A new Caesium quantum gas microscope with precise magnetic field control

Hendrik Paul Carsten von Raven

SIGILL·UNIVERSITAT·LUDOVIC·MAXIM·

München 2022

A new Caesium quantum gas microscope with precise magnetic field control

Dissertation an der Fakultät für Physik
Ludwig-Maximilians-Universität München

vorgelegt von

Hendrik Paul Carsten von Raven

aus Filderstadt

München, den 16. August 2022

Datum der mündlichen Prüfung: 27. September 2022

Erstgutachter: Prof. Immanuel Bloch

Zweitgutachter: Prof. Ulrich Schneider

Weitere Prüfungskommissionsmitglieder:
Prof. Dieter Braun, Prof. Fabian Grusdt

Für Annalena und Nicola

Zusammenfassung

In dieser Arbeit wird der Aufbau eines neuen Experiments mit ultrakalten Quantengasen, sowie der aktuelle Stand der Experimente beschrieben. Das neu aufgebaute Experiment verwendet Cäsium aufgrund seiner günstigen atomaren Eigenschaften. Verschiedene Feshbach-Resonanzen bei niedrigen Magnetfeldern können genutzt werden, um die Streueigenschaften zu kontrollieren. Die große Feinstrukturaufspaltung ermöglicht die effiziente Umsetzung einer neuen Methode zur Erzeugung künstlicher Eichfelder unter Verwendung eines zustandsabhängigen Gitters. Ein integriertes Quantengasmikroskop ermöglicht die ortsaufgelöste Abbildung der erzeugten Systeme und kann zur lokalen Adressierung genutzt werden.

Ein wesentliches Werkzeug für die geplante Methode zur Erzeugung künstlicher Eichfelder ist die präzise Kontrolle über die Magnetfelder und Gradienten. Mit Hilfe mehrerer Spulenpaare entlang aller drei Raumachsen können Magnetfelder bis zu $200\,\mathrm{G}$ und $100\,\mathrm{G\,cm^{-1}}$ in jeder möglichen Orientierung der Quantisierungsachse erzeugt werden. Eine aktive Magnetfeldstabilisierung wurde integriert, um Fluktuationen des Hintergrundmagnetfeldes in den horizontalen Achsen unabhängig vom erzeugten Experimentfeld zu unterdrücken. Das System wurde mittels Mikrowellenspektroskopie charakterisiert, wobei absolute Stabilitäten von $110\,\mu\mathrm{G}$ nachgewiesen werden konnten.

Eine speziell entwickelte Anordnung um die Glaszelle herum kombiniert hochauflösende Bildgebung und Magnetfeldkontrolle in einem kompakten Design und bietet gleichzeitig maximalen optischen Zugang.

Nach einer ersten Kühlungsphase werden die Atome optisch in die Glaszelle transportiert, wo sie weiter abgekühlt werden, bis sie quantenentartet sind. Das erzeugte Bose-Einstein-Kondensat (BEC) wird mit Hilfe eines vertikalen Gitters mit großen Abständen zu einem zweidimensionalen System komprimiert und in ein quadratisches optisches Gitter geladen. Mott-Isolator (MI) Zustände mit homogener Dichte werden als Ausgangspunkt für zukünftige Experimente erzeugt. Zusätzliche Gitterstrahlen mit doppelter Periodizität auf jeder Achse schaffen ein Supergitter und erweitern die Möglichkeiten der Zustandsvorbereitung und Messung.

Das in dieser Arbeit beschriebene Experiment wird die Untersuchung von topologischen Vielteilchenzuständen ermöglichen.

Abstract

In this work the construction of a new ultracold quantum gas experiment and the current status of experiments is described. The newly constructed experiment uses Caesium due to its favourable atomic properties. Various Feshbach resonances at low magnetic fields can be used to control the scattering properties. The large fine-structure splitting will allow for the efficient implementation of a new method to create artificial gauge fields using a state-dependent lattice. An integrated quantum gas microscope allows site-resolved imaging of the created systems and can be used for local addressing.

An essential tool for the planned method to create artificial gauge fields is the precise control over the magnetic fields and gradients. Using multiple coil pairs along all three axes magnetic fields up to 200 G and 100 G cm^{-1} can be generated in any possible orientation of the quantisation axis. An active stabilisation system was integrated to suppress fluctuations of the background magnetic field in the horizontal axes independent of the applied experiment fields. The system was characterised using microwave spectroscopy, demonstrating absolute stabilities of 110 μG root-mean-square (rms).

A specially designed assembly around the glass cell combines high-resolution imaging and magnetic field control in a compact design, while providing maximal optical access.

After a pre-cooling stage the atoms are transported optically into the science chamber, where they are cooled further until they become quantum degenerate. The generated Bose-Einstein condensate (BEC) is compressed into a 2d system using a large-spacing vertical lattice and loaded into an optical square lattice. Mott-insulator (MI) states with homogeneous density are created as starting point for future experiments. Additional lattice beams with twice the periodicity at each axis create a super-lattice, extending the capabilities for state preparation and read-out.

The experiment described in this thesis will allow for the study of topological many-body phases.

Contents

List of Figures

1. Introduction

Quantum mechanics describes the fundamental processes of many physical properties of nature. The closer we look and advance into more miniaturised processes, e.g. in semiconductor, the more relevant quantum descriptions are. The fundamentals have been studied and described about a century ago. Since then our understanding has deepened, yet many more questions emerged. The technical advances in the electronics industry in the 1970s and the following rise of the computers led to significant advances in many fields due to new simulation possibilities. Yet simulations of quantum mechanical effects remain difficult, as the size of the quantum states quickly outgrow any imaginable classical computer system. Instead simulators based on quantum phenomena must be used for simulation.

Over the years quantum simulation has proven to be a valuable tool for the study of fundamental quantum phenomena. Neutral atoms in optical lattices in particular are a good model system due to high level of control, the large system sizes, and the natural analogy to electrons in 'real' crystal lattices. The large scale compared to solid state systems reduces the timescales to technically easily accessible regimes and with the addition of quantum gas microscopes direct measurement of the local lattice occupation became possible [1, 2].

The advances in the field of ultracold quantum gases result in the construction of more and more complex experiments, optimised for the implementation of certain classes of Hamiltonians. Our experiment focusses on the realisation of systems with complex topological phases. In contrast to classical phases these are not related to a spontaneously broken symmetry, and can not be described using a local order parameter, but are characterised by non-local topological invariants [3]. Ultracold quantum gases in optical lattices are an effective tool for the study of these systems and have been used to implement prominent Hamiltonians with topological phases [4, 5]. Examples are the Harper-Hofstadter model [6–9], the Su-Schrieffer-Heeger (SSH) model [10–12] or the Haldane model [13–15].

The physics becomes even richer when interactions between particles are taken into account, creating effects such as the fractional quantum Hall effect (FQHE) [5]. In these systems anyons are predicted to exist, exotic quasi particles with neither bosonic nor fermionic, but fractional particle statistics. The first of them have only recently been observed experimentally for the first time [16–18].

This thesis describes the construction and initial characterisation of a newly constructed Caesium quantum gas microscope. Chapter 2 describes the favourable properties of Caesium for quantum simulation, especially in the context of topological systems. It is followed by a description of the constructed apparatus. Chapter 3 focusses on the fundamental components of the experiment, including the vacuum system, the required lasers and the corresponding first part of the experimental sequence up to the optical transport to the science chamber. The chamber and the surrounding assemblies are described in chapter 4. To complete the description of the setup a short introduction in to the used and extended control software is presented in chapter 5.

A key ingredient for the planned experiments and therefore fundamental feature of the newly constructed apparatus is the generation of ultra-stable magnetic fields. The required components including the active stabilisation are described in chapter 6 and the final stability is measured using microwave spectroscopy.

In chapter 7 the currently available quantum systems and their components needed for their creation are described. The current status of our quantum gas microscope and our machine learning based image reconstruction is described in chapter 8. Lastly an outlook to future developments and experiments is presented in chapter 9.

2. Quantum simulation with ultracold gases

In this chapter our motivation for using Caesium for a new quantum gas experiment is explained, and its relevant atomic properties are described. Subsequently a short introduction into topological systems is given. This class of quantum system is the primary system of interest. The technical and experimental challenges in their realisation strongly influenced the design of the apparatus.

This work assumes a certain level of understanding about quantum optics in general and ultracold quantum gas experiments in particular. A general introduction to quantum optics can be found e.g. in ref. [19]. A description of light-matter interactions relevant for the experiments can be found in refs. [20, 21]. An introduction and overview over the methods used in quantum gas experiments with optical lattices can be found in refs. [1, 22]. Further information about more recent advancements in the field can be found in refs. [5, 23, 24].

2.1. Caesium

The newly constructed experiment initially was planned and designed as an upgraded version of the original 'Bosons' experiment, which operated with ^{87}Rb in optical superlattices and at the time was located in the laboratory next door [25]. Towards the end of the 20 years of its operation experiments focussed on phases in many-body systems with complex topological order (i.e. [8, 11, 26–28]), also demonstrating many measurement techniques for the first time. All of these measurements were performed in the non-interacting or strongly interacting limit, as the scattering properties of ^{87}Rb can not be tuned effectively. ^{133}Cs is a bosonic alkali element as is ^{87}Rb, resulting in similar atomic properties and a Hydrogen like spectrum. In contrast to Rubidium, the scattering properties of Caesium can be tuned with the magnetic field using Feshbach resonances (section 2.1.2).

Another motivation for the construction of a new experiment was the inclusion of a quantum gas microscope, a high resolution imaging technique, allowing for the direct observation of the atom distribution in the optical lattices. The first operating quantum gas microscope was demonstrated in 2009 [29], long after the construction of the original experiment. To achieve the required resolution large optical access is needed, requiring the inclusion of the imaging system into the design and preventing the upgrade of existing experiments.

Lastly Caesium is a prominent candidate to implement an alternate method of creating systems with complex topologies using two internal spin states, trapped in a state-dependent optical lattice [30]. ^{133}Cs has very suitable atomic properties to create

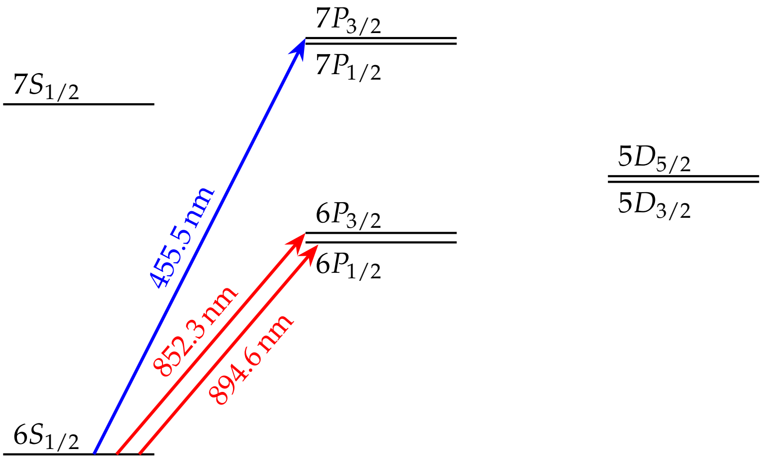

FIGURE 2.1.: State and transition diagram of ^{133}Cs.

the required optical lattice with opposite polarisability for the used internal states as described in section 2.2.1.

2.1.1. Spectral properties

As alkali metal Caesium has the electron configuration of the noble gas Xenon plus an additional free valence electron. At the energy scales present in the experiments the electrons in filled orbitals are not accessible. With the single free valence electron the spectrum can be explained by treating Caesium as a heavy Hydrogen atom.

The spin degree of freedom of the single electron results in two primary optical transitions from the $6S_{1/2}$ ground state to the $6P_{1/2}$ and $6P_{3/2}$ first excited states with wavelengths 894.6 nm and 852.3 nm respectively. These are typically called D1 and D2 lines, from the historic names of the equivalent transitions in Sodium observed in the spectrum of the sun [31]. With the increasing mass of the alkali metals the frequency of these transitions decreases, from the ultraviolet in Hydrogen to the near infrared in Caesium (fig. 2.1).

In the visible range the next pair of transitions from the ground state to the $7P$ exited states is also of potential relevance for experiments. Due to the shorter wavelength of 456 nm it can be focussed well below the diffraction limit of the D1 and D2 lines, improving the possible resolution of a quantum gas microscope. We did consider this transition for imaging (see refs. [32, 33]), but are currently not considering it a viable option for future experiments due to the numerous decay channels. Nevertheless the effect on the polarisabilty can be used for potential shaping or local addressing with high resolution [34].

In addition to the optical transitions a microwave transition at approximately 9.2 GHz can be used to drive transitions between the two hyperfine ground states $F = 3$ und $F = 4$.

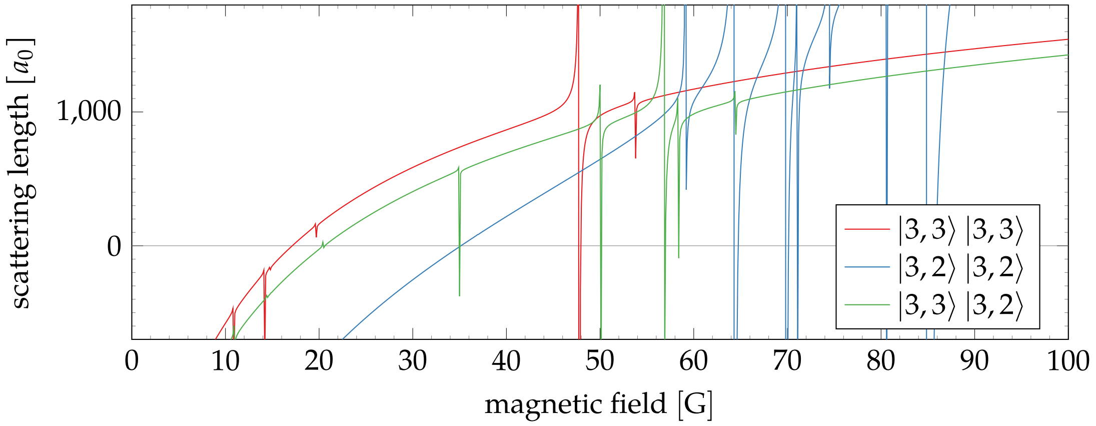

FIGURE 2.2.: Scattering length a_0 of ^{133}Cs for the $F=3$, $m_F=3$ ground state ($|3,3\rangle$) the $F=3$, $m_F=2$ state ($|3,2\rangle$) and between both spin states. Data from [38, 39].

2.1.2. Scattering properties

In the ultracold regime collisions between atoms are purely described using *s*-wave scattering, as all other collisional channels are blocked due to an energy barrier. This allows description of the scattering properties only by the *s*-wave scattering length a [35], commonly given in multiples of the Bohr radius $a_0 = 5.29 \times 10^{-11}$ m.

The scattering length is dependent on the atomic properties of the involved atoms and their internal states. For each state it is given by the state's background scattering length a_{bg}, which is modified by Feshbach resonances, occurring when a bound molecular channel in another scattering channel approaches the energy in the open scattering channel [36]. The energy shift of the involved states can be tuned optically [37], or more commonly magnetically, resulting in a varying scattering length with magnetic field as shown for various states of ^{133}Cs in fig. 2.2. Each magnetic resonance can be described as a function of the magnetic field B by a characteristic width Δ and resonance field value B_0

$$a(B) = a_{bg}\left(1 - \frac{\Delta}{B - B_0}\right). \tag{2.1}$$

The background scattering length of ^{133}Cs in its absolute ground state is unusually high around 2000 a_0 [40], but not directly relevant in experiments due to the presence of numerous and broad Feshbach resonances. These have been studied extensively theoretically [38, 41] and experimentally [42–45] in the past.

One of the primary reasons to build a new quantum gas experiment using Caesium are these Feshbach resonances. At fields up to 100 G the scattering properties are dominated by a broad resonance at −11.7 G [42], resulting in strong attractive interactions at zero field, zero interactions at 17.12 G [46] and large attractive interactions up to 1500 a_0 at 100 G, as shown in fig. 2.2. Additional narrow resonances can be used to change the scattering length quickly, but must be avoided when ramping the fields to avoid excessive losses. In the relevant region two resonances are of interest, a very narrow resonance at 19.84 G with 5 mG width and another at 47.97 G with 120 mG width [36, Table IV].

The complex structure with many resonances also results in strongly fluctuating loss rates with magnetic field. Strong two-body losses at all fields stopped first attempts of condensation of Caesium, and required the usage of the absolute ground state. In this state two-body losses are prohibited, and three-body losses are the dominant loss process. This loss rate is modified from Efimov effects, creating an experimentally relevant minimum with $L_3 \approx 3 \times 10^{-28}\,\mathrm{cm^6\,s^{-1}}$ at 21 G where $a = 210\,\mathrm{a}_0$ [40]. This minimum has a suitable ratio between the elastic and inelastic scattering processes for the creation of Bose-Einstein condensates (BECs). Even at this minimum the three body loss rate is high compared to other alkali elements, e.g. ^{87}Rb with $L_3 = 4.3 \times 10^{-29}\,\mathrm{cm^6\,s^{-1}}$ [25].

2.2. Quantum simulation of topological systems

In the past decades quantum simulation using ultracold neutral atoms has proven to be a valuable tool for the study of many quantum mechanical effects [1, 47]. However the charge-less nature of the neutral atoms at first appears to prevent the study of magnetic effects, which are fundamentally related to the coupling of a charged particle to a electromagnetic field. In the past decade new tools have been designed to emulate coupling between matter and gauge fields, suitable to study the effects of magnetism using charge neutral atoms [5]. More generally these tools allow for the study of topological quantum system. The following introduction for the experimental realisation of topological systems and their relation to magnetic fields follows ref. [48].

To simulate the physics of charge particles in a magnetic field one wants to implement the Hamiltonian of a particle with charge q moving in a magnetic field $\boldsymbol{B}$ with velocity $\boldsymbol{v}$. In this field it experiences the Lorentz force $\boldsymbol{F} = q\boldsymbol{v} \times \boldsymbol{B}$, corresponding to the Hamiltonian

$$\hat{H} = \frac{\left(\hat{\boldsymbol{p}} - q\boldsymbol{A}(\hat{\boldsymbol{r}})\right)^2}{2m}, \tag{2.2}$$

where m is the mass of the particle, and $\hat{\boldsymbol{r}}$ and $\hat{\boldsymbol{p}}$ are the position and momentum operators. $\boldsymbol{A}$ is the vector potential associated with the magnetic field $\boldsymbol{B} = \boldsymbol{\nabla} \times \boldsymbol{A}$. It is important to note that same magnetic field can be generated from an infinite number of vector potentials $\boldsymbol{A}$, as the underlying Maxwell equations are invariant under transformation with a scalar function $f(r,t)$

$$\boldsymbol{A} \to \boldsymbol{A} + \boldsymbol{\nabla} f(r,t). \tag{2.3}$$

Each choice of $\boldsymbol{A}$ is called a *gauge* and $f(r,t)$ a gauge function.

The relation to topology can be seen using the Gedankenexperiment of Aharonov and Bohm [49]. An interferometer between particles passing on either side of an infinite solenoid, shows a change in the interference pattern, depending on the current passing through the solenoid. The particles acquire a complex phase due to the present vector potential $\boldsymbol{A}$, even though the magnetic field outside the solenoid is zero. The detected phase $\Delta\varphi$ is directly dependent on the flux Φ generated by the solenoid enclosed in the two paths

$$\Delta\varphi = 2\pi \frac{\Phi}{\Phi_0}, \tag{2.4}$$

where $\Phi_0 = h/q$ is the flux quantum. The acquired phase difference is a geometric phase, in the sense that it does not depend on the velocity of the particle. It is also a topological phase as it is not dependent on the exact path taken, but only on the amount of flux enclosed within the paths. The acquired phase $\Delta\varphi$ can be derived from a gauge change between the paths and is determined by the contour $\mathcal{C}$ enclosed by the two paths in the vector potential $\boldsymbol{A}$

$$\Delta\varphi = \frac{1}{\hbar} \oint_{\mathcal{C}} q\boldsymbol{A}(\boldsymbol{r})\mathrm{d}\boldsymbol{r}. \tag{2.5}$$

As shown in the following similar geometric phases can be engineered in quantum systems to mimic the Aharonov-Bohm effect, even with charge neutral atoms. They can even been seen as a generalisation, allowing for the study of a wide range of effects.

Generally we consider systems with Hamiltonians $\hat{H}(\boldsymbol{\lambda})$ dependent on an externally controlled parameter $\boldsymbol{\lambda}$ with eigenstates $|\psi_n(\boldsymbol{\lambda})\rangle$ and eigenenergies $E_n(\boldsymbol{\lambda})$. We are interested in the evolution of an arbitrary state vector

$$|\psi(\boldsymbol{\lambda})\rangle = \sum_n c_n(t)\, |\psi_n(\boldsymbol{\lambda}(t))\rangle \tag{2.6}$$

under slow changes of $\boldsymbol{\lambda}$ with time, such that the evolution is adiabatic, i.e. the system initially prepared in an eigenstate remains in a local eigenstate throughout the evolution. In this case the time evolution of the coefficients $c_n(t)$ can be described as

$$i\hbar\dot{c}_n = \Big(E_n(t) - \dot{\boldsymbol{\lambda}}\boldsymbol{\mathcal{A}}_n(\boldsymbol{\lambda})\Big)\, c_n, \tag{2.7}$$

with *Berry's connection* $\boldsymbol{\mathcal{A}}_n(\boldsymbol{\lambda}) = i\hbar\, \langle\psi_n|\boldsymbol{\nabla}\psi_n\rangle$.

Analogous to the Aharonov-Bohm interferometer we consider the evolution of $\boldsymbol{\lambda}$ along a closed contour $\mathcal{C}$ in parameter space, such that $\boldsymbol{\lambda}(T) = \boldsymbol{\lambda}(0)$. Integrating eq. (2.7) yields

$$c_n(T) = \exp\underbrace{\Big(i\,\frac{-1}{\hbar}\int_0^T E_n(t)\mathrm{d}t\Big)}_{\Phi^{\mathrm{dyn}}} \cdot \exp\underbrace{\Big(i\,\frac{1}{\hbar}\oint_{\mathcal{C}} \boldsymbol{\mathcal{A}}_n(\boldsymbol{\lambda})\mathrm{d}\boldsymbol{\lambda}\Big)}_{\Phi^{\mathrm{geom}}} c_n(0), \tag{2.8}$$

with the *dynamic* and *geometric* phase factors. Comparing the geometric phase Φ^{geom} with equation eq. (2.5) shows the equivalence to the Aharonov-Bohm effect, where Berry's connection $\boldsymbol{\mathcal{A}}$ replaces the vector potential $q\boldsymbol{A}$. Equivalent to the magnetic field $\boldsymbol{B}$ one can define *Berry's curvature* $\boldsymbol{\mathcal{B}} = \boldsymbol{\nabla} \times \boldsymbol{\mathcal{A}}$.

To implement these systems in the context of quantum simulation it is therefore necessary to engineer an Hamiltonian with geometric phases, i.e. to engineer a system

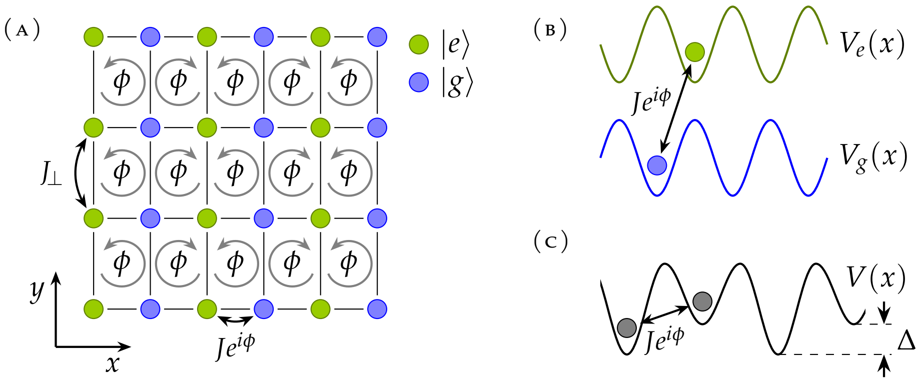

Figure 2.3.: Creation of tunnelling with complex phase factors $e^{i\phi}$ in optical lattices using Raman assisted tunnelling (b) and laser assisted tunnelling (c). Both methods use two-photon transitions to drive the transition, creating complex phase factors from the wavevector difference of the laser beams. Raman assisted tunnelling is performed in an anti-magic lattice, driving transitions between two spin states $|g\rangle$ and $|e\rangle$. Complex tunnelling is only required along one axis, a regular lattice is used in the orthogonal axis (a).

where particles accumulate a non-trivial phase factor ϕ if they move on a closed path. In the context of optical lattices movement corresponds to tunnelling along the lattice rungs, where phases can be accumulated on any tunnelling event. The experimental challenge lies in the creation of these geometric phase factors to create a system with complex topological phases.

2.2.1. Lattice gauge fields using atomic transitions

In the past years various methods for the realisation of systems with complex topological phases have been demonstrated using ultracold atoms in optical lattices, but also using photonic and solid state systems (see ref. [50] for an overview). In the cold atom context most systems were created either using *Floquet engineering* [51, 52], with laser-assisted tunnelling [26, 27], or by using internal degrees of freedom as artificial dimension [53, 54].

In the laser-assisted tunnelling scheme an energy offset $\Delta \gg J$ between neighbouring sites is created (fig. 2.3c), e.g. by applying an magnetic field gradient or using a superlattice (as shown in the figure). This offset prevents tunnelling, which is restored by introducing two additional far-detuned laser beams with relative detuning equal to Δ [55]. The resulting flux can be computed using Floquet theory as $\alpha = \delta k_y a/(2\pi)$, where δk_y is the wavevector difference between the two driving laser beams projected onto the perpendicular lattice axis.

An alternate approach using Raman assisted tunnelling between internal states was proposed by Jaksch and Zoller [30] in 2003 and extended by Gerbier and Dalibard [56], but not implemented in experiments up to date. Similar to laser-assisted tunnelling the phases are generated through the geometric beam orientation during a two-photon transfer, which is coupled to tunnelling along one lattice axis.

In the proposed scheme the flux is generated similarly using a two-photon Raman transition, additionally changing the internal spin states of the atoms. The atoms are trapped in a state-dependent optical lattice with opposite polarisability for the spin states, commonly called an *anti-magic* lattice (fig. 2.3b). Nearest-neighbour tunnelling does not exist because it would involve a spin-flip. This is induced using the Raman beams. Next-nearest-neighbour hopping is controlled independently by the lattice depth.

In this system a flux $\alpha = k_y a/(2\pi)$ is created, where k_y is the projection of the Raman wavevector along the perpendicular axis. Along the perpendicular axis a regular lattice can be used, creating a rectangular configuration with lattice spacings $\lambda/2$ and $\lambda_{\text{anti-magic}}/4$ (fig. 2.3a). The advantage of this approach is the creation of experimentally more favourable energy scale which are defined by the tunnelling rates.

In laser assisted tunnelling experiments the amplitude of the complex valued tunnelling terms J_{eff} is given as [8]

$$J_{\text{eff}} = J_x \mathcal{J}_1\left(\frac{V_0}{\Delta}\right), \tag{2.9}$$

where J_x is the real tunnelling amplitude, V_0 is the potential amplitude of the Raman beams, Δ is the introduced energy shift (fig. 2.3c), and $\mathcal{J}_1(x)$ is the first Bessel function of first kind. This Bessel function can be at most 0.6 and is in the range of 0.1 to 0.3 for typical experimental parameters. To effectively suppress tunnelling the energy shift must be large compared to the real tunnelling amplitude $\Delta \gg J_x$. Typical values of $\Delta/h \approx 5\,\text{kHz}$ limit the maximal values for $J_x/h \approx 500\,\text{Hz}$, resulting in typical values for $J_{\text{eff}}/h \approx 50\,\text{Hz}$.

In the Raman assisted tunnelling scheme tunnelling is not limited by the local energy shift, but by the much larger quadratic Zeeman shift. To suppress off-resonant coupling to the next m_F state the energy difference between the used states Δ_B and the next m_F state Δ_B' must be large compared to the Rabi frequency $\Omega \ll \Delta_B - \Delta_B'$. For moderate magnetic field values with Caesium of 30 G the quadratic Zeeman shift is already $h \times 25\,\text{kHz}$, much larger than any experimentally practical value for Δ in the laser assisted tunnelling scheme. Additionally the effective Raman couplings are larger due to the reduced lattice spacing of $\lambda/4$. We expect to achieve effective complex tunnelling rates of $J_{\text{eff}}/h = 1\,\text{kHz}$.

Without additional terms both methods generate a staggered flux pattern, with alternating flux $\pm\phi$ per unit-cell as shown in fig. 2.3a. The flux direction can be rectified by breaking the degeneracy between the complex tunnelling processes within the enlarged unit cell, e.g. by application of a gradient field [30], by spatially selective coupling [28], or by introduction of an additional super-lattice [56].

The atomic properties of ^{133}Cs make it a suitable candidate for the implementation of the proposed scheme. Not only has it favourable properties for Raman coupling [57], but is also the alkali element with the largest splitting between the D1 and D2 lines. This property is relevant for the experimental realisation of the anti-magic lattice, which can be created using circular polarised light. The vector shift creates spin dependent variations of the polarisability, creating the desired opposite polarisabilities between the

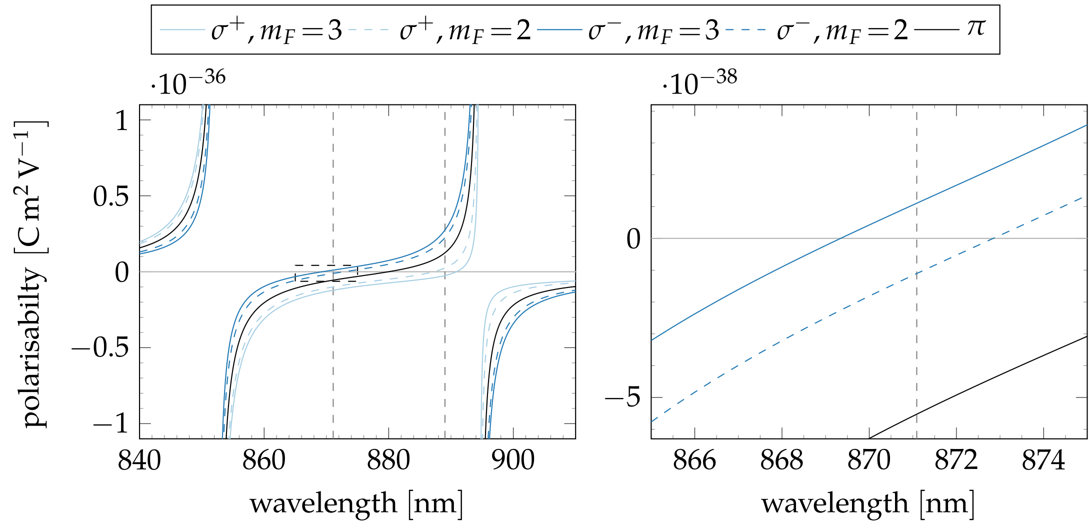

FIGURE 2.4.: Polarisability of ^{133}Cs around the D1 and D2 lines of the $F = 3$, $m_F = 3$ and $F = 3$, $m_F = 2$ states for the various polarisations. The polarisability for π polarised light is identical for both spin states. The marked region in the left plot corresponds to the area shown on the right. The dashed lines denote the anti-magic wavelength for σ^- (σ^+) polarised light at 871.1 nm (889.0 nm).

D1 and D2 lines. In ^{133}Cs optical lattices these frequencies can be far detuned from the transition lines due to the large separation between the lines of about 42 nm (fig. 2.4).

We are planning to use the absolute ground state $F = 3$, $m_F = 3$ together with $F = 3$, $m_F = 2$, for which the desired condition is fulfilled using σ^- polarised light at 871.1 nm, more than 20 nm detuned from either transition. Assuming the beam waists of our currently installed lattices we expect scattering rates below 0.5 Hz at a lattice depth of 60 recoil energies. Sufficient lattice depths for fluorescence imaging (section 7.3.3) can also be reached as it is performed at zero field. This causes the vector light shift to vanish, increasing the polarisability by approximately a factor of five. Vertical separation of the system can be used to image both spin states in a single realisation [58].

An experimental challenge using this approach is the magnetic field sensitivity of the system. The energy difference for the Raman transition is given by the Zeeman splitting between the involved spin states $\Delta E = \mu_B g_F m_F B$, which is linearly dependent on the absolute magnetic field B. In our design we included an active magnetic field stabilisation (section 6.3), to reduce fluctuations as much as possible, and achieve long coherence times in the system.

3. Experimental setup

Significant parts of the work on this thesis have been the design and construction of the necessary technical components for the planned experiments. In the last 20 years experiments with ultracold quantum gases have become more and more elaborate, adding new control parameters and observables. Still large parts of the apparatus, those needed to prepare the degenerate atomic sample, are fundamentally the same, with mainly technical improvements. Due to the technical complexity of the experiments with many interacting components their construction remains challenging, defining their capabilities and possibilities.

This chapter describes the core of our experimental system as it was designed and built. It also covers the construction process where relevant. In the following sections the ultra-high vacuum (UHV) system and different stages of the atomic sample preparation up to their arrival in the science chamber are described. This covers the optical pre-cooling stages, followed by our novel fully optical transport.

The initial design and construction of the experimental apparatus were done by me and Till Klostermann. Parts of the experimental setup, such as the cooling lasers and the optical transport are already described in detail in his thesis [59] and are only summarised in this work.

Coordinate system In the following directions of laser beams and parts of the vacuum chamber will be described in directions in our defined coordinate system. We use a right-handed coordinate system where z is oriented vertically, pointing upwards, x is oriented along the transport axis from the magneto-optical trap (MOT) towards science chamber, and y is correspondingly along the Zeeman slower axis, from the MOT chamber towards the oven. Figure 3.1 on page 13 shows a drawing of the system, including the coordinate system.

The design decisions we made throughout the design of the vacuum system caused us to include two twelve-sided chambers into the system. To properly describe the axes introduced by this design we are using primed and double-primed axis labels for the different axes in the horizontal plane. x' is the axis 30° rotated around z towards y, x'' is the axis rotated by 60°.

As explained later in chapter 6 the magnetic field coils around the science chamber use another rotated coordinate system for historic reasons. These axis are named x_c and y_c to distinguish them from the regular coordinate system.

3.1. Vacuum system

The field of ultracold quantum gases owes parts of its name to the ability of creating the coldest possible temperatures [60]. This is made possible by techniques such as laser and evaporation cooling, which extract energy without classical thermodynamic heat transfer [61]. They allow to only cool the sample itself, relying on perfect isolation from the inevitably hotter environment. This isolation is achieved by performing all experiments in ultra-high vacuum (UHV) environments. Pressures on the order of 10^{-11} mbar are necessary to reduce the rate of collisions with the background gas enough to perform experiments on the required timescales of up to a few seconds.

To achieve these low pressures hermetically sealed chambers are used which are pumped by capturing residual free particles using getter materials. At these extremely low pressures particles no longer follow classical gas theory, as their mean free path is much larger than the size of the system, and they can be described as individual particles moving between the walls of the chamber. Getter materials are designed to capture common particles for a long time and can be used either passively by covering a large surface area, or within ion pumps, which ionise and accelerate particles into the getter material.

An important consequence of the long mean free path of the particles is *differential pumping* which occurs on any bottleneck within the system [62, 63]. The reduced cross-section of the bottleneck reduces the probability of particles passing through it, creating a pressure differential between both sides. This is used to structure the system into various sections by placing narrow, long tubes between them. The conductivity C is dependent on the mean velocity of the particles $\bar{v}(T)$. For a tube of length l and diameter d it is given by

$$C = \frac{\pi d^3}{12 l} \bar{v}(T). \tag{3.1}$$

The pressure difference is dependent on the pumping speed C_{pump}, and given as $p_{out} = {}^{C}/_{C_{pump}} p_{in}$.

Thankfully only the atomic sample itself must be within vacuum. Viewports and a suitable choice of materials allow to place most of the required optics and components for magnetic fields, microwaves, and similar outside of the chamber. Nevertheless the design of the vacuum chamber is a relevant part in the construction of any quantum gas experiment. It has to integrate the requirements of the different steps needed for a quantum gas experiment.

A schematic top down view of our vacuum system is shown in fig. 3.1, its components are described in the following in the order of the experimental cycle.

We assembled the system in two steps. First without the science chamber as the glass cell had not arrived yet at the time. Later we assembled the section of the glass cell and attached it to the main chamber, flooding both chambers with Argon throughout the process. Before sealing the system it was evacuated as low as possible (around 10^{-8} mbar) using turbopumps, while being heated for an extended period to ~150 °C to

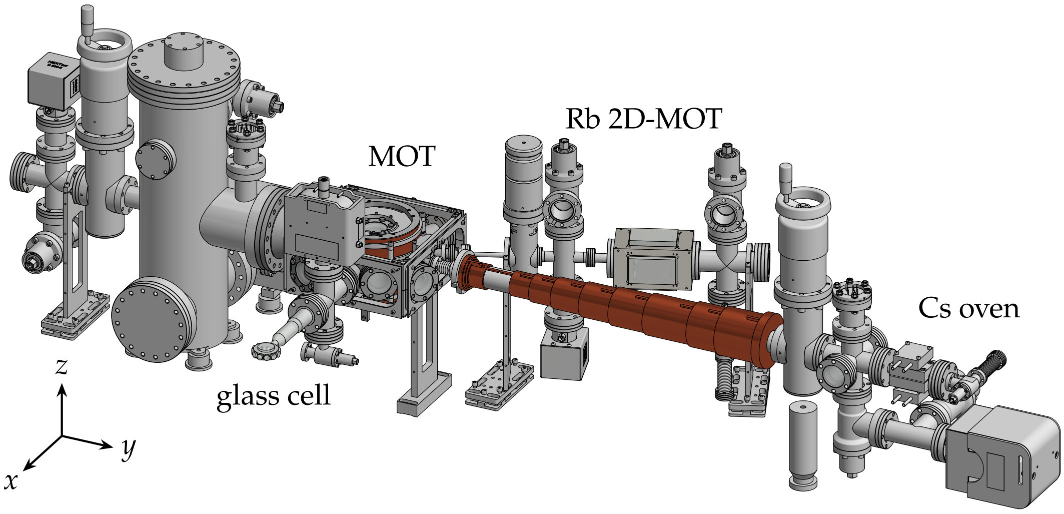

FIGURE 3.1.: Orthographic drawing of the assembled vacuum system with MOT and Zeeman slower coils. Two atoms sources are installed: An oven combined with a Zeemon slower for Caesium and a 2D-MOT for Rubidium. Both feed in the central MOT chamber used for the optical pre-cooling steps. For experiments the atoms are transported optically into the glass cell.

remove most contaminations from water and hydro-carbons. For a description of this process called bake-out and further details on the vacuum system see ref. [59].

3.1.1. Caesium source

Due to the fundamentals of quantum mechanics all measurements of the prepared quantum states, change and ultimately destroy the coherence of the prepared state. To gather statistical data the experiments are performed repeatedly, initialising a new state for every realisation. In experiments using neutral atoms this is usually done by starting with a new atomic cloud, loaded from a hot atom source, which is initially cooled using laser cooling [21, 64], followed by evaporative cooling to achieve quantum degeneracy [61, 65].

The atom source should have a high atom flux, requiring relatively high pressures, but may not compromise the ultra-high vacuum (UHV) system needed for the experiments. For further cooling using a MOT, the atoms emitted from the source must be cooled in at least one axis, to be within the capture velocity range of the MOT.

To meet these requirements one of two design is commonly used: Either an atomic oven, generating a collimated atom beam, in combination with a Zeeman slower for deceleration of the fast moving atoms along their direction of travel (e.g. [66, 67]). Alternatively a reduced MOT operating only in two dimensions, and emitting atoms along the third dimension can be used (e.g. [68, 69]). In these the atoms are loaded directly from an increased background pressure in the source chamber. For our experiment we decided to use an oven and Zeeman slower for Caesium, as it is simple in design, appeared to be the more common choice for Caesium and nominally provides a higher flux. We also included a 2 dimensional magneto-optical trap (2D-MOT) for Rubidium, which is described in section 3.1.3.

To achieve high flux the oven is heated to increase the background pressure, in our experiment to 73 °C, corresponding to a Caesium vapour pressure slightly below 10^{-4} mbar [70]. These high pressures pose a challenge for the vacuum design, to maintain low pressures in the other parts of the vacuum system. The oven section is not pumped as the high pressure is desired and the alkali metals tend to saturate and shorten ion pumps very quickly, hence differential pumping must be used.

The oven we use is adapted from the Chin group in Chicago, originally described by Subulsky [71]. The atomic beam is collimated using a tube with two apertures with 2 mm diameter on either end. This creates an effective differential pumping tube 114 mm in length. The apertures are drilled in 2 mm thick stainless steal disks which are placed in matching bores in CF40 flanges. The diameter of the tube connecting the apertures is much larger (16 mm inner diameter), simplifying production and avoiding the risk of clogging. To reduce the contamination of the vacuum system with Caesium from the oven the tube section is cooled to 7 °C. Cooling is done using thermoelectric coolers transferring heat into a water cooled reservoir[1]. To avoid accumulation of Caesium atoms on the apertures they are heated externally with heating wire. The aperture towards the oven is heated to approximately the same temperature as the atom source in the oven. The aperture towards the UHV section is heated higher to $\sim$100 °C to ensure atoms do not accumulate on it.

The output of the oven is passing through a 6-way cross used to connect a gauge[2], an ion getter pump[3] and two viewports for beam monitoring. It is followed by a straight valve into the Zeeman slower before entering the MOT section.

3.1.2. Magneto-optical trap

The chamber for the magneto-optical trap (MOT) is the first main chamber of the experiment, and is used to capture the atoms from the Zeeman cooler and the subsequent pre-cooling steps before the atoms are transported into the science chamber.

The chamber is manufactured from a single machined body made from stainless steel with low magnetic permeability[4]. Stainless steel tubes with attached flanges for viewports are welded to it. Their length was chosen to be as short as possible, providing just enough space for fasteners to attach the viewports. Most flanges for the viewports are angled by 5° to prevent back-reflections from hitting the atoms. The full chamber was heat treated by the manufacturer to reduce the hydrogen content in the material.

The MOT chamber has a twelve-sided shape, resulting in three sets of orthogonal pairs in the horizontal plane. The first pair is used for the atoms, connecting to the oven with

---

[1] Lower temperatures, even below 0 °C can be achieved using the installed cooling elements. Due to a missing enclosure ice starts to form when decreasing the temperatures significantly further.

[2] Vacom Barion Pro

[3] Gamma Vacuum 25SDI2VSCNN, 25 l/s

[4] 1.4429 ESU

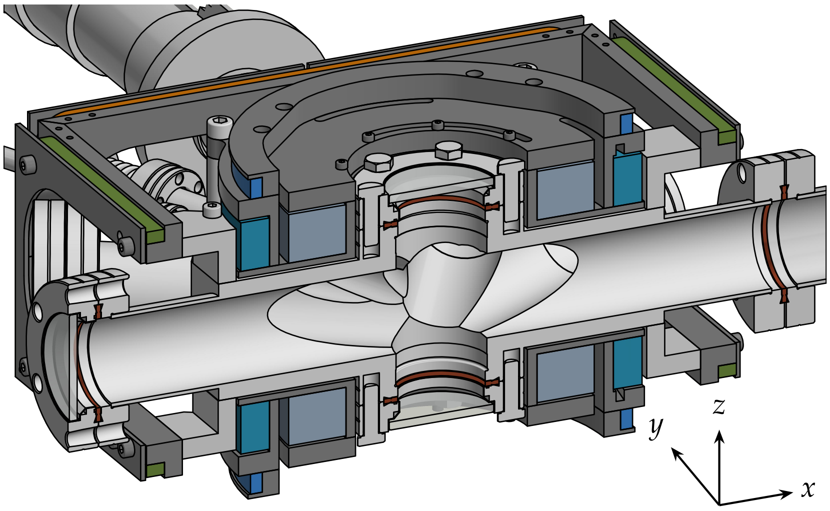

FIGURE 3.2.: Vertical cut through the MOT chamber through the transport axis in the x/z plane, highlighting the magnetic field coils installed around the chamber. Along the vertical axis the main offset (turquise) and gradient (light blue) coil pairs and an additional compensation (dark blue) are installed. Additional rectangular coils are used for compensation in the x (green) and y (orange) axes.

the Zeeman slower and orthogonal to it to the transport axis to the science chamber. The second axes are used for the large MOT beams, the third axes are used for further beams for degenerate Raman sideband cooling (dRSC) and dipole traps, and to connect the Rubidium 2D-MOT (section 3.1.3). This axis uses two CF16 flanges per side, allowing to send in dRSC beams next to the Rubidium section, whereas all other axis use CF40 flanges for viewports. Due to a problem with the initial viewport order we installed viewports from different manufacturers during assembly, most of them without coating (see ref. [59]).

Magnetic fields

We installed two pairs of coils along the vertical axis to generate offset and gradients magnetic fields in the MOT chamber. They were designed in Helmholtz and Anti-Helmholtz[5] configuration respectively, i.e. as symmetric coil pair with current flow within the coils in identical and opposite direction. For both configurations the first order non-uniformity vanishes when the distance between the coils d in relation to the coil radius R is correctly chosen. For the Helmholtz configuration this is true when $d = R$, for the Anti-Helmholtz condition $d = \sqrt{3}R$.

In our configuration both coil pairs are placed the same distance apart, but with different radii, resulting in parameters relatively close to the respective optima. The inner, gradient coil is designed to fit around a standard CF40 flange used for optical access along the vertical axis, the outer offset coil is placed around it. They were wound from

---

[5] also called gradient field Maxwell coils

Coil	Dimensions	Distance	Windings	Generated field
Offset	144 mm ID	55.2 mm	4 × 5	2.261 G/A
Gradient	80 mm ID	55.2 mm	8 × 5	1.328 G/(cm A)
Compensation x,y	45 mm × 100 mm	240 mm	20	1.5 G/A
Compensation z	145 mm ID	113 mm	20	0.5 G/A

TABLE 3.1.: Parameters of the coils placed around the MOT chamber. For round coils the inner diameter (ID) is used.

3 mm × 4 mm hollow core copper wire[6] and glued[7] directly onto the upper and lower parts of the coil mount made from polyether ether ketone (PEEK). After winding the machining of the mounts was finished to remove the central core and the two parts were combined to create the finished coil assembly. Two of these were installed symmetrically directly to the chamber, separated from it by a 1 mm spacer.

In addition to the primary coils we installed compensation along all three axes to compensate the background magnetic field. These coils are wound from 1 mm diameter enamelled copper wire after the initial bake-out of the vacuum system. We wound the horizontal coils in place, as it would not have been possible to place them along the x axis once the vacuum system was assembled. Accordingly, when we had to bake the system for a second time, they stayed in place, but were not visibly impacted by the temperatures up to 150 °C. The compensation coil for the z axis was wound onto separate mounts which are fixed to the main coil assembly. All mounts were manufactured from a combination of PEEK and aluminium, ensuring that no conductive loops were created.

3.1.3. Rubidium 2D-MOT

The experiment built throughout the duration of this thesis is the successor to a similar experiment working with bosonic ^{87}Rb, built about 20 years before this work [25]. This allowed us to use the components which were previously part of this experiment, including the lasers to cool Rubidium. While there were no immediate plans to use Rubidium for experiments in the new machine, there exist interesting ideas and proposals for mixtures. Because of this the decision was made to keep the option open and include Rubidium into the design of the vacuum chamber.

Rubidium and Caesium are both alkali metals, following each other in the periodic table. This results in quite similar optical properties in the ultracold regime and an almost identical cooling scheme. It was only necessary to add an extra pre-cooling stage and chamber. All further stages beginning with the MOT can be done in the same chambers already present for Caesium.

The focus of the work was always to use Caesium. Because of this the design work for this addition was kept to a minimum by re-using existing designs. In contrast to Caesium we did not choose a Zeeman slower, but rather a more compact 2D-MOT for

---

[6] OF-OK oxygen free copper, Hollow Conductor OD 3 mm × 4 mm/ID Ø2 mm, Luvata tool #8329, Luvata

[7] Duralco 128 adhesive expoy

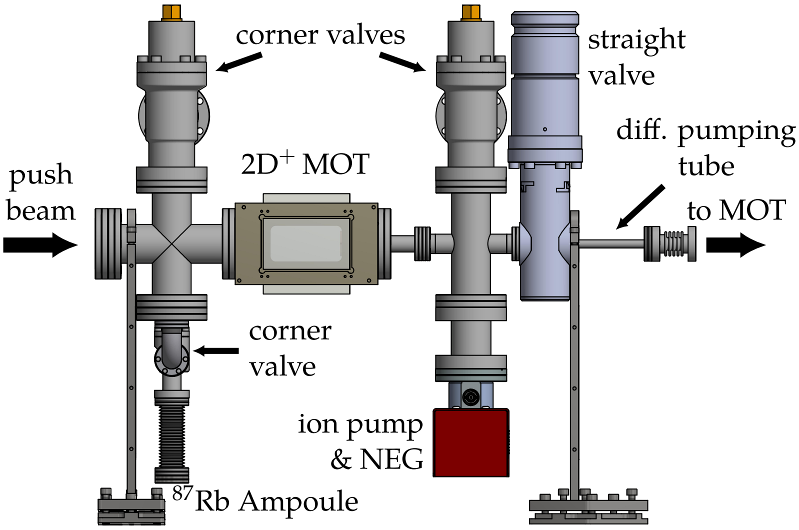

FIGURE 3.3.: Rendering of the Rubidium 2D-MOT section.

pre-cooling. We used a proven chamber design which was developed in the group many years ago and is by now used in several experiments [72, 73]. In the following I will give only a short overview of the design and describe the few minor modifications from the original design. The original design and a full characterisation are described in detail in the theses of Li [74] and Boll [72].

The design goals of the Rubidium 2D-MOT pre-cooling chamber are identical to those of the Caesium Zeeman slower and oven section. It should generate a slowed, collimated and high flux atomic beam to load the main MOT. In the 2D-MOT the atoms are loaded directly from the background pressure into a trapping potential. The separated chamber allows for a pressure differential between the atom source and the MOT chamber. This allows it to maintain UHV conditions in the main parts of the system, whereas the atom source operates at the vapour pressure of Rubidium of roughly 5×10^{-7} mbar[8] [75].

Figure 3.3 shows a side view of the ^{87}Rb pre-cooling section, labelling the different components. It is built up from the atom source, the MOT section and an extra pumping section.

The atom source is a ^{87}Rb ampoule, located in a flexible bellow, behind a corner valve[9]. The bellow can be used to break the ampoule after the bake-out, the angle valve allows to contain the Rubidium to this section while it is not used. The ampoule was broken after the bake-out, while the turbo-pumps were still connected, allowing the protective gas to escape. The valve was then closed and not opened to this day.

The MOT chamber itself is designed to create reduced 2 dimensional MOT with an additional, reduced longitudinal cooling beam, also called a 2D+MOT. Large 20 mm ×

---

[8] Pressure at 25 °C. It would be possible to heat the chamber to further increase the pressure and flux. For Rubidium this is rarely done as the vapour pressure is already relatively high.

[9] VAT, Series 54.1, all-metal angle valve, DN16, 54124−GE02

40 mm circular polarised beams are used for transverse cooling. They are retro-reflected, passing through $\lambda/4$ waveplates twice to create the correct polarisation. In addition two beams of different size along the axial direction are sent in. The larger beam is retro-reflected off an in vacuum mirror with $\lambda/4$ waveplate in front, cooling the atoms along the third direction. The smaller beam pushes the atoms out into the main MOT chamber through a 1.5 mm diameter hole which is drilled into the waveplate, mirror and mount behind it. The thin hole is 20 mm deep before widening to a larger diameter, creating a strong pressure differential between both sides.

In our case the beam then passes through another differential pumping section, containing an additional combined ion and non evaporable getter (NEG) pump[10]. It also contains another corner valve[11] used during bake-out and a straight valve[12] to isolate the whole section from the main chamber. The section is connected to the main MOT section through a custom CF16 tube with a length of 80 mm and an inner diameter of 5 mm, creating an additional pressure differential. It is followed by a short flexible bellow which is connected to the main MOT chamber, allowing to align the atom beam onto the MOT if necessary.

Other experiments using this design are only using a single differential pumping section. In those experiments the exit of the 2D-MOT is only passing through a straight valve, directly into the MOT chamber. The extra section was added as we had to increase the distance between the two MOTs, to have enough space to fit the optics for large MOT beams. As seen in the vacuum overview in fig. 3.1 it is squeezed in between the Zeeman slower and the edge of the table. To maximise the available space around the science chamber we neither wanted to place it towards the science chamber, nor place the MOT further towards the centre of the table. The chosen location was the best compromise we could find at the time[13]. To not only have an unused vacuum tube connecting the chambers, and to ensure good pressure in the main chamber we added the pumping section.

The optical setup for the operation of the 2D-MOT was roughly planned, but never installed. Mounts for the large components (3″ mirrors and large rectangular wave plates needed for the transverse cooling beams) were chosen to fit into the available space and even partially acquired. The same is true for the magnetic field coils which are already wound onto water cooled mounts, but not installed yet.

---

[10] NEXTorr D 200-5, SAES Getters S.p.A.

[11] VAT, Series 54.1, all-metal angle valve, DN40, 54132−GE02

[12] VAT, Series 48.1, all-metal straight valve, DN16, 48124-CE01

[13] Originally it was planned to use the CF16 flange on the MOT chamber next to the one used now. We decided against it only when the chamber was already in production. Because of this the MOT is connected to an angled flange, while the viewport is now placed on a straight flange.

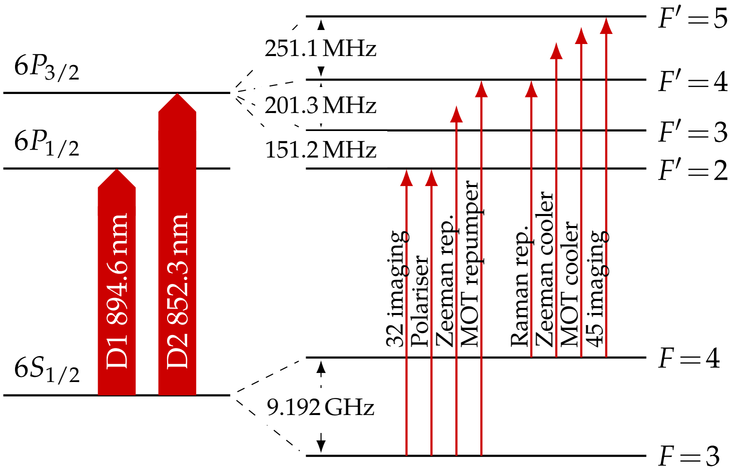

FIGURE 3.4.: Reduced level diagram of ^{133}Cs with used laser transitions.

3.2. Optics setups

3.2.1. Caesium D2 lasers

We have three laser systems operating at, or close to, the D2 line ($6S_{1/2} \rightarrow 6P_{3/2}$) of ^{133}Cs, mostly to drive resonant transitions for imaging or optical cooling. Only a single diode laser called the *master laser* is locked spectroscopically, in our case to the $F=4 \rightarrow F'=5$ transition. It is directly used for imaging and as seed for a tapered amplifier (TA) generating the cooling light for the MOT and Zeeman slower. A second diode laser called *repumper laser* used to drive transitions from the $F=3$ manifold is locked relative to the master laser using a offset lock. The third laser system is only used for degenerate Raman sideband cooling and consists of a free running diode laser, detuned by ~20 GHz from the repumper laser, used as reference for an injection locked slave laser.

The system was designed to reduce the number of required lasers, and with it the number of frequency locks, using acousto-optical modulators (AOMs) to shift the light to the required frequencies. When large frequency shifts or variability in the frequencies are needed we use AOMs in double-pass configuration. The frequency locks for the master and repumper laser are described in detail in ref. [76], drawings and extended documentation of the beam paths are described in ref. [59].

Master laser

The master laser is the primary laser running on the D2 line and used for all transitions from the $F=4$ manifold. It is an external cavity diode laser (ECDL)[14] generating about 50 mW of power. Its frequency is locked spectroscopically to the $F=4 \rightarrow F'=5$ transition using modulation transfer spectroscopy [77–79]. This is a doppler-free spectroscopy

---

[14] DLPro, Toptica

method where sidebands are modulated onto the pump beam using an electro-optical modulator (EOM) with a frequency below the natural linewidth of the transition (3 MHz in our case). On resonance the modulation is transferred to the probe beam using four-wave mixing, where it can be detected on a photodiode. The locking signal is extracted from the phase information by mixing the signal with a copy passing through a delay line.

In contrast to classical frequency modulation spectroscopy the modulation transfer spectroscopy only generates a signal when the sub-Doppler resonance condition is fulfilled, resulting in an intensity independent error signal [79], improving the lock stability. The transfer is most efficient on transitions, hence we are locking the $F=4 \rightarrow F'=5$ transition to lock the master laser. The details of our specific locking setup can be found in ref. [76].

About half of the generated power is frequency shifted by 2×-230 MHz using a double-pass AOM (DPAOM) and used to seed a TA[15], generating 1 W of optical power on its fibre-coupled output. The amplified light is split into two paths, each with another DPAOM, shifting the frequency back and generating the cooling light for Zeeman slower and MOT. The Zeeman slower is operated at a fixed detuning of -100 MHz to the $F=4 \rightarrow F'=5$ transition, the detuning of the MOT light is changed during the different stages from -10 MHz to -115 MHz by changing the radio frequency (RF) input to the AOM. The MOT light is also used for the optical molasses in the science cell during fluorescence imaging.

We are also generating light for the repumper during dRSC on the $F=4 \rightarrow F'=4$ transition and imaging light for the $F=4 \rightarrow F'=5$ transition from this laser. Depending on the applied offset field the frequency of the imaging transition is changed, requiring an adjustment of the imaging frequency. We achieve this by using part of the frequency shifted light also used to seed the TA, passing it through a second DPAOM which shifts the frequency back close to the $F=4 \rightarrow F'=5$ transition. Changing the drive frequency of the AOMs allows us to select the imaging frequency. The imaging light is fibre-coupled to a splitting setup where it is combined with light for the $F=3 \rightarrow F'=2$ transition and then split using $\lambda/2$ on motorised sliders into fibres towards the different imaging axes at the MOT and science chamber.

Repumper laser

The repumper laser is the secondary laser running on the D2 line and used for all transitions from the $F=3$ manifold. It has its name from the generation of the repumper light used during cooling, resonant with the $F=3 \rightarrow F'=4$ transition. The laser itself is locked 50 MHz (-150 MHz) detuned to the $F=3 \rightarrow F'=3$ ($F=3 \rightarrow F'=4$) transition using an offset lock. About 0.2 mW of light from the master laser and repumper laser overlapped on a 50:50 beam splitter, the resulting beat node is measured using a fast photodiode[16]. The signal is amplified[17] and mixed with a 8.79 GHz local oscillator (LO)

---

[15] BoosTA, Toptica

[16] G4176-03, Hamamatsu

[17] KU LNA 8000B, Kuhne Electronic

frequency. The lock signal is generated by mixing the difference frequency with a copy passing through a delay line.

The repumping light for the MOT is shifted to the $F=3 \rightarrow F'=4$ transition using a DPAOM. For the Zeeman slower a single pass AOM is used, shifting the light $-50\,\mathrm{MHz}$ detuned to the transition. The light for MOT is also sent to the science chamber for optical pumping from $F=3$ to $F=4$. We also generate a polariser beam and imaging light at the $F=3 \rightarrow F'=2$ transition using another DPAOM to shift the frequency, also allowing to adjust the frequency for high field imaging. The polariser is used during dRSC and to spin-polarise the sample to the $F=3$, $m_F=3$ ground state. The imaging light is fibre-coupled and merged with the $F=4 \rightarrow F'=5$ imaging light. By now it is rarely used as the excited $F'=2$, $m_F=2$ state can decay to $F=3$, $m_F=1,2,3$, resulting in a non- imaging transition. Instead we typically apply a short repumping pulse, followed by imaging on the $F=4 \rightarrow F'=5$ transition.

Raman lattice laser

The laser setup to generate the lattice for degenerate Raman sideband cooling (dRSC) is a two-stage system of ECDLs. The first stage uses a low-power, narrow linewidth laser diode[18], injected into the second laser using a broad spectrum, high-power diode[19]. This *injection lock* forces the second laser to lase at the same frequency as the seed laser as long as its temperature and current are within a narrow capture window, combining the narrow linewidth of the seed with the high power output of the injected laser [32, 80, 81].

The seed laser is running approximately $-20\,\mathrm{GHz}$ detuned from the repumper laser, far enough to not require a dedicated frequency lock, as small frequency shifts do not influence the cooling performance. We pick off a small amount of light and send it into a wavemeter for constant monitoring of the output frequency, as regular minor adjustments to the current of the injected laser are needed to keep it within the capture window.

3.2.2. Trapping lasers

To trap the Caesium atoms we are using a combination of fibre amplifiers and diode-pumped solid-state (DPSS) systems, all operating at a wavelength of $1064\,\mathrm{nm}$. This frequency is commonly used for optical trapping as it is far red-detuned, i.e. creates an attractive potential with minimal scattering, and can be easily generated at high power, e.g. using diode laser pumped Nd:YAG systems.

Most of our trapping potentials are generated from a system of three fibre amplifiers[20]. Two of these are seeded from the same single low noise, narrow linewidth Nd:YAG based seed laser[21], generating up to $30\,\mathrm{W}$ and $45\,\mathrm{W}$ of output power. The former is

---

[18] EYP-RWE-0860-06010-1500-SOT02-0000, EYP

[19] 22045504, Lumentum

[20] all three from Azurlight Systems

[21] Mephisto S, Coherent

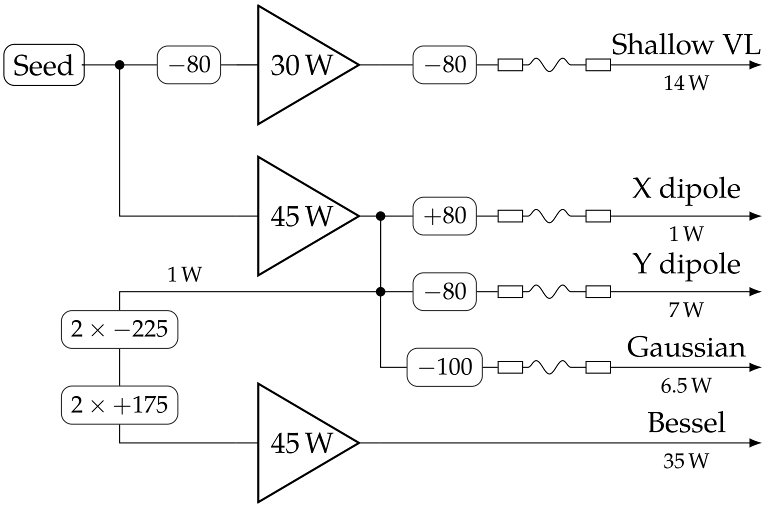

FIGURE 3.5.: Overview of the fibre amplifier setup with the installed AOMs and their respective frequencies.

operated at 26 W and only used for the shallow angle vertical lattice (section 7.2.2), the later is operated at 35 W and used to create the dipole traps along x and y, as well as the Gaussian for the optical transport (section 3.4). Approximately 1 W of this laser is frequency shifted twice using DPAOMs, and used as seed for the third amplifier. It generates up to 45 W and is internally power stabilised[22] as the full output is used directly without further intensity stabilisation to generate the Bessel beam for the optical transport.

All beams except the Bessel beam are sent through AOMs and then coupled into optical fibres to send them to the experiment breadboards. Behind the fibre a small amount of the light is extracted, either using pick-off plates or using a leak through a mirror, and sent on two photodiodes per beam. One of these is used as input to an intensity servo which controls the RF drive power of the AOMs, changing their efficiency and hence the output power. The other photodiode is used for out-of-loop monitoring of the laser power.

To avoid interference between the different trapping beams we have set up the AOMs used for intensity stabilisation such that they shift the light frequency in different directions and added an additional AOM into the seed of the 30 W amplifier. Figure 3.5 shows a schematic overview of the setup with the installed AOMs.

In addition to the fibre amplifiers we have a high power DPSS system[23] with up to 55 W of output power. Currently it is only used to generate the z dipole trap, which is intensity stabilised and fibre coupled as the other dipole traps. The laser is planned to be used for the steep vertical lattice, potentially using frequency doubling to create a super-lattice along the vertical axis similar to refs. [82, 83].

---

[22] in contrast to the other fibre amplifiers, which stabilise the current through the pump diodes.

[23] Mephisto MOPA, Coherent

3.2.3. Lattice lasers

We perform our experiments in optical lattices, created from interfering laser beams. All experiment lattices are created from retro-reflected and self-interfering laser beams. We are using laser light with a wavelength of 767 nm for our regular lattices, combining them with light at double the wavelength of 1534 nm to create so called super-lattices (section 7.2.3). The wavelengths were chosen based on the detuning to the relevant transitions in Caesium, the availability of suitable high power, low noise laser sources, and the resulting lattice spacing. 532 nm and 1064 nm would have been possible, but the very short lattice spacing of 266 nm would not have been resolvable using the primary imaging wavelength at 852 nm. High quality 767 nm light can be created using Ti:sapphire laser (Ti:Sa) lasers, and for 1534 nm high quality fibre amplifiers are available. The resulting lattice spacing of 383 nm is still significantly below the imaging wavelength, but single site-resolved imaging is possible using high resolution objectives and advanced reconstruction algorithms (chapter 8).

767 nm lasers

Our primary lattices are generated from two independent Ti:Sa systems[24], each pumped with a 25 W laser at 532 nm, generating up to 8 W of optical power at 767 nm. We use the full output of one laser per lattice axis to achieve the lattice depths required for fluorescence imaging (section 7.3.3). Due to losses in optical isolators and suboptimal fibre-coupling efficiencies we can currently use up to 3 W at the atoms. Ti:Sa crystals can be tuned over a wide frequency range, and can also be used to generate the light for the planned anti-magic lattice at 871 nm (section 2.2.1).

The beams are intensity stabilised using AOMs with feedback from photodiodes placed on the experiment table behind the optical fibre, similar to the dipole traps discussed in the previous section. We configured the stabilisation system for powers up to 200 mW, to achieve optimal performance in the regime of experimental relevant lattice depths. During experiments a low intensity noise is crucial, as the band gap in the experiment system is typically in the range of 100 Hz to 10 kHz. Noise at these frequencies can directly couple into the system and create unwanted excitations to higher bands.

Stabilisation over the full intensity range would result in reduced performance at the experimentally relevant intensities, as the dynamic range of the photodiodes is limited by electronic noise and ultimately by the photon shot noise [84, 85]. A method to overcome this limitation is the separation of the stabilisation into two regimes with independent feedback loops [86]. Alternatively logarithmic photodiodes can be used to extend the dynamical range [87].

We chose a simpler approach, and only focussed on the low power regime up to 200 mW. Above this range the intensity is no longer controlled and set to the maximum. The maximal intensity is only needed to pin the atoms during fluorescence imaging. During imaging external noise is less critical as the coherent quantum state is no longer preserved. The high lattice depth also increases the energy scales to a few 100 kHz.

---

[24] Matisse CS, Sirah Lasertechnik

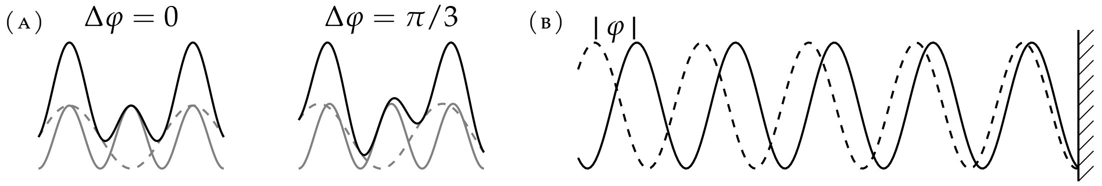

Figure 3.6.: (a) Drawing of a one dimensional super-lattice potential created from two lattice with wavelength ratio 1 : 2 and specified relative phase $\Delta\varphi$. (b) Creation of a local phase difference in a retro-reflected optical lattice. The absolute phase is fixed at the reflecting mirror. On every period a small phase difference is accumulated between beams with slightly different wavelengths.

In normal operation the used intensity servo regulates the AOM efficiency to maintain a photodiode voltage equal to an externally provided reference voltage. We clamp the photodiode voltage slightly below the maximal reference voltage using a 8.5 V Zener diode. This prevents the intensity servo from achieving equal voltages when the reference voltage is chosen above this cut-off. The result is a sudden intensity jump to maximal power as the feedback loop is driven into its limit. The chosen cut-off point allows us to perform all experiments in lattices with stabilised intensity, only jumping to the full lattice power when pinning the atoms for imaging.

1534 nm lasers

The light for the longer wavelength lattice is created using a narrow-linewidth seed laser[25] which is amplified to up to 10 W using a fibre amplifier[26]. A secondary seed[27] and amplifier[28] are available to generate independent frequencies for both axis, but currently not in use.

The intensity of the lasers beams is stabilised identically to the 767 nm lattices providing feedback from a photodiode to the RF power of an AOM[29]. These beams are stabilised over the full power range up to ~150 mW as they are not used for pinning during fluorescence imaging.

To use both wavelengths along a single axis and create a so called *super-lattice* we need to control the wavelength ratio and relative phase of the lasers. In a super-lattice an enlarged unit cell of two sites is created with an energy difference Δ between the sites controlled by the relative phase (fig. 3.6a). To create the new unit cell the wavelength λ_s of the short wavelength to be half the wavelength λ_l of the long wavelength laser.

To control the relative phase of the two lattices we slightly deviate from this condition. The local phase of the lattice beam in distance d from the retro-reflecting mirror is determined by the frequency of the laser light and the refractive index of the material

---

[25] Orion 3065-5-02-1, Rio Lasers

[26] Koheras BOOSTIK HP, NKT Photonics

[27] Koheras ADJUSTIK, NKT Photonics

[28] Nuamp NUA-1534-PB-005-B0, Nufern

[29] G23080-2-1.55LTD, Gooch & Housego

is passes through. A change in either parameter changes the wavelength, resulting in a variation of the local phase at a fixed distance (fig. 3.6b). This allows control of the relative phase between the two lattices by tuning the frequency of one of the lattice lasers relative to the other. The phase shift $\Delta\varphi$ introduced by a frequency shift $\Delta\nu$ is given by

$$\Delta\varphi = 2\pi \frac{d}{c} \Delta\nu. \tag{3.2}$$

In our system the shortest distance d is 20 cm, corresponding to a frequency difference for a π phase shift of $\Delta\nu = c/(2d) \approx 750\,\mathrm{MHz}$. The resulting deviation from the superlattice condition $2\lambda_s = \lambda_l$ is not observable on the relevant scale of the experimental system size up to 100 µm.

We control the frequency by locking the Ti:Sas to the 1534 nm fibre laser using a phase-locked loop (PLL) lock. The long wavelength light is frequency doubled using a periodically poled lithium niobate (PPLN) crystal to get it close in frequency to the short wavelength laser [88]. The frequency doubled light is mixed with the short wavelength light in a 50:50 fibre splitter/combiner[30]. One output is connected to a wavemeter for frequency monitoring, on the second output the beat between the two frequencies is sampled using a fibre-coupled, high-speed photodiode[31]. The signal is amplified[32] and directly processed in a high frequency, PLL based, laser servo[33]. The locking electronics are designed for beat frequencies up to 10 GHz, currently limited by the amplifier to 7 GHz.

The currently installed configuration uses the lock output to stabilise the Ti:Sa lasers to the frequency of the 1534 nm laser. To adapt the lock signal to the interface of the used lasers we separate it using first-order low pass and high pass filters, followed by an attenuator. The experimentally optimised values for the filters are 1 Hz for the low pass and 30 Hz for the high pass. The high frequency signal is used to control a mirror in the cavity path mounted to a piezo. The low frequency signal controls a secondary mirror on a piezo, only designed to keep the primary piezo at the centre of its dynamic range.

We can perform fast phase jumps in lass than 1 ms up to ∼2.4 MHz, corresponding to 15 mrad relative phase. Larger jumps can not be handled by the primary piezo alone and result in a collapse of the locking signal. Larger phase scans must be performed slow enough for the secondary piezo to follow. These can be performed at a rate of 2.5 MHz/ms or 16 rad/s. The linewidth of the locked lasers is 200 kHz.

Larger jumps and tuning ranges could be achieved by inverting the configuration and stabilising the 1534 nm laser to the Ti:Sa lasers. This would require installation of the secondary 1534 nm laser and construction of a second doubling stage, to achieve independent frequency control for both lattice axes.

---

[30] TW805R5A2, Thorlabs Inc.

[31] DX12CF, Thorlabs Inc.

[32] KU LNA BB 0180 A-SMA, Kuhne Electronic GmbH

[33] D2-135, Vescent photonics

3.3. Optical pre-cooling

At the beginning of every experimental cycle we have to prepare a new ultracold atomic sample. This process starts with several optical pre-cooling stages which are performed in the MOT chamber and described in the following. We are loading the MOT from a Zeeman slower, then change parameters for short compression and molasses phases, followed by degenerate Raman sideband cooling (dRSC) before we transport the atoms into the science chamber where we continue with evaporative cooling methods which are described in section 7.3.1.

3.3.1. Magneto-optical trap

The biggest part of the kinetic energy is removed from the atoms using a magneto-optical trap (MOT) as it is common for most ultracold quantum gas experiments. To decelerate the hot atoms emitted from the oven within the capture range of the MOT we are using a Zeeman slower.

In the Zeeman slower the atoms are optical cooled along their direction of travel, using varying magnetic field to compensate the changing frequency shift of the Doppler effect. Our Zeeman slower design is a spin-flip slower [89], i.e. the direction of the magnetic field changes within the slower. This design combines the advantages of an increasing and decreasing slower, as the slowing beam is detuned from the atoms when they exit the slower and does not interfere with the MOT operation, while the magnetic fields towards the MOT are low, requiring only small coils towards the MOT chamber. Our Zeeman slower is 66 cm long and was designed to decelerate atoms from ∼200 m/s down to ∼40 m/s, cooling ∼37 % of the atoms from the oven far enough to be captured in the MOT. A detailed description of the design can be found in ref. [59].

We are using 5 mW of cooling light detuned by −100 MHz from the $F=4 \rightarrow F'=5$ transition, combined with 4.5 mW repumping light resonant on the $F=3 \rightarrow F'=4$ transition. The beams starts with a 9 mm diameter and is slightly focussed onto the oven. Initial alignment was done by observation of scattered light in the oven section for rough alignment, followed by optimising the MOT loading rate.

The MOT is created from four 16 mm diameter beams, two retro-reflected beams along the horizontal axes and two counterpropagating beams along the vertical axis. The vertical axis only has a single viewport which is also used for other beams, e.g. for the Raman cooling, preventing us from using retro-reflection on this axis. Initial alignment of the beams was done by eye first on the retro-mirrors and second back to the fibre coupler. The horizontal retro-beams are slightly misaligned as we found direct reflection to create instabilities within the MOT. When optimising the power balance between the different beams we ended up with a rather asymmetric configuration with 50 mW along x'', 7 mW along y'', and 6 mW and 9 mW along z from top and bottom respectively. We attribute the higher power along x'' to the overlap with the Zeeman slower axis, performing the majority of the capturing into the MOT.

The optics used for the vertical axis are mounted to a vertical, 20 mm thick aluminium breadboard, placed in 50 mm distance from the centre of the MOT to achieve the same beam height as on the horizontal breadboards. Further optics for the horizontal MOT are mounted to the opposite side. The breadboard extends from the table up by 500 mm, and around the MOT chamber using a cut-out, which is at the end using a filler to improve rigidity. The breadboard is resting on four contact points and clamped to the table. Two diagonal support beams, filled with a mixture of fine sand and lead balls are used for stabilisation and vibration dampening.

With our current optimisation and regular sequences we load around 3×10^7 atoms within 3 s. The time is too short to load the MOT fully, but sufficient for our experiments.

Compression and molasses

To increase the density of the loaded MOT we follow it with short compression and molasses phases. At the start of the compression phase we stop the loading by disabling the Zeeman slower beams, and ramp up the magnetic field gradient from $9\,\mathrm{G\,cm^{-1}}$ to $17.5\,\mathrm{G\,cm^{-1}}$ and the detuning of the cooler beams from $-20\,\mathrm{MHz}$ to $-40\,\mathrm{MHz}$. Additionally we reduce the power of the cooler to about one third, and turn off the repumper almost completely. This rapidly compresses the MOT, creating typical densities around $5 \times 10^{10}\,\mathrm{cm^{-3}}$, at the cost of relatively high temperatures around $30\,\mu\mathrm{K}$.

To reduce the temperatures while maintaining the high densities we follow with a short molasses phase during which all fields are turned off. At the same time the repumper is set to full power again, while the cooler is reduced further to 0.2 mW and far detuned to $-100\,\mathrm{MHz}$. This reduces the temperatures to about $10\,\mu\mathrm{K}$ within 10 ms while maintaining densities around $3 \times 10^{10}\,\mathrm{cm^{-3}}$.

3.3.2. degenerate Raman sideband cooling

We finish the optical cooling stage with degenerate Raman sideband cooling (dRSC) to reduce the temperature further while maintaining density. Raman cooling of Caesium works exceptionally well, and is used in most experiments as part of the standard cooling sequence [66, 67].

In contrast to the laser cooling techniques used in the Zeeman slower and MOT stages before, Raman cooling requires the presence of an optical lattice with resolved vibrational levels. Cooling is performed by driving a cycle which transfers atoms into lower vibrational states, decreasing their kinetic energy and cooling them.

The internal spin state m_F is coupled with the vibrational level of the lattice ν. A weak magnetic field shifts the spin state energies, making $|m_F\rangle\,|\nu\rangle$ and $|m_F - 1\rangle\,|\nu - 1\rangle$ degenerate. Two photon Raman transitions are then used to coherently couple the states. A repumper beam, also called polariser, is used to perform the cooling by pumping atoms to the $6P_{3/2}\,F = 2$ state, from which they spontaneously decay back to the ground state. Figure 3.7 shows the process with the involved laser beams.

The repumper beam is almost perfectly σ^+ polarised with only a minor π admixture, predominantly increasing m_F. Throughout this process ν is unchanged, because the system is operated in the Lamb-Dicke regime [90, 91], where vibrational and internal states are decoupled due to the strong lattice confinement. This creates a cycle where

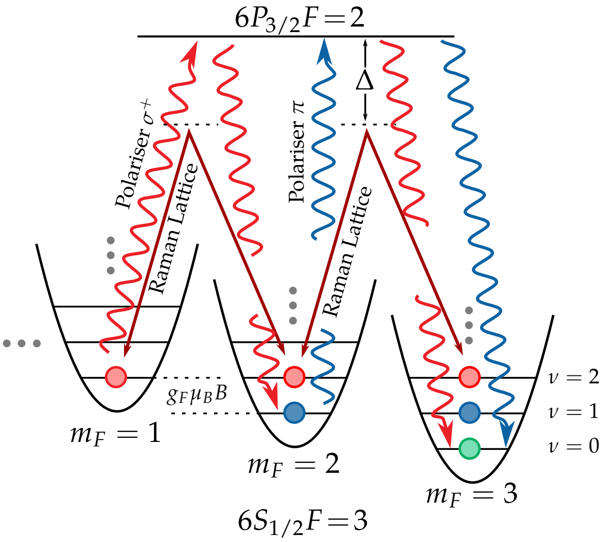

FIGURE 3.7.: Overview of the cooling cycle performed by degenerate Raman sideband cooling (dRSC). A magnetic gradient is used to make the vibrational levels ν of the different m_F spin states degenerate. Raman transitions introduce coupling between the states. The atoms are cooled as the σ^+ polariser drives the atoms to higher m_F states with reduced ν.

the atoms are pumped back into the original internal state, but their vibrational state has been reduced. This cycle repeats until the atoms are in the vibrational ground state and the Raman transitions can no longer resonantly couple to another state. The small π admixture in the repumper pumps atoms from the $|m_F=2\rangle\,|\nu=0\rangle$ state into the absolute ground state $|m_F=3\rangle\,|\nu=0\rangle$, which is dark for all beams, ideally creating a spin polarised sample in the ground state of the lattice.

To create the 3d optical potential we use four interfering near-resonant lattice beams detuned by $-20\,\mathrm{GHz}$ from the polariser beam, which also drive the Raman transitions, similar to refs. [92, 93]. Two beams along the x' and z axis pass through the chamber once and are linearly polarised, such that they do not interfere with each other. The two beams along the y' axis are created from a single retro-reflected beam, also linearly polarised, but under a small angle, creating interference with each other and both other lattice beams, resulting in a 3d optical potential with locally varying polarisation [32, 94]. The resulting potential has asymmetric trapping frequencies, but these are not resolved due to the strong Raman coupling [95].

The polariser for the $F=3 \rightarrow F'=2$ transition is sent into the chamber from the bottom, counter-propagating to the z Raman lattice beam. Its polarisation is carefully tuned to be perfectly circular, for it to be σ^+ polarised when we apply the offset field along the z axis. Any σ^- contribution causes heating, rather than cooling, and increases the final temperature of the cooling stage significantly from $400\,\mathrm{nK}$ in the ideal case to around $3\,\mu\mathrm{K}$. The π polarisation admixture is created by rotating the quantisation axis slightly by adding a horizontal magnetic field component.

In addition to the Raman beams and polariser we have an extra beam for repumping atoms which were off-resonantly pumped to the $F = 4$ manifold. Its direction is not relevant for the cooling process. We send it in from the top, co-propagating with the lattice beam.

After optimisation of the beam parameters and compensation of the magnetic field we cool the cloud within 6 ms to temperatures below 1 μK, maintaining 2×10^7 atoms and densities around $3 \times 10^{10}\,\mathrm{cm}^{-3}$.

3.4. Optical transport

After the Raman cooling the atomic sample must be transported to the science chamber for the final cooling stages and experiments. We perform this transport using a novel, all-optical transport scheme, moving the atoms over a long distance using a travelling optical lattice. The transport scheme has been published [96] and is described and characterised in depth in the thesis of Klostermann [59] and in parts in Chen [88]. In this section I will only give a short overview over the transport.

The transport is performed by creating a 1d optical lattice from two counterpropagating beams. After loading the frequency of one of the lasers is changed, creating a travelling wave which moves the atoms with it. In contrast to other optical transport schemes using variable focus positions [97, 98] this allows for very fast transport due to the high confinement along the transport axis, without movement of any mechanical objects. This approach has already been demonstrated [99–101], but has never been implemented over such a long distance. Commonly long distance transport is performed using magnetic traps [102, 103], which are not usable for ^{133}Cs as the absolute ground state is anti-trapped.

Achieving a long distance using this scheme is challenging as the transport beams need to trap the atoms throughout the whole transport distance. Maintaining a deep enough trap to overcome gravity is especially challenging with Caesium due to its high atomic mass of 133 u. To maintain a deep enough trap throughout the transport we use a Bessel beam, which has an approximately constant core diameter. The lattice is created by a counter-propagating TEM_{00} beam (Gaussian) with up to 6.5 W.

The transport beam is generated from a 1064 nm beam with up to 35 W using an axicon [104], resulting in a beam with a central spot size of 100 μm over the full transport distance of 43 cm. The relative frequency between the beams is controlled by shifting the RF input to two DPAOMs as described in section 3.2.2. Stable frequencies are generated using direct digital synthesis (DDS) boards which are locked to the same 10 MHz reference.

To load the atoms into the Bessel beam after the Raman cooling stage we first load the atoms into a large crossed dipole trap, adding the transport beam at lower power to act as an optical dimple, where the atoms accumulate due to thermalisation [105]. The reservoir beams have a diameter of 1 mm to match the size of the cloud after dRSC. They are created from the laser also used for the Bessel beam by controlling the reflected power on a polarising beam-splitter (PBS) using a motorised waveplate. The beam is passing through the chamber along the x' and is then redirected along the y' axis. A $\lambda/2$

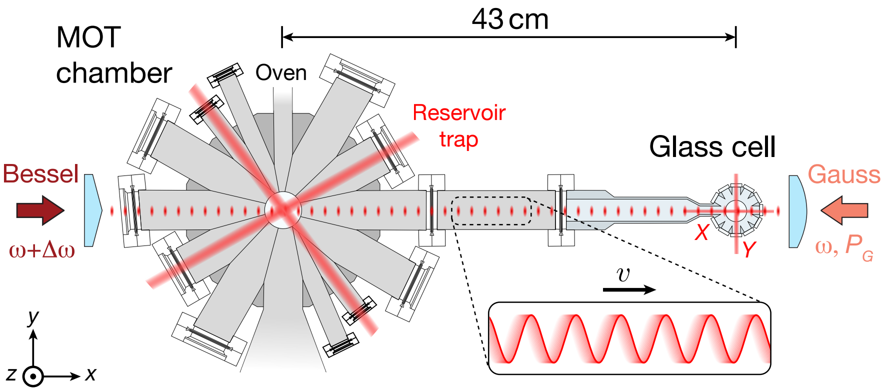

FIGURE 3.8.: Overview of the optical transport from the MOT chamber to the glass cell, also showing the dipole traps in both chambers. Figure from [96].

waveplate is used to rotate the polarisation to avoid interference between the beams. At the first pass the beam has 14 W, on the second pass about 10 W, resulting in trap depths of 8 μK and 6 μK and trapping frequencies $\omega_{x'} = 2\pi \times 12\,\text{Hz}$, $\omega_{y'} = 2\pi \times 15\,\text{Hz}$, and $\omega_z = 2\pi \times 19\,\text{Hz}$. The trap is already turned on during loading of the MOT. To support the atoms against gravity a magnetic field gradient of $31.3\,\text{G}\,\text{cm}^{-1}$ is applied, which also spin-purifies the sample as other spin states are not levitated and drop out of the trap. For loading the transport beam is ramped up to 20 W, causing atoms to accumulate in the dimple. During the process the atoms are heated as density increases and they have to maintain phase-space density (PSD). They are cooled again by collisions with colder atoms from the reservoir, which are in turn heated and lost, similar to evaporative cooling. To increase the efficiency of this process we apply a strong offset field of 120 G, increasing the scattering length to $a \approx 1500\,\text{a}_0$.

Using the dimple we load up to 4×10^6 atoms at 10 μK within 650 ms. Afterwards the dipole trap is removed and the Bessel is ramped up to full power simultaneously. We then transport the atoms to the glass cell in 25.5 ms, using accelerations $a = 2.9\,\text{km/s}^2$ and a peak velocity of $26.6\,\text{m}\,\text{s}^{-1}$. Throughout the transport we maintain polarisation by applying constant offset fields along z in the MOT and science chambers. After transport we typically have around 1.2×10^6 atoms at 2 μK in the glass cell where we load them into another crossed dipole trap as discussed in chapter 7.

4. Science chamber

As discussed in the previous chapter we separated our system into a MOT chamber, used for pre-cooling the atoms, and a science chamber to perform the experiments in.

This chapter describes the design of our science chamber and its surroundings, which were largely designed and assembled in parallel to make best use of the available space. This includes the glass cell itself (section 4.1), the high-resolution objectives and their mounts (section 4.2.1), the components to drive RF and microwave transitions (sections 4.3 and 4.4), as well as the integrated mount for the vertical lattice (section 4.5). The magnetic field coils are also a central part of this assembly. They are described in conjunction with the integrated active magnetic field stabilisation in chapter 6. The laser beams installed around the chamber are described at the beginning of chapter 7.

4.1. Glass cell

As described in section 3.1 we decided to use two independent vacuum chambers to separate pre-cooling from the experiments. The separation maximises optical access around the science chamber and allows to design each chamber to match its requirements. For all experiments with quantum gas microscope the required optical access for the objective(s) is a limiting factor. To capture as many emitted photons as possible one wants to maximise the covered solid angle, requiring a placement of the objective close to the atoms to maintain a reasonable size.

It is possible to reduce the working distance, and with it the size of the objective, by placing at least the first lens in the vacuum system [29] or use the vacuum viewports as lens [106, 107]. Both variants increase the complexity of the vacuum design as the lenses must be present already during vacuum bake-out. Additionally the vacuum lenses can usually not be aligned after installation, making it almost impossible to use a configuration with two objectives, where it is not possible to align the atoms onto the objectives.

The alternate, more common approach is to place the complete objective outside of the vacuum chamber. In this case the vacuum chamber must have viewports with suitable size close to the atomic sample to maintain a reasonable working distance. Common approaches to achieve this is by using stainless steel chambers with recessed viewports [108–110] or cells made from glass only [107, 111, 112]. For our experiment we decided to use a glass cell, partly due to its compact size compared to steel vacuum chambers, and partly due to the lack of metal in the construction, avoiding any possible complications with the generated magnetic fields.

The glass cell is a dodecagon shaped chamber made from a machined glass body

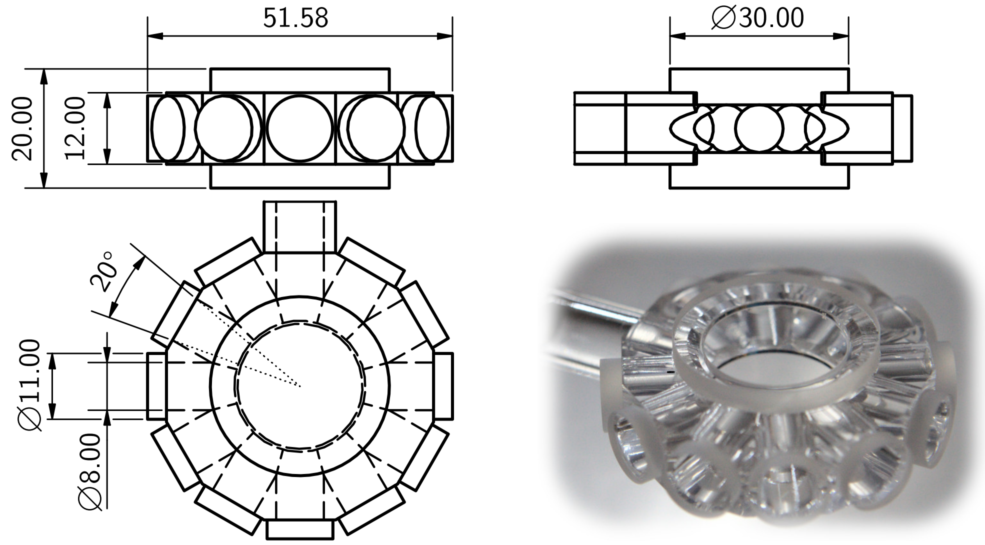

FIGURE 4.1.: Technical drawing and photograph of the installed twelve-sided glass cell with relevant dimensions (all dimensions given in millimetres). The top and bottom viewports used for high resolution imaging are recessed into the main body by 2 mm.

onto which viewports are bonded using optical contacting[1]. Two stepped viewports on the top and bottom are recessed into the chamber to be used for the high resolution objectives. The step increases the material thickness in the area needed for imaging. This was done to reduce possible deformations induced by the pressure difference, while not further increasing the height of the chamber. Horizontally eleven viewports provide optical access for the required laser beams. The twelfth opening connects the chamber through a step-wise widening glass tube to the other parts of the vacuum system, where it is connected with a standard CF40 flange.

The dimensions of our glass cell are a result of the goal of having two identical high numerical aperture (NA) objectives on the top and bottom focussed onto the same atomic sample. In this configuration the atomic sample must be in the centre of the chamber, requiring a low overall height to maintain a short working distance. At the same time the chamber must be sufficiently high to not limit the optical access for perpendicular beams too much and to maintain mechanical stability during and after machining. The height of the main body is 12 mm, with the imaging viewports extending outwards another 4 mm on either side. Due to their stepped shape the distance between the viewports in vacuum is only 8 mm.

The dodecagon shape is a compromise between the overall size of the chamber and mechanical stability. A more classical design using an octagon would have been possible, but it would have been necessary to increase its size significantly to maintain mechanical stability[2]. This would have resulted in a cell of similar size as the final design, but with

---

[1] Custom made by Precision Glassblowing of Colorado

[2] An octagon chamber in very similar design was made by the manufacturer shortly before we requested

reduced optical access due to the increased distance of the viewports from the atoms. Larger viewports were not an option as they would have increased the height of the cell. Non-circular viewports were also not possible due to the complexity of the bores needed in the main body. Instead we decided to enlarge the cell enough for a dodecagon, maximising the optical access with twelve circular viewports.

All viewports are made from UV fused silica glass and are anti-reflection coated using a nano-textured coating[3]. These are etched coatings, creating a sub-wavelength surface structure, resulting in a low reflectance over a broad frequency range, which only slowly increases at larger angles of incidence (AOI) [113]. The physical structure of the coating also withstands higher temperatures than most dielectric coatings, which occur during the process of optical contacting of the glass cell. Their disadvantage is a relatively large variation in the coating quality between production runs and a very sensitive surface. The surface structures may not be touched, either by hand or with tissues, resulting in a complex cleaning procedure using only solvents.

We chose two different wavelength regions for the coatings of our viewport. The large viewports used for the objectives use a coating matched to the objective coating, optimised for the range from 450 nm to 850 nm with $<0.1\,\%$ reflectance in this range (at 0° AOI). Above this wavelength the reflectivity increases, but it is still below 0.5 % (3 %) for 1064 nm at 0° (60°) AOI. The coating for the smaller viewports uses a coating optimised for higher wavelengths up to 1550 nm. These windows have a reflectivity $<0.1\,\%$ up to $\sim$1300 nm and about 0.3 % at 1550 nm.

The installed glass cell shows some foggy spots on a few viewports which were visible when looking onto the glass cell from specific angles. We only realised this when we were replacing the temporary coils for the final assembly. We could not tell whether these spots were already present after manufacturing or due to later contamination, e.g. during bake-out. As the cell was already mounted and evacuated we decided that they are not at critical locations as we did not want to take any risk of damaging the coating when attempting to clean the viewports.

4.2. High-resolution objectives

One of the primary design goals when building this experiment was the inclusion of a high-resolution objective to perform in-situ measurements of the atomic cloud within the optical lattice, commonly called a *quantum gas microscope*. The imaging method and its challenges is described in full detail in chapter 8, here the design of the used objectives and inclusion in the experimental setup is described.

The difficulty of in-situ imaging is highly dependent on the chosen atomic species, as it not only defines the available imaging transitions, but also indirectly the spacing of the optical lattice through the atomic mass. Caesium is in this regard a rather unfavourable choice, with a high atomic mass $m = 133\,\mathrm{u}$ requiring a short lattice spacing to achieve

---

ours. Based on the experience from this cell the manufacturer required the increased dimensions to maintain mechanical stability.

[3] TelAztec LLC

reasonable tunnelling coupling strengths, and a long imaging wavelength of 852 nm for the D2 line.

We decided to create optical lattices using light with a wavelength of 767 nm (see section 7.2.3), resulting in a lattice spacing $a = 383.5\,\mathrm{nm}$. Using this spacing and light from the D2 transition it is not possible to resolve individual atoms directly in the lattice. Using the standard Rayleigh criterion the resolution R of a microscope objective is given as $R \approx {}^{1.22\lambda}/_{2\mathrm{NA}}$, or $R \approx 520\,\mathrm{nm}$ for a perfect NA of 1[4], still much larger than the lattice spacing. With reconstruction techniques using additional information about the lattice structure it is possible to extract the occupation even beyond the Rayleigh criterion. Typically these algorithms are used in experiments where $a < R < 1.5a$ [106, 111, 115]. To help with reconstruction as much as possible and leave some tolerance for misalignment we decided to use an objective with a very high NA of 0.8, resulting in $R \approx 650\,\mathrm{nm} \approx 1.7a$.

The installed objectives were custom designed and manufactured externally[5]. They were optimised for diffraction limited performance for the D2 imaging lines of Caesium (852 nm) and Rubidium (780 nm) as well as the blue $6S_{1/2} \rightarrow 7P_{3/2}$ imaging transition of Caesium at 456 nm. To create repulsive potentials 532 nm was included due to the availability of high quality, high power laser sources at that wavelength[6]. All optical surfaces are anti-reflection coated for these wavelengths. We got two identical objectives which are designed to be used in a symmetric configuration from top and bottom.

The chosen high NA, in conjunction with the large effective working distance of $\sim$10.6 mm[7], resulted in large mechanical dimensions with a total length of almost 120 mm and a diameter of 70 mm.

At the time of writing, everything is prepared for the simultaneous installation of both objectives, but so far only the lower objective mount is used. We had both objectives installed at different times, but then we had to find out that they got damaged during transport and needed to be repaired (section 8.1.1). Once we got the first objective repaired we swapped it with the installed objective, shipping it for repair.

4.2.1. Objective mount

To achieve the theoretical capabilities of our objective it must be correctly aligned onto the atomic sample to very low tolerances. The planned symmetric configuration of two objectives requires full control over the objectives angle and rotation, to match the optical axes and foci of both objectives. It is necessary to position the objective in all three spacial

---

[4] Values larger than one can be achieved in classical microscopy using immersion techniques, but these are not compatible with UHV conditions. Similar enhancements are possible when placing the atom very close to the first optical element, at the cost of increased aberrations [114].

[5] by Special Optics

[6] The original plan assumed the usage of a 2 dimensional acousto-optical deflector (2D-AOD) for this application. By now we switched to a broadband laser source at 520 nm as we decided to use digital micromirror devices (DMDs) instead, requiring an incoherent light source for illumination to reduce interference effects.

[7] The real working distance is 12.34 mm, of which 6 mm are the glass cell windows made from UV fused silica glass.

directions and control its rotation perpendicular to the optical axis (commonly called tip and tilt).

As this is a common challenge when aligning objectives commercial five axis mounts are available, specially designed for this task. None of these suited our needs, as they could either not support our large objectives, were not motorised, or were too large to fit into the available space. Instead we designed and built our own 5 axis motorised translation stage. It uses piezo based linear actuators[8] for motion in all axes. The microscope is mounted in a free-moving central part which we call *slider* (fig. 4.2a). It rests on three motorised adjustment screws providing full angle control and translation along the optical axis. The slider is positioned horizontally by one motorised adjustment screw counteracted by springs.

The objectives themselves are placed in a secondary assembly together with piezo stacks (fig. 4.2b). These are intended for rapid and frequent position adjustment along the optical axis to optimise the focus.

Alignment

For horizontal translation of the optical axis we installed two linear actuators horizontally in the mounting plate. They are pushing directly against the body of the slider and are counteracted by two pairs of springs. They are mounted in the x' and y' axes, not directly along the axis observed in our current camera setup.

The vertical position and the angles of the objective are controlled by three more linear actuators mounted in the vertical direction. Their three contact points define a plane with fixed angle and distance from the mounting plane.

The adjustment screws are always mounted such that the main body of the screws is oriented away from the science chamber. We inserted titanium plates of 1 mm thickness on the contact side to prevent material deformations in the plastic frame. Due to the weight of the objectives we are not using additional springs along the vertical axis.

The 80 threads per inch (TPI) fine thread of the motorised screws provides a very fine resolution of approximately 20 nm per motor step. This is precise enough for precise alignment of the objective. The placement of the actuators relatively far away from the focus point causes a significant shift of the focus region when adjusting the angle of the objective, as the optical axis defined by the objective does not rotate around the focus. The resulting displacement of $\sim 110\,\mu\mathrm{m/mrad}$ is on the same order as the vertical travel distance of the screw used to tilt the objective.

Upon initial test of the lower translation assembly we had to find out that the motorised actuators for horizontal alignment were stuck and no longer turned when being controlled remotely. Manual rotation of the screws was possible when applying significant amounts of torque. Upon further investigation it turned out the coil mount assembly was resting on at least one of the motor housings, likely bending the screw and blocking rotation. We could not tell whether this was the case since the assembly, i.e. we lifted the coil assembly when installing the lower objective plate to its final location, or the

---

[8] Newport Picomotor

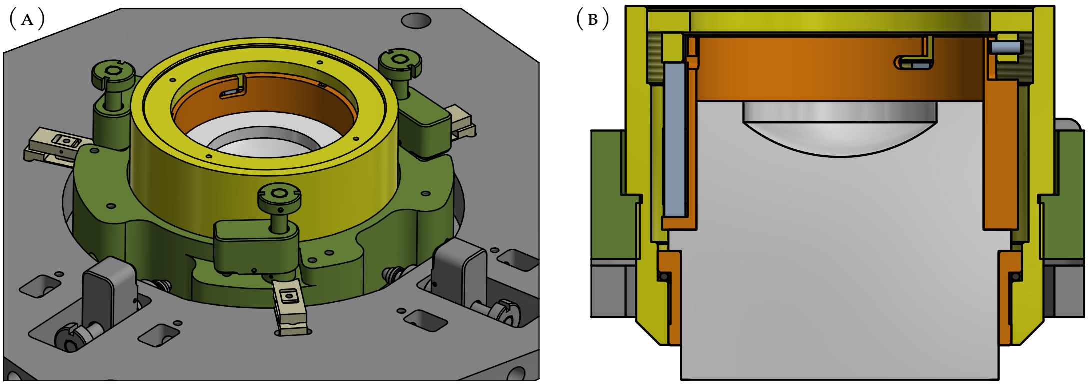

FIGURE 4.2.: (a) False-colour drawing of the upper objective mount. The objective assembly (yellow and orange) is placed in the slider (green) used for 5 axis alignment. Three piezo actuators mounted to the slider provide angle and vertical control, two more actuators mounted in the base plate (grey) are used for horizontal movements. Three clamps (beige) can be used to lock the position of the slider. (b) Cut through the centre of the objective assembly. The outer parts (yellow) are used for mounting to the slider and provide the fixed reference for the piezo stacks (blue). They translate the objective (light grey) which is mounted in the inner assembly (orange), compressing against an O-ring (black).

coil assembly had lowered over time. Before we realised this the coil assembly was only resting on the three screws used for height adjustment. Long term deformations due to the small contact area of the screws resting on a PEEK surface are likely. It was always planned to replace the screws with larger metal spacers of correct height, once the coil mount was aligned, but this was forgotten after assembly.

We installed the spacers once we realised the problem with the blocked screws. We had them machined to the heights we measured while the mount was resting on the actuators, as there were concerns that lifting the coil mount could damage the glass cell. Instead we lowered the objective plate by about 400 µm by placing folded layers of thick aluminium foil between the contact surfaces of the objective plate. With this modification we can barely pass a sheet of paper in between the coil mount and the actuator bodies, corresponding to approximately 100 µm clearance. The increased distance of the objective plate to the atoms can be compensated using the vertical adjustment screws.

Focus plane selection

After the initial alignment of the objective it is also necessary to regularly adjust its focus to compensate for unavoidable drift, i.e. due to temperature fluctuations. Precise fast control can also be used in experiments to image atoms in different focus planes independently [58]. These techniques require fast, precise and repeatable movements of the focus which can not be achieved with the motorised actuators in the slider assembly.

In our design this fast movement is realised using three multilayer piezo stacks of

36 mm length per objective[9] with a nominal travel range of 32 µm when applying 0 V to 100 V. To achieve repeatable and linear motion we control these with closed-loop piezo drivers[10]. The actuators are counteracted by a 2.5 mm rubber ring made from FKM which is compressed for pre-tension by approximately 0.5 mm.

The piezo actuators are located in a tubular assembly constructed around the objectives. It consists of two sub-assemblies moving freely along the optical axis as shown in fig. 4.2b. The inner assembly is axially constrained using two contact friction bearings made from a self-lubricating plastic. Rotation is constrained using three dowel pins.

These pins made from A4 grade stainless steel are the only metallic components in the assembly. All other parts are manufactured from glass fibre filled polyamide. We intentionally avoided metallic components to not include circular conducting structures in which Eddy currents could occur on fast switching of the magnetic field. Temporarily we considered the usage of a machinable glass ceramic Macor for the outer tube structure to increase the rigidity and reduce the risk of material creepage.

In the two years since manufacturing we could not observe any degradations of the moving plastic components. The axial confinement by the sliding bearings is still without any observable play. It was however necessary to re-cut the outer thread in the outer tube and its counterpart (the slider described in the previous sections) two times due to deformations. We attribute those to stress in the parts from machining relaxing over time. When we changed the objective one year after installation the parts had not deformed any further. To reduce the risk of having the objective stuck in the assembly we applied a thin layer of grease[11] on the thread on later assemblies.

To remove the installed objective assembly from the slider it must be locked in place using the three provided clamps. It is then possible to unscrew by hand, even though the long thread requires significant force to be used[12]. For a future design the active thread length should be reduced. It would also be beneficial to include some option to temporarily mount a tool, e.g. by adding some (threaded) holes.

Stability

To validate the stability of the objective mount we performed limited measurements in a test setup, already using the final objective mount. It was assembled using a dummy window of equal size as our glass cell windows, imaging a point pattern on a star target onto a camera. We measured the pattern on the camera for several hours, tracking the movement of the points and the resolution (fig. 4.3). After an initial drift the position fluctuations reduced to approximately 50 nm, about 15 % of the lattice spacing. The resolution shows stronger fluctuations up to 10 %, which we can attribute to temperature fluctuations as they are correlated with the temperature on the optical table.

---

[9] P-885.91 5 mm × 5 mm × 36 mm Multilayer Piezo Actuator with added strain gauge sensor, PICeramic

[10] E-610.S0, Physik Instrumente (PI)

[11] Glisseal HV

[12] On an objective change we had to improve adhesion using double sided tape.

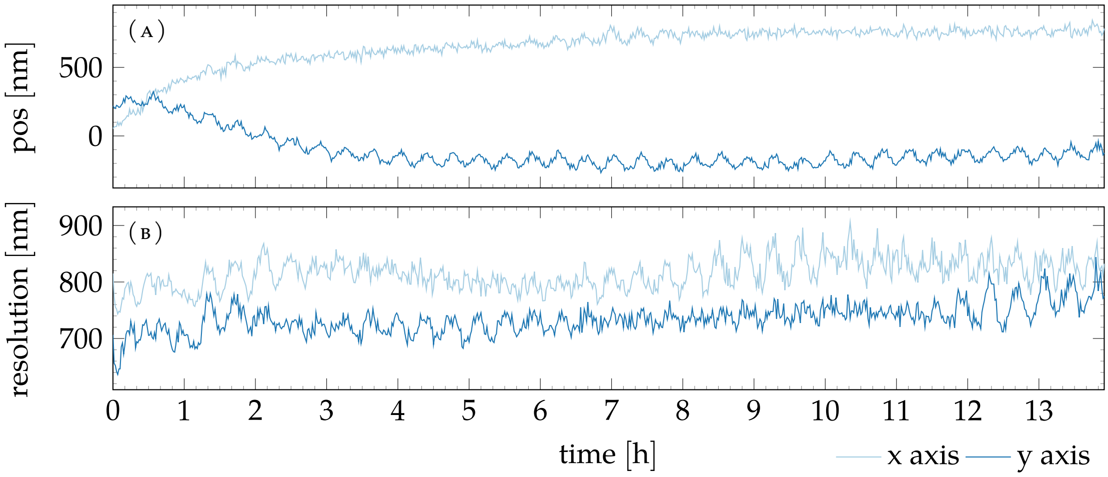

FIGURE 4.3.: Objective stability measurement performed in the test setup with the objective in the original state before repair, measuring position (a) and resolution (b). The short term fluctuations can be attributed to fluctuations in the air temperature introduced by the table air conditioning.

4.3. Microwave

Microwave radiation can be used to drive transitions between the different magnetic states of the $F = 3$ and $F = 4$ ground state manifold. The central transition frequency at 9.192 GHz at zero field is used to define the time in the International System of Units (SI) [116]. In of our experiments magnetic fields up to 50 G are used and magnetic field sensitive states are used. To address all possible transitions a bandwidth of approximately ±150 MHz is needed.

As discussed in section 3.1.3 the vacuum system is already prepared for experiments using Rubidium. We already installed a microwave horn for it as a later addition close to the chamber would be challenging. The relevant hyperfine transition frequency for Rubidium is at 6.835 GHz [75], requiring a bandwidth similar to Caesium of ±175 MHz for magnetic fields up to 50 G.

4.3.1. Microwave amplifiers

To generate the frequencies required to drive hyperfine transitions in Caesium around 9.2 GHz we assembled a heterodyne frequency generator system. We generate the desired base frequencies for pulses or sweeps using commercial function generators in the range of 600 MHz to 700 MHz which are shifted to the desired frequency band using an up-converter[13]. Its local oscillator is operated at a fixed frequency of 8.6 GHz[14]. The generated output signal passes through a fast RF switch[15] with 94 dB suppression and is

---

[13] Aldetec Inc. ALS6066

[14] Kuhne electronic GmbH, MKU LO 8-13 PLL

[15] Qualwave Inc. QPS1.8000.12000

then amplified using a 50 W power amplifier[16]. The amplified signal is passed through a 10 cm semi-rigid cable[17] to a coax-to-waveguide adapter[18]. The adapter is immediately followed by an isolator[19] to protect the amplifier from reflections, before the power is directed to the microwave horn described in the following section.

To allow for fast switching between various frequencies or sweep ranges a 4-port switch[20] is included on the input side before the up-converter. It is followed by another switch[21] used to increase the suppression on the input side. Currently only a single function generator[22] is connected to the system, which we program remotely depending on the application.

To characterise the system we measured the gain and saturation parameters for the complete system, beginning from the input switch stage to the main output amplifier. The gain of the system varies strongly from 44.3 dB to 52.8 dB with frequency from 9.18 GHz to 9.24 GHz as shown in fig. 4.4a, but the saturation power and behaviour towards saturation is identical at all frequencies. The system has a very sharp transition from the linear regime into the saturation regime, resulting in a small separation between the 1 dB compression point (P1dB) at 47.0 dBm and the 3 dB compression point (P3dB) at 48.1 dBm, which is almost identical with the saturation power (PSat) of 48.2 dBm (fig. 4.4c). This matches the curve of the power amplifier, indicating that the up-converter is always operating in the linear regime.

The full up-conversion and amplification system for both the Caesium and Rubidium microwaves, including the required power supply, are placed in a 3U 19″ rack module. The microwave components are mounted on a water-cooled aluminium plate for thermal control. The microwave stages for Rubidium were included and installed because we already had them at hand from the predecessor experiment, but they remain untested. The architecture is similar to the Caesium microwave, combining an up-converter[23] using a 6.3 GHz local oscillator[24], followed by a fast RF switch[25] and a two stage power amplifier[26]

4.3.2. Microwave horns

To deliver the generated microwave frequencies to the atoms an antenna or horn is used. As the atoms are practically concentrated in a single spot strong directional emission from a closely placed antenna is desirable.

---

[16] AM43-9-9.4-47-47, Microwave Amps Ltd

[17] Huber+Suhner minibend-4

[18] A-Info 112WCASM, 7.05 GHz to 10.0 GHz, right angle SMA to WR112 adapter

[19] Qualwave Inc. QWIB-8500-9500

[20] Mini-Circuits ZSWR-3-30DR

[21] Mini-Circuits ZSDR-230

[22] Rohde & Schwarz SMC100A

[23] Kuhne electronic GmbH KU 68 G2

[24] Kuhne electronic GmbH KU LO 6300 - 788

[25] Miteq N136ADM1

[26] Kuhne electronic GmbH KU 681 AL followed by KU PA 6800 C - 274

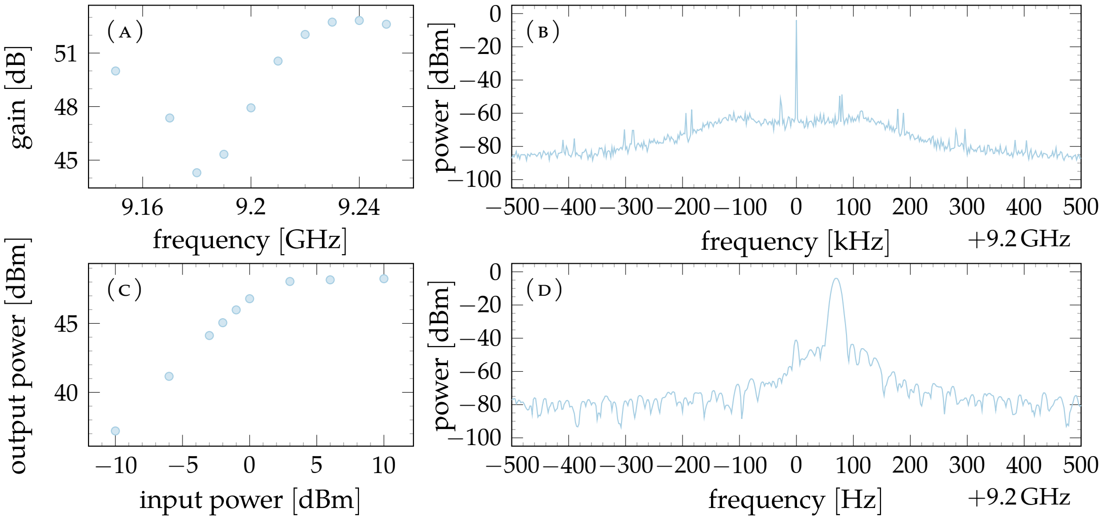

Figure 4.4.: Characterisation of the caesium microawave amplifier system. The gain of the amplifier is strongly frequency dependent (a), however the maximal output power (c) is unaffected. The frequency spectrum of the complete system (b;d) shows spurious freqeuncies at −50 dBc and a slightly increased background due to the phase noise of the PLLs. The spectrum analyser was not locked to the generator system, resulting in a minor deviation from 9.2 GHz.

For our application we decided to use a microwave horn. Horn antennas are reasonably well directed and can withstand very high power levels. A regular horn antenna consists of a coaxial cable to waveguide adapter followed by a section with expanding cross-section. The length of this expanding section determines the beam shaping and the effective gain of the antenna.

One of the biggest concerns when designing the microwave antenna was the limited space around the experiment chamber, as optical access must not be reduced. As the microwave frequencies are blocked by metals only a few openings are usable. Either vertical passing through the microscope, or in the horizontal plane through the gap between the z coils. We decided against the vertical access due to the long distance between the atoms and a possible horn location. To place the horns as close as possible to the atoms we decided to not use an expanding section, but use an ending waveguide instead. This reduces the directional properties, but allows us to place the horn as close as possible to the atoms.

With the ends of the antenna already being a waveguide we decided to use waveguides for as long as possible, minimising the length of coaxial cables. At the required frequencies around 9 GHz even semi-rigid coaxial cables have cable losses of ∼1.5 dB/m. To connect the antenna directly to a closely placed power amplifier we would have needed at least 2 m of cable. The resulting wire loss would have been about 50 %, not only reducing the available power, but also considerably heating up the cable.

As cable losses at higher frequencies increase waveguides become more viable as their cross-section is determined by the wavelength. Standard waveguides are rectangular with a 2 : 1 ratio with values chosen for overlapping frequency bands. These and many common RF components are available with standardised connection flanges, allowing to design and build a waveguide system similar to a vacuum system.

The matching waveguide size to drive the Caesium hyperfine transition at 9.2 GHz is WR112 / WG15. It has fairly compact inner dimensions of 28.5 mm × 12.6 mm. We use four waveguides with a total length of about 8 m, two rigid assemblies along the ceiling of the lab, and two flexible waveguides[27] on the ends to connect to the movable amplifier and the floating optical table. This allows us to place the power amplifier in a better suited location where it is easier to cool. At the same time the losses are considerably lower with only ∼1 dB in the waveguides and additional 0.3 dB in the 100 mm semi-rigid cable from the power amplifier to the adapter.

3D printed horns

To place the microwave horns in the available space we created a waveguide with a sequence of bends to fit around other components (fig. 4.5b). While it would have been possible to manufacture a similar shaped horn it would have been a complex assembly with many solder joints. It would also result in a large, conducting structure with a loop shape directly next to our magnetic field coils, most likely creating problems with Eddy currents when switching the fields.

Instead we had the horns 3d printed from plastic and coated with copper. 3d printing can easily deal with the complex shape of the bent waveguide, even including the mounting structure. Special surface treatment allows galvanisation of the printed parts, creating an 50 µm layer of copper on the outside. In a waveguide the field is contained inside the conducting tube, only entering the material due to the skin effect. For copper at these frequencies the skin depth is below 1 µm, much smaller than the thickness of the material.

Other groups focussing on microwave technology have already reported the successful creation of simple waveguides and even waveguide filters using 3d printed parts at similar frequencies [117, 118], even manufacturing procedures are described [119]. In contrast to these reports our setup has the additional complication that the resulting shapes are curved. When the parts are 3d printed they are assembled in thin layers, producing tiny steps instead of perfectly flat surfaces. To reduce losses due to surface imperfections the parts were smoothed manually by the manufacturer before applying the copper plating.

The tube had to be split into two parts for manufacturing as the surface treatment before galvanisation requires direct access to the surfaces. We split the horn into a main part containing three walls of the tube and a secondary lid close the profile. The two parts were glued together using silver-filled epoxy, curing at room temperature[28]. The two parts were aligned and pressed together during curing using rows of M2 screws in regular intervals on either side.

Two similar microwave horns were built, one matching WR112 dimensions for the Caesium frequencies, another matching WR135 for the lower Rubidium frequencies at

---

[27] Flexguide FT15MM-XXXX-N-FK

[28] Creative Materials GPC-251

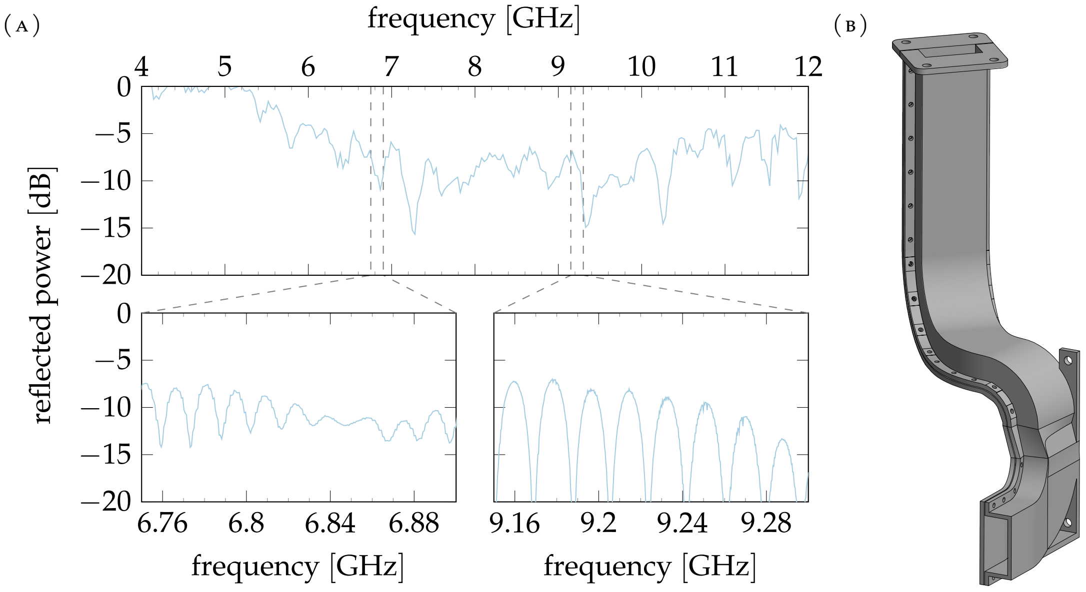

Figure 4.5.: (a) Measured reflected power from the Caesium microwave waveguide and horn. Data is the difference between a reference measurement with open cable and the cable attached to the horn assembly. The resonance peaks match a cavity of the length of the wave guide assembly. The circulator was removed for this measurement. (b) Rendering of the installed 3d printed and galvanised microwave horn.

6.8 GHz. They are placed on opposite sites of the experiment, facing each other. At the moment the Rubidium horn is installed, but not connected to any amplifier.

Characterisation

Due to limited measurement equipment available in our group suitable for the range of 9.2 GHz, we could only measure the reflected power of the installed horn, also called the S11 parameter. The measurement was done using a separate function generator[29] and spectrum analyser[30]. The reflected power was extracted using a directional coupler[31]. To get a meaningful reading of the emitted power from the antenna the circulator was removed for this measurement[32].

The measurement system was calibrated by measuring the reflected power from the cable assembly with unconnected SMA connector, representing an open ending to the transmission line. For the measurement this cable was connected to the coaxial to waveguide interface. Figure 4.5a shows the acquired broadband scan over the accessible

---

[29] Agilent N5183A

[30] Rhode & Schwarz FSP-40

[31] Mini-Circuits ZUDC10-02183-S+

[32] The third port of the circulator used has an integrated 50 Ω terminator to dump the reflected power. This prevented us from directly measuring the reflected power extracted by the circulator.

frequency range and finer scans around the Rubidium and Caesium transitions.

The scans with higher resolution show regular intervals where the reflected power drops strongly. Their spacing correspond to a resonator with a length of about 8 m, matching the length of the waveguide assembly. These will not be present in the assembled system as the circulator strongly suppresses the reflected power, breaking the cavity. The real reflected power can be approximated by connecting the maxima of the measured reflectivity. These also match the values acquired in the broad frequency scan.

The primary Caesium frequencies are unfortunately at a frequency range with relatively high reflected powers in the range of $-7\,\mathrm{dB}$ to $-8\,\mathrm{dB}$, corresponding to 80 % to 85 % emitted power.

To validate the measured data in the experimental setup we performed comparison measurements using microwave spectroscopy as described in section 6.4.2. We compared the 3d printed horn for Caesium against a standard 10 dB horn antenna[33]. The horn antenna was temporarily connected to the same flexible waveguide usually connected to the 3d printed horn and pointed at the atoms through the opening for the upper objective. For both configurations we measured the Rabi frequency using a thermal cloud in the crossed dipole trap with identical measurement parameters. The Rabi frequencies measured using the horn antenna were only 58 % of the frequencies measured using the 3d printed horn, corresponding to approximately 3 times lower microwave power at the atoms. This reduced coupling is within the expectations as the temporarily placed horn was further away from the atoms. We interpret the comparable order of magnitude of both horns as confirmation that the 3d printed horn does emit the power as expected and is not introducing significant absorptive losses to the assembly.

With the system we are able to achieve Rabi frequencies $\Omega_0 = 2\pi \times 20\,\mathrm{kHz}$ between $F = 3$, $m_F = 3$ and $F = 4$, $m_F = 4$, currently limited by coupling of the microwave frequency into the offset lock of the repumper laser. Without that limitation we expect to reach $\Omega_0 \approx 2\pi \times 60\,\mathrm{kHz}$.

4.4. RF coils

In addition to the microwave system described in the previous section we installed coils for the generation of radio frequencies (RFs) to drive direct transitions between hyperfine spin states. Their energy shift is only given by the magnetic field due to the Zeeman effect $\Delta E_{Zee}/h \approx \Delta_{mF} \times 350\,\mathrm{kHz/G}$ for ^{133}Cs. The resulting frequencies are low MHz range, with up to 20 MHz for the standard magnetic fields up to 50 G and $\Delta_{mF} \leq 3$.

Due to the long wavelength of these frequencies the antenna is operating with in the near-field regime and the coupling is dominated by the oscillating magnetic field. The primary factor to maximise the coupling strength is a close proximity to the atoms [73].

In our case we installed two coils for RF generation as close as possible to the atoms around the top and bottom viewport as shown in fig. 6.3 on page 64. We decided to install two coils with different windings into the available space, one with only two, another with ten windings.

---

[33] A-Info LB-112 7.05 GHz to 10.0 GHz WR112 standard gain horn antenna

The coils are intended to be driven using a matching network to increase the coupling efficiency, but can also be used without one over a broad frequency range at lower efficiency. Our amplifier system can combine up to two frequency sources using a frequency combiner[34] before sending the signal through a power amplifier[35] and to the RF coil. The system was not tested since the installation as it was not strictly necessary, and was found to be not worth the effort at the time.

4.5. Vertical lattice

To create strong vertical confinement during fluorescence imaging we included a compact vertical lattice mount in the design of the experiment chamber. We can not use a retro-reflected lattice, as the vertical axis is blocked by the objectives. Instead we use a running wave lattice under the largest opening angle available in our setup. Using a half-opening angle $\alpha = 30^\circ$ beams can pass by very close to the objectives and enter the glass cell through the large viewports on top and bottom (see fig. 4.6), resulting in a lattice with spacing

$$a_{lat} = \frac{\lambda}{2\sin(30^\circ)} = \lambda. \tag{4.1}$$

We designed and built a compact mount to create a vertical confinement lattice using this configuration. Matching the shallow vertical lattice we installed (see section 7.2.2) we call it steep vertical lattice (steep VL). Both lattices are running wave lattices, hence their phase is given by the path length difference of the two interfering beams in the interference region. This configuration tends to be more less stable than retro-reflected lattices, as the relevant beam paths are often longer and contain more components. This increases the influence of temperature fluctuations — either through refractive index changes or mechanical movement due to thermal expansion — and provides more possibilities for mechanical vibrations to couple into the beams.

Active phase stabilisation can be used to suppress external noise by applying feedback on a control element in one of the paths. This technique has been demonstrated to work in certain applications (i.e. [120, 121]), but is generally limited due to the inability to measure the phase at the atoms, necessarily introducing errors. Instead most experiments are designed to be passively as stable as possible [83, 122].

We decided to follow the same route and designed and built a monolithic mount which splits an incoming beam into two arms and directs them to the atoms on the shortest possible path in a mostly mirrored construction. The incoming beam is first passing through a 250 mm lens, focussing a collimated beam onto the atoms. After two mirrors the beam is split into the two paths using a PBS. A zero-order $\lambda/2$ waveplate rotates the polarisation of the passing beam from $\parallel$ to $\perp$ and is the only element present in only one of the paths. Afterwards the polarisation is additionally cleaned in both paths using one PBS each. These are rotated by 90° relative to the splitting PBS, transmitting $\perp$ light

---

[34] ZMSC-2-1, Mini-Circuits

[35] ZHL-3A, 1 W, 0.4 MHz to 150 MHz, Mini-Circuits

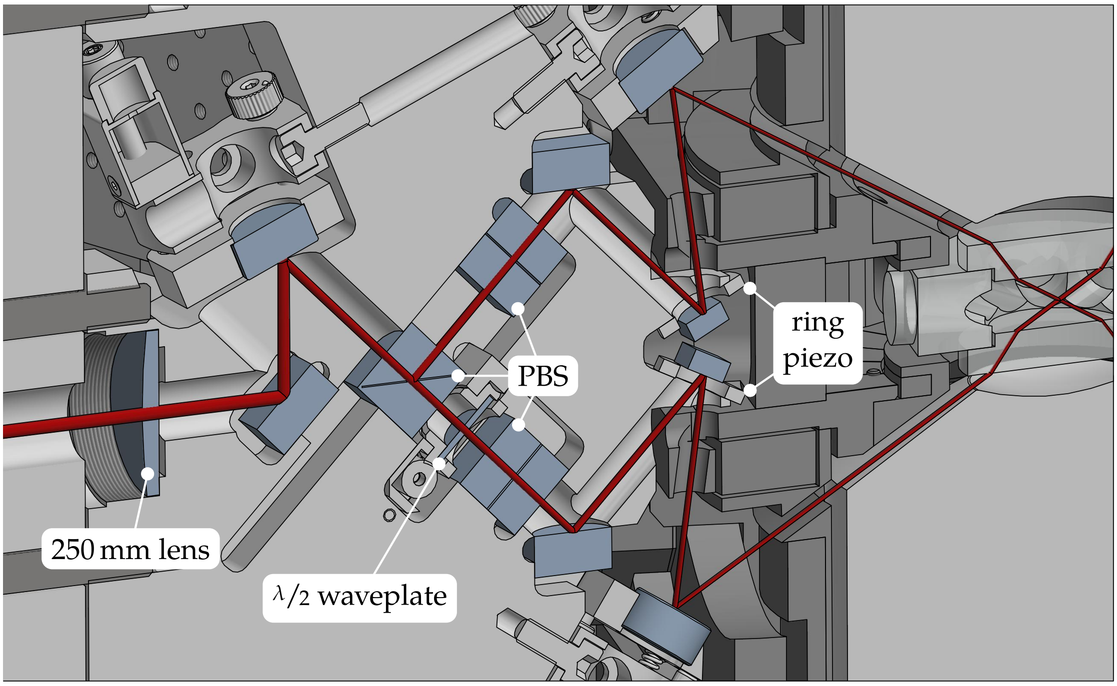

FIGURE 4.6.: Beam path of the steep vertical lattice beams. The single beam is focussed on entry using a 250 mm lens. A PBS is used to split the paths, followed by a zero-order λ/2 waveplate and further PBSs for polarisation cleaning. The second to last mirrors are glued to ring piezo actuators for lattice phase control. The lattice mount is not in contact with the coil mount drawn in dark grey.

and reflecting $\parallel$ out to the side of the mount (polarisations defined with respect to the propagation plane of the beams). The beams are then directed into the chamber using three mirrors each. One of these is mounted onto a piezoelectric ring, to adjust the path length, and set the lattice phase. The length of the separated path length between the splitting PBS and the atoms is only ~180 mm. Similar to the z dipole trap described below the beams are intended to be elliptical with a $2:1$ ratio to compensate the AOI, creating a round potential in the horizontal plane.

To avoid possible phase shifts due to thermal expansion the main body of the mount is manufactured from a single piece of Macor, a machinable glass ceramic with very low thermal expansion coefficient ($8.1 \times 10^{-6}\,°\mathrm{C}^{-1}$ at lab temperatures). All components except three mirrors mounted on motorised mirror mounts[36] are glued into the mount using a two component epoxy resin[37]. Most of the beam path is within the mount to protect it from air drafts.

The optical components in the assembled and installed lattice mount are all anti-reflection coated and optimised for 1064 nm light. The original plan was to use 532 nm,

---

[36] Newport Picomotor 8885

[37] Uhu plus schnellfest

we decided to increase the wavelength and with it the lattice spacing to allow for easier loading of the atomic sample into a single plane. We chose 1064 nm as we had the lasers at this wavelength available and it would be one of the wavelengths used in a future upgrade to a vertical super-lattice [83].

The change to the longer wavelength has the side-effect of an increased reflectivity of the glass cell windows, as it is outside of the specified range of the nano-structured coating. The measured reflectivity at 0° AOI is only 0.5 %, but these beams enter the chamber under a very shallow angle of 60°. When passing the 1064 nm beam through the glass cell we can achieve a maximal transmission of 80 %, corresponding to a reflectivity of ~5 % per surface. Parts of the reflected beam are scattered within the objectives. So far we could not observe any negatives effects from this.

The complete assembly is mounted directly to one of the vertical main pillars using screws, the incoming beam passes through a prepared hole in the pillar. The intent was to be able to place the mount with the pillars, coils and breadboards in place, but the tolerance turned out to be insufficient by less than a millimetre, and the mount can not be mounted or removed without moving the front breadboard.

Lattice alignment The compact size of the mount makes alignment of the beams onto the atoms relatively complicated as it is not possible to reflect the beams out of the plane, except through the polarisation cleaning cubes. To allow for alignment of the beams when the table is closed and thermalised three mirrors are mounted in motorised mirror mounts. Two of these are the last mirror in each beam path before they are directed onto the atoms. The third is a shared mirror directly before the beam splitter.

For optimal alignment it is necessary to overlap the beams at the atoms under the correct angles to create lattice planes perpendicular to gravity. This requires both beams to enter the chamber under the same angle from the horizontal plane, and within a perfectly vertical plane. These constraints must be fulfilled by the alignment of the beams onto the last mirrors, which can then be used to overlap the beams on the atoms. The vertical separation is defined by the position of the beam on the splitting PBS, which is fixed by the height and vertical angle of the incoming beam. Its alignment is more challenging due to the focussing lens at the beginning of the mount, but can be corrected using the third motorised mirror mount. The horizontal alignment of the incoming beam is less relevant, only changing the horizontal position on both last mirrors simultaneously, rotating the vertical plane of incidence. It is not possible to change the angle of this vertical plane, as it would require rotation of the splitting PBS, which is glued into place. It is referenced to the mount, which in turn is referenced against the vertical pillars using a long reference edge. The alignment should be sufficient due to the manufacturing tolerances, but this could not be verified so far as the lattice beams were never fully aligned yet, and it is not possible to resolve the ~1 µm spacing using any commercially available camera.

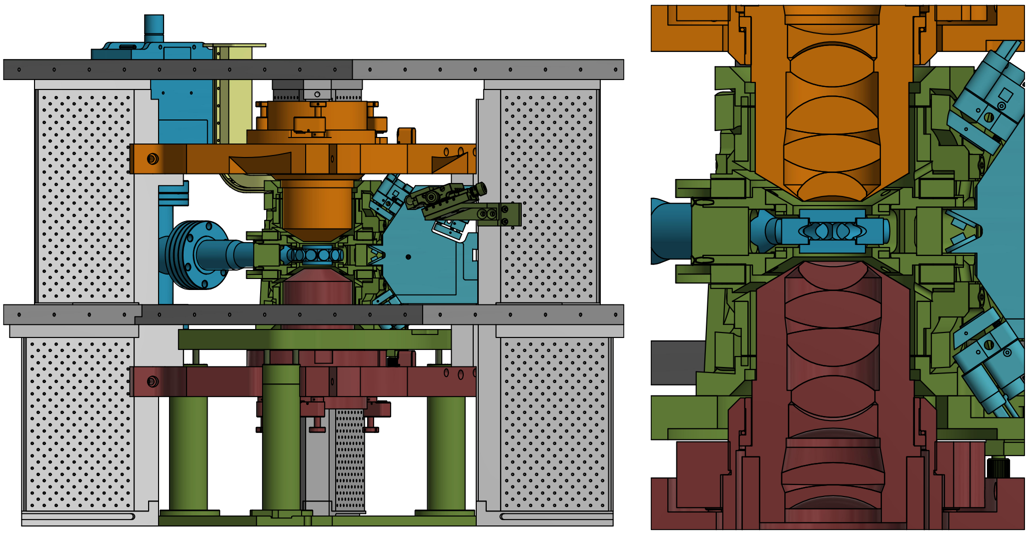

FIGURE 4.7.: Drawing of components around the science chamber, illustrating the independent mounting blocks. The objective mounts (red & orange), the microwave horns (yellow), and steep vertical lattice (light blue) are independent of each other and mounted directly to the pillars (white), which also support the breadboards (grey). The magnetic field assembly (green) is mounted independently and fixed to the table, except the X sensor mount (dark green), which is also mounted to the pillars.. The glass cell (dark blue) is attached to the MOT chamber and supported on the four-way cross.

4.6. Mount

The various components described in this chapter were all designed together to allow for compact placement around the science chamber. Throughout the design process we ensured rigid mounting of the optical components, while decoupling them from the glass cell and magnetic field coils.

Figure 4.7 highlights the various independent sections which are mounted around the science chamber. Most components are mounted to four triangular vertical pillars, which are manufactured from multiple pieces of fibre-glass reinforced epoxy material[38]. They are glued together to form a hollow structure, which we filled with a mixture of sand and lead balls to improve vibration dampening. The sides of the pillar can be used as breadboard to mount optical components horizontally on a 30 mm beam height.

The pillars act as reference for the various components holding optics, namely the two objectives, the steep vertical lattice and the middle and upper breadboard layer. The 20 mm thick breadboards were manufactured from the same material as the pillars, again to reduce the amount of metals around the science chamber. They are supported further away from the chamber using post made from 38 mm outer diameter tubes[39], filled with the same lead and sand mixture as the pillars. The top breadboard is made from a single

---

[38] EP GC 203

[39] compatible with Thorlabs 1.5″ posts.

800 mm × 1100 mm plate, the breadboard for the horizontal level is assembled from three individual pieces, sliding into the triangles formed by the pillars. These are fixed to all pillars and are overlapping each other and screwed together, to create a single rigid structure.

In contrast the magnetic field coil assembly is not touching any of the optic mounts and is directly mounted to the experiment table using three independent legs made from PEEK, also filled with a lead and sand mixture. This is done to avoid direct coupling of mechanical vibrations from the cooling water or due to fast coil switches into the atoms through the optical components.

Within the assembly the glass cell is freely floating, only supported through the vacuum connections on the MOT chamber and a support post below the 4-way cross in front of the glass cell. We can observe vibrations of the glass cell, shifting the observed interference pattern from the imaging beam along the vertical axis from image to image. We hope to remove these by adding some dampening to the support post on the 4-way cross in the future.

5. Experiment control

Performing experiments using ultra-cold quantum gases requires the correct interplay of a multitude of different components, many of which are described in the previous chapters. Not only must they be synchronised to each other for a single experimental realisation, but their relative timing must remain identical for many realisations to gather the needed data for statistical evaluations.

Throughout the years various groups implemented their own experiment control systems, some of them publicly documented and available [123–125] (see ref. [126] for an overview and comparison). In our group a similar software based on Python called *QControl* was developed. It was rewritten for Python 3 with some modifications to the design as *QControl3*.

Throughout the construction of the experiment we implemented significant software changes, improving the overall code quality, while also extending the capabilities of the software. All changes had to be made in the context of the already existing software. We attempted to maintain compatibility with existing installations and packages (foremost the driver packages) as much as possible. In many cases we achieved this goal, but not always, as the original 'specifications' sometimes turned out to be incomplete, neglecting important cases, or to be not present at all. Missing unit tests throughout the code base increased the required time for refactoring, as regressions could not be detected efficiently.

The information provided in this chapter refers to the latest version, providing a general introduction into the software and its architecture. It is intended to be used as a starting point for users and developers alike, giving an overview over the software, its different components and their interactions. Further documentation is available within the codebase as documentation strings for modules, classes and functions. The software with its various components uses an open-source license[1]. However the code is currently not publicly available.

Once the software was usable, coordinated development of the core software came to a rest. The remaining work focussed on new device drivers where necessary, and infrequent bug fixes. Further developments were primarily done in our laboratory, with valuable feedback and tests from the laboratory next door. To reduce the influence of the ongoing software development on the regular lab operations we significantly extended the test coverage and introduced a complete changelog with tagged releases following the semantic versioning standard. We primarily extended the server and hardware control code, ensuring compatibility with existing driver packages. Only limited modifications to the client applications were made and the software is still lacking a complete user

---

[1] GNU General Public License Version 3

interface. This would greatly benefit from a development plan and in general a larger user base.

The software as described in this thesis is currently in use in two different experiments, with at least three other experiments using older versions of the code base. Hardware support for many standard devices is present and some basic user interfaces have been created.

5.1. Design principles

The primary task of the control software is the orchestration of the sequence execution. Operation is organised in units of independent experiment realisations, called *shots*. These are grouped in larger units called *tasks* for organisation. The execution of a single shot is also called a *run*. To maximise the measurement time preparation and finalisation of shots are done in parallel to the execution of shots. A queue is used to keep track of the shots to be performed. Shots are prepared and compiled prior to their execution, and acquired measurement data is stored once a shot has finished. The user interface provides an application programming interface (API) for external software to query and control the general system status, and to modify the state of the sequence queue.

Every parameter of the attached devices must be expressed through a channel with a fixed type. Experiment sequences are programmed as series of value changes on these channels at user specified timestamps. The execution of these changes is performed by the matching device driver. Sequences are programmed in scripts using Python code where the defined channels are made available to the user, described in appendix A.1. The software imports and executes the code and checks the requested channel changes for errors, e.g. limit violations. Channels expressing physical quantities can be created with an associated unit[2], allowing automatic unit checking and conversion. As most users of the software are Physicists with typically limited programming experience the goal was to provide an easy to use sequence interface.

A modular software design is used to support the various devices present in experimental laboratories. Individual device driver packages are used to implement the hardware specific functionality on top of the fixed set of channel types. In general the software only implements the core functionality, only providing limited help or templates for new device drivers or sequence files. The few restrictions provide for many possibilities to use the software, but also pose a challenge to inexperienced users.

The software is built using 'classical' quantum gas experiments in mind, where sequence programming and execution are independent, i.e. the sequence is not dependent on measurements made during execution. This assumption is true for our and most other experiments and allows for many simplifications and optimisations within the software. It is however not true for experiments with live feedback, such as experiments with optical tweezers performing detection and rearrangement to create defect-free arrays [23, 128]. These parts of the sequence require additional logic outside the scope of the control software.

---

[2] implemented using `qunits`, a custom fork of the unmaintained library `misu` [127]

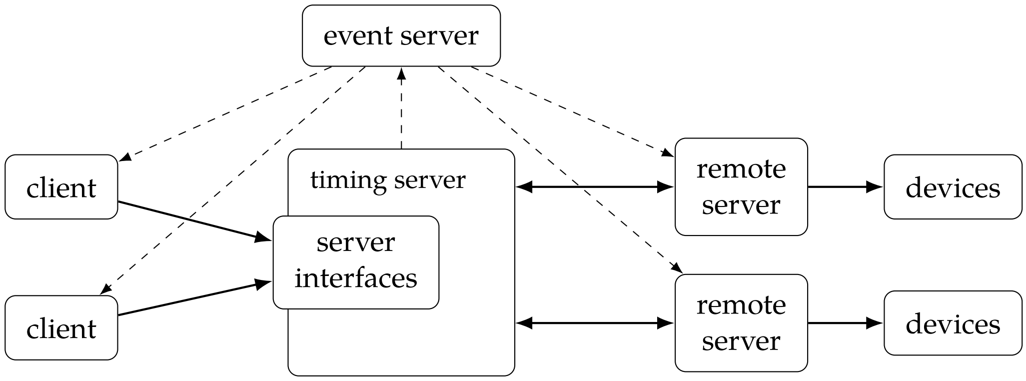

FIGURE 5.1.: Overview of the components of a QControl3 installation. The timing server as central component communicates with client programs and the configured remote servers using network protocols. Remote servers load the driver packages and control the hardware. A single event server is used to communicate updates about the server status to the clients and remote servers.

5.2. Architecture

For flexible installation in the lab environment and to circumvent Pythons general limitations with multithreading the software is split into into multiple independent processes. A single, central server process is used for coordination of the various steps of the sequence processing and execution, also providing the API for the user interface. Device drivers are run in independent processes called remote servers or simply *remotes*[3]. They communicate with the server using a network protocol and can be started on any computer with a network connection to the server. This separation allows to start remote servers on computers where hardware is physically connected, avoiding additional network protocol implementations on the driver level.

An additional event server process is used to communicate status events between the various server components, but also to provide information about system events to attached user interfaces. It is running independently from all other components, and is only responsible for event distribution. The event system is built on top of the asynchronous messaging library ZeroMQ [129].

The central server and remote servers are written using Pythons `asyncio` support. It provides an efficient method of writing concurrent code in environments which are not limited by processor performance. The execution is not truly parallel, but appears parallel. Whenever a function needs to wait for a potentially long running operation it must do so using an asynchronous wait (`await`). The framework then pauses the function and executes other code, until the long running operation has finished. At this point execution continues in the original function. Using this approach functions appear to be executed sequentially and allow for local reasoning like regular functions, but they

---

[3] Originally device drivers could run either in the main server process or on a remote process, hence the name.

can be executed efficiently in parallel. In the context of this software `asyncio` primarily handles the network communication with the different server components and user interfaces and simplifies the parallel processing of sequence preparation, execution and finalisation.

For optimal operation all functions called in the `asyncio` context must be quick, or using the asynchronous wait functionality of the framework. Blocking or long running functions not using the asynchronous wait prevent the execution of other functions, breaking the core assumption of parallel execution. In the current implementation this is true for most functionality, with the exception of the sequence parsing and compilation. These are planned to be moved to dedicated processes, but are currently still executed in the server and remote servers respectively.

Even though the core functionality of the server is using the `asyncio` framework it is not available to device drivers yet. Most drivers were written before the introduction of the asynchronous framework and only a few would directly benefit as the most hardware drivers do not support asynchronous communication by default. Instead threads are used where required.

5.2.1. Configuration

The different server and client process of the software share a common configuration file. In addition the server processes use Python files to configure the attached hardware and the configured user channels.

The main configuration file uses YAML [130] and is located in an operating system dependent standard configuration folder[4]. It is used to define standard options used by all programs, such as the event server address and the logging configuration. It is also used extensively by the server to configure data paths, the modules to load and the remote servers to connect to.

Packages can extend the configuration with their own entries by providing a description of the entries used for validation and default values. By convention each package should use its name as sub path to avoid conflicts, e.g. `qcontrol3_interface_http` adds configuration entries for the host and port in `interface` / `http`.

In addition to the central configuration file, each remote server has an additional configuration file to configure the attached hardware, its location is defined in the central configuration. It is a Python file which must define a function `hardware_setup`. This function gets called with the timing system, the file must then configure and attach the timing controller and attached devices. To allow access by the process the file must be available to the remote server, the hardware configuration is automatically mirrored to the central server when the connection is established.

The user channels on the main server are configured similar to the hardware configuration. The Python configuration file must provide a `channel_setup` function, which receives the timing system with the full timing tree (section 5.3.1).

---

[4] `~/.config/qcontrol3/` on Unix and `%APPDATA%\qcontrol3\` on Windows. Executing `qcontrol3-config` displays the file location and the currently active configuration.

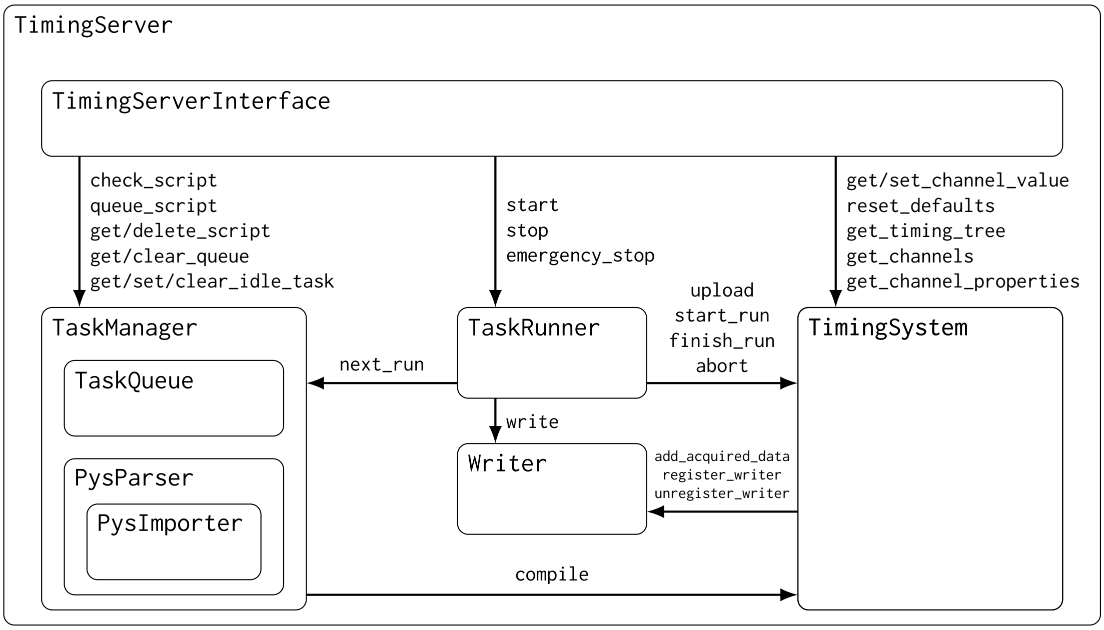

Figure 5.2.: Diagram of the *QControl3* server components and their interaction. The arrows denote function calls from the caller to the callee. The `TaskRunner` coordinates the sequence execution, requesting shots to execute from the queue managed by the `TaskManager`. Execution is performed using the `TimingSystem`, which communicates with the remote servers. Clients communicate with the different components through the `TimingServerInterface`.

5.3. Server components

The largest part of the server code is needed to handle sequences provided by the user and to make sure they get executed as intended. All code for the server can be found in the module `qcontrol3.server` in the package `qcontrol3_server`.

The central class of the timing server is the `TimingServer`. It is responsible to start and stop the system and prepares the different server components as shown in fig. 5.2: The `TaskManager`, used for sequence importing and preparation, the `TaskRunner`, used for sequence execution, and the `TimingSystem`, representing the hardware nodes and used for communication with the remote servers. It also provides user access to relevant functions of all three components through the `TimingServerInterface`. This class only provides internal functions to be used by interface modules which in turn provide user access through various protocols, e.g. a HTTP based REST interface. Multiple interface modules can be loaded in parallel and provide access to the same server. This was done to allow introduction of a new interface, while maintaining compatibility with existing installations.

The primary class for every sequence related task is the `TaskManager`. It provides the functions for the user interface to check scripts for correctness, to add them to the queue and to retrieve the next for execution. Received sequence files are imported and translated into an event table using the `PysParser`. It imports the sequence code as special Python modules with the help of the `PysImporter`, before it executes the script to generate the event table. The queued tasks are stored and organised in a class internal instance of the `TaskQueue`, which maintains a queue with priorities for tasks and another execution queue with the list of shots to be executed next. A background job prepares the upcoming shots for execution, ensuring their event table is generated and compiled before they are needed for execution.

Execution of the shots is organised by the `TaskRunner`. It requests individual shots for execution from the `TaskManager` and coordinates the different steps of execution, until the shot is fully finalised and any acquired data has been processed. Data processing is performed by extra modules and data can be processed multiple times using different methods, e.g. stored to disk in HDF5 format and sent out via the event system.

To ensure correct clean up of the sequence data and deallocation of the acquired resources, references to shots are only held when a component is responsible for it. All other locations are using weak references from the `weakref` module. Initially the queue is responsible for all stored shots until they are retrieved for execution. At this moment responsibility is passed over to the task runner. Once the shot is finalised, i.e. all acquired data was written, the runner releases its handle, allowing clean up of the shot by the garbage collector. Similarly references to tasks are only held by its shots, resulting in clean up as soon as all shots are finished.

The `TimingSystem` is the root node of the timing tree, representing the configured hardware devices and their channels. It provides information about the configured hardware and its channels and allows direct control over the channels when no sequence output is running. When sequences are running it starts and monitors their execution. To do so it handles the communication with the remote server, each providing parts of the timing tree.

5.3.1. Timing tree

The configured hardware devices are organised in a tree structure with five defined layers of node types.

The top layer with the single root node is the timing system. It has the logic to initially configure the tree and acts as a link between the devices and the other server components, namely the interface and the task runner. It also collects data acquired during sequence execution and passes it on to the writer.

On the second layer an arbitrary number of timing controllers can be installed, each running in an independent remote server process. Timing controllers represent special devices with the capability to run commands at user requested times within the sequence. Most devices do not have this capability, and are instead programmed or triggered from some other component with this capability. In all experiments using this software only a single real-time capable hardware[5] implements a special timing controller and synchronises the experiment using trigger pulses. All other devices use the included `SoftwareTimingController`, which uses the computers system clock.

---

[5] based on National Instruments data acquisition cards, or the Adwin system from Jäger

The third layers represents individual devices controlled by its parent timing controller. This layer is the only layer with a variable depth, as devices can have sub devices, e.g. a multi-channel function generator would be represented as a single device with sub devices for each output channel. Devices are the standard components to be implemented in device driver packages.

Each (sub) device can have an arbitrary number of timing channels, representing the different configuration parameters of the device. Each channel has a type and associated limits which are defined by the hardware and checked by the software when sequences are created. The supported types are integer and floating point numbers (optionally with a unit), boolean values, strings, and enumerations, or more generally a set of integers. To represent a physical output channel often multiple channels are needed, as every channel only represents a single parameter. E.g. a function generator output could be described by three channels for frequency, amplitude, phase and an additional channel to enable the output.

In contrast to levels two to four the last level is not provided by device drivers, but by the user configuration. Each timing channel can have an associated user channel on this level. The user channel can be used to transform values before they are passed to the timing channel, e.g. apply an output calibration, and to define new default values. Only user channels are accessible and can be programmed in sequence script files.

5.4. Sequence processing

When a sequence script file is passed to the timing system it is first imported as a Python module during which its code in global scope is executed. A single shot is also fully parsed to provide an immediate feedback on sequence errors. Depending on the import mode a single or multiple shots are created. The information is stored in a single `Task` for the full import, with `Shots` for each shot, and passed to the task queue for storage.

When a shot from the task is soon[6] to be executed it is prepared for execution. During preparation the shot information is parsed, i.e. the `main()` function of the script is executed and the event table with all the channel state changes and their timestamps is generated. The changes are grouped by timing controller and transmitted to the respective remote servers for compilation. In the compilation step each timing controller translates the event information into a custom data format needed for execution. The generated data is stored in a cache on the remote server to reduce the amount of data transferred between main server and the remote servers.

The next shot to execute is retrieved from the queue by the task runner during the execution of the previous shot. While the execution of the previous shot is running, `upload` passes the previously compiled data to the timing controllers to prepare the hardware in advance as much as possible. Once the previous shot has finished the output, the internal signal `on_run_about_to_start` is called to allow for final preparations which can only be performed when the hardware is idle. This function also receives the internal

---

[6] the number is configurable, by default the 5 upcoming shots are prepared.

identifier of the upcoming shot, which can be used to send acquired data, even after the output has finished.

Output is started by calling `start_task`. It only starts the output, allowing the timing system to monitor the start up time of the different timing controllers. For each timing controller data returned from this function is passed to the following call of `finish_task`, which only returns once the sequence output on this controller has finished or created an error. When all timing controller are finished `on_run_ended` is called to allow for final clean up before the output of the next shot is started.

Writing occurs once every device has transmitted its acquired data, or when the wait time (by default 15 s) has passed. To keep track of the transmission process devices must register themselves while the sequence is running and unregister once all data is transmitted. The data is then passed to each configured writer for processing, storing the data on disk in HDF5 format by default. Afterwards some internal clean up is performed and the handle to the shot is released. Once all of the shots of a task have finished all handles are released and the garbage collector removes their data. To reduce the memory footprint of tasks with many shots memory expensive data structures, such as the event list and compiled data, are removed from every shot if they are no longer needed.

6. Magnetic field control

Magnetic field control is an essential tool for all quantum gas experiments. All atomic quantum states are affected by the Zeeman effect [131, 132], shifting their energy levels with the absolute magnetic field. Additionally the direction of the magnetic field determines the quantisation axis, changing the interaction between light and atoms.

For experiments with Caesium magnetic field control is an essential tool to create BECs, to adjust the scattering properties throughout the process. Additionally, to realise the planned experiments described in section 2.2.1, the magnetic field stability is an important factor to maintain coherence throughout the experiments.

In the following the requirements for the magnetic field design are sketched and the resulting coil and mount design is explained (section 6.1), and some custom components for the current control are described (section 6.2). To achieve the required magnetic field stability an active stabilisation system has been integrated which is described in section 6.3, followed by a characterisation of the magnetic field stability in section 6.4.

6.1. Magnetic field coils

6.1.1. Design goals

Due to the chamber with twelve viewports and two lattice sets we want to be able to freely rotate the quantisation axis in the horizontal plane to match the different lattice axes and apply gradients if desired. We also might need fields along the vertical direction for levitation. In all cases we need offset fields not only to maintain the quantisation axis, but also to control the scattering length. Because of this we included multiple coil pairs on at least some axes to generate gradients in addition to the offset fields.

The primary application of offset fields is to control the scattering and three-body loss properties using Feshbach resonances. Compared to most other isotopes used in ultracold quantum gas experiments, the required fields are very low due to a very broad resonance at $-11.7\,\mathrm{G}$, influencing the scattering properties up to $\sim 100\,\mathrm{G}$ [36]. As shown in fig. 2.2 a field up to $50\,\mathrm{G}$ is sufficient to tune the scattering length of the absolute ground state from the strongly attractive regime up to approximately $1000\,\mathrm{a}_0$. The negative scattering rate at zero field also forces us to apply an offset field throughout most of the sequence to maintain repulsive interactions. Higher field values are potentially interesting for condensation in other hyperfine states [38], but were not a goal for the design of the coils.

We have chosen two target values for the magnetic field gradients: We want to be able to perform magnetic levitation of ground state atoms in the glass cell, requiring a gradient of $31.5\,\mathrm{G/cm}$ along the vertical axis. To include some tolerance and to allow for over-levitation, e.g. for evaporative cooling [133], we chose $40\,\mathrm{G\,cm^{-1}}$ as the minimal gradient along the vertical axis.

In the horizontal plane the design goal was to achieve a strong enough energy shift ΔE between neighbouring sites in our regular lattices using a magnetic field gradient along the quantisation axis. The shift is given by

$$\begin{aligned} \Delta E &= \mu_B g_F m_F \Delta B && (6.1)\\ &= \mu_B g_F m_F \frac{\partial B}{\partial x} \times d && (6.2)\\ &\approx \frac{\partial B}{\partial x} \times 40.3 \frac{\mathrm{Hz}}{\mathrm{G\,cm^{-1}}}, && (6.3) \end{aligned}$$

with the Bohr magneton μ_B, the hyperfine Landé g-factor g_F ($= 1/4$ for $F = 3$), the hyperfine quantum number $m_F = 3$ in the absolute ground state, and the lattice spacing $d = 767\,\mathrm{nm}/2$. We chose a minimal energy shift $\Delta = h \times 2\,\mathrm{kHz}$ between neighbouring lattice sites in the short space lattice. This allows the shift to be large compared to the relevant energy scale defined by the tunnelling rate J, at reasonably large values of J. A gradient of $\sim 50\,\mathrm{G\,cm^{-1}}$ is necessary to generate this shift for the $F=3$, $m_F=3$ ground state.

6.1.2. Coil design

During the design process we prioritised a single horizontal axis which we intend to use as quantisation axis for the experimental scheme described in section 2.2.1. The Raman coupling introduced in this scheme is not only sensitive to the absolute magnetic field, but also to local deviations of the magnetic field. These generate a locally varying detuning, resulting in inhomogeneous tunnelling rates. The standard configuration to generate homogeneous fields using a coil pair is the Helmholtz configuration. In this configuration the distance of the coils d is equal to the radius of the two coils R, both carrying currents in the same direction. In this configuration the first non-zero, non-uniformity term $\partial^2 B/\partial x^2$ vanishes. Similarly an ideal configuration exists at $d = \sqrt{3}R$, if the current in one coil is inverted to generate a gradient field. Here again the first non-zero non-uniformity term $\partial^3 B/\partial x^3$ vanishes.

In practice, it is rarely possible to fulfil either condition perfectly, as external constraints prevent the placement of the coils at the optimal location. Also most coil pairs are intended to generate offsets or gradients depending on the application, requiring a trade-off between the two ideal configurations. Additionally, our glass cell geometry using a twelve sided body is not compatible with coils in the Helmholtz condition. If the Helmholtz condition is fulfilled the coils intersect the horizontal plane at ±60° from the coil axis, blocking access to the viewports. For the final design a mixture of round and rectangular coils was used. Some coils were made rectangular to place them closer together, reducing the required absolute field. Table 6.1 lists the target coil parameters chosen for the final design.

The coil assembly is rotated by 30° around the vertical axis with respect to the standard coordinate system. This allows us to align the axes of the horizontal coils with the axes

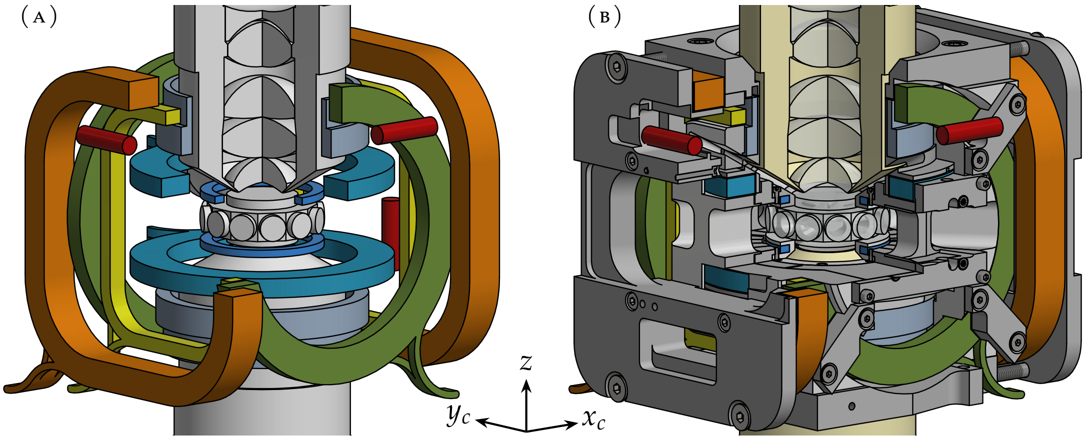

FIGURE 6.1.: Rendering of the glass cell with objectives and magnetic field coils, with and without the coil mount assembly. The coil pairs are coloured per axis for better orientation: Along x_c the large (orange) and small (yellow) coil, the y_c coil (green) and along z the primary coil (turquise), the gradient coil (light blue), as well as the coils for RF generation (dark blue). The red cylinders show the magnetic field sensors, measuring along the cylinder axis.

for a lattice configuration. The coordinate system of the coils is rotated by $-60°$ around the vertical axis. Correspondingly both systems share the z axis, but the coil x and y axes are along the $-y'$ and x' axes of the primary coordinate system. In the following the axes of the coil coordinate system will be labelled x_c, y_c, and z, whenever explicit references to the axes are made.

Current and powers

A fundamental design decision when designing the coils for experimental fields is the available maximal current. We chose an upper current limit of 200 A, which is also within the limits of single power field-effect transistors (FETs) and diodes in SOT-227 package[1] which we use to control the current flow through the coils (see section 6.2.2), allowing us to build simple control circuits without parallelised components.

---

[1] Also called SOT-227-4 or ISOTOP.

Coil	Shape	Dimensions	Distance	Windings
Large X	rectangular	115 mm × 115 mm	110	4×4
Small X	rectangular	73 mm × 93 mm	90	2×2
Y	round	97 mm ID	87	2×5
Z	round	81 mm ID	30.8	2×5
Z gradient	round	77 mm ID	60.4	$6 + 3$

TABLE 6.1.: Physical parameters of the coils around the experiment chamber.

Choosing a high current limit, even though we only need relatively low magnetic fields, allows us to design compact experimental coils with only few windings. This results in a very low inductance L, which scales quadratically with the number of windings, allowing for fast switching of the currents in the coils.

The disadvantage of coils with few turns is the high current required to create the experiment fields, resulting in a high power dissipation. For a given current I the magnetic field B of a coil with n windings is proportional to $I \times n$. The power P dissipated by the coil due to resistive losses is given by

$$\begin{aligned} P &= UI \\ &= RI^2 \\ &= nl_0 A \rho I^2 \end{aligned}$$

with the length of a single coil loop l_0, the wire cross section area A, and the electric resistivity ρ. Assuming a fixed total cross-section of the coil (due to geometric constraints when designing the coil) $nA = \text{const}$, hence $P \propto I^2$.

To minimise the dissipated power it is thus beneficial to use coils with many turns and low currents. Nevertheless we decided to optimise the coils primarily for a low inductance, favouring fast switching speeds. We are using the same large cross section hollow core copper wire already used for the MOT coils[2]. The hollow core allows to cool the coils from inside using water cooling, removing the need for a cooling solution in the external mount. The main disadvantage of these wires is the low flow rate due to the small tube diameter, requiring a high pressure for effective cooling. As we are using these wires already for the MOT coils we already have a cooling water circuit with 10 bar pressure installed with enough capacity to also cool the experiment coils.

Horizontal coils

The main application of the coils in the horizontal plane is the generation of an offset field, controlling the scattering properties using the Feshbach resonance. We installed coil pairs in both horizontal axis to allow for arbitrary rotation of the quantisation axis. We planned for two different square lattice configurations, rotated by 30° relative to each other. Having the option of arbitrary rotations allows to align the quantisation axis with either lattice configuration as needed.

We chose the y_c axis along which we intend to build the beam path for the anti-magic lattice required for the generation of topological systems (section 2.2.1). It not only requires very stable magnetic fields, but to create the required σ^- polarisation of the lattice light the quantisation must be parallel to the laser beam. During the coil design we prioritised this axis, optimising the parameters of this coil pair, then designing the other coils around it. The resulting coil pair is a round coil which is placed as close to the Helmholtz condition as possible and designed to be as small as possible, maintaining sufficient space for the objectives and vertical coils. The estimated deviations for a

---

[2] OF-OK oxygen free copper, Hollow Conductor OD 3 mm × 4 mm/ID Ø2 mm, Luvata tool #8329, Luvata

coil	offset [G/A]	gradient [G/cm/A]
Large X	0.907	0.264
Small X	0.402	0.127
Y	0.990	0.264
Z	2.087	0.443
Z gradient	0.956	0.344

TABLE 6.2.: Computed field parameters for the coils around the experiment chamber.

uniform field are maximally 5×10^{-9} between neighbouring sites, corresponding to an negligible absolute offset of 0.2 μG at 50 G applied field. If the cloud is displaced from the coil centre an additional gradient is added, but the magnitude of the non-uniformities is unchanged.

Around the initially designed coil pair along the y_c axis we designed the coils for the x_c axis. On this axis we added two coil pairs, allowing for the combined generation of offset and gradient fields along this axis. Both coil pairs are rectangular with rounded corners. The straight sections of the coils allowed us to place them closer around the existing coils, reducing the distance from the atoms, hence reducing the number of windings needed per coil. The coil pairs are placed in between and around the coil pair on the y_c axis. The primary coil pair is called *large X* and is mounted around the coil on the y_c axis. Another secondary coil pair intended for gradient generation called *small X* is placed closer to the atoms.

Vertical coils

In vertical direction two coil pairs were installed for independent control of offset and gradient fields. As described in the chapter introduction their main application is magnetic levitation, i.e. the generation of a gradient counteracting the gravitational force. Magnetic levitation can also be used for fast and efficient evaporation to BEC [133], which was used in our experimental sequence before the rebuild of the experiment section to its final configuration.

To allow for magnetic levitation while controlling the scattering properties using an offset field we installed two independent pairs of coils along the vertical direction. The primary coil pair with 10 windings per coil is placed as close as possible around the glass cell. It is designed to clear the small opening angle next to the objectives providing further optical access to the atoms. A secondary, elongated coil pair with 9 windings per coil is placed further away from the atoms, covering the space around the objectives. It has a stepped profile with an inner layer of 6 windings and an outer layer with 3 windings, optimised to the available space defined by the coils on the y_c axis. It is intended to be used for gradient field generation while the primary coils generate an offset field. They are usually referred to as the Z *gradient coils*, while the primary coils are simply called Z *coils*.

The vertical coils are also intended to generate gradients when the quantisation axis is aligned in the horizontal plane. The energy shift of the atoms experienced due to the Zeeman effect is only dependent on the absolute magnetic field $B = \sqrt{B_x^2 + B_y^2 + B_z^2}$. When a dominant offset field is applied in one axis, field deviations and gradients in orthogonal axes are suppressed quadratically, allowing the generation of magnetic gradients using coil pairs orthogonal to the quantisation axis.

The magnetic field gradient generated by a coil pair in the Anti-Helmholtz condition along the orthogonal axes is only half compared to the gradient along the coil axis to fulfil Maxwell's second equation $\nabla \boldsymbol{B} = 0$. Due to the close proximity of the coils gradients up to $100\,\mathrm{G\,cm^{-1}}$ can be generated. This gradient can be applied for any rotation of the quantisation axis in the horizontal plane.

Coil manufacturing

All final coil pairs, except the Z gradient coils, were manufactured externally[3]. We decided to do this as the external manufacturer has the capabilities to change the winding direction during manufacturing, allowing to create coils with two layers such that both connection wires leave the coil on the outside. Usually when the coil is wound in only a single direction the beginning of the wire is on the inside of the finished coil and has to be directed to the outside next to the finished coil. Having both wires end on the outside of the coils simplifies the design of the mount and reduces the distance where the wires can not be run in parallel, reducing the loop area which environmental noise can couple into.

The externally manufactured coils were designed such that this method of production was possible, requiring each coil to have two layers. The large coils on the x_c axis are produced by combining to coils with two layers each. Aiming to increase the symmetry of the system and reduce the possible deviations from the computed, ideal design, the coils for each coil pair were wound in mirrored configuration. For mechanical stability the finished coils were wrapped with $150\,\mu\mathrm{m}$ fibre glass tape and then potted using vacuum casting with a two component epoxy[4].

The Z gradient coils could not be designed to be manufacturable using the same method and were therefore produced in house. They were wound directly onto the final mount and potted using a two component, room temperature curing epoxy[5].

Due to a miscommunication when considering the extra space needed for the fibre glass wrapping and potting with the manufacturer the inner dimensions of the produced coils are 0.5 mm larger on either side. The dimensions given in table 6.1 already include this correction. It leads only to minor deviations of the field from the expected value and is not critical. When mounting the coils onto the mount we filled the gap on three or four points per coil with layers of Kapton tape to centre them correctly on the mount.

At some positions we removed parts of the potting material using an axial grinder to make space for wires from other coils. The planned space for connection wires in

---

[3] by Krämer Energietechnik GmbH

[4] Huntsman Araldite CW 229-3 & Aradur HW 229-1

[5] Delo Duopox AD840

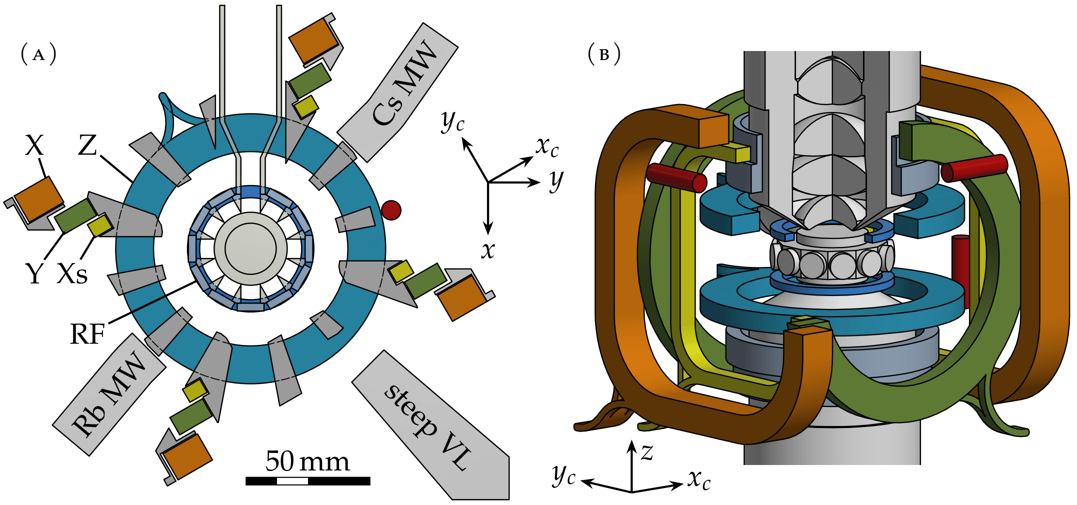

Figure 6.2.: Section of the coil mount assembly along the $-z_c$ axes (a; view from top) and the overview of the coils in the same colours (b). The mount is drawn in grey, as are microwave horn and the steep vertical lattice intersecting the horizontal plane. The red circle is the magnetic field sensor for z at the computed location.

the original design included only minimal tolerances which turned out to be not large enough.

6.1.3. Mount

To place the various coils rigidly and in the correct orientation around the science chamber we designed a mount, which we assembled around the previously installed glass cell. Figure 6.1 shows the coils inside the full mount with all parts.

We designed the mount to fill as much of the available space as possible to maximise mechanical stability, while not blocking any optical access to the science chamber. The final design has openings aligned with the horizontal viewports as shown in fig. 6.2. The size of the openings was chosen to not limit optical access to the atoms beyond the limitations set by the glass cell geometry. The size of the complete coil assembly including mounts is 165 mm × 135 mm × 150 mm.

In addition to the horizontal openings we included curved slots for optical access under ±60° from the vertical axis, passing through the large viewports directly next to the objectives. These are visible in fig. 6.3, which shows sections through the vertical planes. It also shows how the objectives are located relative to, but not in contact with, the coil assembly.

The coil assembly is mounted to a 20 mm thick plate, which is resting on three 1.5″ posts made from PEEK. These are filled with a mixture of lead balls and fine sand to dampen vibrations. To adjust the angle and height of the assembly the plate is resting

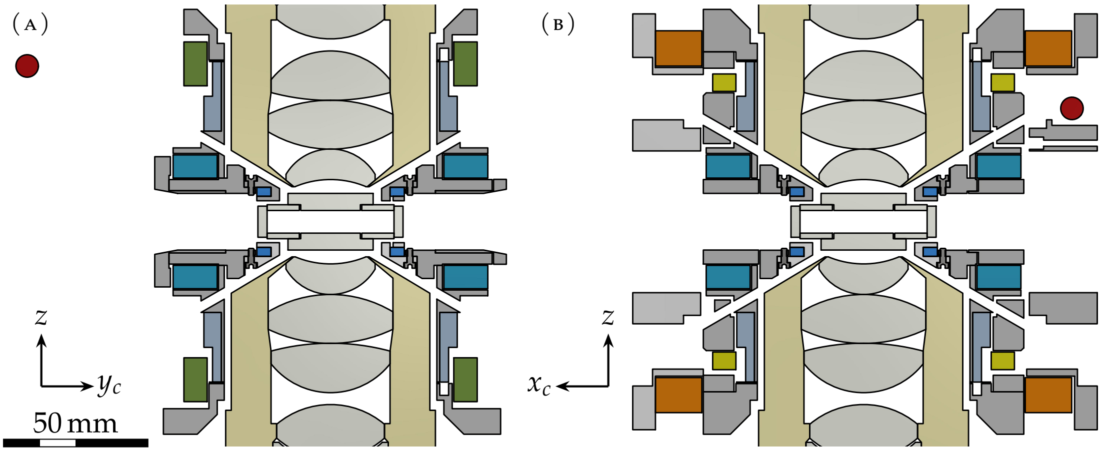

FIGURE 6.3.: Sections of the coil mount assembly through the $y_c z$ (a) and $x_c z$ (b) planes. The colours of the coils match those in fig. 6.2. The red circles are the magnetic field sensors for x_c and y_c at their computed locations. The openings next to the objectives allow optical access under 60° from the vertical axis and are used for the molasses beams, the z dipole trap (section 7.2.1), and the steep vertical lattice (section 4.5).

on a single screw per post. Two other screws next to it can be tightened to clamp the plate rigidly in place. We later replaced the screws with larger spacers, machined to the measured height, to increase the contact area and avoid slow material deformations.

The initial alignment of the coil assembly was done by eye during the assembly process. Halfway through the assembly — once the lower coils were in place, but before the upper coils were placed — we aligned the assembly relative to the glass cell. To help in this process 24 lines pointing outwards from the centre were machined into the surface of the innermost plate, which was the top of the assembly at the time. We used these for horizontal alignment, orienting the assembly to have them point on the corners of the glass cell and the centres of the viewports. To adjust the angle we attempted to achieve a constant distance between the glass cell body and the mount of 1 mm, which we knew would centre the assembly vertically. Using this approach we achieved a reasonable alignment relative to the glass cell. We later found out that the coil assembly is aligned a bit lower than it should be, but did not change it, fearing damaging the glass cell in the process.

In hindsight it would have been a good idea to include temporary mounting fixtures for proper measurement equipment. Using a dial indicator in a rotating mount aligned with the centre axis we could have probed the vertical distance and angle against the body of the glass cell and the horizontal alignment against the edge of the glass cell window. Tiny scratches from a ball on the tip would not have been a problem as both surfaces are not used for any beams.

A challenge in the design was the requirement of assembly around the previously installed and baked glass-cell. We decided to use a design which fully encloses the glass cell, rather than being pushed over it. This maximises mechanical stability and allows to place coils very close to the chamber, but makes the assembly process more complicated and risky, as the various parts need to be placed around the fragile glass cell. With the experience from this assembly I would tend to not follow this approach again. It not only increases complexity of the coil mount, but also makes pre-alignment of the objective almost impossible as it must be installed independently once the coils are installed.

6.2. Current control

6.2.1. Ultra-stable current sources

To achieve the lowest possible magnetic field noise it is essential to use stable and low-noise current sources to supply the experiment coils. As described in section 6.1 we chose a maximal field of 50 G for Feshbach fields. At the same time we want to reduce the magnetic field noise as much as possible to values below 100 µG, requiring relative stabilities on the order of 10^{-6}. This level can only be achieved by specialised low-noise current sources.

Most current sources used today are using switching-mode topologies, due to their high efficiency and compact size. Opposed to a classical, linear power supply they do not control the current flow continuously using an analogue feedback. Instead they switch the input on and off in rapid succession (hence their name), using inductors to provide a smooth output voltage or current. This allows them to reach high efficiencies by control of the duty cycle, but the fast transients from the switching operation are difficult to filter, causing an increased output noise. The quality of switching-mode supplies has increased significantly over the years, replacing linear supplies in most applications. Their noise level however is still higher than comparable linear supplies, achieving relative current noise figures around 10^{-4}.

The continuous architecture of linear power supplies allows them to have generally lower noise figures. Even for those the required noise figures are difficult to achieve as one approaches the limits of the electronic components. An example for this is the required feedback sensor used to maintain a constant output current. State-of-the-art current sensors, such as the ones used in the installed power supplies[6], have a root-mean-square (rms) noise figure around 10^{-6}. This measurement noise fundamentally can not be corrected by the feedback loop, setting a lower boundary on the achievable noise floor. In addition drifts with temperature are on a similar or even higher level, e.g. the current sensors having a temperature coefficient of $2.5 \times 10^{-6}\,\mathrm{K}^{-1}$. Similarly only high-end voltage-references have temperature coefficients and long term stabilities on a comparable level.

One option to create current sources with sufficiently low noise figure by using a good power supply, followed by a custom feedback loop [134, 135]. We decided not to pursue this approach, instead we purchased commercially available supplies with suitable noise figures, specially designed as current sources for magnetic field coils[7].

---

[6] IT 60-S ULTRASTAB and IT 200-S ULTRASTAB, LEM

[7] HighFinesse UCS50/10 (50 A/10 V) and UCS200/20 (200 A/20 V), both with the ultra-high stability and

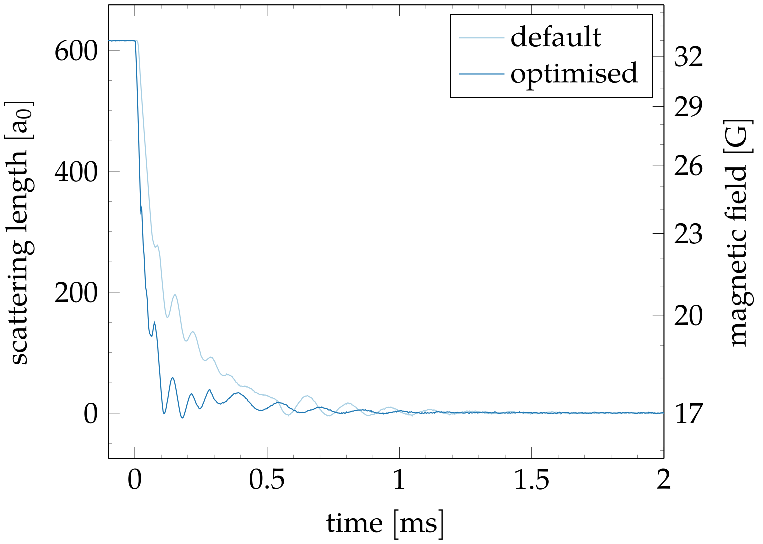

Figure 6.4.: Magnetic field jump from a typical field for an Mott-insulator (MI) at 33 G to zero interactions at 17 G, using the ultra-stable 50 A power supply, switching to the new value at $t = 0$ (default), or using an optimised drive sequence. Data was measured using a current clamp on the coil wires, the calibrated magnetic field values were used to calculate the corresponding scattering length in the used $F = 3$, $m_F = 3$ state. With the default (optimised) sequence the fluctuations are within $\pm 10\,a_0$ after 1 ms (0.7 ms) and within $\pm 2\,a_0$ after 1.6 ms (1 ms).

They are linear power supplies, specified with a current noise below 10^{-5} and a stability better than $5 \times 10^{-6}\,\mathrm{K}^{-1}$, while also having a fast response time on the order of 100 µs (depending on coil inductance). The digital control option integrates an 18 bit digital-to-analog converter (DAC) which can be programmed either via USB or Serial Peripheral Interface (SPI). We chose this option to remove the analogue control generator and lines as potential noise source.

The current limit of the two power supplies is optimised to the tasks we plan to use the supplies for. The higher current supply matches the current limit we chose when designing the coils, allowing us to reach maximal field strength, mostly required for strong magnetic gradients. The lower current supply matches the current range needed for scattering length control using the Feshbach resonance. Using a supply with lower maximal current not only reduces the power loss, but also results in a higher absolute current stability. The relative accuracy of the supplies is the same, correspondingly the absolute accuracy is improved by a factor of four and according to the manufacturer the noise performance of the supplies is best in a medium to high load regime, with an ideal load of 80 %. The combination of the two power supplies allows us to cover a broad current range with low noise by using either of the two supplies. To be able to use them flexible on any coil we included extensive current switching capabilities described in the following section.

As a speed reference measurement for a typical usage of the power supply we tested

---

digital control options

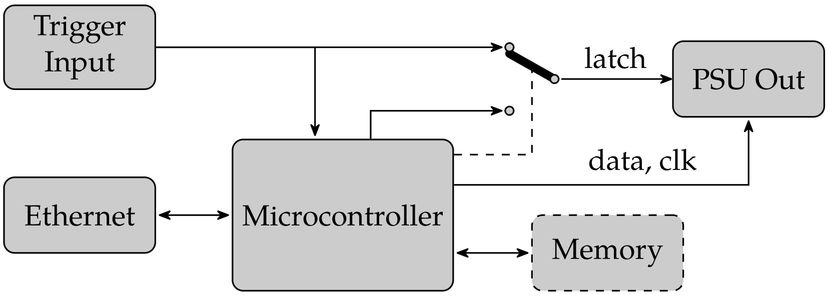

Figure 6.5.: Block diagram of the digital control board for the HighFinesse power supplies. The microcontroller is responsible for programming the power supply using the 'data' and 'clk' lines. The latch channel is used to activate the updates and can be controlled from the microcontroller and through an external trigger input. Communication with the microcontroller to program sequences is performed using ethernet.

jumps from a standard field value in the MI state of 33 G to the zero-crossing of the scattering resonance at ∼17 G. Figure 6.4 shows the measured current in the coil, rescaled to the generated magnetic field and the corresponding scattering length of $F=3$, $m_F=3$. The default ramp is generated by reprogramming the supply with the new value at $t = 0\,\mathrm{s}$. The optimised sequence disables the supply and shorts the coils using metal–oxide–semiconductor field-effect transistors (MOSFETs) for 20 µs to accelerate the decay of the field, afterwards the current is ramped to the final value within 260 µs, updating the target value every 20 µs to reduce overshoots and stabilise the current faster.

Digital control board

The digital control option of the power supplies can either be programmed through USB or through SPI. The USB interface is easy to use using Standard Commands for Programmable Instruments (SCPI) commands, but it does not allow for precise timing and limits the update rate to ∼1 kHz. Instead we use the SPI interface which allows to directly program the DAC. This allows for an update rate up to ∼100 kHz and by controlling the latch line it is possible to trigger each update precisely[8].

To use the interface we built a custom board using a 32 bit ARM microprocessor[9] with ethernet connectivity and some level shifting components to communicate with the power supply. It is based on a design for a DDS board developed during my master's thesis [136, 137] for the usage in ultracold quantum gas experiments. The design relies on the presence of a real-time timing component to provide precise trigger pulses, only focussing on fast programming of the DDS chip between updates.

The derived board uses the same approach, updating the DAC inside the power supply instead of a DDS. In both cases communication between the microcontroller and the integrated circuit (IC) is done using SPI. Similar to the DDS chip the used DAC[10] allows

---

[8] Limited by the timing resolution of our digital output channels of 100 ns.

[9] STM32F407VG, STMicroelectronics

[10] AD5781, Analog Devices

the asynchronous programming of register changes, independent fro their activation. A digital input to the control board allows direct control over the trigger line. When the control board is in sequence execution, playing the programmed sequence, the microcontroller only monitors this line for changes, programming the next set of changes as soon as the current have been activated externally. The microcontroller is programmed via ethernet and set to execution mode at a time earlier in the sequence where timing is not critical.

Using a SPI clock speed of 5.25 MHz the design has been tested successfully with only 15 µs delay between trigger pulses, of which 4.6 µs are required for the data transfer. The resulting update rate is fast enough to program current ramps where the discretisation of the updates is no longer observable. Updates can be sent as needed, also in irregular intervals, as long as the minimal update delay of 15 µs is met.

The currently used board is a repurposed prototyping board of the mentioned DDS board design with modified firmware [138]. It can only use parts of the 192 KiB internal memory of the microcontroller to store data for sequences[11]. With the current implementation only $2^{14} = 16\,384$ samples can be stored at a time. A derived board using a more capable variant of the used microcontroller with integrated memory controller[12] and 8 MiB external memory was designed [139], but not built as sourcing the microcontroller was not possible due to the global chip shortage[13]. If needed it would be possible to increase the buffer size by $\sim$78 % by packing the sequence values, instead of storing them individually in 32 bits each.

6.2.2. Current switching

To allow for flexible magnetic field control along all axis we installed two switch boxes between the experiment coils and the power supplies. One containing a matrix arrangement of switches to connect any of the four connected power supplies to any experiment coil. Another to control the direction of current through the coil pairs, allowing to switch between offset and gradient configurations.

Both boxes use MOSFETs to switch the current, combined with diodes where needed to ensure correct current flow. They are only intended to switch the current, always operating the FETs in the saturation regime. They include digital drive electronics, using integrated direct-current to direct-current (DC-to-DC) converters to enable high-side switching of the used N-Channel MOSFETs.

Current direction control

To maximise flexibility when using the magnetic field coils for experiments we connected all coil pairs on the experiment side to H-bridges to control the current flow per coil. Depending on the coil we installed only a single H-bridge to switch the second coil, or

---

[11] Using the SRAM. The board also has 1 MiB of flash memory. It is not used as it has a finite number of write cycles (10^4) and would be worn down quickly.

[12] STM32F427ZG, STMicroelectronics

[13] If a supplier would provide a delivery date at all it was in the excess of one year.

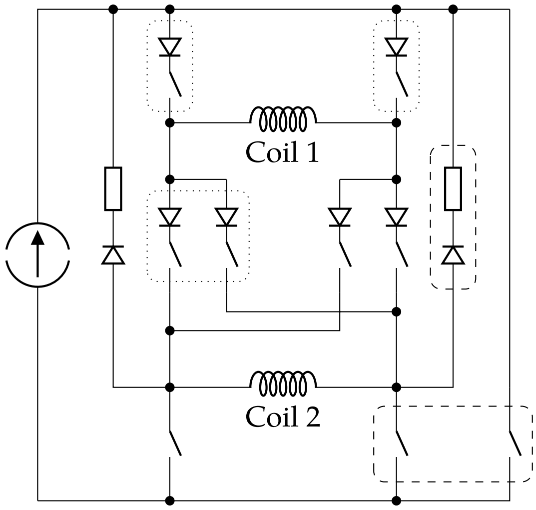

Figure 6.6.: Schematic drawing of the coil direction switching of a single coil. Each switch symbolises a MOSFET with drive electronics. The components within the dashed boxes build the short and are present in all configurations. The components in dotted boxes are only present in the dual H-bridge configuration, but not in the single H-bridge where only the current direction in coil 2 can be switched. In this case the current source is directly connected to the left side of coil 1.

a dual H-bridge switching both coils. The former configuration only allows to switch between offset and gradient fields, while the latter additionally allows to switch the field direction at the cost of more components and higher losses.

We designed a compact 19″ rack mounted box, 3U in height. It can control up to 8 coils, up to 4 of those in dual H-bridge configuration. The power electronics in SOT-227 packages are installed on two water cooled plates, the necessary driver electronics are mounted on the front panels. In the back connecting plates to the coils and current sources are installed as strain relief and to simplify connections.

Figure 6.6 shows a simplified schematic drawing of a single coil stage within the box. Depending on the needs one of three different configurations can be installed. The simplest configuration used for the MOT coils only includes two switches to direct the current through the coils or through a bypass for rapid switch-off. A free-wheeling diode with resistor provides a return path for the current stored in the coils. The next larger configurations include more switches for H-bridges, either only around the second coil, or around both coils. The H-bridge configuration allows to control the current direction through a coil by pairwise enabling a switch on the input and output side. Several diodes are used to block reverse current flow through the MOSFET body diode when switching them under load.

All switches are implemented using N-channel MOSFETs[14], which allow for fast current switching and have a very low resistance ($<2.5\,\mathrm{m}\Omega$) when active. To generate the necessary floating gate voltage between gate and source of each MOSFET we are

---

[14] IXFN360N10T, IXYS

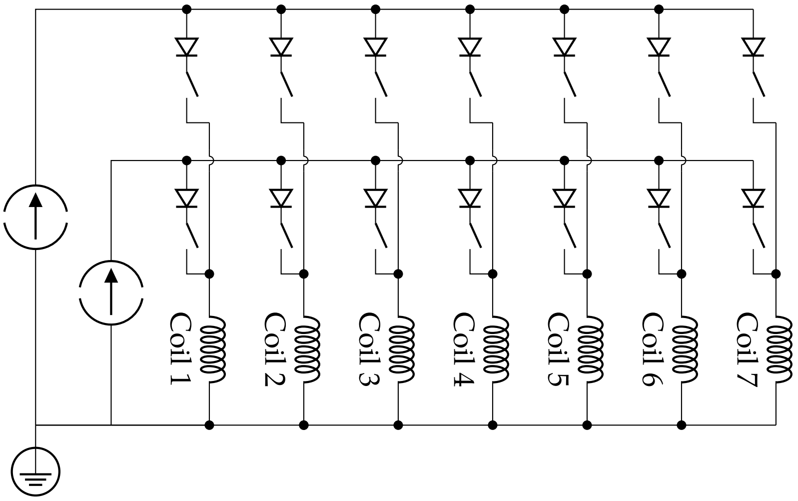

Figure 6.7.: Simplified schematic drawing of the high current connections in the power supply selection box. Each switch represents a N-channel power MOSFET with isolated driver IC. For simplicity only two of the four power possible power supplies are shown. When combined with the coil switch box it replaces the individual coils drawn here. The common ground defined by this box is used as central grounding point for the coil system.

using isolating DC-to-DC converter, followed by a low pass and linear post regulation stage to reduce ripple and noise. The MOSFETs are driven with a high gate voltage deep into saturation mode, reducing noise coupling into the primary circuit as the resistance changes with varying voltage are minimal. MOSFET driver ICs are used, which have high current drive capabilities and can be controlled using signals on a voltage level isolated from the drive voltage. Each control board has up to three opto-coupled inputs to enable the coil, switch between offset and gradient, and flip the overall current direction, which are translated into the corresponding MOSFET states using standard logic ICs.

All design files for the mechanical and electrical components of the box, including assembly instructions, are publicly available [140].

Power supply selection

In addition to the current control through the coils, we installed another box to control the current flow from the power supplies to the coils. It connects up to four power supplies and seven coils, allowing any supply to be connected to any coil (fig. 6.7). This was primarily done to avoid connecting the ultra-stable current sources to a single coil pair and increase the flexibility in their usage.

The power supply selection box also acts as central grounding point for all experiment coils and power supplies. Only the positive sides are controlled by MOSFETs. The negative side of all power supplies and coils are connected together on a central copper bar. This bar is also connected to protective earth, providing a common grounding point. All power supplies connected to this system must have galvanically isolated outputs to prevent the creation of ground loops.

This control box uses a similar construction and design, mounting all components within a 3U 19″ rack box. It has four inputs for power supplies which be switched to any one of seven outputs, connected to the current inputs of the other box. Internally the same MOSFETs and diodes are used to control the current flow and adapted driver boards are included. As these switches are less time sensitive a single board computer is included, allowing to control the switches through ethernet. The I²C GPIO expanders used on the custom driver boards [141] are directly supported by the Linux kernel, and can be controlled remotely via HTTP through a lightweight Python daemon [142].

To monitor the current flow through all coils, 16 cost effective current clamps[15] are installed in both boxes, monitoring the current flow from each power supply and of the individual coils. A central drive board in the power supply selection box provides a common voltage reference and output buffers to our analogue monitoring inputs for permanent monitoring and front connectors for temporary monitoring next to the devices. Their primary goal is the detection of the overall system status and mistakes in the sequence programming. To measure the high accuracy of the ultra-stable current sources larger and significantly more expensive current clamps with better support electronics are required.

6.3. Active field stabilisation

For our experimental setup magnetic field stability is especially important for the planned method to create topological systems described in section 2.2.1. It requires the usage of different spin states, necessarily creating a system which is linearly dependent on magnetic field fluctuations. To create systems with long coherence times the magnetic field noise should be as low as possible. The stability of the power supplies is expected to be around 100 µG as described above. The stability of the environment field should be on the same order or better to not limit the performance of the experiment.

At the beginning of construction we measured the background field in our lab (fig. 6.8). The measurements showed frequent jumps in the background field up to 20 mG, especially along the vertical direction. Fluctuations along the horizontal axes are only 2 mG to 5 mG, still more than an order of magnitude too large. We attribute these mainly to the nearby underground lines, as we could see correlation between the field jumps and the underground schedule. We also observe a more stable field environment during the night.

In addition to the drifts of the background field the magnetic field varies by up to 4 mG synchronous with the mains frequency of 50 Hz (fig. 6.8b). These fluctuations are very consistent between periods and over a long period of time, allowing for feed-forward compensation [134, 143].

Feed-forward compensation would allow for the suppression of the periodic components, but not of the fluctuations of the quasi-random field jumps. Common approaches to suppress field fluctuations are shielding the system from the environmental field,

---

[15] HO 50-S-0100, LEM

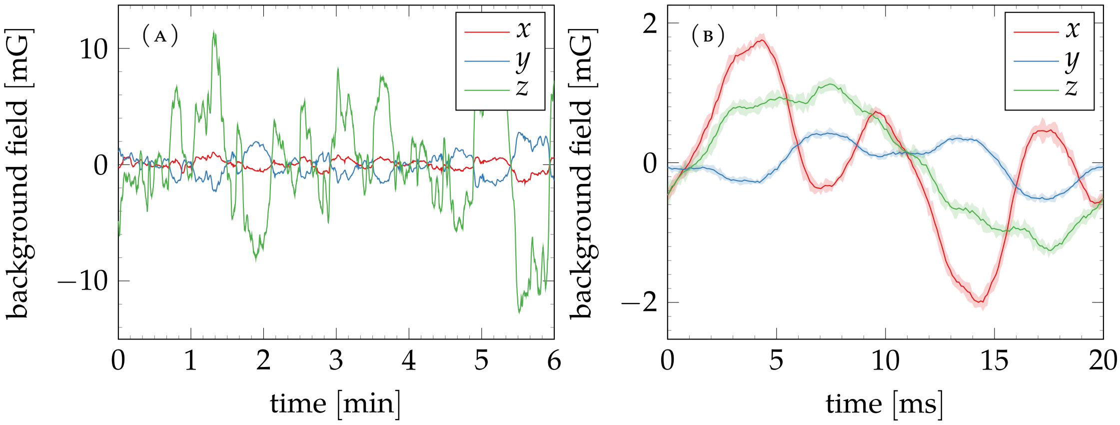

Figure 6.8.: (a) Magnetic field fluctuations of the unstabilised background field along the different axes. Data was averaged over multiple mains cycles to exclude mains synchronous fluctuations. (b) Mains synchronous magnetic field fluctuations over a full period. The shaded areas denote the 3σ interval after subtracting the mean value to exclude shot to shot drifts. Data was averaged over 140 periods.

usually using mu-metal [144–146], or measurement and active cancellation of field fluctuations [134, 143, 146, 147].

We decided against magnetic shielding, as integration of the required metal enclosure, while maintaining the required optical access and flexibility for later extensions is very difficult. Finite element simulation can be used to optimise the shielding design, taking the optical access holes into account, but requires a precise knowledge of the required access in advance. Additional problem can arise due to magnetisation of the shielding when strong magnetic field gradients are created from the coils placed inside the cage.

Instead we decided to include a commercially available active magnetic field stabilisation system[16], designed to stabilise the magnetic field from DC up to 3 kHz. Active stabilisation does not limit the optical access to the chamber, and can not directly be magnetised, but optimal placement of the required magnetic field sensors is challenging. For these applications fluxgate sensors are suitable, as they combine bandwidths of a few kilohertz with relative accuracies below 10^{-5}. Sufficiently low absolute accuracies can be reached when using sensors with low measurement range. We use a sensor with ±700 mG measurement range and a bandwidth of 3 kHz[17]. To avoid saturation of the sensor it must either be placed far from the experimental coils, or very precisely to be not affected by the experiment fields.

A prior attempt of our group in a laboratory in the same building achieved a relative stability of 3×10^{-6} using multiple sensors placed around the experiment table [148], but failed on integration in the experiment as the stability could not be replicated at the atoms. Similar difficulties have been reported by ion experiments [149]. To maximise the correlation between the magnetic field measured by the sensors and the relevant

---

[16] MK5, Integrated Dynamics Engineering GmbH

[17] MAG-03IE70, Bartington Instruments

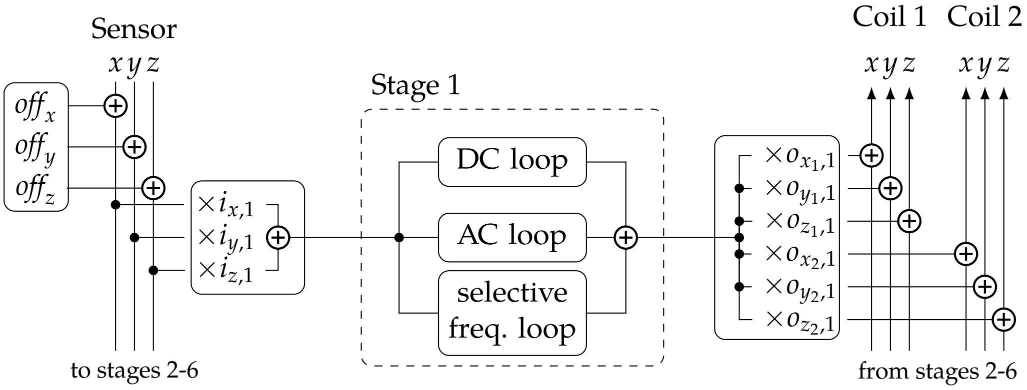

FIGURE 6.9.: Block diagram of the signal processing performed inside the magnetic field stabilisation box. Only one of the six parallel stages is shown. Each has its own input and output matrices, which can be used to control their influence on the different axes. The global offset is applied to the sensor input to set the stabilisation value per axis. A second sensor input is supported, but not used.

magnetic field at the atoms we included the sensors from the beginning into the coil design, locating them as close as possible to the atoms, while being unaffected from the experiment fields as described in the following.

6.3.1. Stabilisation system

The installed commercial stabilisation system integrates different analogue and digital feedback loops, which can be controlled digitally using a provided software. An external fluxgate sensor is used to measure the environment field. Feedback is applied by driving large compensation coils enclosing the experiment and the sensors.

Figure 6.9 shows a reduced block diagram of the stages of the feedback path. At startup or on user request the feedback loop is disabled and the measured field is measured for several seconds to detect the mean background field per axis and derive the set points of the feedback, called offset values in this system. These values are subtracted from the sensor reading, which is then stabilised to zero from the feedback stages. Up to two different three axis sensor inputs and up to two output stages can be used. We use only a single three axes sensor, but both output stages by driving the coils in each coil pair individually.

The system has six identical feedback stages, intended for up to two sensors inputs with three axes each. An input and output matrix allows to assign arbitrary weights for each sensor axis and each coil output for every stage. We use this feature to compensate for a rotation between the axes of the sensors and compensation coils, and the offset of the sensors from the centre of the compensation cage.

Three different feedback loops can be used in parallel. An analogue AC loop and a digital DC loop share the main stabilisation work. Both can be configured by only a single gain parameter. An additional digital selective frequency loop operates on up to

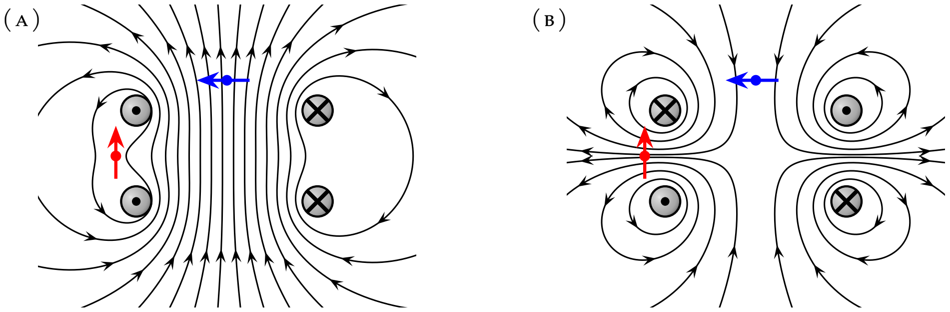

Figure 6.10.: Magnetic field lines generated from a coil pair in Helmholtz configuration (a) and the same coil pair in Anti-Helmholtz configuration (b). The position and orientation of the sensor aligned with the coil axis is shown in red. In blue a sensor for an orthogonal axis unaffected by the field of this coil pair is shown.

seven configurable, very narrow frequency bands and can be used to suppress selective frequencies, including DC.

In the following our integration of the system in the experiment is described. Further information about the integration and the stabilisation system in general can be found in appendix B.

Placement of sensors

To measure the environmental field independent from the generated fields we use independent single-axis vector magnetometers per axis. They can be positioned such that in theory the vectorial field at the sensor is independent from the experiment magnetic fields of all five coil pairs. In practice imperfections in the manufacturing process of the coils create minor deviations from this ideal situation, requiring a compromise on the coupling strength of the different axes.

To place the sensors in the presence of coil along all axes only a few spots close to the atoms can be used. Gradient fields created by coils in the Anti-Helmholtz configuration along the measurement axis can be suppressed by placing the sensor in the plane between both coils (fig. 6.10b). Offset fields generated by the same coil pair in Helmholtz configuration can be suppressed by placing the sensor in the correct distance from the centre where the direction of field changes (fig. 6.10a). The former constrains the sensor location to the plane between the coils, the latter additionally to a single loop around the atoms.

When multiple coil pairs are installed on the same axis it is in general not possible to suppress the field from both coils when they are used in Helmholtz configuration, as the field direction change will occur at different distances from the centre. For our design we chose a primary coil pair per axis, which we exclusively use to generate offset fields along this axis.

Fields generated by coils on axes orthogonal to the measurement axis are suppressed when placing the sensor in the symmetry plane of the coil normal to the measurement axis, i.e. in the middle plane of the coil through the measurement axis (fig. 6.10). Assuming

	Offset [mm]			
Axis	x_c	y_c	z	absolute
x_c	0	102.9	55	116.7
y_c	65.3	0	40	76.5
z	55.7	15.3	0	57.8

TABLE 6.3.: Computed magnetic field sensor distances from the centre. The axes of the 30° rotated coordinate system of the magnetic field coils are used. The final positions can not be measured, and are most likely a few mm further away from the atoms.

coils in an orthogonal configuration this constraint is already fulfilled from the constraint for the gradient fields.

For each axis we chose a convenient position on the loop with minimal influence on other parts of the experiment, not blocking any optical access to the science chamber. The location of the sensors and the distances from the atoms can be found in table 6.3.

As explained it is theoretically possible to place individual sensors such that they can measure only the environmental field close to the atoms. This however requires very precise placement of the sensor on the computed positions as the sensor must be placed at the zero crossing of a steep magnetic field gradient. A displacement of a micrometer towards the atoms already results in a measurement of $\sim$30 ppm of the experiment field by the sensor. Errors of comparable strength are introduced along the other axes when generating gradient fields and from the other coil pairs.

The chosen positions described in the previous paragraphs assume perfect alignment of the measurement axis with the axis of the magnetic field coils. If the sensor is rotated off the ideal axis by an angle α a false reading $B_{\text{err}} = \sin(\alpha) B_\perp$ is introduced, where $B_\perp$ is the magnetic field component on the axis perpendicular to the desired measurement axis. This field can be large, e.g. when an offset field in the orthogonal axis is generated. The relative measurement error is $\sim 10^{-6}\,\mu\text{rad}^{-1}$ with absolute fields comparable to the displacement discussed before, requiring very precise angle alignment on the order of 10 µrad.

In practice alignment is even more complex due to imperfections in the coil manufacturing process. The real coils slightly deviate in the location of the wound wires and their position to the atom centre. Theoretically it would be possible to compensate most of these errors by precise alignment of the coils relative to the atoms. We did not add these alignment options as this would have made the coil mount even more complex. The coil mount described in section 6.1.3 was designed to provide good distance and angle alignment of the coil pairs relative to each other. All remaining imperfections have to be compensated using the sensor alignment.

Sensor mounts To mount the sensors with the necessary degrees of freedom for fine alignment we designed mounts for each of the sensors. Each mount combines linear translation stages in two axis with methods of angle alignment. The mounts for the Y

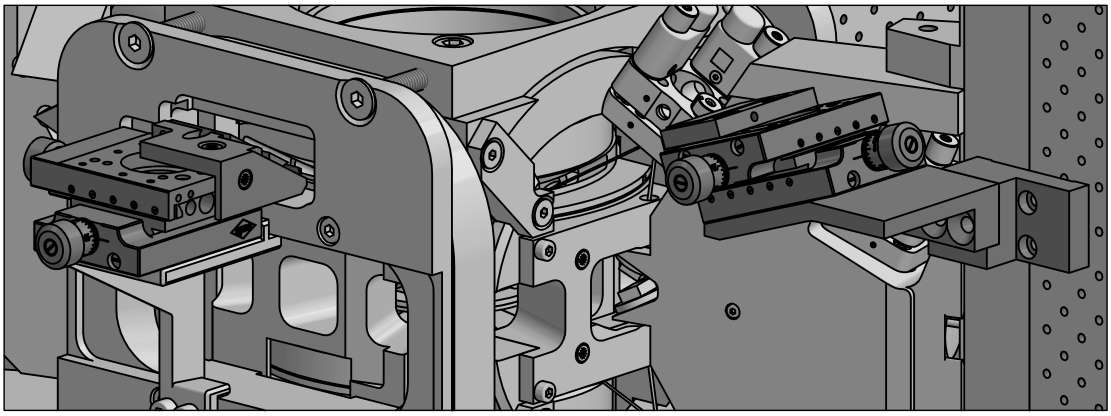

FIGURE 6.11.: Rendering of mounts for the y_c axis (left) and x_c axis (right) magnetic field sensors at their respective positions. The sensors are placed within the mounts on top of the 2d translation stages towards the coils and are not visible in the shown perspective.

and Z axis are part of the coil mount assembly, to reduce the risk of position changes, e.g. due to thermal expansion. The mount for the x sensor had to be mounted onto one of the vertical pillars supporting the upper breadboards and objectives. Figure 6.11 shows the mounts for the x_c and y_c axis sensors.

For position adjustments the mounts use non-magnetic translation stages[18] to avoid influences on the measurement. Every mount has two linear translation stages, one of these is aligned with the sensor axis. The sensor mounts were all manufactured from engineering plastics[19] and include flexure bearings for angle adjustment. Flexure bearings are very simple and compact, at the cost of narrow adjustment range and the risk of material fatigue. These are not a problem in our application as they are only needed to compensate for manufacturing imperfections in the range of ±5°[20] on installation of the sensor. To further simplify the design we use a simple screw without counteracting component to deflect the sensor and adjust the angle. It is still possible to compensate for any manufacturing errors by rotating the sensor in the mount prior to alignment to ensure the deviation is aligned in the correct direction.

At the time of writing of this thesis only two out of the three sensors are installed. The sensor for the Z axis could not be placed at its intended position due to a collision with the bottom primary Z coil. To reach the intended position where the offset field cancels it would be necessary to place the sensor closer towards the atoms. Replacement of the single sensor with a smaller one might be possible in the future. Stabilisation along the vertical axis is not critical for the planned experiments which have the quantisation axis aligned in the horizontal plane.

---

[18] Owis MVT 40B-D15-Z-SK-UM & Standa 7T34NM-20

[19] polyether ether ketone (PEEK)

[20] Most of this deviation is from the sensor itself which has a specified sensor to body angle tolerance of <3.5°.

Compensation coils

To compensate the field deviations measured by the sensors, pairs of large compensation coils are installed around the optical table on a frame from aluminium extrusions. It was built to be as large as possible to create a large homogeneous field, minimising the difference between the sensor and atom positions, with approximate dimensions $1.6\,\mathrm{m} \times 1.5\,\mathrm{m} \times 1.2\,\mathrm{m}$. The coils are made from a single loop of 26 AWG ($0.13\,\mathrm{mm}^2$) flat ribbon cable with 26 individual wires.

As the glass cell is not exactly in the centre of the table, and the coils had to be placed around other parts of the experiment the centre of the coils and the centre of the glass cell are shifted by 8 cm along the x direction, -13 cm along the y direction and -4 cm along the z direction. In addition to the offset the axes of the compensation coils align with the table, whereas the sensors are aligned onto the 30° rotated coil axes. Both factors are compensated in the configuration of the stabilisation system, by driving the coils individually and optimising the output matrix parameters.

In addition to the active compensation coils connected to the stabilisation system we installed another coil pair on the vertical axis in the same frame. This pair is used to compensate the earth's magnetic field along the vertical axis during molasses cooling (see section 8.1.2 for details). Each coils has 35 windings, made from enamelled copper wire with 1 mm diameter, generating a field of $\sim$105 mG/A. When compensating earth's magnetic field and we operate them close to their rated current limit of 2.8 A.

6.3.2. Configuration & Optimisation

Sensor optimisation procedure

As described above the sensor must be placed in the correct orientation and position to minimise coupling of the experiment fields. For optimisation we used special experiment sequences which generated short magnetic field pulses in individual coil pairs for 20 ms to 100 ms. The time was chosen such that a stable reading could be observed, but as short as possible as we found fast feedback to be important during the manual sensor alignment. In the following our rough alignment strategy is presented. Final optimisations were performed using an iterative procedure, minimising the observed coupling of the different coils.

We started by orienting the sensor in the mount to ensure we are in the correct quadrant to align the sensor. For this we tilted the mount manually by a large amount and rotated the sensor in the mount. The sensors orientation was fixed at the rotation where the observed signal was minimal. Subsequently the angle was adjusted to observe a minimal reading from the coil pair on the same axis.

Once the orientation is roughly aligned we optimised the position. The first step is to move the sensor in the centre of the coil pair on the same axis, by minimising the observed signal of the coil in gradient configuration. Subsequently the orthogonal position can be set by observing the signal from the same coil pair in offset configuration.

After the rough alignment further optimisations were performed iteratively to reduce the coupling observed from all coils. We could not find an ideal optimisation procedure as

	X1	X2	Y1	Y2	Z1	Z2
Sensor X	0.93	0.46	−0.70	−0.24	0.07	0.07
Sensor Y	0.29	0.43	1.00	0.58	0.11	0.14
Output X	0.49	1.00	−0.19	−0.53	0	0
Output Y	0.49	0.58	0.33	1.00	0	0

Table 6.4.: Measured normalised coupling amplitude for excitations from each individual coil to the installed sensors and derived output scaling factors for each output to the compensation coils.

angle and position of the sensor are coupled. We achieved the best results by optimising the angle to reduce the influence of the orthogonal coils, especially those in vertical direction.

Feedback configuration

A powerful feature of the chosen stabilisation system is the option to drive the stabilisation coils based on a weighted sum of the outputs of all feedback stages shown as the output matrix in fig. 6.9. During setup of the stabilisation system the correct configuration of these output parameters turned out to be the most critical parameter.

The final parameters used now take into account the rotation of the sensor axes relative to the axes of the compensation coils, as well as the different distances of the compensation coils to the sensors. Initially we only included the factor for a 30° rotation. With this configuration we were able to successfully lock the system and stabilise the field, but frequently observed strong fluctuations in the resonance frequency up to 1 MHz as the feedback loop began to oscillate. Adjusting the coil parameters to include the different coupling strength of the compensation coils removed these fluctuations.

The currently used parameters were derived by applying sinusoidal modulations on a single coil at a time and measuring the resulting signal on the installed sensors. A 21 Hz modulation was injected using the excitation feature of the stabilisation system which can add sinusoidal waveforms, broadband white noise or an externally supplied signal to any coil output, even while the feedback loops are active. The modulation frequency was chosen to be incommensurate from the 50 Hz mains frequency, allowing to use averaging to remove periodic noise components. Table 6.4 lists the measured peak-peak values, normalised to the maximal amplitude. Negative values are assigned for readings with a 180° phase shift between excitation and measurement.

From these measurements the factors for the output matrix were computed using the follow procedure: First the ratio within each coil pair was computed per sensor. The coil with the weaker amplitude was set to 1, the stronger to the inverse ratio. Next the same procedure was applied using the X and Y compensation coil pairs per output. For this the ratio between the weaker coils in each pair was computed. Both previously derived values for the weaker coil pair were then reduced by this factor to generate the final amplitudes. Negative signs were added as before based on the measured phase information. The resulting values from the latest calibration are listed in the second half of table 6.4. The procedure was only applied to the X and Y axes as for these a significant overlap due to the rotation between measurement and compensation is expected.

Further improvements to the system performance might be possible by carefully tuning the computed values to shift the centre closer to the atom position, rather than the sensor position.

6.4. Field characterisation

We performed various measurement to characterise the magnetic field stability. First characterisations are performed using the sensor of the magnetic field stabilisation system (section 6.4.1). These can be performed very quickly, but have only reduced information value as they rely on the actively stabilised in-loop sensor reading. To gain reliable information about the stability we measured the magnetic field at the atoms using microwave spectroscopy. We characterised the field fluctuations at various time scales by driving Rabi oscillations in thermal clouds and MI states (section 6.4.3). We also characterised and compensated the magnetic field gradients using Ramsey interferometry (section 6.4.4).

6.4.1. Sensor measurements

All optimisation and initial characterisation steps were done using the monitoring outputs of the magnetic field stabilisation system. When the stabilisation is active this is a measurement using the in-loop sensor. It would be advantageous to use an extra out-of-loop sensor, however due to the placement difficulties described in section 6.3.1, it is not possible to place it sufficiently close to the atoms. The measurements were only performed using the sensor along the x_c axis. We focussed on this axis initially as it aligns with the currently used quantisation axis. With the currently installed sensors the field can be stabilised simultaneously along both axes of the horizontal plane.

Figure 6.12 shows the magnetic field measured using the x_c sensor, separated into drifts of the DC offset field (fig. 6.12a) and mains synchronous fluctuations (fig. 6.12b). Without stabilisation the DC components fluctuate strongly by approximately 2 mG peak-to-peak. With active stabilisation these are suppressed to 70 µG peak-to-peak. Similarly the mains synchronous fluctuations are suppressed from 4 mG to ~200 µG.

We also computed the magnetic field spectrum using a FFT (fig. 6.13). It shows the dominant magnetic field noise from the 50 Hz mains frequency and its higher harmonics, which are suppressed by more than an order of magnitude. The spectrum also shows continuous suppression of the noise up to ~2 kHz.

6.4.2. Microwave spectroscopy

All further stability measurements were performed using microwave spectroscopy on the cold atomic sample. Using the atoms as a magnetic field probe increases the measurement accuracy and allows us to probe the magnetic field at the relevant position.

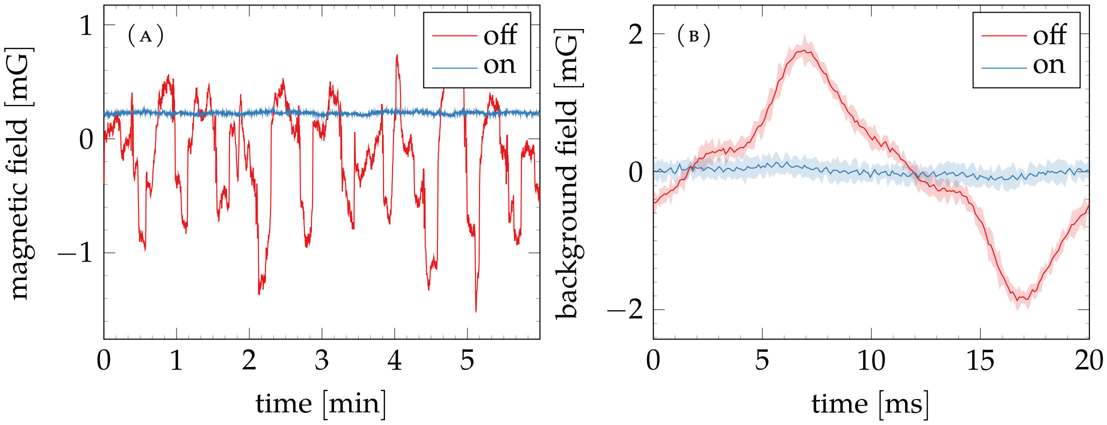

Figure 6.12.: Comparison of the magnetic field along the x_c axis with and without active stabilisation, measured using the sensor of the magnetic field stabilisation. (a) Magnetic field trace over 6 min. Data was binned and averaged over 20 ms to remove mains synchronous fluctuations. (b) Mains synchronous fluctuations. Data was averaged over 300 periods after subtracting the mean value to exclude shot to shot drifts. The shaded area denotes the 3σ interval.

All measurements discussed in this chapter are performed using a microwave transfer from the $F = 3$, $m_F = 3$ ground state to the $F = 4$, $m_F = 4$ state. The energy splitting between the states is magnetic field sensitive due to the Zeeman effect. In first order the energy shift of each state is given by

$$\Delta E_{|Fm_F\rangle} = \mu_B g_F m_F B. \tag{6.4}$$

The used hyperfine states have opposite Landé hyperfine g-factors $g_F = -1/4$ and $g_F = 1/4$, resulting in total energy splitting between the used states of

$$\begin{aligned}\Delta E &= \Delta E_{hfs} + \frac{7}{4}\mu_B B \qquad (6.5)\\ &\approx h\,(9.192\,\mathrm{GHz} + B \times 2.45\,\mathrm{MHz/G}) \qquad (6.6)\end{aligned}$$

For all measurements the expected energy splitting is computed using the exact *Breit-Rabi formula* [150]. Even though the energy shifts due to the used magnetic fields are small compared to the hyperfine splitting, the corrections to the simple first-order Zeeman equation are already significant. For the used $F = 3$, $m_F = 3$ to $F = 4$, $m_F = 4$ transition at a usual field of 20 G the approximation error is already 93 kHz, corresponding to a magnetic field difference of 38 mG.

Using the $F = 4$, $m_F = 4$ excited state results in the largest possible frequency sensitivity, ideal for characterisation of the magnetic field stability. It can also be imaged independent of the atoms in the $F = 3$, $m_F = 3$ ground state by using σ^+ polarised light on the closed transition from $6^2S_{1/2}$ $F = 4$ to $6^2P_{3/2}$ $F = 5$.

The measurements shown in the following were performed by applying fixed frequency microwave pulses with controllable duration. We are limited to square pulse shapes

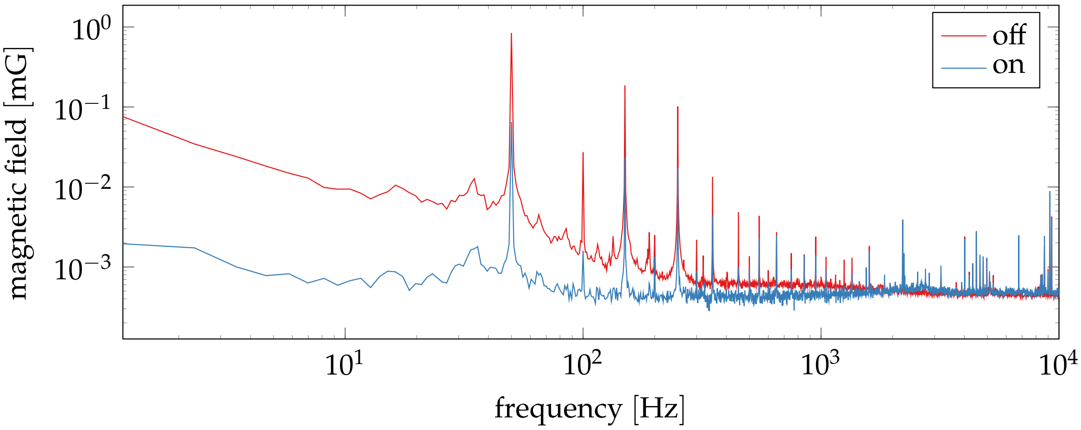

FIGURE 6.13.: Frequency spectrum of the environmental magnetic field along the x_c axis from a fast fourier transform (FFT), measured using the (in-loop) sensor of the magnetic field stabilisation with and without active stabilisation. Twenty-six measured spectra were averaged to reduce the measurement noise.

without phase control as the pulses are generated using a fast RF switch in the amplifier setup (section 4.3). More elaborate pulse shapes and control over the relative phase between pulses could be implemented in the future by improving the programming of the function generators. Besides pulses we also use Landau-Zener (LZ) transfers with a width up to 1 MHz when scanning the frequency over a wide range to determine the resonance position.

6.4.3. Rabi oscillations

Rabi oscillations in thermal clouds

First measurements and calibration of the magnetic field were performed with thermal atoms loaded into a crossed dipole trap. These measurements have the advantage of a shorter cycle time as no evaporation is performed. Using absorption imaging we can also measure both spin components, removing signal variations due to atom number fluctuations.

For these measurements we start with the cloud loaded into the optical traps (section 7.3.1) and prepare the magnetic field amplitude and rotation as desired. We synchronise the experimental sequence to a mains zero crossing before we perform the microwave transfer to suppress mains synchronous noise. To measure the individual spin components within a single realisation we first measure only the transferred atoms in $F=4$ using absorption imaging. We then remove any remaining atoms using a 10 ms blow-out pulse of the same resonant light. Finally we transfer the remaining atoms from $F=3$ to $F=4$ using a short repumper pulse and capture them on a second image. From the extracted optical densities the atom numbers per state and the resulting transfer efficiencies are computed.

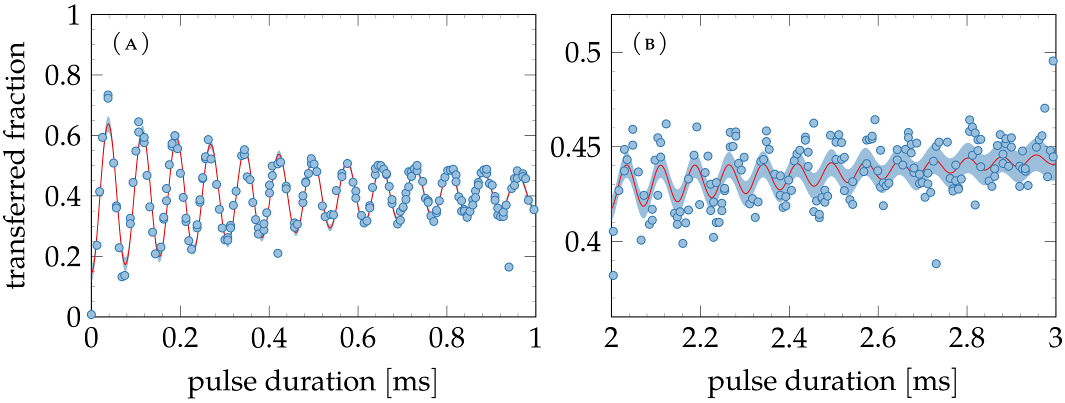

Figure 6.14.: Rabi oscillations measured using a thermal cloud in the crossed dipole traps between the $F = 3$, $m_F = 3$ and $F = 4$, $m_F = 4$ states. Both plots show data from the same measurement. An exponentially damped sine was fitted to the data, measuring the Rabi frequency $\Omega \approx 2\pi \times 13\,\mathrm{kHz}$ and the $1/e$ decay time $\tau = 658\,\mu\mathrm{s}$. An additional linear term was added to capture the increase of the transferred fraction which we attribute to reduced coupling due to the high densities in the dipole trap.

The minimal resonance we can measure in this system is approximately 13 kHz full width at half maximum (FWHM), most likely limited by interaction effects. The pulse duration was chosen long enough to exclude Fourier limitations. Also we can not observe any sinc side lobes which would be expected if the peaks would be Fourier limited. We also excluded power broadening by reducing the pulse intensity until no further reduction in the peak width could be observed.

We use this system to calibrate the magnetic field generated by our offset coils. We measure the resonance position at multiple magnetic field values, and compute the calibration from the generated frequency shift.

We can also use this system to established a lower bound on the magnetic field stability. We drive resonant Rabi oscillations and measure the transferred fraction in relation to the pulse duration. The measured trace is shown in fig. 6.14. To fit the measured data we use an exponentially damped sine with $1/e$ decay constant τ:

$$f(t) = a + bt + \sin(\omega t - \pi/2)e^{-t/\tau}. \tag{6.7}$$

We add the initial linear term to capture an observed slow increase of the transferred occupation of $2\,\%\,\mathrm{ms}^{-1}$. We attribute this increase, as well as the reduced initial transfer efficiency of $\sim$80 % to reduced coupling of the microwave due to the high densities in the cross dipole trap. We verified that the reduced initial contrast is not due to a detuning of the microwave frequency.

From the fit we measure a Rabi frequency $\Omega = 2\pi \times 13.03\,\mathrm{kHz}$ and a decay constant $\tau = 658\,\mu\mathrm{s}$. Assuming the decay to be dominated by normally distributed magnetic field fluctuations the rms noise σ_B can be computed from the coherence time by [151]

$$\sigma_B = \frac{1}{\tau(7/4\mu_B/\hbar)}. \tag{6.8}$$

From the measured decay time the rms noise is computed to be $\sigma_B = 98.8\,\mu\text{G}$.

Rabi oscillations in Mott-insulators

To eliminate the influence of interaction effects and to resolve local inhomogeneities all further measurements were performed using Mott-insulator (MI) states in a 2 dimensional optical lattice. In this state tunnelling is suppressed and the atoms are isolated from each other on the timescales of the experiment. After the state preparation described in section 7.4 we synchronise the experiment with the mains phase and perform the microwave transfer. We use fluorescence imaging to measure the local atom distribution (section 7.3.3). In contrast to the absorption imaging used in the dipole traps we can not detect both states as the spins are mixed during the imaging process. Instead we remove the atoms transferred to $F=4$ from the lattice using a 10 ms blow-out pulse, before we image the remaining atoms in the $F=3$ ground state.

As before we use microwave pulses of fixed frequency to drive the $F = 3$, $m_F = 3$ to $F=4$, $m_F=4$ transition. Our high-resolution imaging allows us to perform locally resolved analysis of the atomic cloud. We measured the oscillations in different intervals up to ~10 ms pulse duration. To analyse the data we detect the mean signal within a 10 px (1.6 µm) wide stripe placed at the centre of the cloud. We determine the local Rabi frequency Ω and the exponential decay constant τ by fitting the acquired data with

$$f(t) = A\sin(2\pi\Omega t)e^{-t/\tau} + c. \tag{6.9}$$

The measured data and fit result are shown in fig. 6.15a. The measured Rabi frequency is $\Omega = 2\pi\times 12.025(2)$ kHz and the fitted decay constant $\tau = 9(1)$ ms.

The observed oscillations are initially homogeneous throughout the cloud (fig. 6.15b), but at longer timescales an apparently random stripe pattern appears (fig. 6.15c-e). The stripe pattern can be explained with the presence of a background gradient, creating a locally varying detuning $\delta(x)$. Due to the detuning the evolution of the atoms is described by a position dependent generalised Rabi frequency

$$\Omega(x) = \sqrt{\Omega_0^2 + \delta(x)^2}, \tag{6.10}$$

where Ω_0 is the resonant Rabi frequency. We measure the local variation of the Rabi frequency throughout the cloud, by repeating the analysis for different centre locations of the analysis region (fig. 6.15g). The measured data roughly follows the expected parabola, with a deviation from the expected shape on one side of the cloud.

Assuming the frequency deviations can be explained with a linear detuning gradient Δ_δ, e.g. due to a magnetic field gradient, we fit the Rabi frequencies using

$$\Omega(x) = \sqrt{\Omega_0^2 + (\Delta_\delta(x - x_0))^2}, \tag{6.11}$$

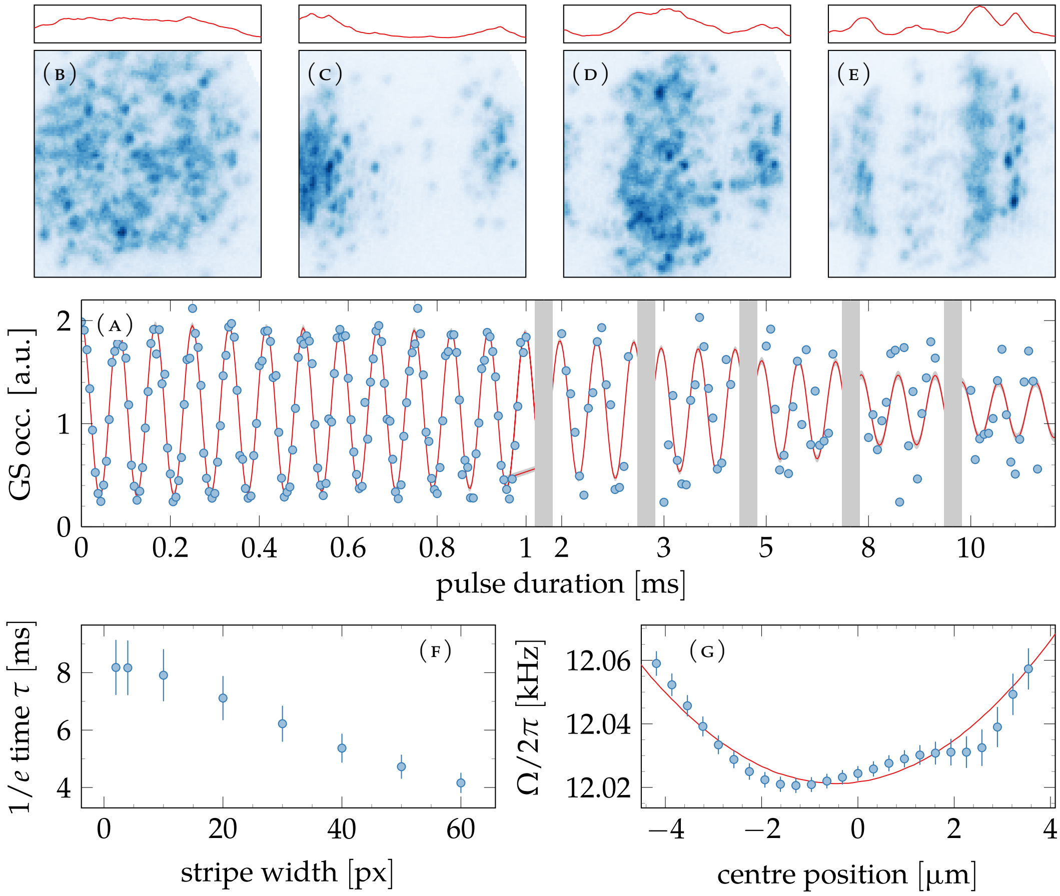

Figure 6.15.: Rabi oscillations measured in the Mott insulator state. (a) Oscillations of a single stripe with 10 px width in the centre of the cloud for varying microwave pulse duration and a fit to the data. The fit is an exponentially dampened sine wave. (b-e) Example images column profiles (red on top) of states after 140 µs (b), 2.05 ms (c), 3.02 ms (d), and 10.15 ms (e) pulse duration. (f) Reduction of the fitted decay time with increasing analysis stripe width. The increased averaging region washes out the stripe pattern, decreasing the measured contrast. (g) Varying local Rabi frequencies measured by analysing stripes of 10 px width at different centre positions in the cloud. Each point is the result of a fit to a time trace as shown in (a). The zero position is set to the centre of the analysis window used in the other traces.

where x_0 is the resonance position in the cloud. The fitted resonant Rabi frequency is $\Omega_0 = 2\pi \times 12\,024.3(8)$ Hz, and the detuning gradient is $\Delta_\delta = 2\pi \times 234(10)$ Hz/µm. From the known background magnetic field gradient $\Delta_B = 1.33\,\mathrm{G\,cm^{-1}}$ (see section 6.4.4) a linear detuning gradient of $2\pi \times 325$ Hz/µm is expected.

The difference of ∼40 % and the deviation of the expected parabola shape could be due to the differential light shift between $F=3$ and $F=4$ of the lattice beams and dipole traps [152, 153]. The expected shift from the lattices and dipole traps of $2\pi \times \sim 20$ Hz per beam is of the same magnitude as the observed shifts. Absolute statements are non-trivial as the shifts of the blue detuned lattices and red detuned dipole traps have opposite sign and the exact local intensity distribution is not known. Further measurements are necessary to verify this hypothesis.

Depending on the width of the analysed region we observe a faster decay of the signal (fig. 6.15f). This increase is expected as the magnetic field noise is artificially increased by averaging over a range of detunings generating the stripe pattern. For the chosen measurement time the decay time saturates for stripe widths of approximately 10 px. We chose 10 px wide stripes for all further analysis.

Independent of the chosen stripe width the decay time of this measurement is significantly longer than the measured decay time in the crossed dipole trap shown above. The resulting value for the magnetic field noise of $\sigma_B = 7.4\,\mu\mathrm{G}$ appears to be too low and disagrees with further measurements performed using Ramsey spectroscopy described below. Yet even for pulses with 10 ms the phase of the oscillations appears to have only minor fluctuations. Currently the reason for these observed long decay times is not clear. Systematic drifts appear unlikely as the data was acquired within almost two hours with randomised sampling. Simulations of the expected decay showed an agreement between the measured decay and the calculated magnetic field noise below 10 µG. Time constraints prevented further studies and an extended verification of this measurement.

π stripe in Mott-insulator

To measure slow variations of the magnetic field we track the spatial position of a single stripe in the MI, generated by applying a single π pulse. Without any compensation a residual magnetic field gradient with a strength of $1.33\,\mathrm{G\,cm^{-1}}$ is present (see calibration on the following page), resulting in a locally varying detuning with a slope of $326\,\mathrm{Hz\,\mu m^{-1}}$. If the stability is high enough all atoms within a small band are transferred to the $F=4$ state. Before imaging we remove the transferred atoms using a blow-out pulse, creating an empty stripe in the cloud (fig. 6.16b).

The width and location of the generated stripes is determined by the stability of the magnetic field during the excitation. To measure fluctuations of the fields between different experiment realisations we generate only a single stripe and analyse its position. The microwave power is reduced to extend the duration for a π pulse to approximately 1 ms. With these settings the width of the stripe is 1.7 µm FWHM. This value is in agreement with the expected Fourier limit of 500 Hz for the chosen pulse duration in the background gradient. The value was chosen to create a narrow stripe with sufficient contrast for a Gaussian fit used for data analysis.

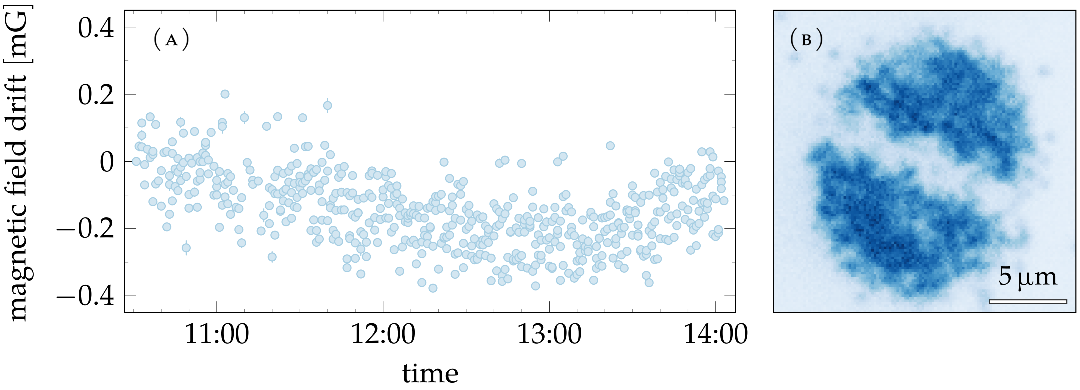

Figure 6.16.: (a) Long term magnetic field drift and shot to shot fluctuations measured using microwave spectroscopy in the MI state. The drift is computed from the position shift of the π stripe in the MI in the background gradient field of 1.33 G cm^{-1}. The standard deviation of is 83 µG, the maximal drift rate is 140 µG h^{-1}. (b) Example in situ image of the stripe in the MI state, generated using a 1 ms microwave π pulse. The FWHM of the stripe is 1.7 µm, corresponding to frequency difference of 500 Hz.

To measure the magnetic field stability we repeat the identical measurement for several hours, analysing the centre position of the stripe within the cloud. For each acquired image the stripe position is determined using a Gaussian fit to the rotated and column averaged image. The measured positions are translated in magnetic field values using the independently calibrated strength of the background gradient of 1.33 G cm^{-1} (fig. 6.16a). From the measured data we determine a standard deviation of 83 µG and a long term drift of maximally 140 µG h^{-1}.

To verify that the stripe width is Fourier limited and not broadened by magnetic field fluctuations we measured the peak width for different π pulse durations (fig. 6.17a). We observe the expected linear relation between the pulse duration and the measured FWHM of the generated stripe down to a stripe width of ~1.1 µm (fig. 6.17b). Only when we extend the pulse duration further to almost 5 ms we observe a deviation, measuring a line width of 800 nm (fig. 6.17c). This width corresponds to 260 Hz in the known gradient. The same width would be expected from a infinitely long pulse in a field with a stability of 106 µG.

6.4.4. Ramsey interferometry

To characterise the magnetic field stability on shorter timescales we are using Ramsey spectroscopy [154]. An initial $\pi/2$ pulse prepares the system in a superposition of the two measurement states. The system then evolves for a variable time t, before another $\pi/2$ is applied to project into the initial states. The system evolves depending on the detuning δ between the pulse and transition frequencies, resulting in oscillating state population with varying detuning.

Without additional compensation a background gradient field is present in our exper-

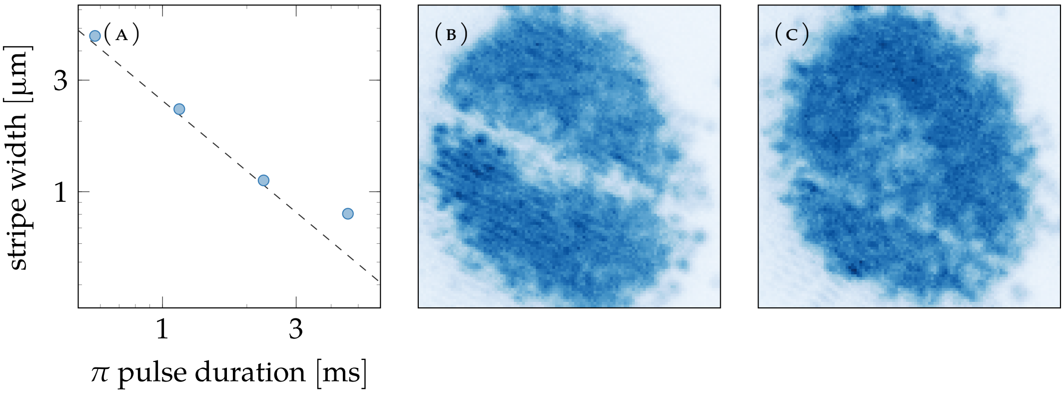

Figure 6.17.: Measured width of a stripe generated in the MI state by applying a single π pulse with varying duration on a double logarithmic plot (a). The dashed line shows the expected peak width based on the Fourier limit of the pulse and the detuning due to the magnetic field gradient. Example images of the smallest Fourier limited stripe with a pulse duration of 2.3 ms (b) and the smallest measured stripe with a pulse duration of 4.6 ms (c) are shown.

iment, creating an approximately linear energy shift along the quantisation axis. The locally varying detuning generated from this shift creates the characteristic Ramsey fringes within each individual image (fig. 6.18).

For this measurement we acquired ∼300 images with a variable evolution time from 400 µs to 1.5 ms in the uncompensated background gradient. The timescale was chosen to create fringes suitable for fitting, generating a period smaller than the size of the system and larger than resolution limit of our imaging system. The images were analysed by computing the average along the stripe axis and fitting the measured data using a sine with variable amplitude, frequency and phase.

From the decay of the fringe contrast with increasing evolution time the magnetic field noise can be estimated. The expected shape of the magnetic field decay depends on the type of magnetic field noise and its correlation time τ_t relative to the experiment times [151, 155].

The dominant components of the magnetic field noise are typically on much longer timescales than the observed coherence times up to a few milliseconds [134]. This is also true for our experiment, even with active stabilisation, as shown with the spectrum in fig. 6.13 on page 81. For these types of noise a Gaussian decay of the amplitude c is expected

$$c = e^{-t^2/\tau_g^2}. \tag{6.12}$$

The magnetic field rms noise σ_B is related to the measured coherence time τ_g by [155]

$$\sigma_B = \frac{\sqrt{2}}{\tau_g(2\pi \times 2.45\,\mathrm{MHz/G})}. \tag{6.13}$$

For white noise and other noise with correlation times shorter than the experiment

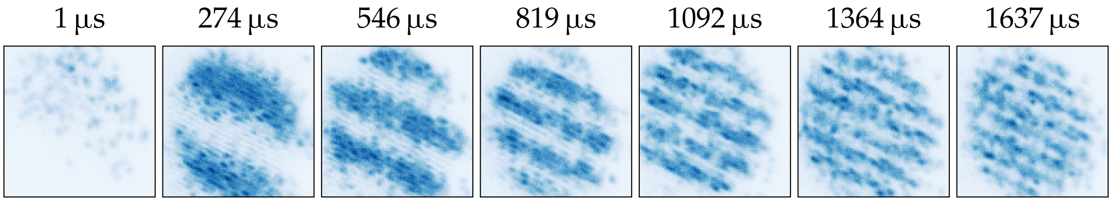

FIGURE 6.18.: Example images of Ramsey fringes imprinted in the MI state using microwave pulses with increasing evolution time 2 shown on top of each image. The uncompensated background gradient creates a locally varying detuning. Before taking the image the transferred fraction was removed using resonant light.

time scales a simple exponential decay

$$c = e^{-t/\tau_e} \tag{6.14}$$

is expected. In this case the magnetic field rms noise σ_B is related to the measured coherence time τ_e by

$$\sigma_B = \frac{1}{\tau_e(2\pi \times 2.45\,\mathrm{MHz/G})}. \tag{6.15}$$

To analyse all magnetic field fluctuations we average the acquired images of same evolution time and determine the contrast and spatial frequencies from the sine fit (fig. 6.19a,c). The quality of the acquired data is not sufficient to determine the type of noise from the measured decay of the contrast. As described above we expect the decay to be dominated by noise with long correlation times and therefore we fit the data using eq. (6.12) (green curve in fig. 6.19a). We measure $\tau_g = 820(90)\,\mu\mathrm{s}$ and $\sigma_B = 113(13)\,\mu\mathrm{G}$. For comparison we fit the same data using eq. (6.14), measuring $\tau_e = 470(70)\,\mu\mathrm{s}$ and $\sigma_B = 140(20)\,\mu\mathrm{G}$.

As we are measuring a portion of the Ramsey stripe pattern with every realisation we can determine the stripe amplitude for each experimental realisation and exclude noise from shot-to-shot fluctuations of the magnetic field. We analyse each image individually and then average the determined fit results (fig. 6.19b,d). We fit the amplitudes using eq. (6.14) and determine a coherence time $\tau_e = 1.22(4)\,\mathrm{ms}$, corresponding to $\sigma_B = 53.3(3)\,\mu\mathrm{G}$. Combining this value with the fluctuations between different realisations on page 85 of 83 μG, results in a total stability of ~100 μG, compatible with the measured total stability of 113 μG shown above.

The same measurement can also be used to determine the strength of the magnetic field from the evolution of the stripe periods (fig. 6.19c,d). A linear fit without offset through the individual data points measures $\Delta_B = 1.334(5)\,\mathrm{G\,cm^{-1}}$.

Gradient compensation

As shown above the residual gradient in our glass cell can be determined from the spacing of the observed Ramsey fringes. To compensate the background gradient we

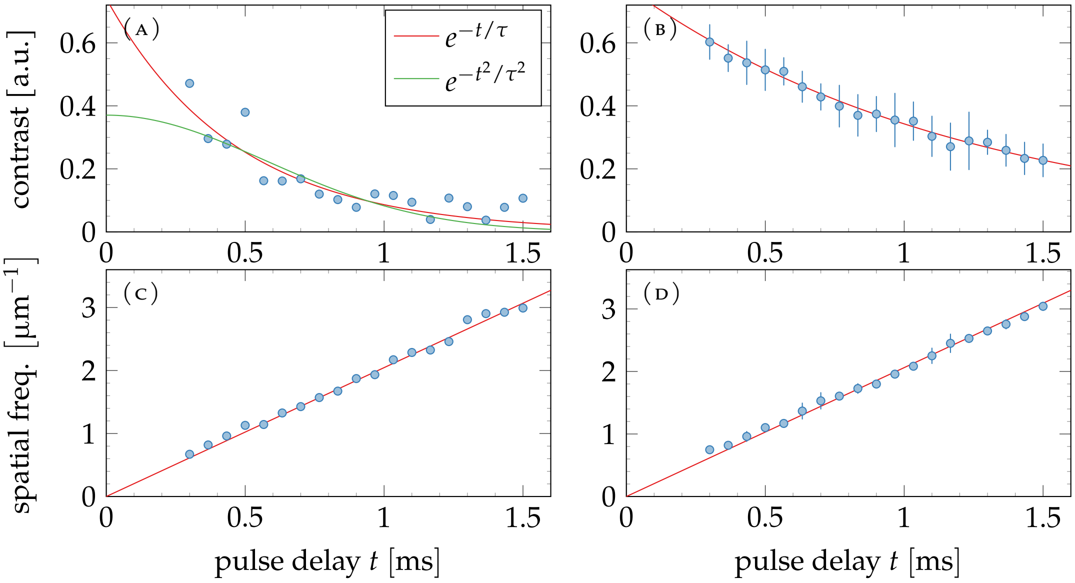

Figure 6.19.: Analysis of the measured contrast and period of Ramsey fringes measured in the MI state for varying time between the Ramsey pulses t. The data of ~300 acquired images (16 per data point) was analysed twice by fitting a sine to the data. Once the images with same delay time were averaged before the analysis (a,c), once they were analysed individually before and the fitted data was averaged (b,d). The latter excludes fluctuations of the magnetic field between shots as the phase of the different fits is not included in the analysis. The averaged data is fitted using an exponential and Gaussian decay to capture noise with different correlation times (see text). From the measured spatial frequencies the strength of the background gradient is determined to be 1.334 G cm^{-1} (a,b).

apply an opposite gradient of equal strength using the primary z coils, followed by fine adjustment of the strength to minimise the residual gradient. Using this method we can compensate the gradient along the quantisation axis (x_c/y'). The stripe pattern is not fully removed due to the residual gradient orthogonal to the quantisation axis. When the gradient along the quantisation axis is scanned the stripe pattern starts to rotate as the influence of the remaining orthogonal components increases (fig. 6.20). The symmetric signal of the stripes angle and period can be used to precisely determine the required compensation field to compensate the gradient. Limited by the calibration of the gradient coils this can also be used to determine the strength of the background gradient.

The compensation gradient was generated using the z coil pair in Anti-Helmholtz configuration. To determine the value at which the field is compensated we transformed the acquired images using a two dimensional FFT and determined the angle and period from the peaks (images in fig. 6.20). To detect the centre position we fitted an arctan to the angle and a Gaussian to the measured period.

The symmetric point of the angle is found to be at 1.334(3) G cm^{-1}, the maximum of the stripe period at 1.332(2) G cm^{-1} in perfect agreement with the gradient measured using the period of the Ramsey stripes on page 86. Note that the measurement error is

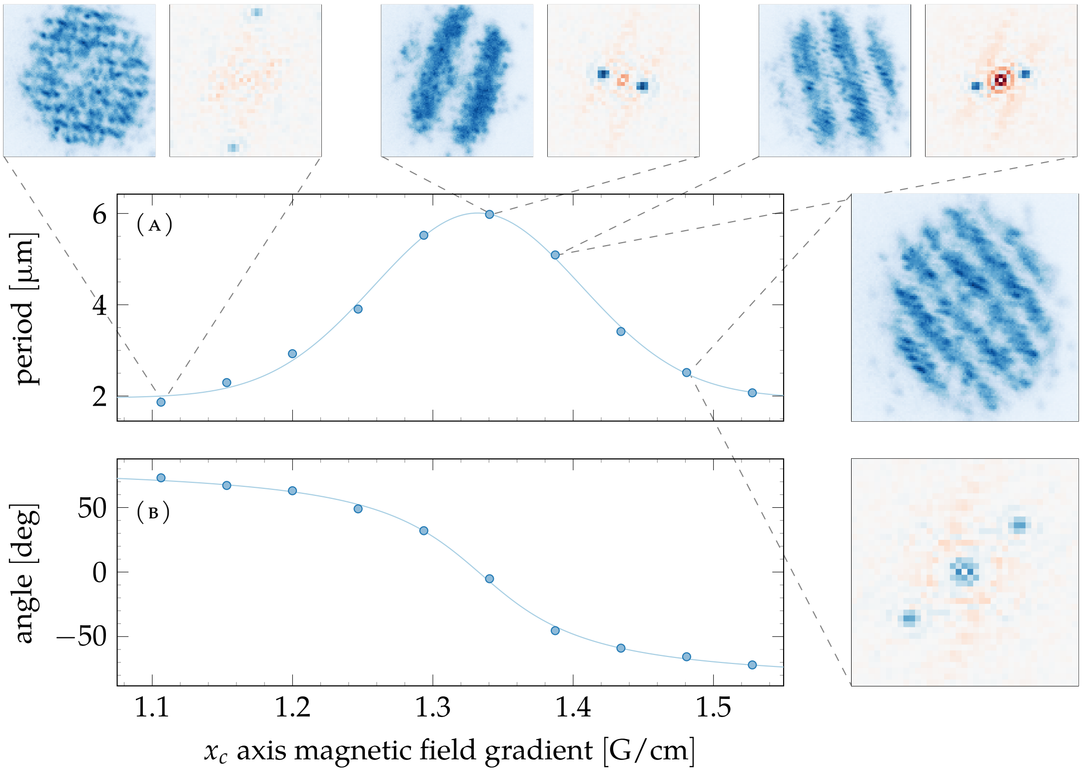

Figure 6.20.: Measurement of the residual magnetic field gradient using Ramsey microwave spectroscopy in the MI state, using 36 µs $\pi/2$ pulses and 10 ms interaction time between the pulses. A magnetic field gradient with varying strength is created using the z coil pair orthogonal to the quantisation axis. The much stronger offset field of 29.8 G suppresses orthogonal gradient components, determining the direction of the stripe pattern. When the gradient along the quantisation axis is fully compensated only orthogonal components remain, rotating the stripe pattern by 90°. The surrounding pairs of images are *in situ* images (left) and their 2D FFT (right). The mean of all FFTs was subtracted to reduce the peak in the centre. (a) Measured period of the strip pattern for the different compensation values, fitted using a Gaussian. At its maximum at 1.332(2) $\mathrm{G\,cm^{-1}}$ the total gradient is minimal as the gradient along the quantisation axis is fully compensated. (b) Measured angle of the stripe pattern for the different compensation values, fitted using an arcus tangens. The centre of the fit where the rotation is orthogonal is at 1.334(3) $\mathrm{G\,cm^{-1}}$.

also comparable to the previously shown measurement, even though 30 times less data points were acquired. The sensitivity of this measurement could be further increased by partially compensating the orthogonal gradient along the y_c axis.

6.4.5. Conclusion

We have constructed a system to control the magnetic fields and gradients in arbitrary rotations, combined with an active stabilisation system achieving magnetic field stabilities of 110 µG. Locally resolved imaging was used to characterise the magnetic field stability, separated into fast fluctuations within and slow fluctuations between experimental realisations. The achieved stability is expected to be sufficient for future experiments with magnetic field sensitive transitions, including the realisation of topological systems in a state dependent lattice (section 2.2.1).

All measurements in this chapter were deliberately performed using the maximally magnetic sensitive transition. For the future experiments with atoms in state-dependent lattices we are planning to use the less sensitive combination of the $F=3$, $m_F=3$ ground state and the $F=3$, $m_F=2$ state . The system is also less sensitive to fluctuations of the magnetic field between realisations, such that the expected energy fluctuations of the used states are less than $h \times 20\,\mathrm{Hz}$.

Further measurements using Ramsey spectroscopy measurements and spin-echo sequences [156], could be used to continue the investigation of the components of the magnetic field noise. These are not part of this characterisation due to the limited measurement time. For the same reason only a single axis of the stabilisation system was used for the characterisation. Stabilisation in both horizontal axis was demonstrated and is expected to provide the same level of stability. With modifications to the sensor hardware active stabilisation in all three axes with simultaneous full field control could be achieved.

The high absolute magnetic field stability combined with full and fast control up to 200 G provides various possibilities for future experiments. The fast reaction times of the power supply allow for fast control of the scattering length, i.e. to quench into a non-interacting system. The high absolute stability also allows to effectively use narrower Feshbach resonances, i.e. the resonance close to 48 G. For this resonance we expect a stability of $\pm 1\,\mathrm{a}_0$ at the zero crossing of the scattering length.

The high absolute stability will also allow to selectively address doubly occupied sites in the optical lattice. Selective spin changes on one or both atoms can be performed by resolving the difference in the on-site interaction energy due to the different intra- and inter-species scattering properties.

7. Condensation and Mott-insulator

After the atoms have been transported to the science chamber they are collected in a crossed dipole trap and cooled further using evaporative cooling methods until we create a Bose-Einstein condensate (BEC). To perform experiments the degenerate cloud is subsequently compressed vertically to form a 2 dimensional plane and loaded into an optical lattice, where we typically prepare a Mott-insulator (MI) state as foundation for the state preparation.

This chapter summarises the installed beams for optical dipole traps, lattices, and optical cooling (section 7.2) and describes the methods and sequences we use for state preparation and imaging (section 7.3). Using these techniques we show the creation of MI states (section 7.4) and show first measurements performed in our optical super-lattices (section 7.5).

7.1. Overview

During the hunt for the first BEC, Caesium was a promising candidate as were the other Alkali elements. But whereas condensation was achieved for most Alkali elements in 1995 [157–159], the first BEC of Caesium was only reported in 2002 by Weber et al. [105].

Condensation of ^{133}Cs was found to be difficult due to a number of atomic properties, requiring adjustments to the cooling and evaporation procedures. The initial condensates were all generated using magnetic trapping, a method not possible when using the absolute ground state of Caesium. Magnetically trappable states such as $F=3$, $m_F=-3$ and $F=4$, $m_F=4$ were found to be unsuitable due to strong losses from inelastic two-body collisions [160]. Optical condensation in a dipole trap allowed to overcome this limitation by using the absolute ground state, which is protected from two-body collisions. It was however found that even for this state three-body loss rates are significantly higher than in other alkali atoms, increasing the difficulty when condensing Caesium. A minimum of the loss rate is found relatively close to the zero crossing of the Feshbach resonance around 21 G ($a = 210\,\mathrm{a}_0$) due to Efimov physics [161]. To successfully condense Caesium the magnetic field has to be controlled to adjust the scattering and loss properties during the evaporation to improve the efficiency.

Since then various groups have demonstrated the creation of Caesium BECs [133, 162, 163] and alternate methods for condensation have been shown, either combining magnetic gradients and optical traps [133], or using sympathetic cooling [163]. We are using a pure optical evaporation scheme (section 7.3.1), combined with a second radial evaporation stage when preparing the final experimental samples (section 7.3.2).

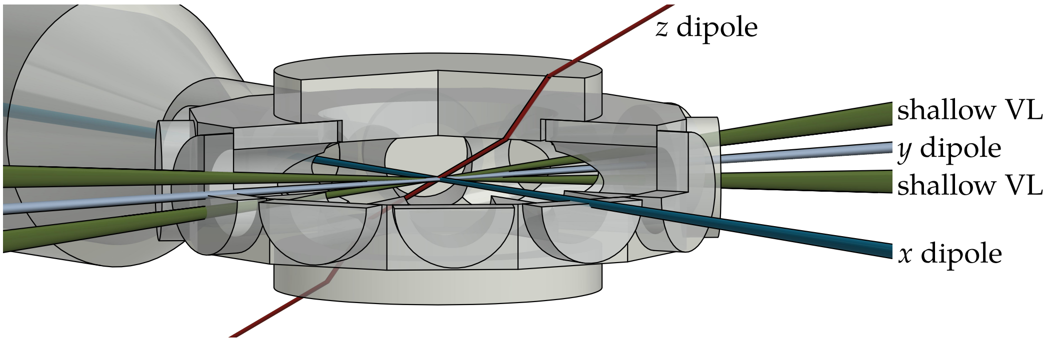

FIGURE 7.1.: Orientation of the 1064 nm trapping potentials in the science chamber. The x and y dipole beams create the crossed dipole trap used for optical evaporation. The lattice created by the two beams of the shallow vertical lattice (shallow VL) is used for vertical compression into a 2d system. Lattice experiments are performed in the shallow VL combined with the angled z dipole trap for horizontal confinement.

7.2. Beam geometry

7.2.1. Dipole traps

Throughout all steps of the sequence in the science chamber the atoms are trapped optically using dipole trap beams. All beams are using light with a wavelength of 1064 nm. In the following the beam properties and their application is described. The laser sources used to generate the light are described in section 3.2.2.

Horizontal dipole traps

For optical evaporation we transfer the atoms from the transport beam from the elongated transport potential formed by the Bessel beam into a elliptical crossed dipole trap. The trap is formed by two beams, propagating along the x and y axis, called *x dipole* and *y dipole* respectively.

The x dipole trap is created using a circular beam, co-propagating with the Gaussian transport beam in the $-x$ direction, shown dark blue in fig. 7.1. We use powers up to 1.2 W at the atoms, focussed to $100\,\mu\text{m} \times 100\,\mu\text{m}$ $1/e^2$ waist, resulting in trap depths $U_0/k_B = 18.0\,\mu\text{K}$ and trapping frequencies of $\omega_y = \omega_z = 2\pi \times 107\,\text{Hz}$.

The y dipole trap is propagating perpendicular to it along the $+y$ direction. It is co-propagating with the shallow VL described in section 7.2.2, passing through the same 100 mm cylindrical lens, creating an elliptical trap with $370\,\mu\text{m} \times 110\,\mu\text{m}$ spot size. The larger axis is oriented along the x axis increasing the overlap with the atomic cloud after optical transport. To achieve comparable trapping parameters to the x dipole trap we use powers up to 8 W, resulting in trap depths up to $U_0/k_B = 28.5\,\mu\text{K}$, and trapping frequencies $\omega_x = 2\pi \times 36\,\text{Hz}$, and $\omega_z = 2\pi \times 119\,\text{Hz}$.

The potential generated by the y dipole trap is not purely harmonic, but has an additional fringe pattern, observable once we create a BEC and even more pronounced in the

MI state. We suspect interference with a weak reflection under a shallow angle due to the large fringe spacing around 4 µm, but could not find their source. Our current sequence avoids this problem by performing the final radial evaporation (see section 7.3.2) without the y dipole trap.

The beams for both dipole traps are generated from the same 1064 nm fibre amplifier (see section 3.2.2). They are passing through independent AOMs used for intensity stabilisation before they are coupled into high power fibres. The modulators shift the frequency of the two beams in opposite directions, resulting in a frequency difference of 160 MHz to avoid interference effects. Additionally we made sure that the beams have perpendicular polarisations.

Vertical dipole trap

To simplify the generation of a uniform MI state we included another dipole trap to confine the atoms horizontally once they are loaded into the vertical lattice. Originally we were planning to use the two horizontal dipole traps at lower power. We had to abandon this plan due to significant corrugations on the y dipole trap, observable as a pattern on the atomic cloud once the atoms were cold enough, as mentioned above. Using only the x dipole and the shallow VL the generated MIs were strongly elliptical due to the asymmetric horizontal confinement.

To decouple vertical and horizontal confinement we installed an extra vertical dipole trap using light from our DPSS laser system. It enters the glass cell under a 60° angle from the z axis, as the direct access is blocked by the high-resolution objectives.

The beam is shaped to be elliptical with a 110 µm × 53 µm waist at the atoms. This ratio compensates for the shallow angle of incidence, creating an almost round trapping potential of 110 µm × 106 µm in the horizontal plane. The resulting horizontal trapping frequencies for typical powers of 500 mW are $\omega_{\parallel} = 2\pi \times 71\,\mathrm{Hz}$ along the beam propagation, $\omega_{\perp} = 2\pi \times 74\,\mathrm{Hz}$ orthogonal to it, and $\omega_z = 2\pi \times 122\,\mathrm{Hz}$ in the vertical direction. When measuring trapping frequencies in the combined potential with the shallow VL we find an almost symmetric configuration of $\omega_x = 2\pi \times 55\,\mathrm{Hz}$ and $\omega_y = 2\pi \times 59\,\mathrm{Hz}$.

Potential distortions When aligning the new independent beam path of the vertical dipole trap we found that depending on the entry position of the beam to the glass cell the potential showed strong distortions. These were only weakly visible when imaging the cloud directly in the trap, but showed very clearly when inducing oscillations for trapping frequency measurements. The strength of the distortions is clearly dependent on the entry position into the glass cell, cleaning the components in the beam path did not have any effect.

Due to the limited access to the glass cell we can not determine if there are visible damages on the glass cell. We managed to find an entry position without noticeable distortions, which we use since then.

7.2.2. Shallow angle vertical lattice

To perform experiments in the two dimensional regime we require strong vertical confinement, which we create using an optical lattice along the vertical direction. We are using a running wave lattice, created by two beams entering the chamber on the y axis under half opening angle $\alpha = 3.6^\circ$. They create a lattice with spacing $a = \lambda/2\sin(\alpha) = 8.4\,\mu\text{m}$ using 1064 nm light. In contrast to the planned vertical lattice under 30° described in section 4.5 it is called shallow vertical lattice (shallow VL).

The lattice is created using 1/2 inch PBS cubes stacked on top of each other, horizontally offsetting part of the beam by the size of the cubes (see refs. [59, 164]). The two parallel beams are passing through the same cylindrical 100 mm lens, converting the vertical displacement of the parallel beams into the opening angle. The cubes are glued into a temperature stabilised mount to maintain a constant distance between the beams and with it a stable lattice phase, the construction and alignment are described in ref. [59]. Since then we decreased the beam size to $1/e^2$ waists $\omega_x = 170\,\mu\text{m}$ and $\omega_z = 35\,\mu\text{m}$, increasing the possible trapping frequencies to $\omega_z \approx 2\pi \times 13\,\text{kHz}$ and potential depths $U/k_B \approx 500\,\mu\text{K}$ at the maximal combined power of 14 W in both beams.

7.2.3. Experiment lattices

The twelve-sided shape of our science chamber allows us to install two independent lattice configurations for experiments, and still maintain independent viewports for the optical transport and dipole traps. We intend to use this flexibility by installing one configuration with an anti-magic lattice at 871 nm to perform measurements with complex topologies as described in section 2.2.1, and another configuration using ‘regular’ lattices. This will allow us to switch between lattice configurations quickly without replacing optical components.

We started by installing the regular lattice configuration along the x'' and y'' axes, as it is simpler and allows for many experiments to characterise the machine. It is a square lattice created from two retro-reflected laser beams with a wavelength of 767 nm. Both lattice axis have been upgraded to *super-lattices* by installing co-propagating beams with twice the wavelength of 1534 nm. All lattice beams are intensity stabilised and fibre coupled to the experiment as described in section 3.2.3.

The beam paths of both lattice axes are almost identical (fig. 7.2) and only differ in the path between atoms and retro-reflecting mirror. After the fibre the polarisation is cleaned and a small fraction of the light is split off for intensity monitoring and stabilisation. Optical isolators are used to block reflections at the fibre tips. For both wavelengths telescopes are used for beam shaping. A cylindrical telescope made from lenses with $f_1 = 300\,\text{mm}$ and $f_2 = -100\,\text{mm}$ is used to create an elliptical profile on the 767 nm beams, the 1534 nm beams are reduced in size using $f_1 = 150\,\text{mm}$ and $f_2 = -75\,\text{mm}$. Both beams are then overlapped using a dichroic mirror before they are focussed on the atoms with a 150 mm lens. Irises are installed in between the last mirrors and the atoms to block stray light, e.g. from secondary mirror surfaces.

The $1/e^2$ beam waists at the atoms are $\omega_\perp = 150\,\mu\text{m}$ and $\omega_z = 45\,\mu\text{m}$ for the short

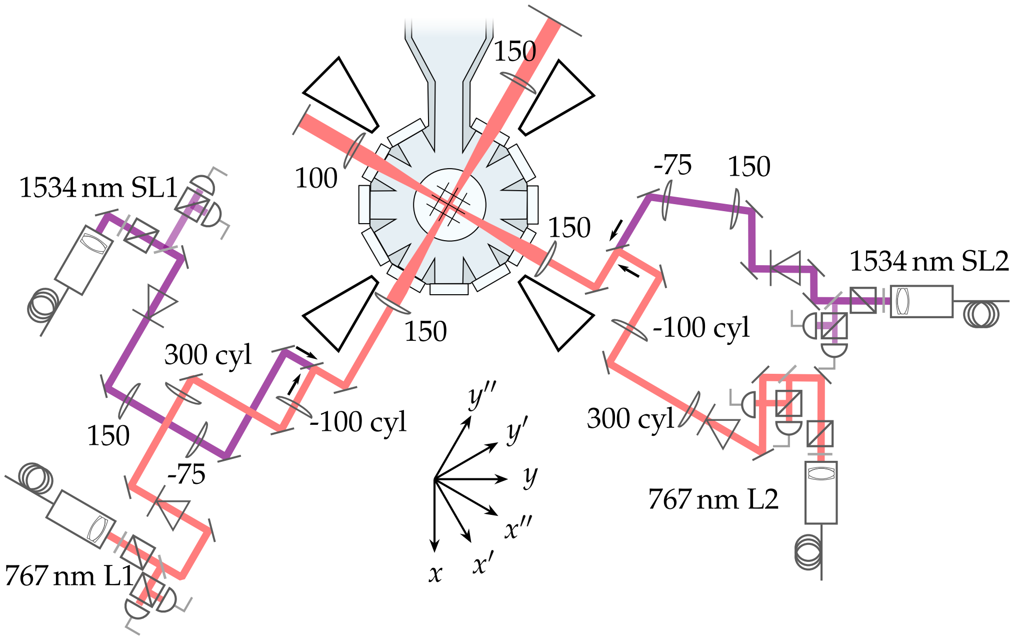

FIGURE 7.2.: Paths of the horizontal lattice beams through the science chamber. Both axes are built identically with the exception of a different lens in the retro path. The triangular objects next to the glass cell indicate the location of the pillars supporting the breadboards and objectives.

spaced 767 nm lattices and $\omega_\perp = \omega_z = 140\,\mu\text{m}$ for the long spaced 1534 nm lattices. We use elliptical beam profiles for the short spaced lattices to increase the intensity, which are also used as pinning lattices during fluorescence imaging (section 7.3.3). Their vertical extend is not experimentally relevant as all experiments are performed in two dimensional systems. The chosen ratio is a compromise between the achievable peak intensity and the implementation difficulty, as smaller beams are more difficult to align initially and are more sensitive to pointing fluctuations.

Using 3 W of 767 nm light we achieve trapping frequencies along the lattice direction up to $\omega = 2\pi \times 280\,\text{kHz}$ and potential depths $U \approx k_B \times 380\,\mu\text{K} \approx 3000 E_{r,767}$, using the recoil energy $E_{r,\lambda} = (2\pi\hbar/\lambda)^2/2m$. During experiments we work at considerably lower powers up to 100 mW in the regime up to $60 E_{r,767}$ for the short spaced lattices and $30 E_{r,1534}$ for the long spaced lattices.

Alignment procedure

To perform the initial lattice alignment we place all components in the correct distances as accurately as possible. We then install cameras behind the retro-reflecting mirrors which we focus on the atoms. Each camera has a single focussing lens and is pre-aligned to infinity before installation. We then adjust the position of the lens in the retro-beam path to focus the atoms on the camera. This correctly positions the lens up to a small deviation due to the wavelength difference between the imaging and lattice wavelengths.

On the now aligned camera we record the position of the atoms and move the incoming lattice beam to this position. To optimise the alignment we block the reflected path and measure the force induced from the repulsive laser beam. The atoms in the cloud are illuminated with the lattice beam for 1 ms before all beams are turned off and the cloud is measured using time-of-flight imaging. When scanning the position of the beam from one side of the cloud to the other the position of the cloud follows a dispersive line shape, which can be used to centre the beam on the cloud. We perform this measurement using the cloud in the compressed y dipole trap (section 7.3.2). The increased confinement after compression reduces the gravitational sag of the cloud and changes the vertical position by a significant amount.

Finally the retro mirror is unblocked and the retro reflected beam is overlapped with the incoming beam for rough alignment. This is typically good enough to observe a signal of Kaptiza-Dirac scattering due to the lattice using the same sequence as before. For optimisation the pulse duration is reduced to a few µs, such that only the first diffracted order is visible. In this regime the population of the nth diffracted order is proportional to J_n^2, where J_n are the Bessel functions of the first kind [165]. We choose the pulse duration to be on the initial decaying slope of the 0th order. We then optimise the alignment of the retro-reflecting mirror to reduce the intensity of the 0th order, decreasing the pulse times further if necessary.

Alignment of the super-lattice beams requires initial alignment of the regular lattices. Afterwards we overlap the beams at two positions far away from each other. We use a flip mirror in front of the focussing lens and the faint leak through the retro-mirror. If the short spaced lattices were properly aligned initially we typically achieve 90 % of the expected trap depth using this procedure.

Lattice calibration

To perform experiments and perform meaningful evaluations we need to calibrate our lattice potentials in the relevant parameters U and J of the Hubbard-Hamiltonian

$$\hat{H} = -J \sum_i \left(\hat{a}_i^\dagger \hat{a}_{i+1} + \text{h.c.} \right) + \frac{1}{2} U \sum_i \hat{n}_i (\hat{n}_i - 1) - \mu \sum_i \hat{n}_i. \tag{7.1}$$

We determine the potential depth of our system by measuring the band gap to the third energy band of our system using parametric heating measurements. We modulate the intensity of one of the lattices by ∼2 % for 50 ms with a varying modulation frequency. When the modulation frequency matches the energy spacing atoms are transferred to higher bands and are lost from the trap. To determine the resonance position we fit a Lorentzian profile and extract the centre position. We perform this measurement in deep lattices with typical lattice frequencies of 40 kHz. We use band structure calculations to determine the lattice potential depth V_0 matching the measured excitation frequency in units of lattice recoil. The tunnelling matrix element J is only dependent on the potential depth and is computed using further band structure calculations [1].

Subsequently we determine the on-site interaction energy U using lattice modulation spectroscopy. We prepare an MI state and perform the same modulation as for the parametric heating measurements, using lower frequencies around 1 kHz. When the modulation frequency matches the frequency related to the interaction energy U/h atoms can tunnel on top of each other, restoring previously suppressed tunnelling in the system. During fluorescence imaging doubly occupied sites are projected to empty sites, reducing the measured number of atoms at resonance. To determine U we scan the modulation frequency, fitting a Lorentzian to the measured atom number.

7.3. Experimental sequence

This section describes the different steps needed to prepare our experimental images and our main imaging method.

After the optical transport described in section 3.4 we typically have 1.2×10^6 atoms in the transport lattice at typical temperatures of 2 μK in the transport lattice. To finish the transport we transfer the atoms into a crossed dipole trap which we use to perform optical evaporation to create BECs (section 7.3.1). To perform measurements we load the BEC into a vertical lattice, creating a two dimensional system (section 7.3.2), suitable for *in situ* fluorescence imaging (section 7.3.3).

7.3.1. BEC creation

Dipole trap loading

The sequence used to load the atoms into the crossed dipole trap is depicted in fig. 7.3. In 250 ms we ramp up the y dipole trap linearly to 25 μK, while ramping down the Gaussian transport beam. Removing the Gaussian releases the atoms from the transport lattice along the x axis, where they are collected in the y dipole trap.

We then ramp up the x dipole trap linearly to 18 μK in 600 ms before removing the Bessel transport beam within 500 ms, fully loading the atoms into the crossed dipole trap. As we are using a rotating waveplate to control the intensity of the Bessel it is ramped down in a sine shape, before we use a mechanical shutter to block it completely. At this stage we have a thermal cloud with 6×10^5 to 8×10^5 atoms at 1 μK to 2 μK.

Optical evaporation

After loading into the crossed dipole trap we perform the evaporation process to create a BEC. We are using an all optical evaporation process by reducing the intensity of the dipole traps. Additionally we change the magnetic field to control the scattering length and three-body loss rate.

At the start of evaporation we reduce the magnetic field from 40 G during transport to 26 G, resulting in a scattering length $a = 410\,\mathrm{a}_0$. At the same time we start the evaporation process. Within the first 3 s we reduce magnetic field exponentially to 23.5 G ($\sim 320\,\mathrm{a}_0$) and the y dipole trap linearly to 9.6 μK. At the same time we start reducing the x dipole trap exponentially to 1.8 μK within 6 s. We then hold the created BEC in this configuration

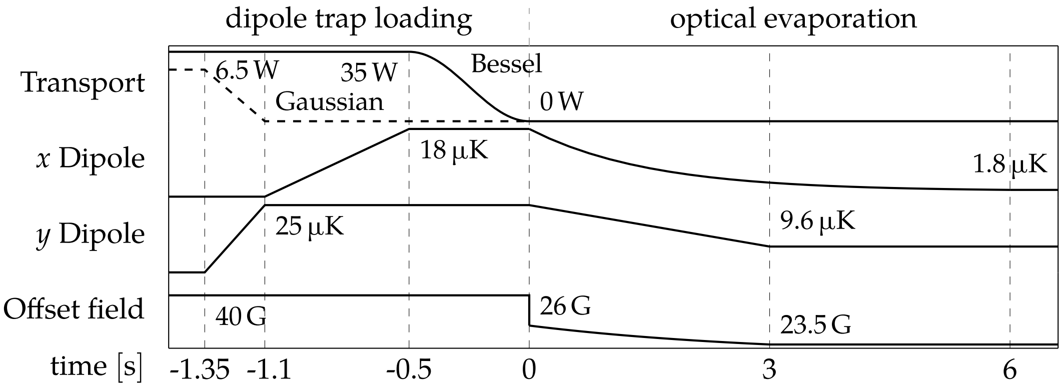

Figure 7.3.: Sequence diagram of the crossed dipole trap loading, followed by the optical evaporation to create a BEC.

for 500 ms before continuing with the next steps in the sequence. We create BECs with typical atom numbers $n \approx 2 \times 10^4$ atoms using this sequence.

7.3.2. Single plane loading

Before we prepare our experimental system we increase the vertical confinement to create a quasi two dimensional system, suitable for *in situ* imaging using our objectives. To reach this regime the trapping frequencies along the vertical direction must be much larger than any energies set by the temperature ($\hbar\omega_z \gg k_B T$) and chemical potential ($\hbar\omega_z \gg \mu$). To create the required strong confinement we are loading the atomic cloud in a single plane of the shallow VL.

Before we begin the loading procedure we reduce the magnetic field further to the three-body loss minimum at 20.25 G in 100 ms. The exact value is experimentally determined, maximising the lifetime of the compressed cloud in the shallow VL. We then increase the depth of the y dipole trap to 14.5 µK in 300 ms to compress the cloud vertically and subsequently load >95 % of the atoms into a single plane by ramping up the shallow VL to 2.7 µK in 400 ms.

We transfer the atoms from the crossed dipole trap into the vertical dipole trap by ramping up the z dipole to 1.3 µK and simultaneously turning off the y dipole trap in 225 ms. 100 ms before the end of these ramps we start the 520 ms ramp down of the x dipole trap. The times and overlap of these ramps were optimised experimentally to reduce the total duration of the sequence, the full loading procedure is performed in 1.45 s.

7.3.3. Fluorescence imaging

To acquire images using fluorescence imaging we pin the atoms in a deep 3d lattice, created from the horizontal 767 nm lattices and the shallow VL. At the end of the experiment stage we start the imaging by driving the intensity regulations into their limits, increasing

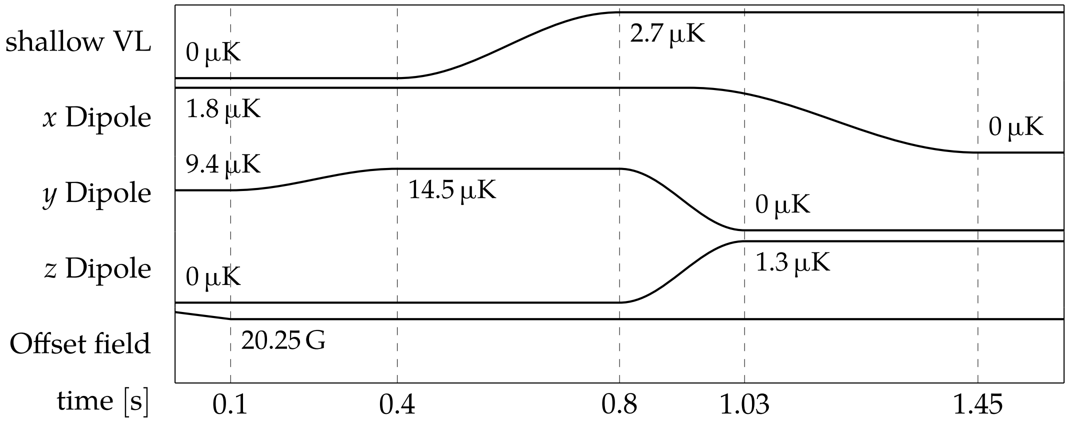

FIGURE 7.4.: Sequence diagram of the single plane loading from the crossed dipole trap into the shallow VL. The cloud is compressed along the vertical direction using the y dipole, to load into a single plane of the shallow VL. Afterwards the atoms are transferred from the crossed dipole trap into the z dipole trap. The magnetic field is reduced to minimise atom loss due to three-body collisions.

the intensity of the lattice beams to their maximum value within a few microseconds. The resulting trapping potentials have a depth of approximately $k_B \times 400\,\mu\mathrm{K}$, or $\sim 4000\,\mathrm{E}_{\mathrm{r},852}$, varying by up to 20 % for longer exposure times as the intensity is no longer stabilised. At the same time we turn off all magnetic fields and compensate residual field using the experiments x and y coils, and the large additional compensation coil along the z direction.

After 10 ms wait time to ensure all fields are stable we turn on the optical molasses beams and repumper in 2 ms to avoid intensity overshoots. We wait for 30 ms before beginning to collect photons to ensure light-assisted collisions have removed all doublons from the lattice. Subsequently we acquire one or more images, using typical imaging times of 200 ms or 400 ms. We acquire up to 10 images of the same atomic cloud, e.g. when scanning the objective distance to optimise the focus.

7.4. Mott-insulator states

The proposal [166] and ultimately the successful realisation of the quantum phase transition from the superfluid to the Mott-insulator (MI) phase by Greiner et al. [167] in 2002 can be seen as the starting point for the field of ultracold quantum gases in optical lattices.

The MI state is an incompressible phase in the phase diagram of the Bose-Hubbard Hamiltonian

$$\hat{H} = -J \sum_{\langle i,j \rangle} \left(\hat{a}_i^{\dagger} \hat{a}_j + \mathrm{h.\,c.} \right) + \frac{1}{2} U \sum_i \hat{n}_i (\hat{n}_i - 1), \tag{7.2}$$

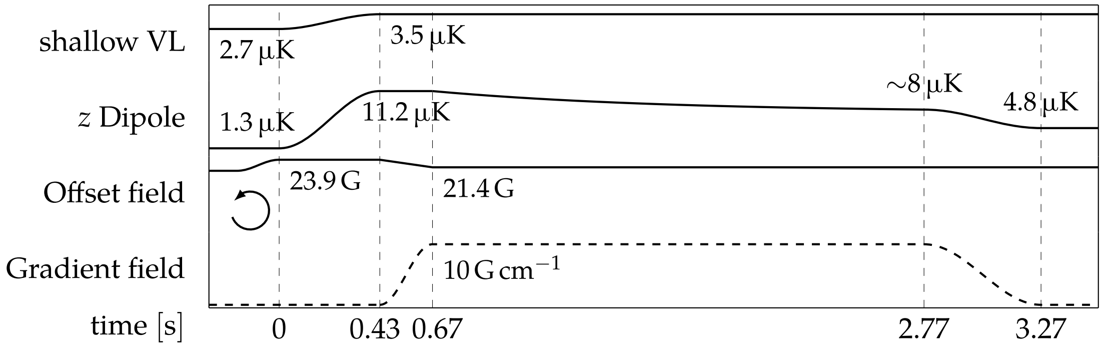

FIGURE 7.5.: Sequence diagram of the radial evaporation in the shallow VL. To prepare the final system for experiments additional evaporation in the 2d system is performed. The quantisation axis is rotated in the horizontal plane and the trapping potentials are increased. A magnetic gradient tilts the trapping potential along the quantisation axis, and atoms are subsequently evaporated along this axis. At the end of the evaporation the gradient is removed and the trapping potentials are reduced for lattice experiments.

where U is on-site matrix element and J the tunnelling matrix element [168]. As the Hamiltonian favours tunnelling the 'standard' phase of this system is the superfluid (SF) phase. However at integer filling the system can transition to the MI phase. When the relevant ratio U/J is large enough tunnelling is suppressed due to the large energy required to create a higher occupied site and an insulator is created.

By now the creation of MI states has become a standard procedure for bosonic ultracold quantum experiments. Due to their homogeneous density distribution they are a useful tool for the controlled preparation of quantum states.

7.4.1. Experimental sequence

To prepare cold MIs we perform an additional evaporation step once the atoms are loaded into the single plane. We control the value of the chemical potential μ in this stage by varying the final depth of the trapping potential. Subsequently we prepare the MI state by ramping up the optical lattices.

Radial evaporation

We perform the second evaporation stage using radial evaporation. A magnetic field gradient tilts the trapping potential, lowering it in one direction. By lowering the horizontal trapping potential atoms are evaporated along this axis.

Before we can perform the evaporation we continuously rotate the quantisation axis from the vertical axis to the y' axis in 80 ms to allow usage of a gradient in the horizontal direction. During this procedure we increase the magnetic field to 23.9 G. We then first increase the trapping potentials in 430 ms by ramping the vertical dipole trap to 11.2 µK and the shallow VL to 3.5 µK. We then simultaneously ramp the magnetic field to 21.4 G

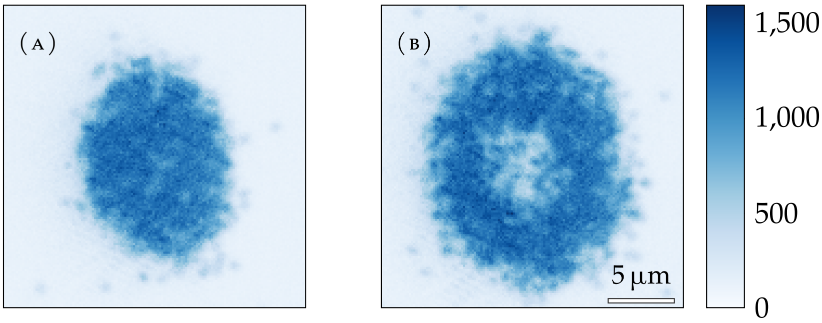

FIGURE 7.6.: Unprocessed *in situ* images of Mott insulators with filling $n = 1$ (a) and $n = 2$ (b), created in a 767 nm/2 square lattice. The $n = 1$ ($n = 2$) MI has approx. 800 (1500) atoms and a temperature $k_B T/U = 0.13$ (0.18).

and tilt the potential by ramping up a gradient of $10\,\mathrm{G\,cm^{-1}}$ in 225 ms. To perform the evaporation we exponentially decrease the power of the vertical dipole trap in 2.1 s to values around 8 µK. The final value of this ramp is varied to control the chemical potential of the prepared system. Lastly we remove the gradient and lower the vertical dipole trap to 4.8 µK in 500 ms.

Due to the additional horizontal confinement of the shallow VL the trapping frequencies in the combined potential are slightly asymmetric in the horizontal plane $\omega_x = 2\pi\times55.8\,\mathrm{Hz}$, $\omega_y = 2\pi\times49\,\mathrm{Hz}$, and $\omega_z = 2\pi\times1260\,\mathrm{Hz}$.

Mott-insulator transition

To cross the MI transition we simultaneously ramp up the optical lattices and the magnetic field, loading in the lattice and increasing the scattering length in parallel. In the first ramp we increase the lattices to $\omega_x = \omega_y = 2\pi\times20\,\mathrm{kHz}$ and change the magnetic field to 29.8 G in 50 ms, ending at tunnelling rates $J \approx h \times 33\,\mathrm{Hz}$. These ramps end at $U/J = 26$, crossing the MI transition. The ramp duration is set to approximately $2\,\mathrm{s}/hJ$, to ensure the atoms have enough time to redistribute into the density distribution of the MI state. The magnetic field ramps increases the scattering length a from $\sim$200 a_0 to 530 a_0, proportionally increasing the on-site interaction energy $U \approx h \times 1\,\mathrm{kHz}$.

To reach the final state we increase the lattice intensities further to $\omega_x = \omega_y \approx 2\pi\times 33\,\mathrm{kHz}$, freezing out tunnelling by reducing the tunnelling rate $J \approx h \times 0.7\,\mathrm{Hz}$, while increasing the on-site interaction energy further to $U \approx h \times 1.5\,\mathrm{kHz}$. Figure 7.6 shows exemplary MIs with $n = 1$ and $n = 2$.

The generated MI states are primarily used for state preparation, utilising the homogeneous density distribution of the $n = 1$ insulator shell. We optimise the MI state by varying the final power of the radial evaporation described on the facing page, adjusting the final atom number and with it the chemical potential μ.

To prepare systems with different values for the relevant parameter U/J we vary the final intensity value of the second lattice ramp, primarily changing J.

We can directly create MI states with $U = h \times \sim 2\,\mathrm{kHz}$ by increasing the vertical

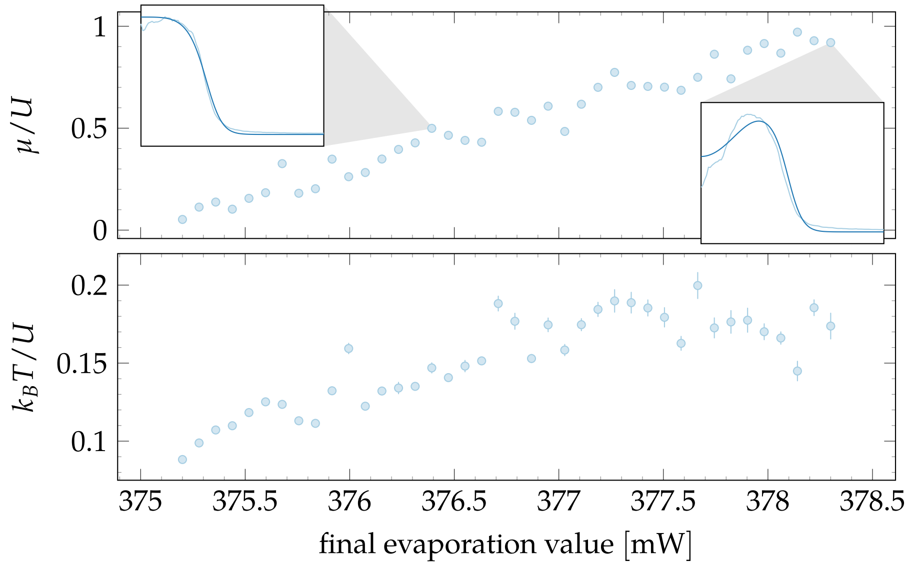

FIGURE 7.7.: Chemical potential μ and temperature T in units of the on-site interaction energy U measured for an evaporation scan. The final evaporation value determines the trap depth at the end of the radial evaporation (section 7.4.1). The insets show exemplary radial profiles and fits for the $n = 1$ and $n = 2$ parameters.

confinement to $2\pi \times 3$ kHz using the shallow vertical lattice and increasing the scattering length further up to ∼650 a_0. The resulting system is elliptical with a 2:3 aspect ratio due to the increased horizontal confinement from the shallow vertical lattice beams.

7.4.2. Thermometry

To determine the temperatures of our MIs we perform *in situ* thermometry, following the method described in chapter 5 of ref. [169]. The temperature is determined by fitting the expected densities $\bar{n}(r)$ on the measured densities of the radially averaged images of the MI states, extracting the information from the 'sharpness' of the edges of the MI.

Considering the local Bose-Hubbard-Hamiltonian in the atomic limit, including an external harmonic confinement

$$\hat{H}(n,r) = \left(\mu - \frac{1}{2}m\omega^2 r^2\right) n - \frac{U}{2} n\,(n-1) \tag{7.3}$$

$$= \mu_{loc}(r) n - E_n, \tag{7.4}$$

one can derive the expected densities at different radii $\bar{n}(r)$ as

$$\bar{n}_{det}(r) = \frac{1}{Z(r)} \sum_n \mathrm{mod}_2(n) e^{(\beta(\mu_{loc}(r)n - E_n))}, \tag{7.5}$$

with the grand canonical partition function $Z(r) = \sum_n \exp(\beta \hat{H}(n,r))$ and $\beta = 1/(k_B T)$. The modulo operation is due to the parity projection process during fluorescence imaging (see section 7.3.3).

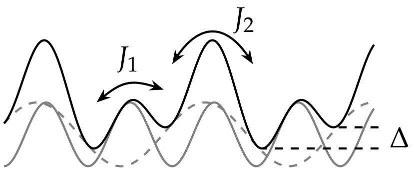

FIGURE 7.8.: Drawing of a super-lattice potential created from an optical lattice overlapped with a secondary optical lattice with twice the lattice spacing. The relative phase between the lattices introduces an energy shift Δ between the two sites in the unit-cell. The tunnelling is characterised by the tunnelling matrix elements J_1 and J_2.

To determine the temperature and chemical potential in the system the acquired images of the MIs are radially averaged and then fitted in relation to U, extracting k_BT/U, μ/U and a radius $r_0 = \sqrt{2U/(m\omega^2)}$ for the external confinement.

We cannot follow this procedure directly, since we cannot directly access the measured densities in the images as required due to limitations in the reconstruction fidelity (section 8.2). Instead we radially average and fit the acquired images, subtracting a background and rescaling the signal amplitude by a factor extracted from multiple averaged $n = 1$ MI shells. Using these measurements we determine the temperature of the generated $n = 1$ MIs to $k_BT/U \approx 0.1$, limited by the uncertainty introduced from the rescaling. Better temperature measurements will be possible in the near future once the reconstruction algorithms are fully integrated.

To obtain good fit results for low values of μ/U it is necessary to reduce the number of fit parameters by fixing r_0. It can be determined from images at higher chemical potential for which the fit converges. We fit each evaporation scan twice, once to determine r_0 and a second time with fixed r_0 to determine μ/U and k_BT/U.

7.5. Optical super-lattices

We recently upgraded the installed lattices to so-called super-lattices by overlapping beams with twice the wavelength. This configuration extends the size of the lattices unit cell and number of possible lattice configurations. By varying the relative phase φ between the short and long spaced lattice beams and their intensities the energy difference Δ between the neighbouring sites and the two different hopping rates J_1 and J_2 can be controlled independently (fig. 7.8).

In the tight-binding limit this system is described as an extended version of the Bose-Hubbard Hamiltonian called the Rice-Mele-model [170]

$$\hat{H}(\varphi) = -\sum_m \left(J_1(\varphi)\hat{b}_m^\dagger \hat{a}_m + J_2(\varphi)\hat{a}_{m+1}^\dagger \hat{b}_m + \text{h. c.} \right) \tag{7.6}$$

$$+ \frac{\Delta(\varphi)}{2} \sum_m \left(\hat{a}_m^\dagger \hat{a}_m - \hat{b}_m^\dagger \hat{b}_m \right), \tag{7.7}$$

where $\hat{a}_m^\dagger$ ($\hat{a}_m$) and $\hat{b}_m^\dagger$ ($\hat{b}_m$) are the creation (annihilation) operators on the even and

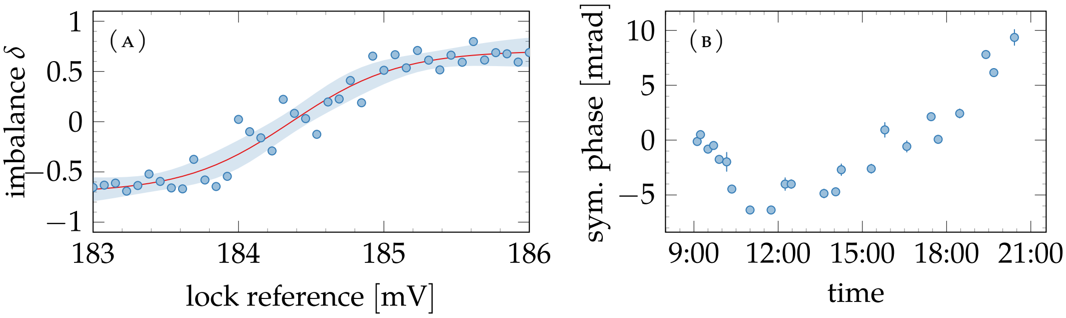

FIGURE 7.9.: (a) Calibration measurement for the relative phase between short-spaced and long-spaced lattices of a super-lattice. For each data point the imbalance between evend and odd rows was measured. The displayed lock reference voltage is used as input to the frequency lock. (b) Stability measurement of the super-lattice phase. Each data point is the symmetric phase extracted from a measurement as shown in (a).

odd sites of the mth unit cell, respectively. Without detuning ($\Delta = 0$) it is known as the SSH model [10].

All parameters of the model can be tuned by adjusting laser intensities and frequencies, providing a flexible tool for state preparation and readout. Dynamical control can be used to implement interesting systems, e.g. topological charge pumps [171].

7.5.1. Super-lattice phase

To calibrate the relative phase within a super-lattice axis we prepare a sample in the MI state in the long wavelength lattice of the measurement axis and the short lattice in the perpendicular axis. We then adiabatically split the system along the measurement axis by ramping up the short wavelength lattice in 300 ms, before measuring the local population using fluorescence imaging. From the local atom distribution acquired from the reconstructed images we compute the imbalance between sites on even and odd rows $\delta = (n_{\text{even}} - n_{\text{odd}})/(n_{\text{even}} + n_{\text{odd}})$, averaging along the long spaced axis. The phase for a symmetric double well configuration is extracted by fitting a sigmoid ramp to the extracted imbalance versus lock reference. The symmetry of the measurement around the double well configuration allows for a very precise calibration of the symmetric phase with fit errors below 0.1 %, even in presence of the stronger shot-to-shot fluctuations.

By repeating this measurement around the next symmetric phase we calibrate the scaling of lock reference voltage to the generated frequency shift and relative phase. The measured frequencies for a π phase shift are 753 MHz for the lattice along the x'' axis, and 637 MHz for the lattice along the y'' axis. From these values the precise distance of the retro-mirrors to the atoms can be determined to be 200 mm and 235 mm respectively.

The lattice phase is determined not only by the relative laser frequency, but also by the refractive index of the medium. Temperature and humidity variations change the refractive index of the air and translate to shift of the lattice phase. To characterise the phase stability of our lattice setup we repeated the phase calibration measurements

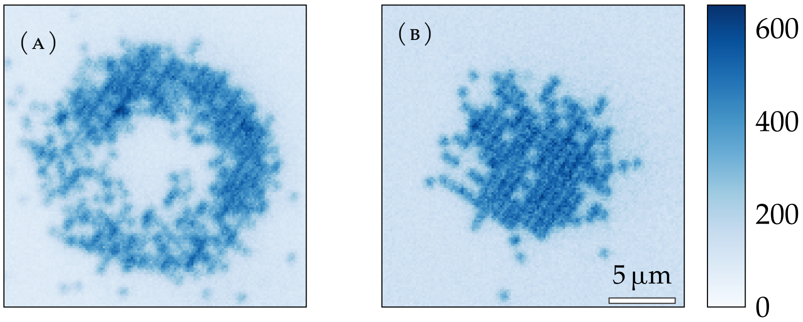

Figure 7.10.: (a) MI produced in a rectangular lattice using a combination of 767 nm and 1534 nm beams. Using the regular MI sequence we load about 1500 atoms into a system with reduced site density, creating an $n = 2$ MI and an overall larger system. (b) MI produced in a long spaced square lattice made from 1534 nm light. The evaporation parameters were adjusted to load fewer atoms to create an $n = 1$ MI.

regularly over a full day, measuring the shift of the symmetric phase (fig. 7.9b). We measure absolute variations of the phase of 15 mrad with a maximal drift of 2 mrad/h.

7.5.2. Mott-insulators

An application of the longer spaced lattice is the creation of MI states, to prepare systems with reduced mean density compared to the regular lattices. An $n = 1$ MI created using the longer spaced lattice in one or both axes has a reduced mean density of 1/2 or 1/4, when it is released in the short spaced lattice. The preparation of systems with fractional mean density is a crucial ingredients for future experiments with topological systems with fractional charge [172–174].

To prepare MIs in the super-lattice we use the standard sequence for the generation of MI states described on page 101. and ramp up the super-lattices in parallel to the standard lattices. We choose the relative phase between the lattices such that a strongly tilted potential is created and only a single site of each double well is populated. As we use both lattices in parallel the lattice parameters are comparable to the short spaced system, and we can use the sequence without further modifications.

By selecting the beams to use we can create MI states in a rectangular lattice created by short and long wavelength beams (fig. 7.10a), and in an enlarged square lattice using the super-lattice along both axes (fig. 7.10b). The created systems are larger and have a larger chemical potential for the identical evaporation parameters, as the number of atoms is unchanged, but the density of available lattice sites is decreased (fig. 7.10a). The horizontal confinement is identical in both cases as we do not adjust the trapping parameters. Because of this the system size of the pure $n = 1$ MI is approximately the same for all configurations (fig. 7.10b).

8. Quantum gas microscope

From the beginning the inclusion of a quantum gas microscope was an integral part of the design of the newly constructed experiment. The included high-resolution objectives allow the direct observation of the atomic cloud, extending the number of available measurement parameters, and can be used for local control.

The high resolution allows the direct imaging of individual atoms in the optical lattice by illuminating them with resonant light, collecting the scattered photons on a camera sensor. The general simplicity of the method comes with a number of technical challenges to capture an image from which the occupation of the lattice sites can be detected.

The imaging resolution of the system is determined by the imaging wavelength λ and the numerical aperture (NA) of the used objectives. A typical resolution criterion is the Rayleigh criterion, specifying the radius of the first minimum of the observed Airy disk $r_A = 0.61\lambda/\mathrm{NA}$. Typically the lattice spacing is smaller than the imaging wavelength, requiring the usage of high NA objectives. Especially in Caesium high NA objectives are necessary, due to the large wavelength of the standard imaging transition of 852 nm, combined with the high mass of caesium, favouring short spaced lattices to achieve fast dynamics. The objectives included in our experiment have a NA of 0.8, and are designed to be used for imaging and local potential shaping. Their mechanical integration is described in section 4.2. In the following the associated imaging system is described and characterised.

In almost all cases the achievable resolution is too low to directly observe individual atoms in the optical lattice. Instead the lattice population must be reconstructed from the acquired images, using the additional knowledge of the possible locations of the atoms in the periodic lattice structures. Due to the short spacing of our lattice and the relatively large imaging wavelength of Caesium we are using a machine learning (ML) based approached for the lattice reconstruction described in section 8.2.

To achieve a sufficient signal-to-noise ratio (SNR) for the reconstruction process we collect about 10^5 photons per atom, within 400 ms. The repeated scattering of photons increases the temperature of the trapped atoms, requiring continuous cooling to avoid detection errors when atoms are lost or tunnel to neighbouring sites. Generally all optical cooling methods can be used and quantum gas microscopes using optical molasses cooling [29, 111], Raman sideband cooling [106, 112], and EIT cooling [175, 176] have been constructed. For our experiment we chose to use an optical molasses, described in the following in section 8.1.2.

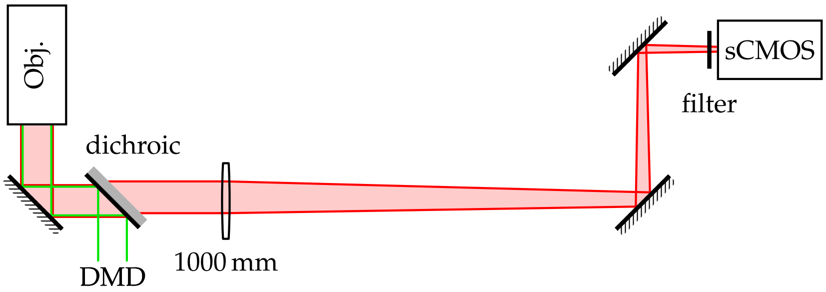

FIGURE 8.1.: Simplified schematic of the imaging beam path. The first mirror behind the objective redirects the beam from the vertical axis into the horizontal plane. A dichroic mirror transmitting 456 nm and 852 nm is installed to project repulsive potentials on the atoms. A single 1000 mm tube lens focuses the image on a sCMOS camera. A 852 nm bandpass filter mounted to the camera filters stray light.

8.1. Imaging system

The experiment and objectives were designed for symmetric installation of two objectives on top and bottom of the glass cell. At the time of writing the objective on the bottom of the atoms is installed and aligned to the atoms.

The attached imaging system was designed to reduce the number of optical components, which could potentially degrade the imaging quality. When the objective is correctly aligned a collimated beam with 42 mm diameter beam is generated. It is redirected into the horizontal plane using a rectangular mirror, before it is focussed using a single 1000 mm achromatic tube lens[1], mounted in a 4 axis stage. Two more mirrors far from the tube lens are used to position the beam on the sensor of a sCMOS camera[2]. A bandpass filter[3] is mounted directly on the camera to block stray light. The distance between the back of the objective and the tube lens is approximately 300 mm, resulting in 1.3 m total length of the imaging path.

All components within the beam path have dielectric coatings optimised for the used wavelengths. Based on the specifications of the components we expect a total efficiency of 98 %. The used camera has a quantum efficiency of 58 % at 852 nm. Our objectives with NA 0.8 cover a solid angle of $2\pi \times 0.22$. Considering all losses we expect to detect ~6.2 % of all emitted fluorescence photons using a single objective.

The objectives have an effective focal length of 25 mm, therefore the imaging system has a magnification of 40, corresponding to ~160 nm/px and 2.36 px/site. During lattice reconstruction (section 8.2.1) we derive the exact magnification of 40.41 and a pixel size 160.9 nm/px from the known lattice spacing.

The chosen magnification is derived from the Shannon-Nyquist sampling theorem and the available focal lengths for the tube lens. To maximise the SNR a small magnification should be chosen to reduce the number of pixels per lattice site and with it the amount

---

[1] AC508-1000-B, Thorlabs

[2] Kinetix 22, Teledyne Photometrics

[3] HC Laser Clean-up MaxLine 852/3,2, Semrock

of background and readout noise. The smallest possible magnification is determined by the requirement to resolve the underlying lattice structure for reconstruction, requiring at least two pixels per lattice site according to the Shannon-Nyquist sampling theorem. In our system 192 nm/px are required, corresponding to a minimal magnification of 33.9. To include some tolerance and avoid mapping the lattice peaks to the edge of the Fourier spectrum during lattice reconstruction a slightly higher magnification should be chosen. Considering only commonly available focal lengths the chosen magnification of 40 with a $f = 1000\,\mathrm{mm}$ tube lens is the smallest possible magnification.

8.1.1. Point spread function

We characterise the quality of our imaging system by measuring the point spread function (PSF), describing the measured signal for an ideal point source. In the ideal case when the system is diffraction limited it is a symmetric Airy disk. The first minimum of the Airy disk is commonly used to measure the size.

To measure the PSF of our system we acquire images of a dilute atomic cloud. In these images we detect isolated atoms using a peak finder algorithm, selecting only atoms with a sufficiently large distance to the next atom. The precise centre location is determined using a fit with an elliptical 2d Gaussian. The detected signals are upsampled without interpolation and overlapped and averaged to determine the average PSF (fig. 8.2a,b). The upsampling allows precise alignment not limited by the pixel size of the camera. To visually highlight deviations from the expected symmetric Airy disk we fit the signal using a 2d Gaussian and compute the difference (fig. 8.2c,d).

To quantify the size of the PSF we fit each detected atom individually and generate a histogram of the measured peak widths per axis. The average PSF size is given as the centre position of a normal distribution fitted to the histogram.

After a long period of frequent re-alignment and characterisation of the objective in the attempt to achieve a PSF size below 900 nm we had to find out that our objective had been damaged during shipping. The procedure to fix the lenses inside the housing used by the manufacturer at the time could fail, allowing lenses to shift off-axis. We were made aware of this potential problem by discussions with the manufacturer, which were later confirmed when we sent the objective back for repair. In hindsight it is also possible to observe the damage in the shape of our PSFs. They always were elliptical and showed pronounced rings in one direction, which we were never able to fully compensate (fig. 8.2a). Most likely we optimised the alignment of the objective to a high tilt angle, to reduce the artefacts introduced by shifted lenses, reducing the effective NA at the same time.

Unfortunately were are not able to achieve the expected resolution, even with the repaired objective. The optimised PSF is still approximately 850 nm, about 25 % larger than expected (fig. 8.2b). This indicates that we are now limited by at least one other components of the imaging system. At the time of writing it is not clear where in the imaging system the resolution is lost. Several attempts to improve the resolution by rebuilding parts of the imaging system were not successful. At optimal alignment the shape of the PSF when scanning through the focus does not show any clear signs of

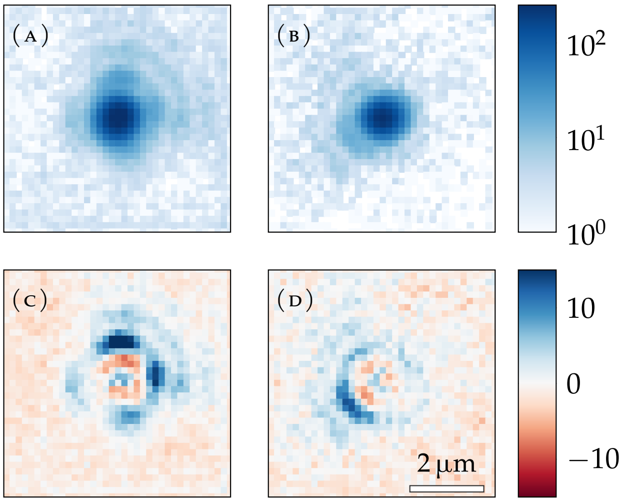

Figure 8.2.: Averaged experimental PSFs before (a) and after (b) the repair of the objective and difference from a fit of a 2d Gaussian to the PSF (c,d). The upper images show the logarithm of the signal to highlight the ring structure. All images were taken using the same camera and use the same pixel scaling. The measured PSF width along both axes is 875 nm/930 nm before the repair and 820 nm/900 nm after the repair.

specific aberrations.

A possibility is limited confinement in the vertical direction, causing the atoms to extend too far outside the focus plane. A first test performed by increasing the intensity of the shallow VL used for vertical confinement showed an improvement of the PSF size by approximately 50 nm along one axis. Further increases are not feasible as the additional horizontal confinement from the shallow VL already results in elliptical systems. To increase the confinement further a different trapping setup is needed, e.g. usage of the steep angle vertical lattice. This was not fully installed and aligned yet due to the challenging alignment and time constraints.

8.1.2. Optical molasses

For continuous cooling during imaging we use an optical molasses on the D2 line, simultaneously generating the fluorescence photons. We decided to use an optical molasses instead of Raman sideband cooling for the simpler setup, the achievable low temperatures below 10 µK and the availability of the required cooling light. We use part of the light also used for the optical cooling in the MOT and split it into three using fibre beam splitters. One path is in the horizontal plane along the x' axis, two others enter the chamber under 30° angle along the y' axis (fig. 8.3). The intensity is stabilised using a photodiode in one of the arms for detection and a common AOM in front of the fibre splitters for feedback.

All beams are retro-reflected, using $\lambda/4$ waveplates in front of the retro mirrors to maintain the correct circular polarisation. The mirrors are mounted on ring piezos,

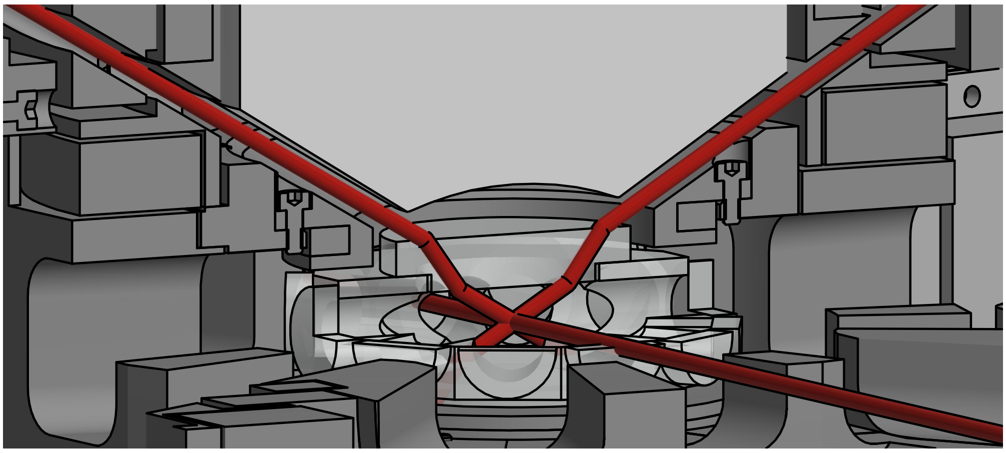

FIGURE 8.3.: Orientation of the molasses beams passing through the experiment chamber. View on the chamber along the transport axis. The horizontal beam passes through the chamber along the x' axis. The angled beams in the y'/z plane enter the chamber under $\pm 60°$ from the top and are retro-reflected on the bottom.

which are periodically driven at 31 Hz, 51 Hz, and 71 Hz to wash out the interference pattern between the beams [169]. Without this modulation we observe strong local variations in the number of photons collected per atom. We use about 3 mW per axis, creating relatively large beams with 2 mm diameter. We aligned them using short pulses of individual beams, detecting the atom loss due to heating from the coherent coupling on a BEC. We then aligned the retro-mirrors to reflect the beams back into the fibre coupler. The polarisation was aligned using a self-built polarisation tester, made from a $\lambda/4$ waveplate and a PBS.

8.1.3. Imaging parameters

The imaging duration is the primary tuning parameter for fluorescence imaging. Longer exposure times reduce the influence of background noise and increase the SNR. They also increase the probability of imaging errors due to atom loss or tunnelling events. To optimise our imaging parameters we optimise the efficiency of the optical molasses to reduce loss and tunnelling events. The imaging time is then chosen as a compromise between losses and a high SNR to achieve a good reconstruction fidelity.

Molasses efficiency

To measure the cooling efficiency of the optical molasses we acquire multiple images of the same atomic cloud, and compare their differences. We illuminate the cloud continuously for approximately 1.7 s, acquiring up to four images, each with 400 ms acquisition time.

The molasses parameter we optimise are the powers and detuning of the involved cooling and repumper beams, and the magnetic field along all three axes. The magnetic

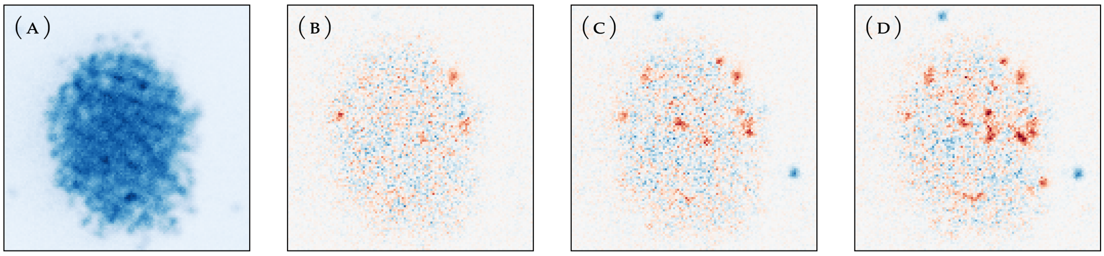

FIGURE 8.4.: Multiple images acquired from the same sample, used to measure tunnelling and loss rates. (a) First acquired image using 400 ms of exposure time. (b-d) Difference between the first acquired image and subsequent images taken after 0.4 s (b), 0.8 s (c), and 1.2 s (d). Red colour indicate a decrease in signal and atom loss, blue indicates an increase in signal from tunnelling events.

field must be fully compensated to achieve optimal cooling efficiency. During optimising we visually analyse the difference between the images (fig. 8.4).

From these images tunnelling events can not be reliably detected, as atoms with enough energy to tunnel can tunnel almost freely, before they get cooled and captured again. If they are cooled into an already occupied site both atoms will be lost due to parity projection, resulting two emptied sites. The lack of correlation between both sites prevents detection and correction for these events. Instead we compute the tunnelling rate by counting the number of newly occupied sites n_{+1} as n_{+1}/n, where n is the number of detected atoms in the first image. From the number of emptied sites n_{-1} the loss rate is computed as $(n_{-1} - n_{+1})/n$. To minimise the influence of reconstruction errors we use dilute clouds, as the mostly isolated atoms in these images are easy to reconstruct. The latest quantitative analysis of the rates measured tunnelling and loss rates of $0.062\,\mathrm{s}^{-1}$ and $0.064\,\mathrm{s}^{-1}$ respectively. Since then optimisations in the molasses parameters reduced this rate qualitatively (see fig. 8.4). However quantitative evaluation was not possible due to the ongoing work on the reconstruction algorithms.

Signal-to-noise ratio

The fidelity of the lattice reconstruction is heavily influenced by the signal-to-noise ratio (SNR). For our application we define it as the separation in the signal of an empty (background; bg) and an occupied (occ) site over the sum of the respective peak widths

$$\mathrm{SNR} = \frac{\bar{n}_{occ} - \bar{n}_{bg}}{\sigma_{occ} + \sigma_{bg}}. \tag{8.1}$$

We include the standard deviation of both peaks, to capture the different physical processes involved. The background is given by the camera dark noise, whereas the signal of occupied sites additionally includes intensity fluctuations, e.g. due to inhomogeneous illumination.

To measure the SNR we acquire images of dilute clouds and compute the histogram over the signal of isolated atoms and randomly chosen empty sites (fig. 8.5a). We observe

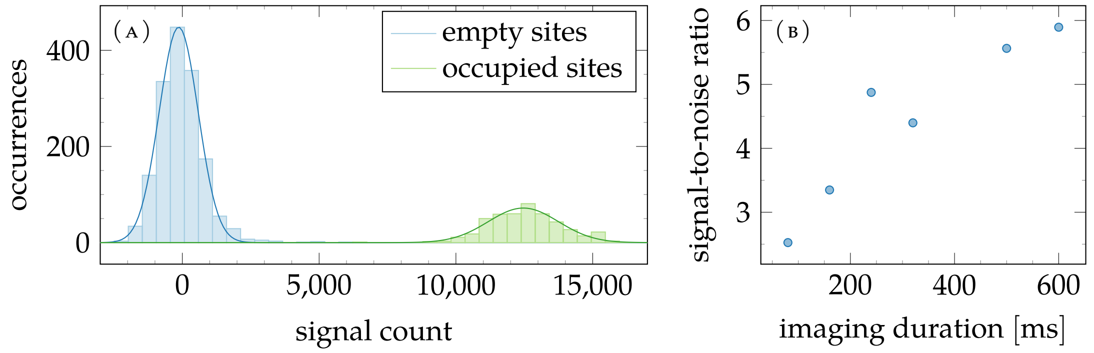

Figure 8.5.: (a) Histogram of the signal count distribution for 400 ms imaging duration to determine SNR. Occupied sites are isolated single atoms from dilute images. Empty sites are randomly sampled areas without atoms from the same images. The computed SNR is ∼6.3. (b) Measured scaling of the SNR with imaging duration. The data was acquired earlier than the histogram shown on the left with slightly stronger intensity fluctuations due to less optimal cooling parameters.

an almost linear increase in the SNR for imaging durations up to 600 ms (fig. 8.5b). Based on the measured SNR and tunnelling and loss rates we are currently using a standard imaging duration of 400 ms. With these parameters measure 12 500(1250) counts per atom, corresponding to 3000 photons per atom and typical SNR $\approx$ 6.3.

8.2. Reconstruction

Reconstruction algorithms are used to extract the *in situ* distribution of the atoms in the lattice from the acquired images. For experiments with a small PSF size r_A compared to the lattice spacing a, where neighbouring atoms can be distinguished by eye, a deconvolution using the measured PSF can be enough to reliably determine the lattice occupation from the photon count on the individual sites [112].

The large imaging wavelength of Caesium of 852 nm combined with the high mass, preventing the use of long lattice spacings, forces us to apply more elaborate reconstruction methods that rely on the additional information from the lattice structure. Various reconstruction algorithms have been implemented throughout the years [106, 111, 115], typically operating in the regime $a < r_A < 1.5a$.

Due to problems with our imaging system described section 8.1.1 our resolution is currently worse than anticipated during design. Optimally the resolution of our objective would be $r_A = 650\,\text{nm} \approx 1.7a$. The currently achieved resolution is significantly worse with $r_A \approx 850\,\text{nm} \approx 2.2a$. Initial tests using the Richardson-Lucy algorithm [115, 177, 178] showed fidelities >90 % on simulated data [59], but were less effective on experimental images. To achieve higher fidelities we implemented a reconstruction algorithm using machine learning (ML) for detection. At the time of writing work on the reconstruction is still ongoing, but so far we already achieve overall detection fidelities of 95 % on experimental images, comparable to other quantum gas experiments. Most of the work was done as a masters project and will be described in detail in ref. [179].

During development of the method various network architectures were implemented and tested. Initially we implemented binary classifications algorithms where a neural network (NN) was trained to detect whether the central site in a cut-out was occupied (section 8.2.3). We later switched to an autoencoder, which can be trained in an unsupervised fashion (section 8.2.4). We also have promising results using a projector based deconvolution approach, where we used a NN to optimise the projector for our system (section 8.2.5).

8.2.1. Lattice reconstruction

Before we can start any reconstruction process we need to determine the precise lattice orientation, spacing and global phase from the acquired images. We use images of dilute atomic clouds, in which we detect individual atoms using a peak finder and measure their position by fitting 2d Gaussian functions. The extracted positions are processed using a 2 dimensional FFT, translating the positions into k space. The lattice vectors create sharp amplitude peaks, from which we extract the position, again fitting 2d Gaussian functions. To help the fit we increase the image size by padding with zeros to the next power of 2, increasing the sharpness of the peak in Fourier space.

From the fitted positions we derive the lattice angles and calibrate the pixel size of our imaging system using the distance with the known lattice spacing $a = 767\,\mathrm{nm}/2$. These parameters do not change frequently and are only calibrated after significant work on the optical table. To perform the calibration we acquire typically around 50 images of very dilute clouds, which we also use to determine the PSF.

The lattice phase is detected for every image, as it changes directly when the relative position of the atoms to the objective changes, i.e. through thermal fluctuations. We use typically 5 to 10 isolated atoms at the edge of the atomic cloud for this calibration. We extract their position using Gaussian fits and determine the phase of the lattices by fitting every atom individually. The total lattice phase is determined by averaging to 5 % of the lattice vectors.

8.2.2. Simulated training data

To train the neural networks many images with known atom distribution are needed. Experimental images can not be used, as we do not know the underlying distribution, rendering them unsuitable for supervised learning and verification. Instead we use simulated images, which we can generate quickly, with known distribution. The generated images must be as close as possible to the real experimental images to ensure a good fidelity of the trained network when it is applied to experimental data.

Starting point for these images is the experimental PSF, generated by averaging around 200 individual PSFs from isolated atoms, randomly distributed over the measurement area. We found that it is necessary to use a large PSF matrix to include signal from the outer wings. Currently we are using a minimum of $28\,\mathrm{px} \times 28\,\mathrm{px}$ ($\sim 5.4\,\mu\mathrm{m}$), or about 6 times the PSF diameter. For smaller sizes the background in high-densities regions was

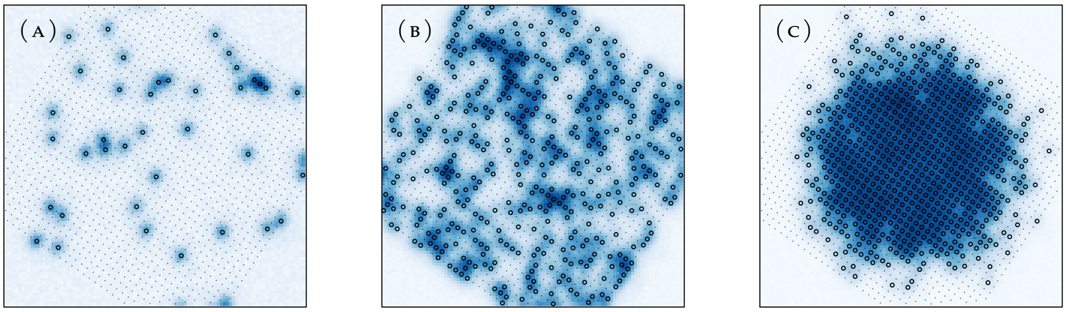

FIGURE 8.6.: Simulated images with a uniform atom distribution and mean filling $\bar{n} = 3\,\%$ (a), $\bar{n} = 30\,\%$ (b) and a simulated MI with a radius of 15 sites and $k_B T/U = 0.1$ (c). The overlaid points and circles denote empty and occupied sites respectively.

significantly weaker when compared to experimental images, resulting in poor detection fidelities when applying the trained network to experimental images. Figure 8.2 shows our currently used PSF and the deviations to a fitted 2d Gaussian distribution.

To generate the image we start with the desired distribution, typically a uniform distribution of atoms for a given mean particle number. To capture shot noise we randomise the number of photons per atom, sampling the exact number from a Poisson distribution with a mean determined by our experimental images. For each 'detected' photon we determine a continuous location by sampling from the extracted PSF distribution, with the centre of the distribution shifted to the atom position. This allows us to include the experimentally measured lattice vectors in the simulation. To capture additional noise in our imaging system we add more counts by sampling from a Gaussian, centred on the atom position, with parameters fitted to experimental images. We collect a list of all sampled points and bin them into sectors on a grid, with sizes matching the measured pixel size of our camera and imaging system, simulating the discrete sampling of the camera sensor. Lastly we add normally distributed random noise to every bin to account for electronic camera noise.

8.2.3. Binary classification networks

To reconstruct the lattice occupation from our images we started with binary classification. Individual cut-outs around each lattice site are generated and classification algorithms are used to determine the occupation of the central site within each snippet. Local optimisation methods have been implemented to perform this classification [180], but were tested to be not usable with our resolution. Instead we use NNs to perform the classification. Two standard network architectures were implemented. A fully connected NN with four hidden layers, and a convolutional neural network (CNN). For a general introduction to deep learning and neural networks see e.g. ref. [181].

The NNs are trained to derive a single binary output value, the occupation of the central site, from the provided snippet. The networks are trained via backpropagation using a training data set of simulated images paired with their correct occupation information.

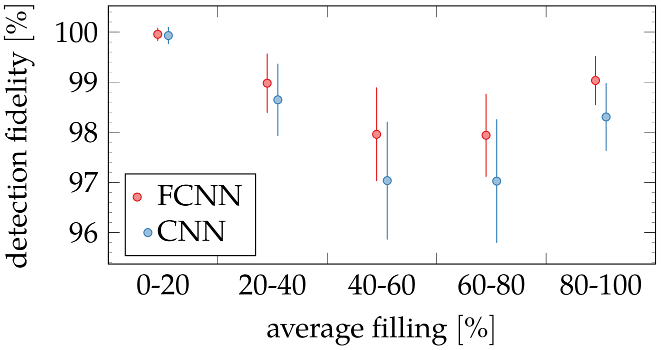

FIGURE 8.7.: Detection fidelities of the reconstruction on simulated images using binary classification with neural networks at different average fillings. Fully connected NNs (FCNN) and convolutional NNs (CNN) were implemented and tested. The data points were offset to improve visibility.

For training approximately 10^6 snippets were used. To test the detection fidelity of the trained networks we use a different set of simulated snippets. Figure 8.7 shows the achieved fidelities on simulated images in dependence of the mean filling of the lattice.

We found that this approach is limited by the quality of the simulated images. The fidelity of the networks is reduced when they are applied to experimental images due to the limitations of the image simulation. We improved the simulation to capture more of the imperfections observed in real images, but concluded that it is not viable to create a sufficiently good simulation. Training of the neural network on experimental images is not possible, as the supervised training requires labelled training data, i.e. snippets for which the occupation is already known.

8.2.4. Autoencoder

To overcome the limitations of the supervised learning needed for the binary classification networks we implemented an autoencoder. Autoencoders are systems for unsupervised learning consisting of an encoder, transforming the input data to a central layer with reduced dimension, commonly called bottleneck. The second stage, called decoder, is trained to reconstruct the original image only from the information contained in the central layer. This structure allows for unsupervised learning, training the encoder and decoder together, by minimising the error between input and output, only imposing a binary distribution on the central layer. The reduced dimension of the central layer forces the network to learn the underlying structure of the image, therefore the encoder is optimised to determine the lattice occupation.

For our application we use upsampled cut-outs from the full acquired images for 16×16 lattice sites, corresponding to an input layer size of 65 536. The central layer has a size of 256, one entry for every lattice site, forcing the NN to assign each entry the occupation of a lattice site. During training this information is used by the decoder to reconstruct an image of the size of the input image. The training process is unsupervised,

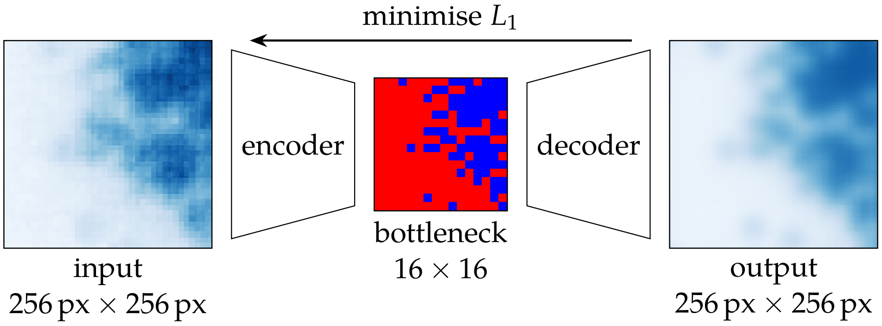

FIGURE 8.8.: Structure of the autoencoder based NN used to extract the lattice occupation. The network operates on image cut-outs of 16 × 16 lattice sites. The encoder is trained to reduce the information to a single central layer (bottleneck) of size 16 × 16. A second neural network reconstructs the original image only from the information contained in this layer. Both networks are trained in parallel, optimising the L_1 difference between the original and reconstructed image. The reduced dimensionality of the bottleneck forces the combined network to learn the underlying lattice occupation.

i.e. only the simulated images are passed to the network without information about the underlying atom distribution. The loss function used to optimise the system is chosen to minimise the L_1 difference between the input and simulated image. An additional regularisation term is added to ensure the output values in the bottleneck are either 0 or 1.

To extract the lattice occupation the measured image patches of matching size are cut from the full image and passed as input to the NN. The values in the central layer then directly correspond to the occupation information per lattice site. We generate patches starting on every lattice site, determining the occupation 256 times per lattice site, the final occupation value is determined using a majority vote. Using this approach we are able to achieve very high fidelities around 99.2 % on simulated images and 97 % on experimental images. The encoder and decoder are using CNN architectures internally. An overview of the structure with exemplary images is shown in fig. 8.8.

An advantage of the unsupervised training of the autoencoder is the option to use transfer learning. The NN is initially trained as usual using only simulated images. Afterwards the system is trained on a smaller set of experimental images with reduced learning parameters to improve the detection fidelity on experimental images. This approach is well suited for our application, as we can easily generate large amounts of simulated images for initial training, but only limited numbers of experimental images, due to the time needed for data acquisition.

8.2.5. Machine learning optimised projector

As an alternative to direct extraction of the lattice population from the acquired images using a NN we also implemented a projector based method, similar to ref. [182]. For each lattice site a kernel is applied to the image, projecting the information of the surrounding

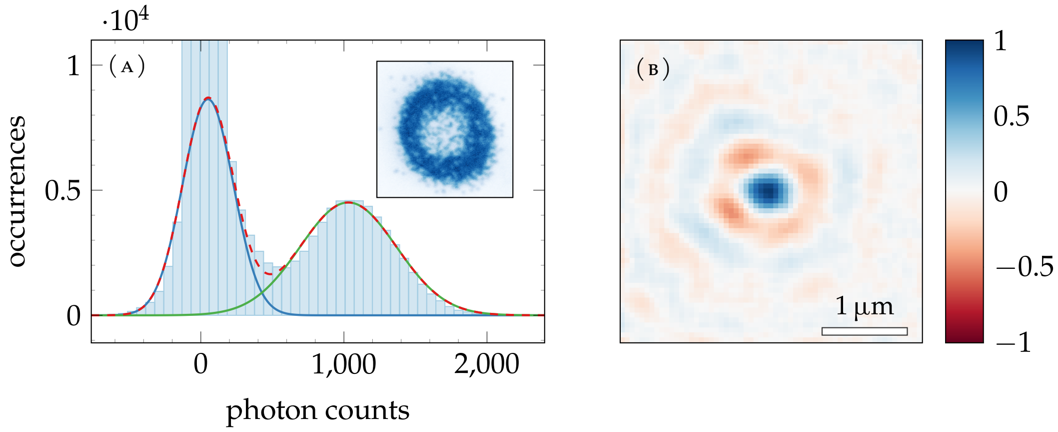

Figure 8.9.: (a) Histogram of the photon counts per site after application of the projector kernel. Two Gaussian functions were fitted to the background peak (blue) and the occupied sites (green) to determine the cut-off. The number of occurrences was clipped to get a good fit of the wings of the background peak. (b) Machine learning optimised projector kernel covering 9 lattice sites.

pixels to a single signal amplitude per site. From a histogram of the amplitudes a cut-off can be determined and be used to determine if a site is occupied or not (fig. 8.9a).

The kernel should have maximal overlap with signal from the same site, while assigning negative weights to the neighbouring sites to compensate for spill-over. Due to our large ratio between PSF size and lattice spacing it was not sufficient to use the PSF as kernel. Instead we use a simple NN to generate an optimised kernel. To achieve a meaningful output we use the values of the experimental PSF as starting point, only optimising the values using ML. The NN only has a single dense layer corresponding to the kernel matrix, followed by a single node layer determining whether the site is occupied. It is trained using supervised learning on simulated images, minimising binary crossentropy.

Using this method we generated a 61×61 kernel covering 9×9 lattice sites (fig. 8.9b) and implemented deconvolution using this kernel. Before application of the projector we upsample the images by a factor $20/7$ to improve the alignment with the lattice position computed from the lattice vectors.

Figure 8.9a shows an extracted histogram of the photon counts per pixel. Even though the two distributions visibly overlap significantly, the number of wrongly classified sites due to the overlap is only 3 % when using the intersection of the two fits as cut-off.

The classification based on the histogram data improves the robustness of the reconstruction against intensity fluctuations. The direct reconstruction approaches discussed before are trained on a specific signal intensity and require the input data to be correctly normalised to achieve high detection fidelities. Using the projector we can recompute the histogram and derive the cut-off for every image, cancelling intensity fluctuations.

9. Outlook

In this work I have presented the current status of the newly built Caesium quantum gas microscope, optimised for future experiments in systems with complex topological phases. The atomic properties of Caesium allow for the study of many body phases with tunable interactions, allowing to probe systems with single site resolution in regimes inaccessible in previous experiments. We have demonstrated the creation of MI states and implemented super-lattices in both horizontal axes. With those the constructed experiment now features all the necessary ingredients to perform first quantum simulation experiments. With the promising results from the ML based reconstruction we expect the first novel experiments with single site resolution to be performed within this year.

A compact assembly combines the necessary components to control and observe the atomic sample using light, radio frequencies, microwaves and magnetic fields. Ultra-stable current sources combined with an active magnetic field stabilisation allow full control over the magnetic fields and gradients in all axes. The magnetic field stability was characterised and found to have a residual noise level of 110 µG, achieving relative stabilities on the order of 10^{-6}. We have demonstrated this stability over several hours using microwave spectroscopy with locally resolved imaging in MI states.

During the writing of this thesis a DMD illuminated with broadband 520 nm light was installed in the experiment system and the first signal on the atoms was observed (fig. 9.1). The initial application will be the creation of a box trap with sharp walls and simultaneous flattening of the harmonic potential from the dipole traps. The flat potential can be used to study systems with homogeneous density distribution outside of the MI regime, whereas the sharp walls of the potential will enable the study of edge and corner modes present in topological systems [174].

In parallel the laser setup is being extended with additional 1534 nm beams along the lattice axes to create the Raman beams required for the implementation of laser assisted

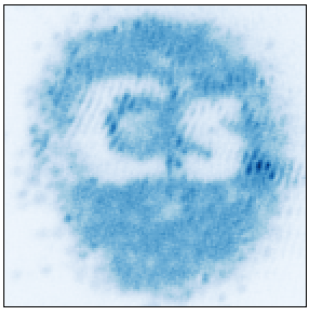

FIGURE 9.1.: MI state created with a local repulsive potential pattern, projected using a DMD.

tunnelling [8]. These could be used to study the system in the Meissner phase and the local currents therein [27]. The capabilities of the quantum gas microscope would allow to measure the chiral currents locally for the first time. In the future the same beams could also be used to extend the available lattice geometries to Lieb lattices [183].

In the future the state-dependent optical lattices will be added to the experiment to complete the designed novel platform for the study of topological systems and their exotic effects, such as the fractional quantum Hall effect [184].

A. QControl3

A.1. Sequence writing

The files defining the sequences to be processed by the software are written using regular Python code, using the `.pys` to mark them as sequence files. The software provides two methods to write sequences: The original sequence interface uses global variables to provide access to the defined channels and to declare the shots within a single task. A new interface was designed and in large parts implemented to reduce this dependence, helping with the creation of modular sequences. It uses a single interface class, providing access to the server information. To maintain compatibility both interfaces are supported and can be used in parallel.

In both cases sequences must be defined as subclasses of the `Script` class provided by the software. For every shot an instance of this class is processed, by calling the three functions `pre`, `main`, and `post` in sequence. These functions do nothing by default and are intended to be overwritten in the subclass to define the sequence. Functionally the three functions are identical. `pre` and `post` are intended to be defined in a more general subclass, defining code to be used by all sequences run in this experiment. `main` can then be used to implement the different sequences in subclasses of the general subclass.

The `Script` class, the more general `ScriptObject` class and the units and constants from the included unit library are imported by default into every sequence file. The units can be accessed as `u`, e.g. `u.s` for seconds. Effectively the following code is prepended to every imported sequence file

```
from qcontrol3.server.script import Script, ScriptObject
from qcontrol3.tools.units import const, u
```

To define the sequence all changes must be expressed as changes of channel values at defined timestamps using the channels `set` function[1]. At the beginning of the sequence all channels are automatically set to their default values, which are also restored at the end of the sequence. As a safeguard from programming mistakes, especially when channels are set from nested functions, the changes must be defined time-ordered per channel. Setting a channel at a time less than or equal to the time of the last value will result in a parser error. As an exception from the rule it is possible to change the value at 0 s from the default value.

```
volt_channel.set(0*u.s, 2*u.V)
```

---

[1] the same function can be accessed using the add name, as the changes are added to the event list

Setting a channel is associated with a small overhead to record the changes, and if enable store the calling location to help with debugging. When many changes should be applied on a single channel, i.e. to create ramps, `numpy` arrays should be used instead.

```
import numpy as np
times = (1 + np.arange(5)) * u.s
values = np.linspace(-5, 10, 5) * u.V
volt_channel.set(times, values)
```

A.1.1. Legacy interface

The classical interfaces exclusively uses global variables to provide access to the channels and to define the sequence execution. Channels are available using their name defined in configuration in global scope.

To determine which script should be parsed an instance of a script class must be assigned to a global variable named `timing_script`. If the script was imported to execute a single shot using the `queue_script_single` function the script is used as is.

Alternatively an iteration can be queued using the `queue_script_iteration` function. For iteration mode the iteration variables must be assigned to the global variable `iter_list`. For every entry in this list the script class will be parsed once. The state of the iteration can be accessed from the global variables `iter_value` containing the current entry from the list and `iter_step`, containing the index of the current value in the list[2]. Another global variable `iter_enabled` can be used to check if iteration mode is enabled. When a script is queued for iteration it is parsed once as if it would be queued for a single execution to check for sequence errors[3].

To add data from the sequence file to the writer data for storage it must be added as class or instance variables on the script. When such a variable is another `Script` or `ScriptObject`[4] a subfolder with the variables name is created and its variables are recursively saved into it. Data stored this way is placed in the `task` subfolder.

A.1.2. Server script interface

The goal of the new server script interface was to convert the sequences to regular Python files by removing 'magic' global variables. All interactions between the server and the script are wrapped by a single instance of the `ServerScriptInterface` class. Currently it is only available during parsing using the global `qc3` name, but it was designed such that it could be created anywhere on a client, only requiring the server to provide the channel information[5]. The channels can be accessed on the `channels` property, returning a dictionary with all user defined channels.

---

[2] To maintain compatibility with older versions both values are always defined and set to -1 when the script is parsed for single shots.

[3] again to maintain compatibility with older versions.

[4] or defines a `to_hdf5_dict` function returning the required data

[5] Direct parsing and sequence validation on the client side was a goal, as it would improve the user experience and reduce the server load, but is incompatible with the user provided converter functions.

To define the individual shots to be executed individual `Script` instances are passed to `queue_script`. Iterations or more complex structures can be created by calling this function repeatedly in a loop, possibly from helper functions. One of these can be defined as idle task by setting `idle=True`.

Data can be stored with the `save_data` function, by providing a path and the data. To avoid collisions with data saved using the legacy interface data all data is stored in the `script` subfolder.

Internally all scripts added using the legacy interface are translated into the new interface, allowing handling of both cases using a single code path.

These are implemented as function overloads on the user channels, preventing serialisation. When breaking changes are accepted these could be replaced by a set of pre-defined calibrations. For all currently known applications linear translation and spline interpolation are sufficient.

B. Magnetic field stabilisation system

B.1. Sensor inputs

As described in the main text the sensor for the *z* axis could not be installed as planned. To prevent regular resets of the stabilisation system due to unavoidable sensor overloads on this axis we needed to disable the sensor. As it turns out it is not enough to disable the usage of the sensor in the software configuration, as it will still trigger a system reset on overload. Instead we installed a small adapter box in between the sensor and the stabilisation system. This sensor passes all connections as they are, but the sensor line for the *z* sensor, which is left open on both ends. The stabilisation system appears to provide internal pull down resistors as the sensor input is now permanently at zero, effectively preventing any resets due to sensor overloads.

B.2. Monitoring outputs

The stabilisation system provides three different monitoring outputs, configurable in software to show various monitoring outputs. We primarily use them to monitor the measured field from the sensors, or the generated output to the compensation coils.

The outputs support a range of $\pm 10\,\mathrm{V}$. When the sensor readings are shows they are shifted by the compensation value, such that an output of $10\,\mathrm{V}$ corresponds to a perfectly compensated value. When monitoring the outputs the measurement range appears to correspond to the capabilities of the output driver. According to the manufacturer the sensor readings are additionally scaled by a factor of 10, corresponding to a scaling of $7\,\mathrm{mG/V}$.

The primary output 'Diag Out' is a direct analogue output, whereas the outputs 'Diag 0' and 'Diag 1' are digital outputs. Output options on 'Diag Out' are limited to the sensor readings, the others can show the outputs and also intermediate values from inside the feedback loops. Signal on the digital outputs are slightly delayed compared to the analogue output and are inverted.

B.3. Feedback loop optimisation

After configuration of the input and output matrix parameters described in section 6.3.2 we performed a quick optimisation of the feedback loop parameters. We tuned the DC and AC loops manually, and used the automatic optimisation procedure to detect frequencies and gain parameters for the selective frequency loop.

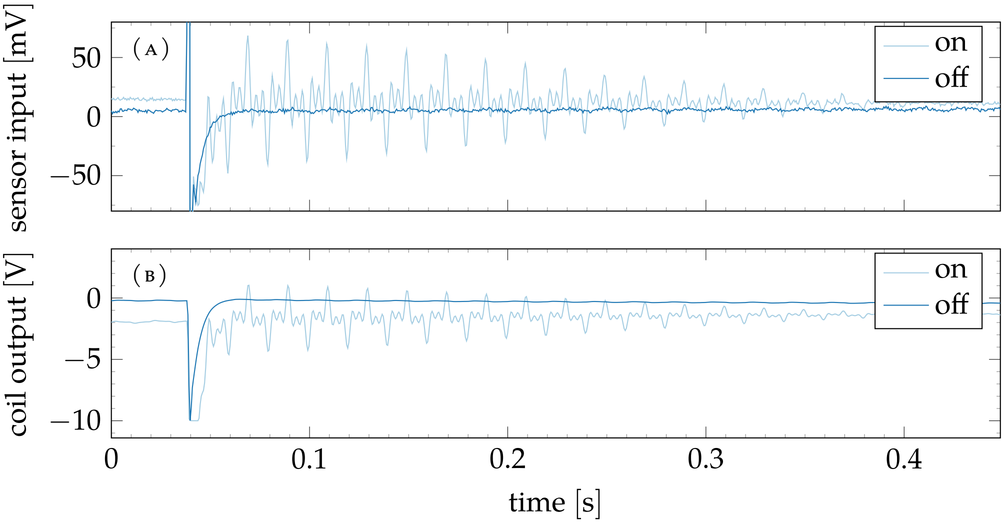

FIGURE B.1.: Oscillations of the feedback loop of the stabilisation system observable in the sensor (a) and the compensation coil output (b) after a short field saturation when the DC loop and the selective frequency loop at DC are active at the same time (on trace). The saturation is generated from induced currents when the offset field is quenched to zero using the MOSFET box. Removing the DC component from the selective frequency loop removes the oscillations (off trace).

Manual optimisation of the loops is very quick as the loops only have a single gain parameter each. Initially we increased either parameter until the system started to oscillate, keeping the other parameter at zero. When activating both systems at the same time the DC gain can be increased a bit further, similar to the integral part in a standard PI loop. We then chose values about 10 % below the oscillation threshold. Improvements to the chosen parameters are most likely possible with a more thorough optimisation, optimising on the reaction to a step response and excitation at specific frequencies.

While testing the system we disabled the DC component of the selective frequency loop, which is enabled by default. We found that the combination of DC loop and selective frequency loop at DC introduced oscillations with ∼1 mG peak-to-peak amplitude at 50 Hz and its higher harmonics when the sensors were saturated for a few ms due to rapid field changes. These oscillations decayed only slowly within 250 ms (fig. B.1), causing atom number fluctuations in the Mott insulator stage.

B.4. Output amplifier current overload

A major practical issue we found with the stabilisation system is a shutdown of the output amplifiers due to a current overload. The specified output capability of the driver is 300 W peak, but only 100 W in continuous operation. When the output amplifier has to provide a high load for extended periods of time it disables the current output, disabling the feedback loop. After a few seconds it activates again, usually disabling on the next experimental cycle when the high load occurs again. The shutdown appears to be controlled by a temperature sensor and/or the microcontroller within the stabilisation system. This leads to a non-deterministic behaviour of the system. The average load of the output amplifier is dependent on the amount of generated fields seen by the sensors within a sequence, hence it is dependent on the sequence duration and the fields generated by it. When changing to a sequence with a too high average load the system operates without problems for some time (minutes up to a few hours) before it starts to activate the output protection. From there on every cycle of the sequence tends to disable the output, however not at a fixed time in the sequence, but with a variation of $\pm 50\,\mathrm{ms}$.

So far we were always able to solve this problem by either optimising the alignment of the corresponding sensor to reduce the coupling of our generated fields, or by adjusting the sequence to reduce the influence on the sensor. Even though we tried to avoided metallic components as much as possible around the chamber, magnetic field gradients are picked up by the sensors and can cause the output overload protection to activate. Even after a strong field gradient of $30\,\mathrm{G/ms}$ the field stabilises within a few milliseconds. Without further modifications the feedback loop requires $25\,\mathrm{ms}$ to recover from the overdrive due to the saturation.

To avoid these issues and avoid the recovery time we are planning to integrate a small circuit which can be used to temporarily disable the sensor input when strong gradients are expected, e.g. on coil shutdown. This will result in a short period where the feedback loop is not and the output will be not stabilised, most likely followed by a short recovery period.

A possible permanent solution for this problem is the usage of an external output amplifier with higher continuous power rating. The stabilisation system has an amplifier pre-out connector which can be used for this application. Care must be taken when choosing the output amplifier to not couple additional noise sources into the stabilisation coils. As the individual coil resistance is approximately $16\,\Omega$ audio amplifiers might be a suitable choice, which are readily available with high power ratings and can have good noise performance.

Alternatively it should be possible to replace the stabilisation coils with new coils with twice the amount of turns, also doubling the wire cross section. The resulting coils would have the same resistance, but generate twice the field, reducing the load on the output amplifier. The increased inductance might reduce the bandwidth of the system, but it should still be high enough for the relevant frequencies up to $500\,\mathrm{Hz}$.

Bibliography

[1] I. Bloch, J. Dalibard and W. Zwerger, *Many-Body Physics with Ultracold Gases*, Rev. Mod. Phys. **80**, 885–964 (2008) 10.1103/RevModPhys.80.885 (cit. on pp. 1, 3, 6, 98).

[2] I. Bloch, J. Dalibard and S. Nascimbène, *Quantum Simulations with Ultracold Quantum Gases*, Nature Phys **8**, 267–276 (2012) 10.1038/nphys2259 (cit. on p. 1).

[3] C.-K. Chiu, J. C. Y. Teo, A. P. Schnyder and S. Ryu, *Classification of Topological Quantum Matter with Symmetries*, Rev. Mod. Phys. **88**, 035005 (2016) 10.1103/RevModPhys.88.035005 (cit. on p. 1).

[4] N. Goldman, G. Juzeliūnas, P. Öhberg and I. B. Spielman, *Light-induced Gauge Fields for Ultracold Atoms*, Rep. Prog. Phys. **77**, 126401 (2014) 10.1088/0034-4885/77/12/126401 (cit. on p. 1).

[5] N. R. Cooper, J. Dalibard and I. B. Spielman, *Topological Bands for Ultracold Atoms*, Rev. Mod. Phys. **91**, 015005 (2019) 10.1103/RevModPhys.91.015005 (cit. on pp. 1, 3, 6).

[6] P. G. Harper, *The General Motion of Conduction Electrons in a Uniform Magnetic Field, with Application to the Diamagnetism of Metals*, Proc. Phys. Soc. A **68**, 879–892 (1955) 10.1088/0370-1298/68/10/305 (cit. on p. 1).

[7] D. R. Hofstadter, *Energy Levels and Wave Functions of Bloch Electrons in Rational and Irrational Magnetic Fields*, Phys. Rev. B **14**, 2239–2249 (1976) 10.1103/PhysRevB.14.2239 (cit. on p. 1).

[8] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes and I. Bloch, *Realization of the Hofstadter Hamiltonian with Ultracold Atoms in Optical Lattices*, Phys. Rev. Lett. **111**, 185301 (2013) 10.1103/PhysRevLett.111.185301 (cit. on pp. 1, 3, 9, 122).

[9] H. Miyake, G. A. Siviloglou, C. J. Kennedy, W. C. Burton and W. Ketterle, *Realizing the Harper Hamiltonian with Laser-Assisted Tunneling in Optical Lattices*, Phys. Rev. Lett. **111**, 185302 (2013) 10.1103/PhysRevLett.111.185302 (cit. on p. 1).

[10] W. P. Su, J. R. Schrieffer and A. J. Heeger, *Solitons in Polyacetylene*, Phys. Rev. Lett. **42**, 1698–1701 (1979) 10.1103/PhysRevLett.42.1698 (cit. on pp. 1, 106).

[11] M. Atala, M. Aidelsburger, J. T. Barreiro, D. Abanin, T. Kitagawa, E. Demler and I. Bloch, *Direct Measurement of the Zak Phase in Topological Bloch Bands*, Nature Phys. **9**, 795–800 (2013) 10.1038/nphys2790 (cit. on pp. 1, 3).

[12] E. J. Meier, F. A. An and B. Gadway, *Observation of the Topological Soliton State in the Su–Schrieffer–Heeger Model*, Nat Commun **7**, 13986 (2016) 10.1038/ncomms13986 (cit. on p. 1).

[13] F. D. M. Haldane, *Model for a Quantum Hall Effect without Landau Levels: Condensed-Matter Realization of the "Parity Anomaly"*, Phys. Rev. Lett. **61**, 2015–2018 (1988) 10.1103/PhysRevLett.61.2015 (cit. on p. 1).

[14] G. Jotzu, M. Messer, R. Desbuquois, M. Lebrat, T. Uehlinger, D. Greif and T. Esslinger, *Experimental Realization of the Topological Haldane Model with Ultracold Fermions*, Nature **515**, 237–240 (2014) 10.1038/nature13915 (cit. on p. 1).

[15] K. Wintersperger, C. Braun, F. N. Ünal, A. Eckardt, M. D. Liberto, N. Goldman, I. Bloch and M. Aidelsburger, *Realization of an Anomalous Floquet Topological System with Ultracold Atoms*, Nat. Phys. **16**, 1058–1063 (2020) 10.1038/s41567-020-0949-y (cit. on p. 1).

[16] H. Bartolomei, M. Kumar, R. Bisognin, A. Marguerite, J.-M. Berroir, E. Bocquillon, B. Plaçais, A. Cavanna, Q. Dong, U. Gennser, Y. Jin and G. Fève, *Fractional Statistics in Anyon Collisions*, Science **368**, 173–177 (2020) 10.1126/science.aaz5601 (cit. on p. 1).

[17] L. W. Clark, N. Schine, C. Baum, N. Jia and J. Simon, *Observation of Laughlin States Made of Light*, Nature **582**, 41–45 (2020) 10.1038/s41586-020-2318-5 (cit. on p. 1).

[18] M. Aghaee et al., *InAs-Al Hybrid Devices Passing the Topological Gap Protocol*, 7th July 2022, 10.48550/arXiv.2207.02472 (cit. on p. 1).

[19] C. Gerry and P. Knight, *Introductory Quantum Optics* (Cambridge University Press, Oct. 2004) (cit. on p. 3).

[20] R. Grimm, M. Weidemüller and Y. B. Ovchinnikov, *Optical Dipole Traps for Neutral Atoms*, in *Advances In Atomic, Molecular, and Optical Physics*, Vol. 42, edited by B. Bederson and H. Walther (Academic Press, 1st Jan. 2000), pp. 95–170, 10.1016/S1049-250X(08)60186-X (cit. on p. 3).

[21] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping of Neutral Atoms*, in *The Optics Encyclopedia* (John Wiley & Sons, Ltd, 2007), 10.1002/9783527600441.oe005 (cit. on pp. 3, 13).

[22] T. Rom, *Bosonische und fermionische Quantengase in dreidimensionalen optischen Gittern*, PhD thesis, Ludwig-Maximilians-Universität München (2009) (cit. on p. 3).

[23] A. Browaeys and T. Lahaye, *Many-body Physics with Individually Controlled Rydberg Atoms*, Nat. Phys. **16**, 132–142 (2020) 10.1038/s41567-019-0733-z (cit. on pp. 3, 50).

[24] C. Gross and W. S. Bakr, *Quantum Gas Microscopy for Single Atom and Spin Detection*, Nat. Phys. **17**, 1316–1323 (2021) 10.1038/s41567-021-01370-5 (cit. on p. 3).

[25] M. Greiner, *Ultracold quantum gases in three-dimensional optical lattice potentials*, PhD thesis, Ludwig-Maximilians-Universität München (2003) (cit. on pp. 3, 6, 16).

[26] M. Aidelsburger, M. Atala, S. Nascimbène, S. Trotzky, Y.-A. Chen and I. Bloch, *Experimental Realization of Strong Effective Magnetic Fields in an Optical Lattice*, Phys. Rev. Lett. **107**, 255301 (2011) 10.1103/PhysRevLett.107.255301 (cit. on pp. 3, 8).

[27] M. Atala, M. Aidelsburger, M. Lohse, J. T. Barreiro, B. Paredes and I. Bloch, *Observation of Chiral Currents with Ultracold Atoms in Bosonic Ladders*, Nature Phys **10**, 588–593 (2014) 10.1038/nphys2998 (cit. on pp. 3, 8, 122).

[28] M. Aidelsburger, M. Lohse, C. Schweizer, M. Atala, J. T. Barreiro, S. Nascimbène, N. R. Cooper, I. Bloch and N. Goldman, *Measuring the Chern Number of Hofstadter Bands with Ultracold Bosonic Atoms*, Nature Phys. **11**, 162–166 (2015) 10.1038/nphys3171 (cit. on pp. 3, 9).

[29] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling and M. Greiner, *A Quantum Gas Microscope for Detecting Single Atoms in a Hubbard-regime Optical Lattice*, Nature **462**, 74–77 (2009) 10.1038/nature08482 (cit. on pp. 3, 31, 109).

[30] D. Jaksch and P. Zoller, *Creation of Effective Magnetic Fields in Optical Lattices: The Hofstadter Butterfly for Cold Neutral Atoms*, New J. Phys. **5**, 56–56 (2003) 10.1088/1367-2630/5/1/356 (cit. on pp. 3, 8, 9).

[31] J. von Fraunhofer, *Bestimmung des Brechungs- und des Farbenzerstreungs-Vermögens verschiedener Glasarten, in Bezug auf die Vervollkommnung achromatischer Fernröhre*, Ann. Phys. **56**, 264–313 (1817) 10.1002/ANDP.18170560706 (cit. on p. 4).

[32] J. F. Wienand, *Multi-Level Fluorescence Imaging and Degenerate Raman Sideband Cooling for a Caesium Quantum Gas Microscope*, MA thesis, Ludwig-Maximilians-Universität München (2019) (cit. on pp. 4, 21, 28).

[33] S. Hubele, *Potential Shaping Using a DMD and High-Resolution Imaging of Cesium Atoms in Optical Lattices*, MA thesis, Ludwig-Maximilians-Universität München (2022) (cit. on p. 4).

[34] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch and S. Kuhr, *Single-spin Addressing in an Atomic Mott Insulator*, Nature **471**, 319–324 (2011) 10.1038/nature09827 (cit. on p. 4).

[35] J.-L. Basdevant and J. Dalibard, *Quantum Mechanics* (Springer, 2005) (cit. on p. 5).

[36] C. Chin, R. Grimm, P. Julienne and E. Tiesinga, *Feshbach Resonances in Ultracold Gases*, Rev. Mod. Phys. **82**, 1225–1286 (2010) 10.1103/RevModPhys.82.1225 (cit. on pp. 5, 6, 57).

[37] L. W. Clark, L.-C. Ha, C.-Y. Xu and C. Chin, *Quantum Dynamics with Spatiotemporal Control of Interactions in a Stable Bose-Einstein Condensate*, Phys. Rev. Lett. **115**, 155301 (2015) 10.1103/PhysRevLett.115.155301 (cit. on p. 5).

[38] M. D. Frye, B. C. Yang and J. M. Hutson, *Ultracold Collisions of Cs Atoms in Excited Zeeman and Hyperfine States*, Phys. Rev. A **100**, 022702 (2019) 10.1103/PhysRevA.100.022702 (cit. on pp. 5, 57).

[39] M. D. Frye and J. M. Hutson, private communication, 2022 (cit. on p. 5).

[40] T. Kraemer, *Few-Body Interaction in an Ultracold Gas of Cesium Atoms*, PhD thesis, Leopold-Franzens-Universität Innsbruck (2006) (cit. on pp. 5, 6).

[41] P. J. Leo, C. J. Williams and P. S. Julienne, *Collision Properties of Ultracold ^{133}Cs Atoms*, Phys. Rev. Lett. **85**, 2721–2724 (2000) 10.1103/PhysRevLett.85.2721 (cit. on p. 5).

[42] C. Chin, V. Vuletić, A. J. Kerman and S. Chu, *High Resolution Feshbach Spectroscopy of Cesium*, Phys. Rev. Lett. **85**, 2717–2720 (2000) 10.1103/PhysRevLett.85.2717 (cit. on p. 5).

[43] C. Chin, V. Vuletić, A. J. Kerman, S. Chu, E. Tiesinga, P. J. Leo and C. J. Williams, *Precision Feshbach Spectroscopy of Ultracold Cs_2*, Phys. Rev. A **70**, 032701 (2004) 10.1103/PhysRevA.70.032701 (cit. on p. 5).

[44] M. Berninger, A. Zenesini, B. Huang, W. Harm, H.-C. Nägerl, F. Ferlaino, R. Grimm, P. S. Julienne and J. M. Hutson, *Feshbach Resonances, Weakly Bound Molecular States and Coupled-channel Potentials for Cesium at High Magnetic Fields*, Phys. Rev. A **87**, 10.1103/PhysRevA.87.032517 (2013) 10.1103/PhysRevA.87.032517 (cit. on p. 5).

[45] M. Mark, F. Meinert, K. Lauber and H.-C. Nägerl, *Mott-insulator-aided Detection of Ultra-narrow Feshbach Resonances*, SciPost Phys. **5**, 055 (2018) 10.21468/SciPostPhys.5.5.055 (cit. on p. 5).

[46] M. Gustavsson, E. Haller, M. J. Mark, J. G. Danzl, G. Rojas-Kopeinig and H.-C. Nägerl, *Control of Interaction-Induced Dephasing of Bloch Oscillations*, Phys. Rev. Lett. **100**, 080404 (2008) 10.1103/PhysRevLett.100.080404 (cit. on p. 5).

[47] C. Gross and I. Bloch, *Quantum Simulations with Ultracold Atoms in Optical Lattices*, Science **357**, 995–1001 (2017) 10.1126/science.aal3837 (cit. on p. 6).

[48] J. Dalibard, *Introduction to the physics of artificial gauge fields*, in Quantum Matter at Ultralow Temperatures, edited by M. Inguscio, W. Ketterle and S. Stringari (2016), pp. 1–55 (cit. on p. 6).

[49] Y. Aharonov and D. Bohm, *Significance of Electromagnetic Potentials in the Quantum Theory*, Phys. Rev. **115**, 485–491 (1959) 10.1103/PhysRev.115.485 (cit. on p. 6).

[50] M. Aidelsburger, S. Nascimbene and N. Goldman, *Artificial Gauge Fields in Materials and Engineered Systems*, C. R. Phys., Quantum Simulation / Simulation Quantique **19**, 394–432 (2018) 10.1016/j.crhy.2018.03.002 (cit. on p. 8).

[51] L. Duca, T. Li, M. Reitter, I. Bloch, M. Schleier-Smith and U. Schneider, *An Aharonov-Bohm Interferometer for Determining Bloch Band Topology*, Science **347**, 288–292 (2015) 10.1126/science.1259052 (cit. on p. 8).

[52] C. Weitenberg and J. Simonet, *Tailoring Quantum Gases by Floquet Engineering*, Nat. Phys. **17**, 1342–1348 (2021) 10.1038/s41567-021-01316-x (cit. on p. 8).

[53] Y.-J. Lin, R. L. Compton, K. Jiménez-García, J. V. Porto and I. B. Spielman, *Synthetic Magnetic Fields for Ultracold Neutral Atoms*, Nature **462**, 628–632 (2009) 10.1038/nature08609 (cit. on p. 8).

[54] M. Mancini, G. Pagano, G. Cappellini, L. Livi, M. Rider, J. Catani, C. Sias, P. Zoller, M. Inguscio, M. Dalmonte and L. Fallani, *Observation of Chiral Edge States with Neutral Fermions in Synthetic Hall Ribbons*, Science **349**, 1510–1513 (2015) 10.1126/science.aaa8736 (cit. on p. 8).

[55] A. R. Kolovsky, *Creating Artificial Magnetic Fields for Cold Atoms by Photon-assisted Tunneling*, EPL **93**, 20003 (2011) 10.1209/0295-5075/93/20003 (cit. on p. 8).

[56] F. Gerbier and J. Dalibard, *Gauge Fields for Ultracold Atoms in Optical Superlattices*, New J. Phys. **12**, 033007 (2010) 10.1088/1367-2630/12/3/033007 (cit. on pp. 8, 9).

[57] R. Wei and E. J. Mueller, *Magnetic-field Dependence of Raman Coupling in Alkali-metal Atoms*, Phys. Rev. A **87**, 042514 (2013) 10.1103/PhysRevA.87.042514 (cit. on p. 9).

[58] J. Koepsell, S. Hirthe, D. Bourgund, P. Sompet, J. Vijayan, G. Salomon, C. Gross and I. Bloch, *Robust Bilayer Charge Pumping for Spin-and Density-Resolved Quantum Gas Microscopy*, Phys. Rev. Lett. **125**, 010403 (2020) 10.1103/PhysRevLett.125.010403 (cit. on pp. 10, 36).

[59] T. M. Klostermann, *Construction of a caesium quantum gas microscope*, PhD thesis, Ludwig-Maximilians-Universität München (2022) (cit. on pp. 11, 13, 15, 19, 26, 29, 96, 115).

[60] C. Deppner, W. Herr, M. Cornelius, P. Stromberger, T. Sternke, C. Grzeschik, A. Grote, J. Rudolph, S. Herrmann, M. Krutzik, A. Wenzlawski, R. Corgier, E. Charron, D. Guéry-Odelin, N. Gaaloul, C. Lämmerzahl, A. Peters, P. Windpassinger and E. M. Rasel, *Collective-Mode Enhanced Matter-Wave Optics*, Phys. Rev. Lett. **127**, 100401 (2021) 10.1103/PhysRevLett.127.100401 (cit. on p. 12).

[61] W. Ketterle, D. S. Durfee and D. M. Stamper-Kurn, *Making, Probing and Understanding Bose-Einstein Condensates*, 1st Apr. 1999 (cit. on pp. 12, 13).

[62] N. F. Ramsey, *Molecular Beams*, Oxford Classic Texts in the Physical Sciences (Oxford University Press, USA, 1956) (cit. on p. 12).

[63] F. Scazza, *Probing $SU(N)$-symmetric orbital interactions with ytterbium Fermi gases in optical lattices*, PhD thesis, Ludwig-Maximilians-Universität München (2015) (cit. on p. 12).

[64] S. Chu and C. Wieman, *Laser Cooling and Trapping of Atoms: Introduction*, J. Opt. Soc. Am. B, JOSAB **6**, 2020–2022 (1989) 10.1364/JOSAB.6.002020 (cit. on p. 13).

[65] W. Ketterle and M. W. Zwierlein, *Making, Probing and Understanding Ultracold Fermi Gases*, 16th Jan. 2008, 10.1393/ncr/i2008-10033-1 (cit. on p. 13).

[66] M. Gustavsson, *A Quantum Gas with Tunable Interactions in an Optical Lattice*, PhD thesis, Leopold-Franzens-Universität Innsbruck (2008) (cit. on pp. 13, 27).

[67] C.-L. Hung, *In Situ Probing of Two-dimensional Quantum Gases*, PhD thesis, University of Chicago (2011) (cit. on pp. 13, 27).

[68] J. R. Kellogg, D. Schlippert, J. M. Kohel, R. J. Thompson, D. C. Aveline and N. Yu, *A Compact High-efficiency Cold Atom Beam Source*, Appl. Phys. B **109**, 61–64 (2012) 10.1007/s00340-012-5220-5 (cit. on p. 13).

[69] A. Z. Lam, C. Warner, N. Bigagli, S. Roschinski, W. Yuan, I. Stevenson and S. Will, *Compact Two-Dimensional Magneto-Optical Trap for Ultracold Atom Setups*, 1st July 2021, 10.48550/arXiv.2012.06688 (cit. on p. 13).

[70] D. A. Steck, *Caesium D Line Data*, 23rd Dec. 2010 (cit. on p. 14).

[71] D. O. Subulsky, *A Cesium-133 Effusive Oven for Ultracold Atomic Experiments*, BA thesis, University of Chicago (2014) (cit. on p. 14).

[72] M. Boll, *Spin and density resolved microscopy of Hubbard chains*, PhD thesis, Ludwig-Maximilians-Universität München (2016) (cit. on p. 17).

[73] K. G. H. Viebahn, *Quasicrystalline Optical Lattices for Ultracold Atoms*, PhD thesis, University of Cambridge (2019) (cit. on pp. 17, 43).

[74] T. Y. Li, *An Apparatus for Probing Fermions in a Quasi-Two-Dimensional Geometry*, MA thesis, Ludwig-Maximilians-Universität München (2011) (cit. on p. 17).

[75] D. A. Steck, *Rubidium 87 D Line Data*, 13th Jan. 2015 (cit. on pp. 17, 38).

[76] A. Reetz, *Laser Cooling, Transport and Imaging of Caesium Atoms*, MA thesis, Ludwig-Maximilians-Universität München (2018) (cit. on pp. 19, 20).

[77] R. K. Raj, D. Bloch, J. J. Snyder, G. Camy and M. Ducloy, *High-Frequency Optically Heterodyned Saturation Spectroscopy Via Resonant Degenerate Four-Wave Mixing*, Phys. Rev. Lett. **44**, 1251–1254 (1980) 10.1103/PhysRevLett.44.1251 (cit. on p. 19).

[78] H. N. Rutt, *A Heterodyne Frequency Offset Locking Technique for Pulsed or CW Lasers*, J. Phys. E: Sci. Instrum. **17**, 704–709 (1984) 10.1088/0022-3735/17/8/017 (cit. on p. 19).

[79] D. J. McCarron, S. A. King and S. L. Cornish, *Modulation Transfer Spectroscopy in Atomic Rubidium*, Meas. Sci. Technol. **19**, 105601 (2008) 10.1088/0957-0233/19/10/105601 (cit. on pp. 19, 20).

[80] I. Petitbon, P. Gallion, G. Debarge and C. Chabran, *Locking Bandwidth and Relaxation Oscillations of an Injection-locked Semiconductor Laser*, IEEE J. Quantum Electron. **24**, 148–154 (1988) 10.1109/3.108 (cit. on p. 21).

[81] P. Unterwaditzer, *Aufbau eines vollständigen Diodenlasersystems zur Laserkühlung und Detektion von gespeicherten Cs-Atomen*, Diploma Thesis, Leopold-Franzens-Universität Innsbruck (2005) (cit. on p. 21).

[82] D. Bourgund, *High-Stability, Tunable, Bichromatic Superlattice*, MA thesis, Ludwig-Maximilians-Universität München (2019) (cit. on p. 22).

[83] J. Koepsell, *Quantum simulation of doped two-dimensional Mott insulators*, PhD thesis, Ludwig-Maximilians-Universität München (2021) (cit. on pp. 22, 44, 46).

[84] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, 1st, Wiley Series in Pure and Applied Optics (Wiley-Intersience, 1991) (cit. on p. 23).

[85] N. Darkwah Oppong, *Towards a Degenerate Fermi Gas of Strontium-87 in a 3D Optical Lattice*, MA thesis, ETH Zurich (2015) (cit. on p. 23).

[86] S. Blatt, A. Mazurenko, M. F. Parsons, C. S. Chiu, F. Huber and M. Greiner, *Low-noise Optical Lattices for Ultracold* ^{6}*Li*, Phys. Rev. A **92**, 021402 (2015) 10.1103/PhysRevA.92.021402 (cit. on p. 23).

[87] J. G. Hinney, *An Adaptive Intensity Stabilizer for Optical Lattices*, MA thesis, University of Strathclyde (2014) (cit. on p. 23).

[88] J. Chen, *Fast Long-distance Transport of Caesium Atoms and Bichromatic Superlattice*, MA thesis, Ludwig-Maximilians-Universität München (2020) (cit. on pp. 25, 29).

[89] C. J. Dedman, J. Nes, T. M. Hanna, R. G. Dall, K. G. H. Baldwin and A. G. Truscott, *Optimum Design and Construction of a Zeeman Slower for Use with a Magneto-optic Trap*, Rev. Sci. Instrum. **75**, 5136–5142 (2004) 10.1063/1.1820524 (cit. on p. 26).

[90] R. H. Dicke, *The Effect of Collisions upon the Doppler Width of Spectral Lines*, Phys. Rev. **89**, 472–473 (1953) 10.1103/PhysRev.89.472 (cit. on p. 27).

[91] S. Chu, *Nobel Lecture: The Manipulation of Neutral Particles*, Rev. Mod. Phys. **70**, 685–706 (1998) 10.1103/RevModPhys.70.685 (cit. on p. 27).

[92] A. J. Kerman, V. Vuletić, C. Chin and S. Chu, *Beyond Optical Molasses: 3D Raman Sideband Cooling of Atomic Cesium to High Phase-Space Density*, Phys. Rev. Lett. **84**, 439–442 (2000) 10.1103/PhysRevLett.84.439 (cit. on p. 28).

[93] A. Flir, *Implementierung und Untersuchung von Raman-Seitenbandkühlung zur Erzeugung eines ultrakalten Cäsiumgases*, Diploma Thesis, Leopold-Franzens-Universität Innsbruck (2006) (cit. on p. 28).

[94] G. Grynberg and C. Robilliard, *Cold Atoms in Dissipative Optical Lattices*, Physics Reports **355**, 335–451 (2001) 10.1016/S0370-1573(01)00017-5 (cit. on p. 28).

[95] A. J. Kerman, *Raman Sideband Cooling and Cold Atomic Collisions in Optical Lattices*, PhD thesis, Stanford University (2002) (cit. on p. 28).

[96] T. Klostermann, C. R. Cabrera, H. von Raven, J. F. Wienand, C. Schweizer, I. Bloch and M. Aidelsburger, *Fast Long-distance Transport of Cold Cesium Atoms*, Phys. Rev. A **105**, 043319 (2022) 10.1103/PhysRevA.105.043319 (cit. on pp. 29, 30).

[97] J. Léonard, M. Lee, A. Morales, T. M. Karg, T. Esslinger and T. Donner, *Optical Transport and Manipulation of an Ultracold Atomic Cloud Using Focus-tunable Lenses*, New J. Phys. **16**, 093028 (2014) 10.1088/1367-2630/16/9/093028 (cit. on p. 29).

[98] C. Gross, H. C. J. Gan and K. Dieckmann, *All-optical Production and Transport of a Large ^{6}Li Quantum Gas in a Crossed Optical Dipole Trap*, Phys. Rev. A **93**, 10.1103/PhysRevA.93.053424 (2016) 10.1103/PhysRevA.93.053424 (cit. on p. 29).

[99] D. Schrader, S. Kuhr, W. Alt, M. Müller, V. Gomer and D. Meschede, *An Optical Conveyor Belt for Single Neutral Atoms*, Appl Phys B **73**, 819–824 (2001) 10.1007/s003400100722 (cit. on p. 29).

[100] S. Schmid, G. Thalhammer, K. Winkler, F. Lang and J. H. Denschlag, *Long Distance Transport of Ultracold Atoms Using a 1D Optical Lattice*, New J. Phys. **8**, 159 (2006) 10.1088/1367-2630/8/8/159 (cit. on p. 29).

[101] M. Langbecker, R. Wirtz, F. Knoch, M. Noaman, T. Speck and P. Windpassinger, *Highly Controlled Optical Transport of Cold Atoms into a Hollow-core Fiber*, New J. Phys. **20**, 083038 (2018) 10.1088/1367-2630/aad9bb (cit. on p. 29).

[102] M. Greiner, I. Bloch, T. W. Hänsch and T. Esslinger, *Magnetic Transport of Trapped Cold Atoms over a Large Distance*, Phys. Rev. A **63**, 031401 (2001) 10.1103/PhysRevA.63.031401 (cit. on p. 29).

[103] D. Pertot, D. Greif, S. Albert, B. Gadway and D. Schneble, *Versatile Transporter Apparatus for Experiments with Optically Trapped Bose–Einstein Condensates*, J. Phys. B: At. Mol. Opt. Phys. **42**, 215305 (2009) 10.1088/0953-4075/42/21/215305 (cit. on p. 29).

[104] O. Brzobohatý, T. Čižmár and P. Zemánek, *High Quality quasi-Bessel Beam Generated by Round-tip Axicon*, Opt. Express, OE **16**, 12688–12700 (2008) 10.1364/OE.16.012688 (cit. on p. 29).

[105] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl and R. Grimm, *Bose-Einstein Condensation of Cesium*, Science **299**, 232–235 (2002) 10.1126/science.1079699 (cit. on pp. 29, 93).

[106] M. F. Parsons, F. Huber, A. Mazurenko, C. S. Chiu, W. Setiawan, K. Wooley-Brown, S. Blatt and M. Greiner, *Site-Resolved Imaging of Fermionic ^{6}Li in an Optical Lattice*, Phys. Rev. Lett. **114**, 213002 (2015) 10.1103/PhysRevLett.114.213002 (cit. on pp. 31, 34, 109, 115).

[107] C. Robens, S. Brakhane, W. Alt, F. Kleißler, D. Meschede, G. Moon, G. Ramola and A. Alberti, *High Numerical Aperture ($NA = 0.92$) Objective Lens for Imaging and Addressing of Cold Atoms*, Opt. Lett., OL **42**, 1043–1046 (2017) 10.1364/OL.42.001043 (cit. on p. 31).

[108] A. Heinz, *Ultracold strontium in state-dependent optical lattices*, PhD thesis, Ludwig-Maximilians-Universität München (2020) (cit. on p. 31).

[109] J. Yang, L. Liu, J. Mongkolkiattichai and P. Schauss, *Site-Resolved Imaging of Ultracold Fermions in a Triangular-Lattice Quantum Gas Microscope*, PRX Quantum **2**, 020344 (2021) 10.1103/PRXQuantum.2.020344 (cit. on p. 31).

[110] A. Z. H. Lam, *Ultracold Dipolar Gases of NaCs Ground State Molecules*, PhD thesis, Columbia University (2022) (cit. on p. 31).

[111] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch and S. Kuhr, *Single-Atom Resolved Fluorescence Imaging of an Atomic Mott Insulator*, Nature **467**, 68–72 (2010) 10.1038/nature09378 (cit. on pp. 31, 34, 109, 115).

[112] A. Omran, M. Boll, T. A. Hilker, K. Kleinlein, G. Salomon, I. Bloch and C. Gross, *Microscopic Observation of Pauli Blocking in Degenerate Fermionic Lattice Gases*, Phys. Rev. Lett. **115**, 263001 (2015) 10.1103/PhysRevLett.115.263001 (cit. on pp. 31, 109, 115).

[113] W. H. Lowdermilk and D. Milam, *Graded-index Antireflection Surfaces for High-power Laser Applications*, Appl. Phys. Lett. **36**, 891–893 (1980) 10.1063/1.91373 (cit. on p. 33).

[114] M. F. Parsons, *Probing the Hubbard Model With Single-Site Resolution*, PhD thesis, Harvard University (2016) (cit. on p. 34).

[115] T. Hilker, *Spin Resolved Microscopy of Strongly Correlated Fermionic Many-Body States*, PhD thesis, Ludwig-Maximilians-Universität München (2017) (cit. on pp. 34, 115).

[116] BIPM, *Le Système International d'unités / The International System of Units*, Ninth Edition (2019) (cit. on p. 38).

[117] M. D’Auria, W. J. Otter, J. Hazell, B. T. W. Gillatt, C. Long-Collins, N. M. Ridler and S. Lucyszyn, *3-D Printed Metal-Pipe Rectangular Waveguides*, IEEE Trans. Compon. Packag. Manuf. Technol. **5**, 1339–1349 (2015) 10.1109/TCPMT.2015.2462130 (cit. on p. 41).

[118] M. I. M. Ghazali, K. Y. Park, V. Gjokaj, A. Kaur and P. Chahal, *3D Printed Metalized Plastic Waveguides for Microwave Components*, International Symposium on Microelectronics **2017**, 000078–000082 (2017) 10.4071/isom-2017-TP33_096 (cit. on p. 41).

[119] M. Dionigi, C. Tomassoni, G. Venanzoni and R. Sorrentino, *Simple High-Performance Metal-Plating Procedure for Stereolithographically 3-D-Printed Waveguide Components*, IEEE Microw. Wirel. Compon. Lett. **27**, 953–955 (2017) 10.1109/LMWC.2017.2750090 (cit. on p. 41).

[120] C. Robens, S. Brakhane, W. Alt, D. Meschede, J. Zopes and A. Alberti, *Fast, High-Precision Optical Polarization Synthesizer for Ultracold-Atom Experiments*, Phys. Rev. Applied **9**, 034016 (2018) 10.1103/PhysRevApplied.9.034016 (cit. on p. 44).

[121] V. Weyerer, *A Polarization Synthesizer for Sublattice-dependent Imaging of Ultracold Atoms in a Honeycomb Lattice*, MA thesis, Ludwig-Maximilians-Universität München (2021) (cit. on p. 44).

[122] M.-D. Li, M.-D. Li, W. Lin, W. Lin, A. Luo, A. Luo, W.-Y. Zhang, W.-Y. Zhang, H. Sun, H. Sun, B. Xiao, B. Xiao, Y.-G. Zheng, Y.-G. Zheng, Z.-S. Yuan, Z.-S. Yuan, J.-W. Pan and J.-W. Pan, *High-powered Optical Superlattice with Robust Phase Stability for Quantum Gas Microscopy*, Opt. Express, OE **29**, 13876–13886 (2021) 10.1364/OE.423776 (cit. on p. 44).

[123] A. Keshet, *A Next-generation Apparatus for Lithium Optical Lattice Experiments*, PhD thesis, Massachusetts Institute of Technology (2012) (cit. on p. 49).

[124] P. T. Starkey, C. J. Billington, S. P. Johnstone, M. Jasperse, K. Helmerson, L. D. Turner and R. P. Anderson, *A Scripted Control System for Autonomous Hardware-timed Experiments*, Rev. Sci. Instrum. **84**, 085111 (2013) 10.1063/1.4817213 (cit. on p. 49).

[125] G. Kasprowicz, P. Kulik, M. Gaska, T. Przywozki, K. Pozniak, J. Jarosinski, J. W. Britton, J. W. Britton, T. Harty, C. Balance, W. Zhang, D. Nadlinger, D. Slichter, D. Allcock, S. Bourdeauducq, R. Jördens and K. Pozniak, *ARTIQ and Sinara: Open Software and Hardware Stacks for Quantum Physics*, in OSA Quantum 20 Conf. 2020 Pap. QTu8B14 (14th Sept. 2020), QTu8B.14, 10.1364/QUANTUM.2020.QTu8B.14 (cit. on p. 49).

[126] P. T. Starkey, *A Software Framework for Control and Automation of Precisely Timed Experiments*, PhD thesis, Monash University (2019) (cit. on p. 49).

[127] *misu, high-speed physical quantities and dimensions in Python*, https://github.com/cjrh/misu (cit. on p. 50).

[128] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner and M. D. Lukin, *Atom-by-atom Assembly of Defect-free One-dimensional Cold Atom Arrays*, Science **354**, 1024–1027 (2016) 10.1126/science.aah3752 (cit. on p. 50).

[129] *ZeroMQ, an open-source universal messaging library*, https://zeromq.org/ (cit. on p. 51).

[130] *YAML Ain’t Markup Language (YAML™)*, version 1.2.2, (1st Oct. 2021) https://yaml.org/spec/1.2.2/ (cit. on p. 52).

[131] P. Zeeman, *The Effect of Magnetisation on the Nature of Light Emitted by a Substance*, Nature **55**, 347–347 (1897) 10.1038/055347a0 (cit. on p. 57).

[132] A. Corney, *Atomic and Laser Spectroscopy*, Oxford Classics Series (Oxford Univ. Press, 2006) (cit. on p. 57).

[133] C.-L. Hung, X. Zhang, N. Gemelke and C. Chin, *Fast, Runaway Evaporative Cooling to Bose-Einstein Condensation in Optical Traps*, Phys. Rev. A **78**, 10.1103/PhysRevA.78.011604 (2008) 10.1103/PhysRevA.78.011604 (cit. on pp. 58, 61, 93).

[134] B. Merkel, K. Thirumalai, J. E. Tarlton, V. M. Schäfer, C. J. Ballance, T. P. Harty and D. M. Lucas, *Magnetic Field Stabilization System for Atomic Physics Experiments*, Review of Scientific Instruments **90**, 044702 (2019) 10.1063/1.5080093 (cit. on pp. 65, 71, 72, 87).

[135] F.-R. G. H. Winkelmann, *Optical Plane Selection in a Dipole Trap*, PhD thesis, Friedrich-Wilhelms-Universität Bonn (2021) (cit. on p. 65).

[136] H. von Raven, *Direct Digital Synthesis and Integrated Experiment Control for a Novel Lattice Setup*, MA thesis, Ludwig-Maximilians-Universität München (2016) (cit. on p. 67).

[137] *Design files for a AD9910 based DDS board*, https://gitlab.physik.uni-muenchen.de/cs/camdds/pcb-rework (cit. on p. 67).

[138] *Firmware source code for the PSU controller*, https://gitlab.physik.uni-muenchen.de/cs/highfinesse-psu-spi-controller-firmware (cit. on p. 68).

[139] *Design files for the PSU controller PCB*, https://gitlab.physik.uni-muenchen.de/cs/highfinesse-psu-spi-controller-pcb (cit. on p. 68).

[140] *Design files for the current direction box*, https://gitlab.physik.uni-muenchen.de/cs/mosfet-driver-box (cit. on p. 70).

[141] *Design files for the PCBs of the PSU select box*, https://gitlab.physik.uni-muenchen.de/cs/mosfet-matrix-pcb (cit. on p. 71).

[142] *Source code for the control software of the PSU select box*, https://gitlab.physik.uni-muenchen.de/cs/mosfet-gpio-controller (cit. on p. 71).

[143] A. Smith, B. E. Anderson, S. Chaudhury and P. S. Jessen, *Three-Axis Measurement and Cancellation of Background Magnetic Fields to Less than 50 μG in a Cold Atom Experiment*, J. Phys. B At. Mol. Opt. Phys. **44**, 205002 (2011) 10.1088/0953-4075/44/20/205002 (cit. on pp. 71, 72).

[144] A. Öttl, S. Ritter, M. Köhl and T. Esslinger, *Hybrid Apparatus for Bose-Einstein Condensation and Cavity Quantum Electrodynamics: Single Atom Detection in Quantum Degenerate Gases*, Rev. Sci. Instrum. **77**, 063118 (2006) 10.1063/1.2216907 (cit. on p. 72).

[145] T. Ruster, C. T. Schmiegelow, H. Kaufmann, C. Warschburger, F. Schmidt-Kaler and U. G. Poschinger, *A Long-lived Zeeman Trapped-ion Qubit*, Appl. Phys. B **122**, 254 (2016) 10.1007/s00340-016-6527-4 (cit. on p. 72).

[146] X.-T. Xu, Z.-Y. Wang, R.-H. Jiao, C.-R. Yi, W. Sun and S. Chen, *Ultra-low Noise Magnetic Field for Quantum Gases*, Rev. Sci. Instrum. **90**, 054708 (2019) 10.1063/1.5087957 (cit. on p. 72).

[147] C. J. Dedman, R. G. Dall, L. J. Byron and A. G. Truscott, *Active Cancellation of Stray Magnetic Fields in a Bose-Einstein Condensation Experiment*, Review of Scientific Instruments **78**, 024703 (2007) 10.1063/1.2472600 (cit. on p. 72).

[148] H. Lüschen, *An Interpolating Magnetic Field Stabilization for Cold Gases Experiments*, MA thesis, Ludwig-Maximilians-Universität München (2013) (cit. on p. 72).

[149] F. Schmidt-Kaler, S. Gulde, M. Riebe, T. Deuschle, A. Kreuter, G. Lancaster, C. Becher, J. Eschner, H. H. ffner and R. Blatt, *The Coherence of Qubits Based on Single Ca$^+$ Ions*, J. Phys. B: At. Mol. Opt. Phys. **36**, 623–636 (2003) 10.1088/0953-4075/36/3/319 (cit. on p. 72).

[150] G. Breit and I. I. Rabi, *Measurement of Nuclear Spin*, Phys. Rev. **38**, 2082–2083 (1931) 10.1103/PhysRev.38.2082.2 (cit. on p. 80).

[151] S. Brouard and J. Plata, *Internal-state Dephasing of Trapped Ions*, Phys. Rev. A **68**, 012311 (2003) 10.1103/PhysRevA.68.012311 (cit. on pp. 83, 87).

[152] S. Kuhr, W. Alt, D. Schrader, I. Dotsenko, Y. Miroshnychenko, A. Rauschenbeutel and D. Meschede, *Analysis of Dephasing Mechanisms in a Standing-wave Dipole Trap*, Phys. Rev. A **72**, 023406 (2005) 10.1103/PhysRevA.72.023406 (cit. on p. 85).

[153] F. Kindermann, *Towards Coherent Control of a Single Cs Atom in an Ultracold Rb Cloud*, MA thesis, Friedrich-Wilhelms-Universität Bonn (2011) (cit. on p. 85).

[154] N. F. Ramsey, *A Molecular Beam Resonance Method with Separated Oscillating Fields*, Phys. Rev. **78**, 695–699 (1950) 10.1103/PhysRev.78.695 (cit. on p. 86).

[155] P. J. J. O'Malley, J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Korotkov, A. N. Cleland and J. M. Martinis, *Qubit Metrology of Ultralow Phase Noise Using Randomized Benchmarking*, Phys. Rev. Applied **3**, 044009 (2015) 10.1103/PhysRevApplied.3.044009 (cit. on p. 87).

[156] E. L. Hahn, *Spin Echoes*, Phys. Rev. **80**, 580–594 (1950) 10.1103/PhysRev.80.580 (cit. on p. 91).

[157] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman and E. A. Cornell, *Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor*, Science **269**, 198–201 (1995) 10.1126/science.269.5221.198 (cit. on p. 93).

[158] C. C. Bradley, C. A. Sackett, J. J. Tollett and R. G. Hulet, *Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions*, Phys. Rev. Lett. **75**, 1687–1690 (1995) 10.1103/PhysRevLett.75.1687 (cit. on p. 93).

[159] K. B. Davis, M. .-. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn and W. Ketterle, *Bose-Einstein Condensation in a Gas of Sodium Atoms*, Phys. Rev. Lett. **75**, 3969–3973 (1995) 10.1103/PhysRevLett.75.3969 (cit. on p. 93).

[160] T. Weber, *Bose-Einstein Condensation of Optically Trapped Cesium*, PhD thesis, Leopold-Franzens-Universität Innsbruck (2003) (cit. on p. 93).

[161] T. Kraemer, M. Mark, P. Waldburger, J. G. Danzl, C. Chin, B. Engeser, A. D. Lange, K. Pilch, A. Jaakkola, H.-C. Nägerl and R. Grimm, *Evidence for Efimov Quantum States in an Ultracold Gas of Caesium Atoms*, Nature **440**, 315–318 (2006) 10.1038/nature04626 (cit. on p. 93).

[162] C. D. Colquhoun, *A New Apparatus for Experiments with Caesium Bose-Einstein Condensates*, PhD thesis, University of Strathclyde (2019) (cit. on p. 93).

[163] A. Z. Lam, N. Bigagli, C. Warner, W. Yuan, S. Zhang, E. Tiemann, I. Stevenson and S. Will, *High Phase-space Density Gas of NaCs Feshbach Molecules*, Phys. Rev. Research **4**, L022019 (2022) 10.1103/PhysRevResearch.4.L022019 (cit. on p. 93).

[164] C. Cabrera, *Quantum Liquid Droplets in a Mixture of Bose-Einstein Condensates*, PhD thesis, ICFO – The Institute of Photonic Sciences Universidad Politécnica de Cataluña (2018) (cit. on p. 96).

[165] B. Gadway, D. Pertot, R. Reimann, M. G. Cohen and D. Schneble, *Analysis of Kapitza-Dirac Diffraction Patterns beyond the Raman-Nath Regime*, Opt. Express, OE **17**, 19173–19180 (2009) 10.1364/OE.17.019173 (cit. on p. 98).

[166] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner and P. Zoller, *Cold Bosonic Atoms in Optical Lattices*, Phys. Rev. Lett. **81**, 3108–3111 (1998) 10.1103/PhysRevLett.81.3108 (cit. on p. 101).

[167] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch and I. Bloch, *Quantum Phase Transition from a Superfluid to a Mott Insulator in a Gas of Ultracold Atoms*, Nature **415**, 39–44 (2002) 10.1038/415039a (cit. on p. 101).

[168] I. Bloch and M. Greiner, *Exploring Quantum Matter with Ultracold Atoms in Optical Lattices*, in *Advances In Atomic, Molecular, and Optical Physics*, Vol. 52, edited by P. R. Berman and C. C. Lin (Academic Press, 1st Jan. 2005), pp. 1–47, 10.1016/S1049-250X(05)52001-9 (cit. on p. 102).

[169] C. Weitenberg, *Single-Atom Resolved Imaging and Manipulation in an Atomic Mott Insulator*, PhD thesis, Ludwig-Maximilians-Universität München (2011) (cit. on pp. 104, 113).

[170] M. J. Rice and E. J. Mele, *Elementary Excitations of a Linearly Conjugated Diatomic Polymer*, Phys. Rev. Lett. **49**, 1455–1459 (1982) 10.1103/PhysRevLett.49.1455 (cit. on p. 105).

[171] M. Lohse, C. Schweizer, H. M. Price, O. Zilberberg and I. Bloch, *Exploring 4D Quantum Hall Physics with a 2D Topological Charge Pump*, Nature **553**, 55–58 (2018) 10.1038/nature25000 (cit. on p. 106).

[172] A. S. Sørensen, E. Demler and M. D. Lukin, *Fractional Quantum Hall States of Atoms in Optical Lattices*, Phys. Rev. Lett. **94**, 086803 (2005) 10.1103/PhysRevLett.94.086803 (cit. on p. 107).

[173] S. A. Parameswaran, R. Roy and S. L. Sondhi, *Fractional Quantum Hall Physics in Topological Flat Bands*, Comptes Rendus Physique, Topological Insulators / Isolants Topologiques **14**, 816–839 (2013) 10.1016/j.crhy.2013.04.003 (cit. on p. 107).

[174] J. Bibo, I. Lovas, Y. You, F. Grusdt and F. Pollmann, *Fractional Corner Charges in a 2D Super-lattice Bose-Hubbard Model*, Phys. Rev. B **102**, 041126 (2020) 10.1103/PhysRevB.102.041126 (cit. on pp. 107, 121).

[175] E. Haller, J. Hudson, A. Kelly, D. A. Cotta, B. Peaudecerf, G. D. Bruce and S. Kuhr, *Single-atom Imaging of Fermions in a Quantum-gas Microscope*, Nature Phys **11**, 738–742 (2015) 10.1038/nphys3403 (cit. on p. 109).

[176] G. J. A. Edge, R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky and J. H. Thywissen, *Imaging and Addressing of Individual Fermionic Atoms in an Optical Lattice*, Phys. Rev. A **92**, 063406 (2015) 10.1103/PhysRevA.92.063406 (cit. on p. 109).

[177] W. H. Richardson, *Bayesian-Based Iterative Method of Image Restoration*, J. Opt. Soc. Am., JOSA **62**, 55–59 (1972) 10.1364/JOSA.62.000055 (cit. on p. 115).

[178] L. B. Lucy, *An Iterative Technique for the Rectification of Observed Distributions*, Astron. J. **79**, 745 (1974) 10.1086/111605 (cit. on p. 115).

[179] S. Häfele, *Deep Learning for Quantum Gas Microscopy of Cesium Atoms in Optical Lattices – Below the Resolution Limit*, MA thesis, Ludwig-Maximilians-Universität München (2022) (cit. on p. 116).

[180] P. Schauss, *High-resolution imaging of ordering in Rydberg many-body systems*, PhD thesis, Ludwig-Maximilians-Universität München (2015) (cit. on p. 117).

[181] M. Erdmann, J. Glombitza, G. Kasieczka and U. Klemradt, *Deep Learning For Physics Research* (World Scientific Publishing Company, Singapore, SINGAPORE, 2021) (cit. on p. 117).

[182] D. Greif, M. F. Parsons, A. Mazurenko, C. S. Chiu, S. Blatt, F. Huber, G. Ji and M. Greiner, *Site-resolved Imaging of a Fermionic Mott Insulator*, Science **351**, 953–957 (2016) 10.1126/science.aad9041 (cit. on p. 119).

[183] S. Taie, H. Ozawa, T. Ichinose, T. Nishio, S. Nakajima and Y. Takahashi, *Coherent Driving and Freezing of Bosonic Matter Wave in an Optical Lieb Lattice*, Sci. Adv. **1**, e1500854 (2015) 10.1126/sciadv.1500854 (cit. on p. 122).

[184] C. Repellin, J. Léonard and N. Goldman, *Fractional Chern Insulators of Few Bosons in a Box: Hall Plateaus from Center-of-mass Drifts and Density Profiles*, Phys. Rev. A **102**, 063316 (2020) 10.1103/PhysRevA.102.063316 (cit. on p. 122).