



Department of Physics
Cavendish Laboratory

Studies of a Homogeneous Bose Gas

Tobias Fabian Schmidutz

Trinity College

July 2014

This dissertation is submitted for the degree of
Doctor of Philosophy

Tobias Schmidutz, 2014

Cavendish Laboratory
University of Cambridge
J J Thomson Avenue
Cambridge CB3 0HE
UK

Declaration

I declare that this thesis is my own work and is not substantially the same as any that I have submitted or am currently submitting for a degree, diploma or any other qualification at any other university. No part of this thesis has already been or is being concurrently submitted for any such degree, diploma or any other qualification. This thesis does not exceed the word limit of sixty thousand words, including tables, footnotes, bibliography and appendices, set out by the Faculty of Physics and Chemistry.

9th of July 2014

Abstract

In this thesis we describe the first experimental realization of a homogeneous atomic Bose-Einstein condensate in three dimensions. We also outline our first experimental studies of this new system, which exhibit textbook-like thermodynamic and quantum-mechanical behaviour.

First, we describe a simple and compact single-chamber apparatus for robust production of harmonically trapped ^{87}Rb Bose-Einstein condensates. The apparatus is built from off-the-shelf components and allows the production of quasi-pure condensates of $> 3 \times 10^5$ atoms in < 30 s. Our system provides an excellent compromise between compactness, cost and performance.

We then detail the transfer of the atoms into our box potential, formed by a dark optical trap. Next, we characterize the homogeneity of the trapped gas and deduce that its behaviour is very close to that of an ideal homogeneous gas. The trapping potential is well approximated by a power-law of order 15, which constitutes only a minor deviation from a uniform box.

We conduct an in-depth thermodynamic study of the homogeneous gas. We characterize the critical point for condensation and observe saturation of the thermal component in a partially condensed cloud, in agreement with Einstein's textbook picture of a purely statistical phase transition. Further, we observe the quantum Joule-Thomson effect, manifesting itself as isoenthalpic cooling of an (essentially) ideal gas. We extract a Joule-Thomson coefficient $\mu_{JT} > 10^9$ K/bar, which is about 10 orders of magnitude larger than observed in classical gases.

Finally, we study the properties of the condensate ground state, using high-resolution Bragg spectroscopy. We directly observe the Heisenberg-limited momentum uncertainty of the condensed atoms for a range of box sizes, up to $70 \mu\text{m}$. We measure the condensate interaction energy with a precision of $k_B \times 100$ pK and study the dynamics of its free expansion upon release from the box potential. All our measurements are in good agreement with theoretical expectations for a perfectly homogeneous condensate, establishing the uniformity of our optical-box system on a sub-nK energy scale.

The main value of this work lies in adding the uniform trap to the toolkit of cold atom experiments. One of the major goals in the field of cold atoms is to establish a closer link to condensed matter problems. As many of these are inherently in a homogeneous environment, this work should constitute a substantial step towards this goal.

Acknowledgements

I would like to express my gratitude to everyone who helped and supported me through the years that lead to this thesis. Foremost, I would like to thank my supervisor Zoran Hadzibabic, who took me on as a green summer intern and who has been a great mentor throughout my PhD. His support, advice and enthusiasm have guided me through many scientific crises during the last years.

The foundations of our lab were laid by its first students, Naaman Tammuz and Robbie Campbell, whose friendliness and willingness to help was one of my reasons to join this group. I would also like to thank Stuart Moulder and Scott Beattie, who were always in the lab next doors, available for advice or a cup of coffee. Following in their footsteps as the steward of the “old lab”, Richard Fletcher has been a formidable proof-reader in-chief of this group and we have shared a quite a few cups of coffee discussing the lab exclusion principle,¹ Thanks also goes to Nir Navon and Martin Robert-de-Saint-Vincent, the most recent additions to our group, who brought a lot of knowledge and energy to the table. My gratitude extends to everyone else who was part of the team along the way, especially Phil Richman, Alexandre Dareau and Aziza Suleymanzade.

Alex Gaunt’s computer genius and experimental skills contributed a huge amount to this work. Beyond that, his service as the quiz master has done much to increase the history knowledge of the members of this group. I am hugely indebted to Robert Smith, whose knowledge, technical expertise and boundless helpfulness were always there to be called upon when things went wrong.

My largest thanks goes to Igor Gotlibovych, who has been my lab buddy for the last years. Having done all of this work jointly with him, this thesis is as much his as it is mine. He was not only a brilliant colleague, but also a great friend. Our shared sense of humour, the late-night discussions about the meaning of life and our common “affection” for the small machine have kept me sane over the last years. Also, I have learned a thing or two about boats in the process.

Last but not least, I would like to thank our fantastic group secretary Pam Hadder, whose friendliness and competency is unsurpassed in the Cavendish.

¹It is impossible for both labs to be in the same (functional) state at the same time.

Outside the lab, there are many people who made my time in Cambridge a great experience. Unable to mention everyone by name, I would like to thank the “Team Youth Hostel”, the Judo club and the Trinity crowd for many good times.

Special thanks goes to Janina for always being supportive and understanding during the last years, even when I was stressed, struck by graduation panic or working late-night shifts.

Most of all, I would like to thank my family, my parents Maria and Hubert, and my siblings Sonja and Florian, for their unconditional support and endless belief.

Contents

1	Introduction	11
1.1	Bose-Einstein Condensation	11
1.2	The box and the harmonic oscillator	12
1.3	Thesis outline	14
2	Bose-Einstein Condensation	15
2.1	Bose-Einstein Condensation in a homogeneous potential	15
2.2	Thermodynamics in Power-law potentials	16
2.2.1	The local density approximation	17
2.2.2	Thermodynamics in power-law potentials	18
2.2.3	Momentum and spatial distributions	19
2.2.4	The condensate wave function	20
2.3	Expansion and Thermometry of cold gases	21
2.3.1	Thermal gas	22
2.3.2	Degenerate gases	22
2.3.3	Condensed gases	23
2.3.4	Thermometry of partially condensed gases	24
2.3.5	Column densities	25
3	Atom Cooling and Trapping	27
3.1	Atomic properties of Rubidium	27
3.2	Laser cooling	28
3.3	Evaporative cooling	31
3.3.1	Evaporative cooling in Power-Law Traps	31
3.3.2	Crossed-Dipole Trap	34
3.3.3	Hybrid Trap	36
3.3.4	Homogeneous Trap	37
3.3.5	Conclusions	37

4	Experimental Setup and Methods	39
4.1	Vacuum System	39
4.2	Magneto-optical trap	41
4.2.1	Magnetic coils	41
4.2.2	Laser Cooling	43
4.3	Evaporative cooling	46
4.3.1	Transfer to the magnetic trap	47
4.3.2	RF Evaporation	48
4.3.3	Single-Beam Dipole Trap: “Hybrid Trap”	50
4.4	Imaging and Image processing	53
4.4.1	Absorption Imaging	53
4.4.2	Density calibration	54
4.4.3	High intensity imaging	55
4.5	Experiment control and data processing	58
4.5.1	Sequence setup and control	58
4.5.2	Imaging control	59
4.5.3	Analysis software	60
5	Creating a homogeneous Bose-Einstein Condensate	61
5.1	Dark Optical Trap	61
5.2	Magnetic fields	63
5.3	Loading of the Uniform trap	64
5.3.1	Transfer procedure	64
5.4	Evaporation and Condensation	67
5.5	Volume and in situ shape	68
5.5.1	Harmonic density profile	68
5.5.2	Uniform density profile	68
5.5.3	General power-law traps	69
5.6	Thermometry	70
5.6.1	Scaling laws	71
5.6.2	Numerical implementation	72
5.7	How uniform is the trap?	74
5.7.1	Conclusions from a single cloud profile	74
5.7.2	Conclusions from several profiles	75
5.7.3	The critical point	76
5.7.4	Finding the critical point	80
5.7.5	Conclusions	81

6	Thermodynamics in a uniform trap	83
6.1	Critical number	83
6.1.1	Experimental procedure	84
6.1.2	Self-consistency	85
6.1.3	Dependence of volume on temperature	86
6.1.4	Critical density	87
6.2	Saturation of the thermal component	87
6.2.1	Mean-Field theories	87
6.2.2	Experimental Results	93
6.2.3	Discussion	94
6.3	Quantum Joule-Thomson effect	94
6.3.1	Joule Thomson effect in a classical gas	96
6.3.2	Joule-Thomson effect in a condensed quantum gas	97
6.3.3	QJT in a degenerate thermal gas	99
6.3.4	Experimental observation of the QJT	100
6.3.5	Experimental details	101
6.3.6	Conclusions	104
7	Bragg spectroscopy	107
7.1	Theory of Bragg spectroscopy	108
7.1.1	Transfer function	109
7.1.2	Momentum distributions	112
7.1.3	Mean-Field corrections	114
7.2	Experimental realization	115
7.2.1	Setup	115
7.2.2	Heisenberg uncertainty relation	116
7.2.3	Mean-field shift	118
7.2.4	Evolution in ToF	120
7.3	Conclusions	123
8	Conclusions and Future Work	125
A	Polylogarithms and the Gamma function	127
B	IGBT driver	129
B.1	Functionality	129
B.2	Inputs	130
B.3	Outputs	131
B.4	Gate charging and protection	133

C	Design of a high-temperature superconducting magnet	137
C.1	Introduction	137
C.1.1	Trapping fields	137
C.1.2	Water-cooled Magnets	139
C.1.3	HTS Magnet requirements	139
C.2	Cryogenics	140
C.2.1	Current leads	141
C.2.2	Insulation	142
C.2.3	Filling system	148
C.3	Geometric design and fields	150
C.4	Coil electronics	153
C.4.1	Quench protection	153
C.4.2	Switching performance	154
C.5	Conclusions	157
C.6	Technical drawings of the cryostat	160

Chapter 1

Introduction

1.1 Bose-Einstein Condensation

The advent of quantum mechanics in the early 20th century has revolutionized the way we think about particles. Rather than considering them as rigid spheres bouncing around, they can be described as wave functions indicating the probability of finding them in a given location. As the name indicates, the energies quantum-mechanical particles can carry are not continuous, but discrete: they are quantized. A particle with a given energy and having the associated wave function is said to be in a quantum-mechanical state. While those assumptions already lead to surprising behaviour, such as tunnelling through seemingly impassable barriers, it gets even more confusing when considering groups of particles sharing a given space. Particles of the same type are fundamentally indistinguishable and exchanging them cannot alter the physical reality. This requirement leads to the conclusion that the total wave function only can get multiplied by a factor of ± 1 upon exchange of the particles. Particles can now be classified accordingly: Those with a symmetric wave function are called bosons, and those with an antisymmetric wave function are called fermions. This difference in overall symmetry has far-reaching consequences. In simple terms, bosons like to be near each other, while fermions do not. The Pauli exclusion principle states that two fermions cannot occupy the same state. If we now consider the ground state of a system- the collective configuration with the least total energy - with several particles of the same kind, we notice a fundamental difference which is illustrated in Fig. 1.1. In the fermionic case, there is one particle in each energy state, filling the states subsequently starting with the one with the least energy. Bosons, however, like to share and all of them jointly populate the lowest energy single-particle state.

In many-particle physics, we are usually dealing not with a few particles as shown in Fig. 1.1, but with thousands, millions or even more. At this point, it makes more sense to



Figure 1.1: Bosons and fermions in their respective ground states.

use a statistical approach and not track each particle individually. From the temperature of the ensemble, we can deduce a prediction of how many particles we expect in each state. For gases at high temperatures, the probability of finding an atom in a given state is very small and the difference between fermions and bosons becomes negligible. At low temperatures, we start to approach the zero temperature limit sketched in Fig. 1.1. In the case of Bosons¹, once we cool them enough, they start to occupy the lowest energy state in large numbers. Interestingly, we see significant occupation of the lowest state before we reach zero temperature; it occurs already at the so-called critical temperature. Under typical experimental conditions, this critical temperature can be hundred times larger than the energy gap between the lowest and first energy level. This effect is called Bose-Einstein condensation and was predicted as early as 1924 [1,2]. So how cold is cold enough to see this effect? We can estimate this by considering the spatial extent of the wave function of a single atom. For a given temperature T this can be approximated by the thermal wavelength Λ , which is defined as

$$\Lambda = \frac{h}{\sqrt{2\pi mkT}}, \quad (1.1)$$

where m is the mass of the atoms in question. Once Λ gets comparable to the inter-particle spacing, the wave functions overlap. Several particles share the ground state and Bose-Einstein condensation occurs. We can express this condition in terms of the phase-space density D

$$D \equiv n\Lambda^3 \approx 1. \quad (1.2)$$

This crude condition is surprisingly accurate in predicting the onset of condensation and is correct up to a numeric factor.

1.2 The box and the harmonic oscillator

To produce atomic Bose-Einstein condensates (BECs), we usually need to cool atoms into the nano-Kelvin regime. The technical challenges are enormous and it took until

¹While interesting things happen to Fermions as well, we are not concerned with them in this thesis.

1995 [3–5] to overcome them. The usefulness of atomic BECs stems from several factors: Unlike in superfluid ^4He , the atomic interactions are weak and can be neglected in many situations, or at least their effects can be described within simple mean-field frameworks. Even better, Feshbach resonances allow tuning the interactions in many atomic species [6] to the experimenters liking. This makes them a great model system for understanding many-body physics problems, such as superconductivity or two-dimensional systems.

We have not addressed one question yet: in which potentials are the atoms confined? During an undergraduate course in quantum mechanics, the first and simplest potential that is inevitably solved is the particle in a uniform (or homogeneous) box. The next major problem that pops up is the harmonic oscillator, which everybody will admit is much harder to solve. The same is true for BECs: Einstein’s original prediction was made for a homogeneous gas and theorists usually still prefer to solve problems in uniform problems. Ironically, it is experimentally much easier to create a harmonic trap than to create a uniform trap. Nature is generally very fond of smooth potentials. This is especially true when working at extremely low temperatures, where the smallest deviations destroy the homogeneity of the trap. For those reasons BECs have traditionally been produced in harmonic traps.

While the behaviour of an ideal quantum gas is well understood even in a harmonic potential, it gets complicated once interactions are thrown into the mix. Also, most solid-state systems, which cold gas researchers are trying to emulate, are homogeneous in nature. Often, the effect of the harmonic trap can be addressed - sometimes even in a beneficial way - by using the local-density approximation (LDA) [7–15]. Within the LDA, we treat the whole non-uniform trap as a collection of small homogeneous subsections of different depth. However, this is not always possible; integrating over the varying density can smear out or even qualitatively change experimental signatures (see, e.g., [10, 16–19]). Moreover, the LDA breaks down close to phase transitions, where the correlation length diverges [20]. In the drive towards a homogeneous trap, BECs have been loaded into modified geometries [21, 22]. However, those traps are uniform along only one direction; they are still harmonic along the other two directions.

In this thesis, we are going to describe the creation of the first cold atom experiment capable of producing three-dimensional homogeneous BECs. Apart from describing the technical challenges, we will describe the follow-up measurements we have made to quantify the homogeneity of our trap. Further, we have already been able to prove the usefulness of the uniform trap by measuring effects that have been obscured in harmonic traps.

1.3 Thesis outline

The main results of this thesis have been published in several papers. Below, we give an outline of the thesis and the respective publication for reference.

- Chapter 2: We summarize the essential theoretical foundations of the thermodynamics of quantum gases in power-law.
- Chapter 3: We introduce the theoretical foundation of cooling and trapping atomic gases.
- Chapter 4: We give an overview of the experimental setup and procedures we use to produce harmonic BECs. See also [23].
- Chapter 5: We describe the technical details of loading harmonic BECs in a uniform trap. We also discuss the characterization of the trap, proving its uniformity. See also [24].
- Chapter 6: We present our measurements of the thermodynamics in a box. We characterize the critical point for condensation and explore the influence of interactions on the thermal fraction. We also observe the quantum Joule-Thomson effect. For all of those effects, we recover ideal gas scalings despite weak interactions. See also [25].
- Chapter 7: We use Bragg spectroscopy to study the ground state properties of a uniform gas and measure the Heisenberg uncertainty relationship between space and momentum. Also, we measure the mean-field shift due to interactions. Finally, we explore the evolution of the condensate in time-of flight. See also [26].

These results are also documented in Igor Gotlibovych's thesis [27].

Chapter 2

Bose-Einstein Condensation

This chapter will provide a brief overview of the thermodynamics of Bose gases, both degenerate and classical in power-law traps. We consider the static behaviour in trap as well as the dynamics in time of flight (ToF) when released from the trap. Most of these results are well known and detailed derivations can be found in textbooks (e.g. [28]). We will summarize the relevant results as they are instrumental for the thermometry of gases, the analysis of the box trap, as well as for the following thermodynamic studies.

2.1 Bose-Einstein Condensation in a homogeneous potential

Here, we will briefly present the derivation of Bose-Einstein condensation in a box potential. This will mainly serve to introduce the basic concepts and variables, before moving on to power-law potentials. The starting point for most calculations is the Bose distribution $f(\epsilon)$ which predicts the occupation number of a state with energy ϵ for non-interacting Bosons in thermal equilibrium:

$$f(\epsilon) = \frac{1}{e^{\beta(\epsilon-\mu)} - 1}, \quad (2.1)$$

with $\beta = 1/k_B T$. μ denotes the chemical potential, which acts as a normalization constant for the total atom number N . μ must always be equal or smaller than the ground state energy ϵ_0 . Unless stated otherwise, we usually set $\epsilon_0 = 0$. The total atom number N is given by the sum over the occupation number of all states:

$$N = \sum_{\epsilon} f(\epsilon). \quad (2.2)$$

This can be approximated by converting the sum over the energy states to an integral over the density of states

$$D(\epsilon) = \frac{Vm^{3/2}}{2^{1/2}\pi^2\hbar^3}\epsilon^{1/2}, \quad (2.3)$$

where V is the volume of the box the particles are confined in. However, we need to treat the ground state population N_0 separately to avoid divergence for $\mu = 0$:

$$N \approx N_0 + \int_0^\infty f(\epsilon)D(\epsilon) d\epsilon \equiv N_0 + N'. \quad (2.4)$$

Here, we introduced the thermal number N' , counting all atoms in excited states. Analysis of the integral shows that N' reaches its maximum value $N_c(T)$ for a given temperature when $\mu = 0$. So what happens when we increase the atom number further? The atoms all have to accumulate in the ground state and $N_0 = N - N_c(T)$. This sudden onset of macroscopic occupation of the ground state is called Bose-Einstein condensation. When fixing the number instead of the temperature, we refer to the temperature at which this phase-transition happens as the critical temperature T_c . Solving Eq. (2.4) for the critical number yields

$$N_c(T) = N'(T, \mu = 0) = V \frac{\zeta(3/2)}{\Lambda^3}, \quad (2.5)$$

where $\zeta(3/2)$ is the Riemann zeta function, which is defined as the polylogarithm function $g_{3/2}(1)$. Appendix A contains the definitions of those functions, as well as some useful integrals commonly used in the solution of equations such as Eq. (2.4). We remember the estimate Eq. (1.2) which predicts a critical density of

$$n_c(T) = N_c/V \approx \Lambda^{-3}. \quad (2.6)$$

Comparison shows that the estimate is only off by a factor $\zeta(3/2) \approx 2.61$.

2.2 Thermodynamics in Power-law potentials

After having considered the onset of Bose-Einstein condensation in a uniform trap, we will now generalize these methods to derive the thermodynamics in general power-law potentials. Most importantly, we are going to introduce the semi-classical local-density approximation, allowing us to relate the thermodynamics in arbitrary potential to the uniform case. Further, we are going to consider the equilibrium in-trap density and momentum distributions. More detailed information can be found in [28, 29].

2.2.1 The local density approximation

We will discuss a degenerate Bose gas in an arbitrary power-law potential in three dimensions

$$U(\mathbf{r}) = \sum_{i=1}^3 \left| \frac{x_i}{L_i} \right|^{s_i}, \quad (2.7)$$

and derive thermodynamics scalings. For all calculations, we will neglect interactions with the ground state. The occupation of levels with energy ϵ is given by

$$f(\epsilon) = \frac{1}{e^{\beta(\epsilon-\mu)} - 1}. \quad (2.8)$$

To calculate meaningful thermodynamic quantities from this, we need to extract the density of states by solving the Hamiltonian

$$H = \frac{p^2}{2m} + U(\mathbf{r}) \quad (2.9)$$

for the eigenstates. While this can be done easily for harmonic or uniform potentials - in fact they are the first thing solved in most introductory textbooks - this is much harder for arbitrary potentials. Even worse, if we want to evaluate the density distribution $n(\mathbf{r})$ of our atoms in trap, we need to add up all the wave functions:

$$n(\mathbf{r}) = \sum_{\epsilon} f(\epsilon) |\psi_{\epsilon}|^2. \quad (2.10)$$

This is rather cumbersome (see [27] for a quantum mechanical treatment of a box potential). The full quantum-mechanical treatment can be avoided by using the semi-classical approximation, which assumes the gas behaves locally as a bulk gas. Often it is also referred to as the local-density approximation. It holds if the thermal wavelength is small compared to the length scale of the variation of the trap. Assuming the energy of classical particles

$$\epsilon(\mathbf{p}, \mathbf{r}) = \frac{p^2}{2m} + U(\mathbf{r}), \quad (2.11)$$

we can now modify Eq. (2.8) to

$$f(\mathbf{p}, \mathbf{r}) = \frac{1}{e^{\beta(\epsilon(\mathbf{p}, \mathbf{r}) - \mu)} - 1} = g_0(\exp[-\beta(\epsilon(\mathbf{p}, \mathbf{r}) - \mu)]), \quad (2.12)$$

where we introduced the polylogarithm function of order zero $g_0(\cdot)$. Its definition and properties are detailed in Appendix A.

To obtain thermodynamic quantities, we have to integrate over phase-space, e.g. we

can calculate the thermal number N' as

$$N' = \iint d\mathbf{r} \frac{d\mathbf{p}}{h^3} f(\mathbf{p}, \mathbf{r}) \quad (2.13)$$

Often, it is simpler to use the (semi-classical) density of states, which can be evaluated as [29, 30]

$$D(\epsilon) = \frac{1}{h^3} \iint \delta(\epsilon - U(\mathbf{r}) - \mathbf{p}^2/2m) d\mathbf{r} d\mathbf{p} = C_\alpha \epsilon^{\alpha-1}, \quad (2.14)$$

where C_α is a function of the trap dimensions. The scaling parameter α can be derived from the trap-exponents s_i as

$$\alpha = 3/2 + 1/s_1 + 1/s_2 + 1/s_3. \quad (2.15)$$

For a uniform gas, we have $s_i = \infty$ and hence $\alpha = 3/2$. The trap dimensions enter as

$$C_{3/2} = \frac{Vm^{3/2}}{2^{1/2}\pi^2\hbar^3}, \quad (2.16)$$

where V is the volume of the box. In a harmonic trap the potential is given by

$$U(\mathbf{r}) = \frac{1}{2}m \sum_{i=x,y,z} \omega_i^2 r_i^2, \quad (2.17)$$

where ω_i are the trapping frequencies and the scaling parameter $\alpha = 3$. The prefactor is given by

$$C_3 = \frac{1}{2\hbar^3\bar{\omega}^3}, \quad (2.18)$$

where $\bar{\omega}$ is the geometric mean of the individual trapping frequencies. Note that we can get the same results by deriving the density of states quantum-mechanically.

2.2.2 Thermodynamics in power-law potentials

With the semi-classical density of states at hand, we can derive thermodynamic quantities for arbitrary power-laws with the usual methods [28]. For now, we consider a partially condensed gas $T < T_C$, i.e. we can set $\mu = 0$. The critical number can then be calculated as:

$$N_c = \int d\epsilon D(\epsilon) f(\epsilon, \mu = 0) \quad (2.19)$$

$$= \int d\epsilon \epsilon^{\alpha-1} g_0 (e^{-\beta\epsilon}) \quad (2.20)$$

$$= C_\alpha \zeta(\alpha) \Gamma(\alpha) (kT)^\alpha, \quad (2.21)$$

where $\Gamma(\alpha)$ is the Gamma function (see Appendix A). Similarly we can obtain further thermodynamic properties of the partially condensed gas:

Table 2.1: Thermodynamic properties of a (partially) condensed gas in a power-law trap.

Property	Value
Critical number N_c	$C_\alpha \zeta(\alpha) \Gamma(\alpha) (kT_c)^\alpha$
Condensed Fraction N_0/N	$1 - \left(\frac{T}{T_c}\right)^\alpha$
Energy E	$C_\alpha \zeta(\alpha + 1) \Gamma(\alpha + 1) (kT)^{\alpha+1}$
Entropy S	$C_\alpha \frac{(\alpha+1)}{\alpha} \zeta(\alpha + 1) \Gamma(\alpha + 1) \frac{(kT)^{\alpha+1}}{T}$

We see that the parameter $\alpha = 3/2 + 1/s_1 + 1/s_2 + 1/s_3$ reflects the scaling of the density of states and will determine the relevant thermodynamic scalings of the gas. These results can be easily extended to cover thermal gases with $\mu < 0$ and $T > T_c$.

2.2.3 Momentum and spatial distributions

While the LDA is useful for determining thermodynamic properties of power-law potentials, this can be also achieved by solving the Hamiltonian to obtain the quantum-mechanical density of states. As mentioned before, its true strength lies in deriving local properties, e.g. the density and momentum distribution of the atoms in trap. This is applicable for the thermal atoms N' only; the condensate has a macroscopic wave function and behaves fully quantum-mechanically. Therefore, we will address the thermal density $n'(\mathbf{r})$ and the condensed density $n_0(\mathbf{r})$ separately.

The degenerate thermal gas

To calculate the in-trap density $n'(\mathbf{r})$, we integrate over the momentum in Eq. (2.12), using the integrals from Appendix A. We obtain

$$n'(\mathbf{r}) = \int_{-\infty}^{\infty} \frac{d\mathbf{p}}{h^3} g_0 \left(e^{-(\beta/2m)\mathbf{p}^2} A(\mathbf{r}, T) \right) \quad (2.22)$$

$$= \frac{g_{3/2}(A(\mathbf{r}, T))}{\Lambda^3}, \quad (2.23)$$

where we set $A = \exp(\beta\mu - \beta U(\mathbf{r}))$. We notice that each component of the momentum p_i^2 increased the power of the polylogarithm by (1/2) and scales the equation with the thermal

wavelength. This equation also underlines that condensation will occur locally as soon as the local phase-space density reaches ≈ 2.61 . Similarly, we can sequentially integrate the spatial components of Eq. (2.13) to obtain the in-trap momentum distribution:

$$n'(\mathbf{p}) = \int_{-\infty}^{\infty} g_0 \left(e^{-\beta \sum_{i=1}^3 \left| \frac{x_i}{L_i} \right|^{s_i}} A(\mathbf{p}, T) \right) d\mathbf{r} \quad (2.24)$$

$$\propto (kT)^\delta g_\delta (A(\mathbf{p}, T)), \quad (2.25)$$

where $A = \exp(\beta\mu - \beta\mathbf{p}^2/2m)$ and $\delta = \alpha - 3/2 = 1/s_1 + 1/s_2 + 1/s_3$. The prefactors were dropped for readability. We see that trap -power anisotropy gets ‘‘averaged-out’’ when considering the scaling of momentum distribution. This explains why a single parameter α captures the interesting physics; it tells us all we need to know about the momentum distribution.

2.2.4 The condensate wave function

The behaviour of the ground state of an interacting Bose gas can be described with the Gross-Pitaevskii equation (GPE):

$$\frac{\hbar^2}{2m} \nabla^2 \Psi(\mathbf{r}, t) + U(\mathbf{r}) \Psi(\mathbf{r}, t) + g |\Psi(\mathbf{r}, t)|^2 \Psi(\mathbf{r}, t) = i\hbar \frac{\partial \Psi(\mathbf{r}, t)}{\partial t}. \quad (2.26)$$

This is a non-linear Schroedinger equation for the wave function $\Psi(\mathbf{r}, t)$ of the condensate. The non-linearity is introduced by a mean-field term, which is proportional to the local condensate density

$$n(\mathbf{r}) = |\Psi(\mathbf{r}, t)|^2, \quad (2.27)$$

where the interaction parameter

$$g = 4\pi\hbar^2 a/m \quad (2.28)$$

scales with the s-wave scattering length a . In the static case, this can be written similar to the time-independent Schroedinger equation

$$\frac{\hbar^2}{2m} \nabla^2 \Psi(\mathbf{r}) + U(\mathbf{r}) \Psi(\mathbf{r}) + g |\Psi(\mathbf{r})|^2 \Psi(\mathbf{r}) = \mu \Psi(\mathbf{r}), \quad (2.29)$$

where the chemical potential $\mu = \partial E / \partial N$ takes on the role of the energy eigenvalues.

Thomas-Fermi approximation

We usually operate in a regime where the mean-field energy dominates over the kinetic energy. Neglecting the kinetic energy term is referred to as the Thomas-Fermi approxi-

mation and allows us to determine the density as

$$n_0(\mathbf{r}) = \max\left(\frac{\mu - U(\mathbf{r})}{g}, 0\right), \quad (2.30)$$

which means that the condensate “fills up” the potential up to the chemical potential. In a harmonic trap, the condensate has thus a parabolic profile, where the half-widths R_i are related to the frequencies w_i as

$$R_i^2 = \frac{2\mu}{mw_i^2}, \quad (2.31)$$

where μ can be found by equating the integral over the density to the total atom number.

In the uniform trap, the density is particularly simple to obtain:

$$n_0(\mathbf{r}) = \mu/g \quad (2.32)$$

and integration gives $\mu = gN_0/V$. Hence, within the TF-approximation, the density of an interacting condensate is homogeneous in a uniform trap at $T = 0$. This is in stark contrast to a non-interacting gas, where the wave function of the condensate takes the sine-like solution of a single particle in a box.

At the edges of the box, the density of the condensates jumps from zero to the finite value in Eq. (2.32). Therefore, the Thomas-Fermi approximation is clearly not applicable at the boundaries of the box as the kinetic energy $\propto \nabla|\Psi(\mathbf{r})|^2$ is infinite there. We have to therefore introduce a region at the boundary where the density varies smoothly from zero to its final value. This length scale is called the healing length ξ and can be estimated by equating the kinetic energy with the mean-field energy:

$$\frac{\hbar^2}{2m\xi^2} = ng. \quad (2.33)$$

The healing length also allows us to estimate the validity of the Thomas-Fermi approximation: once the healing length is on the order of the length of the box, we need to solve the full GPE (usually numerically) to obtain a density profile.

2.3 Expansion and Thermometry of cold gases

If we switch off our trap, the atomic gas expands. From the density distribution after a given time of flight, we can extract the in situ momentum distribution. This not only allows us to determine the temperature of the gas, but also gives us information about the trap shape. Assuming a classical ballistical expansion of the atoms, we can determine

from the in-trap distribution, $f(\mathbf{r}', \mathbf{p}, t = 0)$, the density $n(\mathbf{r}, t)$ at any time t as

$$n(\mathbf{r}, t) = \int \delta(\mathbf{p}t/m - (\mathbf{r} - \mathbf{r}')) f(r', p, t = 0) d\mathbf{r}' d\mathbf{p}. \quad (2.34)$$

In the following we will derive the appropriate expressions in various trap geometries in the thermal, degenerate and condensed regimes.

2.3.1 Thermal gas

Generally, we are only interested in thermal clouds for the optimization of the experimental sequence on the road towards degeneracy. The distribution of a thermal gas obeys Maxwell-Boltzmann statistics, i.e.

$$f(\mathbf{r}, \mathbf{p}) \propto \exp(-[p^2/2m + U(\mathbf{r})]\beta). \quad (2.35)$$

Using this distribution, it is easy to see that for long time of flight the density in Eq. (2.34) is only determined by the momentum distribution, which is Gaussian. We see that for thermal gases, the cloud shape loses the initial trap-shaped aspect ratio and expands isotropically. If the trap can be approximated as an anisotropic harmonic trap, we can derive an analytic form for any time of flight:

$$n(\mathbf{r}, t) = \frac{e^{\mu/k_B T}}{\lambda^3} \prod_i (1 + \omega_i^2 t^2)^{-0.5} \exp\left[\frac{-m\omega_i^2 r_i^2}{2k_B T(1 + \omega_i^2 t^2)}\right]. \quad (2.36)$$

2.3.2 Degenerate gases

In degenerate gases, we have to use the Bose-Einstein distribution in Eq. (2.34) and obtain

$$n'(\mathbf{r}, t) = \int g_0 \left(-\beta \left[\frac{m(\mathbf{r} - \mathbf{r}')^2}{2t^2} + U(\mathbf{r}') - \mu \right] \right) d\mathbf{r}'. \quad (2.37)$$

It is obvious that the limiting case for $t = 0$ is the in situ distribution of Eq. (2.23), while for large ToF the distribution is given by the momentum distribution in Eq. (2.25). The intermediate regime, where the initial size and the distance travelled are comparable, is not trivial though. While it is tempting to write the density distribution at any time as the convolution of the in-trap density distribution with the momentum distribution, the following equation is generally **not** true:

$$n'(\mathbf{r}, t) \propto \int n'(\mathbf{p} = (\mathbf{r} - \mathbf{r}')m/t) \times n(\mathbf{r}') d\mathbf{r}'. \quad (2.38)$$

The reason for this is that momentum and position of a single atom in trap are correlated in a power-law potential, and we cannot treat the expansion of the cloud as the sum of identical point-sources. There are two notable special cases: we can find a closed form for the harmonic trap which greatly simplifies fitting:

$$n'(\mathbf{r}, t) \propto \frac{1}{\Lambda^3} g_{3/2} \left(\exp \left[\mu - \sum_{i=1}^3 \frac{x_i^2}{R_i^2} \right] \right), \quad (2.39)$$

where the thermal radii are

$$R_i = \sqrt{\frac{1 + w_i^2 t^2}{\beta m w_i^2}} \quad (2.40)$$

determined by the trapping frequencies w_i and the time of flight. It is no coincidence that the harmonic equation can be solved analytically; we exploit the fact that momentum and potential enter the Hamiltonian with the same square power-law.

Second, in a uniform trap, the in-trap density is homogeneous. We use a potential

$$U(\mathbf{r}) = \begin{cases} 0 & \mathbf{r} \in V \\ \infty & \mathbf{r} \notin V \end{cases} \quad (2.41)$$

Therefore Eq. (2.38) does apply in this special case such that

$$n'(\mathbf{r}, t) \propto \int_V g_0 \left(-\beta \left(\frac{(\mathbf{r} - \mathbf{r}')^2 m}{2t^2} - \mu \right) \right) d\mathbf{r}' \quad (2.42)$$

$$= \int g_0 \left(-\beta \left(\frac{(\mathbf{r} - \mathbf{r}')^2 m}{2t^2} - \mu \right) \right) \times n'(\mathbf{r}, t = 0) d\mathbf{r}' \quad (2.43)$$

and so we can write the in-ToF distribution as the convolution of the momentum distribution and the in situ distribution.

2.3.3 Condensed gases

The situation is more complicated for BECs. Without interactions, their expansion can be calculated by propagating the ground state wave function with the free-space Schroedinger equation. Once we include repulsive interactions (for ^{87}Rb), we need to use the Gross-Pitaevskii equation (Eq. (2.26)) in the same fashion. For most potentials, these solutions have to be solved numerically.

Even though we never actually produce a non-interacting condensate, it is instructive to consider the expansion from a harmonic potential without interactions. The ground

state is simply a Gaussian

$$\Phi(r)_0^2 \propto \prod_i e^{-x_i^2/2\Delta x_i^2} \quad (2.44)$$

with $\Delta x_i^2 = \hbar/(mw_i)$. The distribution in momentum space is given by the Fourier transform, which is again Gaussian. The width in momentum space is hence given by $\Delta p_i \Delta x_i = \hbar/2$. In an anisotropic potential, the tight directions will therefore have the greater momentum width. Expansion will thus happen anisotropically, inverting the aspect ratio of the cloud after a long time of flight. The width σ_i of the Gaussian-shaped cloud in ToF is [31]

$$\sigma(t)_i^2 = \frac{2\hbar}{mw_i}(1 + w_i^2 t^2). \quad (2.45)$$

This aspect ratio inversion is a common feature for non-interacting, anisotropic condensates irrespective of the trapping potential. In the ground state, the uncertainty principle dictates $\Delta x_i \Delta p_i \gtrsim \hbar/2$. Hence, if we compress one direction in situ, it will necessarily expand more in ToF.

In our experiments, we cannot neglect interactions and we need to numerically invoke the full GPE to calculate the expansion. An exception is the harmonic trap, for which the cloud retains its parabolic shape during expansion [32,33]. Conveniently, the aspect ratio of the rather common cigar-shaped, harmonic traps ($w_1 = w_2 \gg w_3$) inverts as well. However, this time the inversion is not driven by the ground-state momentum, but by the interaction energy. In simple terms, the gradient in density and hence acceleration during ToF will be strongest along the short directions cause the inversion after long enough ToF.

We conclude that the inversion of the aspect ratio is a strong indicator of Bose-Einstein condensation; it can be driven by either the Heisenberg uncertainty of the ground state or by the mean-field acceleration originating from atomic interactions.

2.3.4 Thermometry of partially condensed gases

So far, we have only considered the thermal and condensed gases separately. When considering partially condensed gases, we usually operate under the assumption that the thermal and condensed components of the gas do not affect each other. We should note that this assumption is not strictly true. Especially in harmonic traps, the expected atom numbers deviate from the ideal gas prediction due to the mean-field energy of the condensate. However, we can still use the wings of the thermal component to deduce the temperature of the sample. These effects of the condensate on the thermal atoms have been scrutinised in a series of papers from our group [10, 16, 34].

Harmonic Trap

In a harmonic trap, we obtain a bimodal distribution in ToF where the thermal polylogarithmic distribution (Eq. (2.39)) and the parabolic condensate profile (Eq. (2.30)) are superimposed on top of each other. Using a multi-pass fitting algorithm, we can use the thermal wings to deduce the temperature and the shape of the thermal and condensed fractions to determine the respective atom numbers. The procedure is well-established and is described in depth in [27, 31, 35, 36].

General Power-Law Traps

For general power-law traps, the shape of the expanding condensate has no known analytical solution. While numerical solutions can be obtained using the GPE, using them for fitting is highly impractical. We therefore resorted to fitting the thermal part of the wings and extrapolating them into the region with a condensate. The condensate number can then be obtained by counting the total number and subtracting the fitted thermal part. We will discuss this procedure in detail in Chapter 5 for quasi-homogeneous traps.

2.3.5 Column densities

Experimentally, we cannot observe three-dimensional densities $n(x, y, z)$; instead we observe the two-dimensional column density $\tilde{n}(x, y)$, which is defined as

$$\tilde{n}(x, y) = \int n(x, y, z) dz. \quad (2.46)$$

Fortunately, integration is pretty straightforward for most density profiles. Often they are separable along the trap axis; in the case of polylogarithms they can be evaluated using the integrals in appendix A. The experimental details of measuring the column densities with absorptive imaging can be found in Section 4.4.

2.3. EXPANSION AND THERMOMETRY OF COLD GASES

Chapter 3

Atom Cooling and Trapping

Often, BECs are created using alkali atoms, due to their simple hydrogen-like atomic structure. An atom source is placed in a vacuum system and the atoms are trapped with a magneto-optical trap (MOT). Laser cooling allows cooling down to the μK range. Further cooling is achieved by forced evaporation down to the nK range until condensation occurs. We will briefly review the theory for the cooling mechanisms with the focus on the relevant quantities for our experiments. In-depth explanations can be found in [28,31,37,38].

3.1 Atomic properties of Rubidium

Alkali energy levels As ^{87}Rb has only one electron in the outermost shell, it has a hydrogen-like level structure. Still, analytical calculation of its energy levels is not possible and exact results rely on numerical analysis and experimental measurements. Unlike in hydrogen, the energy levels in alkalis are not degenerate in the angular momentum L , even before taking fine structure corrections into account. Taking the interaction between the spin S and angular momentum L into account, the energy levels are best described using

$$\mathbf{J} = \mathbf{L} + \mathbf{S}, \quad (3.1)$$

where $|\mathbf{J}| = \sqrt{J(J+1)}\hbar$ and similarly for \mathbf{L} and \mathbf{S} . Taking the atomic nuclear magnetic momentum \mathbf{I} into account, we obtain the hyperfine splitting which is parametrized by $\mathbf{F} = \mathbf{J} + \mathbf{I}$. The hyperfine levels are degenerate in the eigenvalues $m_F = -F, -F + 1, \dots, F - 1, F$ of the z-component of \mathbf{F} , where we assume the z-axis is the quantisation axis of the system. In the presence of small magnetic fields B_z along this quantisation axis, the Zeeman effect breaks this degeneracy. The energy difference ΔE from the original hyperfine level is given by

$$\Delta E = \mu_B g_F m_F B_z, \quad (3.2)$$

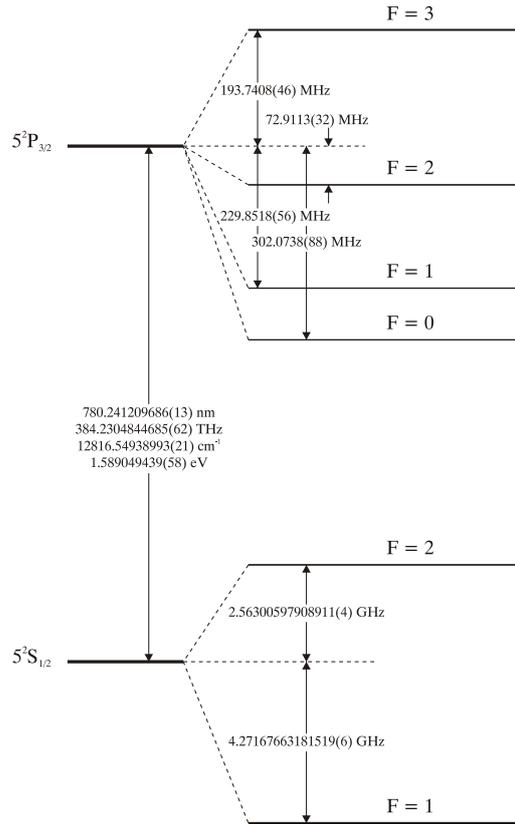


Figure 3.1: The ^{87}Rb D_2 energy level scheme. Data taken from [39], drawing taken from [40].

where $\mu_B = 9.27 \times 10^{-24}$ J/T is the Bohr magneton g_F are the numerical Lande factors.

^{87}Rb levels For experimental purposes, we are mainly interested in the D_2 line ($5^2S_{1/2} \rightarrow 5^2P_{3/2}$) at approximately 780 nm. The D_1 line ($5^2S_{1/2} \rightarrow 5^2P_{1/2}$) is located at 795 nm and is experimentally irrelevant. For most intents and purposes we can consider our atoms as a two-level system, split only at the hyperfine structure level. For simplicity, we refer to the levels of $|5^2S_{1/2}, F = 1, 2\rangle$ simply as $F = 1, 2$ and add a dash $F' = 0, 1, 2, 3$ to refer to $|5^2P_{3/2}, F = 0, 1, 2, 3\rangle$. The D_2 energy levels of ^{87}Rb are shown in figure 3.1. An excellent compilation of data on ^{87}Rb can be found in [39].

3.2 Laser cooling

The first cooling step in BEC experiments is laser cooling. A photon with momentum $\hbar\mathbf{k}$ transfers momentum to an atom when it is absorbed. If we consider a photon beam, we obtain a net time-averaged force $\mathbf{F}_{SC} = \hbar\mathbf{k}\Gamma_{SC}$, where we used the scattering rate Γ_{SC} and \mathbf{k} is the wavevector of the light. The atom will decay back to its ground state by spon-

taneous emission at a rate Γ . Since the emitted photon will leave in a random direction, this contribution will average out over time and only the absorption will contribute to the average force. By tweaking this scattering force so it opposes the direction of movement, one can slow the atoms. This technique is called optical molasses. Adding a carefully designed spatial dependency, the atom can be confined in a trap, as photons will push them back when they try to escape. This is done in a magneto-optical trap (MOT). For simplicity, we will only consider an imaginary two-level $F=0$ to $F'=1$ energy scheme.

Optical molasses The absorption rate and thus the force of a low intensity laser beam on an atom can be derived by using the optical Bloch equations. We assume that the laser light has a wavenumber k and the corresponding angular frequency $\omega = ck$. If the resonance frequency of the $F=0$ to $F'=1$ transition is ω_0 , the detuning is given $\delta = \omega - \omega_0 + \mathbf{k}\mathbf{v}$, where v is the atom velocity and the term $\mathbf{k}\mathbf{v}$ accounts for the Doppler effect. The scattering force is given by [37]:

$$\mathbf{F}_{SC} = \hbar\mathbf{k} \frac{\Gamma}{2} \frac{I/I_{sat}}{1 + I/I_{sat} + 4\delta^2/\Gamma^2}, \quad (3.3)$$

where the saturation intensity is given by $I_{sat} = 2\pi^2\hbar\Gamma c/3\lambda^3$. In one dimension, overlapping two red-detuned (i.e. $\omega_0 - \omega > 0$), counter-propagating beams ensures that the beam opposing the velocity of the atom gets absorbed more due to the Doppler shift. An atom moving through the beams will experience a velocity-dependent damping force, hence the name optical molasses. The principle can be easily extended to three dimensions by using six beams.

The Magneto-Optical Trap (MOT) In optical molasses, atoms are slowed down, but they are not trapped and will eventually escape the area exposed to the laser beams. We therefore introduce a restoring force, which pushes the atoms back towards the centre of the trap. This is done by applying a magnetic quadrupole field, which has a magnetic zero at the trap centre and increases linearly along the axes: $B_x \propto x$, $B_y \propto y$ and $B_z \propto z$ (see Fig. 3.2a). Again, we will consider only the one-dimensional case for simplicity. The excited level $F'=1$ splits due to the Zeeman effect (Eq. (3.2)) into non-degenerate $m_F = -1, 0, +1$ states. The energy splitting ΔE between the states increases linearly with the distance from the trap centre, $\Delta E \propto z$, this is illustrated in Fig. 3.2b. We use circularly polarized light σ^+ and σ^- , which addresses only transitions with $\Delta m_F = \pm 1$ respectively. Away from the trap centre, predominantly the polarization corresponding to the m_F state which is closer to resonance will be absorbed. By choosing the polarization of the beams correctly (see figure Fig. 3.2a for more details), we make sure that predominantly the

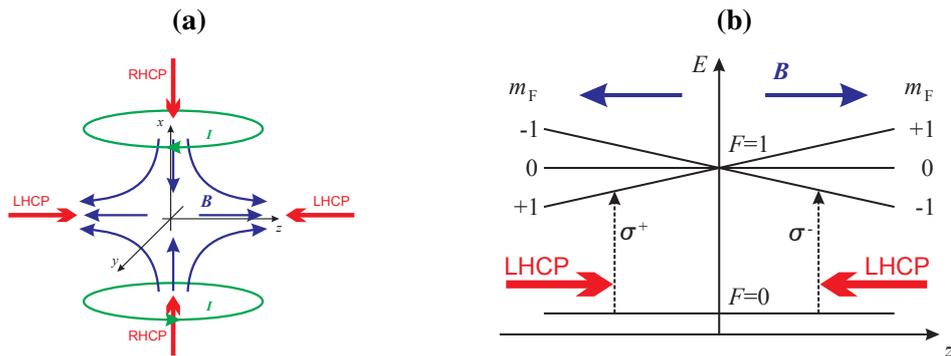


Figure 3.2: Operation principle of a MOT. (a) A pair of coils (current direction in green) in anti-Helmholtz configuration creates a quadrupole magnetic field (blue) at the atoms. We show the correct polarizations of the laser beams (red) needed to create a MOT. The fields and polarizations along the y-axis are identical to those shown along the z-axis. Drawing taken from [40]. (b) The operation principle of a MOT. The magnetic field splits the energy levels according to the Zeeman effect. Circularly polarized light σ^\pm selectively drives $m_F = \pm 1$ transitions. Combined with the spatial dependence of the energy levels, this results in preferential absorption of the beam which pushes the atoms back towards the trap centre. Drawing taken from [40]. Note that 'left/right-hand circularly polarized' (LHCP/RHCP) refers to the direction of propagation of the *light*; σ^\pm is defined with respect to the *quantization axis* of the atoms (here z).

beams pushing the atoms back get absorbed. Thus, we have created a trap.

If the experimental parameters are chosen right, we can combine the cooling of the molasses with the magneto-optical trapping. Mathematically, we describe this by including the Zeeman shift in Eq. (3.2) in the detuning δ in Eq. (3.3). Near the trap centre, this results in a restoring force linear in the displacement z and opposed to the atom velocity v :

$$F_{MOT}(v, z) \approx -\alpha v - \beta z, \quad (3.4)$$

where α and β absorb experimental parameters [37]. Again, the principle is easily extendible to three dimensions.

In practice, the MOT beams and fields are switched on in a vacuum with a small, residual ^{87}Rb background pressure. Slow enough atoms straying into the area where Eq. (3.4) applies are cooled and trapped. Over time, the number of atoms in the MOT grows until the capture rate is balanced by loss processes (e.g. collisions with hot non-trapped atoms).

Limits of cooling The theoretical limit of the laser cooling mechanism as described so far is given by the Doppler temperature T_D ,

$$k_B T_D = \frac{\hbar \Gamma}{2} \quad (3.5)$$

and is caused by the statistic nature of the absorption and emission momentum kicks [28]. For Rubidium, this temperature is $T_D = 144\mu\text{K}$. To the surprise of the first experimenters and in violation of Murphy's law, temperatures below this were achieved, thanks to a process called Sisyphus cooling. The process relies on having a Zeeman structure in the ground state and cannot be explained in the simple $F=0$ and $F=1$ energy scheme employed so far. We are not going to go into detail here, see [28, 37, 41] for more information. For us, it is important that this process allows us to beat the Doppler temperature and is (theoretically) limited by the recoil energy of a single photon. This corresponds to the so-called recoil temperature T_r , with

$$T_r = \frac{(\hbar k)^2}{2mk_B}. \quad (3.6)$$

For ^{87}Rb , $T_r \approx 362\text{ nK} \ll T_D$. The Doppler limit gives a good estimate for experimentally viable MOT temperatures. Including a brief molasses stage after cooling the atoms in a MOT, i.e. switching off the magnetic confinement, allows us to take advantage of the Sisyphus cooling process. While experimentally we do not even get close to T_r , we do cool significantly below T_D .¹

3.3 Evaporative cooling

3.3.1 Evaporative cooling in Power-Law Traps

After the MOT stage, the atoms are loaded into a new trapping potential for further cooling. A range of different trapping potentials exist, most of them based on optical or magnetic confinement. As we are still far from degeneracy (where the phase-space density $D = n\Lambda^3 \approx 1$), we need to cool the atoms further. A trapping potential of finite depth U_0 allows the most energetic atoms with energy $k_B T > U_0$ to escape, truncating the Boltzmann distribution. Elastic collisions then rethermalize the remaining atoms and the average temperature drops. This process is called evaporative cooling.

However, once $k_B T \ll U_0$, the cooling process will come to an asymptomatic standstill. Hence, to keep the cooling process going (and produce BECs in a finite time), we lower the trapping potential continuously. The slower we reduce U_0 , the fewer atoms we need to expend to cool by a given amount. In a perfect system we could reduce the trap height infinitely slowly, allowing us to cool the gas while losing only a negligible number of atoms. However, unwanted atom loss and heating processes compete with the cooling and continuously reduce the phase-space density. Hence, to increase D , we are forced to

¹Typically, we achieve $T \approx 15\ \mu\text{K}$ post-molasses; see Chapter 4.

3.3. EVAPORATIVE COOLING

cool at finite rate. We will derive a simple model following [28, 42] to derive the principles which govern the process of evaporative cooling in a range of power-law trapping potentials. Let us consider an ideal, non-degenerate gas in a power-law potential:

$$U(\mathbf{r}) = \sum_{i=1}^3 U_i |r_i|^{s_i} \quad (3.7)$$

We recall the semi-classical density of states for a power-law potential from Eq. (2.14)

$$D(\epsilon) = C_\alpha \epsilon^{\alpha-1}, \quad (3.8)$$

as well as the power-law parameter $\alpha = 3/2 + 1/s_1 + 1/s_2 + 1/s_3$. For a classical gas it is convenient to consider the single-particle partition function

$$Z_1 = \iint e^{(-U(\mathbf{r}) - p^2/2m)/k_B T} d\mathbf{r} d\mathbf{p} \quad (3.9)$$

$$= \int D(\epsilon) e^{-\epsilon/k_B T} d\epsilon \quad (3.10)$$

$$= C_\alpha \Gamma(\alpha) (k_B T)^\alpha. \quad (3.11)$$

Alternatively we can evaluate the spatial and momentum integrals separately and write

$$Z_1 = V_0 / \Lambda^3, \quad (3.12)$$

where we define the effective volume as

$$V_0(T) = \int e^{-U(\mathbf{r})/k_B T} d^3 r \quad (3.13)$$

and we find that the momentum integral can be expressed in terms of the thermal wavelength Λ

$$\int \exp[(-p^2/2m)/k_B T] d\mathbf{p} = \Lambda^{-3}. \quad (3.14)$$

We now consider the in situ density distribution for N particles

$$n(\mathbf{r}) = N \frac{\int \exp[(-U(\mathbf{r}) - p^2/2m)/k_B T] d\mathbf{p}}{Z_1} \quad (3.15)$$

$$= \frac{N}{V_0} e^{-U(\mathbf{r})/k_B T}. \quad (3.16)$$

The peak density in the centre of the trap is trivially

$$n(0) = N/V_0. \quad (3.17)$$

The energy E of the gas can be calculated as

$$E = \frac{N}{Z_1} \int \epsilon D(\epsilon) e^{-\epsilon/k_B T} d\epsilon \quad (3.18)$$

$$= \alpha N k_B T. \quad (3.19)$$

We consider now a trap containing particles with an average energy $\bar{\epsilon} = E/N$. The trap is truncated in a way that particles with energies $U_0 = (\beta + 1)\bar{\epsilon}$ can escape, where $\beta \gg 1$ ². We assume that β is large enough to justify the assumption that all escaping particles have energy U_0 . We can write for the change in total energy

$$\dot{E} = \dot{N}U_0 = \dot{N}\bar{\epsilon}(\beta + 1). \quad (3.20)$$

We can also write

$$\dot{\epsilon} = \frac{\dot{E}}{N} - \frac{\dot{N}E}{N^2}. \quad (3.21)$$

Combing both equations and eliminating the time dependence yields

$$\frac{d \ln \bar{\epsilon}}{d \ln N} = \beta. \quad (3.22)$$

Eq. (3.19) states that the temperature is proportional to $\bar{\epsilon}$ and hence

$$\frac{d \ln T}{d \ln N} = \beta. \quad (3.23)$$

To achieve degeneracy, we aim to optimize the peak phase-space density,

$$D_0 \equiv n(0)\Lambda^3 = \frac{N}{Z_1} \propto NT^{-\alpha} \quad (3.24)$$

We can derive a scaling constant γ for the evolution of the phase-space density

$$\gamma = \frac{d \ln D_0}{-d \ln N} = \beta\alpha - 1. \quad (3.25)$$

We see that a low power-law trap is advantageous to achieve large gains in phase space density. Experimentally, γ is used as a figure of merit for the evaporation process, even though the derivation above uses very simplified assumptions. So far, we have not taken losses into account and evaporating slowly maximizes efficiency.

To tackle this issue, we introduce the half-life time of the atoms due to evaporative cooling τ_E , and a non-energy-selective loss-lifetime τ_L . Doing so modifies Eq. (3.23)

²Note, that here β is a number and **not** $1/k_B T$.

to [28]

$$\frac{d \ln T}{d \ln N} = \beta \frac{\tau_L}{\tau_L + \tau_E} = \beta'. \quad (3.26)$$

Similarly we have to replace $\beta \rightarrow \beta'$ in Eq. (3.25). We see that losses reduce our evaporation efficiency. A further requirement to keep in mind is that the evaporation speed is limited by the rate at which atoms are scattered into the untrapped states with $\epsilon > U_0$. This rate is comparable to the rate of elastic collisions τ_C in the trap centre

$$1/\tau_C = n(0)\sigma v_{rel} \propto n(0)T^{1/2}, \quad (3.27)$$

where we used that the relative speed between atoms scales as $v_{rel} \propto \sqrt{T}$. We also assume a temperature independent s-wave scattering rate, which is applies in ^{87}Rb for $T < 25 \mu\text{K}$ [43].

A more demanding limitation on the evaporation efficiency is imposed by requiring that the scattering rate increases throughout the evaporation process. In this regime, which is called runaway evaporation, our efficiency γ increases throughout the evaporation process. The change in elastic scattering rate can be quantified by

$$\gamma_{coll} = \frac{\ln d(1/\tau_C)}{-d \ln N} = \beta' (\alpha - 2) - 1. \quad (3.28)$$

Runaway evaporation occurs for $\gamma_{coll} > 0$. If we cannot achieve runaway evaporation, the evaporation efficiency γ is going to decrease over time and the process will become more and more inefficient. From Eq. (3.28), we see that in three dimensions we cannot enter the runaway regime for $\alpha < 2$, i.e., for power-law potentials $U(\mathbf{r}) \propto \mathbf{r}^s$ with $s > 6$. However, we should note this is not prohibitive to achieving condensation for higher-order traps. We simply have to make sure that we start with a sufficiently high collision rate and atom number.

In the following sections, we consider a few traps relevant to our experiment. We derive their respective coefficients β and γ_{coll} , which allows us to predict the performance of the different trap types for evaporative cooling.

3.3.2 Crossed-Dipole Trap

We will now consider the optical crossed-dipole trap. Commonly formed by two overlapping Gaussian beams, the AC stark shift introduces a potential to the atoms proportional to the beam intensity [28]. This potential is approximately harmonic in the beam centre. However, it has a finite depth and lowering it necessarily decompresses the trap. We will show that runaway evaporation is not possible in a simple isotropic configuration. However, as mentioned above, this is not a deal-breaker; condensation has been achieved in

all-optical trap configurations [44–46]. Often they are used as a final ‘storage’ trap after evaporation in a more efficient configuration. They offer many benefits, especially allowing us to use magnetic fields freely, e.g. for tuning interactions with Feshbach resonances. The derivation below follows mainly [47], although with the notation in line with [28]. An optical crossed-dipole trap can be described as

$$U(\mathbf{r}) = U_0 \exp[-2(x^2 + y^2 + z^2)/w_0^2], \quad (3.29)$$

where we neglected the weak longitudinal confinement along the beams and anisotropies. The depth of this trap is given by U_0 and the waist by w_0 . For $k_B T \ll U_0$ the potential is approximately harmonic

$$U(\mathbf{r}) = \frac{1}{2} m \omega^2 r^2, \quad (3.30)$$

with $\omega = \frac{2}{w_0} \sqrt{\frac{U_0}{m}}$. With Eq. (3.30), the effective Volume can be calculated as

$$V_0 = \left(\frac{2\pi k_B T}{m} \right)^{3/2} \frac{1}{\omega^3}. \quad (3.31)$$

In an infinitesimal evaporation step, two things happen at once. The total energy is reduced due to the evaporated atoms leaving the trap, and also decreased due to the adiabatic decompression of the trap. The energy change due to evaporation is given by Eq. (3.20) with $\alpha = 3$

$$\dot{E}_1 = \dot{N}(\beta + 1)3k_B T. \quad (3.32)$$

The decompression is adiabatic, and conserves entropy S and also peak phase space density D_0 [48]

$$\left(\frac{dD_0}{dt} \right) = 0. \quad (3.33)$$

From the definition of D_0 in Eq. (3.24) follows

$$D_0 = \frac{N}{V_0} \Lambda^3 \propto \left(\frac{w}{T} \right)^3 \propto \frac{U_0^{3/2}}{E}. \quad (3.34)$$

Using Eq. (3.33), we can derive the change in energy due to decompression

$$\dot{E}_2 = \frac{E}{2} \frac{\dot{U}_0}{U_0}. \quad (3.35)$$

We assume that the temperature is proportional to the trap depth $U_0 \propto T$. The total energy change rate is given by

$$\dot{E} = \dot{E}_1 + \dot{E}_2. \quad (3.36)$$

Solving this for N and T yields

$$\frac{d \ln(T)}{d \ln(N)} = 2\beta. \quad (3.37)$$

We can now calculate the change in phase space density as in Eq. (3.25), but we need to take into account that the effective trap volume depends on ω . Lowering the trap depth U_0 throughout the evaporation process according to $U_0 \propto k_B T$ decompresses the trap as $\omega \propto U_0^{1/2} \propto T^{1/2}$. The inherent temperature dependence of the effective volume now cancels with the temperature dependence of the trapping frequency:

$$V_0 \propto \frac{T^{3/2}}{\omega(T)^3} \propto T^0. \quad (3.38)$$

The gain in phase space density by lowering the number is then

$$\gamma = -\frac{d \ln D_0}{d \ln N} = -\frac{d \ln \Lambda^3 N V_0^{-1}}{d \ln N} = 3\beta - 1, \quad (3.39)$$

which is the same result as in a harmonic trap without decompression. Given that adiabatic decompression conserves phase-space density, this is not surprising. However, the collision rate is reduced to

$$\gamma_{coll} = -\frac{d \ln (1/\tau_C)}{d \ln N} = -\frac{d \ln (T^{1/2} N V_0^{-1})}{d \ln N} = -\beta - 1. \quad (3.40)$$

We see that γ_{coll} is negative and runaway evaporation is not possible.

3.3.3 Hybrid Trap

We consider now a harmonic trap, which decompresses only in two directions, but retains its trapping frequency along the third direction. This is a good approximation to a hybrid optical-magnetic trap, consisting of a single beam optical-dipole trap with magnetic confinement along the third direction. In our setup, we use this kind of trap (see Section 4.3.3) as the last evaporation stage to achieve a harmonically trapped BEC. As the decompression only occurs along two axes, Eq. (3.31) is modified to

$$V_0 \propto T^{3/2} \omega^{-2} \propto T^{3/2} U_0^{-1}, \quad (3.41)$$

where $U_0 \propto T$ is again the trap depth. Decompression is again adiabatic and γ does not change

$$\gamma = -\frac{d \ln D_0}{d \ln N} = 3\beta - 1. \quad (3.42)$$

The coefficient γ_{coll} is trivial to calculate

$$\gamma_{coll} = -\frac{d \ln (T^{1/2} N V_0^{-1})}{d \ln N} = -\frac{d \ln N}{d \ln N} = -1 \quad (3.43)$$

and does not depend on the evaporation efficiency. Obviously run-away evaporation is still not possible.

3.3.4 Homogeneous Trap

In a homogeneous trap, runaway evaporation is not possible. The trap volume of a box does not depend on temperature and the atomic density cannot increase due to cooling. The gain in phase-space density is given by Eq. (3.25)

$$\gamma = (3/2)\beta - 1 \quad (3.44)$$

and the change in collision rate by Eq. (3.28) .

$$\gamma_{coll} = -\beta/2 - 1 \quad (3.45)$$

While run-away evaporation is not possible, the change of collision rate is somewhere in between the crossed-dipole trap and hybrid trap scenarios.

3.3.5 Conclusions

It is beneficial to evaporate in low power-law traps, such as linear magnetic traps, as much as possible to maximize gains in phase-space density. When using optical traps, there are clear benefits to using the hybrid solution, which decompresses only along two directions and maintains higher collision rates. While the homogeneous trap intuitively seems like a bad choice, it is not much worse than optical traps, especially in terms of collision rates. One should note that the considerations in this chapter are only valid for classical gases.

3.3. *EVAPORATIVE COOLING*

Chapter 4

Experimental Setup and Methods

This chapter will give an overview of the experimental setup and cooling methods to produce ^{87}Rb BECs. The structure of the chapter mimics the order of the experimental steps taken on the road to degeneracy. A more compact documentation of the setup has been published previously [23]. Compared to most state of the art ^{87}Rb systems, our vacuum system is exceptionally compact. It features a single vacuum chamber, where both trapping of the atoms in a magneto-optical trap (MOT) and evaporative cooling take place. While often simplicity comes at the cost of reduced atom numbers, we manage to produce quasi-pure BECs of over $> 3 \times 10^5$ atoms in < 30 s. Here, we only discuss the preparation of BECs in a harmonic trap. The transfer of the cold atoms into a uniform trap is detailed separately in Chapter 5.

The prerequisite for trapping and cooling atoms is an ultra-high vacuum (UHV) environment. Section 4.1 describes the characteristics of the single chamber system used in our experiment. At the start of the sequence we create a magneto-optical trap (MOT) to trap atoms from the background vapour and cool them to around $15 \mu\text{K}$. We detail the necessary magnetic coils and laser setup in section 4.2. The next stage is evaporative cooling and is described in Section 4.3. We transfer the atoms to a magnetic quadrupole trap, where they are further evaporatively cooled using radio-frequency (RF) induced transitions. Before spin-flip (Majorana) losses become significant, we transfer the gas into a hybrid trap [49], where we achieve condensation. At the end of our experimental sequence, the atoms are probed by absorptive imaging (Section 4.4). Finally, in Section 4.5, we describe the experiment control and analysis soft- and hardware.

4.1 Vacuum System

We inherited the vacuum system from previous students of the group [35, 36, 50]. In its previous incarnations, it was capable of producing BECs of around 10^5 atoms. However,

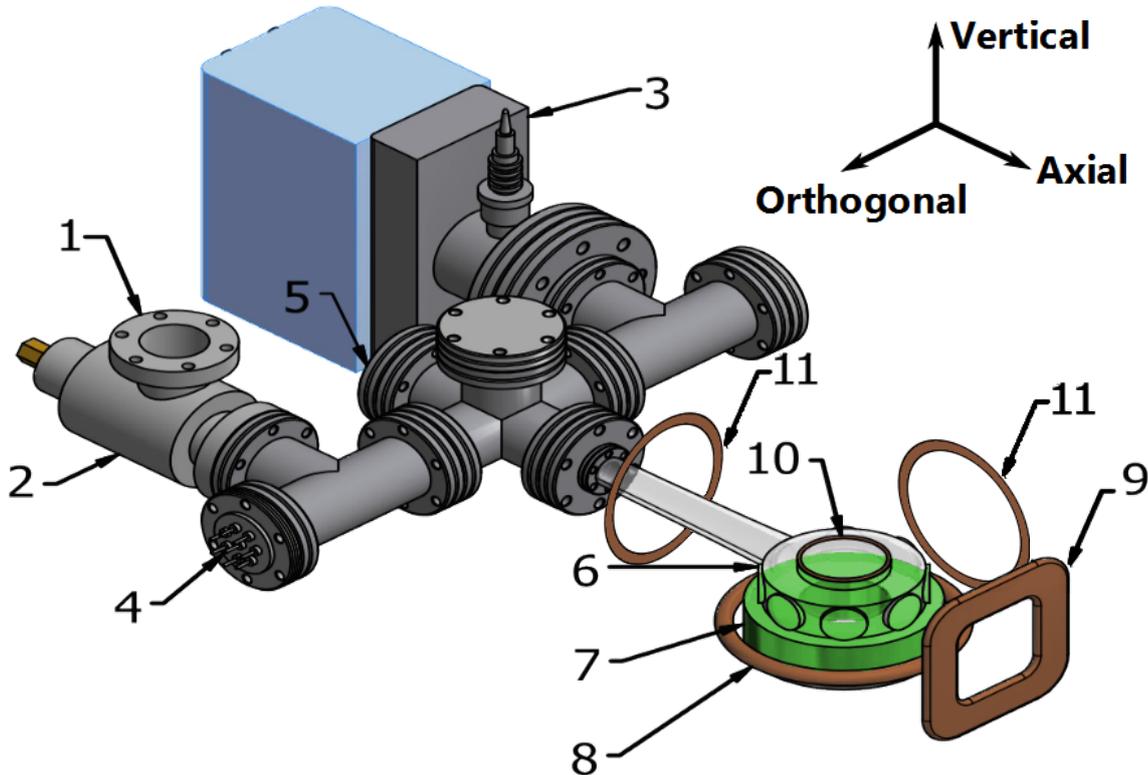


Figure 4.1: Schematic view of the vacuum system. 1 - turbo pump port, 2 - valve, 3 - ion pump, 4 - ^{87}Rb dispensers, 5 - viewport, 6 - glass cell. We also show the magnetic coils for creating: 7 - quadrupole field, 8 - bias field, 9 - guide field, 10 - RF field, 11 - gradient compensation. All coils except 9, 10 and 11 are paired (only one coil per pair is shown).

a full reassembly and thorough scrubbing was needed to revive it. The single-chamber vacuum system is rather unusual as many contemporary experiments are based on multi-chamber geometries [51–53], often in conjunction with Zeeman slowers [49,52]. Separating the MOT spatially from the trap in which evaporative cooling is done allows maximizing atom numbers and life times. Also, it increases optical access for more sophisticated experiments. While a single chamber system is more compact and also cheaper, it suffers from balancing two contradicting requirements: while a higher ^{87}Rb vapour pressure facilitates loading of the MOT, a lower background gas pressure increases the lifetime and evaporation efficiency during later cooling stages. While compact, simple systems are nothing new [44, 54–59], we feel our system excels in the crossover region between compactness, cost and performance.

Our vacuum system with its surrounding magnetic coils is shown in Fig. 4.1. A quartz cell¹ allows optical access via two 50 mm (top and bottom) and eight 25 mm anti-reflection coated windows. The vacuum is maintained by a single 45 l/s ion pump². In

¹Triad Technology

²Oerlikon IZ50

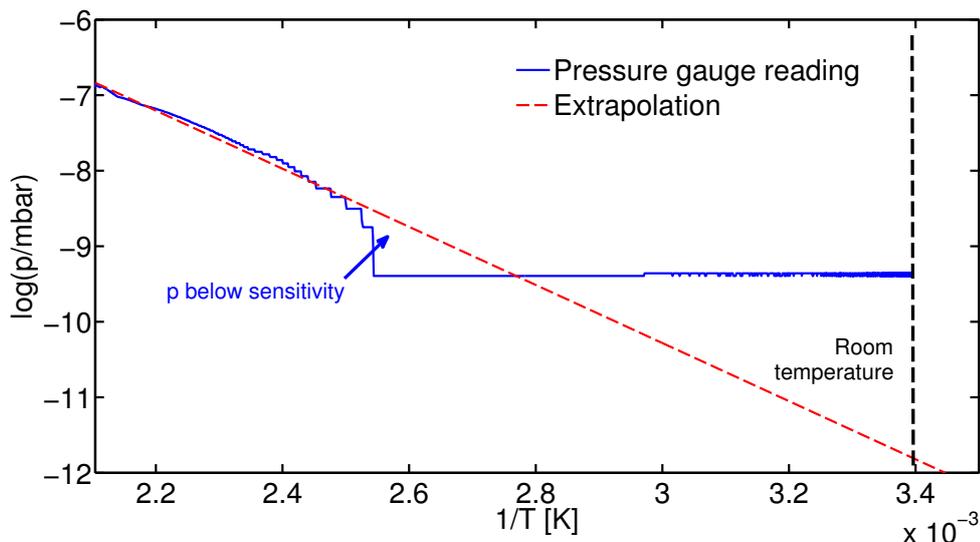


Figure 4.2: Final pressure during bakeout. The red line is an extrapolation beyond the range of the pressure reading of our ion pump assuming $1/T \propto \log(p/\text{mbar})$ [35].

order to remove residual outgassing of contaminants, the system was baked at 200 °C for a week, following the procedures outlined in [35, 36]. We estimated that the pressure of unwanted background gases is below 10^{-11} mbar (see Fig.4.2).

We use vapour sources³ to release pure ^{87}Rb into the system, rather than the natural $\approx 3 : 1$ ^{85}Rb - ^{87}Rb abundance. This increases lifetimes of trapped atoms, as it reduces the rate of background-gas collision for a given ^{87}Rb pressure. We achieved magnetic trap lifetimes of over 20 s for very low ^{87}Rb pressures. We found an optimum for BEC production at a pressure corresponding to atomic lifetimes of ≈ 10 s. Comparing this to the low-pressure lifetime shows that we are barely (if at all) limited by unwanted gas in the system. Rubidium is released by heating the dispensers by running a current through them. Probing the resistance (and hence temperature) of the dispensers with a four point measurement and varying times and currents showed that Rubidium is released above a threshold temperature in a very non-linear fashion. We optimized the ^{87}Rb pressure in our system by trial and error, finding that activating the dispensers once every evening for 80 – 85 s with 6 A gives best results.

4.2 Magneto-optical trap

4.2.1 Magnetic coils

The magnetic quadrupole field for both the MOT and the magnetic trap is created by an anti-Helmholtz pair of 16-turn water-cooled coils wound from 4 mm copper tubing (“7”

³Alvatec

in Fig. 4.1). These create an axial gradient of $B' = 400 \text{ G/cm}$ when carrying 200 A. The coil is switched using an IGBT, which is driven by a driver built by Igor Gotlibovych [27]. A more compact version of the driver is described in the Appendix B. The quantisation axis for optical pumping and imaging is provided by a magnetic field along the imaging axis. We use a low-inductance planar 10 turn coil wound from 5 mm copper wire (“9” in Fig. 4.1), providing a field of 5 G at the atoms. A radio-frequency (RF) field for forced evaporative cooling (see Section 4.3.2) is created by a three-turn coil mounted against the top viewport of the cell. For transitions between neighbouring m_F states separated by up to 15 MHz, we achieve Rabi frequencies above 20 kHz using 600 mW of RF power. A vertical and a horizontal compensation coil allow us to fine tune the magnetic field gradient along the “orthogonal” and “axial” directions (as defined in Fig. 4.1).

Gravity compensation during time-of-flight measurements as well as during later trap stages (i.e. hybrid and uniform) is achieved by combining a quadrupole field (created by the MOT coils) with a homogeneous vertical bias field. The magnetic gradient field B_g required to cancel the gravitational acceleration g is given by $B_g = mg/(\mu_B g_F m_F)$. Typically, our atoms are in the $|F = 2, m_F = 2\rangle$ state, for which $B_g \approx 15.3 \text{ G/cm}$. We apply a quadrupole field with our anti-Helmholtz pair near the atoms

$$\mathbf{B} = B_g \begin{pmatrix} z \\ -r/2 \end{pmatrix}, \quad (4.1)$$

where z is the distance from the magnetic zero along the “vertical” direction and r the respective distance along the radial (i.e “orthogonal” or “axial”) direction. We then add a homogeneous field $\mathbf{B}_0 = B_0 \hat{\mathbf{e}}_z$ along the vertical axis to suppress the radial gradient. The absolute value of the field is given by

$$B = \left| B_g \begin{pmatrix} z \\ -r/2 \end{pmatrix} - B_0 \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right| \approx B_0 - B_g z + \frac{1}{2} \frac{B_g^2}{4B_0} r^2. \quad (4.2)$$

We use two 50-turn coils in a Helmholtz configuration (“8” in Fig. 4.1) to provide the bias field B_0 ; up to a maximum value of 70 G when run at 10 A. The residual curvature corresponds to a trapping frequency of $\omega = 0.5 B_g \sqrt{\mu_B g_F m_F / B_0 m}$. To minimize ω , it is intuitively best to increase B_0 as much as possible. We found, however, that this introduces an additional harmonic potential along x , which can be explained by not perfectly centred bias coils. We therefore settled on running the bias coils at $\approx 5 \text{ A}$ or 30 G. This minimizes any curvature along z , but leaves $\omega = 2\pi \times 1.6 \text{ Hz}$ (measured with oscillating clouds) trapping frequency along r .

4.2.2 Laser Cooling

Our MOT operates on the D_2 $F = 2 \rightarrow F' = 3$ transition of ^{87}Rb (see Fig. 4.3). The spontaneous decay back to $F = 2$ is strongly favourable, which creates a nearly closed transition. Still, off-resonant transfer to $F' = 2$ occurs and can be followed by spontaneous decay into $F = 1$. To prevent atom accumulation in the $F = 1$ dark state, a repump beam continuously drives the $F = 1 \rightarrow F' = 2$ transition.

We generate D_2 light with two grating-stabilized external-cavity diode lasers⁴. Typically, these lasers can produce up to 80 mW of power. We generate cooling light close to the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cycling transition and repumping light resonant with the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. They are referenced and frequency-locked to separate vapour cells, using Doppler-free spectroscopy. This stabilizes the laser frequencies to < 1 MHz, which is smaller than the natural line-width $\Gamma \approx 2\pi \times 6$ MHz. These techniques are well-established, for reference, see, e.g., [27, 35, 37].

The cooling light is amplified with a tapered amplifier⁵, with a maximum output power of 1.5 W. At various stages, light passes through acousto-optic modulators (AOMs), which allow to dynamically shift frequencies and regulate intensities. An overview of laser frequencies with respect to the atomic levels is given in Fig. 4.3. The corresponding setup is shown in Fig. 4.4.

The cooling and repumping light are combined and split into six MOT beams using a free-space fibre cluster, built in-house from slot-in components.⁶

The six MOT beams are delivered to the glass cell via optical fibres (Fig. 4.5). A total of 120 mW of cooling light and 10 mW of repumping light reaches the atoms. With this large amount of power, we can use large diameter (25 mm) beams while still being able to operate near the saturation intensity $I_{sat} \approx 1.67$ mW/cm² in Eq. (3.3). Using large beams like this facilitates alignment of the beams; typically realignment is not needed for months. The light is linearly polarized in the fibres and is converted to the appropriate circular polarization by quarter-wave plates. Final condensate numbers are extremely sensitive to the correct polarization. Occasional realignment (\sim months) of the wave plate is necessary due to long-term drifts of the cooling light polarization in the fibres.

The light for optical pumping and absorption imaging of the atoms is derived from the cooling light. The pumping and imaging beams are controlled by additional AOMs and enter the chamber via the “back” viewport (“5” in Fig. 4.1).

We load the MOT with an axial field gradient of $B' = 13$ G/cm and a cooling-light

⁴Toptica Photonics, DL Pro

⁵Toptica Photonics, BoosTA

⁶Thorlabs, ‘Fiber Bench’ range. The fibre cluster is notoriously hard to adjust, as components creep for a long time after adjustment. In hindsight, it is probably not worth the savings compared to an integrated, commercial solution.

4.2. MAGNETO-OPTICAL TRAP

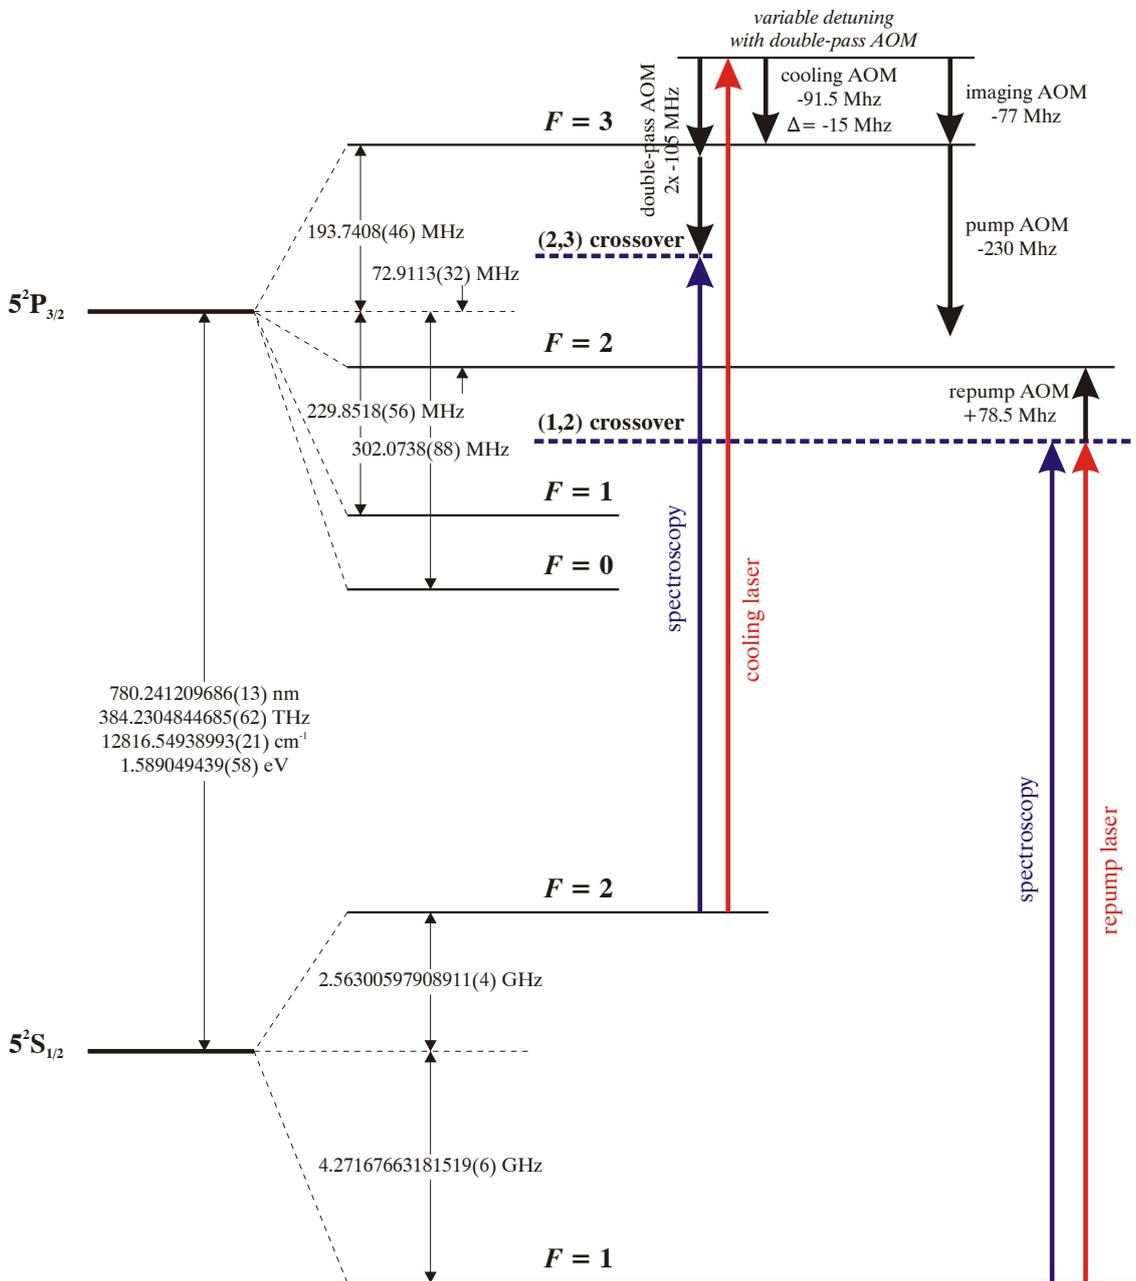


Figure 4.3: The Rb87 level scheme, including the AOM transition frequencies. Data from taken from [39], drawing taken from [40].

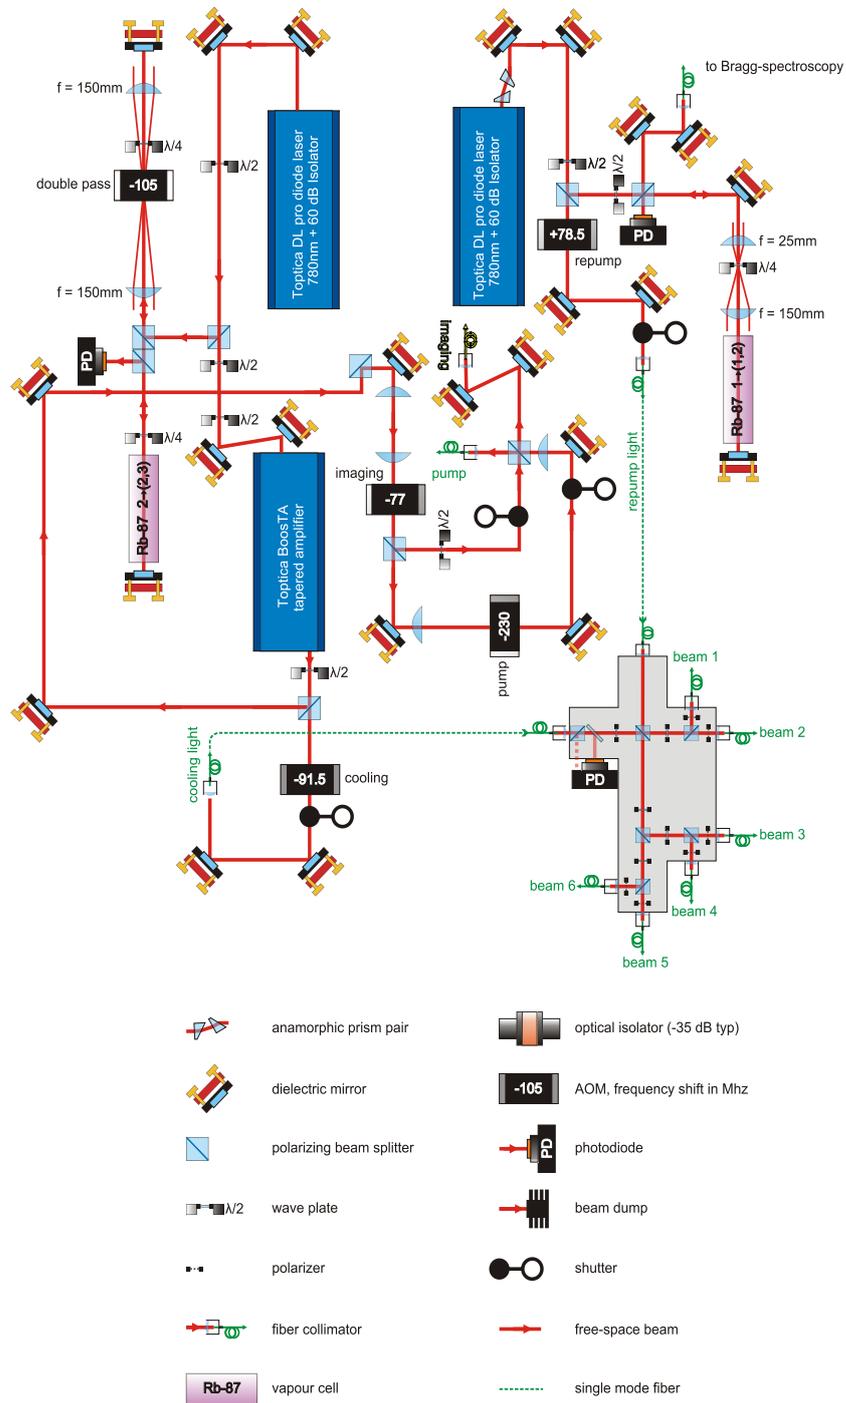


Figure 4.4: The optical setup. AOMs are used to shift the light frequencies if needed. The transition frequencies are given in figure 4.3. Unfortunately, the splitting between repump and cooling transition is 6.8 GHz. This is too large to be covered by AOMs, which is why two separate lasers are used. The drawing is an update to [40].

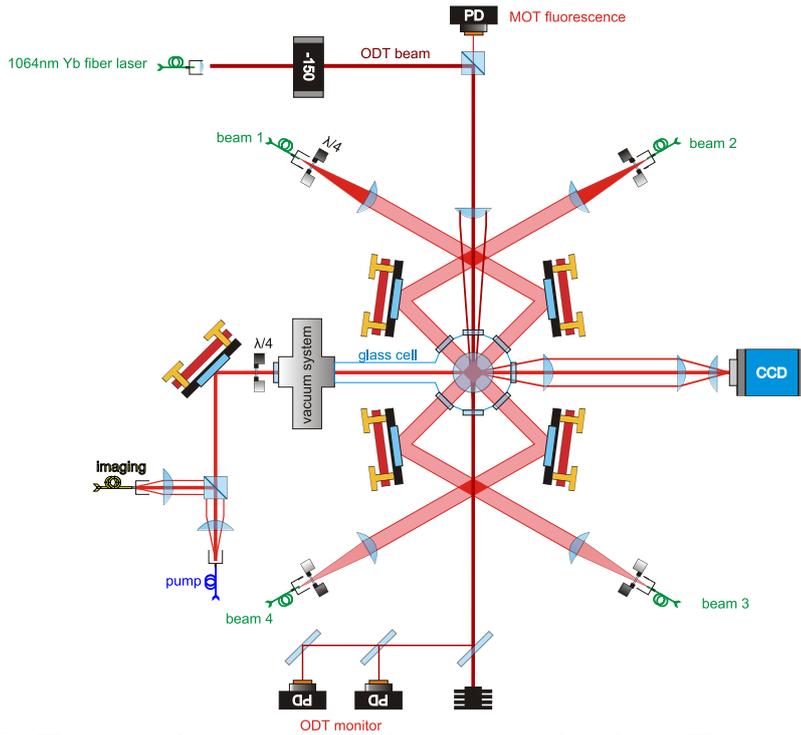


Figure 4.5: The optical setup around the vacuum chamber. The top and bottom MOT beams are not shown. The MOT fluorescence photodiode picks up a signal proportional to the atom number in the MOT. Also shown is the optical dipole trap (ODT). The ODT monitor photodiodes pick up a signal proportional to the optical power and allow us to stabilize the beam intensity with a feedback circuit.

detuning of -25 MHz. Typically, we load the MOT to $\approx 10^9$ atoms in 15 s. Note, however, that 5 s of loading is sufficient to produce BECs with $> 10^5$ atoms. We use the fluorescence signal on a photodiode to monitor the loading of the MOT. Following the initial loading stage, we compress the MOT by increasing B' to 64 G/cm for 20 ms, and then detune the cooling light by -68 MHz for 1 ms [60]. We found that these steps reduce the temperature and improve transfer into the magnetic trap. They are followed by 2 ms of optical molasses during which the magnetic field is switched off and the cooling-light detuning of -68 MHz is maintained. This cools the cloud to ≈ 15 μ K.

4.3 Evaporative cooling

The evaporative cooling to condensation in a harmonic trap can be divided into four stages:

- Transfer to the magnetic trap
- RF evaporation
- Hybrid trap transfer

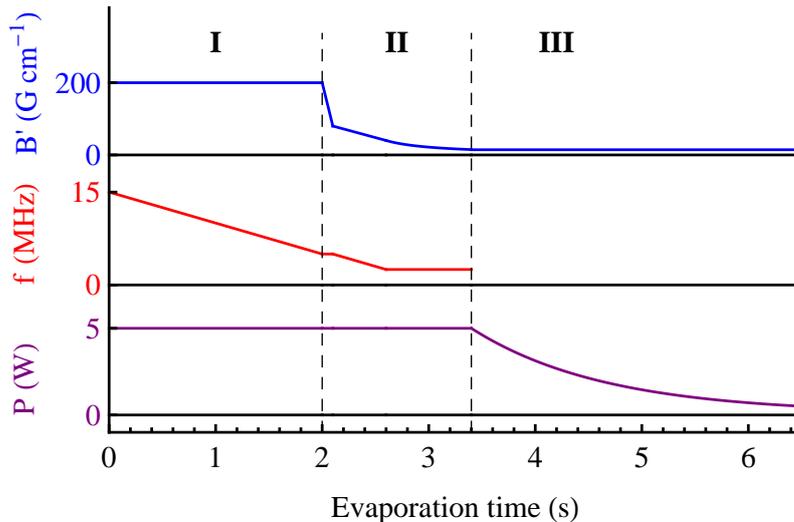


Figure 4.6: Evaporation sequence in the hybrid trap. We show the evolution of (top to bottom) the magnetic field gradient B' , RF frequency f , and ODT power P . The sequence has three stages: **(I)** RF evaporation, **(II)** transfer into the hybrid trap, and **(III)** ODT evaporation.

- ODT evaporation

Fig. 4.6 summarizes the evolution of the relevant experimental parameters during the evaporation sequence, while Fig. 4.7 displays the evolution of the atom number N , the temperature T , and the calculated phase-space density D . N and T are measured using time-of-flight (TOF) absorption imaging and D is calculated from a semi-classical model [49], using an analytical expression for the hybrid trapping potential. We verified the parameters of the model (such as beam waist) by measuring trapping frequencies and comparing them with theoretical predictions. We omit the calculated values of D in stage **II**, where they are unreliable because the trapping potential varies slowly over large volumes and V_0 diverges.

4.3.1 Transfer to the magnetic trap

After the molasses, we apply a σ^+ polarized, resonant $|F = 2\rangle \rightarrow |F' = 2\rangle$ light pulse. This pulse drives only $\Delta m_F = +1$ transitions. Spontaneous decay transfers the atoms back to $|F = 2\rangle$, with $\Delta m_F = \pm 1, 0$. After several of those cycles, most atoms end up in the $|F, m_F\rangle = |2, 2\rangle$ state, which is a dark state for σ^+ light. We also add repump light to prevent accumulation in the $|F = 1\rangle$ dark state. We optimize the pump time to maximize phase-space density, as heating due to photon absorption competes with transfer efficiency. We pump $> 80\%$ of the atoms into the $|2, 2\rangle$ state in $40 \mu\text{s}$, while heating the cloud to $40 \mu\text{K}$.

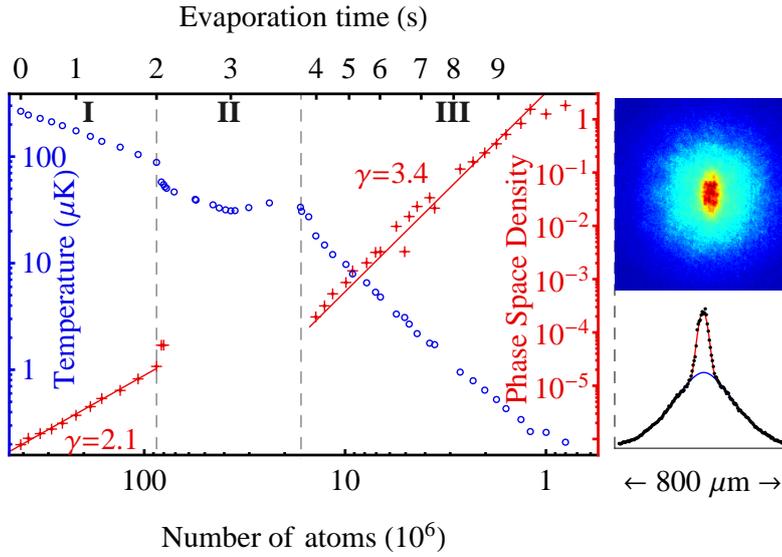


Figure 4.7: BEC production. We plot the temperature T (blue open circles, left axis) and phase-space density D (red crosses, right axis) versus atom number N (bottom) and time t (top) throughout the cooling sequence. See Fig. 4.6 and text for details of stages **I** - **III**. Also shown are an absorption image of a partially condensed cloud after 50 ms of time-of-flight expansion [optical density varies between 0 (blue) and 3 (red)] and its integrated profile. Fits to the thermal and condensed components are indicated.

The atoms are then captured in a magnetic quadrupole field $\mathbf{B} = B' \times (z, -x/2, -y/2)$ produced by the MOT coils. The Zeeman effect creates a potential $U(\mathbf{r}, m_F)$ in which the atoms are trapped

$$U(\mathbf{r}, m_F) = \frac{1}{2} m_F \mu_0 |\mathbf{B}(\mathbf{r})| \quad (4.3)$$

where μ_0 is the Bohr magneton and we used and $g_{F=2} = 1/2$. In practise, we suddenly turn on the quadrupole field with $B' = 64 \text{ G/cm}$ (matching the gradient of the CMOT) and then ramping B' to 80 G/cm over 200 ms. Finally, the trap is further compressed by raising B' to 200 G/cm over 500 ms.

4.3.2 RF Evaporation

We now evaporate by transferring the atoms from the trapped $|F = 2, m_F = 2\rangle$ state to the untrapped $|F = 2, m_F \leq 0\rangle$ states. We generate an RF-field with frequency ω_{RF} with the RF coil above the cell. The resonance condition is met on a ellipsoid around the origin, where the Zeeman splitting between the m_F -states $\Delta E = U(\mathbf{r}, m_F) - U(\mathbf{r}, m_{F-1})$

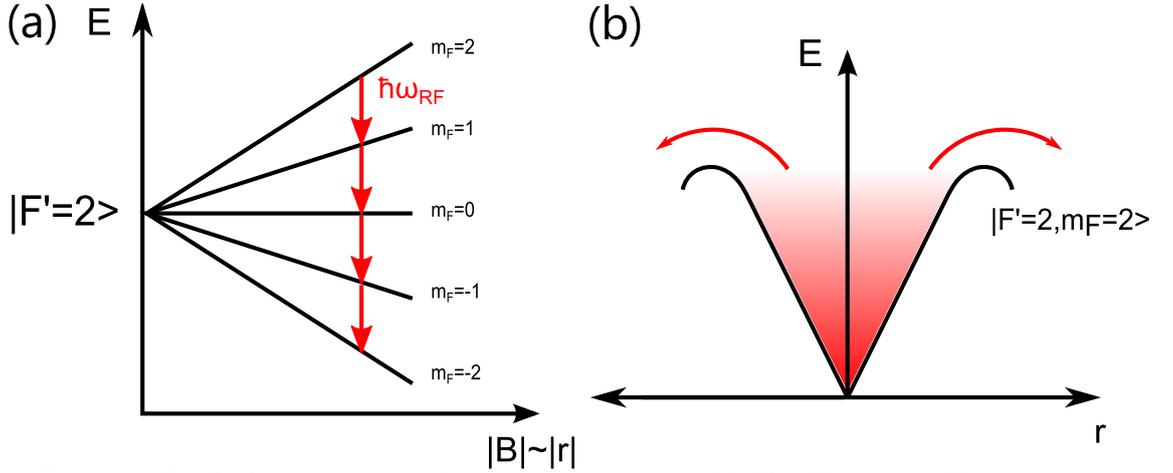


Figure 4.8: Radio-frequency driven evaporation. (a) The Zeeman effect causes a linear splitting in the m_F manifold. A resonant RF-field induces transition between the states. (b) Sketch of potential energy in the trap. Where the splitting is resonant with the applied RF-field, atoms are transferred to untrapped m_F states and ejected from the trap.

matches the transition frequency:

$$\hbar\omega_{RF} = \Delta E = \frac{1}{2}\mu_0 |\mathbf{B}(\mathbf{r})|. \quad (4.4)$$

If initially in the $|F = 2, m_F = 2\rangle$ state, only atoms with energy $E > 2 \times \Delta E = 2\hbar\omega_{RF}$ can reach the transition area. The atoms are then shifted down in steps of $\Delta m_F = -1$ and are expelled. Fig. 4.8a illustrates this process. The resulting trap is approximately a truncated linear power law (Fig. 4.8b). This method allows to control both collision rate by (de-)compressing the trap via tuning B' , as well as the trap depth by ramping ω_{RF} .

During the initial evaporation stage in the compressed trap (**I** in Figs. 4.6 and 4.7), we ramp the RF frequency linearly from 15 MHz to 5 MHz in 2 s. We achieve a 30-fold increase in the peak phase space density, with evaporation efficiency

$$\gamma = -d[\ln D]/d[\ln N] = 2.1. \quad (4.5)$$

At the end of this stage the cloud typically contains $N = 50 \times 10^6$ atoms at $T = 90 \mu\text{K}$. At this point further evaporation is rendered inefficient by spin-flips at the trap centre (Majorana losses), where the field is low. Here, the different Zeeman states become degenerate and the quantization axis is not well-defined. Non-adiabatic transfer to other, possibly untrapped, states occurs and atoms are lost. Even worse, this affects mostly low energy atoms, as they spend more time near the trap centre.

To prevent these losses, we modify the trapping geometry so it does not include the magnetic zero. Various approaches exist, such as a TOP trap in which the first BEC was

produced [4, 61], the Ioffe-Pritchard trap [62], or the use of optical trapping methods. Another common method is shooting a blue detuned laser through the centre of the trap, repelling the atoms. This technique is called an optical plug. Stuart Moulder achieved successful condensation with this method with as little as 70 mW of plug-laser power [23, 50] in this setup.

4.3.3 Single-Beam Dipole Trap: “Hybrid Trap”

We use a hybrid approach, where we retain the linear magnetic trap, but also superimposed a red-detuned laser trapping beam below the magnetic zero. This approach is described in length in [49] and the derivations here closely follow it.

The total potential $U(\mathbf{r})$ is given by the sum of a quadrupole trap, an optical dipole trap (ODT) and gravity

$$U(\mathbf{r}) = \mu_0 B' \sqrt{x^2/4 + y^2/4 + z^2} - U_0 \exp \left[\frac{-2[x^2 + (z - z_0)^2]}{w_0^2} \right] + mgz, \quad (4.6)$$

where z_0 is the distance of the beam centre from the magnetic zero. During the initial RF evaporation (stage I) described above, the ODT is already on, but its effect on the potential is minor (see Fig. 4.9). We can neglect its contribution to Eq. (4.6) and the potential is well approximated by a simple quadrupole trap, where gravity leads to minor corrections. During stage (II), reducing the magnetic field gradient while cooling transfers the atoms gradually into the centre of the ODT beam below the magnetic zero, preventing spin-flip losses. For low temperatures and magnetic gradient forces similar to g (stage III), the trapping potential is dominated by the ODT along the x- and z-axis, with trapping frequencies $\omega_{x,z} = 2\sqrt{U_0/mw_0^2}$. Along the y-axis, the atoms are still trapped magnetically, however the displacement from zero leads to a smooth potential. This can be approximated as a harmonic trapping potential with frequency $\omega_y = 1/2\sqrt{\mu B'/mz_0}$. Like in an all-optical crossed-dipole trap (CDT), the trap depth is given by U_0 . Evaporation is now driven by lowering U_0 , as atoms can escape downwards aided by gravity. As calculated in Section 3.3, a big advantage compared to a conventional CDT is that only two out of the three trapping frequencies are lowered during evaporation, allowing for higher collision rates.⁷

The cross-over between the two regimes happens in the transition region (II). Here, neither the linear nor the harmonic approximations hold and the trap shape is best described numerically. Also, the phase space density will be affected due to the change in effective volume V_0 . At some stages, $U(\mathbf{r})$ varies very slowly over large regions and the numerical values for V_0 and hence D are unreliable.

⁷An extension of this is the tilted trap approach, where the magnetic field gradient is used to lower the trap depth without changing the ODT trapping frequencies [58].

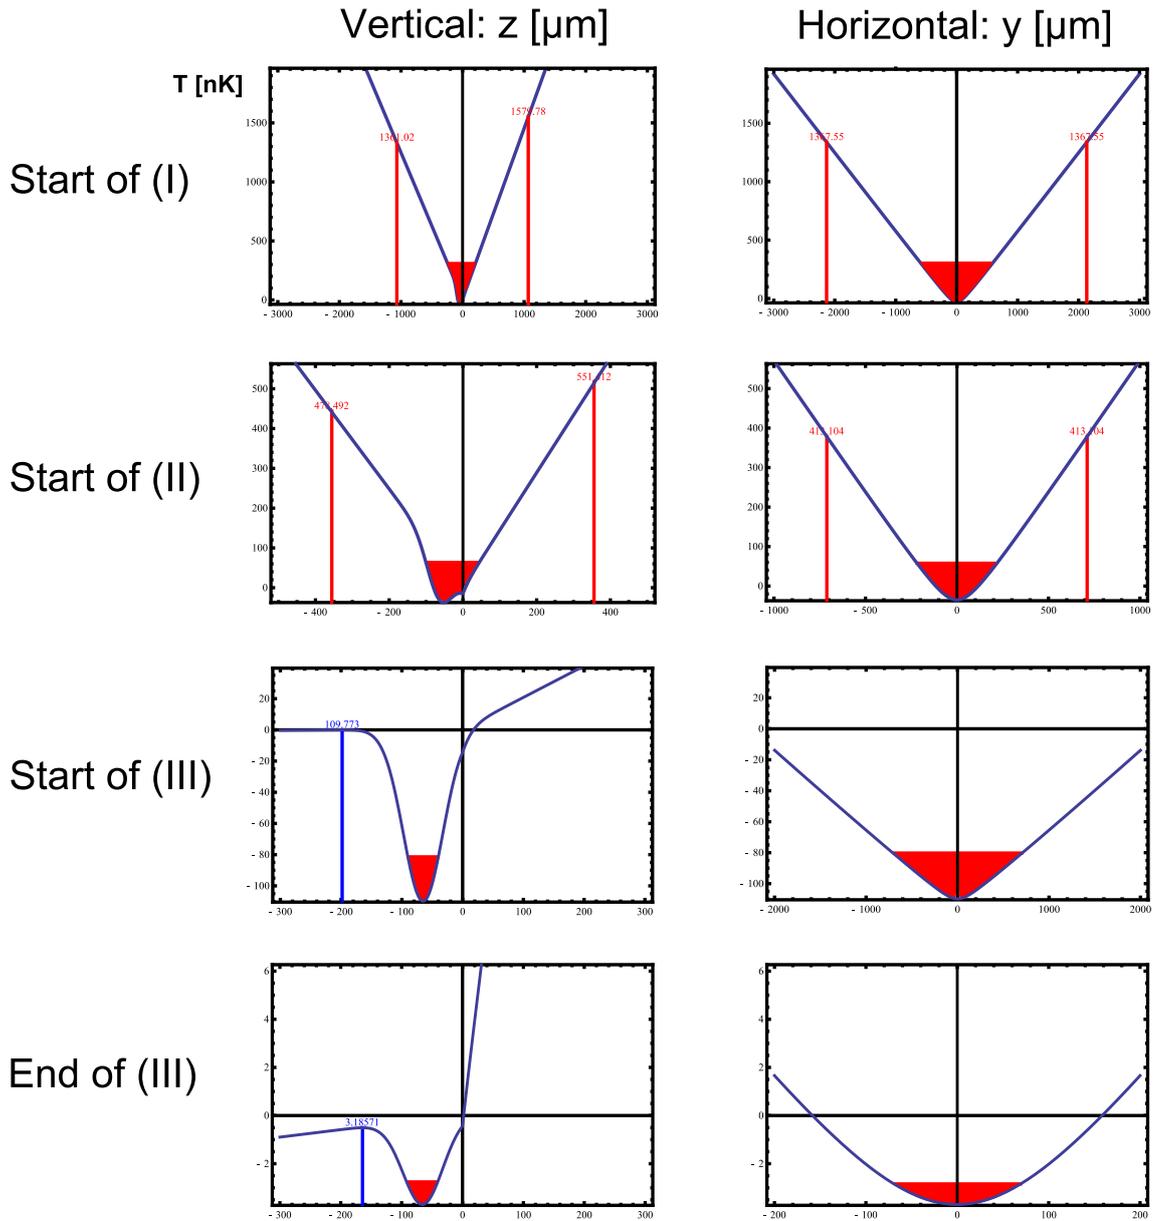


Figure 4.9: Trapping potential evolution throughout the evaporation process. The left column shows the vertical direction which is affected by gravity. The right column shows the horizontal direction along the beam, where the confinement is predominantly magnetic at all times. The red bars signify the resonance of the RF knife, that is where the potential is truncated. The “red liquid” level in the potential well corresponds to T . We see that during the initial stages of evaporation, the trap is predominantly linear and magnetic. Later, the confinement along z is optical and atoms can evaporate by “falling out” thanks to gravity. The potential along y is still magnetic, however the confinement is harmonic, as the minimum of the trap is displaced from zero. The potential along x (in the horizontal plane orthogonal to the ODT) is very similar to the z -potential, except that at later stages gravity does not truncate the trap.

Optical dipole trap laser

The optical-dipole-trap beam used for the hybrid trap is derived from a 20 W ytterbium fibre laser⁸. After passing through an AOM, the beam is focused by a concave-convex pair of lenses, allowing the adjustment of both waist size and position. We typically use a beam waist of $65 \mu\text{m}$ and a maximum power of 5 W at the atoms. Higher powers lead to excessive atom losses, which can be attributed to two-photon transitions between $F = 1$ and $F = 2$ [63]. These transitions are caused by the large linewidth of the trapping laser (1.8 nm in our case). The focal point is positioned approximately a beam waist below the zero of the magnetic quadrupole field. The beam power is controlled in the range 0.01 – 5 W using a feedback loop consisting of a photodiode, a proportional-integral-derivative (PID) controller and an AOM. Throughout this range, we achieve an RMS power noise below 1 mW over a bandwidth of 1 kHz.

Loading of the Hybrid Trap

In stage **II**, the magnetic field gradient is decompressed to just below the gravity-compensating value, $B'_g = mg/\mu_B = 15.3 \text{ G/cm}$. From this point on, the atoms are supported against gravity by optical forces only, while the confinement along the ODT axis is dominated by the magnetic forces [49]. Note that the ODT beam is on from the beginning of the magnetic trapping (see Fig. 4.6), but initially does not have a dominant role.

We decompress B' to 80 G/cm over 100 ms while keeping the RF frequency constant, then decompress further to $B'_1 = 40 \text{ G/cm}$ while sweeping the RF frequency linearly to 2.5 MHz in 0.5 s; this provides a good gain in phase-space density despite the weak magnetic confinement. Finally, the magnetic field is decompressed from B'_1 to $B'_2 = 14.8 \text{ G/cm}$. In this step we sweep the field gradient according to $B'(t) = B'_1 \times (1 + t/\tau)^{-1}$, where $\tau = t_2 B'_2 / (B'_1 - B'_2)$ and the sweep duration is $t_2 = 800 \text{ ms}$. Sweeping B' in this manner changes the RF-limited trap depth approximately linearly with time, with the approximation being exact in the case of no ODT and $B'_2 = B'_g$. At the end of this stage we typically have 20×10^6 atoms at $30 \mu\text{K}$.

Optical Evaporation

For $B' < B'_g$, evaporation is no longer driven by the RF field. Instead, the trap depth is limited by the potential at the saddle point vertically below the magnetic-field zero [58]. Evaporation can be forced by lowering the ODT power U_0 , albeit at the cost of reducing the trapping frequencies.

⁸IPG Photonics, YLR-20-LP

We ramp U_0 exponentially to its final value with a time constant of 1.25 s and total ramp duration of 7 s. Both parameters have been optimized empirically and the optimum depends on the current vapour pressure. We achieve an evaporation efficiency of $\gamma = 3.4$, leading to condensation at $P = 0.16$ W.

At the critical point, $N \approx 10^6$ and $T \approx 250$ nK. At this point, the trapping frequency is about 30 Hz along the ODT axis and about 90 Hz radially.

Lowering the ODT power further results in quasi-pure⁹ condensates of $> 3 \times 10^5$ atoms. The lifetime of the quasi-pure BEC is ~ 3 s, consistent with the expected 3-body losses [64, 65].

4.4 Imaging and Image processing

We image the atoms absorptively, using a circularly polarized beam resonant with the $|F = 2, m_F = 2\rangle \rightarrow |F' = 3, m_{F'} = 3\rangle$ cycling transition and measure the resulting absorption pattern on a CCD camera¹⁰. We start by considering the basic principles of absorptive imaging in the low intensity regime. After this we discuss the deviation of the measured absorption cross-section from its theoretical value, which are primarily due to imperfections in the light polarisation. We calibrate this deviation by comparing the measured atom number at the BEC critical point \tilde{N}_c with the theoretically predicted critical number N_c . \tilde{N}_c was measured by extrapolating the non-saturation slope of the thermal component to the condensation point, as described in [16, 34]. Finally, we describe high-intensity imaging [66], which allows us to measure higher optical densities. This technique also reduced the density correction factor.

4.4.1 Absorption Imaging

The absorption of a beam of intensity I propagating along the z-axis through a cloud of atoms is given by [37]

$$\frac{dI}{dz} = -n(z)\sigma_0 \frac{I}{(2\Delta/\Gamma)^2 + (1 + I/I_{sat})}, \quad (4.7)$$

where n is the density, $\sigma_0 = 3\lambda^2/2\pi$ is the resonant cross-section, $I_{sat} = (\pi/3)hc\Gamma/\lambda^3$ the saturation intensity and Δ the angular detuning from resonance. We consider now the simplest case with no detuning $\Delta = 0$ and low intensity $I \ll I_{sat}$. After passing through an atomic cloud with density $n(x, y, z)$, we measure the following intensity pattern on the

⁹We can only measure thermal fractions accurately down to about 50%.

¹⁰PCO, Pixelfly USB

CCD screen:

$$I(x, y) = I_0(x, y) \exp[-\sigma_0 \tilde{n}], \quad (4.8)$$

where we introduced the the column density $\tilde{n}(x, y) = \int n(x, y) dz$. To analyse a cloud, we first measure the intensity $I_A(x, y)$ in the presence of the atoms. Then we measure the original light intensity $I_L(x, y)$ without absorption from atoms. Last we take one background image $I_{BG}(x, y)$ without atoms or light to measure the dark count and other offset effects. We can extract the two-dimensional column density from the optical density

$$OD(x, y) = \sigma_0 \tilde{n}(x, y) = -\ln \left[\frac{I_A(x, y) - I_{BG}(x, y)}{I_L(x, y) - I_{BG}(x, y)} \right]. \quad (4.9)$$

In practise, we apply a magnetic bias field for the shot with I_A and tune the light frequency accordingly to compensate for the Zeeman shift. Before measuring $I_L(x, y)$, the field is then quickly switched off, tuning the atoms out of resonance and thus making them transparent to the imaging light. The background shot is less critical and fairly reproducible. We also measure and correct for fluctuation in global light intensity between I_A and I_L (corresponding to an offset in OD) by comparing an illuminated region far from the atoms. We usually also remove a weak offset (a camera artefact) from the OD by fitting a wide parabolic profile to the region of the image not containing atoms.

4.4.2 Density calibration

In most experiments it is necessary to introduce an empirical factor F (often referred to as the “fudge factor”) to take deviations from the theoretical cross section into account. These deviations are usually caused by polarization imperfections, but also other effects can reduce the imaging efficiency.¹¹ We thus modify Eq. (4.9) to

$$OD(x, y) = \sigma_0 \tilde{n}(x, y) = -\ln \left[\frac{I_A(x, y) - I_{BG}(x, y)}{I_L(x, y) - I_{BG}(x, y)} \right] \times F. \quad (4.10)$$

F can be calibrated by measuring the apparent critical number according to Eq. (4.9) in a harmonic trap and comparing it to the predicted critical number for the appropriate temperature and trapping frequencies. To do so, we need to take into account that the ideal gas model does not fully capture the behaviour of the condensed gas. We need to take the non-saturation of the thermal fraction into account, i.e., the thermal number

¹¹For instance, F increases if the frequency-lock of the imaging beam is not tight enough, i.e., if the linewidth is too broad.

behaves as [16, 34]:

$$N' = N_c + S_0 N_0^{2/5} + S_2 N_0^{4/5} \quad (4.11)$$

$$S_0 = \frac{\xi(2)}{2} \left(\frac{k_B T^2}{\hbar \bar{\omega}} \right) \left(\frac{15a}{a_{ho}} \right)^{2/5} \quad (4.12)$$

$$S_2 = (3 \pm 0.7) \frac{\xi(3)}{4} \left(\frac{k_B T}{\hbar \bar{\omega}} \right) \left(\frac{15a}{a_{ho}} \right)^{4/5}. \quad (4.13)$$

Hence, the thermal number N' varies in a nonlinear fashion near the condensation point. A more detailed treatment of this effect is given in Section 6.2.1. We measure a range of condensate sizes for a fixed temperature and use Eq. (4.11) to extrapolate to the condensation point. Multiplying the observed numbers N' and N_0 with F allows us to relate them to the theoretical critical number N_c

$$N' = F^{-1} (N_c + S_0 N_0^{2/5} F^{2/5} + S_2 N_0^{4/5} F^{4/5}). \quad (4.14)$$

We now can fit Eq. (4.14) to experimental data, allowing us to extract F . We did this calibration whenever we had reason to believe that the factor might have changed, such as after modifying imaging beam optics. We obtained values between 1 and 2, which is fairly typical [59, 67]. Figure Fig. 4.10 shows a calibration example. The main errors are induced by the uncertainty of temperature and trapping frequencies. We measured the trapping frequencies ω_y and ω_z with a precision of better than 1 Hz by exciting a small amplitude oscillation. The main uncertainty in frequencies is due to ω_x , as we cannot image oscillations along the imaging axes. Assuming an isotropic ODT, it is similar to ω_y . We estimate the difference due to the magnetic field using Eq. (4.6). Depending on the trapping potential and temperature, deviations from the harmonic trap shape need to be taken into account when calculating the effective volume. However, these corrections typically are only on the order of a few percent.

4.4.3 High intensity imaging

Using low imaging intensities comes with some drawbacks. Foremost, the dynamic range of counts per pixel is limited (12 bit in our case), leading to an upper cap in OD value. The theoretical cap of the optical $OD < \ln(12) \approx 9.7$, is further reduced by background counts and other noise, rendering images with $OD > 3$ questionable. This limits us to measuring dilute clouds. While reducing the cross-section by increasing Δ helps, this can distort images severely by lensing effects [66].

We now drop the requirement that $I \ll I_{sat}$. We include a fudge factor \tilde{F} in Eq. (4.7)

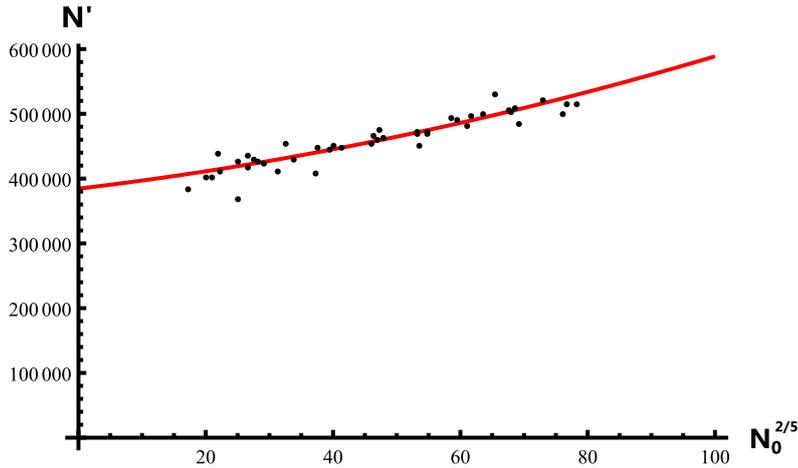


Figure 4.10: Example of a fudge factor calibration. Using the theoretical variation of the thermal number with condensate number, as predicted by Eq. (4.13), we find the best fit for the fudge factor $F = 1.7 \pm 0.2$. The error is dominated by systematic uncertainties in the trapping frequencies. Here, $\omega_{x,z,y} = 2\pi \times (106, 30, 102)$ Hz. We also leave the prefactor of S_2 as a free fit parameter and obtain 3.9 ± 0.6 in accordance with the experimental value 3 ± 0.7 given in [34]. We corrected for a linear temperature drift between 245 and 255 nK using the correction suggested in [16]. The critical number is corrected for finite-size effects [28] and the mean-field shift [10,59].

and integrate to obtain the corrected column density

$$\sigma_0 \tilde{n} = \tilde{F} \left\{ -\ln \left[\frac{I_A(x, y) - I_{BG}(x, y)}{I_L(x, y) - I_{BG}(x, y)} \right] + \frac{I_L - I_A}{\tilde{F} I_{sat}} \right\}. \quad (4.15)$$

We can identify the first term as Eq. (4.9), which holds for low intensities. The second term dominates at high intensities when the transition is saturated. We distinguish the new over-all fudge factor \tilde{F} from the corresponding low-intensity fudge factor F . This is necessary as F often includes a contribution from the high-intensity term, if $I \ll I_{sat}$ is not fully met.

[66] uses Eq. (4.15) to find \tilde{F} by varying I at constant atom number. For the correct value of \tilde{F} there is minimal variation in the column density. However, this requires an accurate calibration of I_{sat} via other means, which we found difficult to achieve. Instead we use Eq. (4.15) to find effective saturation intensity $\tilde{F} I_{sat}$ and then calibrate the fudge factor using the non-saturation method.

We use the following relation between experimentally measured counts per pixel N and intensity I

$$N = It \frac{4qA}{hf}, \quad (4.16)$$

where $A = (2.51 \mu\text{m})^2$ is the pixel area, $q \approx 10 - 20\%$ the quantum efficiency of the

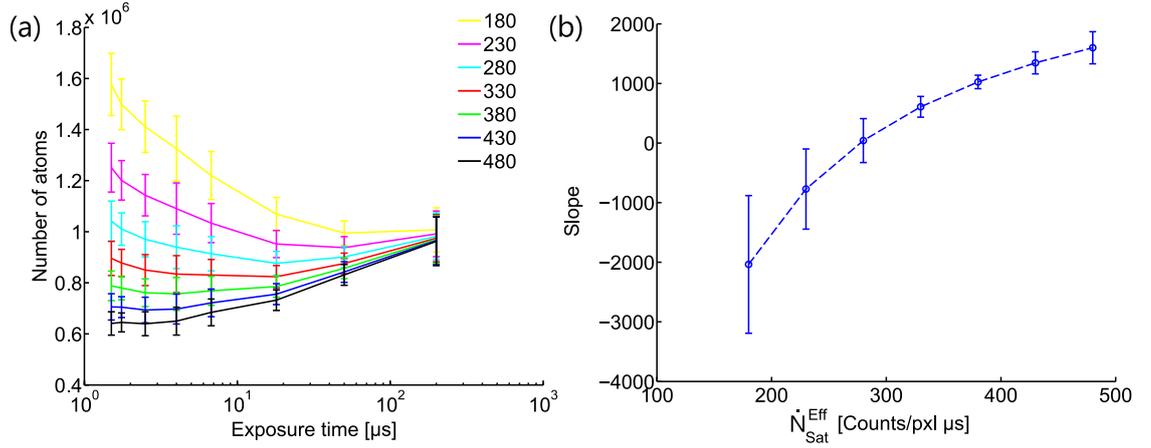


Figure 4.11: Calibration of the imaging intensity. (a) We vary the imaging intensity I and exposure time t while keeping the total photon count constant. The images are evaluated with a different values for the saturation count \dot{N}_{sat}^{eff} . For the correct value, there should be no variation in of the measured number N with I (or equivalently t). (b) The average slope of the different curves in (a). We find minimum variation for $\dot{N}_{sat}^{eff} \approx 280 \mu\text{s}^{-1}$.

camera¹² and t the exposure time. For reasons only known to the camera manufacturer, the camera internally multiplies the counts by 4. We rephrase equation Eq. (4.15) by replacing all the intensities with their respective counts

$$\sigma_0 \tilde{n} = \tilde{F} \left\{ -\ln \left[\frac{N_A(x, y) - N_{BG}(x, y)}{N_L(x, y) - N_{BG}(x, y)} \right] + \frac{N_L - N_A}{t \dot{N}_{sat}^{eff}} \right\}, \quad (4.17)$$

where we introduced the effective saturation count rate $\dot{N}_{sat}^{eff} = \tilde{F} \frac{d}{dt} N_{sat}$, which is the counterpart of the effective saturation intensity $\tilde{F} I_{sat}$. We then varied the intensity over two orders of magnitude while adjusting the exposure time to keep the average number of imaging photons constant. The pictures were post-processed using different \dot{N}_{sat}^{eff} (see Fig. 4.11a). For the correct value, we expect a flat line, as the total atom count should not depend on imaging intensity when processed correctly. In the example shown, we found $\dot{N}_{sat}^{eff} \approx 280 \mu\text{s}^{-1}$ meets this criterion best (see Fig. 4.11b).

Repeating the density calibration with Eq. (4.17) with typical imaging intensities, we found that $\tilde{F} \approx 0.8F$, i.e., 20% of our low-intensity imaging fudge factor could be attributed to operating too close to I_{sat} .

Note that this value for \dot{N}_{sat}^{eff} is roughly consistent with the estimated counts:

$$\dot{N}_{sat}^{eff} \approx \tilde{F} I_{sat} \frac{4qA}{hf} \approx (240 \pm 80) \tilde{F} \mu\text{s}^{-1}. \quad (4.18)$$

¹²Some measurements we conducted caused us to have some doubts concerning the efficiency quoted by the manufacturer.

It is important to consider the maximum amount of photons absorbed per atom during imaging. The absorption of too many photons (followed by spontaneous reemission) can lead to acceleration of the atom towards the camera. If the resulting Doppler-shift is too large, this can tune the atom out of resonance. It can be shown that if we keep the total photon count per pixel N_{max} constant, the maximum number of photons is absorbed for $I \ll I_{sat}$. In our case $N_{max} = 6000$, and we obtain a maximum of 500 photons per atom. This corresponds to a Doppler-shift of around $3 \text{ MHz} \ll \Gamma/2$ at the end of the imaging pulse. Hence the scattering cross section at the end of the light pulse is halved for low intensities; however the effect on the total counts is less than that. Still, in retrospect we should have used fewer photons.

The calibration of the \dot{N}_{sat}^{eff} allows us to increase I , which allows us to image dense clouds reliably. This is essential in the medium-intensity range, which is useful for resolving thermal distributions originating from homogeneous traps, which are inherently “peaky” in the centre.

In the limit $I \gg I_{sat}$, Eq. (4.15) is especially simple and we obtain a column density profile, which is directly proportional to $(I_L - I_A)/(I_{sat})$. Even when I_{sat} is not properly calibrated, this allows us to extract the correct density profile, e.g., when considering in situ distributions.

4.5 Experiment control and data processing

4.5.1 Sequence setup and control

We control our experiment with a National Instruments system consisting of a chassis (PXIe-1082), two digital output cards (PXI-6536) and two analogue output cards (PXI-6713, PXI-6733). This gives us a total of 64 digital TTL-compatible output channels, which are isolated from each other with optocouplers to prevent ground loops. The analog cards provide 16 output channels in total. These digital channels allow us to interface with the various devices in our experiment in an on-off fashion. Devices which need gradual control, such as the quadrupole power supply, are controlled by the analogue controls. All of these channels are synchronized with a common clock.

The hardware is interfaced using Aviv Keshets “WordGenerator” software [68], which allows us to design experimental sequences. Besides the digital and analogue channels, it adds the possibility of outputting serial commands on the computers port via RS232 and GPIB. Besides communicating with instruments with the appropriate ports, we use this capability to interface with custom-made software, e.g. to implement functionalities not native to “WordGenerator”. This is achieved by creating a pair of virtual serial commu-

nication ports with the free “com0com” software. If WordGenerator sends a command to one of those ports, it can be read out by any other program from the other port. We use this for two purposes. First, this allows us to save complementary files for each experimental run with the most relevant parameters. This is immensely useful for evaluation and documentation. Second, by adding an additional software layer between WordGenerator and the experiment, we can implement feedback. This means we can automatically modify future measurements to compensate for drifts, without the need of manually adjusting the experiment settings in WordGenerator.

To reduce the influence of noise from the mains power, we use WordGenerator’s built-in feature to synchronize the start of each experimental run with the phase of the mains power. Using a simple zero-crossing circuit, we detect the phase of the 240 V AC to ≈ 1 ms accuracy and use a digital trigger to synchronize the experiment with it.

4.5.2 Imaging control

The CCD camera is controlled by a custom software written in LabView. Images are taken when the camera is triggered with a digital signal from the control hardware. The imaging parameters (such as exposure time) are set within the imaging software. The software has several operation modes:

- **Absorption imaging:** Three subsequent images are taken and the OD is calculated according to Eq. (4.9).
- **Absorption imaging with double-shuttering:** Double shuttering allows us to take two images within only $10\mu\text{s}$. This allows us to minimize intensity fluctuations between the images with and without atoms, reducing noise.
- **Single shot:** Takes a single image, e.g. to image the MOT for fluorescence imaging.
- **“Live mode”:** Continuously records images without waiting for a trigger and displays them as a live video. This is particularly useful for observing the MOT or aligning beams.

The camera has a resolution of 14-bit, but registers a single photon as 4 counts, resulting in a maximum count of 2^{16} . We can thus save the unprocessed images as 16-bit grey-scale “.tiff” files. When generating OD-images, we stretch their values according to $\tilde{OD} = 2^8 \times (OD + 8)$, using the full bit depth to ensure efficient and lossless storage. When high-intensity imaging is used, we use additional programs to generate optical density images according to Eq. (4.15) using the unprocessed images.

While the exposure time of the camera is set by the software, we set it to be larger than the actually desired exposure time. The effective exposure time is then determined by the duration of the imaging pulse. This time is controlled to $0.5 \mu\text{s}$ precision by an Arduino microcontroller, allowing us to improve beyond the time resolution of the WordGenerator software.¹³ The required exposure time is sent to the Arduino via serial communication at the beginning of the sequence; the pulse is digitally triggered when required.

4.5.3 Analysis software

Igor Gotlibovych has programmed an analysis software allowing us to view the images in a graphical user interface. Once loaded, various fits can be applied, extracting information such as temperature. These fits can be saved to disk and contain a link to the original images, fitting results and optional extra information from log files. More information on the software architecture and its implementation can be found in [27].

¹³Alternatively, this can be implemented with a variable time base as described in the WordGenerator documentation.

Chapter 5

Creating a homogeneous Bose-Einstein Condensate

The creation of homogeneous Bose-Einstein condensates has been a long-standing goal in the field. They offer a closer connection to conventional many-body systems, while simplifying theoretical problems. In this chapter, we explain how we created the first quasi-homogeneous BEC. We cover the technical challenges encountered and discuss the first thermodynamic studies we conducted to confirm the uniformity. The main results of this chapter have been published in [24].

We use blue-detuned light to confine our atoms in a steep-walled dark optical trap, while cancelling gravity with a magnetic gradient (Sections 5.1 and 5.2). The trap is loaded with a gas precooled in a harmonic trap, followed by evaporation in the dark trap (Sections 5.3 and 5.4). We characterize the trap by examining the in situ density distribution in Section 5.5. In Section 5.6, we approximate our trap by a high power-law potential and derive the resulting momentum distribution. This model allows us to measure thermodynamic properties such as temperature and number. Comparing the predictions with the measurements in Section 5.7 allows us to quantify the uniformity of the potential: we find the trap is well approximated by a power-law of order 15, which is essentially uniform. Finally, we confirm the theoretically expected critical point in a homogeneous gas.

5.1 Dark Optical Trap

Like attractive optical traps, dark optical traps operate on the principle of the AC Stark effect [69]. The light frequency is blue-detuned from resonance and the atoms are repelled by the light. We enclose our atoms with walls of light, while minimizing the amount of light inside the enclosed trapping volume. This technique has several advantages over

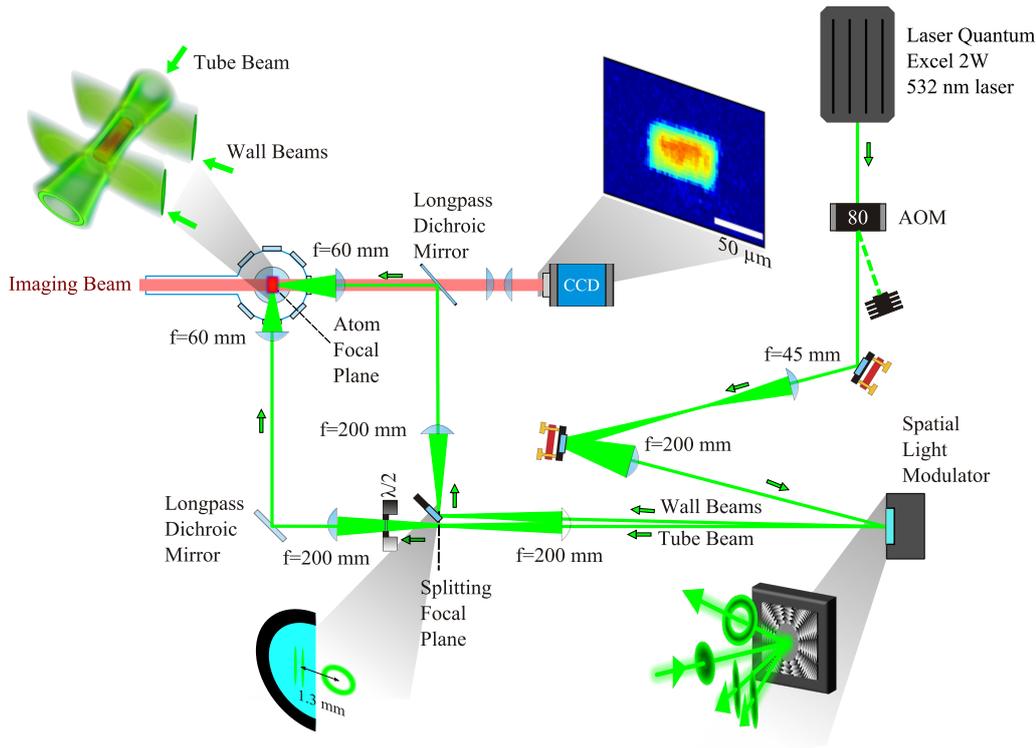


Figure 5.1: The beam configuration for the production of the green trap. Many of the optics are dichroic and are shared with the MOT and ODT beams (compare with Fig. 4.5).

attractive optical traps: First, there is only negligible heating due to light scattering, as the “contact area” between atoms and the walls is small. Second, it requires far less light, as the wall thickness can be very thin (it is limited by the diffraction limit). Most importantly, the trap is *dark* inside, so that intensity variations do not cause variations of the trapping potential.

We use a hollow-tube beam (pseudo-Laguerre-Gaussian [70]) and intersect it with two Gaussian sheet beams, creating a can-shaped dark region enclosed by repulsive walls. To create those beams, we reflect a 532 nm Gaussian laser beam¹ off a spatial-light modulator (SLM)². The spatial profile of the laser beam is very close to a true Gaussian and does not require cleaning up by passing it through a fibre. The SLM allows us to imprint an arbitrary phase pattern onto the beam. In the far-field plane, which we emulate by focussing the beam, this yields approximately the Fourier-transform of this phase pattern, allowing us to create arbitrary, two-dimensional patterns. In the focal planes, the beam profile is approximately constant over a Rayleigh length ($\sim 50 \mu\text{m}$), allowing us to extrude these profiles into three-dimensional shapes. Details of the beam-shaping and the digital holography in our experiment can be found in [70, 71].

¹Quantum Laser, Excel 2W

²Hamamatsu, X10468-04

Our setup to create the green beams is shown in Fig. 5.1. A Gaussian beam is incident on the SLM, where a phase pattern corresponding to the wall and tube beams are imprinted. We also diffract the beams into different directions using phase “wedges”. We focus the beams to create a first far-field image in the splitting focal plane, where the beams are separated by 1.3 mm. This allows us to use a pick-off mirror to deflect only the wall beams by 90 degrees. Further lenses recollimate the beams and focus them again onto the atoms, creating the dark trap.

Alex Gaunt has assessed in [24] the steepness of the beam shaping apparatus, both theoretically and on the bench, parametrizing the radial and axial directions as power-laws. The tube beam can be parametrized by a radial power-law $\propto r^s$ with an exponent $s \gtrsim 7$, while the sheet beams have a profile with $s \gg 10$. As the actual shape of the beams are not power-laws, these numbers are only approximations to the actual profile at a given trap depth. Therefore these values should be taken with a pinch of salt and we should not expect that these values give more than a ballpark estimate of how close our trap is to an ideal box. Instead, we will use thermodynamics later on to extract an overall effective power-law that describes the box.

5.2 Magnetic fields

Even if the trap is perfectly dark inside, we still need to compensate the gravitational potential, which introduces a linear gradient along z . We address this by adding a linear magnetic field gradient $B'_g \approx 15.3 \text{ G/cm}$, which induces a force compensating gravity as described in Section 4.2.1. We assess the necessary flatness by considering the density $n(\mathbf{r})$ of fully condensed BEC in the Thomas-Fermi approximation (Eq. (2.30))

$$n(\mathbf{r}) = [\mu - U(\mathbf{r})]/g, \quad (5.1)$$

where μ is the chemical potential, $U(\mathbf{r})$ is the potential and g our usual interaction strength parameter. We require that any potential variation across the trap satisfies $\Delta V \ll \mu$, to make sure there are no significant variations in density.

Let us consider the restrictions this imposes on the accuracy of our magnetic field gradients in a typical trap with length $L = 63 \mu\text{m}$ and radius $R = 15 \mu\text{m}$, loaded with 50 000 atoms. We require that the variation of the magnetic energy according to Eq. (3.2) across the trap is below $\mu \lesssim k_B \times 1 \text{ nK}$, which means we need to stabilize the gradient to within $10^{-4} B'_g$. We accomplish this by releasing the cloud in ToF with the same gravity compensating field. If gravity is perfectly compensated, the centre of mass of the cloud does not move. We use up to 0.2 s ToF and can localise the centre of mass with a precision

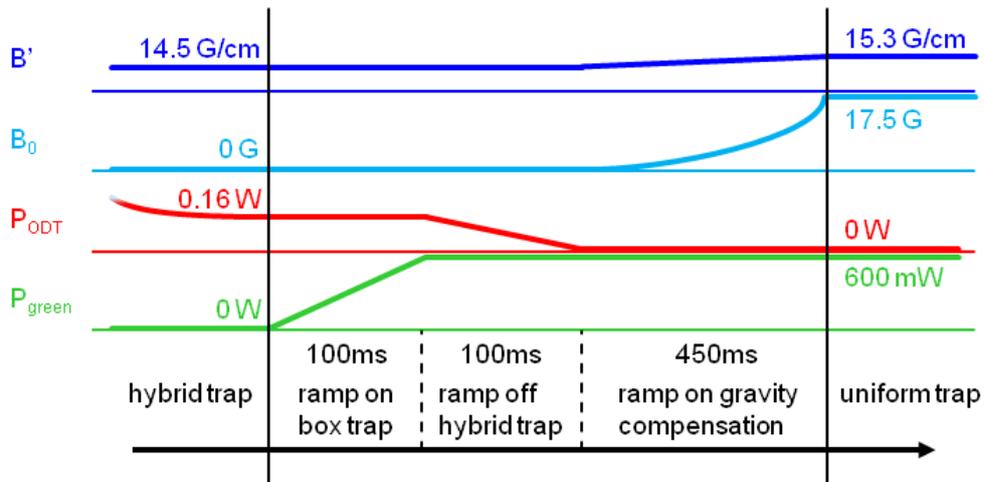


Figure 5.2: Laser power levels and magnetic fields during the transfer of the atoms from the harmonic to the uniform trap. The compensation fields are not included. Figure taken from [27, 70].

of several pixels ($2.5 \mu\text{m}$) on the image. This allows us indeed to detect gradient variations of $< 10^{-4} B'_g$. Another correction to take into account is the curvature of the magnetic field in the horizontal plane (see Section 4.2.1). The curvature of $\omega = 2\pi \times 1.6 \text{ Hz}$ gives a potential variation of $\approx 0.5 \text{ nK}$ across the box and is negligible for thermodynamics.

Note that we have been playing devil's advocate in this calculation by choosing μ as our energy benchmark. For the behaviour of thermal clouds (such as the critical point), $T \approx 50 \text{ nK}$ is the more appropriate and also more forgiving energy level. This changes in Chapter 7, where we consider the ground state properties of our condensed gas.

5.3 Loading of the Uniform trap

After having considered the tools to create a uniform trap, we will now discuss the experimental procedure of loading atoms into the uniform trap. First, we produce degenerate clouds in the harmonic trap as described in Chapter 4. We then transfer those precooled atoms into the uniform trap.

5.3.1 Transfer procedure

Typically, we evaporatively cool the gas in the harmonic (hybrid) trap down to $T \approx 120 \text{ nK}$, when the cloud size is similar to the size of our optical box (see Fig. 5.3). We then ramp up the green walls, followed by switching off the ODT beam. Typically, we use a total laser power of $P = 600 \text{ mW}$, corresponding to a trap depth $U_0 > 1 \mu\text{K}$. The power level is not critical; using excessive power, however, leads to heating. Too little

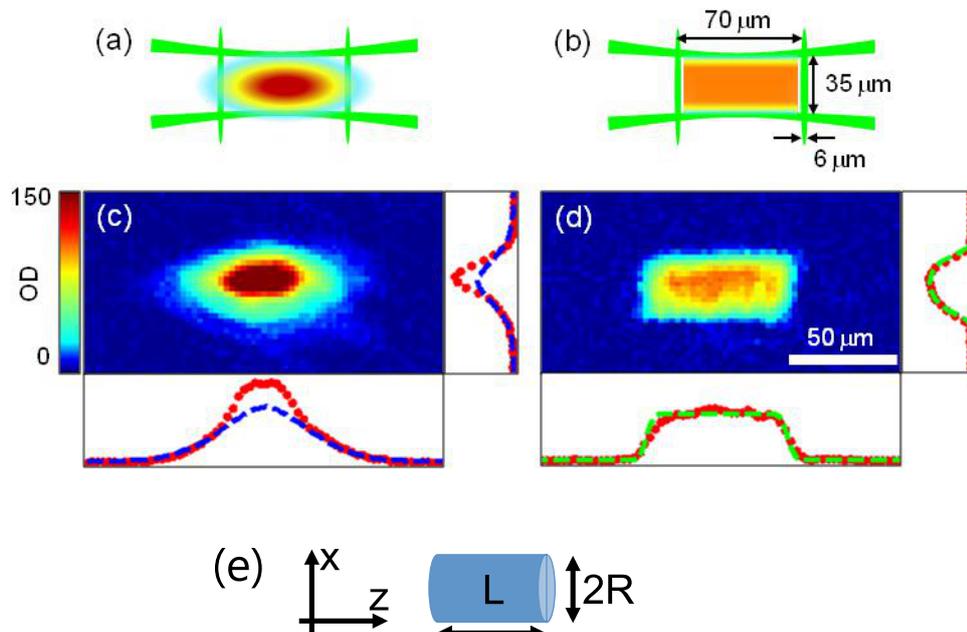


Figure 5.3: Transfer of the cloud from the harmonic to the uniform trap. (a) Sketch of the harmonic trap surrounded by the uniform walls. (b) The cloud in the dark trap after switching off of the harmonic confinement. (c) Experimental column density of the harmonic trap before transfer. The line-densities (red) are shown next to the column density intensity image. The dashed blue lines are fits to the thermal component. (d) Experimental column density of the cloud in the uniform trap. The line-densities (red) are shown below next to the column density intensity image. The dashed green lines are fits to the line density assuming a uniform density distribution. (e) The coordinate system used for the rest of the chapter. Figure taken from [24].

power leads to incomplete transfer, as it is effectively truncating the thermal distribution. Typically, we capture around 80% of the atoms. After the transfer into the dark trap is complete, we ramp up the magnetic gradient to its gravity compensating value B'_g . Simultaneously, we ramp up the vertical bias field to cancel the magnetic gradient in the horizontal plane. In addition to that, we use additional compensation coils (item “11” in Fig. 4.1) to fine tune the magnetic fields. The magnitude and variation of the magnetic fields produced by the compensation coils during transfer is optimized empirically. An overview of the laser powers and magnetic fields during transfer is shown in Fig. 5.2.

Phase-space density during transfer

Before the transfer into the uniform trap, the gas is usually partially condensed, but the BEC is lost during the transfer into the box trap. Even perfect adiabatic transfer does not retain phase space density, if the power-law is changed [48]. This effect has been exploited to achieve condensation solely by deformation of the trap [72]. We will now derive the change in phase-space density when transferring the gas between different power-law

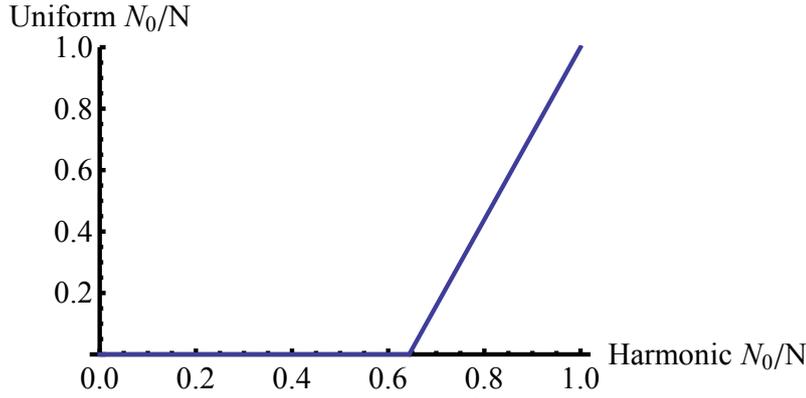


Figure 5.4: Change in condensed fraction N_0/N when transferring a partially condensed gas from a harmonic to a uniform trap.

traps. We consider both a classical gas as well as a condensed gas, and apply our results to the transfer from a harmonic to a uniform trap.

Classical gas Let us first consider a hot gas at $T \gg T_c$. The peak phase space density D_0 can be related to the entropy S as [48]

$$D_0 = \exp\left(\alpha + 1 - \frac{S}{Nk}\right). \quad (5.2)$$

α is the power-law parameter introduced in Chapter 2. As a reminder, in a power-law $U(\mathbf{r}) \propto \mathbf{r}^s$, we have $\alpha = 3/2 + 3/s$. Transferring adiabatically ($S = \text{const.}$) from a harmonic trap with $\alpha_1 = 3$ to a uniform trap $\alpha_2 = 3/2$ leads to a significant decrease in peak phase-space density:

$$\frac{D_2}{D_1} = e^{\alpha_2 - \alpha_1} \approx 0.22. \quad (5.3)$$

Since we also know that collision rates and evaporative cooling efficiency are worse in a uniform trap than in a hybrid trap (Section 3.3), this suggests we should transfer as late as possible.

Condensed gas For a partially condensed gas in a power-law potential, we can express the entropy in terms of the condensed fraction (Table 2.1)

$$\frac{S}{Nk} = (\alpha + 1) \frac{\zeta(\alpha + 1)}{\zeta(\alpha)} \left(1 - \frac{N_0}{N}\right) \equiv f(\alpha) \left(1 - \frac{N_0}{N}\right), \quad (5.4)$$

where we introduced the function $f(\alpha)$ for brevity. Changing the power-law from α_1 to α_2 with constant entropy yields a new condensed fraction:

$$\left(\frac{N_0}{N}\right)_2 = 1 - \frac{f(\alpha_1)}{f(\alpha_2)} \left(1 - \left(\frac{N_0}{N}\right)_1\right). \quad (5.5)$$

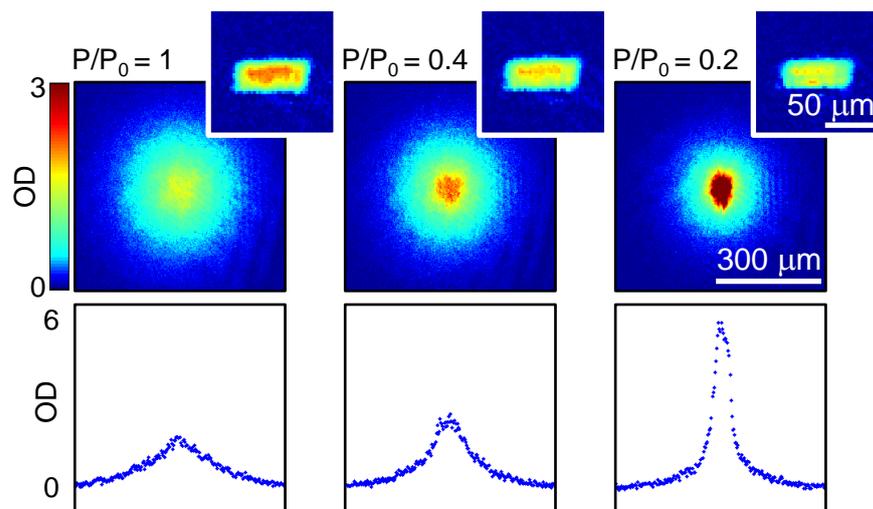


Figure 5.5: Condensation by forced evaporation in a uniform trap. The power P of the trapping potential is reduced to a fifth of its initial value P_0 . Top panels: While the in situ shape (small insets) remains relatively unchanged, the momentum distribution as observed in ToF shows a clear bimodality. Bottom panels: We show the cuts through the respective column density distributions in ToF from the top panels. Figure taken from [24].

This means that the condensate is only retained if the initial condensed fraction is large enough, i.e. if

$$\left(\frac{N_0}{N}\right)_1 > 1 - \frac{f(\alpha_2)}{f(\alpha_1)} \approx 64\%. \quad (5.6)$$

We see in Fig. 5.4, that we need large condensed fractions in the harmonic trap to get a notable condensed fraction in the uniform trap. Due to residual heating and instability at low powers, we found that large condensed fractions in the harmonic trap ($> 50\%$) are hard to achieve and even harder to measure. Hence, further evaporation after transfer in the uniform trap is necessary to achieve condensation.

Also, early transfer into the uniform trap allows us to minimize three-body losses in the harmonic trap. The three-body loss rate scales $\propto n^3$, and is therefore strongly enhanced in dense regions. This explains why these losses are prominent in the parabolic condensate of a harmonic trap, but virtually non-existent in a uniform trap with similar spatial extent and atom number.

5.4 Evaporation and Condensation

After the transfer into the box trap the cloud contains typically $N \approx 6 \times 10^5$ atoms at $T \approx 130$ nK. From this point, we cool the gas below T_c by forced evaporative cooling in the box trap, lowering the trapping power P in an exponential ramp with a 0.5 s time

constant. Fig. 5.5 qualitatively illustrates the effects of evaporation and condensation in the box trap. We show images of the cloud both in situ and after $t = 50$ ms of ToF expansion from the trap.

While in a harmonic trap cooling results in a real-space bimodal distribution, as shown in Fig. 5.3, condensation in a uniform trap does not induce any dramatic change in real space. Condensation occurs exclusively in momentum space (i.e., in ToF), and the emerging bimodality in the expanding density profile allows us to detect the formation of a BEC. Besides the bimodality, we observe how the aspect ratio inverts, again a typical smoking gun of condensation. Typically, we achieve condensates with up to 2×10^5 atoms.

5.5 Volume and in situ shape

The initially most striking signature of the transfer into the box is the change of the in situ density profile. Fig. 5.2 shows the density profiles before and after the transfer. The clouds were imaged using high intensity imaging as described by Eq. (4.15).

5.5.1 Harmonic density profile

Integrating the density of the thermal component of a partially condensed gas in a harmonic trap Eq. (2.23), gives us the line density of the thermal component along any direction (here the x-axes)

$$n(x) \propto g_{5/2} \left(\exp \left[-(1/2)mw^2x^2 \right] \right). \quad (5.7)$$

Unsurprisingly, we find this describes our measured line density in Fig. 5.3c in the wings, while a small condensate is visible in the centre.

5.5.2 Uniform density profile

For an ideal uniform box, we expect the same uniform profile for the thermal and condensed component (neglecting the healing length at the edges). We use the coordinate system shown in Fig. 5.3, in which the x-axis is vertical (parallel to gravity) and the y-axis runs in the horizontal plane parallel to the rotation axis of the cylindrical box trap. Integrating our can-shape trap of length L and radius R along z to obtain the column density, we expect a rectangle distribution along y . Along x we obtain $\sqrt{1 - (x/R)^2}$ due to the integration across the circular can-profile, and we can write the column density as

$$n(x, y) \propto \Pi \left(\frac{y}{L} \right) \sqrt{1 - (x/R)^2}, \quad (5.8)$$

where $\Pi(y/L)$ is the rectangle function, which is one for $|y| < L/2$ and zero otherwise. This is separable and the line density profiles are trivially given by the respective components of above equation. A measurement of our real box will deviate from this for two reasons. First, the box potential is not perfect due to the diffraction-limited steepness of the walls and other experimental imperfections. These imperfections are what we want to quantify. Second, the measurement of the density profile is not perfect. We need to take into account the resolution of our imaging system which smears out the edges. Mathematically this can be modelled by convolving Eq. (5.8) with an assumed Gaussian point spread function (PSF)

$$PSF(x, y) = \frac{1}{2\pi\sigma^2} \exp\left[-\frac{x^2 + y^2}{2\sigma^2}\right], \quad (5.9)$$

with a waist $w_0 = 2\sigma$. The measured density $\tilde{n}(x, y)$ is then given by the convolution $n(x, y) * PSF(x, y)$

$$\tilde{n}(x, y) \propto \int \int \Pi\left(\frac{y'}{L}\right) \sqrt{1 - (x'/R)^2} \exp\left[-\frac{(x - x')^2}{2\sigma^2}\right] \exp\left[-\frac{(y - y')^2}{2\sigma^2}\right] dx' dy'. \quad (5.10)$$

Again, this expression is separable in x and y and can be easily integrated to obtain the corresponding line density profiles. The uniform profile in Fig. 5.3d fits this prediction quite well assuming a PSF with $\sigma = 2.5 \mu\text{m}$.

5.5.3 General power-law traps

The description above is appropriate for an ideal box. Given that no box in nature is ideal, we return to the power-law approximation in Eq. (2.23) to predict the in-trap thermal density for a degenerate gas at $T = T_c$. We will discuss to what extent measuring the in-trap density profile allows us to extract a power-law estimate. In the spirit of the characterization of the trap on the bench mentioned earlier [70], we assume a radial power-law of power n and infinitely steep walls for the sheets. We obtain a thermal density

$$n'(\mathbf{r}) \propto \Pi(y) g_{3/2} \left(-\sqrt{(x^2 + z^2)^n} \right). \quad (5.11)$$

As we are only interested in the profile, we use dimensionless variables x, y, z . Calculating the line density profile along x is analytically not trivial, but can be done numerically. In Fig. 5.6 we show the resulting profiles for various power-laws compared with the ideal circular profiles. We see that for the range of power-laws in question, it becomes very hard to differentiate them for $n > 6$, especially when including the imaging resolution blur by convolving with the PSF. Experimentally, shot noise and other imperfections make it very

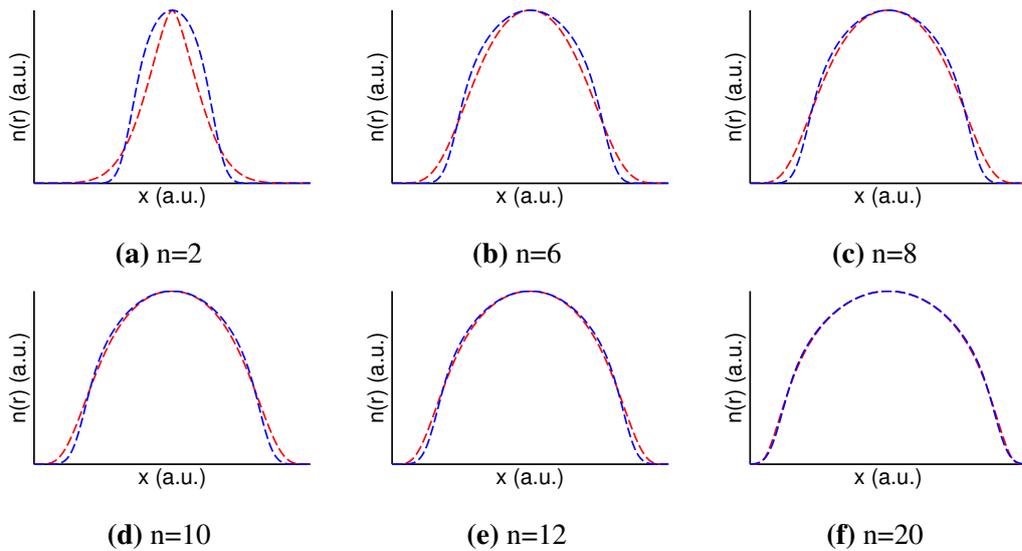


Figure 5.6: The theoretical column/line density profile (red) of various power-laws r^n at T_c along the radial direction of the box. They are compared to the ideal prediction in an infinitely steep box (blue). Both are convolved with the same imaging resolution PSF to simulate the observed profile. The densities are scaled to unity and the widths of the power-laws and ideal box are matched for simplicity. We see that we can easily distinguish power-laws $\lesssim 6$ from an ideal box profile, but it gets increasingly hard for higher exponents. For very high power-laws, the profiles are identical to the ideal box profile as expected.

hard to differentiate those profiles reliably. To add to the difficulties, we are never exactly at T_c . This means, we have either a uniform density offset from the condensed density n_0 for $T > T_c$, or we need to evaluate Eq. (2.23) for $T < T_c$ with an unknown $\mu < 0$. Given these predictions, we see that obtaining a good estimate for the power-law exponent from the in situ density is not practical. For now, we approximate the real in situ density distribution as identical with the density in an ideal uniform trap. We will see from further analysis that the power-law is steep enough to justify this assumption.

5.6 Thermometry

We now assess the flatness and uniformity of our potential by considering the thermodynamics of the trapped gas. By contrasting our findings with the prediction for an ideal box and a harmonically trapped gas, we quantify the uniformity of our gas.

5.6.1 Scaling laws

We recall from Section 2.2.2, that the thermodynamics of a power-law potential

$$U(\mathbf{r}) = \sum_{i=1}^3 \epsilon_i \left| \frac{x_i}{L_i} \right|^{s_i} \quad (5.12)$$

scale with the parameter

$$\alpha = 3/2 + \sum_{i=1}^3 1/s_i. \quad (5.13)$$

For a uniform box, we have $s_i = \infty$ and thus $\alpha = 3/2$; for a harmonic trap with $s_i = 2$ we have $\alpha = 3$. From theoretical limitations and bench measurements of the trap, [70] predicts power-laws of $s \gg 10$ axially and $s \gtrsim 7$ radially. Hence, we expect $\alpha \lesssim 3/2 + 1/\infty + 2/7 \approx 1.8$. We have seen that the measurement of thermal in situ profiles agrees with these predictions, but does not allow great accuracy. Instead, we consider the in-trap momentum distribution from Eq. (2.25)

$$n_{PL}(\mathbf{p}) \propto (kT)^\delta g_\delta \left(\exp(-\beta(\mu - \mathbf{p}^2/2m)) \right), \quad (5.14)$$

which also depends directly on $\alpha = \delta + 3/2$.

The traditional way to extract the momentum distribution is to image the cloud in ToF. Unfortunately, even with a long time of flight ($t > 50$ ms), we cannot use Eq. (5.14) to describe the density distribution. At typical temperatures ($T \sim 100$ nK), the distance travelled and the initial size of the trapped cloud are comparable. While we can evaluate Eq. (2.37) to describe the expansion at any ToF numerically, this is very unwieldy for fitting purposes.

We have established in Section 2.3.2 that we cannot generally write the distribution in ToF $n(r, t)$ as the convolution of the in situ density with the momentum distribution - unless the in situ distribution is uniform. We have established earlier that the in-trap *density* distribution is de-facto uniform. Hence, we approximate the ToF-distribution $n(\mathbf{r}, t)$ as the convolution of the density of an ideal box $n_{box}(\mathbf{r}')$ with the power-law momentum distribution $n_{PL}(\mathbf{p})$ from Eq. (5.14)

$$n(\mathbf{r}, t) \propto \int n_{PL}(\mathbf{p} = (\mathbf{r} - \mathbf{r}')m/t) \times n_{box}(\mathbf{r}') d\mathbf{r}' \quad (5.15)$$

$$\propto g_\delta \left(\exp \left(-\beta\mu - (\mathbf{r} - \mathbf{r}')^2 \frac{\beta m}{2t^2} \right) \right) \times \Pi \left(\frac{y'}{L} \right) \Pi \left(\frac{\sqrt{x'^2 + z'^2}}{R} \right) d\mathbf{r}'. \quad (5.16)$$

Integration along the line of sight z yields the column density \tilde{n} , which is observed directly

in our experiments

$$\tilde{n}(x, y, t) \propto \int g_{\delta+1/2} \left(\exp \left(-\beta \left[\mu - (x - x')^2 \frac{m}{2t^2} + (y - y')^2 \frac{m}{2t^2} \right] \right) \right) \quad (5.17)$$

$$\times \Pi \left(\frac{y'}{L} \right) \Pi \left(\frac{x'}{R} \right) \sqrt{1 - (x'/R)^2} dx' dy'. \quad (5.18)$$

Further integration along y yields the line density \hat{n}

$$\hat{n}(x, t) \propto \int g_{\delta+1} \left(\exp \left(-\beta \left[\mu - (x - x')^2 \frac{m}{2t^2} \right] \right) \right) \quad (5.19)$$

$$\times \Pi \left(\frac{x'}{R} \right) \sqrt{1 - (x'/R)^2} dx'. \quad (5.20)$$

Both the column density and line density can be easily evaluated numerically. The convolution adds a small initial size correction to the ToF-expansion of the momentum distribution.

5.6.2 Numerical implementation

To implement this into a numerical fitting algorithm for the column density, we introduce the following reduced variables: the reduced chemical potential $\tilde{\mu} = \beta\mu$, the thermal widths $R_{x,y} = \sqrt{\frac{kT_{x,y}t^2}{m}}$, the centres of the cloud x_c, y_c and a scaling parameter $A(T)$. Additionally, we use a normalized convolution matrix (kernel) $K(x, y)$, which is proportional to the in situ column density of an ideal box given by Eq. (5.8). For simplicity we assume that the measured optical densities have been converted to column densities according to Eq. (4.9). Eq. (5.18) can now be written as the convolution

$$\tilde{n}_{fit}(x, y) = \frac{A}{R_x R_y} g_{\delta+1/2} \left(\exp \left[\tilde{\mu} - \frac{(x - x_c)^2}{2R_x^2} - \frac{(y - y_c)^2}{2R_y^2} \right] \right) * K(x, y). \quad (5.21)$$

We have introduced different thermal widths for x and y to take a small harmonic magnetic trapping potential into account, which compresses the distribution along y during ToF. We therefore extract two different temperatures T_x and T_y . As T_x is not compressed, it is more reliable. Unless mentioned otherwise, we will set $T = T_x$.

Usually, we deal with a bimodal distribution including both a thermal component described by Eq. (5.21) and a condensate of unknown shape occupying the centre of the image. To obtain a fit for the thermal component, we ignore the region in the centre containing the condensate and consider the thermal wings only. Numerically, we implement this using a mask excluding an elliptical region around the centre of the condensate

(x_0, y_0)

$$M(x, y) = \begin{cases} 1 & \text{for } \frac{(x-x_0)^2}{r_{0x}} + \frac{(y-y_0)^2}{r_{0y}} > 1 \\ 0 & \text{for } \frac{(x-x_0)^2}{r_{0x}} + \frac{(y-y_0)^2}{r_{0y}} < 1 \end{cases}. \quad (5.22)$$

We then use Matlab's least-square fitting algorithm to find the best match of $\tilde{n}_{fit}(x, y)$ to the measured column density \tilde{n}_{exp} outside the masked area by minimizing the sum-of-squares of the matrix

$$|\tilde{n}_{fit}(x, y) - \tilde{n}(x, y)_{exp}| \times M(x, y). \quad (5.23)$$

Usually, we vary $A, R_x, R_y, x_c, y_c, \tilde{\mu}$ as free parameters. The procedure can be easily modified to work in one dimension on line densities.

Divergence issues

A fine point that can cause some numerical headaches is that $g_z(x)$ diverges for $z \leq 1$ and $x \rightarrow 0$. While this is not an issue for the line density $g_{\delta+1}(x)$ with finite power-laws with $\delta > 0$, the column density $g_{\delta+1/2}(x)$ diverges if $\delta < 1/2$. It is important to keep in mind that we do not sample a continuous distribution, but an averaged value for each pixel. This is irrelevant for slowly varying distributions far away from the singularity. Near the divergence, we need to integrate x and y over the pixel area. The integration along two spatial directions increases the order of the polylogarithm by one to $\delta + 3/2$ and ensures a finite value.

Even when considering continuous functions on a theoretical level, the convolution with a finite size kernel has the same effect. As we always deal with a finite-size, real BEC and not a point source, this prevents an unphysical singularity in the density.

Extracting thermodynamic properties from the fit

After finding the best fit, we can easily extract thermodynamic properties from the fit parameters. The total thermal number N' can be extracted by integrating Eq. (5.21)

$$N' = 2\pi A g_{\delta+3/2}(\tilde{\mu}). \quad (5.24)$$

The condensate number N_0 is calculated by summing over excess atoms within the condensate region:

$$N_0 = \int \tilde{n}_{exp}(x, y) - \tilde{n}_{fit}(x, y) \times (1 - M(x, y)) dx dy. \quad (5.25)$$

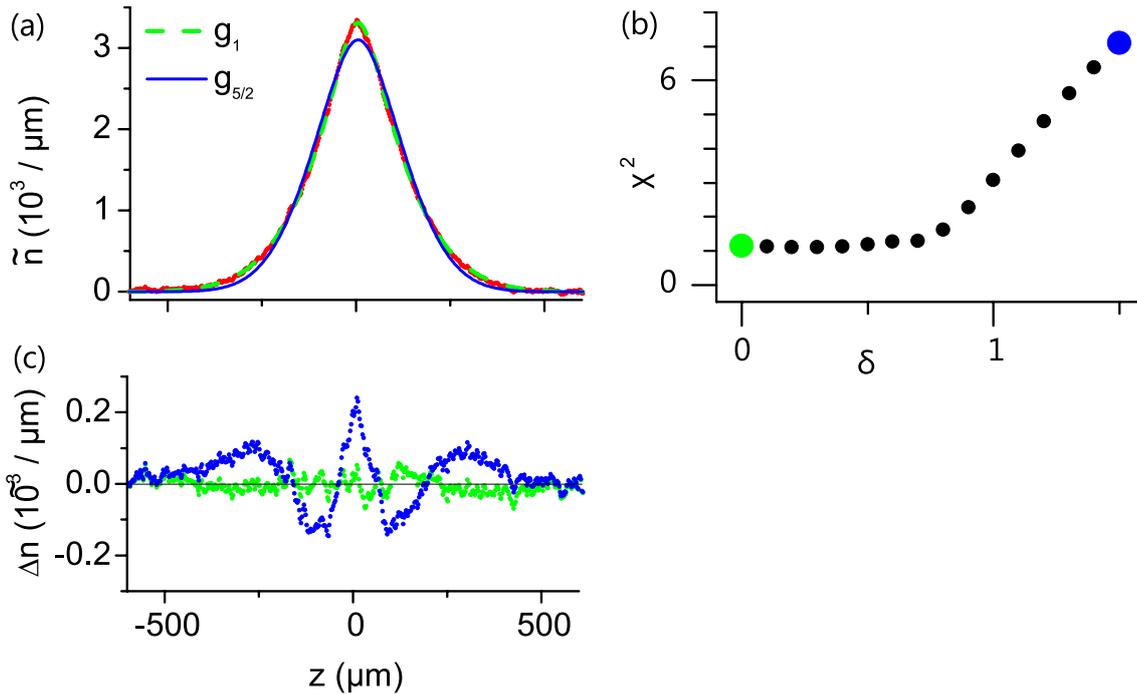


Figure 5.7: We search for the best fit to the line-density \hat{n} of a thermal, degenerate cloud for various power-law parameters δ . In (a), we compare the experimental line-density along x (red) to the uniform fit (green dashed) and the harmonic fit (blue). In (c), the fitting residues are shown. (b) shows the reduced fitting χ^2 against δ . Figure adapted from [24].

The temperature $T = T_x$ is then

$$T = \frac{R_x^2 m}{k t^2}, \quad (5.26)$$

while trivially

$$\mu = \tilde{\mu} k T. \quad (5.27)$$

In practise, all these continuous functions and integrals need to be converted to discrete matrices and sums, which is here omitted for readability.

5.7 How uniform is the trap?

5.7.1 Conclusions from a single cloud profile

As an initial test of our fitting procedure we scrutinize the line density of a non-condensed degenerate cloud in ToF. We apply various fits (using the line density version of Eq. (5.21)) starting with a uniform profile $\delta = 0$ to a harmonic profile $\delta = 3/2$. We find that the uniform profile fits the experimental data substantially better than the harmonic profile. Fig. 5.7 illustrates this in various ways. The difference $\Delta \hat{n}$ between the fitted profiles and

the experimental line densities shows a clear preference for the uniform profile. While the harmonic profile has clear systematic deviations of $\Delta\hat{n}$ from zero, the uniform profile fulfills $\Delta\hat{n} \approx 0$. To quantify this, we fitted the profile for various value of δ and calculated the reduced χ^2 . We find that χ^2 is approximately constant for values of $0 \leq \delta \leq 0.7$, while values above that are strongly excluded.

This result is in line with our previous boundaries from theoretical predictions and in situ volumes. While it excludes a harmonic trap strongly, it does not place very strict bounds on the effective power-law ($s > 4$).

5.7.2 Conclusions from several profiles

Looking at a single image is clearly not enough to determine the effective power-law accurately. We therefore consider a series of images across a range of temperatures $90 \text{ nK} < T < 135 \text{ nK}$, which is above or just below condensation. We exploit the fact that A as defined in the previous section is not an arbitrary scaling constant - even though we treated it as a free fitting parameter. Evaluation of the *column* density with the relevant prefactor yields

$$A = \left(\frac{T}{\theta}\right)^\alpha. \quad (5.28)$$

Here, we switched again to $\alpha = 3/2 + \delta$ for compactness, as it is the natural variable for many thermodynamic relations in power-laws (see table 2.1). θ is a temperature independent constant that absorbs various numerical and physical prefactors. While we are not interested in the absolute magnitude of θ , there should be no variation with temperature, i.e. $d\theta/dT = 0$ for the correct power-law.

We fit the series of pictures with power-laws of $1.5 \leq \alpha \leq 2$ (corresponding to power-law exponents between $s = \infty$ and $s = 6$) and extract $\theta(\alpha)$ for all pictures. As we can see in Fig. 5.8a, $\theta_\alpha(T)$ is quite flat for $\alpha \leq 1.8$, while for larger power-laws there is a clear temperature dependence. We quantify the deviations in first order using linear fits. The slope of the fits are shown in Fig. 5.8b, confirming that the effective power-law is $\alpha \leq 1.8$, with the best candidate being $\alpha = 1.7$ ($s = 15$).

To underline that this analysis finds a good fit and not only selects the best out of a bunch of bad options, Fig. 5.9 shows a two-dimensional cloud profile near condensation fitted with $\delta = 0.2$ and $\delta = 1.5$. We can see by looking at the fitting residues that $\delta = 0.2$ yields a good fit with no major deviations from the experimental data. $\delta = 1.5$, however, shows significant deviations near the cloud centre, and does not represent the data well.

In the original publication [24], we used a slightly different approach. First of all, we used the 1-D line-density for the analysis. The advantage of this is simplicity and fitting speed. On the down side, a lot of information can be lost due to averaging prior to fitting.

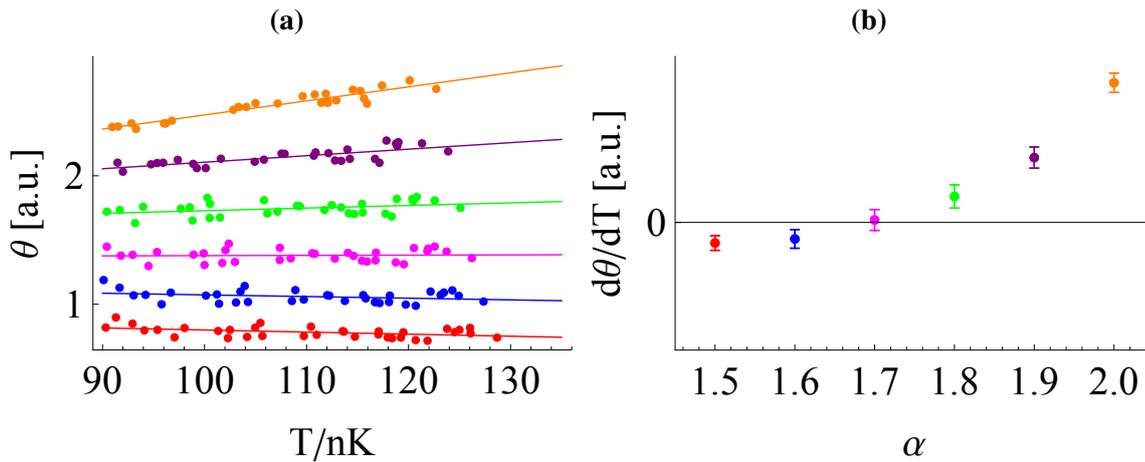


Figure 5.8: Considering the momentum distribution over a range of temperatures allows us to assess the uniformity of our trap. (a) The ideal gas scaling in Eq. (5.28) predicts that the parameter $\theta(T)$ is constant for a range of temperatures for the correct value of α . We show $\theta(T)$ for $\alpha = 1.5$ (bottom red) up to $\alpha = 2$ (top orange). The data is approximated by linear fits. (b) The slopes of the fits in (a) are plotted against the respective value of α .

Notably, to avoid the condensate region, we have to exclude a whole strip of the image, rather than just an elliptical region around the condensate. Also, we included temperatures far below condensation, which has several drawbacks. First of all, the effective power-law is an approximation and we cannot assume that one power-law value is appropriate over an arbitrary range of temperatures. In terms of Fig. 5.8a, this would imply that there is no “correct” α to ensure $\theta(T) = \text{const}$ for all temperatures.

More importantly, we made the inherent assumption that the condensate does not influence the thermal cloud and that we can use ideal gas theory to describe the gas below the critical temperature T_c . While we are going to show in Chapter 6 that this is a good approximation, we should not draw any premature conclusions at this stage. For instance, it has been shown [16], that the presence of a condensate alters the thermal distribution in harmonic traps.

5.7.3 The critical point

Using $\alpha = 1.7$ to fit the clouds, we will compare the observed critical temperature T_c with the prediction for a uniform box. We will show that the quantum mechanical treatment yields the same transition temperature as the semi-classical treatment in Section 2.2.2. In addition to that, we consider corrections to the transition temperature due to the finite size of the trap as well as interactions.

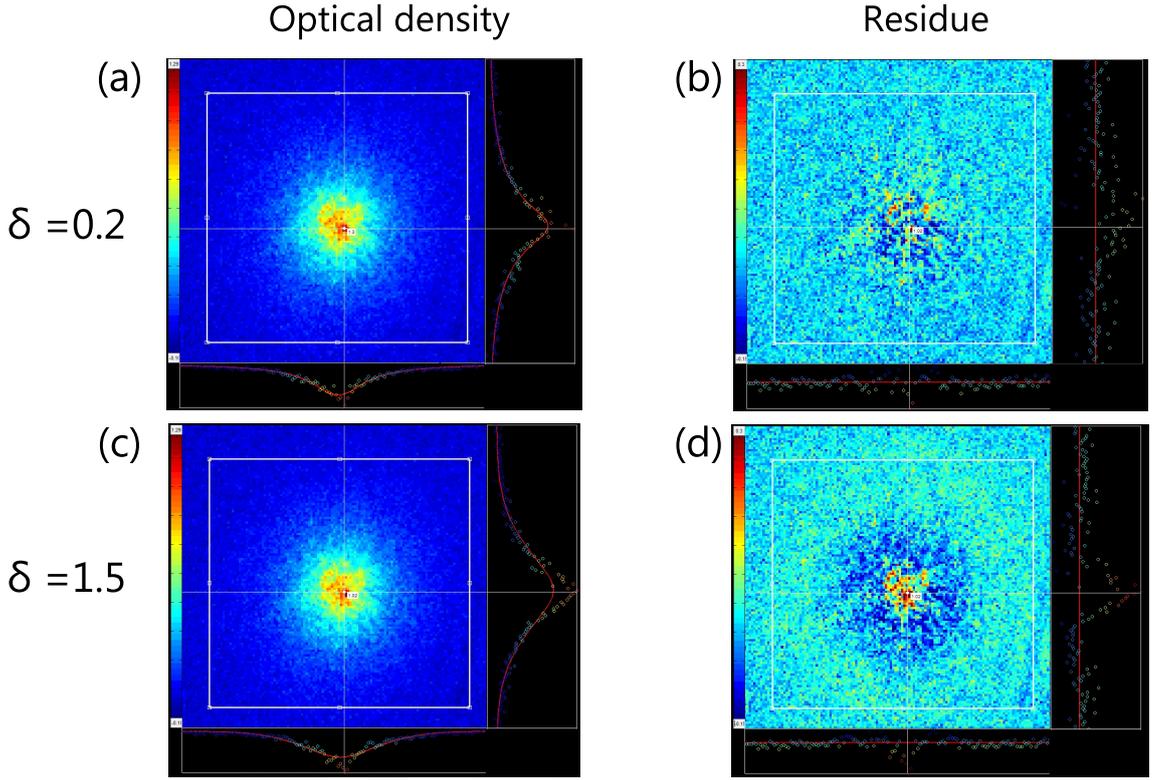


Figure 5.9: A two dimensional optical density image fitted with different values for δ . The cloud is near the critical point. The top row is fitted with the best value $\delta = 0.2$, while the bottom row shows the same cloud fitted with $\delta = 1.5$, which would be the right parameter for a harmonic trap. (a) An intensity plot of measured cloud profile. The panels below and to the right show cross sections superimposed with the respective fit (red). The colour map ranges from -0.1 to 1.29. (b) The fitting residue. There are no significant deviations from zero, apart from the cloud centre where the BEC starts forming in various locations. Note that the colour map ranges now from -0.15 to 0.3. (c) Major deviations between the fitted and measured cross sections can be seen in the cloud centre which is most sensitive to the chosen value of δ . (d) These deviations are clearly visible as coloured rings in the residue plot; the fit does not describe the cloud shape well.

Quantum mechanical density of states

Using the semi-classical description, we have already derived the critical number N_c for a given temperature as:

$$N_c = V \frac{\zeta(3/2)}{\Lambda^3} = V \zeta(3/2) \left(\frac{mkT}{2\pi\hbar^2} \right)^{3/2}. \quad (5.29)$$

To do so, we approximated the density of states by the semi-classical distribution

$$N \approx \int g(\epsilon) f(\epsilon) d\epsilon \approx \int f(\mathbf{r}, \mathbf{p}) / h^3 d\mathbf{r} d\mathbf{p}. \quad (5.30)$$

We quickly recapitulate the derivation of the quantum-mechanical density of states in the usual form, before considering additional corrections which give rise to finite size corrections to N_c . The energy eigenvalues of a three-dimensional cube with length L are given by

$$\epsilon(n_x, n_y, n_z) = \frac{\pi^2 \hbar^2}{2mL^2} (n_x^2 + n_y^2 + n_z^2) \equiv E_1 (n_x^2 + n_y^2 + n_z^2), \quad (5.31)$$

where $n_i = (1, 2, 3, \dots)$. We can consider those eigenvalues as points in a three-dimensional space with coordinates (n_x, n_y, n_z) . The total number of states $G(\epsilon)$ with energy $E < \epsilon(k, l, m)$ is approximated by the volume of a sphere with radius $\sqrt{\epsilon/E_1}$

$$G(\epsilon) = \frac{1}{8} \left[\frac{4\pi}{3} \left(\frac{\epsilon}{E_1} \right)^{3/2} \right], \quad (5.32)$$

where the factor $1/8$ makes sure that we only count the states in the positive octant of the coordinate system. Differentiation yields the density of states $g(\epsilon)$:

$$g(\epsilon)d\epsilon = \frac{\pi}{4} E_1^{-3/2} \epsilon^{1/2} d\epsilon. \quad (5.33)$$

If we use this in Eq. (5.30) to calculate N_c we obtain the same result as the in Eq. (5.29), which has been derived semi-classically.

Finite size corrections

While the above derivation is correct in the thermodynamic limit, we need to take finite size effects into account when evaluating the density of states of a finite sized box. So far, we have ignored the states with $n_i = 0$ in Eq. (5.32). Quantum-mechanically, these are not accessible as all quantum numbers need to fulfill $n_i \geq 1$. We therefore need to subtract the states on the coordinate planes to obtain a better estimate for the total number of states

$$G(\epsilon) = \frac{1}{8} \left[\frac{4\pi}{3} \left(\frac{\epsilon}{E_1} \right)^{3/2} - 3\pi \frac{\epsilon}{E_1} \right], \quad (5.34)$$

yielding the corrected density of states

$$g(\epsilon)d\epsilon = \left[\frac{\pi}{4} \frac{\epsilon^{1/2}}{E_1^{3/2}} - 3\pi \frac{1}{E_1} \right] d\epsilon. \quad (5.35)$$

The observant reader will have noticed that we are still neglecting over-counting on the axes and the origin - these effects are of higher order and negligible [73].

Unfortunately, we cannot integrate this new density of states in a straightforward man-

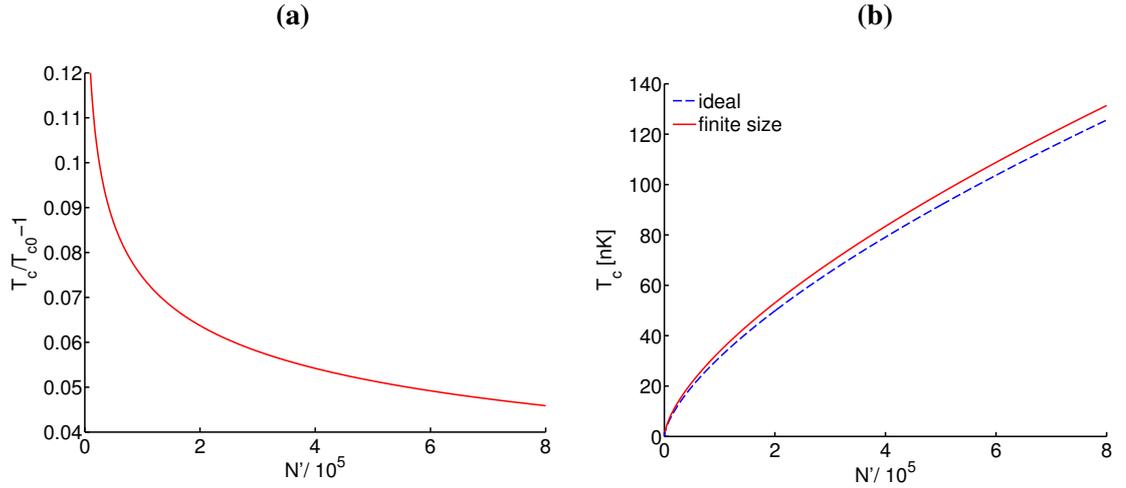


Figure 5.10: Critical temperature in a box with $L = 63 \mu\text{m}$ and $R = 15 \mu\text{m}$. (a) Relative finite size shift $T_c/T_c^0 - 1$. (b) Critical temperature for a range of common atom numbers.

ner, i.e. we cannot use

$$N_c = \int_0^{\infty} f(\epsilon, \mu = 0)g(\epsilon) \quad (5.36)$$

to obtain a revised N_c , as the correction to the density of states causes the integral to diverge. This can be solved by introducing a lower cut-off energy on the order of the ground state energy $\epsilon_0 = 3E_1$. The details are rather intricate and are described in [73]. The final estimate for the critical temperature T_c is

$$T_c = T_c^0 \left[1 - \frac{2}{\sqrt{3\pi}} \frac{1}{\zeta(3/2)} \sqrt{\epsilon_0/kT_c^0} \ln(\epsilon_0/kT_c^0) \right], \quad (5.37)$$

where T_c^0 is the critical number in the thermodynamic limit as obtain by inverting Eq. (5.29).

So far, we have assumed that our trap is an isotropic cube. The generalization of Eq. (5.37) to an anisotropic box with the same aspect ratio as our trap has been done in [27]. This anisotropy scales the correction term in Eq. (5.37) with a shape factor of order 0.95, which is experimentally irrelevant. Also, we should note that exact summing over the states (also done in [27]) introduces slightly larger deviations, which are still below our experimental accuracy. A more details analysis of the finite size corrections (especially in a cylindrical geometry) might become necessary once precision measurements of the phase transition are attempted. For our purposes, we estimate the leading order of the finite size shift with Eq. (5.37), while using the ground state energy of a cylindrical box³

³Its ground state energy is pretty much identical to the ground state energy of a cubic box with the same aspect ratio and volume.

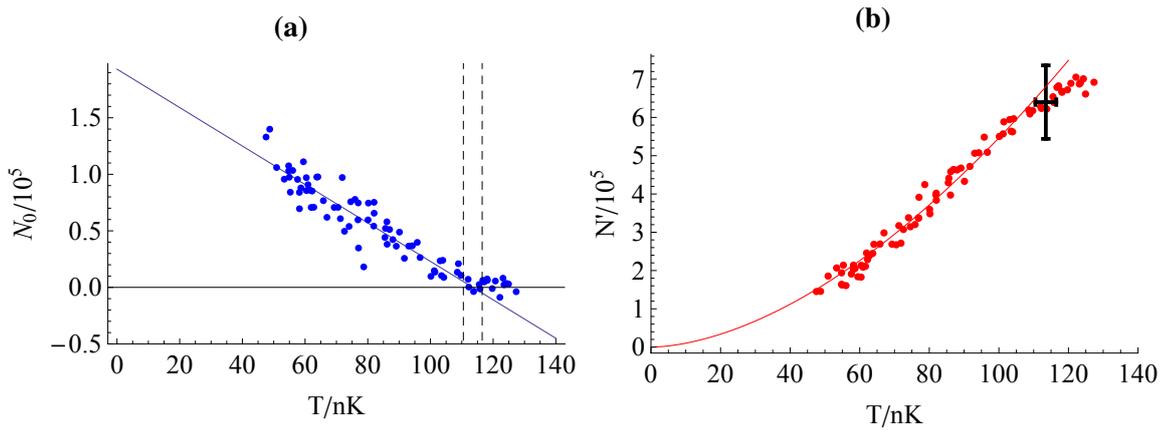


Figure 5.11: Determination of the critical point in a quasi-uniform trap. (a) Extrapolation of the condensed number N_0 to zero to find the experimental critical temperature T_c . (b) The thermal number atom numbers N' (red dots) are plotted against temperature T . Based on the critical temperature $T_c = 113 \pm 3$ nK measured in (a), we determine the theoretical prediction for the critical number $N(T_c)$ (black). We also indicate its errors due to experimental and theoretical uncertainties. The red line is a $T^{1.7}$ fit to the N' below T_c , which is the expected scaling for an ideal gas. Above T_c , $\mu < 0$ and N' departs from this scaling.

as ϵ_0 . Fig. 5.10 illustrates the relative and absolute magnitude of the finite-size T_c shift.

Interaction shift The presence of small interactions in a homogeneous gas are assumed to cause a beyond-mean-field shift of the critical temperature [10]:

$$\frac{T_c}{T_{c0}} - 1 = 1.3an^{1/3}. \quad (5.38)$$

For typical atom numbers, $N < 6 \times 10^5$, this relative shift is smaller than 2.5% and therefore small compared to the other theoretical and experimental uncertainties. For the time being, we ignored the shift. In the future, a precision measurement of T_c using tunable interactions might be used to verify the magnitude of this shift.

5.7.4 Finding the critical point

Now we can calculate T_c accurately, we compare the theoretical prediction with the experimental measurements. We use the same set of data that we have used to calibrate α earlier. In Fig. 5.11a, we extrapolate the condensed number N_0 to zero and find the critical temperature $T_c^{exp} = 113 \pm 3$ nK. From this we deduce the expected critical number N_c using Fig. 5.10b. In Fig. 5.11b, we compare this to the experimental findings. We find that the expected critical number (black marker) matches the experimentally measured thermal number (red) at the transition (i.e., $N_0 = 0$) within error. The error is

mainly introduce due to uncertainties in atom number calibration $\approx 10\%$ and trap volume measurement $\approx 14\%$.

Below T_c , we expect $N' \propto T^\alpha$ with $\alpha = 1.7$ for an ideal gas. The fit to Fig. 5.11b yields good agreement. This is a first indication that the gas behaves like an ideal gas despite interactions, i.e. the thermal gas is not influenced by the condensate. Above T_c , the thermal number N' drops below the $N' \propto T^\alpha$ line, as expected for a non-condensed gas.⁴

A note on θ and N' Now it might be more obvious why we have considered θ only for $T \gtrsim T_c$. If we had applied this method below T_c , we would select the specific α forcing the thermal number to follow the ideal gas scaling $N' \propto T^\alpha$. More specifically, below T_c , the chemical potential is zero, $\mu \approx 0$, and there are only two degrees of freedom in the fits Eq. (5.21): T and α . If we were to select α such that $\theta(T < T_c)$ is constant, we automatically make sure that $N' \propto T^\alpha$. This can be seen by inserting Eq. (5.28) into Eq. (5.24).

In Chapter 6, we will properly scrutinize the saturation of the thermal gas. Hence, here we can only use $T \gtrsim T_c$ to analyse the trap or we risk falling for a self-fulfilling prophecy.

5.7.5 Conclusions

We have created the first truly three-dimensional quasi-uniform trap. By approximating the trap with an effective power-law potential, we have quantified the small deviation from uniformity. All different forms of analysis put a lower bound on the power-law exponent $s \geq 10$, while $s \approx 15$ describes the trap best. We found that the critical point can be well estimated by approximating the trap with an ideal box of appropriate volume.

⁴This illustrates the quantum-statistical character of Bose-Einstein condensation. For a given temperature, the total number of occupiable thermal states scales $N_c \propto T^\alpha$. If the total number of atoms N exceeds N_c , all thermal states are occupied, i.e., $N' = N_c$ and the rest of the atoms occupies the ground state: $N_0 = N - N_c$. If we have fewer atoms than the critical number, $N < N_c$, we can accommodate all the atoms in the thermal states $N' = N < N_c \propto T^\alpha$.

5.7. HOW UNIFORM IS THE TRAP?

Chapter 6

Thermodynamics in a uniform trap

In this chapter, put our uniform trap to use and conduct a range of thermodynamic measurements. We focus on measurements of a range of quantum-statistical effects, which are obscured in harmonic traps, even when using only weakly interacting gases. First, in Section 6.1, we measure the critical number for a range of temperatures. In other terms, we are mapping out the boundary in the T - N phase diagram, which separates the condensed from the thermal gas. Comparing the experimental data to the prediction allows us to finalize the value of the effective trap power-law. In Section 6.2, we show that the thermal component in our trap is saturated, i.e. it follows ideal gas scalings even in the presence of a condensed component. Finally, in Section 6.3, we observe the quantum version of the Joule-Thomson effect. Unlike the interaction-driven classical equivalent, the quantum-mechanical version is purely statistical. The main results of this chapter have been published in [25].

6.1 Critical number

In the previous chapter we analyzed a rather arbitrary trajectory (see Fig. 6.1) in the T - N phase diagram to find α . This trajectory was determined by the evaporation process rather than by theoretical considerations. While offering experimental ease, we have seen that this causes two major theoretical headaches. Below T_c , we have to assume saturation, i.e. $N' = N_c$, and above T_c we rely on correctly measuring the chemical potential μ .

Here, we will instead measure the critical number N_c at different temperatures, solving both of these issues. In terms of the N - T phase diagram, we will trace the boundary between the non-condensed and condensed phase (black in Fig. 6.1). We expect that the critical number follows

$$N_c \propto T^\alpha. \tag{6.1}$$

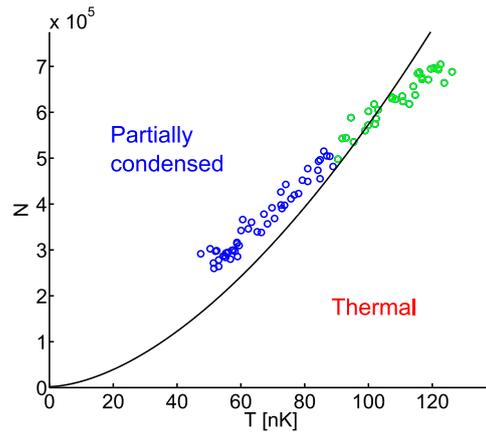


Figure 6.1: The evaporation trajectory from Chapter 5 is shown. The green points were used to extract α , while we discarded the blue points due to their high condensed fraction. The black line represents the phase boundary, which we trace and use for analysis in this chapter.

Assuming that this ideal gas scaling holds, we will determine the power-law parameter α , which fits the experimental data best. We will find that this method offers better accuracy than the one employed in the previous chapter.

6.1.1 Experimental procedure

To measure T and the number of thermal and condensed atoms, we image clouds after 50–70 ms of time-of-flight expansion from the trap. We vary T by changing the final trap depth of the evaporation sequence, while we vary the total atom number N by adjusting the MOT loading time. After forced evaporation we always raise the trap depth to $U_0 \sim k_B \times 0.4 \mu\text{K}$, where evaporation is negligible. Also, ramping back up to a fixed trap depth, i.e. the same trapping laser intensity, ensures the trap shape does not vary between measurements. After changing the depth, we hold the gas for another second before ToF to allow for thermalisation.

We determine the critical point by the appearance of small BECs. All images that have small, but unmistakable, condensates in the range $1000 < N_0 < 5000$ are selected and assumed to have a thermal number $N' \approx N_c$.¹ We then average the data points within small temperature bins (3 nK), and calculate the errors of the mean. The resulting smooth series of points can then be fitted with function Eq. (6.1) to obtain α as shown in Fig. 6.2.

¹For a saturated gas, this equation holds.

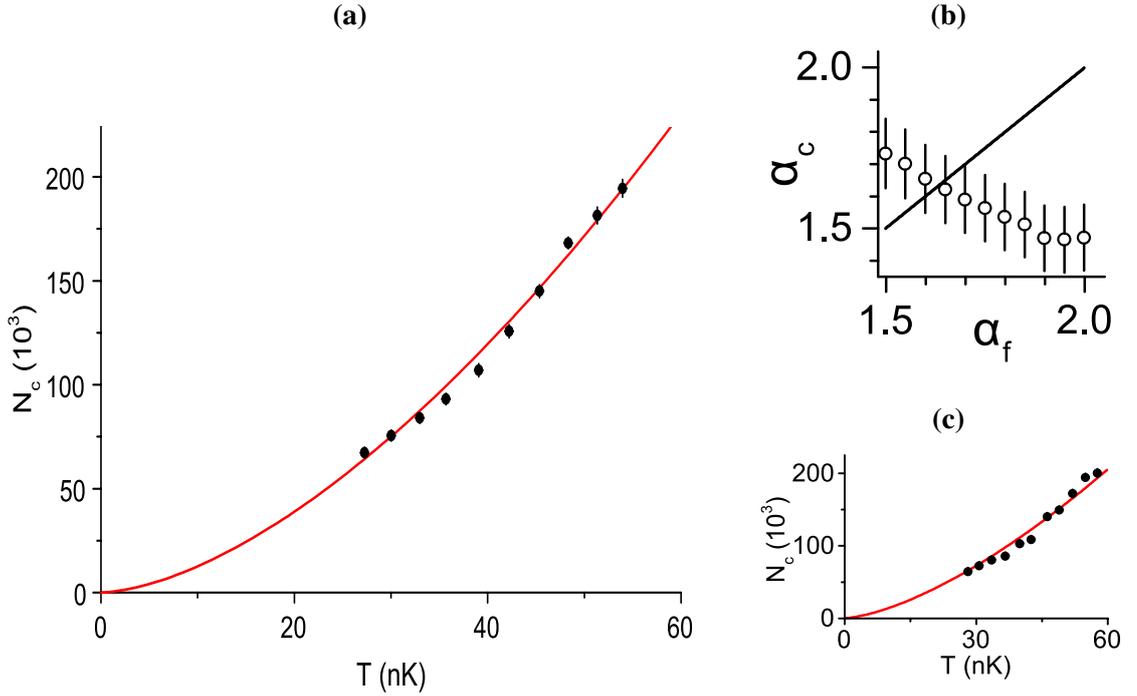


Figure 6.2: Critical point in a quasi-uniform box. (a) A fit to the critical number $N_c \propto T^\alpha$ gives $\alpha = 1.65$. (b) $\alpha_f = \alpha_c$ (see text) is fulfilled within errors for $1.6 \lesssim \alpha \lesssim 1.7$. (c) $\alpha = 1.5$ is still a good fit, illustrating the smallness of the deviation from perfect box.

6.1.2 Self-consistency

Unfortunately, this method has a slight complication: in order to extract N' , N_0 and T from an image, we need to fit the cloud column density with the appropriate fit function $\propto g_{\alpha-1}$. Obviously, this requires prior knowledge of α , which we were hoping to extract by fitting $N_c(T)$. To resolve this issue, we fit all images with a sensible range of α values $1.5 \leq \alpha_f \leq 2.0$. For each α_f value, we therefore obtain a different N_c - T plot similar to Fig. 6.2. To each of these plots, we fit $N_c \propto T^{\alpha_c}$, where α_c is a free parameter. For each assumed cloud profile α_f , we extract a corresponding α_c describing the phase boundary. For an ideal gas, α_c and α_f are identical. Therefore self-consistency requires $\alpha_c = \alpha_f$. For our trap, this requirement yields $1.6 \lesssim \alpha \lesssim 1.7$ (see Fig. 6.2b).

This value is consistent with the estimate in the Chapter 5, $1.5 \lesssim \alpha \lesssim 1.8$, but offers smaller error bounds. It should be noted that the trap was not exactly identical and we operated at lower temperatures, so we do not expect exact agreement. For the further data analysis we fix $\alpha_f = 1.65$, which corresponds to an effective power-law r^{20} .

While in the above we were careful to characterize the small deviation of our trap from a perfect box, we note that simply assuming a perfectly homogeneous gas of constant volume leads to very small errors. This is shown in Fig. 6.2c: we fix both α_f and α_c to

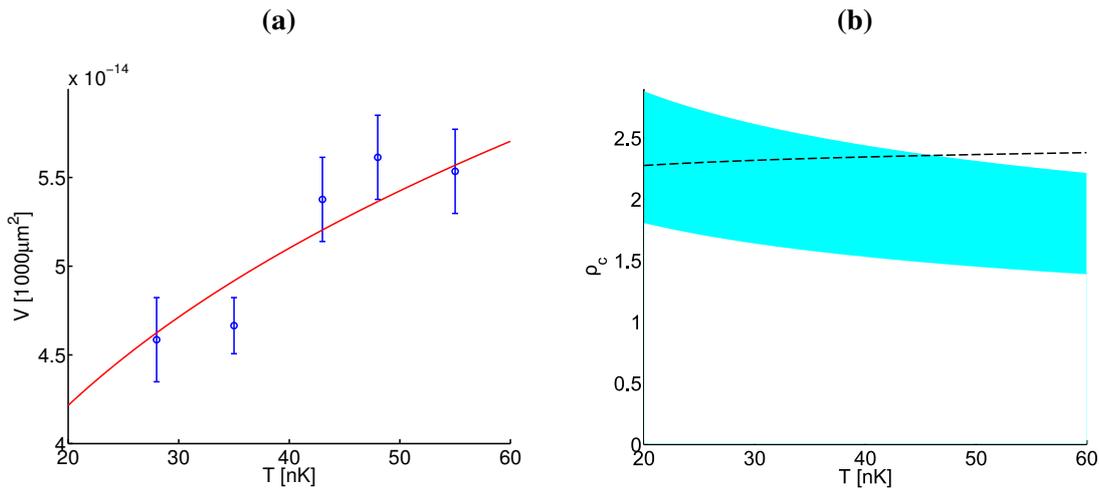


Figure 6.3: (a) Dependence of measured volume on temperature (blue). A power-law fit gives $V \propto T^{0.3 \pm 0.2}$ (red). (b) Comparison of the experimentally measured critical density (blue band) with the theoretical prediction of the critical density of an ideal box including the finite size effect (black dashed line). The blue band is calculated from the power-law fit to the measured volumes in (a) and the critical number measured in Fig. 6.2a. The width of the band represents the experimental systematic errors.

the ideal-box value $\alpha = 3/2$ and obtain a fit of $N_c(T)$, which is still quite good.

Method comparison It is interesting to note that this self-consistent method of finding α is conceptually using the same procedure as the one based on $\theta(T) = \text{const}$, which we used in Section 5.7.2:

- We measure the density distribution in ToF of a degenerate gas at different temperatures and numbers.
- We guess a range of values for the power-law parameter α and calculate the dependent thermodynamic quantities for each value of α .
- We then select the value for α which matches the ideal gas prediction $N'(T, \mu, \alpha)$ (Eq. (2.13)) best in a self-consistent manner. We introduce a single-parameter to quantify the deviation. In Chapter 5.7.2 we required $d\theta/dT = 0$. With the self-consistent method above, the corresponding criterion is $\alpha_f = \alpha_c$.

6.1.3 Dependence of volume on temperature

In the spirit of the local-density approximation, we can also attribute the deviation from the perfect uniform scaling, $N_c \propto VT^{3/2}$, to an increase of the box volume V with temperature T . Small deviations of α from $3/2$ correspond to a slight increase of the effective

trapping volume with temperature, $V_0 \propto T^\delta \propto T^{\alpha-3/2}$ (see Eq. (2.25)). From in situ images, we measure a $(20 \pm 10)\%$ increase in V_0 between $T = 25$ and 50 nK, (shown in Fig. 6.3a). A power-law fit $V(T) \propto T^{\alpha-3/2}$ gives $\alpha = 1.8 \pm 0.2$. While consistent with the previous results, it should be clear from the scatter and systematic errors that this measurement is supplementary at best.

6.1.4 Critical density

Over the whole temperature range in Fig. 6.2a, we can use the effective volumes $V_0(T)$ as interpolated in Fig. 6.3a to calculate the corresponding critical phase space density $\rho_c(T) = N_c(T)\Lambda^3/V_0(T)$. We obtain $\rho_c(T) \approx 2.0 \pm 0.2$ (blue band in Fig. 6.3b), which we compare to the theoretical prediction. In the thermodynamic-limit we expect $\rho_c \approx 2.612$. The inclusion of the finite-size correction as in Fig. 5.10 reduces this value $\rho_c \approx 2.4$ (black dashed in Fig. 6.3b). The remaining discrepancy can be explained within our systematic uncertainties of 10% in N_c and 20% in V_0 . Also, possible disorder, i.e. noise at the bottom of the potential could reduce the effective volume further. The influence of disorder remains to be studied.

6.2 Saturation of the thermal component

So far we have neglected the interactions in the gas. We have assumed that the thermal gas acts like an ideal, non-interacting gas with the total number capped at the critical number: $N' = N_c^{id}$. No interaction occurs with the condensed component, which holds the surplus atoms $N_0 = N - N_c^{id}$. In a harmonically trapped gas this picture is strongly violated even for weak interactions, such as in ^{87}Rb , and can be recovered only by extrapolating to the strictly non-interacting limit [16], where direct measurements cannot be performed due to absence of thermal equilibrium [74]. This strong deviation from ideal-gas behaviour arises due to an interplay of interactions and the non-uniformity of the gas, and can be explained using a mean-field theory that does not require violation of the saturation picture at the level of the local thermal density (LDA) [16,34]. We will now examine the situation experimentally and theoretically in a uniform gas.

6.2.1 Mean-Field theories

To first order, saturation effects are usually discussed in the framework of mean-field theories, instead of considering every single atom-atom interaction. We have already encountered the Gross-Pitaevskii equation, which addresses the influence of the mean-field gn on the ground state wavefunction at $T = 0$. Several frameworks exist to tackle

the influence of the mean-field for $T > 0$ onto the thermal and condensed distributions.

Hartree-Fock approximation

In the harmonic trap, experimental studies revealed a strong non-saturation, $N'(T \lesssim T_c) > N_c$. For temperatures $T \lesssim T_c$, this can be explained within the Hartree-Fock approximation [33]. In this framework, we account for the extra interaction energy between atoms in different energy levels, which arises from the exchange term of the bosonic wavefunction. In a partially condensed gas, only the ground state is macroscopically occupied. We can therefore approximate that adding a ground state atom increases the mean-field energy by $g(N_0 + 2N')$, while adding it in any other state increases the energy by $g(2N_0 + 2N')$. Incorporating the LDA, we obtain the following set of equations for the thermal distribution [28]:

$$\epsilon_{\mathbf{p}}(\mathbf{r}) = p^2/2m + 2n'(\mathbf{r})g + 2n_0(\mathbf{r})g + U(\mathbf{r}) \quad (6.2)$$

$$n'(\mathbf{r}) = \int \frac{1}{\exp[(\epsilon_{\mathbf{p}}(\mathbf{r}) - \mu)/k_B T] - 1} \frac{d\mathbf{p}}{h^3}. \quad (6.3)$$

Within the Thomas-Fermi approximation, the local condensate density is related to the chemical potential as:

$$\mu = U(\mathbf{r}) + [n_0(\mathbf{r}) + 2n'(\mathbf{r})]g. \quad (6.4)$$

Generally, these equations need to be solved self-consistently. Fortunately, in the case of a harmonic trap, we can apply several approximations, which simplify the problem to a point where we do not need numerical calculations. First of all, in harmonic traps, the condensate is concentrated in the centre of the trap and has typically densities $n_0(\mathbf{r}) \gg n'(\mathbf{r})$. We therefore neglect the mean-field of the thermal density, and only consider the repulsion originating from the condensed fraction. The condensed density is then described by a parabolic Thomas-Fermi profile

$$n_0(\mathbf{r}) = \max([\mu - U(\mathbf{r})]/g, 0). \quad (6.5)$$

We can thus approximate $(\epsilon_{\mathbf{p}}(\mathbf{r}) - \mu)$ in Eq. (6.3) as

$$\frac{p^2}{2m} + U_{eff}(\mathbf{r}), \quad (6.6)$$

where the effective potential $U_{eff}(\mathbf{r}) = |\mu - U(\mathbf{r})|$. Hence, the thermal atoms behave as if they were confined in a potential U_{eff} , which is shaped like a Mexican hat. Within the LDA and at a fixed temperature, this effectively increases the number of thermal states as the condensate grows. Therefore, when increasing N_0 , N' grows as well and thus

$N' > N_c$.

In first order, this effect can be approximated as [34]:

$$\frac{N'}{N_c} \approx 1 + 1.37 \frac{gn_0(0)}{k_B T}. \quad (6.7)$$

This equation has been confirmed experimentally for small condensates [16]. For larger condensates, including higher orders of $gn_0(0)/k_B T$ becomes necessary. Even when doing so, the experimental measurements deviate from the MF-predictions and show that the full scope of the problem cannot be captured within MF-theory.

We now apply the Hartree-Fock approximation to a uniform gas. The local densities are trivially related to the total number by:

$$n_0(r) = N_0/V \quad (6.8)$$

$$n'(r) = N'/V \quad (6.9)$$

Condensation occurs for $\mu_c = 2n'g \approx 0.05k_B T$ (for typical numbers) ; the critical number is not affected by this shift in μ . Below T_c , the chemical potential is given by

$$\mu = (n_0 + 2n')g = \mu_c + n_0g \quad (6.10)$$

and the density can be evaluated as:

$$n' = \int \frac{1}{\exp[(\mathbf{p}^2/2m + n_0g)/k_B T] - 1} \frac{d\mathbf{p}}{h^3} \quad (6.11)$$

$$= g_{3/2}(\exp[-n_0g/k_B T]) / \Lambda^3 \quad (6.12)$$

This expression is very similar to the thermal number of an ideal gas above T_c , where $\mu < 0$ in the polylogarithm is replaced by $-n_0g < 0$.

We see that within the HF-theory, we actually expect weak anti-saturation, i.e. a decrease of the thermal number with increasing condensed number. Unlike in harmonic trap, the deviation from the ideal gas is not driven by a change in effective volume, but can be regarded as a result of the modification of the effective dispersion relation by mean-field energies.

Popov approximation

We now consider the situation in a general power-law potential with exponent s by solving the local equations for the density numerically. Intuitively, we expect the saturation slope to vary somewhere in between the limiting cases $s = 2$ and $s = \infty$. As we will solve the

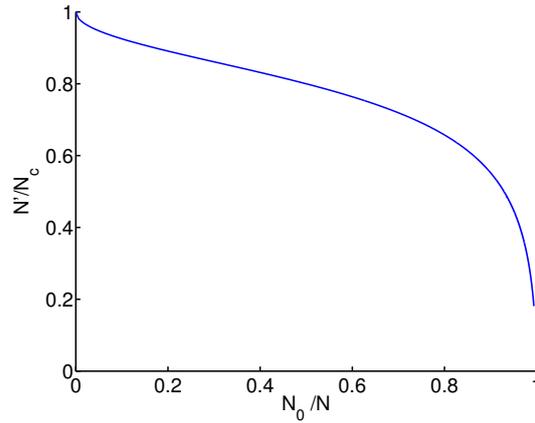


Figure 6.4: Hartree-Fock approximation for a uniform gas. The normalized thermal number is plotted against the condensed fraction.

problem numerically anyway, we will use the slightly more complex Popov approximation [28, 33, 75]. The numeric solutions of these equations have been considered in detail in [76] for the uniform and harmonic trap. For completeness we will quickly reproduce the methods and notation used for solving the uniform problem, before applying them to a range of power-laws. The equations for the uncondensed gas remain the same as in the Hartree-Fock approximation Eq. (6.3). In the presence of a condensate we use [28]

$$\mu = U(\mathbf{r}) + (n_0(\mathbf{r}) + 2n'(\mathbf{r}))g \quad (6.13a)$$

$$n'(\mathbf{r}) = \int \frac{1}{\exp[\epsilon(\mathbf{r}, \mathbf{p})/k_B T] - 1} \frac{\epsilon_a(\mathbf{r}, \mathbf{p})}{\epsilon(\mathbf{r}, \mathbf{p})} \frac{d\mathbf{p}}{h^3} \quad (6.13b)$$

$$\epsilon_a(\mathbf{r}, \mathbf{p}) = \frac{p^2}{2m} + gn_0(\mathbf{r}) \quad (6.13c)$$

$$\epsilon(\mathbf{r}, \mathbf{p}) = \left(\frac{p^2}{2m} \left[\frac{p^2}{2m} + gn_0(\mathbf{r}) \right] \right)^{1/2}. \quad (6.13d)$$

Like in the Hartree-Fock approximation, we take the exchange energies into account. Additionally, the dispersion spectrum Eq. (6.13d) is modified to match the $T = 0$ Bogoliubov dispersion spectrum [77]. This ensures we treat the low-energy phonon-like excitations correctly.

Within the LDA, we can solve the equations locally. For each value of μ , U and T we can calculate the densities n_0 and n' by solving Eqs. 6.13. We introduce the dimensionless variables

$$\alpha = \frac{\mu - U}{k_B T} \quad A = 4 \frac{a}{\Lambda} \quad G_0 = \frac{2gn_0}{k_B T} \quad G' = \frac{2gn'}{k_B T}. \quad (6.14)$$

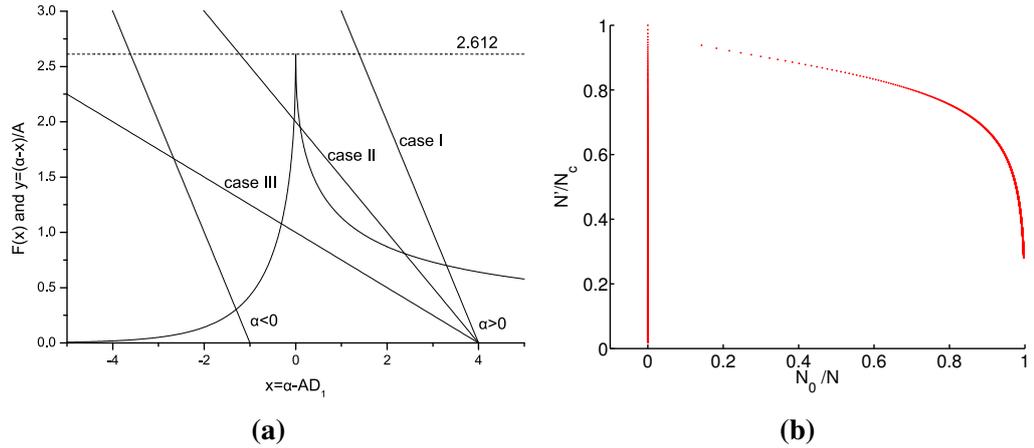


Figure 6.5: Numerical solution of the Popov equations in a uniform potential. (a) Graphical representation of Eq. (6.16) (figure from [76]). The value at the intersection of $F(x)$ and $\alpha - x/A$ gives the phase space density D' of the thermal fraction. For $x > 0$, condensation occurs and $D_0 > 0$. (b) The normalized thermal number N'/N_c is plotted against the condensed fraction for a uniform gas at $T = 45$ nK. The gap at low condensed fractions is caused by the ambiguity between multiple solutions in case II.

The phase-space densities are given by

$$D_i = G_i/A. \quad (6.15)$$

It can be shown that solving Eqs. 6.13 is equivalent to solving the following equation for $x = \alpha - AD'$:

$$\frac{\alpha - x}{A} = F(x), \quad (6.16)$$

where

$$F(x) = \left\{ \begin{array}{ll} g_{3/2}(e^x) & : x \leq 0 \\ \frac{2}{\sqrt{\pi}} \int_0^\infty \frac{[(x^2+Q^2)^{1/2}-x]^{1/2}}{e^{Q-1}} dQ & : x \geq 0 \end{array} \right\}. \quad (6.17)$$

From the value x which solves Eq. (6.16), we can deduce

$$D' = (\alpha - x)/A \quad (6.18)$$

and

$$\left\{ \begin{array}{ll} D_0 = 0 & : x \leq 0 \\ D_0 = 2Ax & : x \geq 0 \end{array} \right\}. \quad (6.19)$$

Fig. 6.5a illustrates the problem graphically. We find solutions by searching for the intersection between the left- and right-hand sides of Eq. (6.16). There are three different

kind of cases:

- Case I: One solution with $x > 0$, a BEC is present.
- Case II: Three solutions. We assume that the purely thermal solution at $x < 0$ is correct.
- Case III: One purely thermal solution with $x < 0$.

Case II is somewhat problematic. We are not aware of any justification for choosing the purely thermal solution. This choice ensures that the transition in n_i from Case III to Case II is continuous. The transition from Case II to Case I when $\alpha/A > \zeta(3/2) \approx 2.612$ introduces a discontinuity in the solution for x and thus n_i .

Ignoring this issue, we see that the solutions for thermal and condensed density are only a function of α and A . In terms of common parameters, it depends on chemical potential, local potential and temperature. For a given temperature, we can now vary μ to obtain the self-consistent solutions for a uniform potential with $U = 0$ as shown in Fig. 6.5b.

With the uniform solutions $n_i(\mu/k_B T)$ at hand, we can calculate the solutions for arbitrary potentials within the LDA. We sample the potential $U(r)$ in small enough volume elements dV to obtain the distribution $f(U)$. Each of these elements contributes a density $n_i(\mu - U(r)/k_B T)$. The total atom numbers N_0 and N' for a given chemical potential are calculated as

$$N_i(\mu) = \sum f(U) n_i(\mu - U, T) dV, \quad (6.20)$$

where we can truncate the sum for marginally populated states $U \gtrsim 2k_B T$. The implementation of this is straightforward; convolution of the arrays containing the tabulated values for n_i and f gives N_i , while circular shifting of n_i allows us to vary μ . In Fig. 6.6 the results for various power-laws are compared. We see that increasing the power beyond $s = 2$ quickly reduces the saturation slope. For $s = 8$, corresponding to $\alpha \approx 1.9$, it is approximately flat. Increasing the power-law further, causes the slope to slowly converge towards the uniform anti-saturation. For our trap, at $s = 20$ or $\alpha = 1.65$, we therefore do not expect major deviations from ideal gas behaviour at low to intermediate condensed fractions.

The influence of random variations of the potential in the trap centre can be treated in a similar manner. The saturation slope is affected in a similar way as when lowering the power-law. Within the LDA, only the total number of volume elements $f(U)$ at a given potential matters; it is not relevant how they are distributed spatially. We should keep in mind that mean-field theories do not always capture the full picture and beyond-mean-field results can be important. In the harmonic trap, deviations from the HF (and Popov)

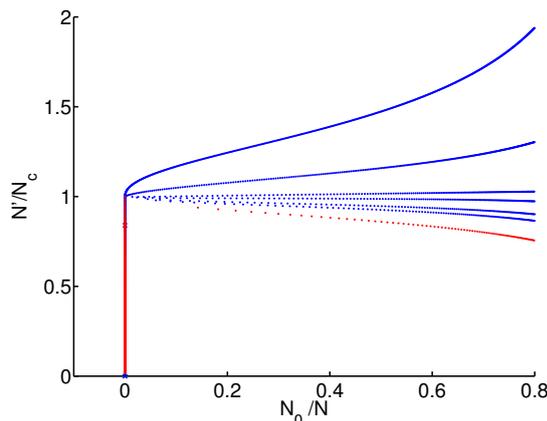


Figure 6.6: Numerical calculations for isotropic three-dimensional power-law traps $U(r) = (k_B T)r^s$. The condensed fraction is plotted versus thermal number normalized to the critical number. The blue curves show, top to bottom, $s = (2, 4, 8, 10, 15, 20)$, while the red curve shows the result for a uniform potential.

approximations at larger condensed fraction have been confirmed experimentally [16,34], while beyond-mean-field shifts of the critical point have been observed [10].

Our theoretical considerations show that a uniform gas is not expected to be fully saturated within mean-field theory. However, the deviations are expected to be minor under typical experimental conditions.

6.2.2 Experimental Results

To study saturation experimentally, we prepared a condensed gas at $T_0 = 45 \pm 1$ nK with a varying total number. We found that the MOT loading time provides good control over the total number at the end of the sequence, while the evaporation depth in the green trap sets the final temperature. In Fig. 6.7a, we plot the condensed fraction N_0/N against the normalized thermal number N'/N_c . We see that the gas is saturated within error.

The data shown is compensated for small temperature variations to first order by correcting the thermal number according to ideal gas scalings:

$$N'_{corr} = N'(T_0/T)^\alpha. \quad (6.21)$$

The data is then averaged within bins of size $\Delta N_0/N = 0.04$. The x-y errorbars are the statistical error of the mean. In Fig. 6.7b, we evaluate the same data at the edge of our confidence interval for α . We see that the variation in the thermal number is small for both values.

For comparison, we also show analogous measurements for the harmonic trap from which our box trap is loaded. Here the non-saturation of N' is prominent and agrees with the prediction from [16, 34] (Eq. (4.13)).

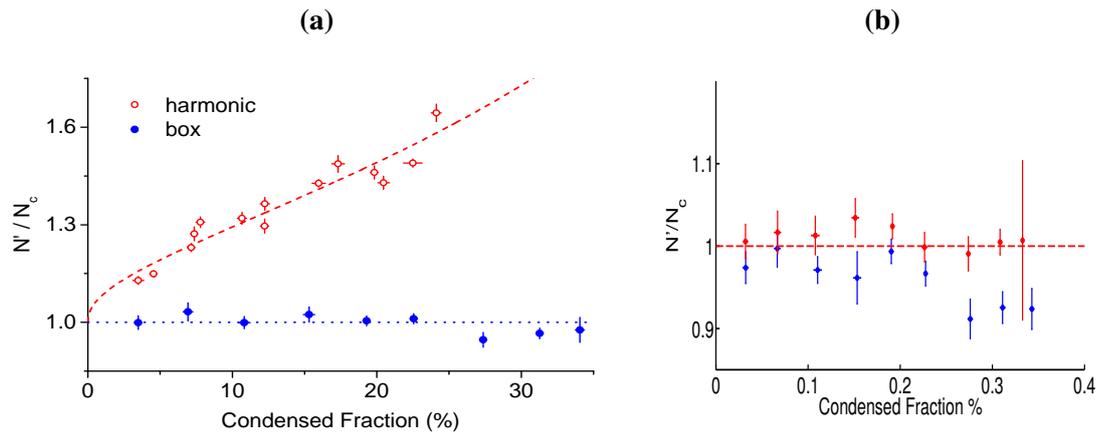


Figure 6.7: (a) Saturation of the thermal component in a partially condensed gas. In the box trap the gas follows the ideal-gas prediction $N' = N_c$, whereas in the harmonic trap the thermal component is strongly nonsaturated. The dashed red line shows the theoretical prediction for the harmonic trap. (For the harmonic trap data, $T \approx 110$ nK and $N_c \approx 65000$). Figure taken from [25]. (b) Variation of the saturation in the uniform trap with α . The same data as in (a) evaluated with $\alpha = 1.6$ (blue) and $\alpha = 1.7$ (red).

6.2.3 Discussion

We have found experimentally that the thermal component in our quasi-uniform trap is saturated. Within experimental uncertainties, the gas follows Einstein’s textbook picture

$$N' = N_c \propto T^\alpha \quad (6.22)$$

for $N > N_c$. This is consistent with theoretical calculations for high, but not infinite power-laws. Further experimental and theoretical studies might be fruitful and could shed light on the behaviour at higher condensed fractions and tunable interactions.

Eq. (6.22) also explains why the inclusion of condensed clouds in the power-law analysis in [24] did not distort the results for α .

6.3 Quantum Joule-Thomson effect

During our saturation measurements, we observed unexpected cooling of condensed gases while holding them in the uniform trap. Since atom numbers were dropping as well, we first explained this by evaporation. However, this explanation did not hold, as we could not stop the cooling by increasing the depth of the trap. Also, the effect was severely reduced or even absent for thermal gases.

The cooling can instead be explained with the quantum Joule-Thomson (QJT) effect.

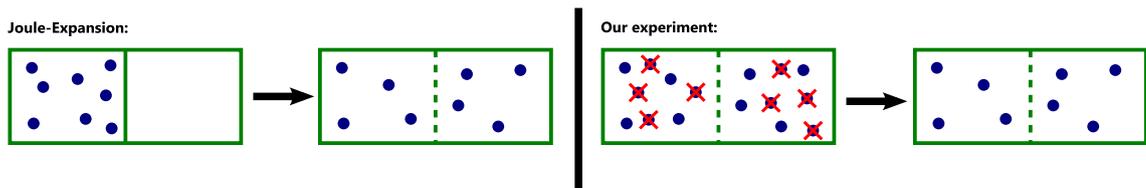


Figure 6.8: Isoenthalpic rarefaction. In a conventional Joule expansion process the volume V is increased, whereas in our experiments the atom number N is reduced. In an ideal gas both processes conserve the specific enthalpy and are thermodynamically equivalent.

Collisions with high-energy background gas atoms remove atoms slowly from the trap. Crucially, collisions with the background gas are independent of the energy of the trapped atoms, so the average energy per trapped particle, E/N , remains constant.

The pressure of an ideal gas is given - irrespective of statistics- by [28]:

$$p = \frac{2}{3} \frac{E}{V} \quad (6.23)$$

The Enthalpy H is then:

$$H = E + pV = \frac{5}{3} E \quad (6.24)$$

Keeping E/N constant throughout the process means that the specific enthalpy $h = H/N$ is conserved², which is the definition of a Joule-Thomson process.

The Joule-Thomson effect causes the cooling or heating of a gas when changing the pressure in an isoenthalpic manner. Its magnitude and sign are quantified by the Joule-Thomson coefficient [78]

$$\mu_{JT} = \left(\frac{\partial T}{\partial p} \right)_h. \quad (6.25)$$

The JT process is exploited on an industrial scale for cooling and liquifaction of cryogenic gases, by repeatedly throttling the gas through a valve. On an everyday level, it causes compressed air to cool when leaving a pressured container.

Unlike in those cases, the effect we observe is not driven by a change in volume, but a change in number. As all relevant thermodynamic properties in Eq. (6.25) are intensive (i.e. n , p , T and h), the processes are thermodynamically equivalent nevertheless. We illustrate this equivalence in Fig. 6.8. In the left panel we illustrate free Joule expansion into vacuum. In the right panel we show the loss-driven rarefaction happening in our experiment.

Before describing our experimental measurements and analysis of the QJT, we are going to derive the theoretical framework to describe the Joule-Thomson effect:

- In a classical (non-ideal) gas, $T \gg T_c$ where the Joule-Thomson process is only

²Sometimes the specific Enthalpy is defined as H/m .

driven by interactions.

- In an ideal, condensed gas, $T < T_c$. Here, the Joule-Thomson process occurs due to quantum-statistical effects.
- In an ideal, thermal but degenerate gas, $T \gtrsim T_c$. Even without a BEC, quantum statistics predicts a QJT effect.

6.3.1 Joule Thomson effect in a classical gas

Here, we will quickly recapitulate the classical Joule-Thomson effect. First, let us consider an ideal gas without interactions. The equation of state is given by

$$pV = Nk_B T. \quad (6.26)$$

It is easy to show that

$$T = \frac{5}{2} \frac{h}{k_B}, \quad (6.27)$$

and therefore

$$\mu_{JT} = \left(\frac{\partial T}{\partial p} \right)_h = 0. \quad (6.28)$$

There is no Joule-Thomson effect in a non-interacting (i.e. ideal) classical gas. Expansion does not change the temperature, as each atom is not aware of the presence of the other atoms.

The Joule-Thomson effect in real gases is driven by interactions between the atoms. In the case of attractive interactions, expanding the gas requires energy to overcome the attractive forces - the gas cools.

The simplest model for a gas with interactions is the Van-der-Waals gas:

$$\left(p + \frac{a}{v^2} \right) (v - b) = RT. \quad (6.29)$$

Here, we introduced the specific Volume $v = VN_a/N$ and the constants a and b . a is a measure of the attraction between particles, whereas b relates to the excluded volume per particle. The gas constant R is related to the Avogadro constant N_a by $R = N_a k_B$. The Joule-Thomson coefficient can be expressed as [78]

$$\mu_{JT} = \frac{1}{C_{p,m}} \left[T \left(\frac{\partial v}{\partial T} \right)_p - v \right], \quad (6.30)$$

where $C_{p,m} \approx (5/2)R$ is the molar heat capacity. Using Eq. (6.29), the Joule-Thomson

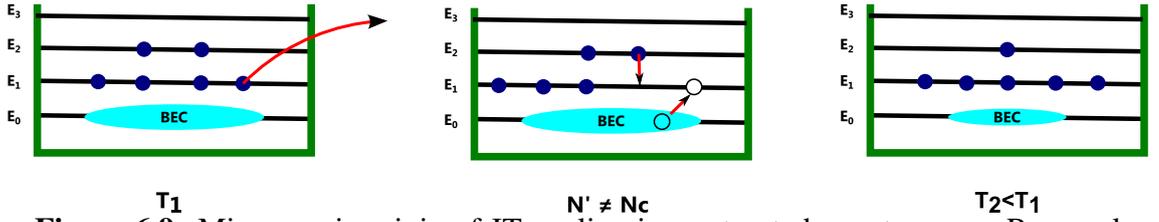


Figure 6.9: Microscopic origin of JT cooling in a saturated quantum gas. Removal of a thermal atom requires a zero-energy particle to come out of the BEC in order to maintain $N' = N_c$, while energy conservation requires redistribution of atoms between energy levels. Cooling is seen in the change of the relative populations of the excited states. In this cartoon we neglect the fact that N_c also slowly decreases with T .

coefficient can now be evaluated as

$$\mu_{JT} = \frac{1}{C_{p,m}} \left[\frac{(2a/RT)(1 - b/v)^2 - b}{1 - (2a/vRT)(1 - b/v)^2} \right]. \quad (6.31)$$

This expression is not very intuitive. For comparative purposes, we are mainly interested in the magnitude of a positive μ_{JT} , which is appropriate for cooling gases. We approximate $C_{p,m} \approx (5/2)R$ and assume $v \gg b$, which allows us to estimate the value of the Joule-Thomson coefficient as

$$\mu_{JT} \approx \frac{a}{R^2 T}. \quad (6.32)$$

a is typically on the order of $< 0.5 \text{ Pa m}^6 / \text{mol}^2$. At room temperature, we expect a Joule-Thomson coefficient on the order of $< 3 \text{ K/bar}$. Typically, experimental values are smaller than this, even at lower temperatures. One should also note that while the Van-der-Waals model reproduces general trends and behaviour, it is of limited usefulness in practice.

6.3.2 Joule-Thomson effect in a condensed quantum gas

We have seen that in a classical gas the Joule-Thomson effect is caused by interactions. This is very different in a quantum gas, where the Joule-Thomson effect occurs without interactions, purely due to quantum statistics. This effect has been predicted as early as 1937 [79].³

The cooling effect can be derived quickly without having to calculate the Joule-Thomson coefficient. Neglecting the interaction energy, the energy can be written as

$$E \propto N' k_B T. \quad (6.33)$$

³The original paper considers the magnitude of the QJT in ^4He at 5 K. They estimate that the QJT coefficient is on order 0.076 K/bar, while the expected Van-der-Waals JT coefficient is $\approx 0.7 \text{ K/bar}$. This would make it difficult, but probably not impossible to measure the QJT coefficient in Helium.

6.3. QUANTUM JOULE-THOMSON EFFECT

We have shown earlier that the gas is saturated, i.e. $N' = N_c$. Also, we have shown that the critical number scales with temperature as $N_c \propto T^\alpha$. We can therefore write

$$E \propto N_c T \propto T^{\alpha+1}. \quad (6.34)$$

As the loss process conserves E/N , the atom number N and temperature T can be related as

$$T \propto N^{1/(\alpha+1)}. \quad (6.35)$$

The microscopic origin of this cooling is qualitatively illustrated in Fig. 6.9. Removal of any thermal atom requires a zero-energy particle to come out of the BEC in order to maintain the saturation of the thermal component. Energy conservation then requires simultaneous redistribution of atoms between energy levels. The net cooling is seen in the change of the relative populations of the excited states. Note that direct removal of BEC atoms does not change the temperature, and also that in this simplified cartoon we neglect the fact that $N_c(T)$ gradually decreases as T drops.

In a uniform gas, it is simple to evaluate the coefficient μ_{JT} ⁴. We rewrite Eq. (6.25) as

$$\mu_{JT} = \left(\frac{\partial N}{\partial T} \right)_h^{-1} \left(\frac{\partial N}{\partial p} \right)_h. \quad (6.36)$$

Expressing Eq. (6.23) as

$$N = \frac{5}{2} \frac{pV}{h} \quad (6.37)$$

allows us to calculate

$$\left(\frac{\partial N}{\partial p} \right)_h = \frac{3}{2} \frac{VN}{E}. \quad (6.38)$$

Now, we write Eq. (6.34) as $E = B k_B T N_c(T)$, where B absorbs constants. Rearranging yields

$$N = \frac{5}{3} \frac{B k_B T}{h} N_c(T), \quad (6.39)$$

and by differentiation we obtain

$$\left(\frac{\partial N}{\partial T} \right)_h = \left(\frac{N}{T} \right) B k_B \left(\frac{5}{2} \right) N_c(T). \quad (6.40)$$

Combining the two derivatives in Eq. (6.36) allows us to calculate the Joule-Thomson

⁴For a calculation in a power-law potential, the volume V is not defined. Therefore we need to redefine the pressure $p = T \left(\frac{\partial S}{\partial V} \right)_{E,N}$ by replacing V with the appropriate trap-shape parameter (e.g. the trap frequency in a harmonic trap). The resulting Joule-Thomson coefficient is not expressed in Kelvin/bar and is not very intuitive.

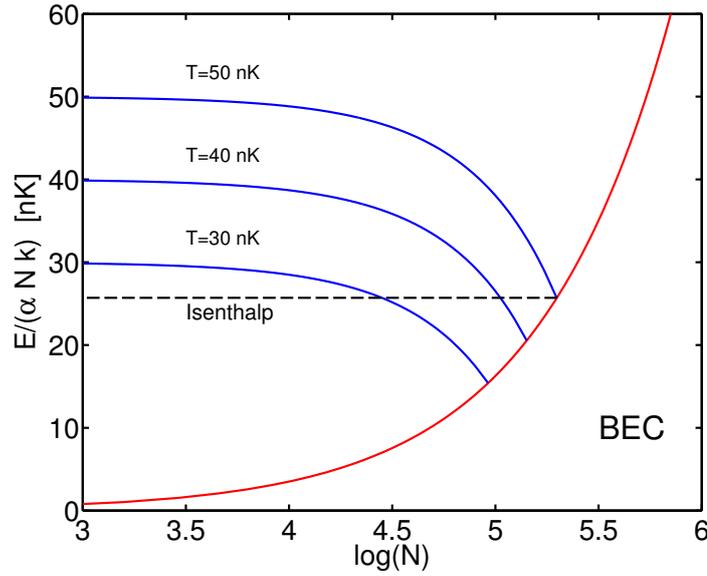


Figure 6.10: The quantum Joule-Thomson effect in a degenerate gas $T > T_c$ in a uniform trap ($\alpha = 1.5$). As $h \propto E/N$, isenthalps are horizontal lines. A range of isotherms is depicted in blue. In this example, the isenthalp starts at the critical point with 50 nK. Diluting the gas by two orders of magnitude, e.g. by loss-driven Joule-Thomson rarefaction, cools it down to $T \approx 26$ nK.

coefficient

$$\mu_{JT} = \frac{2}{5} \frac{\Lambda^3}{\zeta(5/2) k_B}. \quad (6.41)$$

It is interesting to note that in Fermi gases the loss-driven QJT is predicted to cause heating rather than cooling [80].

6.3.3 QJT in a degenerate thermal gas

The quantum Joule-Thomson effect does not only occur in condensed gases at $T < T_c$. It also affects degenerate, thermal gases with $T \gtrsim T_c$. This is actually the effect predicted in [79].

The energy in a degenerate gas can be evaluated (using Eq. (A.11)) from the density of states $D(\epsilon)$ in Eq. (2.14) and the Bose-Einstein distribution with $\mu < 0$:

$$E = \int \epsilon D(\epsilon) f(\epsilon, \mu) d\epsilon \quad (6.42)$$

$$= C_\alpha \Gamma(\alpha + 1) g_{\alpha+1}(\exp[-\mu/k_B T]) (k_B T)^{\alpha+1}. \quad (6.43)$$

In the classical limit, this reduces to [28]

$$E = \alpha N k_B T, \quad (6.44)$$

while at the critical point with $\mu = 0$ we have

$$E = \alpha N k_B T \zeta(\alpha + 1) / \zeta(\alpha). \quad (6.45)$$

The maximum reduction of temperature under isenthalpic conditions, i.e. when conserving E/N , is given by the ratio of the classical and the critical energies, $\zeta(\alpha + 1) / \zeta(\alpha) \approx 51\%$ in a uniform gas. Qualitatively, degeneracy enhances occupation of low energy states, with $\epsilon < |\mu|$, so E/N is lower than the classical value $\alpha k_B T$. In a harmonic gas the ratio is $\zeta(4) / \zeta(3) \approx 0.9$, which would make the QJT effect very hard to observe.

To illustrate the variation of energy per particle according to Eq. (6.43), we plot $E / \alpha N k_B$ against the atom number in Fig. 6.10. We display this variation for a range of temperatures. These isotherms curve smoothly from the classical value in Eq. (6.44) to the critical value in Eq. (6.45). In this diagram, isenthalps with $h = \text{const}$ are simply horizontal lines. The QJT loss-cooling emerges now naturally by following an isenthalp in the direction of lower atom numbers, which intersects the isotherms. The QJT effect persists over a large range of numbers; in the example shown we would have to dilute the gas by an order of magnitude to reduce the cooling due to QJT notably.

6.3.4 Experimental observation of the QJT

To study the QJT, we prepared clouds at $T_i \approx 45 \text{ nK}$ in a trap of depth U_0 , where evaporation is negligible. We constantly lose atoms due to background gas collisions on an exponential timescale $\tau \approx 10 \text{ s}$. Due to low atomic density in our uniform trap ($< 5 \times 10^{12} \text{ cm}^{-3}$), the three-body recombination rate is negligible [65]. As N decays, the elastic collision rate (among the trapped atoms) remains sufficiently high for the gas to continuously re-equilibrate on a timescale $\tau_{\text{eq}} \lesssim 2 \text{ s} \ll \tau$ [43, 81]. Using our previously established thermometry methods, we observed the atom number and the temperature of the trapped cloud and monitored their variation over time. In Fig. 6.11 we show the evolution of T with decaying N , with both quantities scaled to their initial values. We show data for a partially condensed gas, with an initial atom number $N_i \approx 1.7 N_c(T_i)$, and a thermal cloud, with $N_i \approx 0.6 N_c(T_i)$. Note that in both cases we measure down to approximately the same final N value, $\approx 0.3 N_c(T_i)$. At this point N_0 in the initially condensed sample vanishes, and our temperature fits become less reliable.

In the (partially) condensed sample we observe a drop in T by a factor of two, i.e. by $\approx 22 \text{ nK}$. Meanwhile, the interaction energy per particle is always smaller than $(8\pi\hbar^2 a/m)N_i/V$ [28], corresponding to $< 4 \text{ nK}$. The change in T therefore must predominantly be a quantum-statistical, rather than an interaction effect. We estimate that the total change in T due to the change in the interaction energy per particle could at most

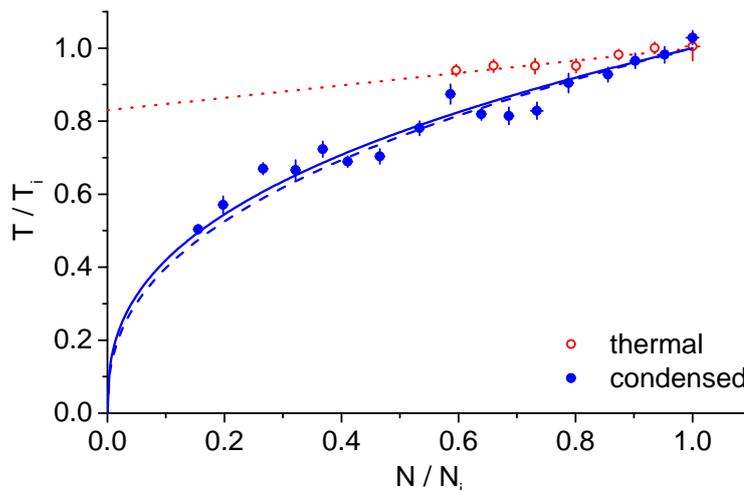


Figure 6.11: Measurements with a partially condensed and a thermal cloud. Solid and dashed blue line are predictions of Eq. (6.35) with $\alpha = 1.65$ and $\alpha = 3/2$, respectively. The dotted red line is a numerical calculation based on the equation of state for the degenerate, thermal atoms.

be ~ 3 nK, which is comparable to our experimental data scatter. Moreover, note that for $a > 0$ the drop in the interaction energy should lead to slight *heating*, rather than cooling of the gas.

Indeed, the scaling of the QJT in a condensed gas (Eq. (6.35)) fits the data very well. We show predictions for $\alpha = 1.65$ (solid blue line) and $\alpha = 3/2$ (dashed blue line). The two are almost indistinguishable, reaffirming that the behaviour of our clouds is very close to that of a perfectly homogeneous gas. We also notice a slight cooling in the thermal gas (red points), which is consistent with a numerical calculation (dotted red line) based on the QJT in a degenerate thermal gas (Eq. (6.43)). Unfortunately, we cannot explore a thermal gas in the classical regime, where no QJT should be present. Diluting the gas further reduces the signal level to a point where the determination of atom numbers and temperatures is unreliable. Increasing the temperature is also not possible, as it is limited by the maximum depth of our uniform trap.

We can now evaluate the JT coefficient μ_{JT} using Eq. (6.41). At our lowest temperature, we observe $\mu_{JT} \approx 4 \times 10^9$ K/bar, which is about ten order of magnitude higher than in classical gases.

6.3.5 Experimental details

For completeness, we describe some additional experimental details and analysis, which were not included in our original paper [25]. We focus on two topics: First, we detail the preparation of the gas by quenching non-adiabatically through the transition. The dynamics of condensate formation after quenching the atoms in a such a manner are very

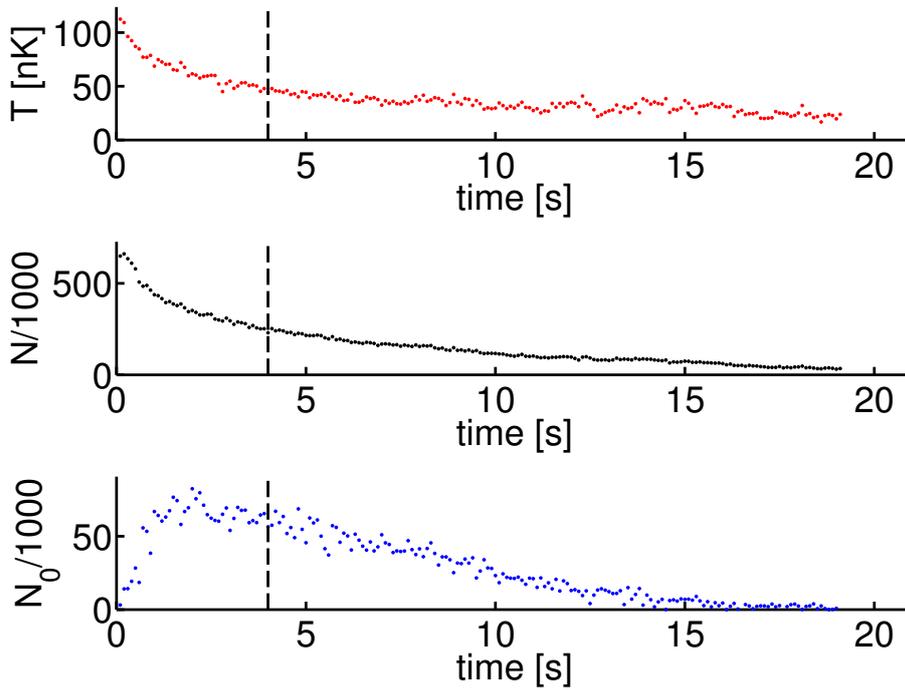


Figure 6.12: Temporal evolution of a gas in a uniform trap after lowering the trap depth abruptly. At first the gas rethermalizes in the condensed state; a condensate forms. After around 2s the condensate has stopped growing. We identify the cooling after 4s (dashed line) as driven by the quantum Joule-Thomson effect.

interesting and are currently investigated in follow-up experiments. Second, we discuss the scaling of the condensed atom number N_0 throughout the QJT process, as opposed to the total number N we have been considering so far.

Preparation of the cloud

We prepared the gas by lowering the depth of the uniform trap non-adiabatically to $U_0 = k_B \times 0.4 \mu\text{K}$, rather than by gentle lowering of the depth. This truncates the high-energy end of the thermal distribution. The following rethermalization cools the gas below its critical temperature and condensation occurs. We show the evolution of the gas in Fig. 6.12. Until 2s after the quench, the condensate grows rapidly. Until then, the gas is not in equilibrium and evaporation-driven cooling is still ongoing. This is consistent with our previous estimate of the collision rate. At this point, the temperature is around $45 \text{ nK} \approx U_0/(9k_B)$, and it is reasonable to assume that evaporation is negligible. For good measure, we waited another 2s before considering data for Fig. 6.11.

Condensed atom number during the QJT

We have found in Fig. 6.11, that the total number and temperature follow the prediction in Eq. (6.35). Combining Eq. (6.22) and Eq. (6.35) allows us also to predict the point where the condensate vanishes. We define the initial total number N_i and initial thermal number N'_i . We define N_t as the number where the condensed fraction vanishes so that $N_t = N'_t$. We obtain

$$\frac{N_t}{N_i} = \left(\frac{N'_i}{N_i} \right)^{\alpha+1}, \quad (6.46)$$

which relates the initial thermal fraction N'_i/N_i to the ratio N_t/N_i at which the condensate disappears.

Our experimental data shows that the condensate persists down to $N_t/N_i \approx 0.25$, which is significantly longer than $N_t/N_i \approx 0.45$, as suggested by the initial thermal fraction $N'_i/N_i \approx 0.74$. This discrepancy can be explained because we underestimated the initial condensate number N_0^i . This undercounting is caused by “blacking out” in the centre of the image for high optical densities, i.e. the cloud is so dense that barely any light can reach the camera. This is illustrated in Fig. 6.13b. While this reduces the measured condensate size, it does not affect the temperature reading, which is deduced from the low-density wings of the thermal distribution. Also, the effect on the total number N' is rather small.

In order to test this hypothesis, we rewrite Eq. (6.46), to predict the condensate size N_0^{theo} for all atom numbers N , based on the atom number N_t where the condensate disappears :

$$N_0^{theo} = N_t - (NN_t^\alpha)^{1/(\alpha+1)}. \quad (6.47)$$

In Fig. 6.13a, we compare this prediction N_0^{theo} with the measured number N_0^{exp} . We find that the prediction is correct for small numbers, but deviates from the measured number N_0^{exp} for $N > 130\,000$. This coincides for the point where the peak optical densities exceed the trusted range, $OD \gtrsim 4$, and the image blacks out. To check that this does not alter the results of Fig. 6.11, we correct the total number for the difference in measured and predicted condensate size

$$N_{corr} = N + (N_0^{theo} - N_0^{exp}) \quad (6.48)$$

and replot Fig. 6.11. The corrected data points are shown in Fig. 6.13c. We see that the corrected data fits the scaling Eq. (6.35) at least as good as original data. It even seems to suggest it is possible to extend the range of data where QJT scaling applies, indicating that the gas is in equilibrium after 3 s rather than after 4 s (not shown in Fig. 6.13c).

There are also a couple of data points which do not have a condensate. Strictly speak-

ing, the condensed QJT scaling does not apply to them; however the effect of the thermal QJT near N_c has a very similar scaling, which allows us to plot them on the same curve [27].

6.3.6 Conclusions

In summary, we have experimentally confirmed the quantum Joule-Thomson effect in a degenerate gas. Unlike in classical gases, interactions are not the driving force of the quantum manifestation of the Joule-Thomson effect. We are able to extract a coefficient $\mu_{JT} \approx 4 \times 10^9$ K/bar, which is ten orders of magnitude larger than observed in classical gases. It is also interesting to note that μ_{JT} in Eq. (6.41) explicitly vanishes in the classical limit, $\hbar \rightarrow 0$, reiterating that this is a purely quantum effect.

Similar theoretical analysis applies to an ideal harmonically trapped gas, but in that case $\alpha = 3$, implying weaker cooling in both condensed clouds and degenerate thermal samples. Moreover, in practice harmonically trapped gases strongly deviate from this picture, because N' is not saturated, much higher typical BEC densities enhance the role of interactions, and three-body recombination continuously increases E/N .

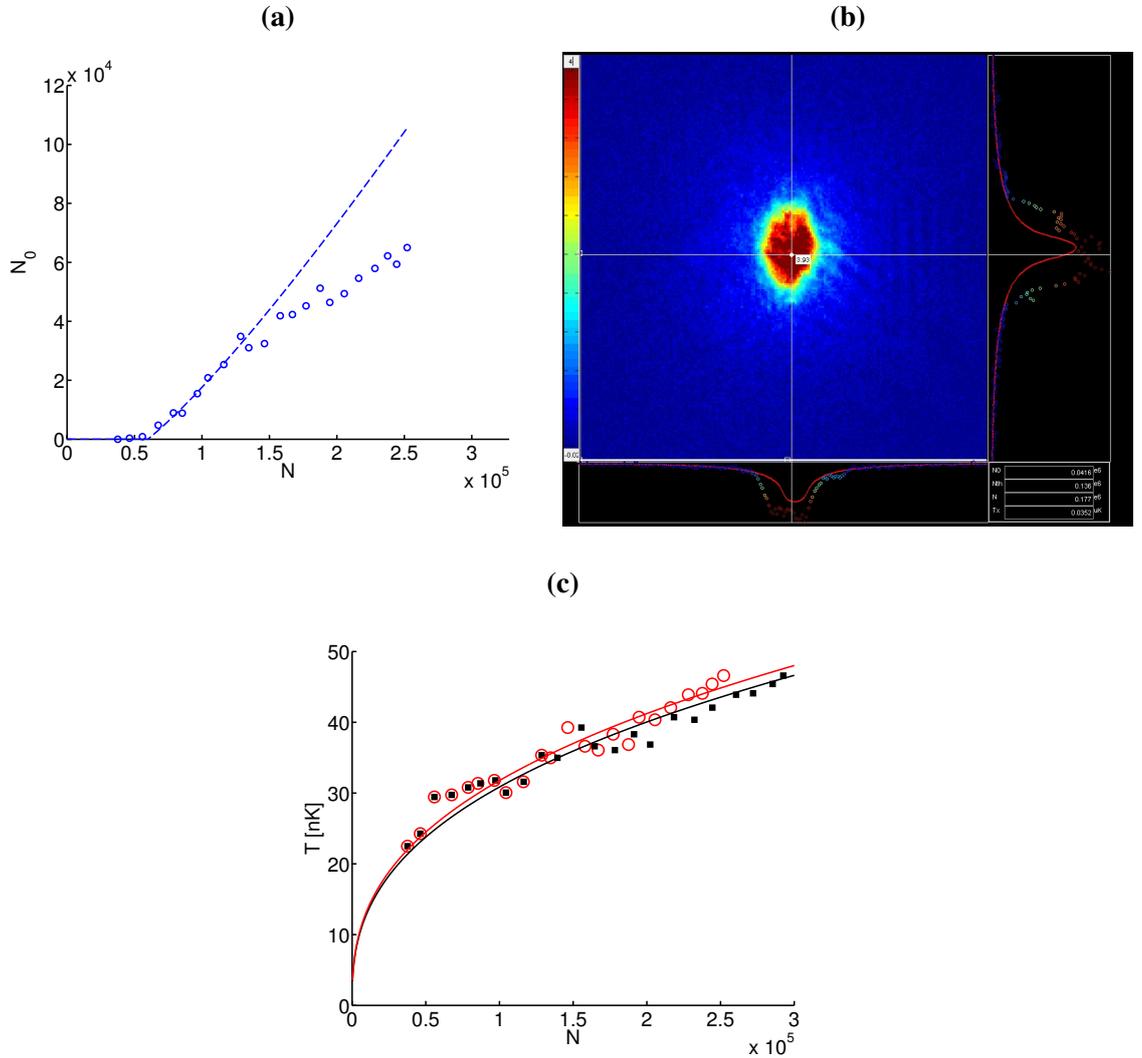


Figure 6.13: (a) Comparison of the expected and measured condensate numbers. The dashed blue line shows the prediction N_0^{theo} given by Eq. (6.47). The blue circles show the measured condensate number N_0^{exp} . (b) An image that is “blacked-out” in the centre. For optical densities $OD > 4$, the condensate number N_0 is underestimated by N_0^{exp} . This occurs at numbers $N > 130\,000$, which is consistent with (a). (c) We correct the total number N for the difference $N_0^{theo} - N_0^{exp}$ observed in (a). We compare the original data (red) with the corrected data (black). The difference is small for most data points. The continuous lines show the respective best fit $N \propto T^{-(1.65+1)}$. The corrected data still agrees with this prediction. For clarity, error bars were omitted.

Chapter 7

Bragg spectroscopy

Our studies of degenerate Bose gases in a box potential have so far focused on the critical point for condensation and the thermodynamics of the gas close to the critical temperature. These previous experiments have established the uniformity of the *thermal gas* from which the BEC forms. In this chapter, we probe a box-trapped Bose gas in the low-temperature regime of quasi-pure *condensates*. These measurements require a precision which can resolve a two orders of magnitude lower energy scale.

We study the coherence and energy of the gas, as well as its free expansion from the box trap, employing two-photon Bragg spectroscopy [82–89]. This method allows us to obtain high resolution measurements of the momentum distribution and interaction energy.

In Section 7.1, we establish the theoretical foundations of Bragg diffraction; see also [83, 86, 90] for more complete treatments. In the experimental Section 7.2, we describe the experimental setup, followed by a description of the measurements we conducted to characterize the condensate. First, for a wide range of box sizes (up to 70 μm), we directly observe the Heisenberg-limited momentum uncertainty of the condensed atoms, corresponding to a fully coherent macroscopic BEC wavefunction spanning the whole box trap. From the interaction shift of the Bragg resonance we deduce the BEC ground-state energy (per atom) with a precision of $k_B \times 100$ pK, and find good agreement with the mean-field theory for a perfectly uniform condensate.

Finally, we study the free time-of-flight (ToF) expansion of a BEC from the box trap. We follow the evolution of the cloud shape and the gradual conversion of the interaction energy into the width of the momentum distribution, and reproduce our observations in numerical simulations based on the Gross-Pitaevskii equation.

The results of this chapter have been published in [26].

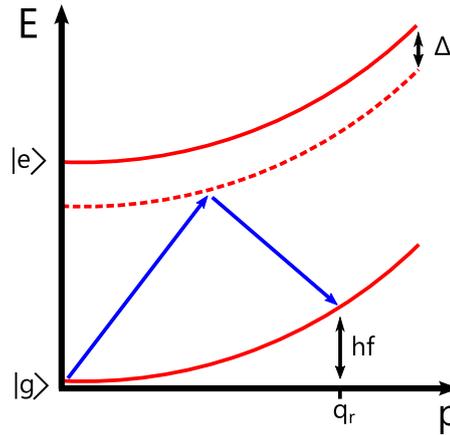


Figure 7.1: Bragg scattering in a two level system (not to scale). The energy of the ground state $|g\rangle$ and the excited state $|e\rangle$ vary with momentum p according to the dispersion relation. The single photons are detuned from the transition by Δ and only two-photon transitions can occur via a virtual level (dashed line). The atom is still in $|g\rangle$ after the transfer, but has acquired momentum q_r and kinetic energy hf .

7.1 Theory of Bragg spectroscopy

Bragg spectroscopy uses two off-resonant laser beams to transfer momentum and energy to the atoms. The frequency dependence of this process allows us to probe the energy levels of the atoms. Unlike in Raman spectroscopy, we are not transferring the atom between different *internal* states (such as hyperfine levels), but we transfer the atom only to a different *momentum* state. We illustrate the process in Fig. 7.1. Two beams with different frequencies ($f_2 < f_1$) are overlapped at an angle 2θ on the atoms. The two beams have wavevectors \mathbf{k}_1 and \mathbf{k}_2 . Absorbing one photon from one beam, and stimulated emission into the other beam causes the atom to absorb energy and recoil momentum. Energy conservation dictates that the atom acquires the energy

$$hf = h(f_1 - f_2) \quad (7.1)$$

and the recoil momentum

$$\mathbf{q}_r = \hbar(\mathbf{k}_1 - \mathbf{k}_2). \quad (7.2)$$

For now, we neglect mean-field effects and approximate the dispersion relation of the atoms by that of free, classical particles

$$E_{kin} = \frac{\mathbf{p}^2}{2m} \quad (7.3)$$

with momentum \mathbf{p} . We assume the momentum transfer \mathbf{q}_r is parallel to the z-component of the momentum p_z . It is now easy to derive the Doppler-shifted resonance frequency

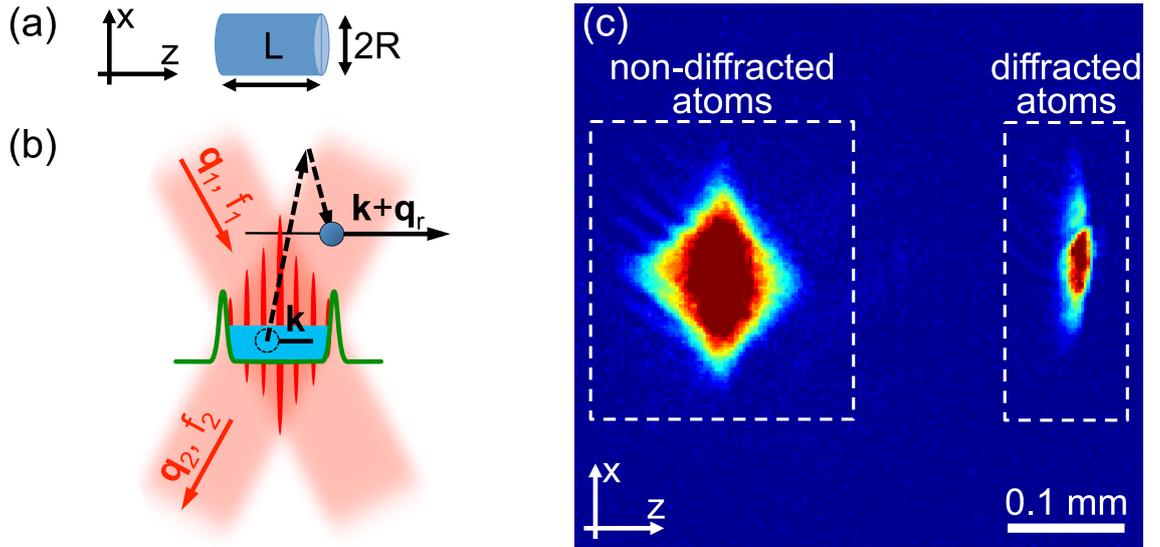


Figure 7.2: Bragg spectroscopy of a uniform BEC. (a) Shape and orientation of our box trap; in the lab frame the x direction is vertical. (b) A trapped atom with initial momentum $\hbar k$ absorbs recoil momentum q_r aligned with z . The trap wall height (green) is lower than the energy of the recoiling atoms. The Bragg beams and lattice are not drawn to scale; in reality the beams are much larger than the box and the lattice spacing is much smaller. (c) Absorption image of the atoms, taken 170 ms after the start of the Bragg pulse. After the end of a 35 ms light pulse, the non-diffracted atoms were released from the trap. Figure adapted from [26].

$f_r(p_z)$

$$hf_r(p_z) = \frac{q_r^2}{2m} + \frac{p_z q_r}{m}. \quad (7.4)$$

Scanning the resonance frequency allows us to measure the momentum distribution of the cloud. Usually, $k_1 \approx k_2 = k$ and we can approximate $q_r \approx 2\hbar k \sin(\theta)$. We will operate in a regime where the recoil velocity $v_r = q_r/m$ imparted on the atom is much larger than the initial velocity of the atom p_z/m . After a long enough ToF, the diffracted atoms will thus separate from the rest of the cloud as shown in Fig. 7.2. From the number of atoms in the diffracted cloud we can deduce the fraction of atoms whose momentum matches the applied frequency.

7.1.1 Transfer function

So far, we have not discussed the time evolution of the Bragg diffraction process. It can be addressed within the same mathematical framework as two-photon Raman transitions. The two overlapping beams in the rotating-wave approximation produce a semi-classical

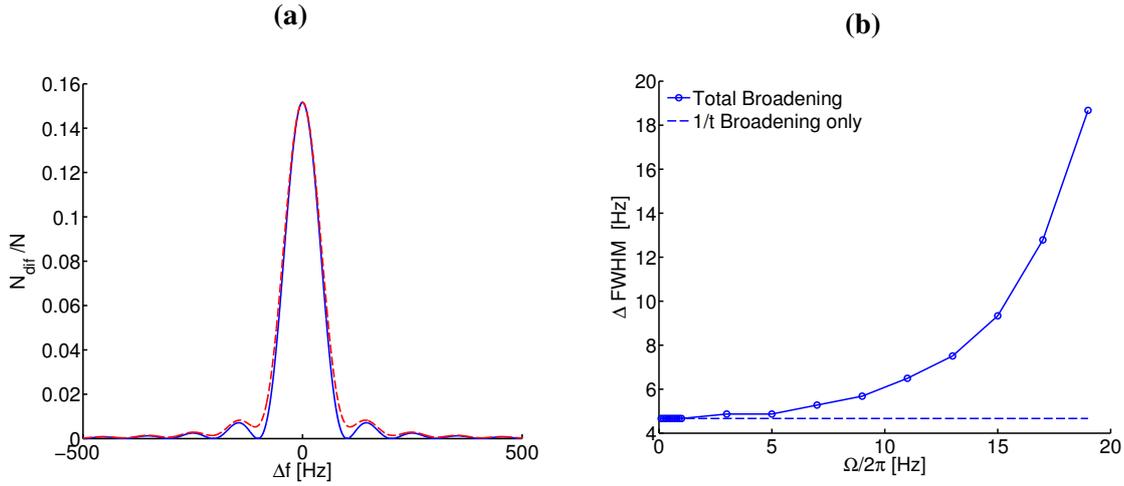


Figure 7.3: (a) Ideal box spectrum (blue) with a zero-to-zero width of 200 Hz compared with the broadened spectrum (dashed red) due to $\Omega = 2\pi \times 8$ Hz and $t = 35$ ms. (b) Broadening of the FWHM (see text) due to Ω with fixed $t = 35$ ms (continuous blue). The dashed line indicates the baseline broadening due to the temporal $1/t$ Fourier-broadening.

light field with amplitude

$$|E| \propto \cos \left(\frac{1}{2} [(\mathbf{q}_r/\hbar)\mathbf{r} - 2\pi ft] \right). \quad (7.5)$$

This can be seen as an optical lattice of wavelength h/q_r moving with velocity $v_r/2$. The name Bragg spectroscopy is derived from this picture. Unlike in the traditional Bragg process, we are not scattering photons (e.g. x-rays) off a lattice made of matter (e.g. crystals), but we are diffracting a matter wave off a lattice made of photons. The perturbation Eq. (7.5) causes a Rabi oscillation between the levels $|g, p_z\rangle \rightarrow |g, p_z + q_r\rangle$ via a virtual level (see Fig. 7.1). The effective Rabi frequency Ω is given by [37]

$$\Omega = \frac{\Gamma^2 I}{4\Delta I_{\text{sat}}}, \quad (7.6)$$

where Γ is the line width of the $|e\rangle$ level and Δ is the detuning from the $|g\rangle \rightarrow |e\rangle$ transition. $I \approx I_1 \approx I_2$ are the beam intensities and I_{sat} is the saturation frequency of the transition. We define the detuning δ as

$$\delta = 2\pi(f - f_r(p_z)) = 2\pi \left(f - \frac{q_r^2}{2mh} - \frac{p_z q_r}{mh} \right). \quad (7.7)$$

The transfer probability $P(p_z)$ for a given momentum p_z is then given by [37]

$$P(\delta, t) = \frac{\Omega^2}{\Omega^2 + \delta^2} \sin^2 \left(\frac{\sqrt{\Omega^2 + \delta^2} t}{2} \right), \quad (7.8)$$

where t is the duration of the Bragg lattice light pulse.

It is convenient to use Bragg frequency units to describe the momentum p_z

$$\tilde{f} = \frac{p_z q_r}{\hbar m}. \quad (7.9)$$

In these units, Eq. (7.8) can be written as

$$P(\tilde{f}) = \left(\frac{\Omega t}{2} \right)^2 \text{sinc}^2 \left(\frac{\sqrt{\Omega^2 + 4\pi^2(f - f_r(0) - \tilde{f})^2} t}{2} \right). \quad (7.10)$$

We assume that all of our atoms N are in the ground state with momentum distribution $\Psi(p_z)$ along the z-axis, or equivalently $\Psi(\tilde{f})$ in frequency units. The diffracted atom number N_{dif} for a given frequency can be calculated by summing the transferred fraction over all momentum states

$$N_{dif}(f) = N_0 \int |\Psi(p_z)|^2 P(p_z) dp_z \quad (7.11)$$

$$= N_0 \int |\Psi(\tilde{f})|^2 P(\tilde{f}) d\tilde{f}. \quad (7.12)$$

This is simply a convolution of the transfer function Eq. (7.8) with the momentum distribution $|\Psi(\tilde{f})|^2$ centred around $f_r(0)$:

$$\frac{N_{dif}(f - f_r(0))}{N} = \left(\frac{\Omega t}{2} \right)^2 \int |\Psi(\tilde{f})|^2 \times \text{sinc}^2 \left(\frac{\sqrt{\Omega^2 + 4\pi^2(f - f_r(0) - \tilde{f})^2} t}{2} \right) d\tilde{f}. \quad (7.13)$$

In the limit of long Bragg times t and low Rabi frequencies Ω (i.e. low beam intensities), $P(\tilde{f})$ turns into a δ -function and we observe the momentum distribution directly. However, finite times t and Rabi frequencies Ω result in broadening of the spectrum, as illustrated in Fig. 7.3a. Considered separately, finite times t with negligible Ω cause a temporal sinc-squared Fourier-broadening with zero-to-zero width $2/t$ in frequency units. Also, finite Rabi frequencies Ω cause additional power-broadening (see Fig. 7.3b). Experimentally, we observe the diffracted fraction, which scales $\propto (\Omega\tau)^2$. Hence, to keep t in an experimentally viable range (especially for measurements which are in ToF rather than in trap), we have to accommodate a certain broadening. However, Eq. (7.13) gives

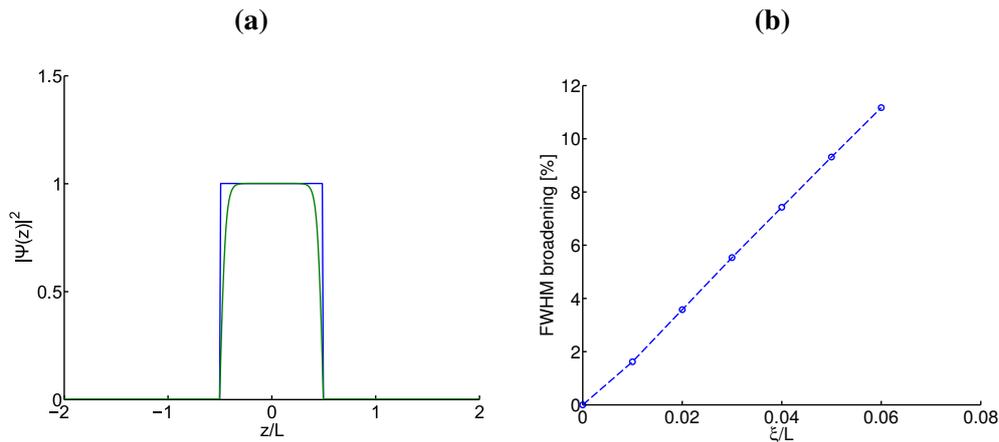


Figure 7.4: (a) Wavefunction of an interacting BEC in a perfect box (blue). The rounding-off effect of the healing length (here for a typical length $\xi \approx 0.03L$) is shown (green). (b) FWHM broadening in momentum space due to the healing length for a range of typical experimental values of ξ/L .

us the analytic tools to predict this broadening.

To quantify the broadening in a simple manner, we use the full-width at half-maximum (FWHM)¹. Analytically, the implementation is trivial. To obtain the FWHM from experimental data, we use appropriate fitting functions to approximate the data; the function used is not relevant as long as it reproduces the data well at the maximum and the half-maxima. Typically, the sum of two Gaussians or a sinc-squared on top of a Gaussian work well.

7.1.2 Momentum distributions

The momentum distribution of a non-interacting BEC in a box of length L is given by the quantum-mechanical ground state of a single particle. The spatial wave function is given by a cosine in the box and is zero elsewhere

$$\Psi(z) \propto \Pi(z/L) \cos(\pi z/2L), \quad (7.14)$$

which we expressed mathematically by truncating the cosine at the box edges using the tophat function $\Pi(z/L)$, where $\Pi(z/L) = 1$ for $|z| < L/2$ and zero otherwise. We obtain the corresponding momentum-space wave function by Fourier transformation

$$\Psi(k_z) \propto \cos(k_z L/2)/(k_z^2 L^2 - \pi^2). \quad (7.15)$$

¹We decided against using the variance as a measure of the broadening, as it diverges for a sinc-squared distribution - which is the the momentum distribution in a perfect box.

The zero-to-zero width in k -space is $6\pi/L$. According to Eq. (7.9), the corresponding Bragg spectral line has zero-to-zero width $3v_r/L$ (in frequency units).

On the other hand, the profile of an interacting BEC in the Thomas-Fermi regime is a simple tophat function of width L (see Eq. (2.32))

$$\Psi(z) \propto \Pi(z/L), \quad (7.16)$$

where we are for now neglecting blurring at the edges of the box because of the healing length (Eq. (2.33)). The corresponding wave function in momentum-space is

$$\Psi(k_z) \propto \text{sinc}^2(kL/2), \quad (7.17)$$

with a zero-to-zero of $4\pi/L$ in k -space and $2v_r/L$ in Bragg frequency units.

We notice that these scalings of the widths in real and momentum space are manifestations of the Heisenberg uncertainty relation

$$\Delta x \Delta k \gtrsim 1. \quad (7.18)$$

Confining an atom in a box gives us information about its position, fixing the uncertainty² to $\Delta x \approx L$. Hence, the momentum uncertainty has to scale as $\Delta k \approx 1/L$.

In practice, the edges of the box are going to be smeared out on the order of the healing length ξ , which is defined as

$$\frac{\hbar^2}{2m\xi^2} = ng, \quad (7.19)$$

where n is the atom density and we recall $g = 4\pi\hbar^2 a/m$ is proportional to the s-wave scattering length a . If $\xi \ll L$, we can approximate the blurred tophat wave function with [28]

$$\Psi(z) = \Psi_0 \tanh\left(\frac{(L/2) - |z|}{\sqrt{2}\xi}\right) \Pi(z/L). \quad (7.20)$$

A comparison of this shape and the original rectangle function is illustrated in Fig. 7.4a. The broadening induced by the healing length can be determined by numerically calculating the Fourier transform. For typical values of $\xi \approx 0.05L \approx \mu\text{m}$, the FWHM broadening is below 10%. The expression for the healing length in Eq. (7.20) is only valid for small healing lengths; if ξ becomes comparable to L the underlying Thomas-Fermi approximation becomes invalid. The effects of the diffraction-limited steepness of the trap are hard to measure, as they cannot be resolved by our imaging resolution of $\approx 5 \mu\text{m}$. Hence, we can estimate the effects to be comparable to typical healing length broadening.

²One should note that numerical factors enter the Heisenberg relation, depending on how the uncertainties Δx and Δk are defined.

7.1.3 Mean-Field corrections

Until now, we have only discussed the effect of the interactions on the momentum distribution. However, they also affect the dispersion relation. In the Bogoliubov approximation, the dispersion relation is given by [28, 77]

$$E = hf(q) = \sqrt{\left(\frac{q^2}{2m}\right) 2ng + \left(\frac{q^2}{2m}\right)^2}, \quad (7.21)$$

where q is the momentum. For large momentum transfers $q_r^2/2m > ng$, we can approximate this as

$$hf(q) = \frac{q^2}{2m} + ng. \quad (7.22)$$

This introduces an additional mean-field shift $\Delta f_{MF} = ng/h$ to the centre frequency $f_0 = q_r^2/2m$ in Eq. (7.4) [83, 87].

$$hf_r(p_z) = \frac{q_r^2}{2m} + \frac{p_z q_r}{m} + ng \quad (7.23)$$

$$\equiv f_0 + \Delta f_{Doppler}(p_z) + \Delta f_{MF}. \quad (7.24)$$

This shift can be explained by the additional mean-field energy gn an atom gains when leaving the ground-state. This is caused by the ‘‘bosonic factor of 2’’, which doubles the interaction energy of an atom with atoms in different states. Thus, if all atoms are in the ground state, the mean-field energy of an atom has to double when leaving it. In our experiments, we are always deeply in the regime where Eq. (7.22) applies. The phonon regime, where the dispersion relation Eq. (7.21) is linear, can be probed by reducing q_r [84, 91]. This is typically achieved by using small angles θ between the beams.

Besides the shift of the resonant frequency Δf_{MF} , any spatial variation of the density n leads to inhomogeneous broadening of the Bragg spectrum. To extract the momentum distribution from the measured spectra, we have to subtract the mean-field broadening. As this can be easily comparable to the momentum width, it introduces quite large uncertainties [83]. In our case, n is essentially homogeneous, the broadening effect is absent and interactions just shift the Bragg resonance. We can thus directly measure the momentum distribution in the BEC.

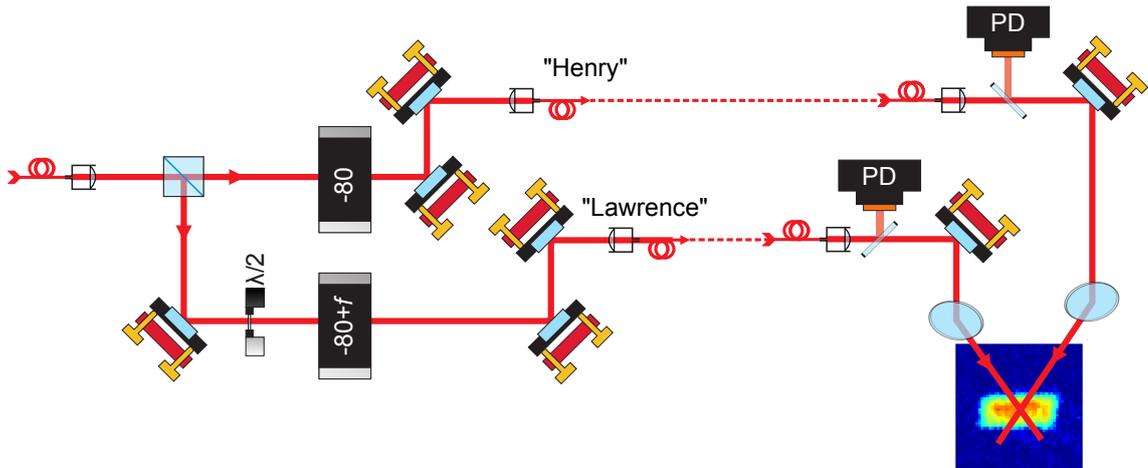


Figure 7.5: The optical setup for Bragg spectroscopy. We split off light from the repump laser via an optical fibre. The beams are then sent through AOMs, shifting the frequencies and stabilizing the intensity. Finally, they are delivered to the cloud via further fibers. Photodiodes (PD) are used for intensity feedback. Figure taken from [27].

7.2 Experimental realization

7.2.1 Setup

In Fig. 7.5, we outline our Bragg-spectroscopy setup. We intersect two collimated (3 mm wide) 780 nm laser beams on the atoms at an angle of $2\theta = 30^\circ$. They are detuned from the atomic $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition by 6.8 GHz. This is several orders of magnitude larger than typical Rabi frequencies, strongly suppressing single-photon transfers to $|F' = 3\rangle$. The beam parameters result in a recoil velocity $v_r \approx 3$ mm/s. The diffracted and non-diffracted atoms are thus well separated in velocity space, as qualitatively seen in Fig. 7.2c.

The beams are branched off from the repump light and split into two branches of equal intensities. In honour of Bragg father and son [92] we named them “Lawrence” and “Henry”. The beams are delivered via the top viewport of the MOT cell. To rule out broadening of the spectra due to incoherences, we adjusted the optical path difference between “Lawrence” and “Henry” to within a few cm.

Acousto-optic modulators (AOMs)³ are used to shift Henry’s and Lawrence’s frequencies by 80 MHz and $(-80 + f)$ MHz respectively. The AOMs are driven by two phase-locked function generators⁴, which are amplified with broadband radio-frequency (RF) amplifiers. The phase-lock makes sure that the two beams are shifted coherently.

³Crystal Tech

⁴Agilent 33250A

Two TTL-controlled RF switches⁵ allow us to switch the AOM with the required μs precision. The frequency of the ‘‘Lawrence’’ function generator is set via the GPIB protocol⁶.

To ensure stable and reproducible beam intensities, we monitor weak reflection off glass plates. As typical light pulses are very short, it is very difficult to stabilize them on-the-go. Instead, we pre-stabilize the beam intensities while the MOT is loading and then switch them back on when needed to the same value. This allows us to counter at least long term drifts. To avoid having to use analogue sample-and-hold feedback devices (such as special PIDs), we use digital feedback instead. The intensity signal from the photodiodes is read by an Arduino microcontroller. The microcontroller forwards the intensity values to a custom-written MATLAB program on the experiment control computer (via serial communication). This program then adjusts the output voltage of the function generators via GPIB until the intensities are stabilized. The target values are received at the beginning of the sequence from the WordGenerator software⁷.

7.2.2 Heisenberg uncertainty relation

First, we measured the experimental Bragg spectrum and assessed how close it is to the sinc-squared momentum distribution $S(f)$ (Eq. (7.17)) of a uniform gas, as well as to the corresponding time- and power broadened- distribution $\tilde{S}(f)$ given by Eq. (7.13). We prepared clouds with condensed fractions $\gtrsim 80\%$ at a temperature $T \lesssim 10\text{ nK}$ by evaporative cooling in the box trap. We then applied a Bragg pulse of length $t = 35\text{ ms}$ and a Bragg Rabi frequency with $\Omega/(2\pi) \approx 8\text{ Hz}$, keeping the fraction of diffracted atoms to $\lesssim 15\%$. We turned off the trap 25 ms after the end of the Bragg pulse and measured the fraction of diffracted atoms after further 110 ms of ToF (see Fig. 7.2c).

The duration of the Bragg pulse was chosen to be $t \gtrsim L/v_r$. This ensures that the Bragg spectrum is not significantly Fourier-broadened. Equivalently, a recoiling atom can traverse the box during the pulse, and we thus probe phase coherence across the whole BEC. To apply very long Bragg pulses to a trapped cloud, we set the trap depth below the recoil energy $hf_r(0) \approx 50\text{ nK}$, so that the diffracted atoms can leave the box without bouncing off the trap walls (see Fig. 7.2b).

In Fig. 7.6a, we show measurements for $L = 30 \pm 1\ \mu\text{m}$. As shown in the left inset, L is determined by fitting the in-trap BEC density profile with a top-hat function convolved with a Gaussian that accounts for our $5\ \mu\text{m}$ imaging resolution. Within experimental errors, the measured L is consistent with the value we calculate from the settings of the

⁵Minicircuits ZYSWA-2-50DR

⁶We use a USB to GPIB adapter (prologix, GPIB-USB Controller 6.0) allowing us to use the simpler RS232 protocol.

⁷Communication between software is implemented with a virtual serial port created by the free com0com software.

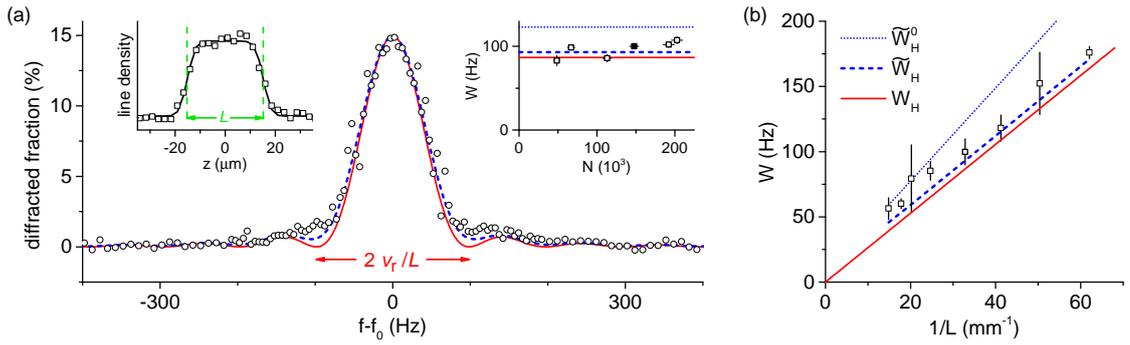


Figure 7.6: Heisenberg-limited momentum spread in a uniform interacting BEC. (a) Main panel: Measured Bragg spectrum of a trapped BEC of length $L = 30 \mu\text{m}$. The solid red and dashed blue lines show the theoretical Heisenberg-limited spectra $S(f)$ and the power-time broadened $\tilde{S}(f)$, respectively. Left inset: L is determined by fitting the in-trap BEC density profile, accounting for the imaging resolution. Right inset: FWHM of the spectrum, W , versus the atom number N , for the same L . The filled square corresponds to the data in the main panel, with $W = (100 \pm 3) \text{ Hz}$. The solid red and dashed blue lines show $W_H = 0.89 v_r/L = 87 \text{ Hz}$ and $\tilde{W}_H = 93 \text{ Hz}$, corresponding to $S(f)$ and $\tilde{S}(f)$, respectively. The dotted blue line shows $\tilde{W}_H^0 = 123 \text{ Hz}$, expected for a non-interacting BEC. (b) W versus inverse box length, $1/L$, showing the expected Heisenberg scaling. Solid red, dashed blue and dotted blue lines show W_H , \tilde{W}_H and \tilde{W}_H^0 , respectively. Figure taken from [26].

spatial light modulator used to create and align the trapping beams. In the main panel, we show the measured Bragg spectrum, centred on $f_r(0) \approx 1 \text{ kHz}$. The solid red and dashed blue lines show theoretical expectations for a Heisenberg-limited momentum spread. The solid line is the sinc-squared function $S(f)$, while the dashed line represents the expected broadening $\tilde{S}(f)$ according to Eq. (7.13). The data is extremely close to the Heisenberg limit, corresponding to a fully coherent BEC wavefunction spanning the whole box. (For corresponding measurements on a harmonically-trapped gas, in momentum and real space, see [83] and [93], respectively.)

We quantitatively compare different (measured and calculated) Bragg spectra using their full-width-at-half-maximum W . In the right inset of Fig. 7.6a, we plot the measured W versus the total atom number N , for $L = 30 \mu\text{m}$. The solid red and dashed blue lines show the two calculated Heisenberg-limited values, W_H for $S(f)$ and \tilde{W}_H for $\tilde{S}(f)$. For comparison, we also calculate \tilde{W}_H^0 (dotted blue line) for the sine-like non-interacting ground state of the box potential. We see that interactions reduce W below \tilde{W}_H^0 . Moreover, the measured W shows essentially no dependence on N , as expected for a BEC of a spatially uniform density.

To confirm the Heisenberg-like scaling of the momentum-width with length as predicted in Eq. (7.18), we repeated the measurement for a range of box lengths, $L = 15 - 70 \mu\text{m}$. While numbers and densities varied across those measurements, we have

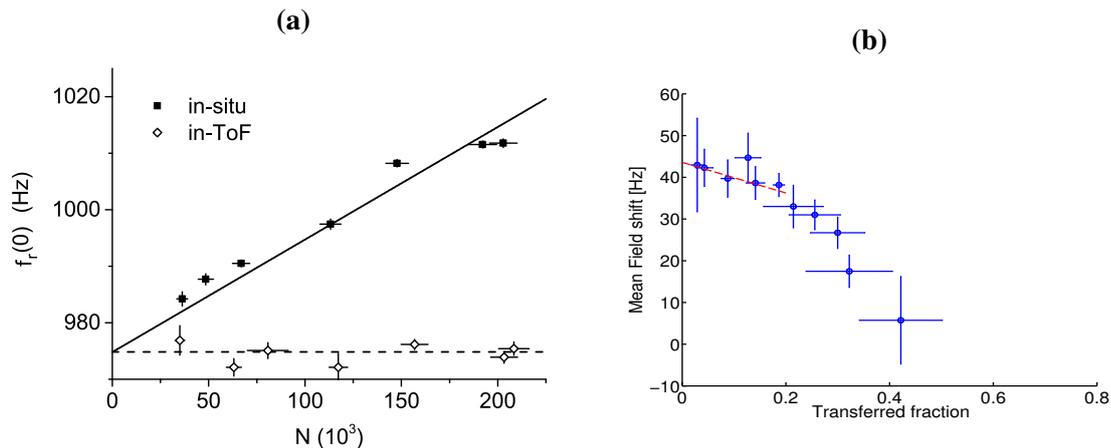


Figure 7.7: (a) Interaction energy in a uniform BEC, for $L = 30 \mu\text{m}$. We plot $f_r(0)$ versus N for Bragg pulses applied in situ (solid squares) and after 50 ms of ToF (open diamonds). Dashed line: $f_r(0) \approx f_0 \approx 975$ Hz, solid line: $f_r(0) = f_0 + \alpha N$, with $\alpha \approx 20 \times 10^{-5}$ Hz/atom. Figure adapted from [26]. (b) Dependence of the mean-field shift $\Delta f_{MF} = f_r(0) - f_0$ on the transferred fraction. To vary the transferred fraction, we scanned Ω at constant $t = 35$ ms.

already established that the momentum-width is essentially number-independent as long as we are deeply in the Thomas-Fermi regime. In Fig. 7.6b, we summarise our data. Plotting W versus $1/L$, we demonstrate the expected Heisenberg scaling of the momentum uncertainty. The deviations from the expected broadened width \tilde{W}_H are below 10% and can be explained by a combination of healing length (see Fig. 7.4b) and diffraction-limited steepness of the walls.

7.2.3 Mean-field shift

We now turn to the study of the ground-state energy of a uniform interacting BEC. From Eq. (7.24), we see that the mean-field shift of the centre frequency of the Bragg spectrum, Δf_{MF} , is directly proportional to the chemical potential $\mu = ng = \partial E / \partial N$. Thanks to the unprecedented narrowness of our Bragg spectra, we measured such shifts with a precision of 2 Hz, corresponding to an energy of $k_B \times 100$ pK.

In Fig. 7.7a we plot the central frequency $f_r(0)$ versus N , for a fixed $L = 30 \mu\text{m}$ and for two sets of Bragg spectra:

- The in situ spectra (solid symbols) were taken as above, with the main cloud trapped during the Bragg pulse. These spectra are obtained in the same way as the Heisenberg spectra above. Varying the atom number N changes the mean-field shift.
- The “in-ToF” spectra (open symbols) were taken by releasing the BEC from the box and letting it expand for 50 ms before applying the pulse. After long ToF the

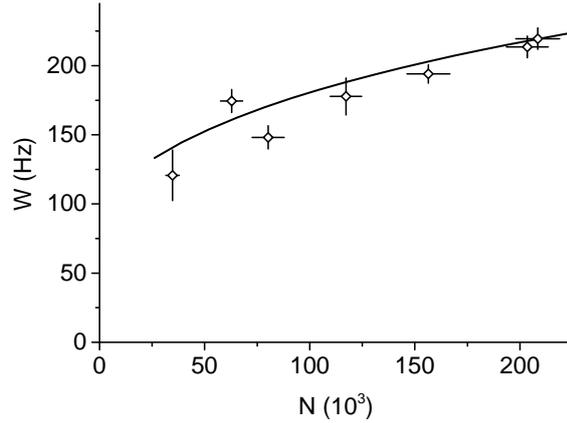


Figure 7.8: Spectral width for Bragg pulses applied after 50 ms of ToF expansion, for $L = 30 \mu\text{m}$. The W values are extracted from the in-ToF spectra in Fig. 7.7a. The solid line shows a numerical simulation based on the GP equation. Figure taken from [26].

atomic density is negligible, and $\Delta f_{MF} \approx 0$. The N -independent central frequency $f_r(0)$ provides a good measurement of the baseline resonance $f_0 \approx 975$ Hz. For the measurement of f_0 , it is quite important to measure both $\pm f_0$ diffraction orders. Small potential gradients during long ToF accelerate the cloud to the left or right, causing a global Doppler shift of the spectrum. Averaging the central frequencies allows us to cancel this effect.

For the in situ spectra, within mean-field theory and for an infinitesimal Bragg transfer ($\Omega\tau \rightarrow 0$), we expect

$$\Delta f_{MF} = f_r(0) - f_0 = \frac{2\hbar a \eta}{Vm} N \equiv \alpha N, \quad (7.25)$$

where η is the condensed fraction and V is the volume of the box. From in situ images we measured $V = (25 \pm 2) \times 10^3 \mu\text{m}^3$. We assess $\eta = 0.8 \pm 0.1$ from the maximal fraction of Bragg-diffracted atoms (15% in Fig. 7.6a). This estimate is further supported by ‘‘BEC filtering’’ introduced in [59].⁸ The uncertainty indicates variations between experimental runs. We thus obtain a theoretical prediction $\alpha_{\text{th}} = (24 \pm 2) \times 10^{-5}$ Hz/atom.

From a linear fit to the in situ data (solid line in Fig. 7.7a) we get $\alpha_{\text{exp}} = (20 \pm 1) \times 10^{-5}$ Hz/atom, slightly below α_{th} . This small difference can be attributed to the 15% depletion of the ground-state population during the pulse; the effective mean-field shift varies with time. To quantify this effect, we took additional measurements with a range of Ω for our largest N , with constant $t = 35$ ms. This corresponds to varying the

⁸We apply a short (4 ms) Bragg π -pulse to the cloud. This pulse has a Fourier-width of 500 Hz zero-to-zero, which transfers the BEC, but barely affects the much broader thermal distribution. In ToF, the two clouds separate and the condensed fraction can be measured.

transferred fraction $\propto (\Omega t)^2$. In Fig. 7.7b, we see a clear reduction of the mean-field shift Δf_{MF} with increasing transferred fraction. By extrapolating the shift to zero transfer (red dashed line), we can estimate that $\Delta f_{MF}(15\%) \approx \Delta f_{MF}(0\%) \times 0.87$. We thus get a slightly revised $\tilde{\alpha}_{\text{exp}} = (23 \pm 1) \times 10^{-5}$ Hz/atoms. The systematic uncertainty in $\tilde{\alpha}_{\text{exp}}$, due to the systematic uncertainty in our absolute atom-number calibration, is 10%.

Complementary to Fig. 7.7a, we plot the measured width W for the in-ToF Bragg spectra in Fig. 7.8. Qualitatively, W now grows with N because during ToF interaction energy gets converted into kinetic energy. Quantitatively, the problem of the expansion of an interacting BEC from a box potential has not been solved analytically (see [32] for the harmonic-trap case). However, we find good agreement between our data and numerical simulations (solid line) based on the GP equation (Eq. (2.26)) done by Alexander Gaunt. Using the GPE, we find the momentum and density distribution in-trap and after expansion respectively; convolution with Eq. (7.24) then predicts the measured Bragg spectrum including mean-field broadening as well as power- and time-broadening. In our simulations, we neglect the small thermal component and use the measured in-trap BEC energy ($\propto \tilde{\alpha}_{\text{exp}} N$) to predict W in ToF. We also take into account the small curvature of our gravity-compensating magnetic field. This has a noticeable effect only for long ToF ($\gtrsim 50$ ms).

7.2.4 Evolution in ToF

Finally, we study the evolution of the BEC in ToF. We probe the cloud after 0 ms to 50 ms free expansion with $L = 30 \mu\text{m}$ and $N \approx 200 \times 10^3$. In Fig. 7.9a, we show absorption images of the expanding cloud (top), and the corresponding GP-based simulations (bottom). The simulations reproduce the characteristic diamond shape that emerges during ToF. Qualitatively, this conversion to a diamond shape is the analogue of the inversion of the aspect ratio of a BEC released from an anisotropic harmonic trap. In both cases the force driving the initially accelerating expansion is given by the gradient of the atomic density (i.e., the gradient of the interaction-energy density) and in both cases the fastest moving wave fronts develop at the points where the curvature of the constant-density surfaces is minimal during the early stages of the expansion.

We measured the Bragg spectrum at different times during the expansion, similarly to the in-ToF spectra above. In Fig. 7.9b we show the gradual decay of Δf_{MF} and growth of W for in-ToF Bragg spectra, again finding good agreement with our simulations (solid lines). Note that in these experiments, and simulations, we reduced τ to 10 ms and increased Ω to $2\pi \times 28$ Hz. While this broadens our spectral resolution, we need to improve the time resolution to be able to probe the distribution at a given ToF.

Alternatively, it is possible to use the position of the diffracted cloud to extract its

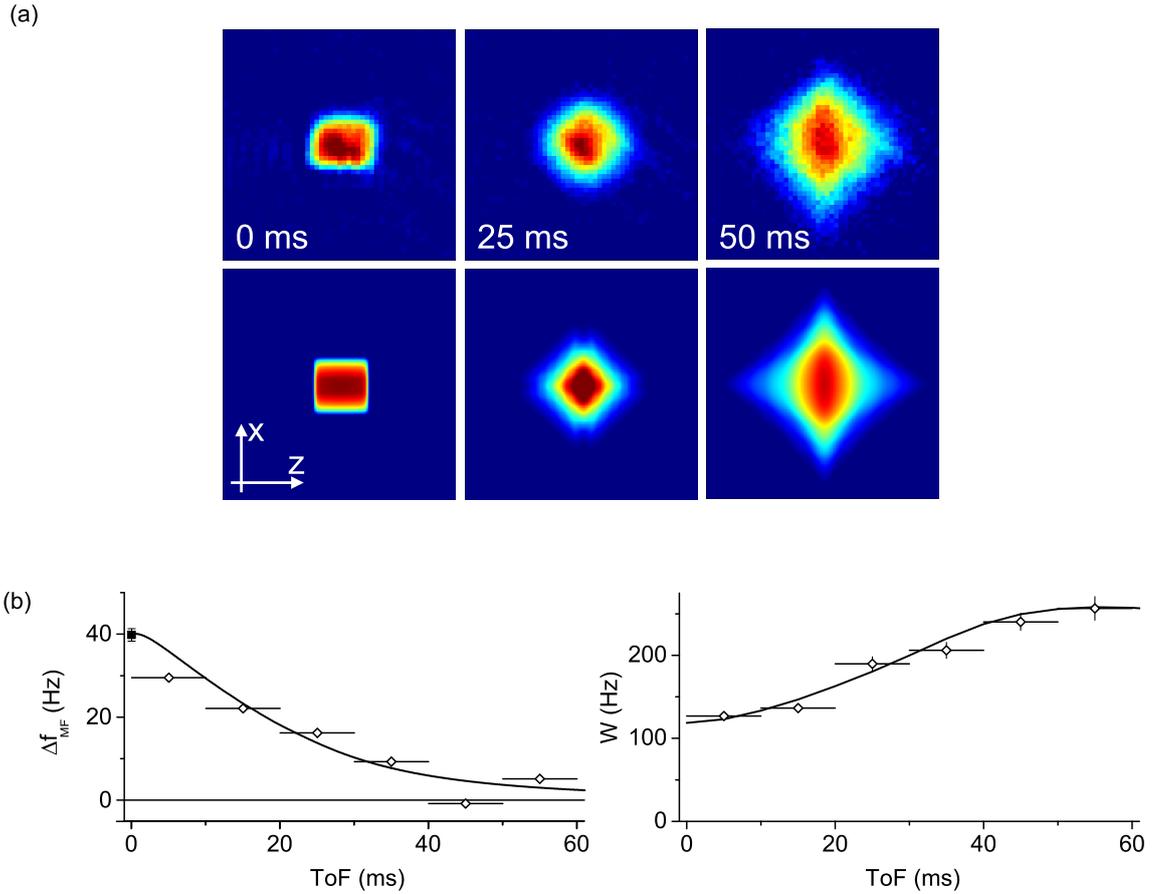


Figure 7.9: Free expansion of a BEC released from the box trap with $L = 30 \mu\text{m}$ and $N \approx 200 \times 10^3$. (a) Absorption images (top) and simulations (bottom) of the atomic distributions. (b) Evolution of Δf_{MF} and W during ToF. The solid lines show numerical simulations based on the GP equation. The solid square corresponds to the in-trap data from Fig. 7.7a. The horizontal error bars indicate our temporal resolution, limited by the Bragg-pulse duration. Figure taken from [26].

momentum p_z at the time of the outcoupling. During the initial ToF, $\text{ToF}_A = 0 - 60$ ms, the cloud moves with a velocity of $v_z \approx p_z/m$. This assumes a negligible (or fast enough) change of the initial velocity due to mean-field energy. The Bragg pulse then increases the velocity from v_z to $v_z + v_r$. The diffracted atoms are much faster than the rest of cloud, and separate quickly. In the process, the atoms “roll” down the mean-field potential of the main cloud and convert $g \times n(\mathbf{r}, t = \text{ToF}_A)$ into additional kinetic energy. The velocity \tilde{v} is now

$$E_{kin} = \frac{(p_z + q_r)^2}{2m} + gn = \frac{1}{2}m\tilde{v}^2. \quad (7.26)$$

If we now wait for another expansion time ToF_B before imaging, the cloud will have travelled a total distance z

$$z(\text{ToF}_A + \text{ToF}_B) \approx v_z \times \text{ToF}_A + \tilde{v} \times \text{ToF}_B. \quad (7.27)$$

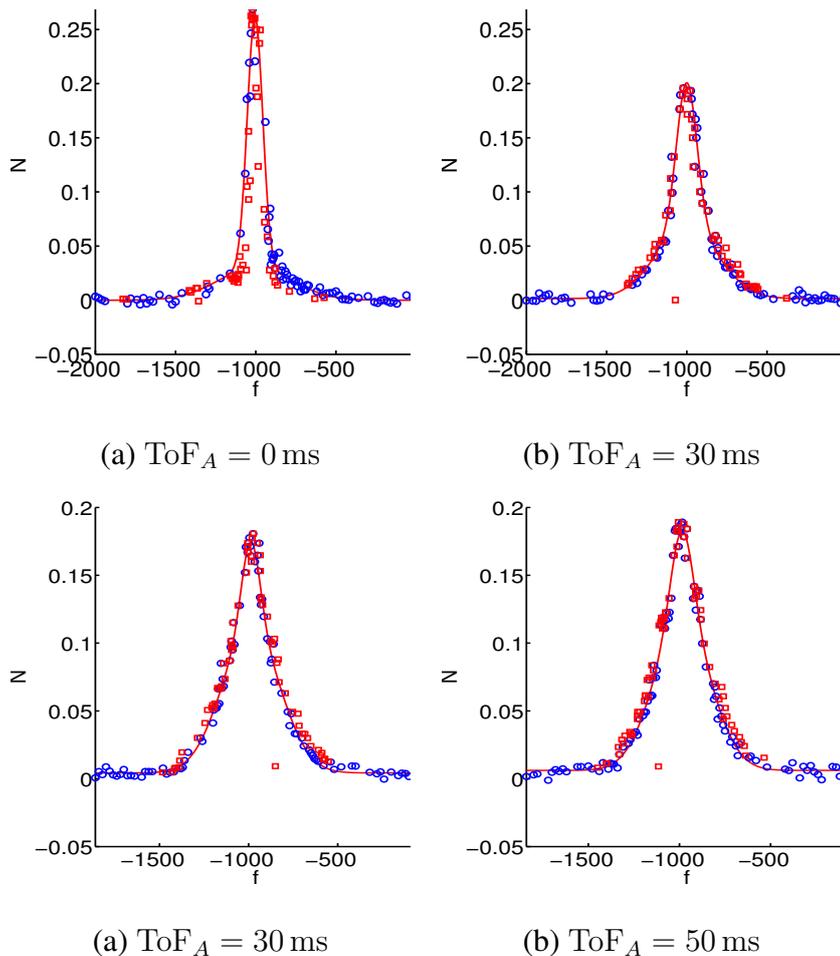


Figure 7.10: Comparison of the Bragg spectra (blue) and position spectra (red) for various ToF_A s. The continuous red lines are bimodal-Gaussian fits to the Bragg spectra. We see that the two methods produce similar results, although the Bragg offers a better signal to noise ratio in the wings.

Clearly this is an approximation, neglecting changes in velocity v_z during ToF_A . The approximation holds in the limit of $\text{ToF}_A \ll \text{ToF}_B$. By fitting the position of the diffracted cloud, we can extract z . As ToF_A is small, we approximate $v_z \approx \tilde{v} - v_r$. Now, we extract \tilde{v} using Eq. (7.26), which allows us to reconstruct a velocity against diffracted number spectrum that is equivalent to the mean-field broadened Bragg spectrum given by Eq. (7.24).

In Fig. 7.10, we overlap the resulting position spectra (red) with their Bragg counterpart (blue). We see that we manage to obtain very similar spectra. However, the noise levels are much higher for the position-based spectra. Especially in the wings, where signal levels are low, one cannot obtain reliable position measurements. Further, the influence of the expansion during ToF_A is unclear. Additionally, we need to calibrate for potential variations along the flight path; in our case the weak harmonic curvature induced by the compensating magnetic field. There are several ways to improve this

method; especially a differential measurement of z for various ToF_B would allow us to extract \tilde{v} without having to worry about what happens during ToF_A . We did not pursue this idea, as the frequency method provided the precision we required.

7.3 Conclusions

In conclusion, we have characterized the ground-state properties of an interacting homogeneous Bose gas, including the Heisenberg-limited momentum distribution, the interaction energy, and the free-expansion dynamics. An important by-result of our measurements is that they place the most stringent bound so far on the spatial uniformity of an ultracold gas produced in our optical box. While earlier (thermodynamic) studies established uniformity on a 30 – 100 nK energy scale [24, 25], all our present measurements indicate that our gas behaves as a homogeneous system down to a sub-nK energy scale (corresponding to 20 Hz in frequency units).

7.3. CONCLUSIONS

Chapter 8

Conclusions and Future Work

In this thesis we have demonstrated the first experimental realization of a homogeneous Bose-Einstein condensate. We have discussed the experimental procedure and equipment for trapping atoms in a uniform potential. Further, we have shown by verifying thermodynamic scalings that the trap is nearly uniform on an energy scale $k_B T \approx k_B \times 100$ nK. In the process, we have observed the quantum Joule-Thomson effect for the first time. We have also seen that the gas follows approximately textbook-like ideal gas scalings despite atomic interactions; most notable, the thermal component is not altered by the presence of the condensate. This is in stark contrast to the situation in a harmonic trap, where interactions strongly modify the thermodynamic behaviour. After scrutinizing the thermal gas, we have probed the ground state of the condensate. Using Bragg spectroscopy, we have measured the momentum and density distribution of the condensed fraction. We have managed to observe the Heisenberg uncertainty on a macroscopic length scale to within a few percent.

Besides the creation of a homogeneous BEC, we have demonstrated the ability of a comparatively simple and compact experimental system to produce BECs suitable for cutting-edge experiments. Still, the shortcomings of the system are not deniable: While the size of the condensate is respectable, the accessible parameter range was often limited by atom numbers. Further, the use of a single chamber for the MOT and evaporation process limits optical access to the cloud. At this stage, we have pretty much exhausted all free space around the cell for additional optics. Also, the high background gas pressure severely limits the life-time of BECs. Finally, the lack of a convenient Feshbach resonance in ^{87}Rb (and the physical space to accommodate coils) means we cannot tune the interactions at will. Still, in the long run, it should be of interest especially for groups on a tight budget or for teaching purposes.

For those reasons, a new experiment similar to the one described in [53] is currently under construction. Together with the now established experimental toolkit for the pro-

duction of uniform gases, it should enable (or at least facilitate) a range of interesting experiments:

- Tunable interactions and a better imaging resolution will enable more accurate studies of the thermodynamics in the box and should allow us to observe beyond mean-field effects.
- Squeezing the trap along one dimension, will allow us the study of homogeneous gases in two dimensional gases.
- Using two separate boxes, coupled with a thin channel should allow us the observation of transport phenomena similar to superfluid helium experiments.
- Currently our group is exploring the dynamics of condensate domain formation depending on the quench rate through the Bose-Einstein phase transition. This should allow us to establish a link to the Kibble-Zurek mechanism.

While we have measured some so-far elusive effects in our work with uniform traps, none of the results presented in this thesis was unexpected from a theoretical point of view. The main value lies in adding the uniform trap to the toolkit of cold atom experiments. One of the major goals of the field of cold atoms is to establish a closer link to condensed matter problems. As those are inherently in a homogeneous environment, this work should be a substantial step towards this goal.

Appendix A

Polylogarithms and the Gamma function

Function Definitions

The polylogarithm function is defined as

$$g_\gamma[A] = \sum_{n=1}^{\infty} \frac{A^n}{n^\gamma}. \quad (\text{A.1})$$

The Gamma function is defined as

$$\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t} dt. \quad (\text{A.2})$$

Relation to the Bose distribution

A very simple case is

$$\frac{1}{A^{-1} - 1} = g_0(A). \quad (\text{A.3})$$

We can thus represent the Bose Distribution function as a polylogarithm

$$f(\mathbf{r}, \mathbf{p}) = \frac{1}{A^{-1} - 1} = g_0(A), \quad (\text{A.4})$$

where we set

$$A = \exp[(\mu - \epsilon(\mathbf{r}, \mathbf{p}))\beta]. \quad (\text{A.5})$$

Power-law integrals

We often have to solve the following integral:

$$\int_{-\infty}^{\infty} g_{\gamma} [Ae^{(-a|x|^s)}] dx = 2 \int_0^{\infty} g_{\gamma} [Ae^{(-ax^s)}] dx. \quad (\text{A.6})$$

Now we use the definition of the polylogarithm to obtain:

$$\int_{-\infty}^{\infty} g_{\gamma} [Ae^{(-a|x|^s)}] dx = 2 \int_0^{\infty} \sum_{n=1}^{\infty} \frac{A^n}{n^{\gamma}} e^{-anx^s} dx. \quad (\text{A.7})$$

The definition of the Gamma function allows us to evaluate the integral and collapse the result into a higher-order polylogarithm:

$$\int_{-\infty}^{\infty} g_{\gamma} [Ae^{(-a|x|^s)}] dx = 2 \sum_{n=1}^{\infty} \frac{A^n}{n^{\gamma}} (an)^{-1/s} \Gamma(1 + (1/s)) \quad (\text{A.8})$$

$$= 2\Gamma(1 + (1/s)) \sum_{n=1}^{\infty} \frac{A^n}{n^{\gamma+(1/s)}} (a)^{-1/s} \quad (\text{A.9})$$

$$= 2a^{-1/s} \Gamma(1 + (1/s)) g_{\gamma+(1/s)}(A). \quad (\text{A.10})$$

This can easily be generalized to include a power-law prefactor

$$\int_{-\infty}^{\infty} g_{\gamma} (Ae^{(-a|x|^s)}) x^p dx = \frac{2}{s} a^{[-(1+p)/s]} \Gamma\left(\frac{1+p}{s}\right) g_{\{\gamma+\frac{1+p}{s}\}}(A). \quad (\text{A.11})$$

Appendix B

IGBT driver

We rely on an IGBT (insulated-gate bipolar transistors) to switch the high currents needed to operate our quadrupole magnet. Crucially, it can also tolerate the inductive voltage spikes created during switch-on and switch-off of the coils.

IGBTs require drivers to (de-)charge the gate of the IGBT to open or close the circuit. Further, these drivers need to have the capability to interface with security interlocks (such as temperature sensors or flow meters) and shut down the IGBT under dangerous conditions. The driver design we used in the experiments in this thesis are described in [27].

Here, we are going to describe an improved design, which is functionally nearly identical, but uses significantly fewer components. This makes it possible to accommodate the whole driver board on top of the IGBT. Also, assembly time and cost are decreased significantly. The system consists of three major components (see Fig. B.1): The “IGBT driver TS V4” is designed to slot onto the matching “contact board V4”, which is bolted onto a Semikron Semix IGBT module (here SEMIX402GB066HDs). The design has been tested and used to drive the superconducting coil in Appendix C.

B.1 Functionality

The driver state is set with a TTL-compatible 5 V input. Depending on the applied voltage, the IGBT opens or closes. The driver also provides four multi-purpose error inputs that can be used e.g. for interfacing with temperature monitors or interlocks. One of those inputs is directly connected to the integrated temperature sensor or the Semix module, preventing overheating of the IGBT. The driver latches-up on error, shutting down the driver until manual reset.

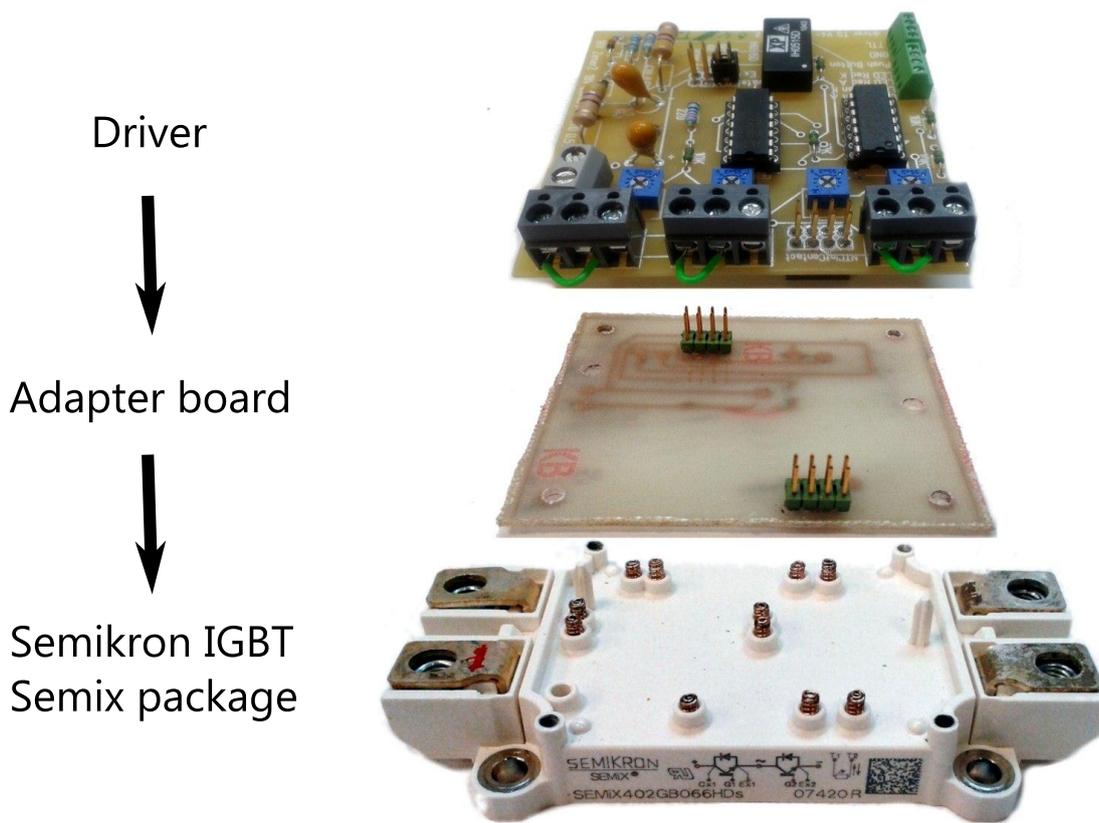


Figure B.1: Assembly of an IGBT, contact board and driver.

B.2 Inputs

Power The driver requires 5V DC ($I \approx 500$ mA), which need to be applied to the “5V in” and the “Gnd in” pins of the dual screw terminal (see Fig. B.6). A cheap plug-in-the-wall supply is good enough, but the power supply should not be shared with other devices. The power is internally converted and regulated to provide ± 15 V to the driver stage.

TTL The ground and signal wires of the TTL are connected to the respective “GND” and “TTL” screw terminals. The ground is shared with the shared with the 5 V input power ground, so only floating ground TTL channels (e.g. isolated with an optocoupler box) should be used. The accepted signal levels are approximately compatible with the TTL standard.

Multi-purpose error inputs There are four multi-purpose error inputs, which can be used to detect a range of errors and shut down the IGBT. They are simple voltage bridges, where the middle terminal voltage is compared against a reference voltage, which is set by a simple adjustable resistor (trimpot). If the reference voltage is lower than the voltage

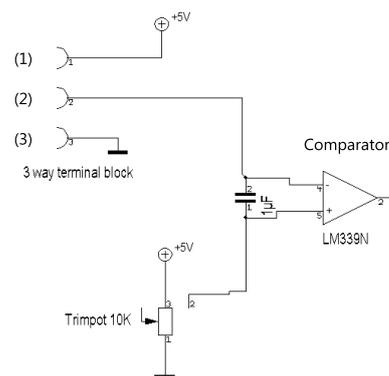


Figure B.2: The input stage of an error input.

on terminal block (2), the IGBT will detect an error and shut down. If any of the error inputs detects an error, a latch shuts down the IGBT and a manual reset is required to resume operation. Capacitors together with the (external) resistors act as a low-pass filter, rejecting noise on the input of the comparators. One of the inputs is already connect to the internal PTC of the IGBT via the contact board. Fig. B.3 explains the use of the error inputs.

Reset button Connect to the “GND” (shared with TTL) and the “Push Button” terminals. The error-latch resets, if the two terminals are short-circuited. It works best using a “push-to-make” button.

Trimpot adjustment Under non-error conditions, short-circuit the reset terminals (e.g. by pressing the button continuously) and adjust the trimpot in question until the red LED just goes out. Note that all the others error terminals need to be high (i.e. no other errors); if in doubt short-circuit their terminals (2) and (3), or turn their trimpots to 5 V.

B.3 Outputs

LEDs The red LED illuminates if the IGBT is shut down due to an error and the latch is activated. A $500\ \Omega$ resistor in line is recommended, as the red LED pin is driven by 5 V. Alternatively, the pin could be used provide a TTL error signal (5 V on error), e.g. to shut down a power supply. The pin is not buffered, so not more than 10 mA should be drawn and external buffering and/or isolation with a relay is recommended.

A green LED illuminates, if the IGBT is in the low-ohmic state. **THE LED NEEDS TO BE CONNECTED** for the driver to work. No resistor should be used in line.

The terminal names “LED A/K” refer to the anode and cathode of the LEDs, respectively.

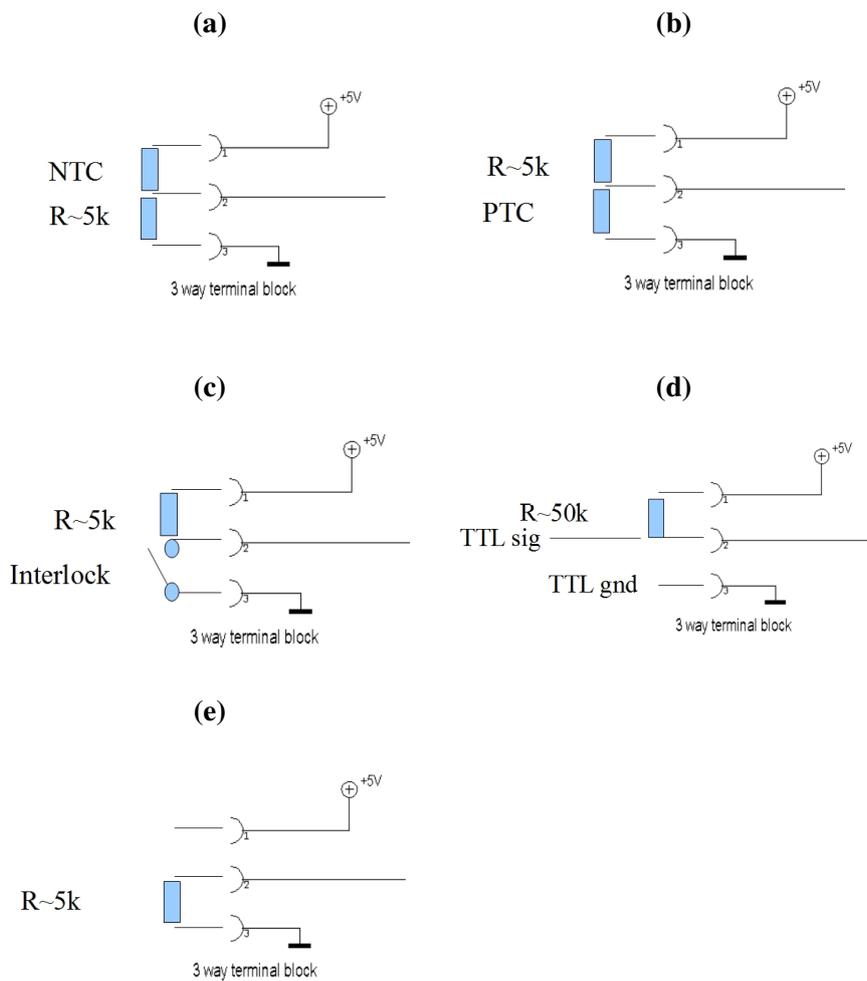


Figure B.3: Suggested uses of the multi-purpose error inputs. (a) NTC thermistor. Assume the negative temperature coefficient (NTC) thermistor has a resistance of $5k\ \Omega$ at room temperature and $500\ \Omega$ when hot. Therefore terminal (2) will be at 2.5 V at room temperature and at 4.5 V when hot. If the reference voltage is set to 3.5 V, a hot NTC will produce an error and shut down the IGBT. The resistor R should be chosen to have a similar value to the NTC under normal conditions so that the reference voltage is somewhere in the middle of the 5V range. Also, it should be not much smaller than a couple of hundred Ohms to prevent excessive current flowing between terminals (1) and (3). (b) PTC thermistor. Use as in the NTC case, but swap the resistor and the thermistor. (c) Interlock. Set V_{ref} to 2.5 V. If the interlock is closed: Terminal (2) is low, no error. If the interlock is open: Terminal (2) is high, an error is detected. Swap interlock and R to invert behaviour. (d) (Floating) TTL. Not tested or recommended. This might even work if swapping the TTL signal and ground. Use a high Pull-up resistor (a floating TTL might already have a pull-up/pull down resistor somewhere, which might limit their current sourcing abilities). (e) Terminal not used. To prevent random charge oscillations, pull down terminal (2) to ground in all unused error input terminals. Set the corresponding trimpots to approximately 2.5 V.

B.4 Gate charging and protection

Gate power supply A ± 15 V DC/DC converter powers the gate driver power supply. The gate driver IC is optically isolated from the inputs. There are power resistors in the line to the gate driver power supplies, limiting the DC current to the gate to a couple of mA. This protects against overheating if the IGBT breaks so that uncontrollable current can flow from the broken gate to the emitter. However, the driver power pins are buffered with 10 μ F capacitors so that the driver can provide 1-2 A peak-current during typical gate charging times. While it is perfectly capable of driving the gate at low switching frequencies, a high switching frequency (> 100 Hz) should not be attempted.

Gate charging The gate is charged to around +13 V (There is a small voltage drop in the driver chip. The voltage can be increased by **carefully** increasing the “5 V” power supply voltage). If the driver is switched off, the gate is charged to -15 V. This helps to discharge the gate to below 0 V quickly, decreasing switching times and reducing heating of the IGBT (negative gate drive). Peak currents are around 1 to 2 A, which helps fast switching and reduces switching power losses. Typical gate charges are 1 μ C. The maximum gate current can be limited by the changing the gate resistor, which is per default 10 Ω (the internal gate resistance of the IGBT is typically a couple of Ω). It can be omitted to speed up switching, but that is usually not necessary. Do not increase the gate resistor beyond 50 Ω , as the IGBT might be damaged due to slow switching. If more information on switching speed is needed, check the homepage of Semikron for the IGBT application notes.

Gate protection The driver chip has (depending on model though) an undervoltage lockout, that shuts down the driver if the supply voltages is too low to drive the gate properly. Two back-to-back 18 V Zener diodes limit the maximum voltage on the gate to ± 20 V. A 50 k Ω pull-down resistor between gate and emitter prevents the gate to float to random voltages if it is not powered. The resistors in the power line to the gate chip prevent it from overheating incase the IGBT breaks and excessive gate-emitter current flows.

Gate selection The semix module has two IGBT modules and therefore also two gates controlling the modules. The driver can only drive one module at a time. To select the appropriate gate/emitter combination, move the jumpers on the board to connect the “driver out” pins (on the far left/ far right) to the IGBT gate, as shown in Fig. B.4. **NEVER** short Gate and Emitter vertically.

B.4. GATE CHARGING AND PROTECTION

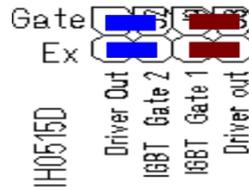


Figure B.4: Selection of IGBT module. The correct jumper setting to select module (1) is shown in red; for module (2) in blue.

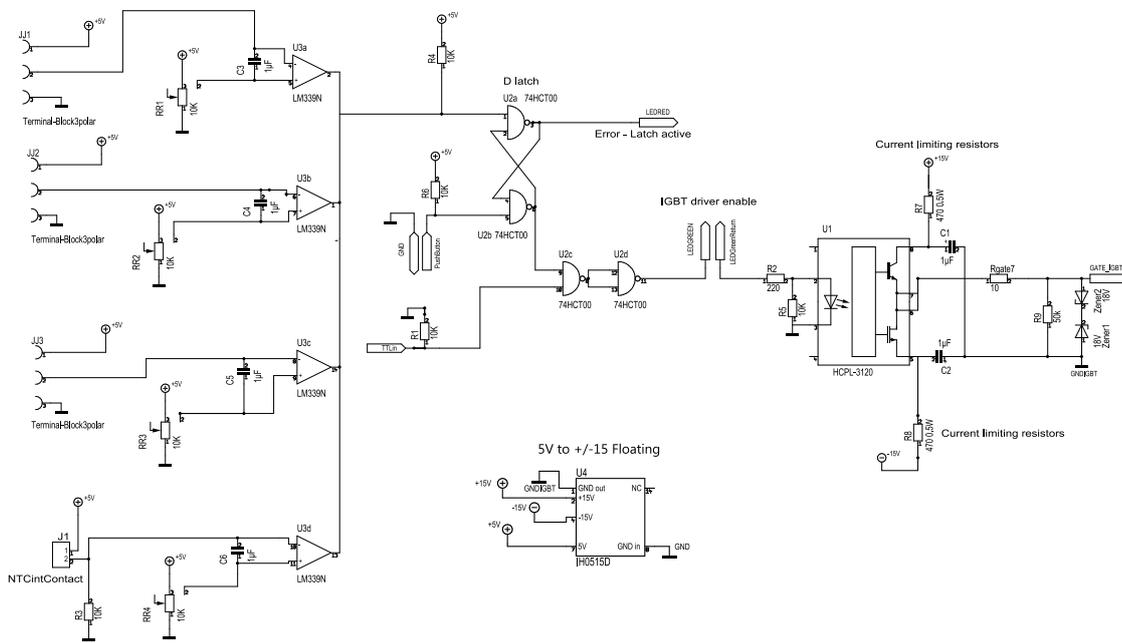


Figure B.5: Schematic of the Analog/Digital circuit of the driver.

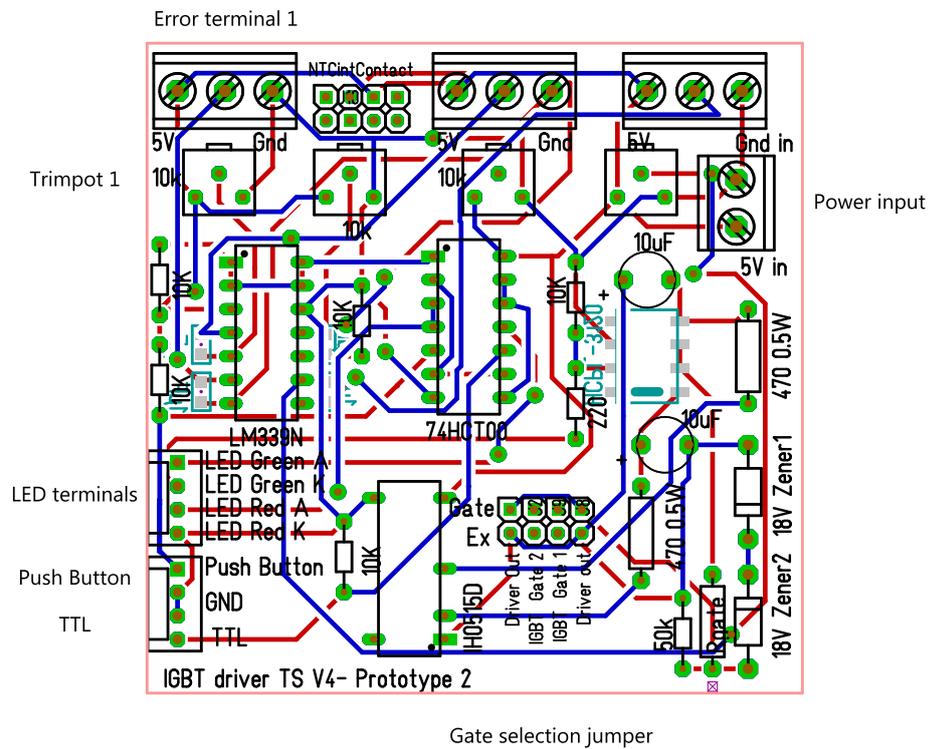


Figure B.6: Circuit board design of the Driver. Red is the top copper layer, blue the bottom layer. Most components are mounted on the top (black), only a few are mounted on the bottom (turquoise).

Appendix C

Design of a high-temperature superconducting magnet

C.1 Introduction

In this appendix, we describe a prototype of a high-temperature superconducting (HTS) trapping magnet. While not relevant to the rest of this thesis, we included it here for documentation purposes.

We initiated this project hoping to achieve higher current densities than in conventional water-cooled copper coils. This would allow us to reduce the size of the magnets or to increase the spacing between the individual coils. Both measures would increase optical access to the science cell between the coils.

While we were able to construct a functional prototype, the increase in performance was not large enough to justify the increased complexity and cost. We therefore did not use it in our experiments.

The first prototype coil was built as my diploma thesis project [94]. The second prototype, which is described here, was then constructed during the first year of my PhD. The following report is based on [95].

C.1.1 Trapping fields

High-strength quadrupole magnets are essential to achieve trapping of atoms during evaporative cooling. They are usually produced by a coil pair in Anti-Helmholtz configuration as shown in figure C.1. We will evaluate how to calculate their fields and consider the implications for coil design.

The field of a current loop (radius a) on the z -axis can be obtained by integrating the

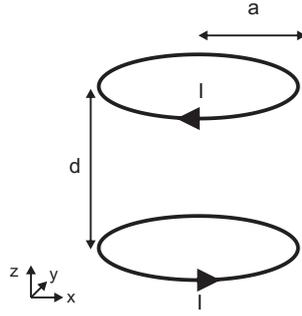


Figure C.1: Two current loops in Anti-Helmholtz configuration.

Biot-Savart law:

$$\mathbf{B} = \frac{\mu_0 I}{2} \frac{a^2}{(a^2 + z^2)^{\frac{3}{2}}} \mathbf{e}_z. \quad (\text{C.1})$$

Differentiating and adding a second coil at distance d yields the following gradient at $\mathbf{B}(z = 0) = 0$ in the centre between the coils:

$$\frac{dB_z}{dz} = -\frac{3\mu_0 I}{2} \frac{a^2 d}{(a^2 + d^2/4)^{\frac{5}{2}}}. \quad (\text{C.2})$$

Using $\nabla \mathbf{B} = 0$ and the rotational symmetry of the system around the z -axis, we obtain

$$\frac{dB}{dx} = \frac{dB}{dy} = -\frac{1}{2} \frac{dB}{dz}. \quad (\text{C.3})$$

The z -axis gradient is usually referred to as the strong axis gradient (as opposed to the weak axis gradient in the x - y -plane). We will use this gradient as a benchmark for coil comparison.

It can be seen that the field gets stronger for smaller d and a . Hence, the coils need to be placed close to the atoms to obtain a sufficiently high gradient. For a given radius, the maximum gradient is given for $d = a$. In this case equation C.2 simplifies to:

$$\frac{dB}{dz} = -0.86\mu_0 I a^{-2}. \quad (\text{C.4})$$

In this simple model, doubling the current allows us to extend the coil spacing by a factor 1.4 while keeping the gradient constant. In a real magnet, we have to take into account the width and height of the coils. This can be done numerically by adding up the contributions of many thin loops. From Eq. (C.2), we can see now that stacking windings provides only limited benefit. Each new winding is one wire cross section further away and with diminishing contribution to the field. Building stronger magnets buy stacking up windings leads to considerable bulk. Therefore, increasing the current density is the only viable way to gain a significant amount of optical access.

C.1.2 Water-cooled Magnets

It is impossible to build conventional high field magnets without proper cooling; without it, the electric heating $P = I^2R$ quickly turns the magnet into a piece of modern art. Generally, hollow copper tubes are used to form the windings of the coil and water is pumped through at high pressure, cooling the magnet.

The pressure drop Δp in a pipe in the turbulent regime is according to the Darcy-Weissbach equation

$$\Delta p = 0.5R_{fl}Q^2, \quad (\text{C.5})$$

where R_{fl} is the flow resistance of the pipe and Q is the volumetric flow rate. The exact expression for R_{fl} is rather complicated, but it is approximately proportional to $L/A_{fl}^{1/2}$, where L is the length of the pipe and A_{fl} its cross section. Assuming that the heat transfer from the copper to the water is perfect, the cooling power is

$$P_{cool} = QC\delta[T_{coil} - T_{H_2O}], \quad (\text{C.6})$$

where C is the heat capacity of water, δ its density and $T_{coil} - T_{H_2O}$ the temperature difference between coil and water. Equating heating and cooling gives us:

$$\Delta p \propto \frac{L^3}{A_{fl}^{0.5}} \frac{I^4}{A_{Cu}^2}, \quad (\text{C.7})$$

where A_{Cu} is the copper area of the tube.

We can see that increasing the pressure allows only moderate current gains. Running the cooling water through parallel layers helps quite a lot, however at the cost of increasing bulk due to the additional connections. Ultimately, current densities are capped by technically viable pressures. In our lab, the total current density $I/(A_{fl} + A_{cu})$ is around 1.5 kA/cm^2 at 2 bar. The extreme end of the scale is used in the labs in Innsbruck [96], with up to 12.8 kA/cm^2 at 35 bar. One should note that a 35 bar cooling system requires a lot of infrastructure and does not come cheap.

C.1.3 HTS Magnet requirements

The currently best commercially available HTS is YBCO (Yttrium barium copper oxide). We bought the superconducting tape for the prototype coil from American Superconductor. We used a tape with a critical current of 90A and a cross section of $4 \text{ mm} \times 0.2 \text{ mm}$. This corresponds to a current density of 11.3 kA/cm^2 at 77K (liquid nitrogen temperature). This is comparable to the best copper cooled systems. However, this is not the full truth. First of all, the self-field of a coil reduces the critical current of the superconductor. In ad-

dition to that, the cooling structure and insulation around the coil are going to reduce the effective current density. On the plus side, lowering the temperature below 77K increases the critical current drastically. Taking all these factors into account, one can currently not expect to surpass water-cooled designs in terms of current densities. However, expecting future progress in the field of HTS made us explore their use in cold atom experiments.

A good design of a HTS coil system for magnetic trappings needs to take care of the following things:

- **Cooling:** Making sure that the superconductor stays cold, and that the outside of the cryostat remains at room temperature to prevent condensation of air humidity.
- **Optimizing the geometry:** The gradients need to be sufficiently large, while reducing bulk to allow optical access to the cell with the atoms.
- **Quench protection:** Superconductors have to be protected against damage from overheating.
- **Fast ramping:** Ramping of the coil needs to be fast enough for initial trapping of the atoms and for time-of-flight switch-off.

These requirements will be discussed in the following sections.

C.2 Cryogenics

The cryogenic requirements in high-temperature superconductors (HTS) are much less stringent than for low-temperature superconductors. Pouring liquid nitrogen over a HTS coil sends it quickly superconducting state. During my diploma thesis we built a cryostat with conventional insulation instead of the customary vacuum insulation. Since nitrogen is cheap, we were not concerned about the extra consumption. However, we realized that the biggest challenge is not to get the liquid nitrogen into the cryostat, but to get the nitrogen gas out of it (expansion ratio ≈ 200). Evaporation rates were around 100 W and bubbles formed in the cryostat. The created gas layer insulated the HTS against the liquid nitrogen severely encumbering cooling. This made the operation of this prototype extremely unreliable and not suitable for experiments. Additionally, the outside of the cryostat got cold as well, leading to condensation air humidity.

For these reasons, we decided to use vacuum insulation for the prototype presented in this report. Vacuum insulation renders direct heat conduction losses through the walls negligible (≈ 1 W), solving the above problems. The other important heat leak (≈ 10 W) is caused by the non-superconducting current leads connecting the coil to the outside

world. However, this only causes evaporation in the pipes leading to the cryostat, which is non-critical to the coil stability.

As manual topping up the coil with liquid nitrogen is not practical in the long run, an automated system has been designed.

C.2.1 Current leads

Metal current leads connect the HTS coil to the outside world. They need to carry the current from the point where superconductivity ceases (around 77 K) to the outside world (300 K). The heat load caused by the leads is the sum of electrical heating and heat conduction. The Wiedemann-Franz Law (WFL) relates resistivity $\rho(T)$ and heat conductivity $\lambda(T)$ for metals as:

$$\rho(T)\lambda(T) = L_0T, \quad (\text{C.8})$$

where $L_0 = 2.44 \times 10^{-8} \text{W}\Omega\text{K}^{-2}$ is the Lorentz number. If we consider a current lead with diameter A , length l carrying a current I , the heat flux \dot{Q} through an element dx along the superconductor is given by

$$\dot{Q} = \lambda(T)A \frac{dT}{dx} + I^2 \frac{\rho(T)}{A} dx = 0. \quad (\text{C.9})$$

Solving this equation with the boundary conditions $T_l = 77 \text{ K}$ and $T_h = 300 \text{ K}$ yields a minimum heat load P_{opt} for a given current (see [97] for the calculation in Helium cryostats, and [94] for HTS). P_{opt} is approximately independent of the metal used :

$$P_{opt} = I \frac{\sqrt{L_0} (T_h^2 - T_l^2)}{\sqrt{T_h^2 - T_l^2}}. \quad (\text{C.10})$$

This heat load is only achieved for an optimum length to area ratio f , which depends on the metal as well as T_l and T_h :

$$\frac{Il}{A} = f(\lambda, T_h, T_l). \quad (\text{C.11})$$

The resulting values are tabulated in Tab. C.1.

For the critical current in the prototype coil of around 55 A, this yields a heat flux of around 4.6 W for two leads. The heat load can be reduced by a factor of two by vapour-cooling the leads with the evaporating nitrogen (see Tab. C.1). We paralleled the leads with HTS tape until the filling pipe, so that the heat load is localized in the filling pipe and does not cause bubbles in the cryostat. Therefore it is not critical for superconductivity stability. It might pay off to make the leads thinner than the optimum steady-state ratio $I_{max}l/A$, as the coil current has typically a small duty cycle ratio in

		Copper	Brass	Stainless Steel
Conduction cooled				
Numerical				
$(IL/A)_{opt}$	[A/m]	3.5×10^6	6.5×10^6	1×10^5
$(P/I)_{opt}$	[10^{-3} W/A]	42.7	45.5	45.5
Semi-analytical				
$(IL/A)_{opt}$	[A/m]	3.5×10^6	5.7×10^5	1.1×10^5
$(P/I)_{opt}$	[10^{-3} W/A]	45.3	45.3	45.3
Vapour cooled				
Numerical				
$(IL/A)_{opt}$	[A/m]	5×10^6	8×10^6	1.5×10^5
$(P/I)_{opt}$	[10^{-3} W/A]	23.2	25	25.5

Table C.1: Comparison of the optimal design parameters for different materials and different sources. Semi-analytical results with thermal conductivities taken from [98] and [99] were done based on the calculations above. They are compared to the results from [100] which take deviations from the WFL into account. The values for vapor cooling are also taken from this reference. Both calculations agree very well, except for brass which can be explained by the variety of different brass alloys (unfortunately the alloy type is not specified in [100]) and a lack of data points for its thermal conductivity. The temperatures used were the boiling temperature of nitrogen $T_l = 77$ K and room temperature $T_h = 300$ K.

cold atoms experiments. However, the deviation should not be larger than a factor of two, otherwise the leads melt during continuous operation [100].

C.2.2 Insulation

After the poor results using polyurethane foam as well as with aerogel, we decided to use vacuum insulation. Vacuum insulation can be made rather compact (e.g. in thermosflasks) if designed carefully. We will show that the vacuum requirements to reduce insulation losses below the losses in the current leads are not very demanding (1×10^{-3} mbar).

Our cryostat, shown in Fig. C.2), features a doubled wall vacuum space. Lids are needed to insert the coils, while indium seals are used to seal the gaps. Multi-layer insulation (MLI) is used to shield to radiation losses.

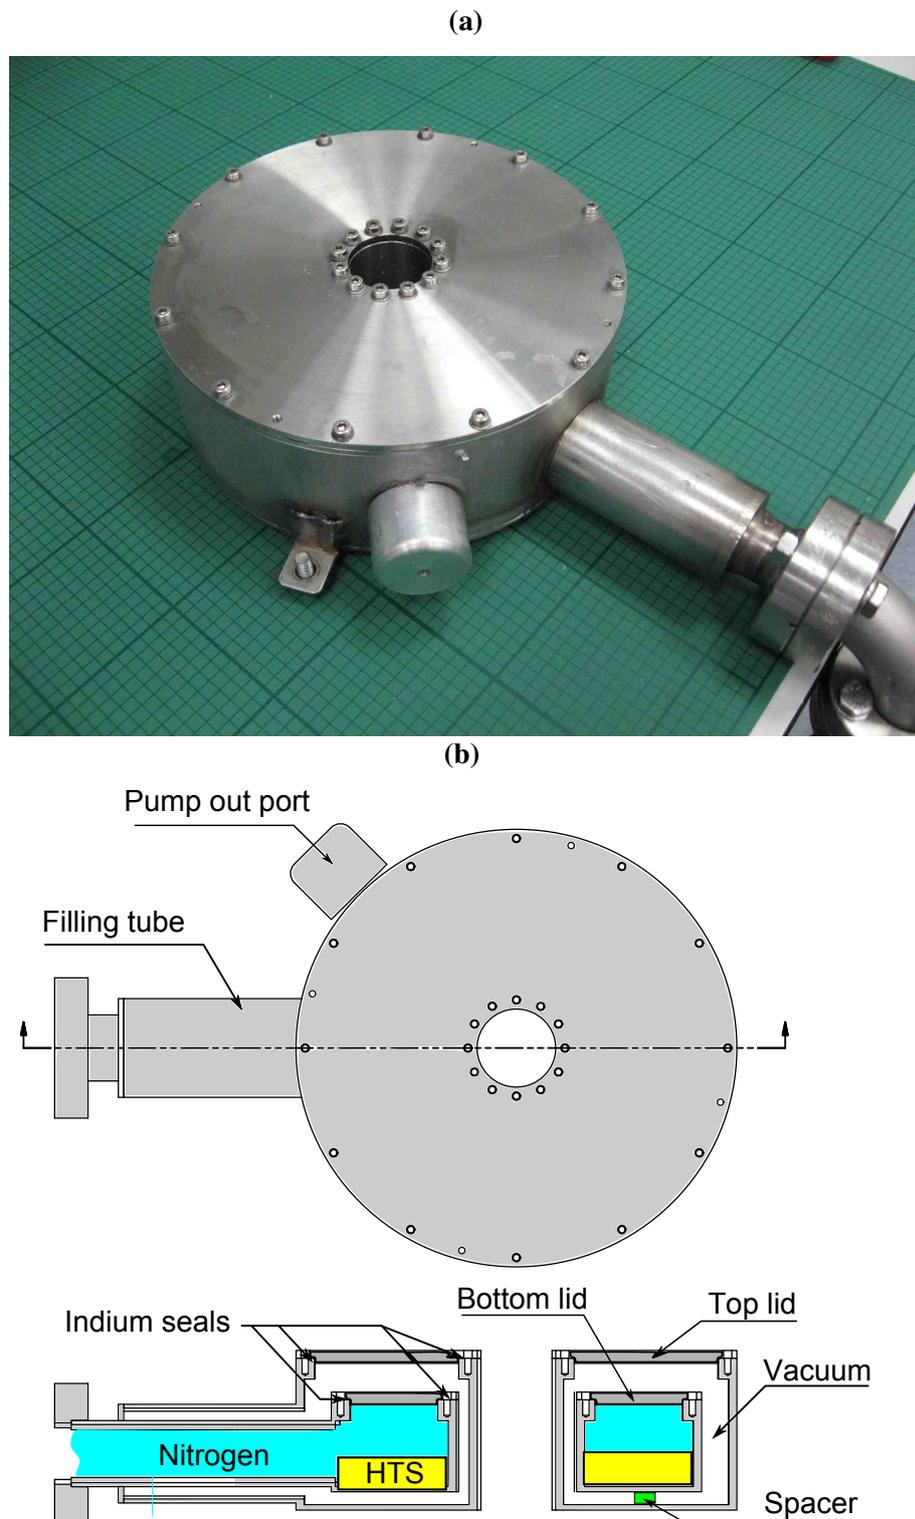


Figure C.2: (a) A photo of the prototype. (b) Schematic cross section of the prototype.

Theoretical predictions

Mean free path Heat conduction through a gas falls either into a low pressure or a high pressure regime, depending on whether the mean free path λ of the particles is smaller or bigger than the distance between the walls [99]. The mean free path is the average distance travelled between collisions of the particles. It can be estimated as

$$\lambda = \frac{kT}{\sqrt{2}\sigma p}, \tag{C.12}$$

where T is temperature, p the pressure and σ the collision cross section of the gas molecules. The cross section is related to the diameter as $\sigma = \pi d^2$. Air mainly consists of N_2 and O_2 , which both have two atomic shells. The Bohr model yields an estimated diameter of 2×10^{-10} m for each atom. Taking the second atom in the molecule into account, one can estimate $d = 3 \times 10^{-10}$ m [101]. The resulting λ for room temperature is plotted in figure C.3. At pressures of around 1×10^{-3} mbar, we can be sure that λ is larger than the spacing between the walls.

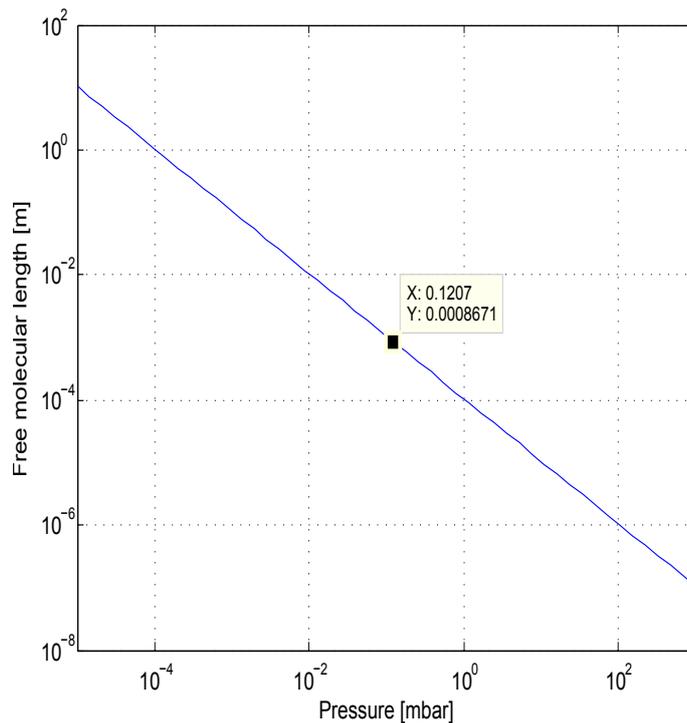


Figure C.3: The mean free path of an average air molecule depending on pressure at 300 K. The crossover between the high pressure regime C.13 and the low pressure regime C.14 occurs when λ is comparable to the wall spacing, which is 5 mm for the prototype coil.

High pressure At high pressures (such as atmospheric pressure), λ is shorter than the distance between the walls. In this case the heat transfer between two plates (a good

approximation for a double walled cryostat) \dot{Q} can be calculated as

$$\dot{Q} = \lambda \Delta T A / d, \quad (\text{C.13})$$

where $\lambda \approx 0.025$ W/m K for air [99]. A is the area between the plates and d their separation, ΔT is the temperature difference. The results for our cryostat are plotted in Fig. C.4. It can be seen that the high pressure regime leads to unacceptably high heat losses.

Low pressure We therefore need to operate in the low pressure regime ($p < 1 \times 10^{-3}$ mbar) to achieve good enough insulation. There, the heat flux depends only on pressure:

$$\dot{Q} = k \Delta T p A, \quad (\text{C.14})$$

where k is a gas and geometry dependent constant on the order of one [99]. Below 1×10^{-3} mbar, our heat load is clearly negligible compared to the current lead contribution (see Fig. C.5).

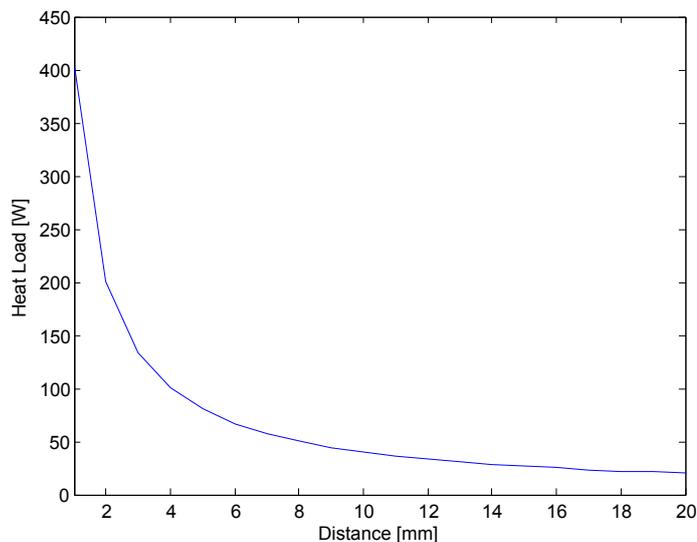


Figure C.4: Heat flux in the high pressure regime, calculated for $\Delta T = 230$ K and a cryostat area $A = 0.07$ m². Note that the flux is not pressure dependent. This is a rather optimistic calculation, since convection currents are likely to increase the heat transfer. To reduce the heat load below 10 W, a huge air gap would be needed.

Radiative Losses In the prototype cryostat, vacuum pressures of 1×10^{-5} mbar were easily obtained using a turbo and scroll pump combination. At these pressures, we are clearly in the free molecular regime and the heat load transmitted by particles is basically zero. However, radiative heat transfer becomes relevant. We estimate it with the Stefan-

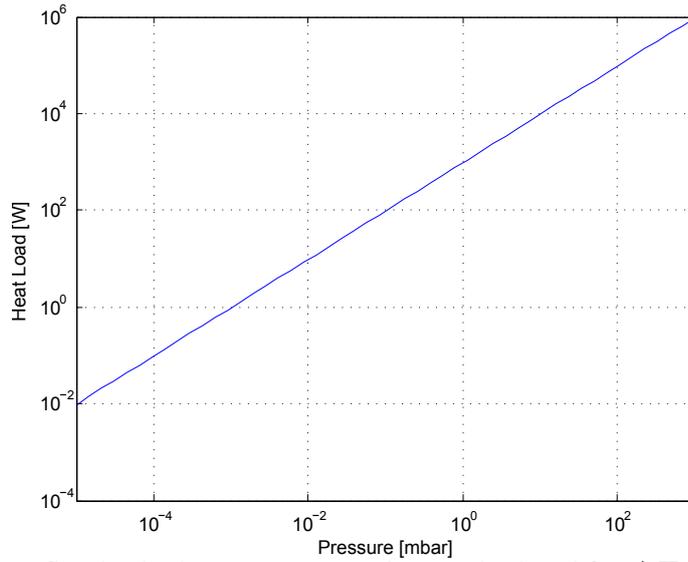


Figure C.5: Heat flux in the low pressure regime, calculated for $\Delta T = 230$ K and the area of the cryostat $A = 0.07$ m². At pressures below $p < 1 \times 10^{-3}$ mbar it becomes negligible.

Boltzman equation:

$$q_r = \epsilon_r \sigma T^4, \quad (\text{C.15})$$

where q_r is the heat flux from a surface at Temperature T , $\sigma = 5.67 \times 10^{-8}$ is the Stefan-Boltzman constant and ϵ_r is the emissivity. For the flux between two walls at different temperatures, an effective emissivity

$$\epsilon_{eff} = \frac{\epsilon_{cold} \epsilon_{warm}}{\epsilon_{cold} + \epsilon_{warm} - \epsilon_{cold} * \epsilon_{warm}} \quad (\text{C.16})$$

can be derived. The flux between two walls is then [102]:

$$q_r = \epsilon_{eff} \sigma (T_{warm}^4 - T_{cold}^4). \quad (\text{C.17})$$

[102] gives tabulated values for ϵ_{eff} and q_r . For polished stainless steel, $\epsilon_{eff} = 0.12$ and $q_r = 55$ W/m² for surfaces at room and liquid nitrogen temperatures. Using the approximate area of our cryostat, $A = 0.07$ m², this causes a radiative heat load of approximately 4 W. This heat load can be reduced further by inserting multi-layered insulation (MLI) between the walls. MLI is a very thin aluminium foil coated with plastic that acts as floating radiation shield and is readily available from cryogenic suppliers. Using flux conservation, we can show that N layers of MLI reduce the heat load to

$$q_r = \frac{\epsilon_{al}}{N + 1} \sigma (T_{warm}^4 - T_{cold}^4), \quad (\text{C.18})$$

where ϵ_{al} is the emissivity of aluminium, which is similar to stainless steel. This effect can be used to get a good emissivity even when the cryostat surface is not well polished. More importantly, the heat load is reduced by more than a factor of N . We used 5 layers, reducing the radiative heat load below 1W.

Spacer and connection pipe losses There are also conductive heat leaks from the places where the inner cryostat wall is connected to the outer wall. These heat leaks \dot{q} can be calculated as

$$\dot{q} = A/L \int_{T_1}^{T_2} \lambda(T) dT, \quad (\text{C.19})$$

where A is the cross-section, $\lambda(T)$ the thermal conductivity at temperature T and L the length and T_i are 77 K and 300 K respectively. Values for the integral are tabulated in [99].

In the cryostat, we have several connections where losses occur. The stainless steel pipe leading to the outside induces losses of only 0.5 W, as stainless steel is a bad conductor and the pipe is fairly thin. Three spacers between the walls prevent movement of the inner cryostat. Using PTFE as a material, we expect another 0.5 W heat load.

Experimental results

To measure the total heat leak, we filled the cryostat with liquid nitrogen. We measure the time until complete evaporation of the nitrogen, as well as monitoring the flow rate of the evaporated gas. Both of these measurements allow us to estimate the heat leak.

Vacuum The vacuum insulated cryostat (no MLI) filled with 0.3 L of liquid nitrogen lasted 2h at an average gas flow rate of 1.5 l/min. [99] lists the liquid (2.25×10^{-2} L/h W) and the gas volume boil-off rates (0.243 L/min W) for liquid nitrogen. Our measurements correspond to heat leaks of 6.7 W and 6.2 W, respectively. Theoretically, we expect a heat leak of 4 W from radiative losses plus 1 W from spacers/pipe conduction, which is compatible with our observations.

Air gap The measurement was repeated with an air-filled vacuum gap, measuring heat loads of around 40 W. Theoretical estimates using Eq. (C.13) range between 40W-80W. The estimate depends strongly on which temperatures are assumed. This is difficult as the outer wall is cooling down over time. Also the heat conductivity of air varies considerably from source to source. Given the uncertainties, this is a rather good agreement; it certainly shows that air gap insulation is not practical.

Conclusions It is certainly advisable to test the vacuum before inserting the coil with a helium leak tester. On the first test, this revealed a huge leak (actually too big for the helium leak tester to resolve). Over-pressuring the cryostat and floating it in ethanol revealed a bubbling gap where a weld connection was not done properly. After mending the gap, the vacuum worked great and no measurable leaks were found with the leak tester.

We can conclude that the vacuum insulation performs as expected. The problems with the previous non-vacuum insulated prototypes are resolved. The consumption of liquid nitrogen is reduced by an order of magnitude (from 100W to below 10 W, corresponding to less than 1 l/h) and we found that the coils can be run continuously for hours at critical current.

C.2.3 Filling system

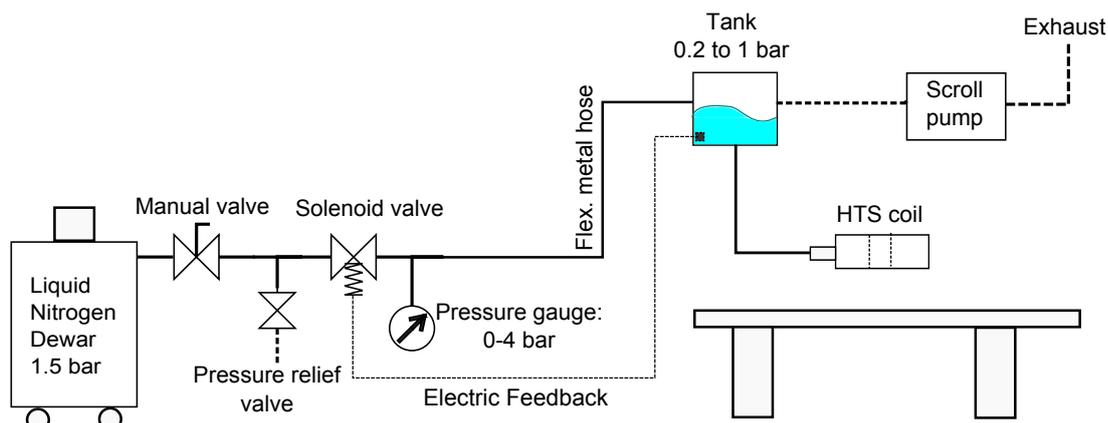


Figure C.6: Liquid nitrogen is pumped into the tank, from where gravity feeds it directly into the coils. Electric feedback keeps the nitrogen level constant. Pumping down the pressure of the gas phase above the liquid in the tank allows to reduce the temperature of the bath.

We built a fully-automated liquid nitrogen filling system (see Fig. C.6). Liquid nitrogen is pumped from the dewar through transfer lines into the top filling tank. From there, it is fed by gravity into the HTS cryostat. The evaporating gas is vented into an exhaust to prevent low air oxygen levels. For safety reasons, an oxygen detector is installed in the proximity of the coil. Connections to the cryostat are made with vacuum CF-Flanges. Standard BSP-threaded fittings also work sufficiently well when enough PTFE tape and excessive amounts of torque are used.

Nitrogen level sensors Home-made nitrogen sensors in the tank sense the filling level. They consist of PTC thermistors (B59721A0070A062, Epcos), which we found exhibit

a 1000-fold increase of resistance between room and liquid nitrogen temperature. They are heated by $20\ \Omega$ resistors with around $0.05\ \text{W}$ each. As the cooling power of liquid nitrogen is much higher than of nitrogen gas, this leads to a steep change in resistance of the PTC, depending if it is below or above the liquid level. A similar method with carbon resistors is described in [103]. In [94], we introduced the dual Resistor-Thermistor system to increase the resistance contrast ratio. Both PTC and resistors are surface mount components, taking up minimal space as shown in Fig. C.7.



Figure C.7: Composite PTC-R nitrogen level sensor covered in epoxy.

Signal processing The sensor signals are processed by a Max32 Chipkit microcontroller (compatible to the Arduino), which is programmed in C++. It features digital I/O channels as well as 10-bit analogue inputs. The analogue sensor readings are converted into digital level signals. These level readings are processed and control the nitrogen flow via an electric valve. The Max32 has pre-built libraries for LCD displays, allowing us to create a user-friendly visual interface. The control box is shown in Fig. C.8. The filling system is fully automated. The filling behaviour can be reprogrammed in-place via a USB port.

Vapor pressure pumping We found that inserting a scroll pump into the vent line speeds up filling due to the increase in differential pressure between the tank and cryostat. More importantly, it lowers the vapour pressure of the liquid nitrogen. Lowering the vapour pressure reduces the temperature of the nitrogen and thus increases the critical current of the coil. The pressure could be pumped down all the way to $0.2\ \text{bar}$, which is close to the triple point of nitrogen, where it solidifies at a temperature of $67\ \text{K}$. This yielded an increase of the critical current from $55\ \text{A}$ to $85\ \text{A}$.

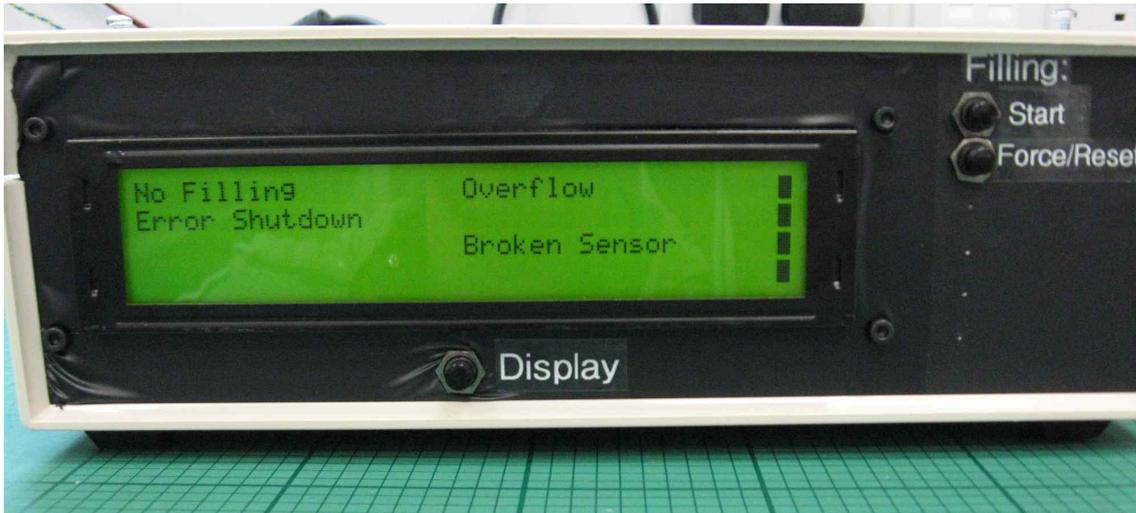


Figure C.8: The control interface of the filling system. It controls the filling valve depending on the nitrogen level and shuts down the magnet if a fault is detected. As the sensors are disconnected in the picture, several error messages are displayed. The user can also manually control the filling process.

C.3 Geometric design and fields

Geometry of the coil The YBCO tape is a flat 4 mm by 0.4 mm tape that is sold in unspliced lengths of around 30 m. The main parameters for coils are their inner and outer radius and the number of layers. Since splicing different tapes together is difficult and can damage the superconductor, the most common shape are coil stacks that are much wider than high ('pancake coils'). The main considerations when choosing the design parameters are:

- To create a large, uniform magnetic field gradient near the atoms.
- To reduce the size of the cryostat.
- To minimize the self-field, i.e. the magnetic field in the superconductor.

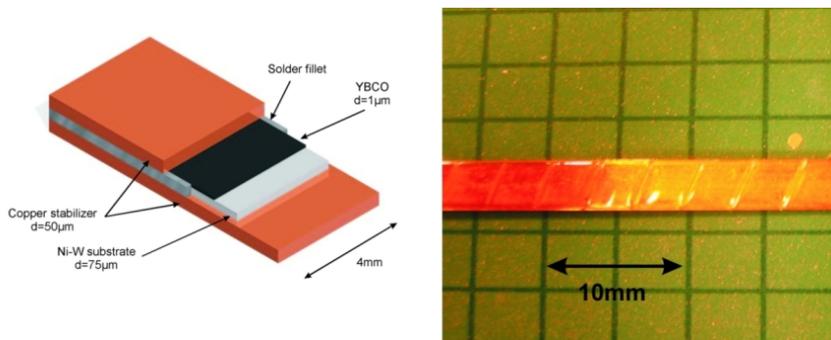


Figure C.9: A photo of YBCO tape next to a schematic cross section. The actual superconductor constitutes only small part of the tape.

The first two aims are obvious; minimizing the self-field not so much. As any superconductor, YBCO has a critical Temperature-Magnetic Field-Current surface. On one side of this surface, YBCO is superconducting; on the other side it is not. Crucially, a high self-field reduces the critical current and thus the field at the atoms. These requirements cannot be maximized at the same time, so a sensible compromise is needed. To calculate the fields for a given configuration, we used calculated the fields for a range of pancake coil configurations by decomposing them into current loops and adding up their contributions using elliptical integrals [104]. By determining the maximum self-field in the coil, we can estimate the critical current from Fig. C.11. An important result from these calculation is that more than two coil layers provide very limited benefits, as the increase in field per current is eaten up by the self-field reduction of the critical current.

Taken these considerations into account, we designed a cryostat with an inner radius of 26 mm, an outer radius of 67 mm and two layers of HTS rated at 90 A . A double wall allows for vacuum insulation. The lids are sealed with 1mm indium wire (0.7 mm turned out to be too thin). Indium seals seal well at cryogenic temperatures, do not take up much space and are not very difficult to use. U-formed PTFE spacers were used to clamp the coils down. A pump out port (similar to the ones on nitrogen dewars) allows us to create the vacuum between the walls. Detailed drawings can be found at the end of this report.

Critical Currents HTS are much more susceptible to parallel fields than to perpendicular ones [105]. Using our numerical simulation of the coils, we obtain a maximum parallel self-field of around 46 G/A. Combining this with the manufacturer's plot of the self-field dependence of the critical current, we estimate a critical current of 45A (see Fig. C.11). Experimentally, we observed 55-60A. Pumping the vapour pressure of the nitrogen close to the triple point down to 0.2 bar (≈ 65 K) raised the current to around 85-90 A. Fig. C.11 predicts a critical current of around 80 A. We see that the predictions are generally around 10 A too low. This might be accounted for by the fact that our model is a worst-case scenario. Only 2% of the coil length exceeds 90% of the critical field. Current sharing as shown in figure C.10 in these regions will suppress heat production and should increase the viable current by a few Amperes.

Magnetic Field Measuring the magnetic field along the z-axis with a Gaussmeter allows to check the numerical predictions. Fig. C.12 shows that they are in excellent agreement, validating the coil model. This allows us to predict a gradient with a second coil of 505 G/cm as shown in figure C.13. Vapour pressure pumping increases this to 780 G/cm.

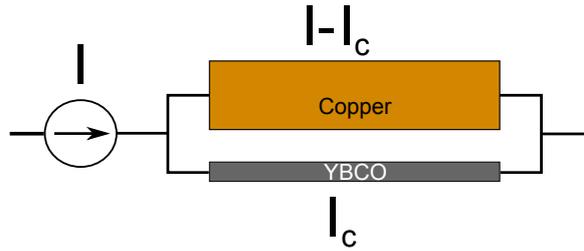


Figure C.10: Once the critical current is exceeded, the copper cladding takes over part of the current.

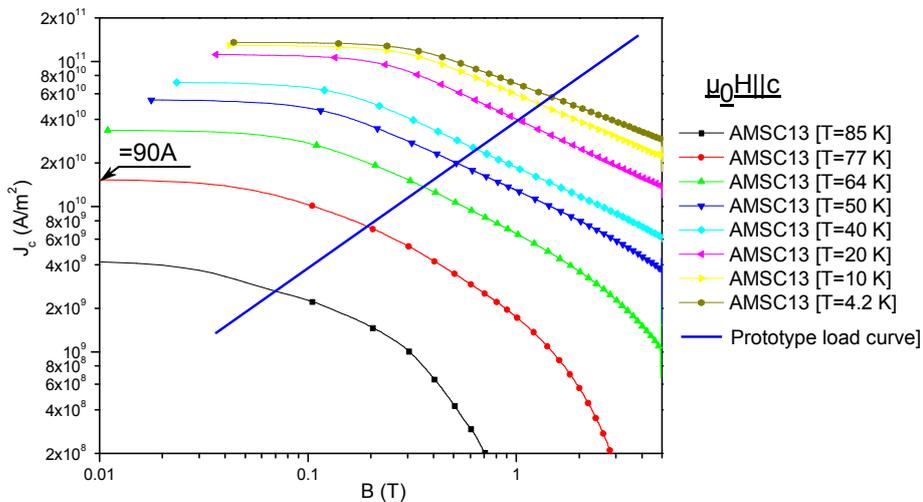


Figure C.11: Intersecting the manufacturer’s background field reduction data [105] and the numerical simulated maximum self-field per current (blue line) allows us to estimate the critical current for different temperatures. The maximum current density at zero field and 77 K of the manufacturers curve has been normalized to the 90 A capacity of our tape. The benefits of lowering the temperature are quite obvious.

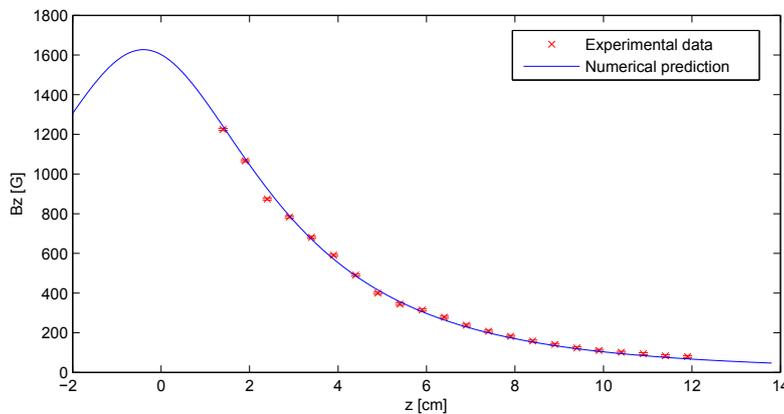


Figure C.12: A comparison of the measured and the numerically predicted field along the z-axis. They agree very well, confirming the validity of model.

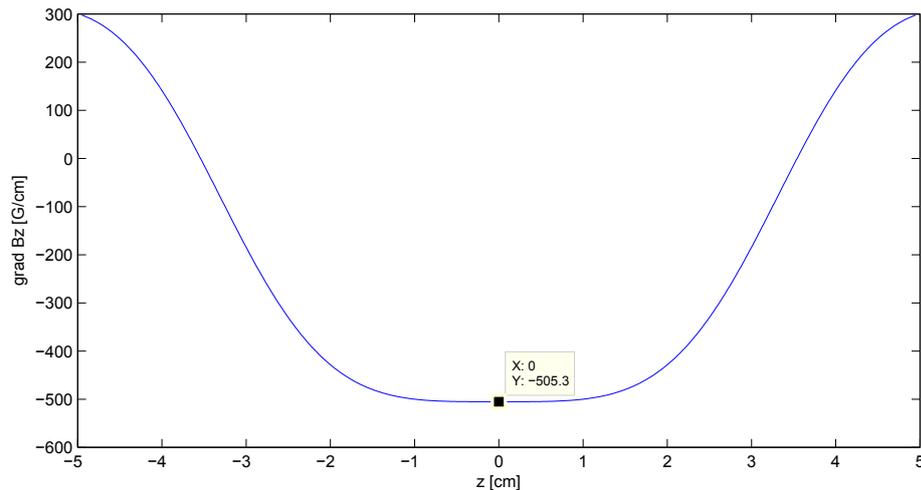


Figure C.13: Predicted Gradient with a second coil. At 55 A, a gradient of 505 G/cm at the trap centre can be expected.

C.4 Coil electronics

C.4.1 Quench protection

Superconductors are very prone to local overheating due to their high current densities and the steep increase of resistance when transitioning to normal state. The process is usually initiated by a local hot spot. Such a hot spot can be induced by a defect in the superconductor or other sources of heating. If the temperature of the hot spot is above the critical temperature of the superconductor, the section will go normal and ohmic heating will start. This will raise the temperature even further and heat up neighboring sections and eventually drive them normal as well. The normal zone will thus spread through the magnet. Such an event is called a quench. In big¹ low- T_c magnets this is a rather spectacular event, as quenches spreads with a velocity of several meters per second [106]. This event will convert the magnets energy into heat immediately and evaporate the liquid helium.

In a small high- T_c magnet, such a quench is rather different. It spreads much slower (mm/s) and the developing heat will not evaporate the nitrogen instantly [97, 107]. Because of this slow propagation, the superconductor will melt locally in a thermal runaway. Therefore a quench needs to be detected and countermeasures need to be taken. In our case, we decided to use the most common approach, a 'detect and dump' protection scheme. If a resistive voltage due to the quench is detected, the current is switched off and the coil energy LI^2 is dumped into a varistor.

The same protection electronics as in [94] were used. Three voltage taps were attached to the coil, connected to an inductive Wheatstone bridge. This cancels inductive

¹The size of an NMR or a particle accelerator magnet.

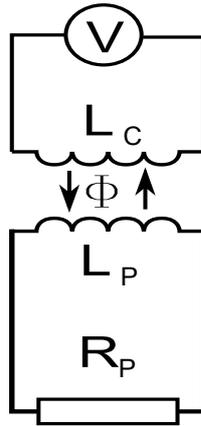


Figure C.14: A simple circuit explaining eddy currents.

voltages as well as noise and detects quenches reliably. We established numerically and experimentally [94] that the experimentally achieved detection voltage of 20 mV can be considered as absolutely safe. Having a good protection system is crucial to prevent damage when running the coil close to its critical current.

C.4.2 Switching performance

Theoretical predictions A HTS coil has typically many more windings n than a conventional coil. This means we get a huge inductance, as it scales as $L \propto n^2$. The coil prototype has $n \approx 2 \cdot 100$ and $L \approx 4$ mH. The energy stored in the coil scales as $E = 0.5LI^2$, so energizing and deenergizing the coil quickly is difficult for large inductances. However, fast switching is crucial for trapping cold atoms. Fast switch-on is required to trap atoms after the molasses stage before. Also, fast switch-off is required to produce accurate time of flight (TOF) measurements. Both situations can be represented within the same mathematical framework, in which a voltage V_0 is applied to an inductance. During switching on, V_0 is the driving voltage of the power supply (Delta SM15-200D) ranging up to 15 V. During switching off, V_0 is the (negative) clamping voltage of the varistor² and is limited by the maximum voltage rating 1200 V of the switching device (Semix 252GB126HDs IGBT). In both cases, the change in current is:

$$\dot{I}(t) = V_0/L \Rightarrow T_{switch} = \frac{I_{max}L}{V_0} \quad (C.20)$$

An anti-Helmholtz coil pair composed of two HTS coils would have $L \approx 8mH$ and would produce a maximum gradient of around 500 G/cm. For switch-on, we expect a

²A Varistor is a surge suppressing device, which nearly fulfils $V(I) = const.$

switching time $T_{switch} = 30$ ms corresponding to a magnetic gradient ramp of 17 G/(cm ms). When switching off, we get (with $V_0 = 1000$ V) a switching time $T_{switch} = 0.44$ ms, corresponding to a gradient ramp of 1100 G/(cm ms). These equations are only valid if no eddy currents are induced in the surrounding structures. Since the cryostat walls are made from stainless steel, eddy currents cannot be avoided. We will therefore derive a model of the coil, which includes eddy currents. We assume that the cryostat walls contributing to the eddy currents act as if they were a coil with a single winding $n_P = 1$, inductance L_P , resistance R_P , and current I_P . Additionally, we assume that the coil and the wall are close together and have a similar geometry. We represent our model in the circuit diagram shown in Fig. C.14. Generally, the magnetic flux through a coil A is related to its inductance by:

$$\Phi_{AA} = L_A I_A / n_A. \quad (\text{C.21})$$

For two close and well coupled coils and a magnetic flux coupling coefficient $0 \leq T \leq 1$, we can write for the flux through coil A caused by coil B:

$$\Phi_{AB} = T L_B I_B / n_B. \quad (\text{C.22})$$

The induction voltage in A is then:

$$U_A^{ind} = n_A \frac{d}{dt} (\Phi_{AB} + \Phi_{AA}). \quad (\text{C.23})$$

Using these equations for the HTS Coil C and the Plate 'Coil' P, we obtain:

$$-V = n T L_P \dot{I}_P + L_C \dot{I}_C \quad (\text{C.24})$$

$$-R I_P = L_P \dot{I}_P + T L_C \dot{I}_C / n \quad (\text{C.25})$$

$$(\text{C.26})$$

Solving these equation gives a change of coil current of:

$$\dot{I}_C = -\frac{V}{L_C} \left[1 + T^2 e^{-\frac{t}{\tau}} \right] \quad (\text{C.27})$$

$$\tau = \frac{L_P (1 - T^2)}{R_P} \quad (\text{C.28})$$

Looking at this, we see that in the limits of low coupling $T \rightarrow 0$ or small eddy current time constant $\tau \rightarrow 0$, we recover Eq. (C.20).

We obtain an upper limit for τ by approximating the relevant cryostat parts with a $d = 5$ mm thick stainless steel plate (resistivity ρ_{SS}) with the same inner and outer radii as the HTS coil. We approximate $L_P \approx L_C / n^2 \approx 1 \times 10^{-7}$ mH and $R_P =$

$\rho_{SS} d \ln(r_{in}/r_{out})/2\pi \approx 6 \times 10^{-4} \Omega$. With $T = 0$, this gives $\tau \approx 0.17$ ms. This timescale is much shorter than the switch-on time; no eddy-current effects are to be expected.

However, the switch-off time is on the millisecond scale and eddy currents are expected to be significant. The decay of the field happens in two steps: First, the coil current I_C decays, charging the plate current I_P in the process. Integrating C.27 with the starting condition $I_C(t) = I_{max}$, we see that the coil current decays quicker than without eddy currents. Once the coil current has dissipated, we can assume $I_C(t) = 0$. After this, the current in the plate I_P will decay exponentially with a time constant τ , causing the magnetic field to decay with a similar decay time.

Experimental results We checked the predictions above by measuring I_C with a current clamp as well as using a pick-up coil to monitor the change in magnetic field. The results are shown in Fig. C.16 for switch-on and Fig. C.15 for switch-off. Our measurements confirm the theoretically predicted behaviour and timescales. We achieve switch-on ramps of 4.5 G/(cm ms) and switch-off times of below 1 ms.

Timescales in cold atom experiments Extrapolating the switching behaviour to a quadrupole coil pair (with twice the inductance), we expect it to switch-on at a rate of 2×10^4 G/(cm s). Switch-off has a nearly linear part which is negligible and a ringing decay time $\tau \approx 0.15$ ms. Switch-off should therefore be complete after ≈ 1 ms. Is this fast enough for trapping and releasing atoms?

When releasing the atoms, this has to happen non-adiabatically, that means that the trap has to be switched off faster than the atoms move in the potential. For a harmonic trap with frequency ω this means:

$$t_{off} \ll \omega^{-1}. \quad (\text{C.29})$$

In typical experiments, the final trap frequency ω is on the order of 100 to 200 Hz [28], i.e. we need to switch faster than 1 ms. This is the magnitude of our switch-off time, which could cause problems. However, this is unlikely, as most of the field is already decayed at that time and thus distortions of the cloud shape should be small.

For trapping, Stuart Moulder has established optimized procedure for ramping the magnetic fields [50]. Immediate switch-on of the coils to 62 G/cm for mode-matching is followed by an adiabatic ramp of 3.5×10^3 G/(cm m). The adiabatic ramp is easy enough to match, but how fast is an 'immediate' switch-on? Assuming that the atoms with temperature $T \approx 0.2$ mK expand freely for a time $t_{on} \approx 3$ ms, we can estimate the

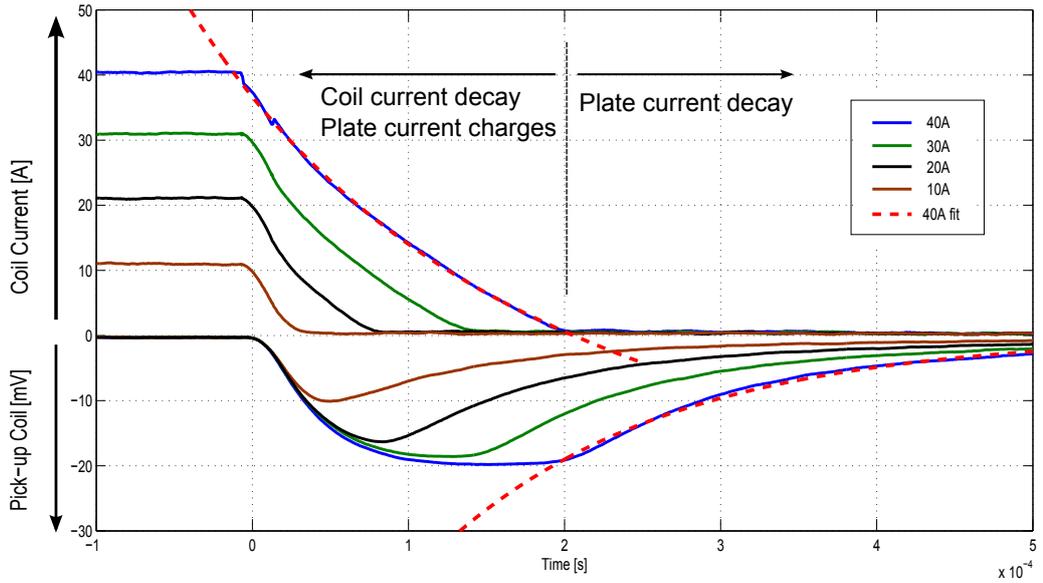


Figure C.15: Switch-off times at different starting currents. The varistor used has $V_{Clamp} = 550 \pm 10V$. The upper half shows the current in the HTS coil measured with a current clamp. The shape of the curve is as expected the integral of equation C.27. Fitting gives reasonable results. However, the relevant timescales: V/L_C and τ vary considerably between curves and fitting methods. The most consistent results for $V/L_C \approx (1.75 \pm 0.5) \times 10^5$ A/s can be extracted by fitting to the initial linear slope of I_C . The theoretical prediction is $(1.4 \pm 0.1) \times 10^5$ A/s. The discrepancy can be attributed to the simple model or that the calculated inductance is a bit too high. The lower half shows the induced voltage in the pick-up coil, caused by the magnetic flux change. This is nearly a pure exponential decay as the varistor is not in the saturation region and $I_C(t) = 0$. This allows easy fitting and extraction of $\tau \approx (0.14 \pm 0.01)$ ms. This is in very good agreement with the estimate of 0.17 ms. One should note, that this excellent agreement is probably a lucky punch given the crude model.

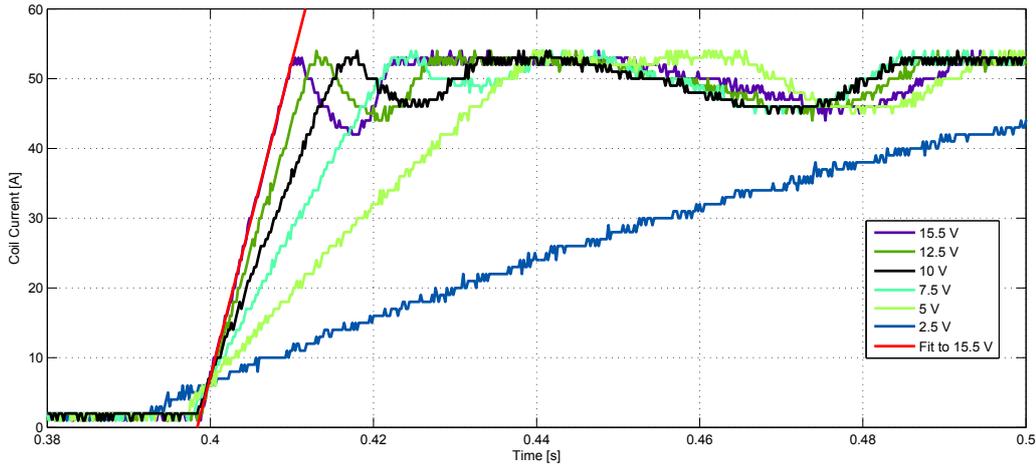
gain in potential magnetic energy:

$$\Delta E \approx \mu B' v_0 t_{on} = B' \sqrt{2kT/m_{Rb}} \approx 3 \times 10^{-3} kT. \quad (\text{C.30})$$

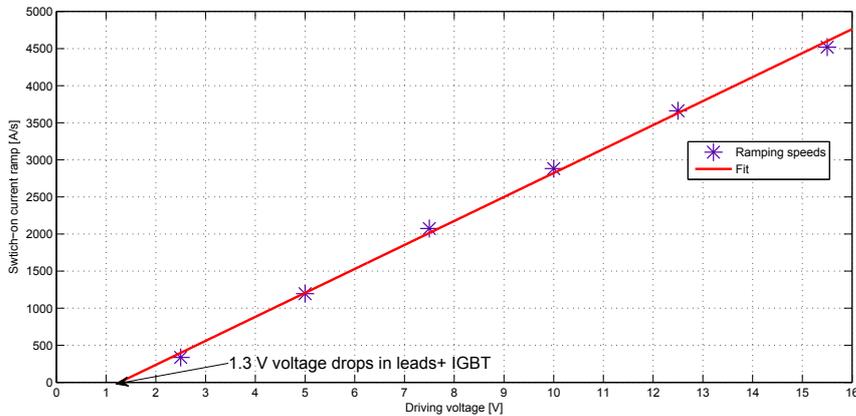
We see that even in this rather pessimistic scenario, the additional energy is negligible. The travelled distance $d = v_0 t_{on} \approx 40\mu\text{m}$ is much smaller than the diameter of a typical MOT (≈ 1 mm). The initial ramp should be fast enough in order not to affect the phase-space density.

C.5 Conclusions

We built a functional HTS prototype coil. We achieve a magnetic field gradient 25% above the water-cooled coils we use in the experiment. When employing vapour-pressure pumping, the gain can reach up to 175%. However, the prototype is still bigger than the water-



(a) A plot of the coil current during switch-on with different driving voltage. The ramps are linear, as predicted by $\dot{I}_C = V/L_C$. One exemplary fit is shown. The wiggles once the coils are ramped up are substantial and need to be taken care of. Using a PID for feedback should allow to fix it.



(b) A plot of the ramping speeds against the driving voltage. As expected, the relation is linear, with an 1.3 V offset that corresponds to the voltage drop in the leads and the IGBT. A fit gives the value of the ramp: $1/L_C = 1/3.1 \text{ mH} = 323 \text{ A}/(\text{V s})$. The inductance seems to be lower than the numerical prediction of 4 mH.

Figure C.16: Switch-on behaviour of the coil.

cooled design used in the experiment. Even when taking into account that the prototype has been over-dimensioned for ruggedness, one should not expect a substantial increase of current density from nitrogen-cooled HTS coils compared to the water-cooled solution. To increase performance further, liquid helium cooling or a close-circuit cryocooler are necessary to lower the temperature. This is viable with the vacuum insulation and should increase the current density by an order of magnitude.

However, it is questionable whether this increase in performance outweighs the increased cost and technological complexity imposed by the necessary cryogenic infrastructure of HTS coils. If better optical access is really needed, it can be achieved in many other ways, e.g. transporting the atoms to another experimental chamber, removing the

need for coils with all-optical-trapping methods or simply increasing the pressure of the cooling water through the coils.

While we originally intended to implement the coils in our current setup, we decided against it due to the reasons listed above.

C.6 Technical drawings of the cryostat

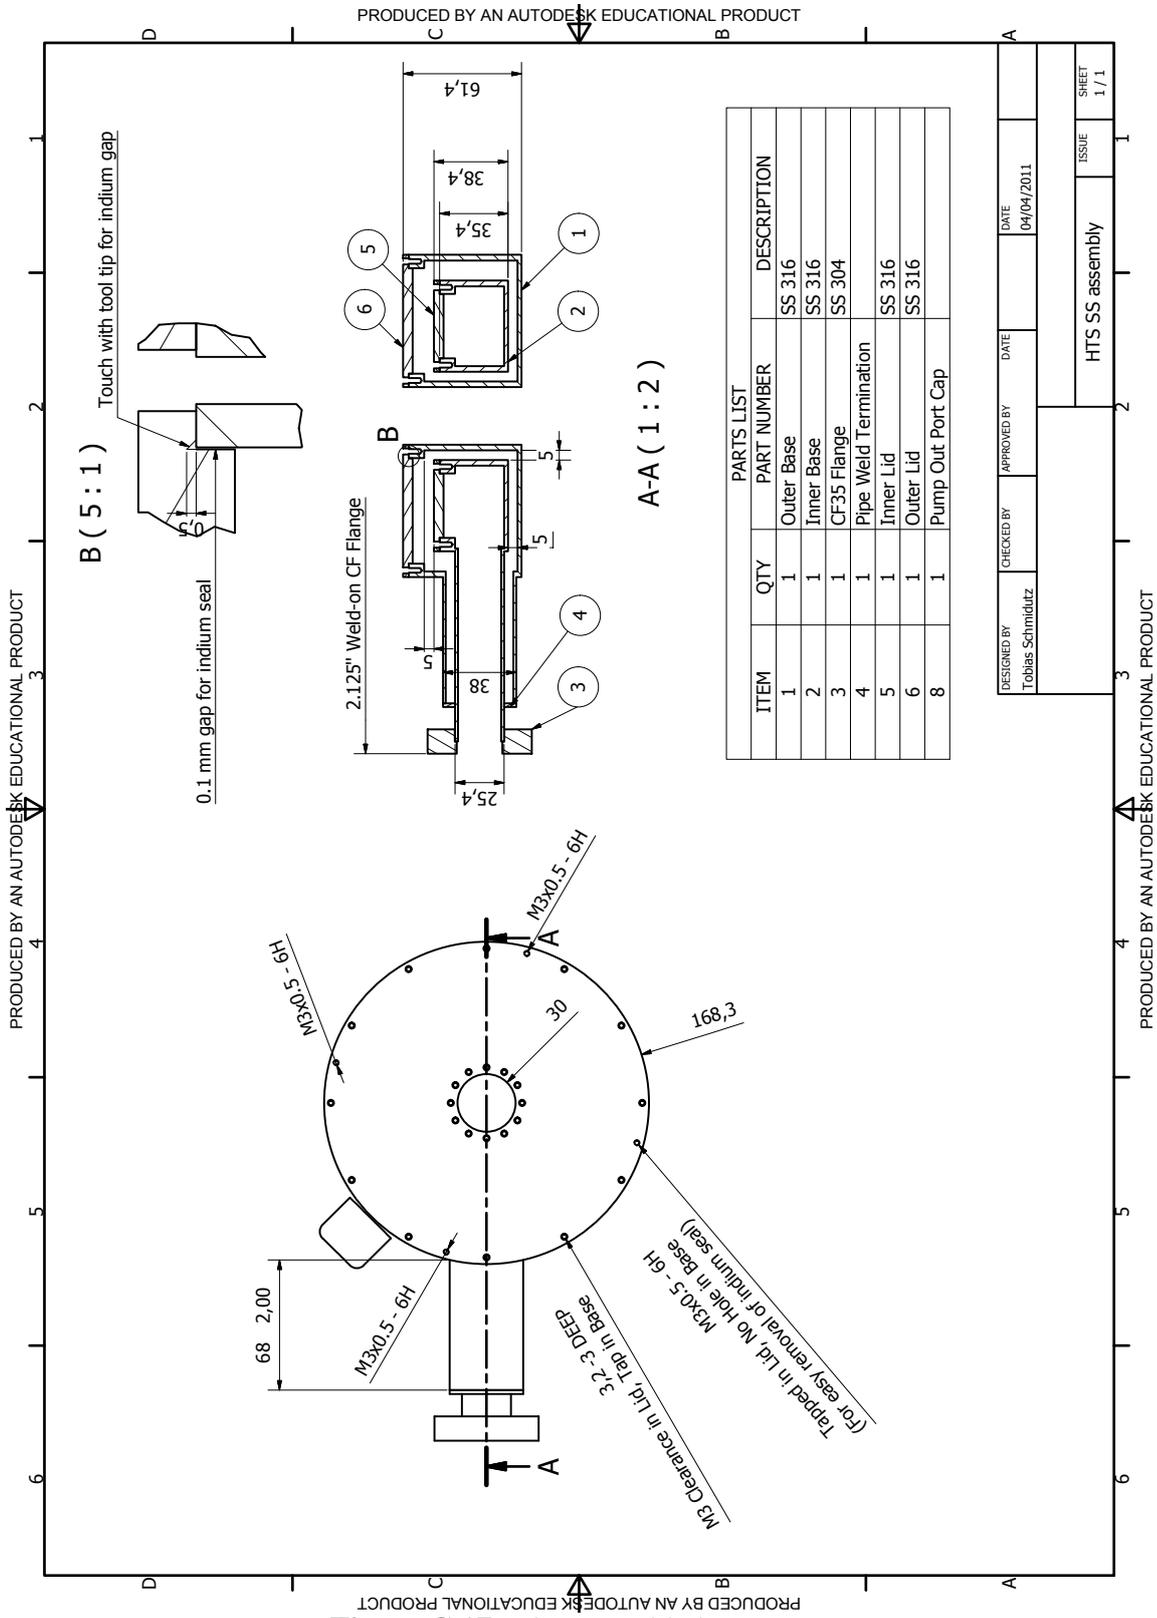
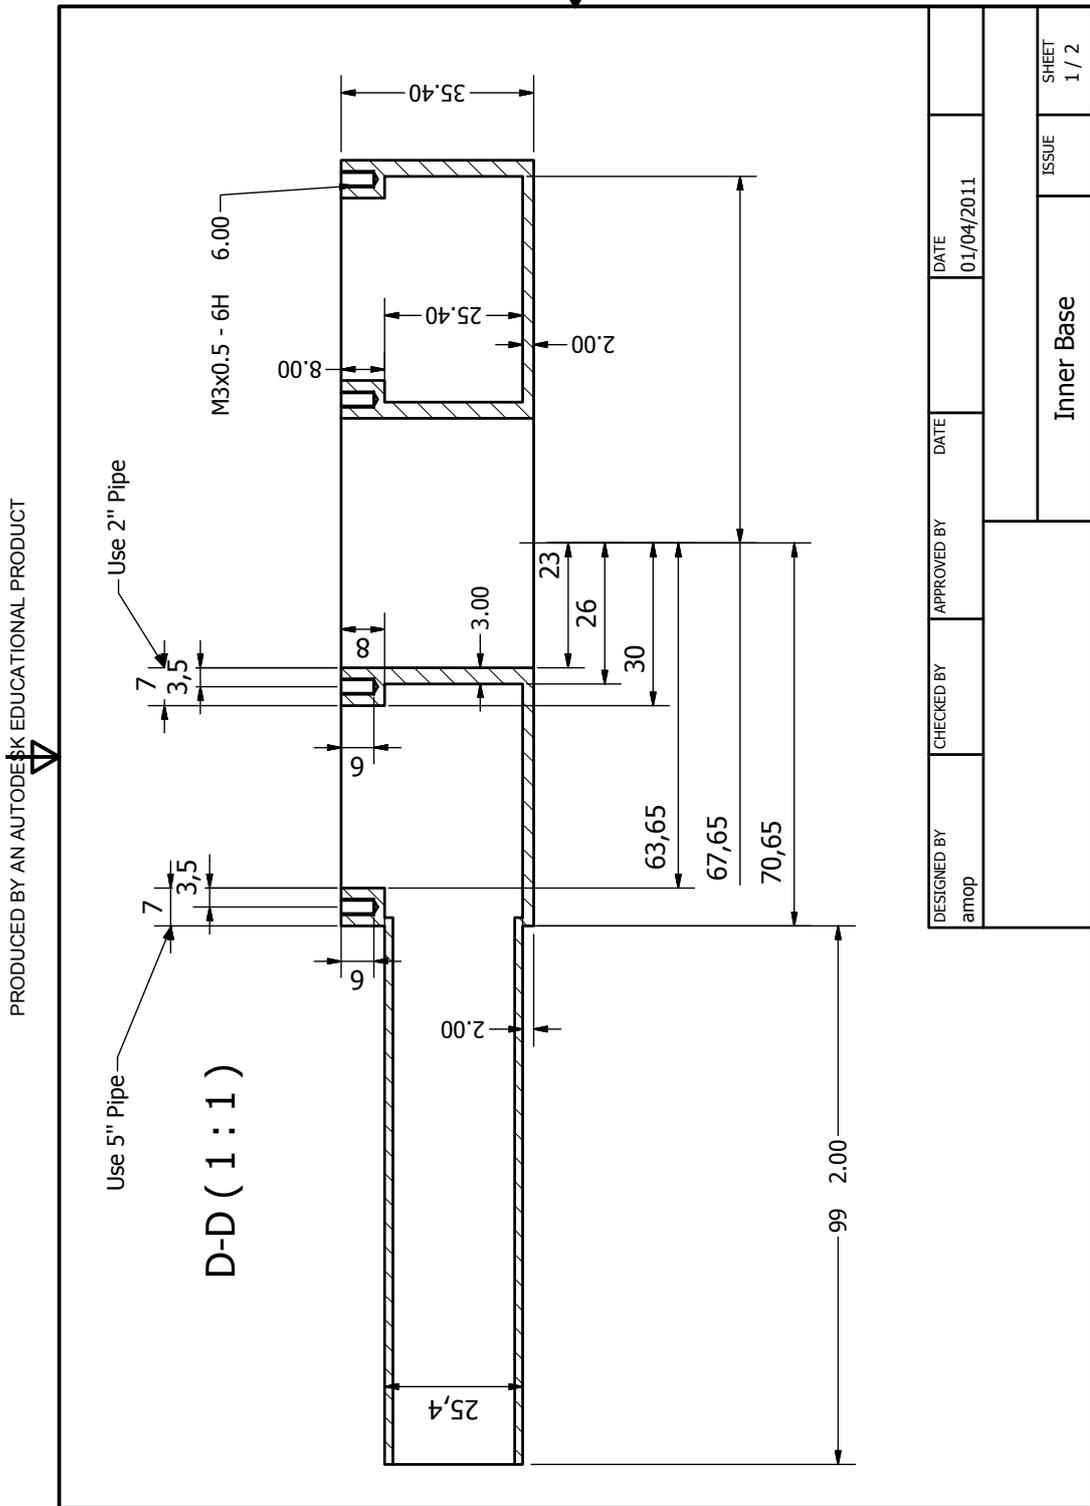


Figure C.17: The assembled cryostat.

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT



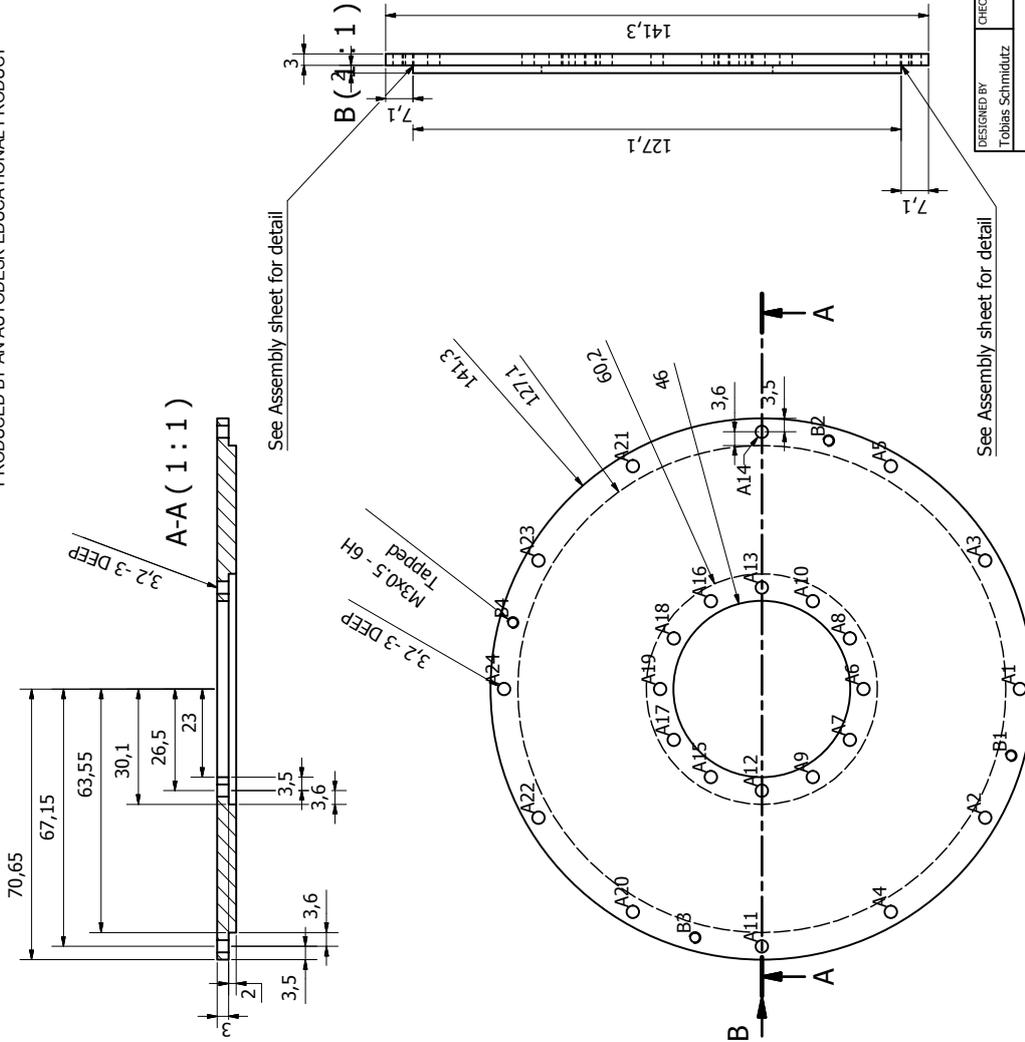
PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT

C.6. TECHNICAL DRAWINGS OF THE CRYOSTAT

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT



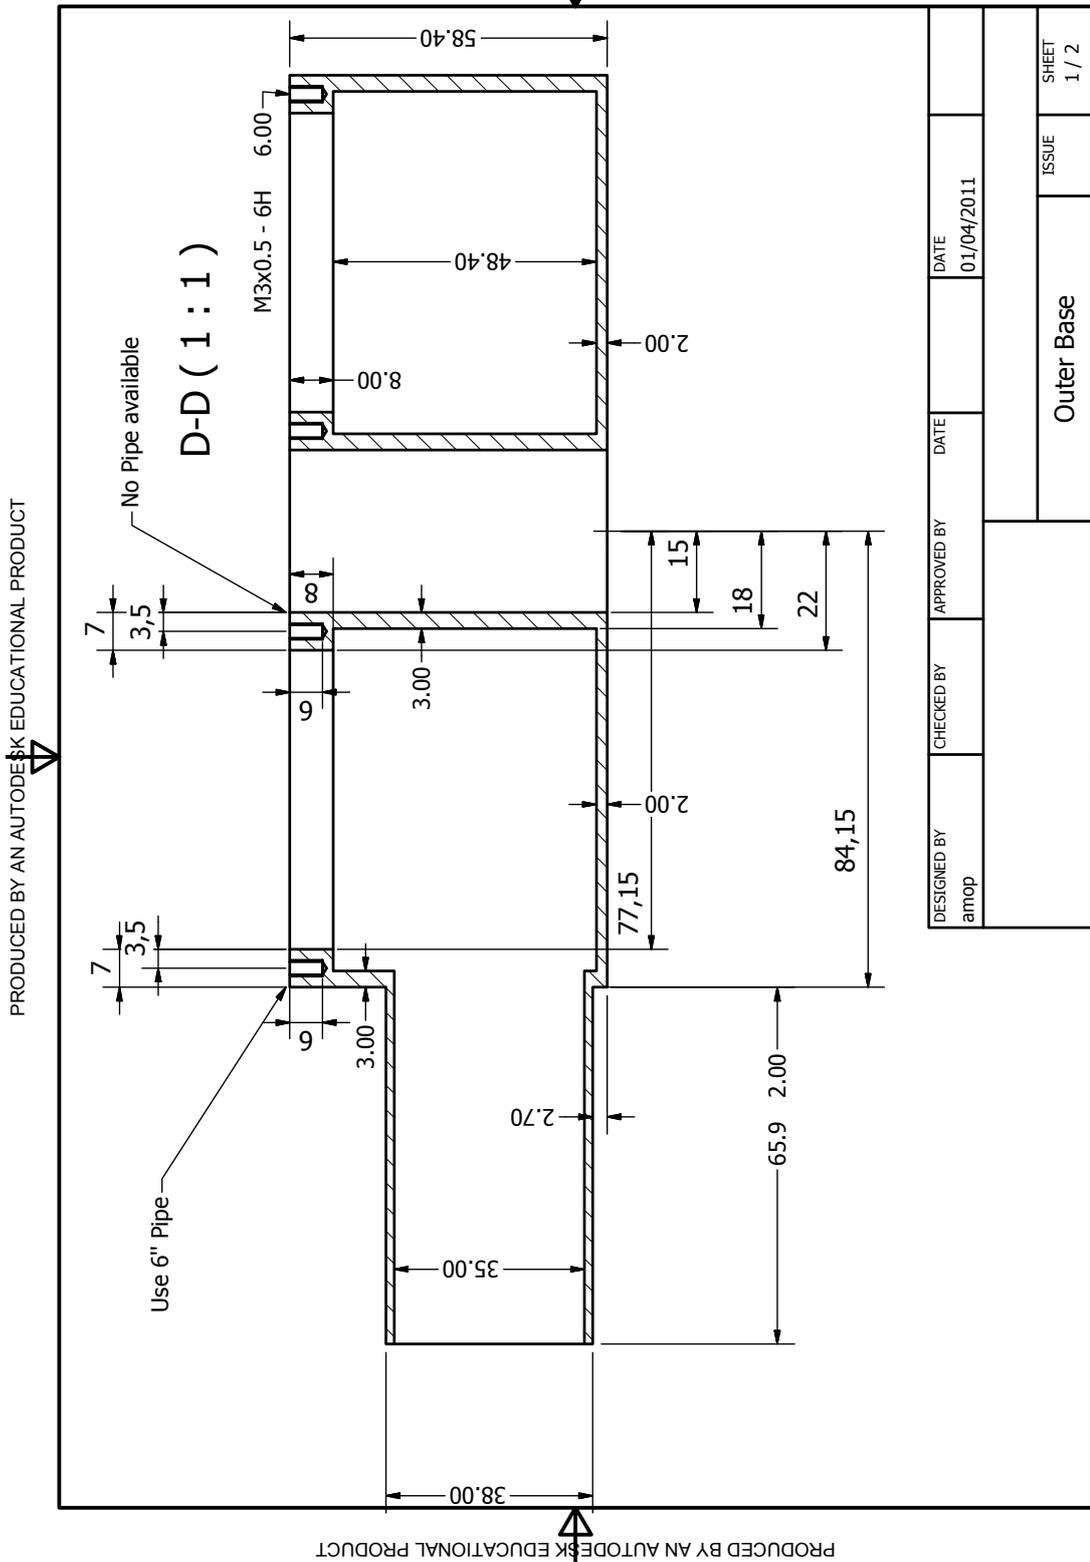
Holes A (1-12 o'clock): Clearance M3
Holes B : Tap M3 through

HOLE TABLE			
HOLE	XDJM	YDIM	DESCRIPTION
A1	0	-130,7	3,2-3 DEEP
A2	-33,58	-121,7	3,2-3 DEEP
A3	33,58	-121,7	3,2-3 DEEP
A4	-58,15	-97,13	3,2-3 DEEP
A5	58,15	-97,13	3,2-3 DEEP
A6	0	-90,05	3,2-3 DEEP
A7	-13,25	-86,5	3,2-3 DEEP
A8	13,25	-86,5	3,2-3 DEEP
A9	-22,95	-76,8	3,2-3 DEEP
A10	22,95	-76,8	3,2-3 DEEP
A11	-67,15	-63,55	3,2-3 DEEP
A12	67,15	-63,55	3,2-3 DEEP
A13	-26,5	-63,55	3,2-3 DEEP
A14	26,5	-63,55	3,2-3 DEEP
A15	-67,15	-63,55	3,2-3 DEEP
A16	67,15	-63,55	3,2-3 DEEP
A17	-22,95	-50,3	3,2-3 DEEP
A18	22,95	-50,3	3,2-3 DEEP
A19	-13,25	-40,6	3,2-3 DEEP
A20	13,25	-40,6	3,2-3 DEEP
A21	0	-37,05	3,2-3 DEEP
A22	-58,15	-29,98	3,2-3 DEEP
A23	58,15	-29,98	3,2-3 DEEP
A24	-33,58	-5,4	3,2-3 DEEP
B1	33,58	-5,4	3,2-3 DEEP
B2	0	3,6	3,2-3 DEEP
B3	-17,38	128,41	M3X0,5 - 6H
B4	17,38	128,41	M3X0,5 - 6H

DESIGNED BY Tobias Schmitz	CHECKED BY	APPROVED BY	DATE	DATE
			05/04/2011	
Inner Lid			ISSUE	SHEET
			1 / 1	1 / 1

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT

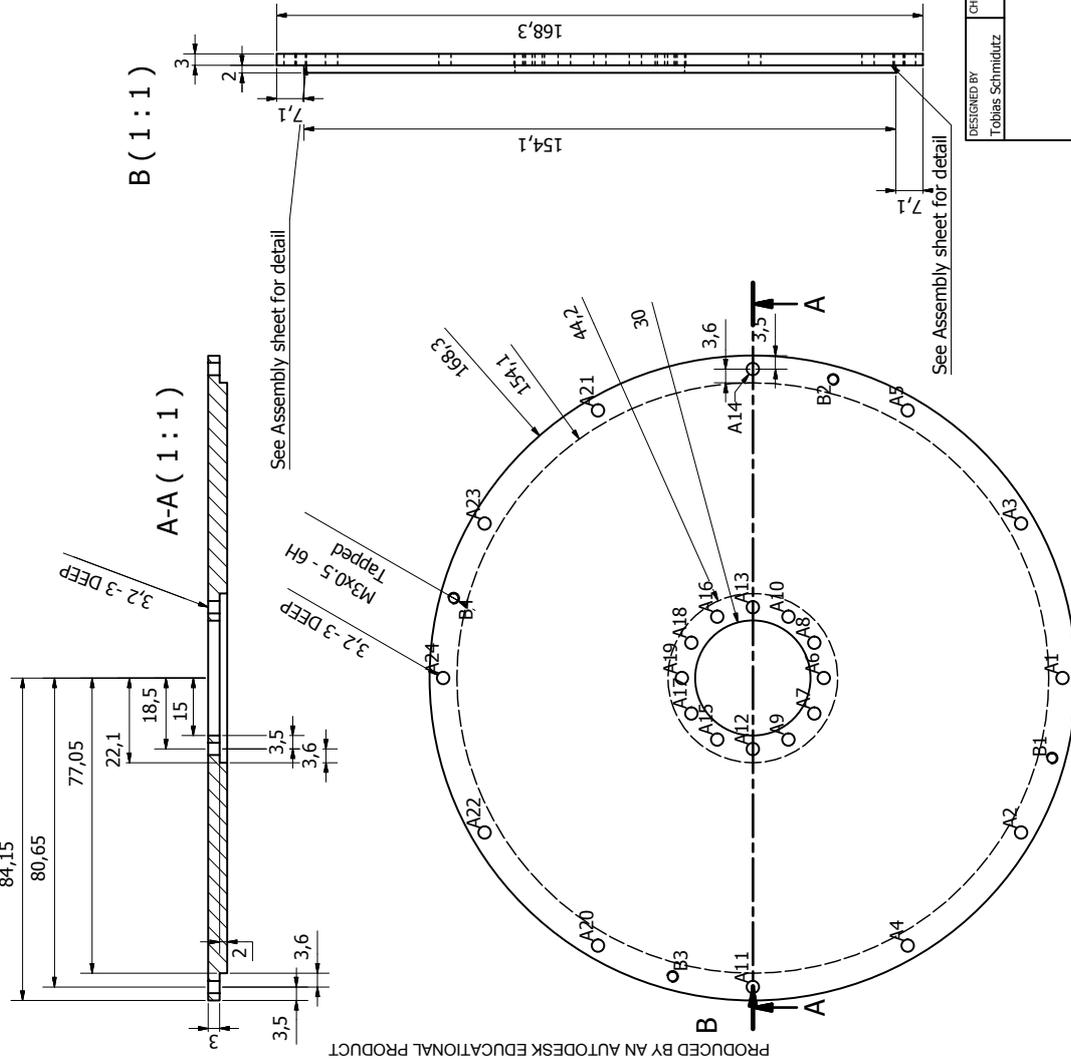
PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT



C.6. TECHNICAL DRAWINGS OF THE CRYOSTAT

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT



Holes A (1-12 o'clock): Clearance M3
Holes B : Tap M3 through

HOLE	XDJM	YDIM	DESCRIPTION
A1	0	-157,7	3,2-3 DEEP
A2	-40,33	-146,89	3,2-3 DEEP
A3	40,33	-146,89	3,2-3 DEEP
A4	-69,84	-117,38	3,2-3 DEEP
A5	69,84	-117,38	3,2-3 DEEP
A6	0	-95,55	3,2-3 DEEP
A7	-9,25	-93,07	3,2-3 DEEP
A8	9,25	-93,07	3,2-3 DEEP
A9	-16,02	-86,3	3,2-3 DEEP
A10	16,02	-86,3	3,2-3 DEEP
A11	-80,65	-77,05	3,2-3 DEEP
A12	-18,5	-77,05	3,2-3 DEEP
A13	18,5	-77,05	3,2-3 DEEP
A14	80,65	-77,05	3,2-3 DEEP
A15	-16,02	-67,8	3,2-3 DEEP
A16	16,02	-67,8	3,2-3 DEEP
A17	-9,25	-61,03	3,2-3 DEEP
A18	9,25	-61,03	3,2-3 DEEP
A19	0	-58,55	3,2-3 DEEP
A20	-69,84	-36,73	3,2-3 DEEP
A21	69,84	-36,73	3,2-3 DEEP
A22	-40,33	-7,21	3,2-3 DEEP
A23	40,33	-7,21	3,2-3 DEEP
A24	0	3,6	3,2-3 DEEP
B1	-20,87	-154,95	M3X0,5 - 6H
B2	77,9	-97,92	M3X0,5 - 6H
B3	-77,9	-56,18	M3X0,5 - 6H
B4	20,87	0,85	M3X0,5 - 6H

DESIGNED BY Tobias Schmitz	CHECKED BY	APPROVED BY	DATE 05/04/2011
Outer Lid			ISSUE 1 / 1

PRODUCED BY AN AUTODESK EDUCATIONAL PRODUCT

Bibliography

- [1] A. Einstein. Quantentheorie des einatomigen idealen Gases. *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, **261** (1924).
- [2] S. N. Bose. Plancks Gesetz und Lichtquantenhypothese. *Z. phys*, **26**, 178 (1924).
- [3] K. Davis, M. Mewes, M. Andrews, N. Van Druten, D. Durfee, D. Kurn, & W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Physical Review Letters*, **75**, 3969–3973 (1995).
- [4] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, & E. A. Cornell. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, **269**, 198–201 (1995).
- [5] C. Bradley, C. Sackett, J. Tollett, & R. Hulet. Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions. *Phys. Rev. Lett.*, **75**, 1687–1690 (1995).
- [6] C. Chin, R. Grimm, P. Julienne, & E. Tiesinga. Feshbach resonances in ultracold gases. *Rev. Mod. Phys.*, **82**, 1225–1286 (2010).
- [7] T.-L. Ho & Q. Zhou. Obtaining the phase diagram and thermodynamic quantities of bulk systems from the densities of trapped gases. *Nat. Phys.*, **6**, 131–134 (2009).
- [8] Y. Sagi, T. E. Drake, R. Paudel, & D. S. Jin. Measurement of the Homogeneous Contact of a Unitary Fermi Gas. *Phys. Rev. Lett.*, **109**, 220402 (2012).
- [9] T. E. Drake, Y. Sagi, R. Paudel, J. T. Stewart, J. P. Gaebler, & D. S. Jin. Direct observation of the Fermi surface in an ultracold atomic gas. *Phys. Rev. A*, **86**, 031601 (2012).
- [10] R. Smith, R. Campbell, N. Tammuz, & Z. Hadzibabic. Effects of Interactions on the Critical Temperature of a Trapped Bose Gas. *Phys. Rev. Lett.*, **106** (2011).

- [11] M. J. H. Ku, A. T. Sommer, L. W. Cheuk, & M. W. Zwierlein. Revealing the superfluid lambda transition in the universal thermodynamics of a unitary Fermi gas. *Science*, **335**, 563–7 (2012).
- [12] T. Yefsah, R. Desbuquois, L. Chomaz, K. J. Günter, & J. Dalibard. Exploring the Thermodynamics of a Two-Dimensional Bose Gas. *Phys. Rev. Lett.*, **107**, 130401 (2011).
- [13] C.-L. Hung, X. Zhang, N. Gemelke, & C. Chin. Observation of scale invariance and universality in two-dimensional Bose gases. *Nature*, **470**, 236–9 (2011).
- [14] N. Navon, S. Nascimbène, F. Chevy, & C. Salomon. The equation of state of a low-temperature Fermi gas with tunable interactions. *Science*, **328**, 729–32 (2010).
- [15] S. Nascimbène, N. Navon, K. J. Jiang, F. Chevy, & C. Salomon. Exploring the thermodynamics of a universal Fermi gas. *Nature*, **463**, 1057–60 (2010).
- [16] N. Tammuz, R. P. Smith, R. L. D. Campbell, S. Beattie, S. Moulder, J. Dalibard, & Z. Hadzibabic. Can a bose gas be saturated? *Phys. Rev. Lett.*, **106**, 230401 (2011).
- [17] S. Nascimbène, N. Navon, S. Pilati, F. Chevy, S. Giorgini, A. Georges, & C. Salomon. Fermi-Liquid Behavior of the Normal Phase of a Strongly Interacting Gas of Cold Atoms. *Phys. Rev. Lett.*, **106**, 215303 (2011).
- [18] A. Perali, F. Palestini, P. Pieri, G. C. Strinati, J. T. Stewart, J. P. Gaebler, T. E. Drake, & D. S. Jin. Evolution of the Normal State of a Strongly Interacting Fermi Gas from a Pseudogap Phase to a Molecular Bose Gas. *Phys. Rev. Lett.*, **106**, 060402 (2011).
- [19] C. Raman, M. Köhl, R. Onofrio, D. Durfee, C. Kukulwicz, Z. Hadzibabic, & W. Ketterle. Evidence for a Critical Velocity in a Bose-Einstein Condensed Gas. *Phys. Rev. Lett.*, **83**, 2502–2505 (1999).
- [20] T. Donner, S. Ritter, T. Bourdel, A. Ottl, M. Köhl, & T. Esslinger. Critical behavior of a trapped interacting Bose gas. *Science*, **315**, 1556–8 (2007).
- [21] T. Meyrath, F. Schreck, J. Hanssen, C.-S. Chuu, & M. Raizen. Bose-Einstein condensate in a box. *Phys. Rev. A*, **71**, 041604 (2005).
- [22] S. Gupta, K. Murch, K. Moore, T. Purdy, & D. Stamper-Kurn. Bose-Einstein Condensation in a Circular Waveguide. *Phys. Rev. Lett.*, **95**, 143201 (2005).

- [23] I. Gotlibovych, T. F. Schmidutz, S. Moulder, R. L. D. Campbell, N. Tammuz, R. J. Fletcher, A. L. Gaunt, S. Beattie, R. P. Smith, & Z. Hadzibabic. A compact single-chamber apparatus for Bose-Einstein condensation of ^{87}Rb . *arXiv:1212.4108* (2012).
- [24] A. L. Gaunt, T. F. Schmidutz, I. Gotlibovych, R. P. Smith, & Z. Hadzibabic. Bose-Einstein Condensation of Atoms in a Uniform Potential. *Phys. Rev. Lett.*, **110**, 200406 (2013).
- [25] T. F. Schmidutz, I. Gotlibovych, A. L. Gaunt, R. P. Smith, N. Navon, & Z. Hadzibabic. Quantum Joule-Thomson Effect in a Saturated Homogeneous Bose Gas. *Phys. Rev. Lett.*, **112**, 040403 (2014).
- [26] I. Gotlibovych. Observing properties of an interacting homogeneous Bose-Einstein condensate: Heisenberg-limited momentum spread, interaction energy, and free-expansion dynamics. *Phys. Rev. A*, **89** (2014).
- [27] I. Gotlibovych. *Degenerate Bose Gases in a Uniform Potential*. Ph.D. thesis, University of Cambridge (2014).
- [28] C. Pethick & H. Smith. *Bose-Einstein condensation in dilute gases* (Cambridge University Press) (2002).
- [29] V. Bagnato, D. Pritchard, & D. Kleppner. Bose-Einstein condensation in an external potential. *Physical Review A* (1987).
- [30] O. Luiten, M. Reynolds, & J. Walraven. Kinetic theory of the evaporative cooling of a trapped gas. *Physical Review A*, **53**, 381–389 (1996).
- [31] W. Ketterle, D. S. Durfee, & D. M. Stamper-Kurn. Making, probing and understanding bose-einstein condensates. *arXiv:cond-mat/9904034* (1999).
- [32] Y. Castin & R. Dum. Bose-Einstein Condensates in Time Dependent Traps. *Phys. Rev. Lett.*, **77**, 5315–5319 (1996).
- [33] F. Dalfovo, S. Giorgini, & S. Stringari. Theory of Bose-Einstein condensation in trapped gases. *Rev. Mod. Phys.*, **71**, 463–512 (1999).
- [34] R. P. Smith & Z. Hadzibabic. Effects of interactions on bose-einstein condensation of an atomic gas. In *Physics of Quantum Fluids*, 341–359 (Springer) (2013).
- [35] R. L. D. Campbell. *Thermodynamic Properties of a Bose Gas with Tuneable Interactions*. Ph.D. thesis, University of Cambridge (2011).

- [36] N. Tammuz. *Thermodynamics of Ultracold ^{39}K Atomic Bose Gases with Tuneable Inter- actions*. Ph.D. thesis, University of Cambridge (2011).
- [37] C. Foot. *Atomic physics*. Oxford master series in physics (Oxford University Press) (2005).
- [38] H. Metcalf & P. V. der Straten. *Laser cooling and trapping* (Springer) (1999).
- [39] D. A. Steck. Rubidium 87 d line data (2001).
- [40] I. Gotlibovych. Single chamber system for bec studies (2011). CPGS report.
- [41] J. Dalibard & C. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. *J. Opt. Soc. Am. B*, **6**, 2023 (1989).
- [42] W. Ketterle & N. Van Druten. Advances in atomic, molecular, and optical physics. *Advances in Atomic, Molecular and Optical Physics* (1996).
- [43] N. Newbury, C. Myatt, & C. Wieman. s-wave elastic collisions between cold ground-state ^{87}Rb atoms. *Phys. Rev. A*, **51**, R2680–R2683 (1995).
- [44] T. Kinoshita, T. Wenger, & D. Weiss. All-optical Bose-Einstein condensation using a compressible crossed dipole trap. *Phys. Rev. A*, **71**, 011602 (2005).
- [45] M. Barrett, J. Sauer, & M. Chapman. All-Optical Formation of an Atomic Bose-Einstein Condensate. *Phys. Rev. Lett.*, **87**, 010404 (2001).
- [46] J.-F. Clément, J.-P. Brantut, M. Robert-de Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel, & P. Bouyer. All-optical runaway evaporation to Bose-Einstein condensation. *Phys. Rev. A*, **79**, 061406 (2009).
- [47] K. O’Hara, M. Gehm, S. Granade, & J. Thomas. Scaling laws for evaporative cooling in time-dependent optical traps. *Physical Review A*, **64**, 051403 (2001).
- [48] P. W. H. Pinkse, A. Mosk, M. Weidemüller, M. W. Reynolds, T. W. Hijmans, & J. T. M. Walraven. Adiabatically changing the phase-space density of a trapped bose gas. *Phys. Rev. Lett.*, **78**, 990–993 (1997).
- [49] Y.-J. Lin, A. Perry, R. Compton, I. Spielman, & J. Porto. Rapid production of ^{87}Rb Bose-Einstein condensates in a combined magnetic and optical potential. *Physical Review A*, **79**, 063631 (2009).

- [50] S. Moulder. Simple single chamber system for the condensation of dilute atomic gases (2010). CPGS report.
- [51] M. Greiner, I. Bloch, T. Hänsch, & T. Esslinger. Magnetic transport of trapped cold atoms over a large distance. *Phys. Rev. A*, **63**, 031401 (2001).
- [52] E. W. Streed, A. P. Chikkatur, T. L. Gustavson, M. Boyd, Y. Torii, D. Schneble, G. K. Campbell, D. E. Pritchard, & W. Ketterle. Large atom number Bose-Einstein condensate machines. *Rev. Sci. Instrum.*, **77**, 023106 (2006).
- [53] R. Campbell, R. Smith, N. Tammuz, S. Beattie, S. Moulder, & Z. Hadzibabic. Efficient production of large ^{39}K Bose-Einstein condensates. *Phys. Rev. A*, **82**, 7 (2010).
- [54] R. Dubessy, K. Merloti, L. Longchambon, P.-E. Pottie, T. Liennard, A. Perrin, V. Lorent, & H. Perrin. Rubidium-87 Bose-Einstein condensate in an optically plugged quadrupole trap. *Phys. Rev. A*, **85**, 013643 (2012).
- [55] S. Du, M. Squires, Y. Imai, L. Czaia, R. Saravanan, V. Bright, J. Reichel, T. Hänsch, & D. Anderson. Atom-chip Bose-Einstein condensation in a portable vacuum cell. *Phys. Rev. A*, **70**, 053606 (2004).
- [56] M.-S. Heo, J.-Y. Choi, & Y.-I. Shin. Fast production of large ^{23}Na Bose-Einstein condensates in an optically plugged magnetic quadrupole trap. *Phys. Rev. A*, **83**, 013622 (2011).
- [57] D. Jacob, E. Mimoun, L. De Sarlo, M. Weitz, J. Dalibard, & F. Gerbier. Production of sodium Bose-Einstein condensates in an optical dimple trap. *New J. Phys.*, **13**, 065022 (2011).
- [58] C.-L. Hung, X. Zhang, N. Gemelke, & C. Chin. Accelerating evaporative cooling of atoms into Bose-Einstein condensation in optical traps. *Phys. Rev. A*, **78**, 011604 (2008).
- [59] F. Gerbier, J. Thywissen, S. Richard, M. Hugbart, P. Bouyer, & A. Aspect. Critical Temperature of a Trapped, Weakly Interacting Bose Gas. *Phys. Rev. Lett.*, **92**, 030405 (2004).
- [60] W. Petrich, M. Anderson, J. Ensher, & E. Cornell. Behavior of atoms in a compressed magneto-optical trap. *JOSA B* (1994).

- [61] W. Petrich, M. Anderson, J. Ensher, & E. Cornell. Stable, Tightly Confining Magnetic Trap for Evaporative Cooling of Neutral Atoms. *Phys. Rev. Lett.*, **74**, 3352–3355 (1995).
- [62] D. Pritchard. Cooling Neutral Atoms in a Magnetic Trap for Precision Spectroscopy. *Phys. Rev. Lett.*, **51**, 1336–1339 (1983).
- [63] T. Lauber, J. Küber, O. Wille, & G. Birkl. Optimized bose-einstein-condensate production in a dipole trap based on a 1070-nm multifrequency laser: Influence of enhanced two-body loss on the evaporation process. *Phys. Rev. A*, **84**, 043641 (2011).
- [64] E. A. Burt, R. W. Ghrist, C. J. Myatt, M. J. Holland, E. A. Cornell, & C. E. Wieman. Coherence, correlations, and collisions: What one learns about bose-einstein condensates from their decay. *Phys. Rev. Lett.*, **79**, 337–340 (1997).
- [65] J. Söding, D. Guéry-Odelin, P. Desbiolles, F. Chevy, H. Inamori, & J. Dalibard. Three-body decay of a rubidium BoseEinstein condensate. *Appl. Phys. B*, **69**, 257–261 (2014).
- [66] G. Reinaudi, T. Lahaye, Z. Wang, & D. Guéry-Odelin. Strong saturation absorption imaging of dense clouds of ultracold atoms. *Opt. Lett.*, **32**, 3143 (2007).
- [67] C. Sanner, E. J. Su, A. Keshet, R. Gommers, Y.-I. Shin, W. Huang, & W. Ketterle. Suppression of density fluctuations in a quantum degenerate fermi gas. *Phys. Rev. Lett.*, **105**, 040402 (2010).
- [68] A. Keshet & W. Ketterle. A distributed, graphical user interface based, computer control system for atomic physics experiments. *Rev. Sci. Instrum.*, **84**, 015105 (2013).
- [69] R. Grimm, M. Weidemüller, & Y. B. Ovchinnikov. Optical Dipole Traps for Neutral Atoms. *Adv. At. Mol. Opt. Phys.*, **42**, 95–170 (2000).
- [70] A. Gaunt. Unconventional Bose-condensed Gases: Tuning Interactions & Geometry (2013). Preliminary thesis writeup for Junior Research Fellowship application.
- [71] A. L. Gaunt & Z. Hadzibabic. Robust digital holography for ultracold atom trapping. *Sci. Rep.*, **2**, 721 (2012).
- [72] D. Stamper-Kurn, H.-J. Miesner, A. Chikkatur, S. Inouye, J. Stenger, & W. Ketterle. Reversible Formation of a Bose-Einstein Condensate. *Phys. Rev. Lett.*, **81**, 2194–2197 (1998).

- [73] S. Grossmann & M. Holthaus. Bose-Einstein condensation in a cavity. *Zeitschrift für Phys. B Condens. Matter*, **97**, 319–326 (1995).
- [74] A. L. Gaunt, R. J. Fletcher, R. P. Smith, & Z. Hadzibabic. A superheated Bose-condensed gas. *Nat. Phys.*, **9**, 271–274 (2013).
- [75] V. N. Popov. *Functional integrals in quantum field theory and statistical physics*, volume 8 (Springer) (2001).
- [76] J. Dalibard. Using the Popov approximation for a Bose gas (2010). Unpublished.
- [77] N. N. Bogoliubov. On the theory of superfluidity. *J. Phys. (USSR)*, **11** (1947).
- [78] S. W. Van Sciver. *Helium cryogenics* (Springer) (2011).
- [79] D. Kothari & B. Srivasava. Joule-Thomson effect and quantum statistics. *Nature*, **140**, 970—971 (1937).
- [80] E. Timmermans. Degenerate Fermion Gas Heating by Hole Creation. *Phys. Rev. Lett.*, **87**, 240403 (2001).
- [81] C. R. Monroe, E. A. Cornell, C. A. Sackett, C. J. Myatt, & C. E. Wieman. Measurement of Cs-Cs elastic scattering at $T=30 \mu\text{K}$. *Phys. Rev. Lett.*, **70**, 414–417 (1993).
- [82] M. Kozuma, L. Deng, E. Hagley, J. Wen, R. Lutwak, K. Helmerson, S. Rolston, & W. Phillips. Coherent Splitting of Bose-Einstein Condensed Atoms with Optically Induced Bragg Diffraction. *Phys. Rev. Lett.*, **82**, 871–875 (1999).
- [83] J. Stenger, S. Inouye, & A. Chikkatur. Bragg spectroscopy of a Bose-Einstein condensate. *Phys. Rev. Lett.* (1999).
- [84] N. Katz, J. Steinhauer, R. Ozeri, & N. Davidson. Excitation spectrum of a Bose-Einstein condensate. *Phys. Rev. Lett.* (2002).
- [85] S. Richard, F. Gerbier, J. Thywissen, M. Hugbart, P. Bouyer, & A. Aspect. Momentum Spectroscopy of 1D Phase Fluctuations in Bose-Einstein Condensates. *Phys. Rev. Lett.*, **91**, 010405 (2003).
- [86] R. Ozeri, N. Katz, J. Steinhauer, & N. Davidson. Colloquium: Bulk Bogoliubov excitations in a Bose-Einstein condensate. *Rev. Mod. Phys.*, **77**, 187–205 (2005).
- [87] S. B. Papp, J. M. Pino, R. J. Wild, S. Ronen, C. E. Wieman, D. S. Jin, & E. A. Cornell. Bragg Spectroscopy of a Strongly Interacting ^{85}Rb Bose-Einstein Condensate. *Phys. Rev. Lett.*, **101**, 135301 (2008).

- [88] P. T. Ernst, S. Götze, J. S. Krauser, K. Pyka, D.-S. Lühmann, D. Pfannkuche, & K. Sengstock. Probing superfluids in optical lattices by momentum-resolved Bragg spectroscopy. *Nat. Phys.*, **6**, 56–61 (2009).
- [89] M. Lingham, K. Fenech, S. Hoinka, & C. Vale. Local Observation of Pair Condensation in a Fermi Gas at Unitarity. *Phys. Rev. Lett.*, **112**, 100404 (2014).
- [90] P. B. Blakie & R. J. Ballagh. Mean-field treatment of Bragg scattering from a Bose-Einstein condensate. *J. Phys. B At. Mol. Opt. Phys.*, **33**, 3961–3982 (2000).
- [91] D. Stamper-Kurn, A. Chikkatur, A. Görlitz, S. Inouye, S. Gupta, D. Pritchard, & W. Ketterle. Excitation of Phonons in a Bose-Einstein Condensate by Light Scattering. *Phys. Rev. Lett.*, **83**, 2876–2879 (1999).
- [92] W. H. Bragg & W. L. Bragg. *X rays and crystal structure* (Bell) (1915).
- [93] E. Hagley, L. Deng, M. Kozuma, M. Trippenbach, Y. Band, M. Edwards, M. Doery, P. Julienne, K. Helmerson, S. Rolston, & W. Phillips. Measurement of the Coherence of a Bose-Einstein Condensate. *Phys. Rev. Lett.*, **83**, 3112–3115 (1999).
- [94] T. Schmidutz. Development of a quadrupole high-temperature superconducting coil for cold atom experiments (2010). Diploma thesis.
- [95] T. Schmidutz. Cold atoms in a single chamber system and high-*t_c* superconducting trapping magnets (2011). CPGS report.
- [96] C. Schnuck. Study of an ultracold cloud of fermionic ⁶Li atoms near a Feshbach resonance. Diploma thesis.
- [97] M. Wilson. *Superconducting magnets* (Oxford University Press, USA) (1987).
- [98] Cryogenic material properties. Cryogenics Technologies Group. <http://cryogenics.nist.gov/MPropsMAY/materialproperties.htm>.
- [99] J. Ekin. *Experimental techniques for low-temperature measurements: cryostat design, material properties, and superconductor critical-current testing* (Oxford University Press, USA) (2006).
- [100] P. Herrmann. *Handbook of applied superconductivity*, chapter D10: Current leads (Institute of Physics Publishing Inc) (1998).
- [101] T. Fließbach. *Statistische Physik* (Bibliographisches Institut) (1993).

- [102] Y. Iwasa. *Case Studies in Superconducting Magnets: Design and Operational Issues* (Springer) (2009).
- [103] C. Monaghan, E. O. Jr, & M. Good. An automatic liquid nitrogen filling system. *Cryogenics*, **21**, 540 – 542 (1981).
- [104] J. C. Simpson, J. E. Lane, C. D. Immer, & R. C. Youngquist. Simple analytic expressions for the magnetic field of a circular current loop. *Technical report*, NASA Center (2001). <http://ntrs.nasa.gov/search.jsp>.
- [105] Euratom ybco report. *Technical report*, TUWien (2007). Available from AMSC upon request.
- [106] K. Meß. *Handbook of applied superconductivity*, chapter C3: Quench propagation and magnet protection (Institute of Physics Publishing Inc) (1998).
- [107] Y. Iwasa. Protection of HTS magnets. *Physica C: Superconductivity*, **426-431**, 1348–1352 (2005).