$ where R_{ph} is the real pinhole radius.
Resonance ω_0	The resonance frequency. Can either be calibrated or fitted.
$\Gamma(\rho)$	The local dependence of linewidth on the density of atoms. The function would be specified as a list of linewidths $\Gamma_i(\rho_i)$ at a list of densities ρ_i, and interpolation is performed in between.

Table 6.1: A list of possible calibration factors to optimize before extracting cloud parameters from spectra.

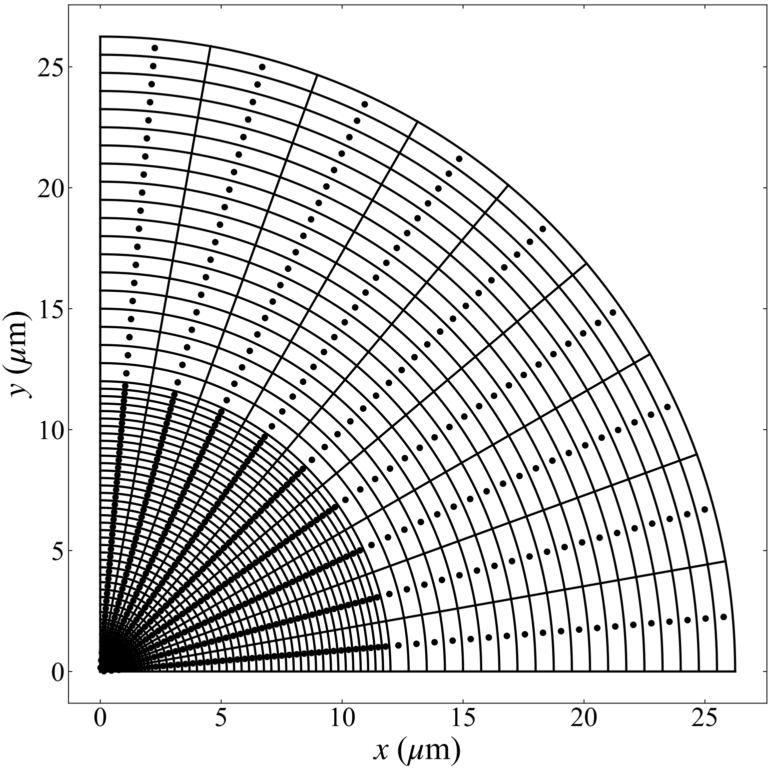

Figure 6.3: The cylindrical grid on which the cloud spectra are calculated. The figure shows only the points on the grid with $z = 0$, but there is a similar grid at each different value of z which takes the same values as r.

r, longitudinal coordinate z along the probe beam, and angle around the z axis θ, and only in the positive quadrant and within the in-trap image of the pinhole with radius R^*_{ph} ($0 \leq r \leq R^*_{\mathrm{ph}}$, $0 \leq \theta \leq \pi/2$, and $0 \leq z \leq R^*_{\mathrm{ph}}$). There are 9 angular points and 58 radial points. The values for z are the same as the values for r. The radial and longitudinal points are spaced closely within the first $12\,\mu$m, which is just outside the region with rings in the phase-shift for all clouds. This radius is kept constant to avoid technical problems with the minimizer function when fitting cloud spectra. Figure 6.3 illustrates the r, θ part of the grid, where the markers are the points used for calculating intensities by integrating through the cloud over the z-variable, and the areas surrounding them are used for integrating the intensities over r and θ to eventually get to $\mathcal{T}$.

The grid is chosen to strike a balance between complexity and efficiency. The simplest grid would be a Cartesian grid, but this is very inefficient since most of the atoms are tightly packed around the center, whereas a cylindrical grid automatically decreases the sample density as the radius increases. The most efficient grid would use ellipsoidal coordinates shaped by the trap frequencies with continuously varying sample density in the radial coordinates. This would, however, complicate the column density calculations. Depending on the speed necessary for future applications, it might be valuable to revisit the grid.

6.1.3 Bright field

The data acquisition sequence for bright field measurements is shown in Figure 6.4, which shows the SPD trace of an example in the top and a labeling of the different sections in the bottom. The sequence is a combination of SPD measurements and TOF measurements, which is possible because the SPD measurements are non-destructive. Before the sequence starts, the probe power is turned on and stabilized, but a shutter keeps it from entering the science chamber. The magnetic field rotates to 3.7 G along the probe propagation direction. The first step is the opening of the shutter, which can take a few milliseconds. Data is then acquired, typically for about 80 ms (the exact duration depends on the FlexDDS sweep and can vary by a few hundred microseconds) during what is called interval **A**. Then, the atoms are released in time of flight for $30 - 38\,\mathrm{ms}$, which varies for different clouds, with the

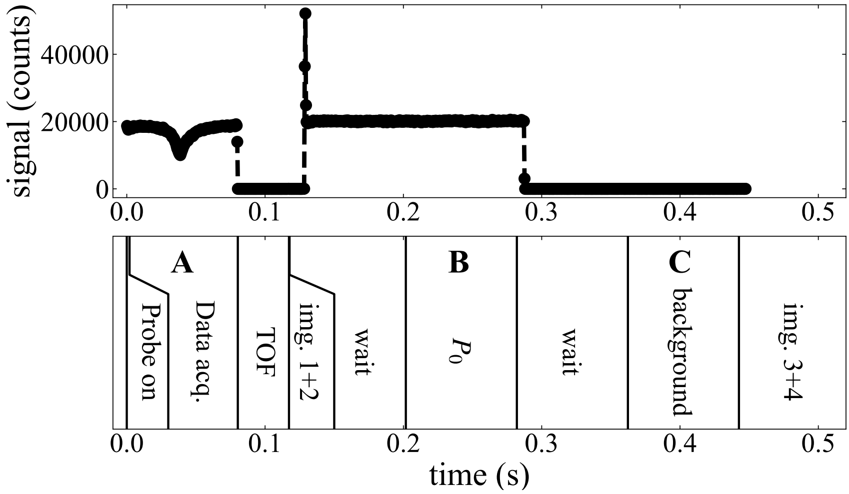

Figure 6.4: The data acquisition sequence for bright field measurements. The top shows an example of the SPD trace during the measurement, and the bottom labels the different sections of the sequence. Three sections are labeled **A**, **B**, and **C** and are the measurement of the spectrum of the cloud, the probe power measurement, and the background power measurement.

coldest ones needing the longest expansion time. Each pixel in an image must be kept below an OD of ≈ 3 to reduce errors in the atom number and temperature estimation, but also well above the noise of the background. During the TOF, the magnetic fields are rotated from the Z-direction along the probe beam to the X-direction along the imaging axis. The atoms are repumped from $|F=1\rangle$ to $|F=2\rangle$ in the last $500\,\mu$s before absorption imaging. The first two images of the imaging sequence (described in Section 5.6.1) follow the TOF. After the atoms are gone, the probe light and all other beams used during the first acquisition are turned on again to measure the probe intensity with the background signal. When the probe is turned on this time, the power stabilization has not yet set in, resulting in an initial burst of the maximum power accessible. This is used to check that the power is not saturated, in which case the power stabilization would not work. A waiting period is inserted to allow the power stabilization to reach a constant value, after which the background plus probe power is extracted. This is called interval **B**. The probe field is then turned off, but all other beams are left on, and the background is extracted after another wait time in what is called interval **C**. Finally, the third and fourth images of the imaging sequence, used to estimate the background in the images, are taken, during which all laser beams are turned off. The wait time at the beginning of intervals **B** and **C** can be shortened in future measurements if needed.

Because bright field data have a high background signal during data acquisition, it is necessary to accumulate the data over many runs. A bright field dataset is taken by repeating the same experiment multiple times and then using the TOF images to estimate the number of atoms in each run by summing over the column densities in each pixel within the thermal and BEC ROIs. The dataset is then sorted into subsets within which a maximum relative variation in atom number of ±10% is allowed. The data displayed in Figure 6.4 has on average 146 k BEC atoms, 290 k thermal atoms, and $T = 239\,\mathrm{nK}$, and the data is taken with a pinhole diameter of $75\,\mu\mathrm{m}$. The dataset includes 20 runs.

The signal during interval **A** is called P_{atom}, the signal during interval **B** is called P_0, and the final signal for background power during interval **C** is called P_{bg}. These three subsets of the dataset are used to calculate the normalized transmission

$$\mathcal{T} = \frac{P_{\mathrm{atom}} - \bar{P}_{\mathrm{bg}}}{\bar{P}_0 - \bar{P}_{\mathrm{bg}}}, \tag{6.1}$$

where the bar indicates an average. The normalized transmission of this particular dataset is shown in the top right panel in Figure 6.2.

6.1.3.1 Calibration of bright field

The calibration factors α and α' are always calibrated to the expected values and are therefore omitted. This can be explained by the low power of the probe [117]. The power factor β must be 1 for the tails of $\mathcal{T}$ to approach 1, which they must because the spectra are normalized to a measurement without atoms. It is therefore not used in bright field.

The aperture radius R_{NA} is necessary to calibrate the bright field spectra. The simulation of the effect caused by a finite aperture is done using the dark field signal, because it is the diffracted light that gets blocked. This gives a transmission of the dark field signal T_{DF} and must be translated into a transmission of the bright field signal T_{BF}. The consequence of a numerical aperture is a larger intermediate image of the cloud through the pinhole, and assuming the pinhole covers all of this image, the loss of energy from the blocked diffracted field must be the same as the loss of energy through the pinhole. This gives the relation

$$T_{\mathrm{BF}} = 1 - \Big(1 - T_{\mathrm{DF}}\Big)\frac{\mathcal{T}_{\mathrm{DF}}}{\mathcal{T}_{\mathrm{BF}}}. \tag{6.2}$$

The final fraction rescales the signal from dark field to bright field.

The R_{NA} parameter is not a perfect representation of the setup, and will function as an effective radius. It will be able to cover both real aperture effects and some of the approximations made in the model, such as the approximation analogous to the thin-lens approximation made in the propagation of light through the clouds, the negligence of bosonic effects in the light-atom interaction, or the assumption that Γ is independent on density.

The in-trap pinhole radius R^*_{ph} is also calibrated, and is a sensitive parameter in bright field. It scales the absorbed light relative to the light making it through the pinhole and would be equivalent to a power

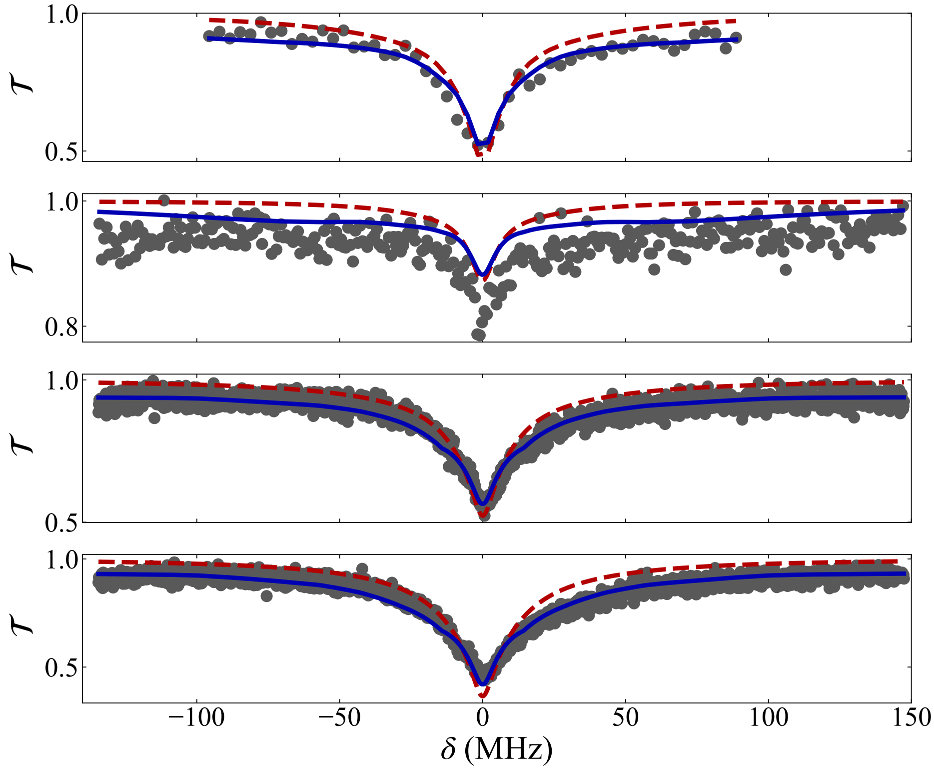

Figure 6.5: Calibration of bright field. The red dashed lines show the expected spectra with $R^*_{\text{ph}} = \frac{1}{2}\frac{140\,\text{mm}}{200\,\text{mm}} \cdot 75\,\text{mm} = 26.25\,\text{mm}$ and $R_{\text{NA}} = \infty$. The blue solid lines show the spectra after calibration of the two parameters such that all four measured spectra match the closest. The second panel from the top shows the coldest cloud and is not perfectly calibrated with the TOF-measured cloud parameters.

factor in dark field, assuming the pinhole covers the entire intermediate cloud image.

The resonance is calibrated once, and separately from the other calibration parameters, because there is no covariance between them.

Finally, $\Gamma(\rho)$ is not calibrated, even though a discrepancy from the 6.065 MHz probably occurs at the high densities in the experiment due to the light-induced dipole-dipole interaction. This is to avoid over-fitting.

The calibration is done on four clouds, by calculating the expected spectrum using TOF images and then varying the two parameters R_{NA} and R_{ph}^{*} to minimize the discrepancy between the expected and measured spectrum. The points are chosen to span different types of clouds, and to specifically contain one cloud with a high number of BEC to get a sensitive measure of the NA. The calibrated clouds are indicated with triangle markers in later figures.

The result of the calibration is depicted in Figure 6.5. The red, dashed lines show the expected spectrum without calibration. The blue, solid line shows the spectra after calibration (but still without fitting the cloud parameters). All four points were simultaneously calibrated. The second panel in the figure still shows a discrepancy, probably because the BEC fraction was too high for TOF imaging to extract the correct numbers. Notice how the red, dashed lines quickly approach $\mathcal{T} = 1$ when detuned, while the blue calibrated lines correctly correct the missing signal at high detuning.

The calibrated value for R_{ph}^{*} was $27.6\,\mu$m, which should be compared to the expected value of $26.2\,\mu$m. It is thus slightly larger, which is the expected direction if the cause is imperfections in the imaging optics.

The NA was calibrated to be $R_{\mathrm{NA}} = 6.77\,$mm, which is smaller than expected from 2 ” optics. Either it points to the presence of misjudged or unnoticed obstacles in the imaging path, or it is a consequence of the effective model that is used to calculate the effects. It should be noted, however, that the value has been very consistent throughout different calibrations and datasets.

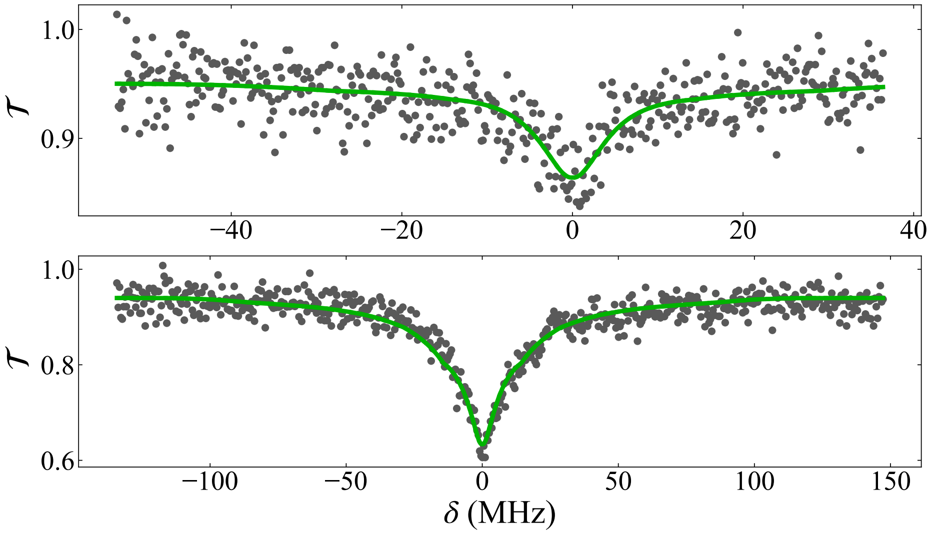

Figure 6.6: Two examples of the fitted spectra to data. The temperature increases from top to bottom. The top panel has $N_0 = 64\,\mathrm{k}$, $N_{\mathrm{th}} = 8.7\,\mathrm{k}$, $T = 46\,\mathrm{nK}$, $f_x = 93.7\,\mathrm{Hz}$, $f_y = 71.8\,\mathrm{Hz}$, and $f_z = 40.5\,\mathrm{Hz}$. The bottom panel has $N_0 = 191\,\mathrm{k}$, $N_{\mathrm{th}} = 104\,\mathrm{k}$, $T = 159\,\mathrm{nK}$, $f_x = 117.5\,\mathrm{Hz}$, $f_y = 114.6\,\mathrm{Hz}$, and $f_z = 57.5\,\mathrm{Hz}$.

6.1.3.2 Spectral fitting of bright field

Each dataset is binned to have roughly 2000 counts per bin, which results in relative Poisson noise of 2%. The value is chosen to be most optimal for the very coldest clouds since they are the most sensitive. The spectra are fitted by calculating the numerical aperture effects using the TOF image parameters first and then fitting the cloud with the numerical aperture transmission curve held constant. The NA transmission curve is then updated by the spectrally fitted parameters. This is repeated four times in total. The NA transmission curve cannot be a direct part of the fitting routine because it is too slow. With four iterations, a fit takes about 1 min to complete.

Examples of two fitted cloud spectra are shown in Figure 6.6. The top panel shows a warmer and larger cloud with a BEC fraction of 87%, and the bottom panel shows a colder and smaller cloud with a BEC fraction of 65%. The rest of the parameters are specified in the figure caption. The model agrees excellently with the measured traces.

The full dataset is taken with a wide range of trap frequencies, atom numbers, and preparation sequences to probe different regions and test the robustness of the spectral fitting routine. TOF images are always evaluated with the method described in Section 5.6.1.1, which fits the chemical potential and the temperature.

The spectrally fitted values of N_0 versus the measured values using TOF imaging are displayed in Figure 6.7. A dashed line shows the expected one-to-one relation, and the markers are colored according to their BEC fraction.

While the data follows the correct trend, the extraction of condensed atom numbers from these spectra is very noisy. The dominant reason is most likely the numerical aperture, which disturbs the signal enough to wash out the small dip in signal caused by the small cross-section of the BEC component. With a better spatial resolution (that is, a larger NA), the pinhole could have been chosen to more precisely match the image of the cloud and thus achieve a better signal-to-noise ratio.

Figure 6.8 shows the comparison of spectrally fitted N_{th} and the value extracted from TOF imaging. Again, a dashed line shows the expected one-to-one relation, and the markers are colored according to their BEC fraction. The points generally follow the expected dashed line over a wide span, but divergences start to happen for the clouds with a condensed fraction above $\approx 70\%$. The reason is expected to be erroneous fits of the TOF images, which can no longer resolve the thermal component well and can be substantially influenced by the repulsive BEC forces discussed in Chapter 3. The resonant light is also not well suited to image the high optical densities of the BEC,

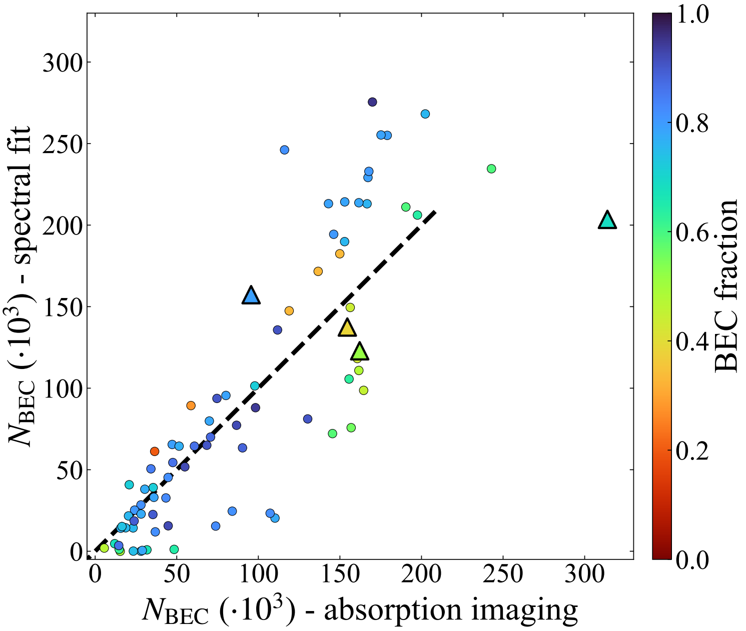

Figure 6.7: Spectrally fitted number of condensed atoms plotted against the TOF imaging extracted number. The color indicates the BEC fraction of each point. The triangle markers are the four points used for calibration. A dashed line follows the 1:1 relation.

even after expansion. Another contribution to the discrepancy is the assumption in the fits of the spectra that Γ is independent on spatial density.

Figure 6.9 plots the spectrally fitted temperature against the value extracted using TOF imaging, with the dashed line showing the one-to-one relation and the color of the markers indicating the BEC fraction. The temperature is, in general, the most stable fitting parameter and is not as affected by calibration as the atom numbers. It follows

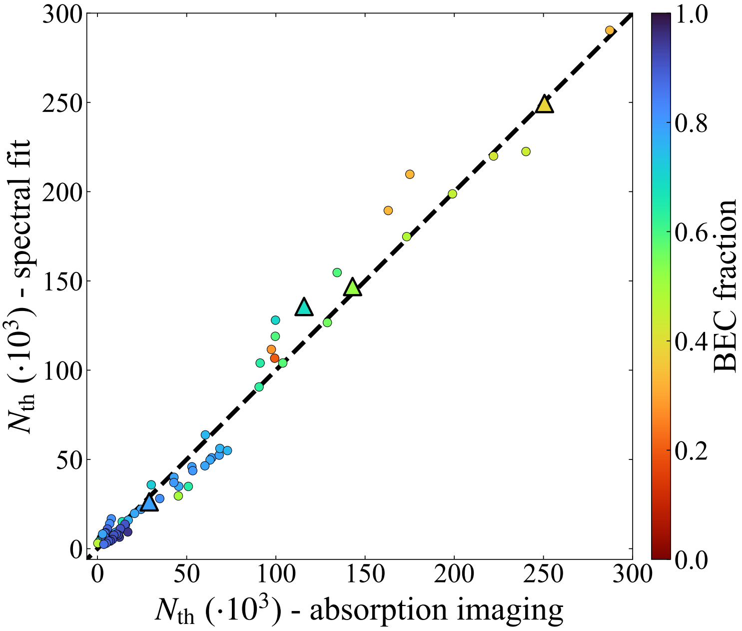

Figure 6.8: Spectrally fitted number of thermal atoms plotted against the TOF imaging extracted number. The color indicates the BEC fraction of each point. The triangle markers are the four points used for calibration. A dashed line follows the 1:1 relation.

the expected curve well, but similar to the thermal atom numbers, some divergence happens for the coldest clouds in the bottom left. When the condensed fraction is above $\approx 70\%$, the disagreement is appreciably larger than for warmer clouds, which is consistent with the results seen for the thermal atom numbers.

It is meaningless to plot the spectrally obtained BEC fraction against the temperature and compare the results to the semi-ideal model in eq. 2.21, because the spectral fit uses that model and would always

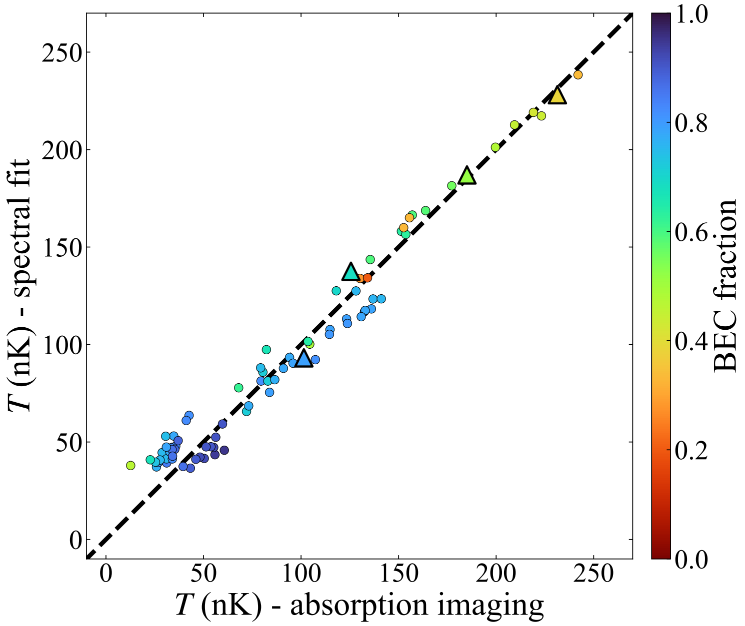

Figure 6.9: Spectrally fitted temperature plotted against the TOF imaging extracted value. The colorbar indicates the BEC fraction of each point. The triangle markers are the four points used for calibration. A dashed line follows the 1:1 relation.

yield results on the theoretical curve. What can be done instead is to use the thermal atom numbers and temperatures extracted from the spectra together with the total number of atoms calculated as the sum over the column densities in the TOF images, which specifically does not involve any fitting, to achieve a measure of BEC fraction versus temperature. The result of that analysis is shown in Figure 6.10, where the gray area between two dark gray curves shows the upper and lower theoretical semi-ideal curves for the range of all atom numbers

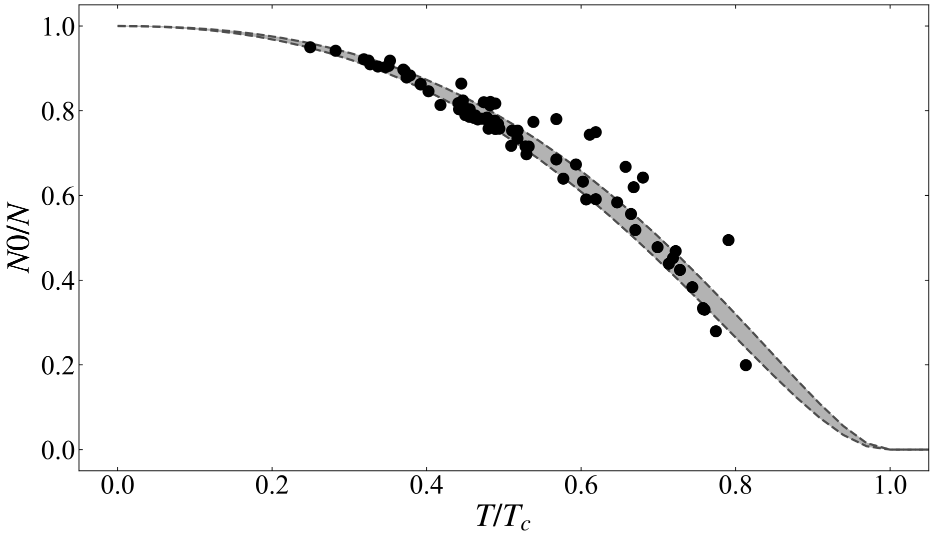

Figure 6.10: The BEC fraction versus temperature. The spectrally fitted temperature and thermal atom number were combined with the TOF imaging extracted total atom number to estimate the BEC fraction. The gray area and dark gray dashed lines indicate the area in which the semi-ideal model predicts all data points lie, given the range of trap frequencies and atom numbers.

and trap frequencies in the dataset.

There is great agreement between the theory and the measurement, with the highest fraction reaching up to 95% BEC. The exception is a couple of unexplained outliers above the curve. This kind of measurement illustrates the useful potential of combining non-destructive and destructive measurements to obtain more knowledge than by using any of them in isolation.

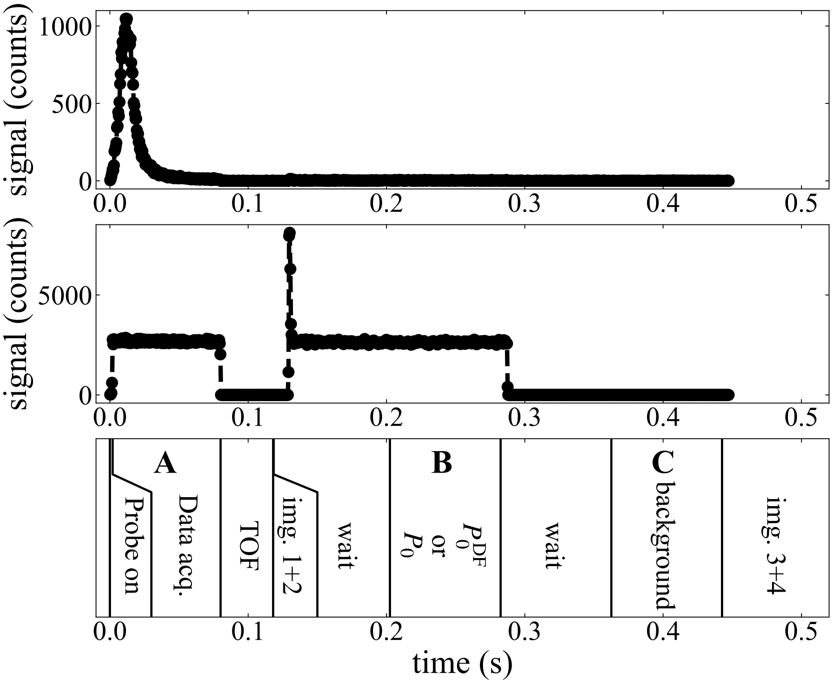

Figure 6.11: The data acquisition sequence for dark field measurements. The top shows an example of the SPD trace during the measurement, the middle shows a trace during the probe power measurement without atoms and a pinhole, and the bottom labels the different sections of the sequence. Three sections are labeled **A**, **B**, and **C** and are the measurement of the spectrum of the cloud, the probe power measurement, and the background power measurement.

6.1.4 Dark field

Dark field data acquisitions follow the same sequence as bright field, except a separate measurement without the dark field target must be performed to get P_0. Dark field datasets are not accumulated because a single run contains enough data to fit the models. The sequence is illustrated in Figure 6.11 with the SPD trace of a pure thermal cloud with 251 k atoms and a temperature of 409 nK shown as an example. The pinhole is always chosen to be much larger than the cloud to avoid any possible complications, but small enough to block most of the remaining probe light after the dark field target. A pinhole diameter of 150 μm is almost always a good choice and was used in all dark field datasets. The background magnetic field was 0.24 G along the probe beam for all dark field measurements. The top panel shows a measurement with dark field, and the middle panel shows the same measurement sequence but without a dark field target and without atoms. The pinhole in the second measurement was 75 μm to avoid saturation of the SPD.

Intervals **A**, **B**, and **C** of the measurement with dark field and atoms (top panel) gives P_{atom}, P_0^{DF}. and P_{bg} respectively. The **B**-interval of the measurement without dark field and atoms (middle panel) gives P_0. The value of $\bar{P}_0$ has to be scaled properly to match the powers of the first measurement because the pinholes were different. With the pinholes used here, the multiplication factor is 4. It is only necessary to measure P_0 once and then use this value for many datasets, unless the probe power was updated. The normalized transmission is calculated from these as

$$\mathcal{T} = \frac{P_{\text{atom}} - \bar{P}_0^{\text{DF}}}{\bar{P}_0 - \bar{P}_{\text{bg}}}. \tag{6.3}$$

The result of this particular analysis can be seen in the lower right panel of Figure 6.2.

6.1.4.1 Calibration of dark field

The calibration factor α was always calibrated to be near unity and was therefore deemed unnecessary. Similarly, α' was always calibrated

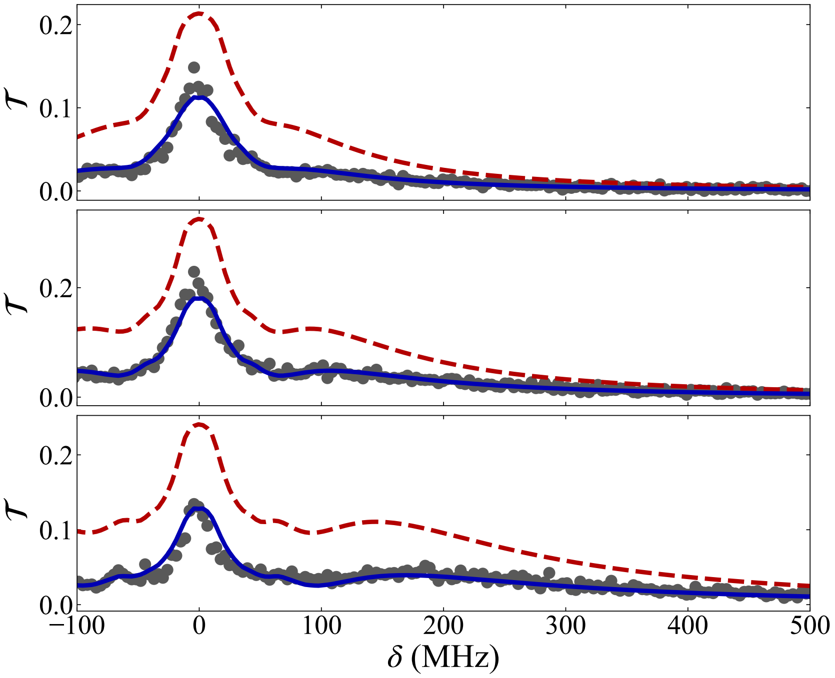

Figure 6.12: Calibration of dark field. The red dashed lines show the expected spectra with $\beta = 1$ and $R_{\mathrm{NA}} = \infty$. The blue solid lines show the spectra after calibration of the two parameters such that all three measured spectra match the closest.

to be near 0 and was also omitted. This can be explained by the low power of the probe [117].

The power factor β seems to be necessary when using dark field, which consistently gives a lower signal than expected from a perfect setup. This factor is calibrated together with the numerical aperture R_{NA}.

As was also the case for bright field, the R_{NA} parameter is not a perfect representation of the setup, and will function as an effective parameter, covering both real aperture effects and some of the approximations made in the model, e.g., the approximation analogous to the thin-lens approximation made in the propagation of light through the clouds, the negligence of bosonic effects in the light-atom interaction, or the assumption that Γ is independent on density.

The in-trap pinhole radius is not calibrated since the only effect of this parameter in dark field would be to scale P_0 and is thus equivalent to β. The resonance is calibrated before anything else and is independent of all other parameters. Finally, while it is expected that the value of Γ changes depending on atom density, this is not calibrated because it is prone to over-fitting, especially when combined with R_{NA}.

The dark field calibration was done on three clouds chosen to span a range in temperature and atom numbers, but most importantly, to include at least one point with a high number of BEC to get a sensitive measure of the NA. The three clouds are shown as triangular markers on the upcoming plots. The parameters β and R_{NA} were varied, and the resulting expected spectrum from the TOF images was calculated. This was used to find and minimize the difference between the expected and measured curves on all three spectra simultaneously.

The result of the analysis is shown in Figure 6.12, where the dashed, red curves show the expected spectra without calibration and the solid, blue curves show the theoretical spectra with the calibrated parameters. The results of the calibration was $\beta = 0.607$ and $R_{\mathrm{NA}} = 7.23\,\mathrm{mm}$.

The power factor is the scaling of the normalized transmission and is therefore not related to the efficiency of the setup directly. Instead, it gives how much more of the diffracted signal is lost compared to the probe light, after some of the loss has been accounted for by the NA. It also takes care of errors in calculating the in-trap probe intensity P_0, which can, for example, be modified by small deviations in the magnification of the setup. Ideally, the factor should be close to 1; however, the calibrated value is somewhat lower than anticipated. After the installation of the 3 ” imaging optics in the near future, this can potentially be revisited.

The calibrated numerical aperture is also smaller than expected. It is about 7% larger than the value found for bright field. Just before taking the data for dark field, the probe field propagation was optimized to exit the vacuum window through the center, which could be connected to the small increase in numerical aperture. Overall, the value for the numerical aperture in bright field and dark field seems to agree on the same, relatively low value.

6.1.4.2 Spectral fitting of dark field

Dark field spectra are fitted with the same routine as bright field spectra. That is, the NA effects are calculated from the TOF imaging extracted cloud parameters, then the normalized transmission $\mathcal{T}$ is fitted to the data, and the NA effects are recalculated. The fit is performed four times in total, with the NA updated each time.

Three examples of fitted dark field spectra are shown in Figure 6.13, where the BEC fraction decreases from top to bottom, and with the specific cloud parameters specified in the caption. The dotted part of the fitted green line indicates the red-detuned region of the data that is excluded from the fit to avoid influence from the $|F' = 1\rangle$-state.

The dark field spectra are much richer in structure than the bright field counterparts because of the phase shift. The model can replicate those features to some extent. Most notable is the ”shoulder” that appears around $100 - 200\,\mathrm{MHz}$ in the two top panels, which contain BEC. The model is, however, not able to replicate the coldest cloud data (top panel) for moderate detunings between the resonance and the shoulder. This is probably a combination of NA and broadening effects on the linewidth, which were not taken into account.

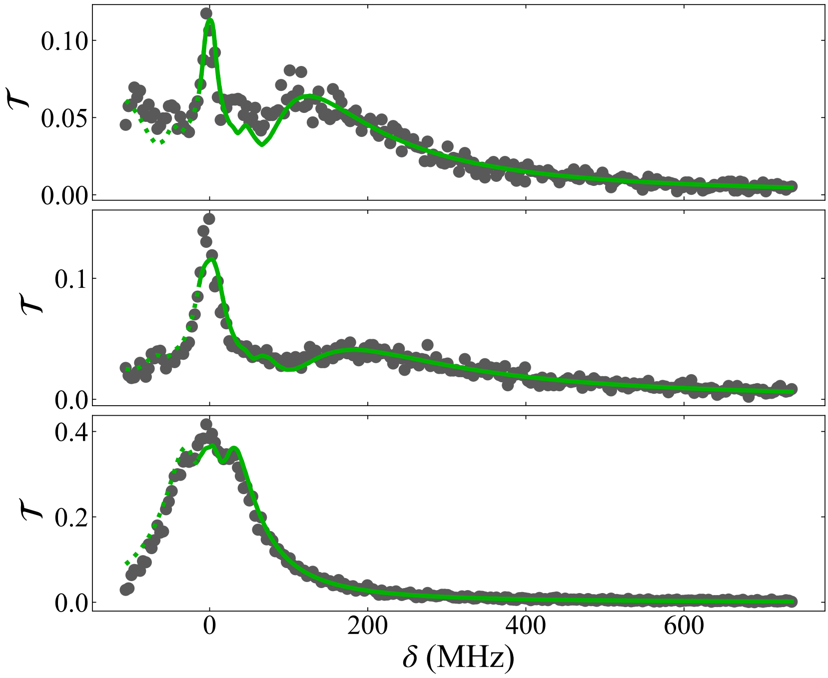

Figure 6.13: Three exemplary fits of dark field spectra. The cloud temperature increases from top to bottom. The top panel has $N_0 = 156\,\mathrm{k}$, $N_{\mathrm{th}} = 31\,\mathrm{k}$, $T = 67\,\mathrm{nK}$, $f_x = 85.8\,\mathrm{Hz}$, $f_y = 55.0\,\mathrm{Hz}$, and $f_z = 48.9\,\mathrm{Hz}$. The middle panel has $N_0 = 148\,\mathrm{k}$, $N_{\mathrm{th}} = 56\,\mathrm{k}$, $T = 121\,\mathrm{nK}$, $f_x = 109.7\,\mathrm{Hz}$, $f_y = 98.6\,\mathrm{Hz}$, and $f_z = 62.6\,\mathrm{Hz}$. The bottom panel has $N_0 = 0\,\mathrm{k}$, $N_{\mathrm{th}} = 259\,\mathrm{k}$, $T = 418\,\mathrm{nK}$, $f_x = 126.9\,\mathrm{Hz}$, $f_y = 124.3\,\mathrm{Hz}$, and $f_z = 67.6\,\mathrm{Hz}$.

TOF images are always evaluated with the method described in Section 5.6.1.1, which fits the chemical potential and the temperature.

Figure 6.14 shows the extracted condensed atom numbers compared to the TOF imaging estimated numbers. The results are significantly

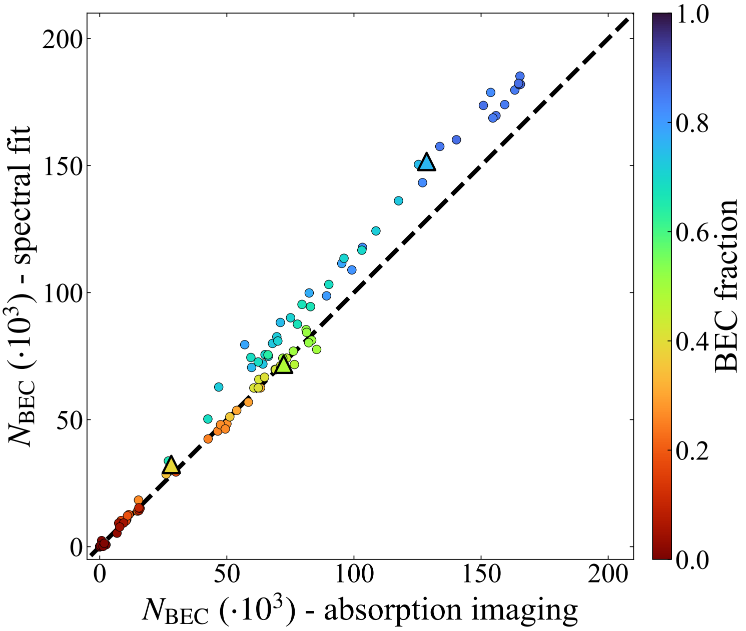

Figure 6.14: Spectrally fitted number of condensed atoms plotted against the TOF imaging extracted number. The color indicates the BEC fraction of each point. The triangle markers are the four points used for calibration. A dashed line follows the 1:1 relation.

better than the outcome from bright field in Figure 6.7. Once again, a disagreement appears around 70% BEC fraction, but it is even clearer here than in any of the previous plots. Notice how the split from the 1-to-1 line is independent of the number of condensed atoms, but is instead given specifically by the condensed fraction. This indicates, as previously, that the main problem is not the spectral fits but the extraction of parameters from very cold clouds using TOF imaging. The assumption that Γ is independent on spatial density is most likely

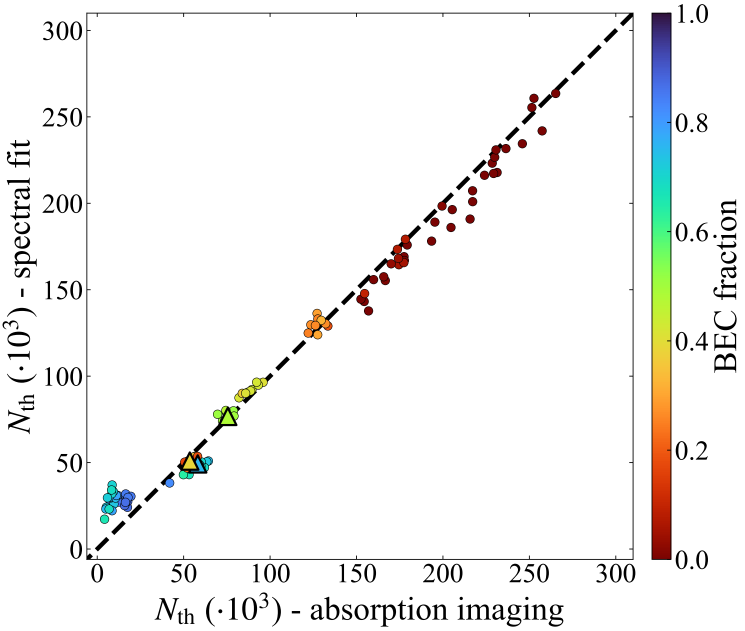

Figure 6.15: Spectrally fitted number of thermal atoms plotted against the TOF imaging extracted number. The color indicates the BEC fraction of each point. The triangle markers are the four points used for calibration. A dashed line follows the 1:1 relation.

also contributing to the discrepancy.

The measured values of N_{th} with the two methods are compared in Figure 6.15 and also shows remarkable results. There is great agreement between the methods over a wide range except, as always, for clouds with more than 70% condensed atoms.

Finally, the temperature measured with spectra is compared to the temperatures measured with TOF images in Figure 6.16. Again, the two methods agree well over a large range of cloud parameters, and

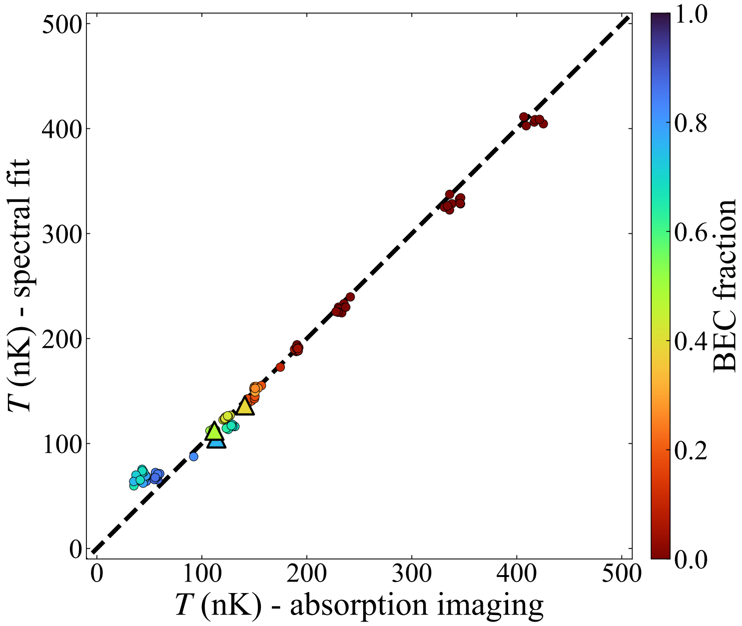

Figure 6.16: Spectrally fitted temperature plotted against the TOF imaging extracted number. The color indicates the BEC fraction of each point. The triangle markers are the four points used for calibration. A dashed line follows the 1:1 relation.

similar to bright field, the temperature is the most stable parameter. Clouds with more than 70% condensed fraction continue to give disagreement between the two methods.

To conclude, the spectral method works excellently for extracting all three parameters N_0, N_{th}, and T in a wide range of scenarios. It appears to work better than TOF imaging for very cold clouds with more than 70% BEC, which is most likely a consequence of three factors: the expansion of the cloud is significantly affected by the repulsive forces of BEC at this fraction, the thermal component is hidden behind the BEC component and is hard to fit, and the high optical density is poorly imaged with resonant light.

6.2 Time resolved measurements

The SPD setup can be used to obtain time-resolved measurements on a single cloud. This was exploited to extract useful information from the experiment much quicker than what is possible with absorption imaging. This section presents the most useful applications of this method that have been realized so far, outside the measurement of spectra. The four use-cases are trap frequency measurements presented in Section 6.2.1, Rabi oscillations presented in Section 6.2.2, state preparation dynamics presented in Section 6.2.3, and parametric sweeps presented in Section 6.2.4. While not exhaustive, this list of applications illustrates the method's potential to enhance a broad spectrum of measurements. See, for example [42] for a recent application of time-resolved measurements to extract complex dynamical parameters from ultracold bosonic clouds.

6.2.1 Trap frequencies

Trap frequencies are mostly measured in the cODT. In this trap, an oscillation can be initiated by pulling the cloud with a magnetic QP field and then quickly releasing it. The direction of the pull can be controlled by moving around the QP center with background magnetic fields by varying the currents in the shim coils. It gives the best data quality if the oscillation is initiated in only one direction at a time, although it is possible to measure several frequencies at once if more oscillations are initiated simultaneously.

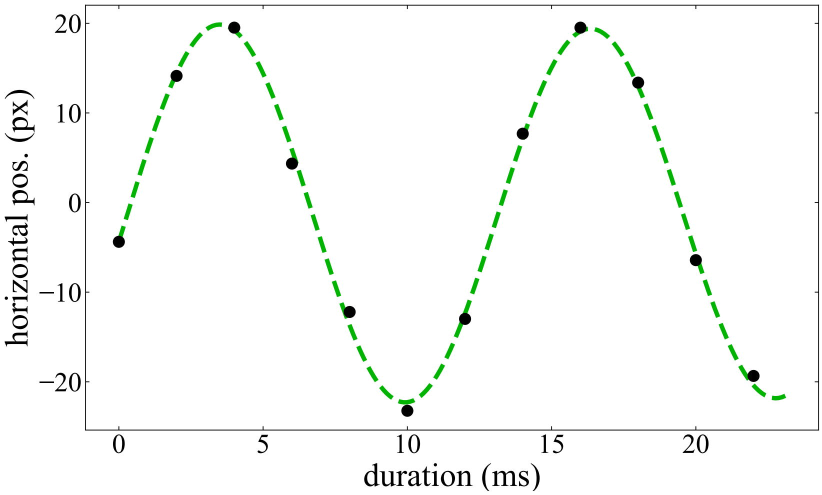

Figure 6.17: Trap frequency measurement in the Z-direction by imaging the position of the cloud after 38 ms TOF. The experiment was repeated with different durations after the oscillation was initiated, and required the production of a new cloud for each point.

The trap depth of a cODT is relatively low, especially in the vertical direction at low powers. If the cloud is pulled too hard initially, it can enter the anharmonic region of the trap, and this will reduce the lifetime of the oscillation. For the lowest powers in a cODT, this is almost impossible to avoid.

It is beneficial to evaporate to a relatively pure BEC and then increase the dipole powers again to the desired value before measuring trap frequencies, because this increases the precision of the measured cloud position.

Trap frequencies can be measured with absorption imaging by initiating an oscillation and then releasing the cloud from the trap at various points in time after the initiation. This technique measures

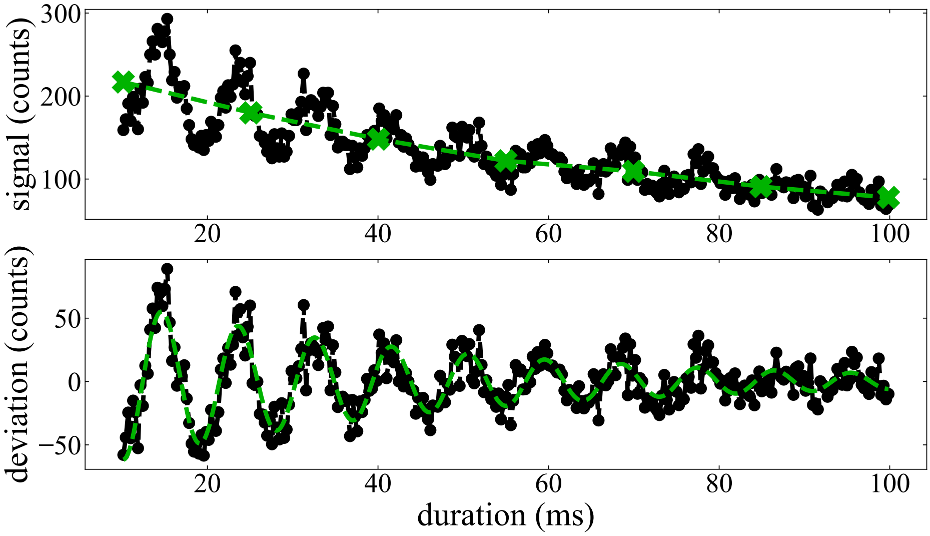

Figure 6.18: Trap frequency measurement with the SPD setup. The cloud is partly blocked by the pinhole, and the oscillation translates to an oscillating signal. The decay of the atoms is subtracted (illustrated on the top), and the frequency of the residual is fitted (bottom). The data was taken on a single cloud.

the oscillation of the momentum of the cloud instead of the position, but it has the same frequency as the in-trap oscillation. An example of such a measurement can be seen in Figure 6.17 with both dipole beam powers set to 500 mW. The data shows the horizontal positions of the cloud relative to the average position, obtained by fitting a Thomas-Fermi distribution to the images taken from the X-direction. The frequency measured here is thus the Z frequency and the result is $2\pi \cdot 77.84\,\mathrm{Hz}$. The dataset is comprised of 12 runs. The main limitation is that the data points have to be taken close enough to get a full resolution of the oscillation, but many oscillations need to be included in the scan to get better accuracy. The measurements could be improved by taking a second set of data separated by a couple of

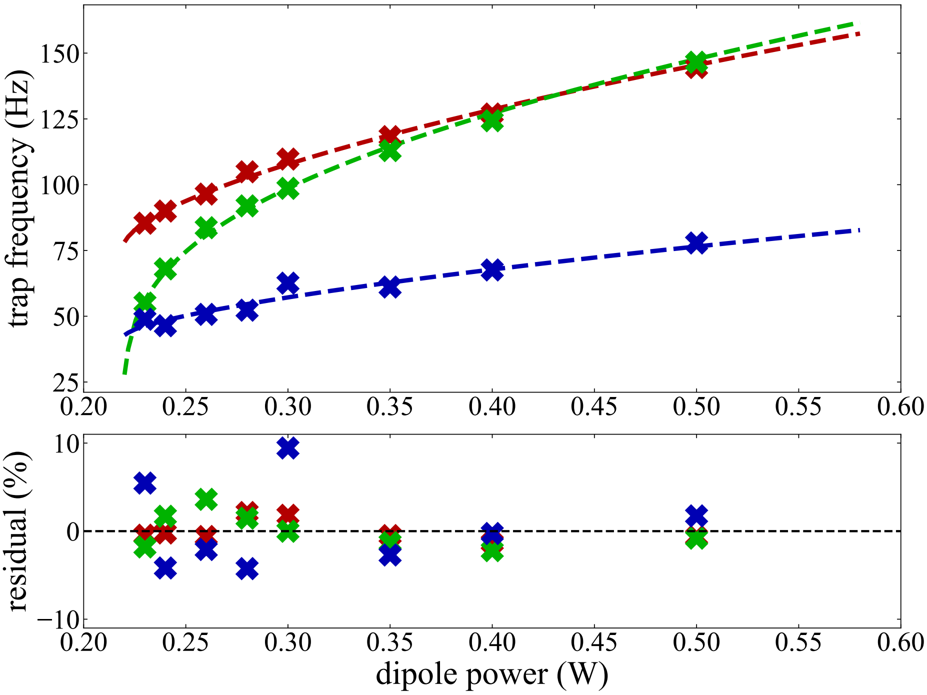

Figure 6.19: Trap frequencies versus dipole beam power. Both dipole beams were set to the same power. Red points are the X frequency, green points are the Y frequency, and blue points are the Z frequency. The data is fitted with a numerical model, and the relative residual from the fit is shown in the bottom panel.

periods.

An alternative way to measure trap frequencies is with the SPD setup. By placing the pinhole edge to cover half the cloud, any oscillation in the position orthogonal to that edge will be translated into an oscillating SPD signal. Figure 6.18 shows such a measurement using a single run, where the oscillation is now in the vertical direction. The dipole powers are the same as those from Figure 6.17. The beat frequency is 7 GHz giving a probe detuning of about 160 MHz.

This section uses the dark field setup. More data can be accumulated with more runs if needed. The top part of the figure shows the raw binned signal from the SPD during the oscillations, where the probe power is set just low enough for the cloud to survive longer than the oscillations. The signal is fitted with a series of straight lines interpolated between 7 points shown with crosses, which is then subtracted to get the pure oscillation seen in the bottom panel. This data is then fitted with an exponentially decaying sinusoidal to extract the trap frequency $2\pi \cdot 111.36\,\mathrm{Hz}$. The SPD data is able to catch the entire lifetime of the oscillation in a single run and is thus both faster and more accurate than the absorption imaging method. Unfortunately, the lifetime of the oscillations in the cODT is relatively short, which is the main limiting factor for improving this measurement.

It is also possible to use the SPD to measure the oscillations along the probe beam direction because the focus of the cloud would vary slightly, which modulates the signal going through the pinhole. The effect is smaller than any of the transversal directions and requires a fit of all oscillation directions at the same time to extract the longitudinal frequency, and this makes it necessary to accumulate more data to be able to separate the three frequencies. It is thus both simpler and just as fast to use absorption imaging to measure the frequency in that direction.

The faster measurement of trap frequencies, at least in two of the three directions, allows for the measurement of trap frequencies at more dipole powers. The top panel of Figure 6.19 plot the trap frequencies measured in the cODT at powers between 230 mW and 500 mW and with the same power in both beams. The X (horizontal and transverse to the probe beam) and Y (vertical) frequencies were measured with the SPD, and the Z direction (along the probe beam) was measured with absorption images. The Y-frequency starts to decrease drastically towards the lowest powers, indicating that gravitational sag is becoming significant, while the X and Z frequencies decrease more slowly throughout the full dataset. At high powers, the Y frequency becomes the largest because both beams are trapping in that direction, while only one beam contributes to the potential in each horizontal direction.

The measured frequencies were fitted with a numerical model of the trap frequency. Section 5.2.3 gives the equations for the potential of a dipole trap, but the trap frequencies are only given without the gravitational potential included because it requires a numerical model. The model used here works in two steps. It first finds the trap minimum position, which is offset by gravity, and then estimates the trap frequencies by a numerical second derivative around the minimum position. The input parameters are the two beam powers and the four waists. The beams are assumed to overlap perfectly and to be perfectly focused on the atoms. By fitting this model to the data, four waists were extracted. The result was $w^x_{0,y} = 82.3\,\mu\mathrm{m}$, $w^x_{0,z} = 79.6\,\mu\mathrm{m}$, $w^z_{0,x} = 68.5\,\mu\mathrm{m}$, and $w^z_{0,y} = 74.0\,\mu\mathrm{m}$. The values are close to, but slightly larger than, the expected waists as discussed in Section 5.2.3.5. The fit is able to mimic the data very well. The relative residuals from the fit are plotted in the bottom panel of Figure 6.19. The absorption imaging trap frequencies have a slightly higher relative noise than the SPD measured trap frequencies, possibly because fewer periods of the oscillation were covered with the absorption images.

6.2.1.1 Maximum likelihood fit to trap frequencies

This section investigates the limit of counts needed to fit trap frequencies if the MLE fit is used to analyze the data (see Section 5.6.2.3 for how MLE fitting works). The dataset is the one displayed in Figure 6.18 with the same fitting routine, which fits the background as an interpolation between points and then fits the oscillations on top.

To analyze the limit of counts needed to fit trap frequencies, a random sample of N points is drawn from the list of timestamps of

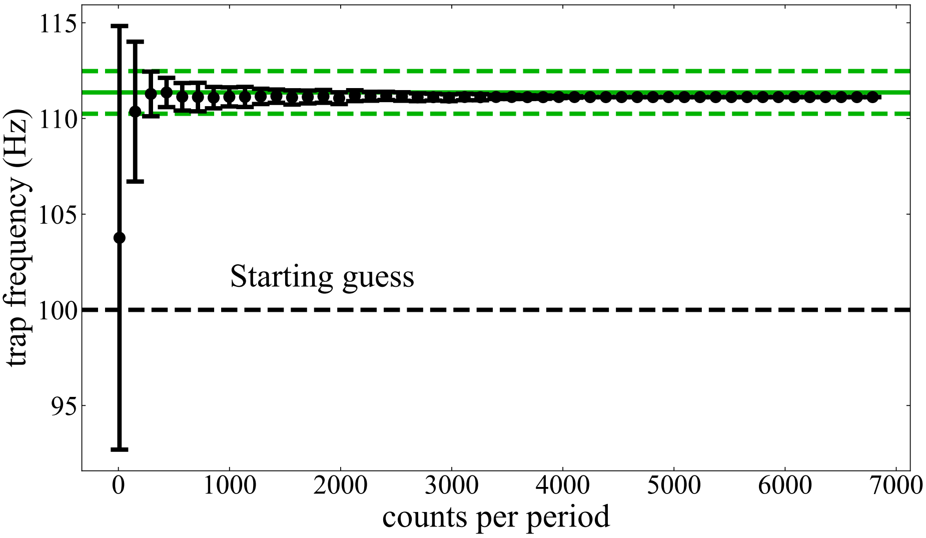

Figure 6.20: Randomly removing counts from the dataset displayed in Figure 6.18 to test the least signal needed to fit trap frequencies with MLE. Each point is an average of 50 repetitions. Dashed lines illustrate the ±1% band around the fitted value in Figure 6.18 (solid line). The dashed lines are crossed below 300 counts per peak.

photon detections. This is done 50 times for each value of N and for 50 different values of N, between 100 counts in total and 69443 counts in total (the full dataset). The mean and the standard deviation of all the trap frequencies from the MLE fits are plotted in Figure 6.20 for each value of N, except the full dataset, which would yield the same result each time, as a function of counts per period. One period is taken to be $T = 1/111.36\,\mathrm{Hz} = 8.98\,\mathrm{ms}$, which comes from the fit shown in Figure 6.18. The starting guess is deliberately set to be off to test the stability and accuracy of the fit and is indicated with a black dashed line. A green solid line shows the 111.36 Hz value and two dashed lines show a ±1% interval around that. At a 300-count

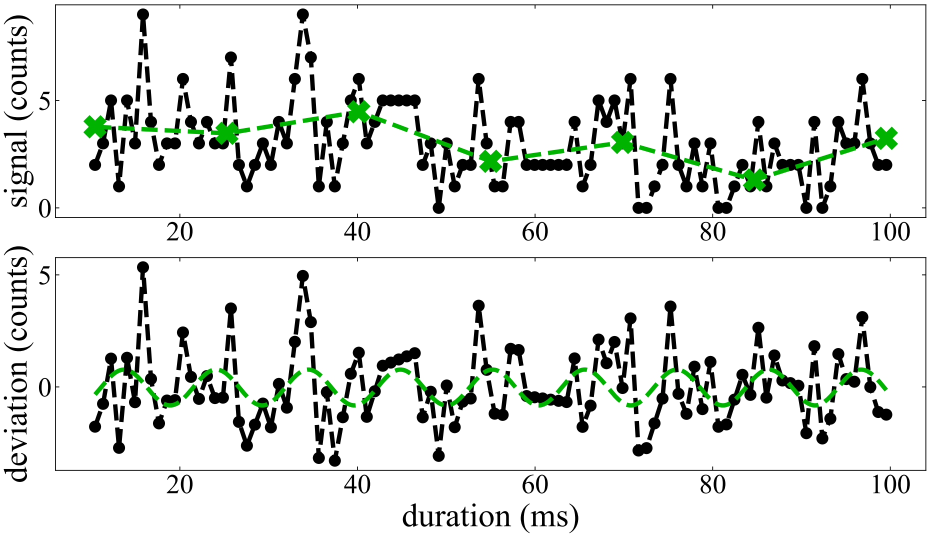

Figure 6.21: An attempt to fit a 300 counts per peak instance with histograms instead of the MLE method.

period, the standard deviation starts to cross the ±1% boundaries.

Figure 6.21 shows an attempt to repeat the analysis with histograms, but with only 300 randomly selected points. As can be seen, there are not enough counts to meaningfully bin the data. By repeating this analysis 50 times as well, the results varied by ±30%. This illustrates how MLE fitting can push the boundary of counts needed in an experiment.

With only 300 counts per period of 8.98 ms, and if it is assumed that 10 counts per lost atom can be achieved, which is well within reasonable limits considering the results in Figure 4.4, the trap frequency in a cloud with 200 k atoms can be measured for a full minute until all atoms are gone because of the probe light. That is, the bottleneck in trap frequency measurements with the SPD is not the destructivity of the probe field, but the lifetime of the atoms in the trap itself

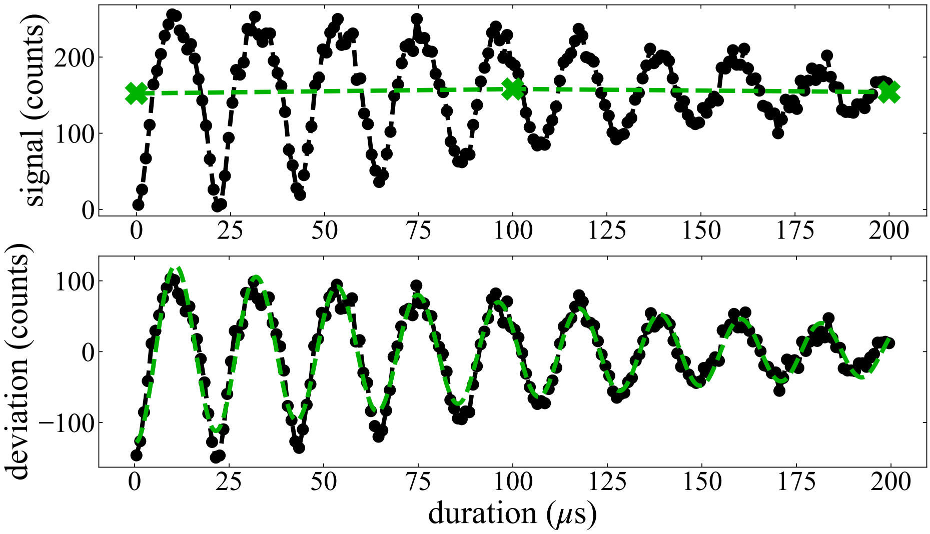

Figure 6.22: A measurement of a Rabi oscillation between the two F-states in the $S_{1/2}$-manifold. Only $|F=1\rangle$ is visible to the detector, and the signal oscillates together with the population in that state. Top: the raw signal and the fitted background. Bottom: the residual from the background and a decaying sinusoidal fit.

or the lifetime of the oscillations. To get around oscillation lifetime, the frequency could be initiated multiple times with the same cloud, perhaps even with changing experimental parameters in between.

6.2.2 Rabi oscillations

Rabi oscillations can be measured live with the SPD if the oscillations are between states such that only one of them is visible to the probe. The results in this section measure oscillations between $|F=1, m_F=1\rangle$ and $|F=2, m_F=2\rangle$ in the $S_{1/2}$ manifold. The MW signal is turned on at the start of the measurement on a constant fre-

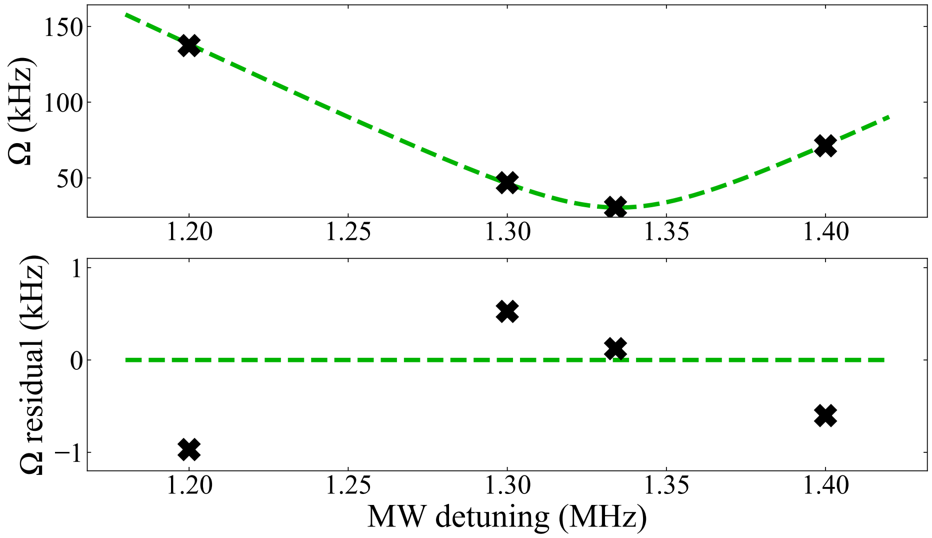

Figure 6.23: The measured Rabi frequencies against various MW detunings from the nulled magnetic field resonance. The data is fitted with the formula for generalized Rabi frequency, shown in the top panel, and the residuals from the fit are shown in the bottom.

quency, and the signal is fitted in the same way as trap frequencies in Section 6.2.1 and a constant magnetic field offset is applied along the probe beam in the Z-direction. In this section, the probe frequency was resonant, which is not optimal for signal-to-noise ratio, meaning more signal could have been acquired with a larger detuning. This section uses the dark field setup.

Figure 6.22 illustrates an example of a Rabi frequency measurement with 42 repeated runs for accumulated data and with an almost pure BEC cloud. The analysis is the same as for trap frequency measurements, where the background is fitted with a set of straight lines (top panel) and the signal on top of the background is fitted with a decaying sinusoidal (bottom panel), but the signal is in principle more complicated and would follow eq. 4.42 for $\mathcal{T}$ with a sinusoidally varying atom number in $|F=1, m_F=1\rangle$. This can be seen clearly in the top panel, where a saturation effect results in a more rounded top of the oscillation signal. However, the extracted frequency will still match the Rabi frequency well.

The decay of the Rabi oscillation is partly due to decoherence effects and partly due to non-classical effects described by the Jaynes-Cummings model [118], and longer acquisitions show revival of the signal in agreement with the Jaynes-Cummings model. This should be investigated further.

Figure 6.23 shows the fitted Rabi frequencies at different MW frequencies on the top and the residuals from a fit with the generalized Rabi frequency $\sqrt{\delta_{\mathrm{MW}}^2 + \Omega_{0,\mathrm{MW}}^2}$, where δ_{MW} is the MW detuning and $\Omega_{0,\mathrm{MW}}$ is the resonant Rabi frequency. The resonance is estimated to be $1334.9 \pm 0.6\,\mathrm{kHz}$, which can be translated to a magnetic field of $635.7 \pm 0.3\,\mathrm{mG}$. More precise measurements can be achieved with more MW frequencies, but the main bottleneck at this point is the temporal drift in the magnetic field, which is possibly comparable to the current precision.

6.2.3 State preparation

The ^{87}Rb atoms in the experiment start in the $|F=2, m_F=2\rangle$ state until they are cold enough to be trapped in a pure cODT, after which they are transferred adiabatically to $|F=1, m_F=1\rangle$ with a MW frequency sweep. This transfer can be monitored with the SPD because the atoms are transferred between a visible and an invisible state. To do this, the probe beam should be turned on at the beginning of the MW, and the probe frequency should be constant and ideally far-detuned for minimal destruction.

Measurements were performed on a very cold BEC with dark field.

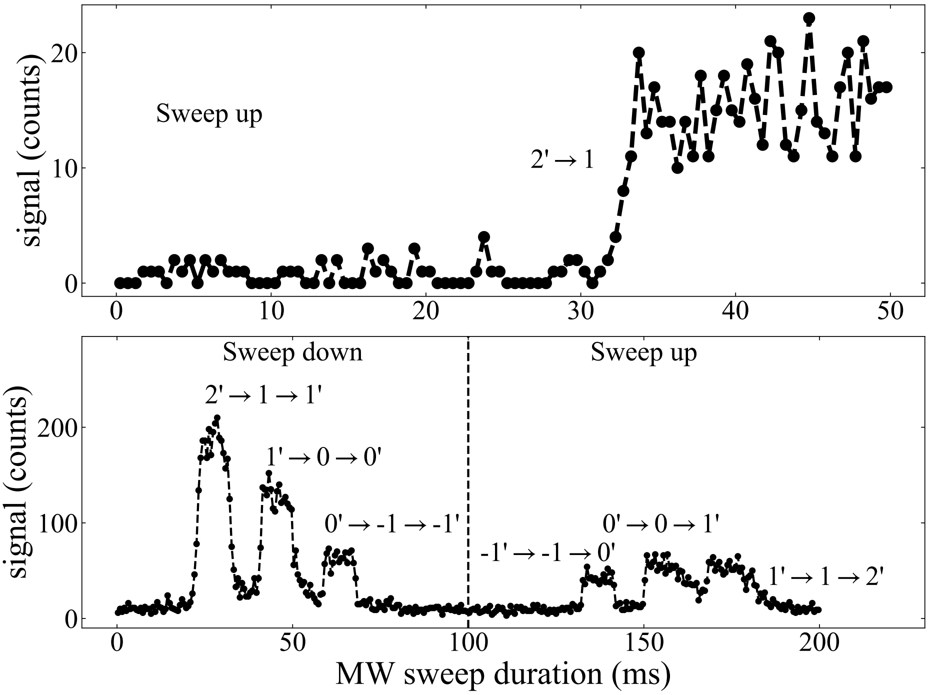

Figure 6.24: Live measurements of state manipulation of atoms on single clouds. All state transfers are labeled, with numbers indicating the values of m_F, a mark indicating the atoms being in the $|F = 2\rangle$-state, and no mark indicating atoms in the $|F = 1\rangle$-state, all in the $S_{1/2}$ manifold. Top panel: a MW sweep crosses the resonance from below, transferring the atoms from an invisible to a visible state of the probe. Bottom panel: the MW frequency sweeps from above all resonances to below all resonances and back. It crosses several transitions, resulting in multiple peaks in the signal.

Non-ideally, the probe frequency was resonant for all measurements in this section. Measurements were performed on a single cloud. The signal on the SPD while sweeping the MW across the $|F = 1, m_F = 1\rangle \rightarrow |F = 2, m_F = 2\rangle$ resonance can be seen in the top panel of Figure 6.24. The MW frequency was swept from $\delta = 0$ to $\delta = 1\,\text{MHz}$ over $50\,\text{ms}$, where $\delta = 0$ is the zero magnetic field resonance for the transition. There was a small bias magnetic field during the measurement along the probe beam. The signal on the SPD is essentially zero at the beginning of the sweep until about 32 ms and a MW frequency detuning of 640 kHz. At 34 ms or a MW detuning of 680 kHz, the transfer is over. The resonance will be somewhere in between the two values, but because the signal does not scale linearly with atom numbers at this probe frequency, it will not be exactly in the middle. The transfer onset and completion frequencies would correspond to a magnetic field of 305 mG and 324 mG, respectively. A more precise value could be obtained with a larger probe detuning or by using eq. 4.42 to model the signal.

The bottom panel of Figure 6.24 shows what happens if the MW sweep is performed from above instead. This time, the sweep lasted 100 ms and went from $\delta = 2.5\,\mathrm{MHz}$ to $-\delta = 2.5\,\mathrm{MHz}$. It was then reversed over another 100 ms back to $\delta = 2.5\,\mathrm{MHz}$. The sweep has now crossed a total of 12 allowed transitions, giving rise to 6 peaks in the SPD signal. Each peak is labeled with the three m_F values the atoms go through, and a mark indicates a state in the $|F = 2\rangle$ manifold. The magnitude of the signal decreases between the peaks because of a change in the resonance frequency of the probe laser, a change in Clebsch-Gordan coefficients, and the lifetime of the atoms.

6.2.4 Parametric sweeps

Non-destructive live detection can be used to measure the effect of time-varying parameters. There are countless ways to achieve this, and this section shows only one.

By performing a MW sweep across the $|F = 1\rangle \leftrightarrow |F = 2\rangle$ resonance from below, the atoms can be transferred from an invisible to a visible state for the probe light. This can be seen in the top of Figure 6.24. The position of the rising edge can be used to calculate the

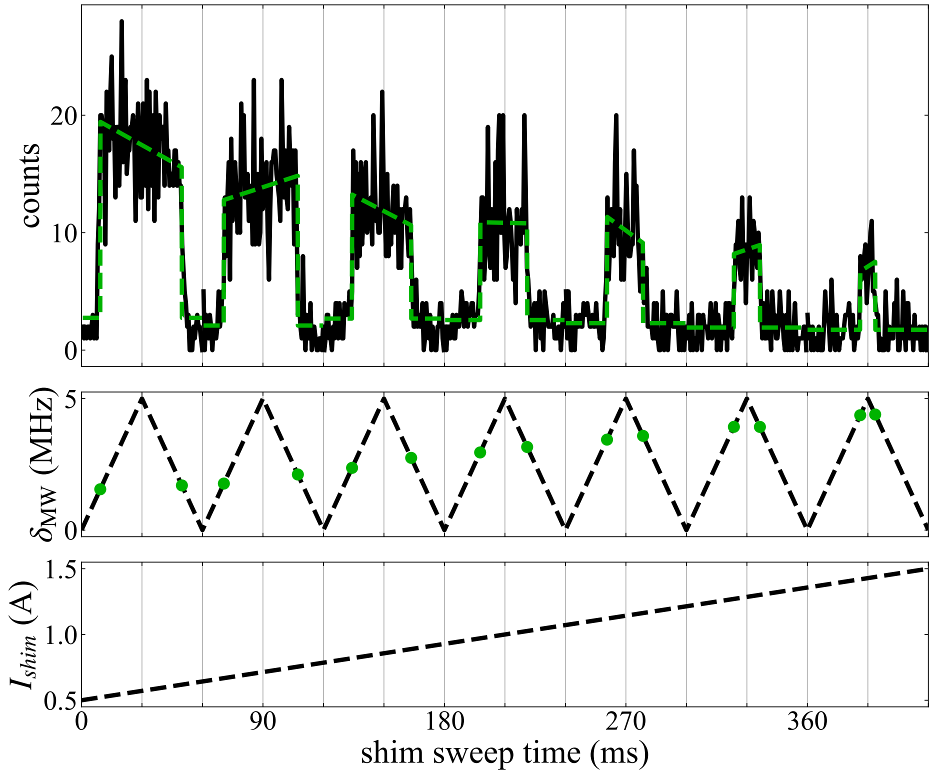

Figure 6.25: Live measurement on a single cloud of a continuous transfer between invisible and visible atomic states to the probe light while slowly ramping the background magnetic field. Top panel: the signal during the data acquisition. On top is a step-wise fit to extract the time of transition and thus the background magnetic field at that moment. Middle panel: the MW frequency sweeps. Green markers show the time and MW frequency when a transition occurs. Bottom panel: the current in the coils generating the magnetic field, which is slowly increased throughout the scan.

resonance and thus the magnetic field. By sweeping the MW frequency down again, the atoms will be transferred back into the invisible state $|F = 2\rangle$, and a second magnetic field measurement is achieved.

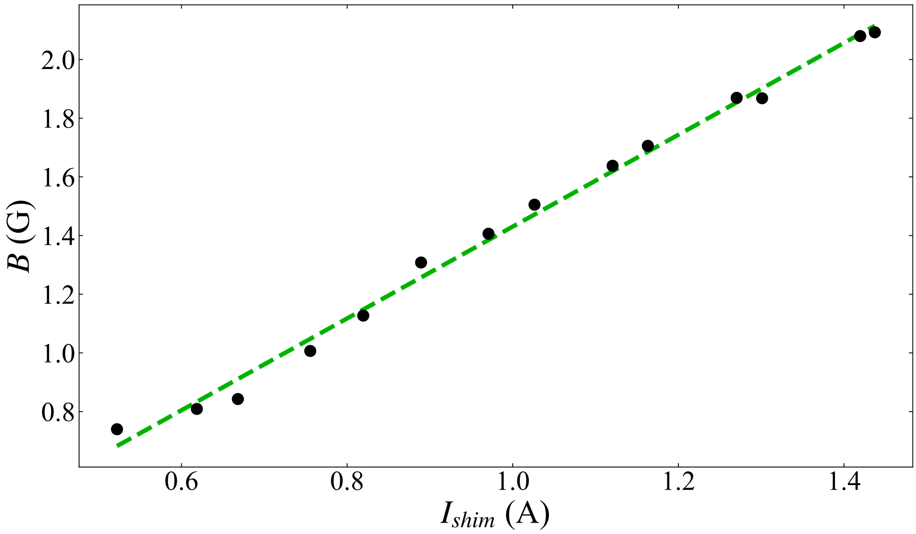

Figure 6.26: The extracted background magnetic fields from the measurement in Figure 6.24 as a function of the current in the shim coils. The data is fitted with a straight line.

If the MW field is continually swept up and down, the magnetic field can be monitored over time. In this example, the current in the shim-Z coils was ramped slowly while the MW-triangle sweep was active, and with a constant probe frequency turned on. The cloud is a pure 50 k BEC cloud, and the sweep is performed once. The result is shown in Figure 6.25, where the SPD signal is displayed in the top panel, the MW frequency scan is displayed in the middle panel as a detuning from the nulled B-field resonance frequency, and the shim-Z current scan is displayed in the bottom panel.

The SPD detector signal is fitted separately for each combined up-down MW sweep interval. The function has 5 parameters, which are the background signal, the signal just after the rising edge, the signal just before the falling edge, and the two positions in time of the rising and falling edges. The function models the edges as instant transfers between the powers and thus effectively fits the middle of the edges. The MW frequencies at the fitted moments of the rising and falling edges are plotted in the second panel as green points on top of the MW sweep curve.

As the magnetic field increases, the MW frequency has to scan further up from below before crossing the resonance, and less down from the maximum frequency to cross the resonance from above. The ”square” peaks in the SPD data will therefore shrink in width as the scan continues, and the dots in the second panel of the figure will move upwards.

Figure 6.26 plots the moments of resonance crossings translated into the magnetic field values versus the shim-Z current. The data is fitted with a straight line, which comes from the low B-field Zeeman shift with zero residual magnetic field in the X- and Y-direction. The slope gives the shim coil calibration value 1.57 G/A.

6.2.4.1 Maximum likelihood fit to magnetic field sweeps

Similar to the analysis done in Section 6.2.1.1, the lower limit of counts needed to measure the magnetic field sweep of Section 6.2.4 was investigated by randomly drawing a smaller sample of the dataset and fitting with the MLE method. For each number of counts, the random selection and fitting were performed 50 times. The mean and standard deviation of the fitted shim coil calibrations are plotted in Figure 6.27, with a solid green line indicating the result from Section 6.2.4 and dashed green lines indicating a ±5% band around the solid line.

In this case, to stay within ±5%, a minimum of about 1500 counts is needed in total. By once again assuming it is possible to get 10 counts per lost atom, taken conservatively from the results in Figure 4.4, this would make it possible to repeat the measurement about 130 times on

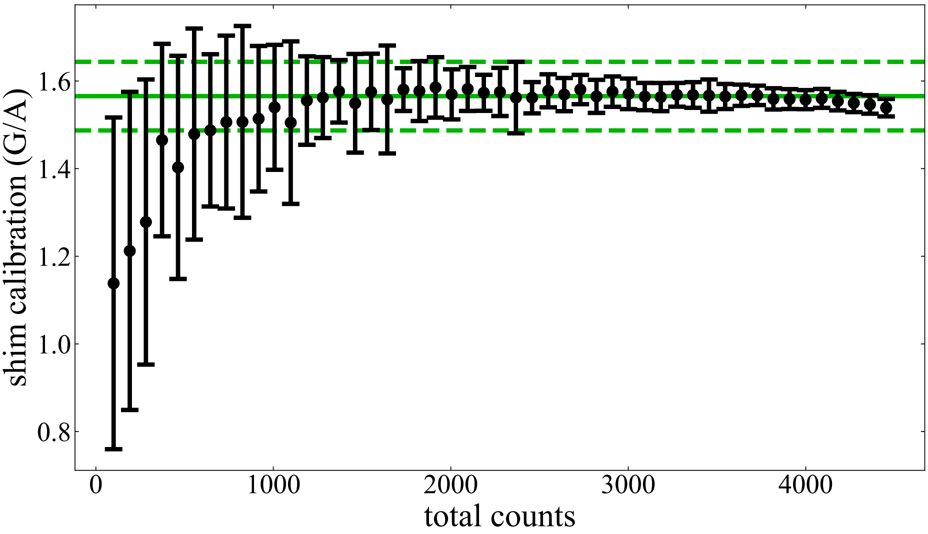

Figure 6.27: Re-analysis of the live magnetic sweep measurements with a random subset of all the photon detections and the MLE method. Each total number of counts is sampled 50 times, and the markers show the average value, while the error bars show the standard deviation. A solid green line shows the previous result from Section 6.2.4 and dashed green lines show a ±5% interval around it.

a single 200 k condensed atoms cloud, although the loss during MW sweeps would reduce this number significantly. This analysis appears to need more data than a trap frequency measurement (Figure 6.20).

Chapter 7

Conclusion and outlook

This thesis presents three main results. First, a detailed analysis of the free-fall expansion of bimodal clouds was carried out, taking into account the repulsive forces between atoms in-trap. This analysis enhances the accuracy of absorption imaging after time-of-flight (TOF) and complements previous work in the group on the statistics of absorption images [12]. Second, a new detection method was developed for extracting cloud parameters, N_0, N_{th}, and T, using spectra instead of spatial resolution. This method, originally conceived as a side-product of a broader goal to measure light-induced dipole–dipole broadening, constitutes an essential step towards future experiments. Third, the experimental apparatus itself has been significantly restored and upgraded, with improved stability, reduced cycle times, and the implementation of a phase-locked laser system.

Cloud expansion

The repulsive interactions between condensed and thermal atoms in a bimodal gas reshape the thermal distribution and alter the expansion profile of the cloud during TOF. If ignored, these interactions introduce systematic errors in the estimation of temperature and atom numbers. Because no analytical expression exists for the expansion, simulations are required. The simulations assume ballistic expansion, neglecting interactions during TOF, and then compare the result to a Bose-enhanced fit to estimate errors in data analysis.

The simulations presented in this thesis reveal errors on the order of 10%, depending on TOF, temperature, and trap geometry, with only weak dependence on the total atom number. The dominant factor determining the error magnitude is the relative size of the BEC compared to the thermal cloud.

The simulation can also be employed directly as a fitting model. When applied to experimental data, it yielded better agreement with the semi-ideal model [29] than Bose-enhanced fits. However, fitting each cloud with the simulation required approximately ten minutes, making it significantly slower and more complex than the Bose-enhanced fit. Thus, while the simulation-based model offers higher precision, it is best suited for cases where accuracy is paramount.

Spectral detection

A detailed model of light propagation through ultracold clouds was developed. While the underlying atom–light interaction theory is well established, the spatial distributions of atoms in thermal, condensed, and bimodal clouds were found to produce richly structured spectra. These spectra can be fitted to extract cloud parameters without the required spatial resolution for resolved images. Because a spectrum requires far fewer data points than a two-dimensional image, the required probe-light exposure is drastically reduced compared to absorption imaging. Moreover, since most spectral regions are far detuned, the number of absorbed photons, and thus the destructiveness of the measurement, is significantly lower than in resonant detection methods.

The non-destructive nature of spectral data acquisition enabled the combination of spectral measurements with TOF absorption imaging, allowing for the extraction of more information than either method may provide alone. This was illustrated with bright-field spectra, which were used to measure the condensed fraction as a function of reduced temperature, showing good agreement with the semi-ideal model up to 95% condensate fraction. A measurement that would be challenging with absorption imaging alone.

Dark-field spectra were also investigated, providing sensitivity to both absorption and phase shifts. The model reproduced these spectra with high accuracy. Because of their sensitivity to phase shifts, dark-field spectra are well suited for detecting light-induced dipole–dipole broadening. In particular, the spectral contributions from condensed and thermal atoms are sufficiently distinct to enable measurements focused on the densest regions of the cloud, with reduced interference from the thermal atoms. Furthermore, the rich spectral structures could provide enhanced sensitivity to a variety of effects, such as atom-number fluctuations between thermal and condensed components, or deviations from the two-level single-atom interaction model, including collective many-particle effects like sub- and superradiance.

The use of weak probe light enables time-resolved spectral measurements. This was demonstrated for well-understood cases such as trap-frequency measurements. In future work, this technique can be built upon to measure more novel effects, such as the lifetime of Raman Rabi frequencies and its dependence on cloud parameters, which may be used as a tool to probe many-body physics.

Additionally, continuous spectroscopic probing of a single cloud could be employed and used to monitor atom numbers and temperatures across the evaporation trajectory, including through the critical temperature.

Electromagnetically induced transparency

A setup with two phase-locked lasers was built to implement electromagnetically induced transparency (EIT). The effect was measured and characterized with high controllability. The system was originally conceived to enhance sensitivity to light-induced dipole–dipole broadening. However, the method was found to require higher precision than measurements of absorption spectra, and the need for longer dephasing lifetimes of the second ground state made this application unfeasible for measuring the highest density regions of the cloud. At present, EIT is therefore not in use for this purpose, but the setup remains a versatile tool for precise control of probe-laser frequencies.

A promising application of EIT is the measurement of optical Rabi frequencies and excited-state lifetimes. Conventional Rabi oscillation measurements are constrained by the requirement that the Rabi frequency exceed Γ, the natural linewidth of the excited state. This limitation can be circumvented by measuring Raman Rabi oscillations instead. In this case, the two-photon oscillations still reflect the excited-state lifetime, but decay much more slowly because only a small fraction of atoms occupy the excited state at any given time [88–90]. This approach could improve Rabi measurements and provide a tool for probing density-dependent many-particle effects [119].

More generally, EIT can enhance sensitivity to perturbations of the two ground states. For example, it could be employed to measure magnetic fields with optical light fields, overcoming the limited precision imposed by the $2\pi \cdot 6\,\mathrm{MHz}$ linewidth of the excited state. Both narrow transmission windows and far-detuned Raman absorption peaks could be utilized for such measurements.

Additional outcomes

The experiment has undergone substantial improvements and upgrades. The cycle time has been reduced from approximately 1.5 minutes to 30 s, enabling faster data collection and experimental progress. The stability of the setup has been significantly improved compared to the start of my PhD, and condensates of up to 300 k atoms can now be produced in a crossed optical dipole trap (cODT). The waists of the cODT can be tuned, with a minimum waist of 15 μm, although the smallest-waist configuration requires daily optimization of the beam overlap. Such tighter waists provide higher atomic densities, facilitating studies of density-dependent effects.

A combined camera and single-photon detector (SPD) setup with the option of inserting a dark-field target was also constructed. This setup has proven invaluable, substantially accelerating data collection by enabling the acquisition of many data points within a single experimental cycle. While originally designed to measure spectra, it has also facilitated other tasks such as trap-frequency and magnetic-field measurements.

Finally, a new three-inch optics system was added to one of the cameras and the SPD. Although not yet thoroughly tested, it has already shown improvements in the acquired data. This enhancement is expected to substantially benefit future measurements of light-induced dipole–dipole effects by improving spectral data quality.

Final remarks

The work in this thesis builds upon the group's previous focus on refining and advancing data-acquisition methods in ultracold atom experiments [12, 78]. The spectral detection method developed here is a significant result in its own right, while also serving as a stepping stone towards the next generation of experiments. In particular, the detailed understanding of dark-field spectra developed in this work will be invaluable for both the design of new experiments and the analysis of future data.

On a personal note, the PhD has provided me with substantial growth in technical, academic, organizational, and physics-related skills, which will undoubtedly be of great value in both my professional career and personal life. The work environment in Jan Arlt's group and at Aarhus University has offered academic freedom, strong support, and stimulating discussions with a diverse and inspiring community.

Bibliography

[1] M. Andersen, F. Nørfjand, and N. T. Zinner, "Real-time correlation function of floquet conformal fields", Phys. Rev. D **103**, 056005 (2021).

[2] A. Einstein, "Quantentheorie des einatomigen idealen Gases. II.", German, Sitzungsber. Preuß. Akad. Wiss., Phys.-Math. Kl. **1925**, 316 (1925).

[3] S. N. Bose, "Plancks gesetz und lichtquantenhypothese", Zeitschrift für Physik **26**, 178–181 (1924).

[4] J. P. Dowling and G. J. Milburn, "Quantum technology: the second quantum revolution", Philos. Trans. R. Soc. London A **361**, `https://doi.org/10.1098/rsta.2003.1227` (2003).

[5] K. Bongs et al., "Taking atom interferometric quantum sensors from the laboratory to real-world applications", Nat. Rev. Phys. **1**, 731–739 (2019).

[6] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, "Optical atomic clocks", Rev. Mod. Phys. **87**, 637–701 (2015).

[7] L. Pezzè, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein, "Quantum metrology with nonclassical states of atomic ensembles", Rev. Mod. Phys. **90**, 035005 (2018).

[8] B. Stray et al., “Quantum sensing for gravity cartography”, Nat. **602**, 590–594 (2022).

[9] I. Bloch, J. D., and S. Nascimbène, “Quantum simulations with ultracold quantum gases”, Nat. Phys. **8**, 267–276 (2012).

[10] S. Subhankar, Y. Wang, T.-C. Tsui, S. L. Rolston, and J. V. Porto, “Nanoscale atomic density microscopy”, Phys. Rev. X **9**, 021002 (2019).

[11] “40 years of quantum computing”, Nat. Rev. Phys. **4**, 1–1 (2022).

[12] T. Vibel, M. B. Christensen, R. M. F Andersen, L. N. Stokholm, K. Pawłowski, K. Rzążewski, M. A. Kristensen, and J. J. Arlt, “Atom number fluctuations in bose gases—statistical analysis of parameter estimation”, Journ. of Phys. B **57**, 195301 (2024).

[13] A. R. Perry, S. Sugawa, F. Salces-Carcoba, Y. Yue, and I. B. Spielman, “Multiple-camera defocus imaging of ultracold atomic gases”, Opt. Exp. **29**, 17029–17041 (2021).

[14] D. G. Fried, T. C. Killian, L. Willmann, D. Landhuis, S. C. Moss, D. Kleppner, and T. J. Greytak, “Bose-einstein condensation of atomic hydrogen”, Phys. Rev. Lett. **81**, 3811–3814 (1998).

[15] 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**, 198–201 (1995).

[16] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, “Evidence of Bose-Einstein condensation in an atomic gas with attractive interactions”, Phys. Rev. Lett. **75**, 1687–1690 (1995).

[17] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, “Bose-einstein condensation in a gas of sodium atoms”, Phys. Rev. Lett. **75**, 3969–3973 (1995).

[18] C. J. Pethick and H. Smith, *Bose-Einstein condensation in dilute gases*, Second Edition (Cambridge, 2008).

[19] D. V. Schroeder, *An introduction to thermal physics*, First edition (Pearson, 2014).

[20] M. B. Kruk et al., *On the fluctuations of the number of atoms in the condensate*, 2025.

[21] H. D. Politzer, "Condensate fluctuations of a trapped, ideal bose gas", Phys. Rev. A **54**, 5048–5054 (1996).

[22] R. M. Ziff, G. E. Uhlenbeck, and M. Kac, "The ideal bose-einstein gas, revisited", Phys. Rep. **32**, 169–248 (1977).

[23] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari, "Theory of bose-einstein condensation in trapped gases", Rev. Mod. Phys. **71**, 463–512 (1999).

[24] E. P. Gross, "Structure of a quantized vortex in boson systems", Il Nuovo Cimento **20**, 454–477 (1961).

[25] L. P. Pitaevskii, "Vortex lines in an imperfect bose gas", Soviet Physics JETP **13**, 451–454 (1961).

[26] S. Giorgini, L. P. Pitaevskii, and S. Stringari, "Condensate fraction and critical temperature of a trapped interacting bose gas", Phys. Rev. A **54**, R4633–R4636 (1996).

[27] M. J. Davis and P. B. Blakie, "Critical temperature of a trapped Bose gas: Comparison of theory and experiment", Phys. Rev. Lett. **96**, 060404 (2006).

[28] F. Gerbier, J. H. Thywissen, S. Richard, M. Hugbart, P. Bouyer, and A. Aspect, "Critical temperature of a trapped, weakly interacting Bose gas", Phys. Rev. Lett. **92**, 030405 (2004).

[29] M. Naraschewski and D. M. Stamper-Kurn, "Analytical description of a trapped semi-ideal bose gas at finite temperature", Phys. Rev. A **58**, 2423–2426 (1998).

[30] M. Kristensen, “Atom number fluctuations in Bose-Einstein condensates”, PhD thesis (Aarhus University, July 2018).

[31] Y. Castin and R. Dum, “Bose-Einstein condensates in time dependent traps”, Phys. Rev. Lett. **77**, 5315–5319 (1996).

[32] P. Arora, S. B. Purnapatra, A. Acharya, R. Kumar, and A. Sen Gupta, “Measurement of Temperature of Atomic Cloud Using Time-of-Flight Technique”, MAPAN **27**, 31–39 (2012).

[33] A. Griffin, “Conserving and gapless approximations for an inhomogeneous bose gas at finite temperatures”, Phys. Rev. B **53**, 9341 (1996).

[34] M. B. Christensen, “Microcanonical fluctuations in interacting Bose-Einstein condensates”, PhD thesis (Aarhus University, July 2020).

[35] M. Gajdacz, P. L. Pedersen, T. Mørch, A. J. Hilliard, J. Arlt, and J. F. Sherson, “Non-destructive faraday imaging of dynamically controlled ultracold atoms”, Rev. of Scient. Inst. **84**, 083105 (2013).

[36] M. A. Kristensen, M. Gajdacz, P. L. Pedersen, C. Klempt, J. F. Sherson, J. J. Arlt, and A. J. Hilliard, “Sub-atom shot noise faraday imaging of ultracold atom clouds”, Journ. of Phys. B **50**, 034004 (2017).

[37] G. Reinaudi, T. Lahaye, Z. Wang, and D. Guéry-Odelin, “Strong saturation absorption imaging of dense clouds of ultracold atoms”, Opt. Lett. **32**, 3143–3145 (2007).

[38] T. Vibel, M. B. Christensen, M. A. Kristensen, J. J. Thuesen, L. N. Stokholm, C. A. Weidner, and J. J. Arlt, “Spatial calibration of high-density absorption imaging”, Journ. of Phys. B **57**, 145301 (2024).

[39] R. W. Boyd, “Chapter 6 - nonlinear optics in the two-level approximation”, in *Nonlinear optics*, Third Edition (Academic Press, Burlington, 2008), pp. 277–328.

[40] J. D. Jackson, *Classical electrodynamics*, 3rd (John Wiley & Sons, New York, 1999).

[41] R. Meppelink, R. A. Rozendaal, S. B. Koller, J. M. Vogels, and P. van der Straten, “Thermodynamics of bose-einstein-condensed clouds using phase-contrast imaging”, Phys. Rev. A **81**, 053632 (2010).

[42] K. Konstantinou et al., *How interacting bose gases scatter light*, 2025.

[43] O. Morice, Y. Castin, and J. Dalibard, “Refractive index of a dilute bose gas”, Phys. Rev. A **51**, 3896–3901 (1995).

[44] K. S. Hardman, P. B. Wigley, P. J. Everitt, P. Manju, C. C. N. Kuhn, and N. P. Robins, “Time-of-flight detection of ultra-cold atoms using resonant frequency modulation imaging”, Opt. Lett. **41**, 2505 (2016).

[45] F. Kaminski, N. S. Kampel, M. P. H. Steenstrup, A. Griesmaier, E. S. Polzik, and J. H. Müller, “In-situ dual-port polarization contrast imaging of faraday rotation in a high optical depth ultracold 87rb atomic ensemble”, The Eur. Phys. Journ. D **66**, `10.1140/epjd/e2012-30038-0` (2012).

[46] A. Ramanathan, S. R. Muniz, K. C. Wright, R. P. Anderson, W. D. Phillips, K. Helmerson, and G. K. Campbell, “Partial-transfer absorption imaging: a versatile technique for optimal imaging of ultracold gases”, Rev. of Scient. Inst. **83**, 083119 (2012).

[47] L. D. Turner, K. P. Weber, D. Paganin, and R. E. Scholten, "Off-resonant defocus-contrast imaging of cold atoms", Opt. Lett. **29**, 232–234 (2004).

[48] L. D. Turner, K. F. E. M. Domen, and R. E. Scholten, "Diffraction-contrast imaging of cold atoms", Phys. Rev. A **72**, 031403 (2005).

[49] M. Pappa et al., "Ultra-sensitive atom imaging for matter-wave optics", New Journal of Physics **13**, 115012 (2011).

[50] A. Reinhard, J. Riou, L. A. Zundel, and D. S. Weiss, "Dark-ground imaging of high optical thickness atom clouds", Opt. Comm. **324**, 30–33 (2014).

[51] M. R. Andrews, M.-O. Mewes, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, "Direct, nondestructive observation of a bose condensate", Science **273**, 84–87 (1996).

[52] M. R. Andrews, D. M. Kurn, H.-J. Miesner, D. S. Durfee, C. G. Townsend, S. Inouye, and W. Ketterle, "Propagation of sound in a bose-einstein condensate", Phys. Rev. Lett. **79**, 553–556 (1997).

[53] C. C. Bradley, C. A. Sackett, and R. G. Hulet, "Bose-einstein condensation of lithium: observation of limited condensate number", Phys. Rev. Lett. **78**, 985–989 (1997).

[54] M. F. Fernandez, P. G. S. Dias, P. H. N. Magnani, M. do A. Martins, M. Hugbart, A. Cipris, P. W. Courteille, and R. C. Teixeira, *Phase-contrast imaging of a dense atomic cloud*, 2025.

[55] S. Kadlecek, J. Sebby, R. Newell, and T. G. Walker, "Nondestructive spatial heterodyne imaging of cold atoms", Opt. Lett. **26**, 137–139 (2001).

[56] P. F. Gao, G. Lei, and C. Z. Huang, “Dark-field microscopy: recent advances in accurate analysis and emerging applications”, An. Chem. **93**, PMID: 33620210, 4707–4726 (2021).

[57] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn, *Making, probing and understanding Bose-Einstein condensates*, 1999.

[58] L. Brossard, “Study of light-induced dipolar interactions in cold atoms assemblies”, PhD thesis (Université Paris-Saclay, June 2020).

[59] S. Jennewein et al., “Coherent scattering of near-resonant light by a dense microscopic cold atomic cloud”, Phys. Rev. Lett. **116**, 233601 (2016).

[60] J. Pellegrino, R. Bourgain, S. Jennewein, Y. R. P. Sortais, A. Browaeys, S. D. Jenkins, and J. Ruostekoski, “Observation of suppression of light scattering induced by dipole-dipole interactions in a cold-atom ensemble”, Phys. Rev. Lett. **113**, 133602 (2014).

[61] M. Maiwöger et al., “Observation of light-induced dipole-dipole forces in ultracold atomic gases”, Phys. Rev. X **12**, 031018 (2022).

[62] C. J. Foot, *Atomic physics*, Oxford Master Series in Physics (Oxford University Press, Aug. 2011).

[63] R. Finkelstein, S. Bali, O. Firstenberg, and I. Novikova, “A practical guide to electromagnetically induced transparency in atomic vapor”, New Journ. of Phys. **25**, 035001 (2023).

[64] M. Fleischhauer, A. Imamoglu, and J. P. Marangos, “Electromagnetically induced transparency: optics in coherent media”, Rev. Mod. Phys. **77**, 633–673 (2005).

[65] E. E. Mikhailov, I. Novikova, Y. V. Rostovtsev, and G. R. Welch, “Buffer-gas-induced absorption resonances in Rb vapor”, Phys. Rev. A **70**, 033806 (2004).

[66] T. Zanon-Willette, E. de Clercq, and E. Arimondo, “Ultrahigh-resolution spectroscopy with atomic or molecular dark resonances: exact steady-state line shapes and asymptotic profiles in the adiabatic pulsed regime”, Phys. Rev. A **84**, 062502 (2011).

[67] P. v. d. Straten, *Atoms and molecules interacting with light* (Oxford University Press, Feb. 2016).

[68] 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”, Nat. **397**, 594–598 (1999).

[69] A. Rastogi, E. Saglamyurek, T. Hrushevskyi, S. Hubele, and L. J. LeBlanc, “Discerning quantum memories based on electromagnetically-induced-transparency and autler-townes-splitting protocols”, Phys. Rev. A **100**, 012314 (2019).

[70] E. Saglamyurek, T. Hrushevskyi, L. Cooke, A. Rastogi, and L. J. LeBlanc, “Single-photon-level light storage in cold atoms using the autler-townes splitting protocol”, Phys. Rev. Res. **1**, 022004 (2019).

[71] E. Saglamyurek, T. Hrushevskyi, A. Rastogi, L. W. Cooke, B. D. Smith, and L. J. LeBlanc, “Storing short single-photon-level optical pulses in bose–einstein condensates for high-performance quantum memory”, New Journ. of Phys. **23**, 043028 (2021).

[72] E. Saglamyurek, T. Hrushevskyi, A. Rastogi, K. Heshami, and L. J. LeBlanc, “Coherent storage and manipulation of broadband photons via dynamically controlled autler–townes splitting”, Nat. Phot. **12**, 774–782 (2018).

[73] H. K. Andersen, “Bose-Einstein condensates in optical lattices”, PhD thesis (Aarhus University, Sept. 2008).

[74] J. F. Bertelsen, “Ultracold atomic gases, mixtures and molecules”, PhD thesis (Aarhus University, Aug. 2007).

[75] M. Gajdacz, “Non-destructive imaging and feedback with ultracold gases”, PhD thesis (Aarhus University, July 2015).

[76] S. S. Mai, “Dynamical control of matter waves in optical lattices”, PhD thesis (Aarhus University, Nov. 2010).

[77] P. L. Pedersen, “Multi-mode spin dynamics of a Bose-Einstein condensate in an optical lattice”, PhD thesis (Aarhus University, Aug. 2014).

[78] T. Vibel, “Statistical analysis of atom number fluctuations in quantum degenerate bose gases”, PhD thesis (Aarhus University, Feb. 2024).

[79] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, “Rapid production of ^{87}Rb bose-einstein condensates in a combined magnetic and optical potential”, Phys. Rev. A **79**, 063631 (2009).

[80] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell, “Simplified system for creating a Bose–Einstein condensate”, J. Low Temp. Phys. **132**, 309–367 (2003).

[81] D. A. Steck, *Rubidium 87 D line data*, http://steck.us/alkalidata/, 2019.

[82] K. B. Davis, M.-O. Mewes, M. A. Joffe, M. R. Andrews, and W. Ketterle, “Evaporative cooling of sodium atoms”, Phys. Rev. Lett. **74**, 5202–5205 (1995).

[83] W. Ketterle and N. J. V. Druten, *Advances in atomic, molecular, and optical physics* (Elsevier, 1996).

[84] A. Marte, T. Volz, J. Schuster, S. Dürr, G. Rempe, E. G. M. van Kempen, and B. J. Verhaar, “Feshbach resonances in rubidium 87: Precision measurement and analysis”, Phys. Rev. Lett. **89**, 283202 (2002).

[85] V. Negnevitsky and L. D. Turner, “Wideband laser locking to an atomic reference with modulation transfer spectroscopy”, Opt. Exp. **21**, 3103 (2013).

[86] J. Appel, A. MacRae, and A. I. Lvovsky, “A versatile digital ghz phase lock for external cavity diode lasers”, Meas. Scien. and Tech. **20**, 055302 (2009).

[87] A. Camacho-Guardian, K. K. Nielsen, T. Pohl, and G. M. Bruun, “Polariton dynamics in strongly interacting quantum many-body systems”, Phys. Rev. Res. **2**, 023102 (2020).

[88] S. Alexander, S. Markevich, and R. Fedaruk, “Multiphoton raman transitions and rabi oscillations in driven spin systems”, Phys. Rev. A **98**, 043814(1–8) (2018).

[89] V. Kumar, S. P. Ram, S. Singh, K. Bhardwaj, V. B. Tiwari, and S. R. M, *On damping of rabi oscillations in two-photon raman excitation in cold* ^{87}*rb atoms*, 2025.

[90] B. Yavkin et al., “Room temperature high-field spin dynamics of nv defects in sintered diamonds”, Appl. Mag. Res. **44**, `10.1007/s00723-013-0476-4` (2013).

[91] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, “Optical dipole traps for neutral atoms”, in , Vol. 42, edited by B. Bederson and H. Walther, Advances In Atomic, Molecular, and Optical Physics (Academic Press, 2000), pp. 95–170.

[92] H. P. Mishra, A. S. Flores, W. Vassen, and S. Knoop, “Efficient production of an 87Rb F = 2, mF = 2 Bose-Einstein condensate in a hybrid trap”, The Euro. Phys. Journ. D **69**, 52 (2015).

[93] D. M. Stamper-Kurn, M. R. Andrews, A. P. Chikkatur, S. Inouye, H.-J. Miesner, J. Stenger, and W. Ketterle, “Optical confinement of a bose-einstein condensate”, Phys. Rev. Lett. **80**, 2027–2030 (1998).

[94] M. Zaiser, J. Hartwig, D. Schlippert, U. Velte, N. Winter, V. Lebedev, W. Ertmer, and E. M. Rasel, “Simple method for generating bose-einstein condensates in a weak hybrid trap”, Phys. Rev. A **83**, 035601 (2011).

[95] L. N. Stokholm, “Implementing two laser based traps for cooling to quantum degeneracy”, MA thesis (Aarhus university, Sept. 2022).

[96] T. Esslinger, I. Bloch, and T. W. H., “Bose-einstein condensation in a quadrupole-ioffe-configuration trap”, Phys. Rev. A **58**, R2664–R2667 (1998).

[97] H. J. Metcalf and P. van der Straten, *Laser cooling and trapping* (Springer New York, 1999).

[98] B. T. Sheard, “Magnetic transport and bose-einstein condensation of rubidium atoms”, PhD thesis (University of Oxford, 2010), pp. 138–140.

[99] J. Dalibard and C. Cohen-Tannoudji, “Laser cooling below the doppler limit by polarization gradients: simple theoretical models”, J. Opt. Soc. Am. B **6**, 2023–2045 (1989).

[100] M. S. Hamilton, A. R. Gorges, and J. L. Roberts, “Influence of optical molasses in loading a shallow optical trap”, Phys. Rev. A **79**, 013418 (2009).

[101] *Flexdds-ng-dual manual*, July 2025.

[102] *Ad9910 data sheet*, July 2025.

[103] D. Genkina, L. M. Aycock, B. K. Stuhl, H.-I. Lu, R. A. Williams, and I. B. Spielman, “Feshbach enhanced s-wave scattering of fermions: direct observation with optimized absorption imaging”, New Journ. of Phys. **18**, 013001 (2015).

[104] K. Hueck, N. Luick, L. Sobirey, J. Siegl, T. Lompe, H. Moritz, L. W. Clark, and C. Chin, “Calibrating high intensity absorption imaging of ultracold atoms”, Opt. Exp. **25**, 8670–8679 (2017).

[105] D. B. Hume, I. Stroescu, M. Joos, W. Muessel, H. Strobel, and M. K. Oberthaler, “Accurate atom counting in mesoscopic ensembles”, Phys. Rev. Lett. **111**, 253001 (2013).

[106] W. Muessel, H. Strobel, M. Joos, E. Nicklas, I. Stroescu, J. Tomkovič, D. B. Hume, and M. K. Oberthaler, “Optimized absorption imaging of mesoscopic atomic clouds”, Appl. Phys. B **113**, 69–73 (2013).

[107] A. Putra, D. L. Campbell, R. M. Price, S. De, and I. B. Spielman, “Optimally focused cold atom systems obtained using density-density correlations”, Rev. of Scient. Instr. **85** (2014).

[108] J. Szczepkowski, R. Gartman, M. Witkowski, L. Tracewski, M. Zawada, and W. Gawlik, “Analysis and calibration of absorptive images of bose–einstein condensate at nonzero temperatures”, Rev. of Scient. Inst. **80** (2009).

[109] R. Veyron, V. Mancois, J.-B. Gerent, G. Baclet, P. Bouyer, and S. Bernon, “Quantitative absorption imaging: the role of incoherent multiple scattering in the saturating regime”, Phys. Rev. Res. **4**, 033033 (2022).

[110] G. Ness, A. Vainbaum, C. Shkedrov, Y. Florshaim, and Y. Sagi, “Single-exposure absorption imaging of ultracold atoms using deep learning”, Phys. Rev. Appl. **14**, 014011 (2020).

[111] M. Christensen et al., “Observation of microcanonical atom number fluctuations in a Bose-Einstein condensate”, Phys. Rev. Lett. **126**, 153601 (2021).

[112] M. A. Kristensen et al., “Observation of atom number fluctuations in a bose-einstein condensate”, Phys. Rev. Lett. **122**, 163601 (2019).

[113] J. Szczepkowski, R. Gartman, M. Witkowski, L. Tracewski, M. Zawada, and W. Gawlik, “Analysis and calibration of absorptive images of Bose-Einstein condensate at nonzero temperatures”, Rev. Sci. Instr. **80**, 053103 (2009).

[114] M. Gajdacz, A. J. Hilliard, M. A. Kristensen, P. L. Pedersen, C. Klempt, J. J. Arlt, and J. F. Sherson, “Preparation of ultracold atom clouds at the shot noise level”, Phys. Rev. Lett. **117**, 073604 (2016).

[115] I. Zebergs, “Progress report”, Progress report (Aarhus University, Aug. 2025).

[116] C. F. Ockeloen, A. F. Tauschinsky, R. J. C. Spreeuw, and S. Whitlock, “Detection of small atom numbers through image processing”, Phys. Rev. A **82**, 061606 (2010).

[117] R. Veyron, V. Mancois, J.-B. Gerent, G. Baclet, P. Bouyer, and S. Bernon, “Quantitative absorption imaging: The role of incoherent multiple scattering in the saturating regime”, Phys. Rev. Res. **4**, 033033 (2022).

[118] C. C. Gerry and P. L. Knight, *Introductory quantum optics* (Oxford University Press, 2012).

[119] T. S. do Espirito Santo, P. Weiss, A. Cipris, R. Kaiser, W. Guerin, R. Bachelard, and J. Schachenmayer, “Collective excitation dynamics of a cold atom cloud”, Phys. Rev. A **101**, 013617 (2020).