Defect-Free Atom Arrays of ^{88}Sr in Optical Tweezers

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

von
Felix Rönchen
aus
Engelskirchen

Bonn, August 2024

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der Rheinischen Friedrich-Wilhelms-Universität Bonn

Gutachter/Betreuer:	Prof. Dr. Michael Köhl
Gutachter:	Prof. Dr. Stefan Linden
Tag der Promotion:	29.10.2024
Erscheinungsjahr:	2024

Abstract

Neutral atoms trapped in optical microtraps, so called *optical tweezers*, have emerged as a platform for controlling large many-particle quantum systems, with applications in many-body physics, metrology, quantum information processing and cavity quantum electrodynamics. A key requirement for building a programmable quantum simulation platform is to gain control over the initial state of the system. It is essential to eliminate uncertainties in the initialization process, especially caused by variations in the number of stored atoms but also by uncertainties associated with the electronic and motional degrees of freedom of a trapped atom itself.

In this work, we report on the design and characterization of a new optical tweezer setup, aimed at preparing, detecting, and manipulating single strontium atoms cooled to their motional ground state. We design a tweezer system around a high-NA microscope objective (NA = 0.7) with different tweezer wavelengths to exploit their respective advantages. We load the tweezers from a MOT and by employing light-assisted collisions, we create highly sub-Poissonian atom number distributions, where a trap is occupied by a single atom at most. In tweezers at 515 nm, we cool atoms using resolved sideband cooling, where we observe a three-dimensional motional ground state fraction of 95^{+2}_{-10} % at a tweezer aspect ratio of 5.1 ± 0.1, improving on existing experimental implementations. Further, we use a sisyphus cooling process to cool atoms in tweezers at 532 nm and 813 nm, where we also confirm temperatures close to the motional ground state. At 813 nm, we confirm array homogeneities on the 1 %-level, measured with up to 60 ultracold trapped atoms. The characterization of the sisyphus cooling process in tweezers at 532 nm is, to our current knowledge, the first time that strontium atoms have been investigated in high NA tweezers at this wavelength. They confirm that a tweezer wavelength of 532 nm is a viable candidate for future research, as commonly available lasers and optical elements can reduce complexity and cost of the experimental setup.

As described above, we can reliably prepare single atoms in their electronic and motional ground state in a tweezer array. To address the remaining initial entropy of the atom array, namely the uncertainty associated with partially filled arrays, we rearrange atoms in the underlying array using a separate, spatially tunable optical tweezer. We develop a control-theory optimal trajectory model for moving an atom between sites, based on the kinematic parameters of the motion. We then present the assembly of defect-free arrays of 16 tweezers, where we determine a single move success probability ("move fidelity") of (97.8 ± 2.2) %, on par with other state-of-the-art implementations.

With this work we lay the foundation of a programmable quantum simulation platform, and in the future, the new experiment can be used to create large many-particle states, enabling exciting studies at the frontiers of experimental quantum simulation research.

Parts of this work have been published in:

- J. W. Lis, A. Senoo, W. F. McGrew, F. Rönchen, A. Jenkins, and A. M. Kaufman, *Midcircuit Operations Using the omg Architecture in Neutral Atom Arrays*, Phys. Rev. X 13, 041035 (2023)

Contents

CHAPTER 1

Introduction

The study of light and matter sparked the development of early quantum theories at the end of the 20th century, and it opened the door to an entirely new world that is counterintuitive to the classically trained human mind. An indispensable tool for exploring and developing new fundamental concepts were *thought experiments*. These thought experiments allowed scientists to investigate single quantum systems in isolation and mathematically study the effects of newly theorized quantum rules. But even up until the 1950s, Erwin Schrödinger argued that manipulating single atoms would forever be impossible [1]. This was a bleak prediction, as atoms and molecules of the same species are, to our current knowledge, indistinguishable, allowing us to predict and replicate the behavior of these systems with high precision. This makes them ideal candidates for ultra-precise measurements to test our hypotheses on.

Fortunately, Schrödinger was wrong about this specific aspect, and the field of experimental physics was on the verge of a radical transformation. This transformation came with the invention of the laser in the 1960s, a prime example for the symbiosis between fundamental research and technological progress. Theoretically proposed during the early development of the new quantum theories in the 1910s, the laser led to one of the most profound technological advancements of the 20th century. Access to the (arguably) most important modern precision measurement tool enabled the detailed study of atomic and molecular structures with unprecedented accuracy, previously inaccessible for conventional, less coherent and spectrally impure, light sources.

Scientists quickly recognized the enormous potential of lasers for the study of atomic and molecular systems which led to the idea of cooling atomic samples solely by scattering laser light. Proposed by Hänsch and Schawlow [2] and independently Wineland and Dehmelt [3] in 1975, it took only three years for Wineland and Dehmelt to realize it experimentally in 1978. Laser cooling, whether in the form of Doppler cooling or one of its more advanced successors, paved the way to the engineering of artificial quantum systems. It became possible to store, manipulate and systematically investigate isolated quantum systems, making the early thought experiments come full circle, as experiments were now able to test and advance our understanding of the quantum world [4].

In the 1980s, the concept of building quantum simulators to investigate the behavior of quantum systems became more refined [5], and these ideas have been further concretized since [6–8]. A quantum simulator, compared to a classical one, exploits the quantum behavior of the system, including effects like superposition and entanglement, to simulate the dynamics of a system or carry out computations. Similar to the classical case, a simulation can be of digital or analog nature, akin to understanding aerodynamics either through digital simulation or in a purpose-built wind tunnel [9]. A long term goal

in the field of quantum simulation is the development of a fully digital quantum simulator, a *quantum computer*, where the state of a system is encoded in quantum bits ("qubits") which are manipulated using elementary quantum gates to create a circuit-based quantum computer. A digital quantum computer is expected to be able to tackle demanding computing tasks, with applications in material science and chemistry, logistics and optimization, encryption and information processing but also artificial intelligence, on a level that is inaccessible with classical super-computers. A key milestone is the point of *quantum advantage* or *quantum supremacy*, where computational capabilities of quantum devices surpass that of traditional supercomputers [9]. However, leveraging quantum effects to solve intricate problems is a highly non-trivial task [10], and consequently, recent efforts primarily focus on solving artificially tailored problems, serving as proof-of-concept demonstrations [11]. Furthermore, large-scale quantum computations require fault-tolerant qubits and error-corrections schemes, which have only been implemented in few-qubit systems, impractical for realizing a general digital quantum computer. The current era is therefore referred to as noisy intermediate-scale quantum (NISQ) computing [12, 13].

A second approach to the topic of quantum simulation are *analog* quantum simulators. This approach involves engineering a specific system to study, resembling the wind tunnel-based study from the analogy above. The object under investigation in this case might be a subsystem of a larger entity, or it might even be a system merely mimicking the behavior of another system. The analog approach offers less flexibility in some areas, for example the tunability of certain parameters, but it has a decisive advantage compared to the digital case: The system itself implements the dynamics of, and the interactions between, its constituents. An example of an analog quantum simulation is the realization of the *Hubbard*-model for studying magnetism in solids, by investigating the effect with single fermions trapped in periodic optical potentials [14]. In addition to neutral atoms [15–18], various other experimental platforms for analog quantum simulators are currently under investigation, including superconducting systems [19, 20], trapped ions [21–23] and photons [24, 25].

A short history of neutral atoms in optical tweezers

The first use of optical gradient forces to trap micrometer sized particles is attributed to Arthur Ashkin[1] and his colleagues at Bell Labs in 1970 [26]. Since then, biologists used optical tweezers in the form of tightly focussed laser beams to trap viruses and bacteria [27] to investigate cell structures and dynamics on the molecular scale. In physics, the idea to use strong, off-resonant laser light to spatially confine atoms was also not new [28, 29], but the required technology to create three-dimensional potentials for single atoms was simply not available. Compared to other atom traps, like magneto optical traps (MOTs) or radiofrequency-traps for ions which work at up to room temperature, optical tweezer are comparably shallow, with depths in the order of one millikelvin. Loading atoms into such shallow potentials needs sophisticated pre-cooling protocols, and to create and maintain such *ultracold* conditions requires a strong decoupling from the environment that only an ultrahigh vacuum (UHV) system can achieve. Combining these experimental constraints with the demands on laser and optical technologies necessary for creating (sub-)micrometer sized potentials took experimentalists time to master. It was therefore only at the beginning of the 2000s that physicists began to investigate single neutral atoms in optical tweezers [30, 31]. Initial research was focused on achieving the first milestone of a quantum simulation platform: Gaining precise control over the initial state. The latter includes the knowledge of the atom number per trap but also control over the internal and external (motional) degrees of freedom of a

[1] Ashkin would go on to receive the 2018 Nobel Prize in Physics for this discovery.

trapped atom. It was demonstrated that the tightly confining potential can be used to induce collisions between atoms which in turn leads to a single atom inside a trap at most [32]. Further important early work includes the development of detection schemes [33], thermometry [34] and later also cooling to the quantum-mechanical ground state of the trap [35, 36].

Extending the system from one to many optical tweezers by using multiple laser beams, established optical tweezers as a promising quantum simulation platform. As compared to other neutral-atom based systems, like quantum gas microscopes [16], dynamically tunable traps can be used to rearrange atoms between tweezers and assemble systems in a *bottom-up* approach [37–39], achieving arbitrarily low entropy [40]. Since then, neutral atoms in optical tweezers have been used for simulating quantum many-body systems [41, 42], for research in quantum information processing [43, 44] and cavity quantum electrodynamics (CQED) [45].

The choice of atom species initially fell on alkali elements like rubidium which had been investigated for decades in other cold-atom experiments and were therefore well known. However, starting from the mid 2010s, experimenters began to also investigate alkaline-earth(-like) elements, like strontium [46–48] and ytterbium [49], in optical tweezers. Alkaline-earth or alkaline-earth-like elements exhibit a rich level structure which opens up new possibilities for fast and efficient cooling and spectroscopy schemes. The interaction of two valence electrons results in a triplet-manifold which connects to the electronic ground-state by one or more dipole-forbidden transitions. These narrow transitions allow for efficient multi-state cooling protocols with achievable temperatures in the region of microkelvin in MOTs, from which atoms can be directly transferred to the tweezers. In addition to the low temperature, these thermal laser-cooling techniques can also cool atoms to the motional ground state of the potential in tens of milliseconds, which is orders of magnitudes faster than methods like evaporative cooling, enabling fast experiment cycle times [50]. In recent years, alkaline-earth(-like) atoms in optical tweezers have been used for research in many-body physics [51, 52], metrology [53–55] and quantum information processing [56, 57] and have thus matured to a state-of-the-art implementation of an analog quantum simulation platform.

Thesis outline

Optical tweezers are a versatile experimental platform for studying various physical phenomena, and the goal of this work was to lay the foundation for a programmable quantum simulator. This requires the ability to prepare and detect single atoms in arrays of optical tweezers, along with single-site addressing capabilities to move atoms between different tweezers as needed, thus reducing the uncertainty associated with the probabilistic loading of tweezers. Ultimately, we aim to create defect-free arrays of single atoms cooled to their motional ground state.

As such a large endeavor cannot be undertaken by a single person within a reasonable time frame, this project has been a team effort. Two key colleagues are Jonas Schmitz and Tobias Kree. Jonas Schmitz designed the vacuum system, developed the experimental control system, and set up the initial cooling stages, detailed in his thesis [58]. The latter also covers aspects like trapping different isotopes and the magnetic trapping of ^{88}Sr that are not directly relevant for this work. Tobias Kree built the 461 nm and 689 nm laser systems during his Master's thesis [59] before joining the project as a PhD student. His work includes detailed descriptions of the electronics and stability measurements for the 689 nm laser. The relevant aspects of their works are summarized in Chapter 3.

My work focused on creating static and dynamic optical potentials as described in Chapter 4, characterizing trapped single atoms in these potentials as detailed in Chapter 5 and assembling defect-free arrays outlined in Chapter 6. To gain insight in these topics, we will discuss the considerations for designing the experimental setup in Chapter 2. The results will be summarized in Chapter 7, together with a discussion of potential future developments.

CHAPTER 2

Experimental considerations for building an optical tweezer experiment

An experiment aimed at the manipulation and detection of single atoms, with control over all degrees of freedom, is a complex machine requiring the precise coordination of many components. This chapter seeks to unravel this complexity by examining the different elements, through the example of a single atom trapped in an optical potential. Subsequent chapters will delve into these components in greater detail. We begin by introducing the light shift, that gives rise to the trapping effect in strongly focused laser beams. Next, we will discuss how to prepare a single atom for further measurements and control its motion within the trap. Following this, we will explore methods for detecting a single atom with sub-micrometer resolution. Finally, we conclude the chapter by presenting a method for achieving single-particle control in arrays of trapped atoms.

A graphical representation of the all relevant components is shown at the end of the chapter in Figure 2.1.

2.1 The dipole force

The dipole force is the result of the coupling of an external oscillating electric field to an induced dipole. This coupling shifts the level structure of a system which is the reason for the effect also being referred to as *light shift*. To quantify the effect, we can model the atom as a driven two-level system with energy spacing $\hbar\omega_0$. The Hamiltonian $\hat{\mathcal{H}}$ describes the interaction between a classical oscillating electric field $\vec{E}$ and the induced dipole moment $\vec{\mu}_e$ [60]:

$$\hat{\mathcal{H}} = -\vec{E}\left(\vec{r},t\right)\cdot\hat{\vec{\mu}}_e = -e\vec{E}\left(\vec{r},t\right)\cdot\hat{\vec{r}} \tag{2.1}$$

$$\vec{E}\left(\vec{r},t\right) = E_0\epsilon\cos\left(kz-\omega_l t\right), \tag{2.2}$$

where we chose the z-axis as the propagation direction of the light field. $\vec{\epsilon} = \left(\epsilon_x, \epsilon_y, \epsilon_z = 0\right)$ is the polarization vector of the light field. This interaction leads to off-diagonal terms in the Hamiltonian and the new eigenstates- and energies are different from the unperturbed ones. It is common to refer to the new states as *dressed*-states, as the light field dresses the energy levels of the atom. We define the detuning $\delta = \omega_l - \omega_0$ as the difference between the driving frequency ω_l and the resonance frequency of

the system ω_0. In order to solve for the new eigenenergies of the system and finally the dipole force, we make use two approximations. The first one is the rotating-wave-approximation (RWA) where we neglect terms of order $1/\omega_l$ compared to $1/\delta$. Secondly, we assume that the spatial variation of $\vec{E}$ over the size of the atom can be neglected, known as the dipole-approximation. Transforming into a frame rotating at ω_0, the Hamiltonian in the unperturbed basis reads:

$$\hat{\mathcal{H}} = \frac{\hbar}{2}\begin{bmatrix} -2\delta & \Omega \\ \Omega & 0 \end{bmatrix} \tag{2.3}$$

leading to the eigenenergies:

$$E_{+,-} = \frac{\hbar}{2}\left(-\delta \mp \Omega'\right) \tag{2.4}$$

where $\Omega' = \sqrt{\Omega^2 + \delta^2}$ is the generalized Rabi frequency. With $\Omega \ll |\delta|$ we arrive at:

$$\Delta E_\mathrm{e} \approx \Delta E_+ \approx \frac{\hbar\Omega^2}{4\delta} \tag{2.5}$$

$$\Delta E_\mathrm{g} \approx \Delta E_- \approx -\frac{\hbar\Omega^2}{4\delta}. \tag{2.6}$$

With the intensity of the light field $I \propto \vec{E}^2 \propto \Omega^2$, the resulting force for the ground state is given by:

$$\vec{F}_\mathrm{dip} = -\vec{\nabla}\left(\Delta E_\mathrm{g}\right) \simeq \frac{\hbar}{4\delta}\vec{\nabla}\left(\Omega\left(\vec{r}\right)^2\right) \tag{2.7}$$

$$\simeq \frac{\hbar}{4\delta}\vec{\nabla} I\left(\vec{r}\right), \tag{2.8}$$

which is proportional to $1/\delta$. The dipole force $\vec{F}_\mathrm{dip}$ competes against the scattering force $\vec{F}_\mathrm{scatt}$, and we have to compare the two to determine where the $\vec{F}_\mathrm{dip}$ dominates.

The scattering force $\vec{F}_\mathrm{scatt}$ is given by [60]:

$$\vec{F}_\mathrm{scatt} = \frac{\hbar k s_0 \gamma/2}{1 + s_0 + (2\delta/\gamma)^2}, \tag{2.9}$$

for a transition with linewidth γ, where $s_0 = I/I_s$ is the saturation parameter with the saturation intensity I_s. As $\vec{F}_\mathrm{scatt}$ is proportional to $1/\delta^2$ and constant for $s_0 \gg 1$, it is possible to trap atoms for a sufficiently large detuning δ and intensity I. In general, one would like the atoms to stay trapped in the electronic ground state, as they spend more time in the ground state than the excited state. To create a confining potential for the electronic ground state, dipole traps therefore commonly operate at a negative detuning compared to the atomic resonance, i.e. by using a laser frequency smaller than the optical resonance frequency of the system (red-detuned). This results in ground-state atoms being pulled towards regions of large intensity ("high-field seekers"), whereas the excited state is anti-trapped ("low-field seekers"). It is also possible to use blue-detuned traps by sculpturing e.g. a hollow-beam potential [61].

In order to create the optical potential, a focused laser beam is used, usually a Gaussian beam in the

TEM_{00} mode. We can describe the intensity profile for such a beam traveling in z-direction by:

$$I(r,z) = I_0 \left(\frac{w_0}{w(z)}\right)^2 \cdot \exp\left(-2r^2/w^2(z)\right) \tag{2.10}$$

where $w(z) = w_0\sqrt{1+\left(z/z_\mathrm{R}\right)^2}$ is the $1/e^2$-radius at axial position z, and w_0 is the minimum waist at $z = 0$. The Rayleigh range z_R is given by $z_\mathrm{R} = \pi w_0^2/\lambda$, where λ is the wavelength of the trapping light. z_R gives a length-scale for the axial direction as it determines the divergence of the beam. Around the focal spot, the potential can be approximated as harmonic and trap frequencies for the radial direction ω_r and axial direction ω_a can be defined. A detailed treatment of the harmonic approximation and the parameters that influence the confinement, for example finite apertures in the optical setup, are investigated at the beginning of Chapter 4 in Section 4.2.

Historically, dipole traps have been used to confine quantum gases and condensates (see e.g. [62] for an overview) and it is instructive to take a look at the differences. Common parameters for these traps are waists on the order of 10 µm and powers on the order of 1 W. This leads to trap frequencies on the order of 100 Hz and depths of 100 µK. Additionally, the large waist leads to a weak axial confinement, and optical setups therefore often combine multiple laser beams to create a three-dimensional confining potential. The difference to an optical tweezer is that the latter creates such a confinement with a single beam. This requires $O\left(w_0\right) \approx O\left(z_\mathrm{R}\right)$ which in turn requires $w_0 \lesssim \lambda$. Together with the harmonic approximation, a single optical tweezer can then be seen as an idealized three-dimensional harmonic oscillator. A schematic of a single atom trapped in a tightly focused optical potential, together with a harmonic approximation close to the center is shown in Figure 2.1(a). For optical tweezers, trap frequencies are on the order of 100 kHz and it is thus possible to resolve the quantized motional states of a trapped particle optically. A sketch of the full state of the atom, consisting of the the internal (electronic) state and the external (motional) state is shown in Figure 2.1(d).

To achieve $w_0 \lesssim \lambda$, the input beam diameter has to be of similar size as the focal length of the focussing lens (high numerical aperture NA), which imposes strict requirements on the optical setup. We will discuss the demands on the vacuum system, which is necessary to isolated the trapped atom from the environment, in Chapter 3.

2.2 Single atom preparation

Dipole traps with depths around $k_\mathrm{B} \times 1$ mK cannot be loaded with atoms at room temperature directly. Therefore, one or more laser-cooling steps are required, often in the form of a magneto-optical trap (MOT) [60, 63]. The MOT combines magnetic field gradients with a three-dimensional optical molasses to trap and cool atoms. It has become a standard component in modern cold atom experiments, where it is often the first step to capture and cool atoms emerging from an atom oven. The final temperature of a MOT is given by the Doppler-temperature $T_\mathrm{D} = \hbar\gamma/2k_\mathrm{B}$, which depends on the linewidth of the transition γ. For an $s \rightarrow p$ transition in alkali elements with $\gamma \approx 2\pi \times 10$ MHz, T_D is about 300 µK. As this is on the order of the tweezer depth, sub-Doppler protocols like polarization gradient cooling are used to lower the temperature further to reliably load atoms into the trap [35] . For alkaline-earth(-like)

elements, the existence of narrow intercombination lines means that MOTs with temperatures $\lesssim 10\,\mu\mathrm{K}$[1] can be created, from which the tweezer can be loaded directly. We will discuss these initial cooling stages to create MOTs at a few microkelvin at the end of Chapter 3.

Loading an optical tweezer from a cold sample of atoms results in the trap being occupied by more than one atom on average. The atom-number follows a Poissonian-distribution $P(n)$, often with an average around $\langle N_{\mathrm{atom}}\rangle \approx 10$ atoms for a MOT[2] and a corresponding standard deviation of $\sigma_{\mathrm{atom}} = \sqrt{\langle N_{\mathrm{atom}}\rangle}$. Our goal to experiment with single atoms requires us to lower the atom number to the single atom level. A mechanism that can be used to remove excess atoms from the trap, is the effect of light-assisted collisions (LACs). Due to the tightly confining potential, two atoms are likely to be excited to a molecular state which can then leave the trap (see Figure 2.1(b), drawn after [50]). This process leads to a pair-wise removal of atoms, given a light field tuned to a molecular resonance is present. LACs reducing the atom number have been observed in optical lattices [16] and optical tweezers [32] where highly sub-Poissonian atom-number distributions with $\sigma_{\mathrm{atom}} < \sqrt{\langle N_{\mathrm{atom}}\rangle}$ have been reached. More precisely, the LAC-process results in $P(n > 1) \approx 0$ and we are therefore assured to experiment with a single atom per trap at most. The probability of preparing a single atom compared to an empty trap depends on the specific LAC-process, but for this work we can assume $P(0) = 1/2$ and $P(1) = 1 - P(0) = 1/2$. We will demonstrate and verify that we can prepare single atoms in Section 5.3.

Inside the tweezer, different cooling protocols can be used to reach the motional ground state of the trap. These mechanisms depend on the details of the light shift which in turn depends on the trapping wavelength. We motivate the choice of the trapping wavelengths in Section 4.1 where we also expand the model of the light shift from a two level system to a real atom. We investigate cooling of a single atom to the motional ground state in Section 5.4.

2.3 Single particle imaging

While the previous section focussed on the trapping and preparation of a single atom in an optical tweezer, this section discusses the detection. In general, there are two ways to take an image of an object: Either by shining light onto it and detecting a shadow, known as absorption imaging, or by shining in light and detecting the scattered light, known as fluorescence imaging. Although single particles have been detected using absorption imaging [64], the canonical way is to adhere to fluorescence imaging. This type of imaging has been used in trapped ion experiments since the 1980's (see e.g. [65, 66] for an overview of ion trap experiments), and the process can be transferred to neutral atoms as well. To understand the signature we can expect from a single trapped atom, it is instructive to investigate the number and the distribution of scattered photons detected by a camera or a photo diode. When illuminating an atom, it scatters photons with the rate [60]:

$$R = \frac{s_0\gamma/2}{1 + s_0 + (2\delta/\gamma)^2}. \tag{2.11}$$

[1] At this point, also the recoil limit of the transition $T_{\mathrm{R}} = \frac{\hbar^2 k^2}{m k_{\mathrm{B}}}$ with $k = 2\pi/\lambda$ becomes important. For the ${}^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ transition in strontium, the recoil temperature of 500 µK is actually larger than the Doppler temperature of 180 nK (see Table 3.4.)

[2] Assuming a center density of $1 \times 10^{13}\,\mathrm{cm}^{-3}$ and a tweezer volume of $1\,\mu\mathrm{m}^3$. This atom number can be significantly larger when loading from e.g. a degenerate quantum gas.

In a time τ, we thus detect $N = \alpha R \tau$ photons *on average*. Here, α is the detection efficiency of the system, combining the finite solid angle we collect photons from and finite efficiencies of optical elements and detectors. However, the true number of scattered photons follows a Poissonian distribution, and we will detect a varying number of photons between repetitions of the experiment.

We can estimate the number of photons reaching the detector for a $s \rightarrow p$-transition for $s_0 \ll 1$, that is for non-destructive imaging. Non-destructive in this case means, that we can image the atom, without heating it out of the trap. For $s_0 \ll 1$, Equation (2.11) simplifies to $R = \gamma/2s_0$. We can estimate the requirements on the detection system, by considering the total number of scattered photons together with the collection efficiency. Assuming $\gamma = 2\pi \times 10\,\mathrm{MHz}$ and $s_0 = 0.005$ this results in $R \approx 20\,000\,\mathrm{s}^{-1}$. For common detection efficiencies around 5 to 10 %, calculated from the solid angle that is covered by the imaging lens, and a detection efficiency of $\approx 50\,\%$[3], we can expect around 100 photons/100 ms *on average*. This estimation shows, that a single photon counting device is required to detect a single atom.

The specific number of photons for a single repetition of the experiment follows a Poissonian distribution with mean N_{avg} if an atom is present. This value has to be significantly larger than the background signal due to stray light, electronic noise and electronic gain. We will see later, that the background can be approximated to follow a normal distribution when we investigate the imaging process in Section 5.5. The two components of this bimodal distribution can be separated by choosing a threshold at a certain photon number N_{cut} that essentially divides the signal into binary no-atom vs. one-atom events. The accuracy with which we can make this distinction is called the *fidelity* of the imaging system which is of great importance for the single atom addressing presented in Chapter 6. We will make use of the imaging fidelity when we optimize the detection process for our system in Chapter 5. An example of a histogram showing the bimodel distribution is shown in Figure 2.1(c).

2.4 Single site addressing

The previous sections explained different tweezer concepts, like trapping and detection, at the example of a single trap. However, a key strength of the tweezer platform is the ability to create multiple traps, often arranged in a grid or an *array*, in which atoms can further be manipulated dynamically during the experiment. Multiple and tunable traps can be created using devices that alter the wavefront of an incident beam. For this work, we will use two of such devices, one to a create static grid to trap multiple atoms, and a second one to address a single site out of this grid for site-selective manipulation. Both of these devices are discussed in detail in Chapter 4.

There are two main reasons why controlling individual particles out of an array is desirable. The first reason is related to the preparation of the quantum system, particularly to the uncertainty associated with the initial state. By employing light-assisted collisions, we can ensure that each site contains at most one atom. Therefore, populating a single tweezer is akin to a Bernoulli experiment, much like flipping a coin. For multiple tweezers, each with a probability p of being populated, the total number of populated sites in an array of N_{sites} sites follows a binomial distribution with a mean of $\langle N_{\mathrm{sites}} \rangle = N_{\mathrm{sites}}\, p$. To decrease the entropy of this distribution, atoms can be transferred between traps to create a smaller, defect-free region [40] (see Figure 2.1(e)). This process involves (a) identifying populated and empty sites with low error rates and (b) transferring atoms between traps with low loss rates. The second reason involves the ability to locally tune light-matter interactions for site-specific manipulation or readout. This local

[3] This efficiency combines multiple factors, for example transmittance/reflectance from optical elements or the quantum efficiency (QE) of the detection device.

tunability is crucial for the design of quantum simulators, enabling partial read-out and error correction [9, 13, 57] (see Figure 2.1(f)). While this work does not focus on this type of single-site addressing, the rearrangement step inherently includes this capability. We will present the assembly of defect-free regions from larger, partially filled arrays in Chapter 6.

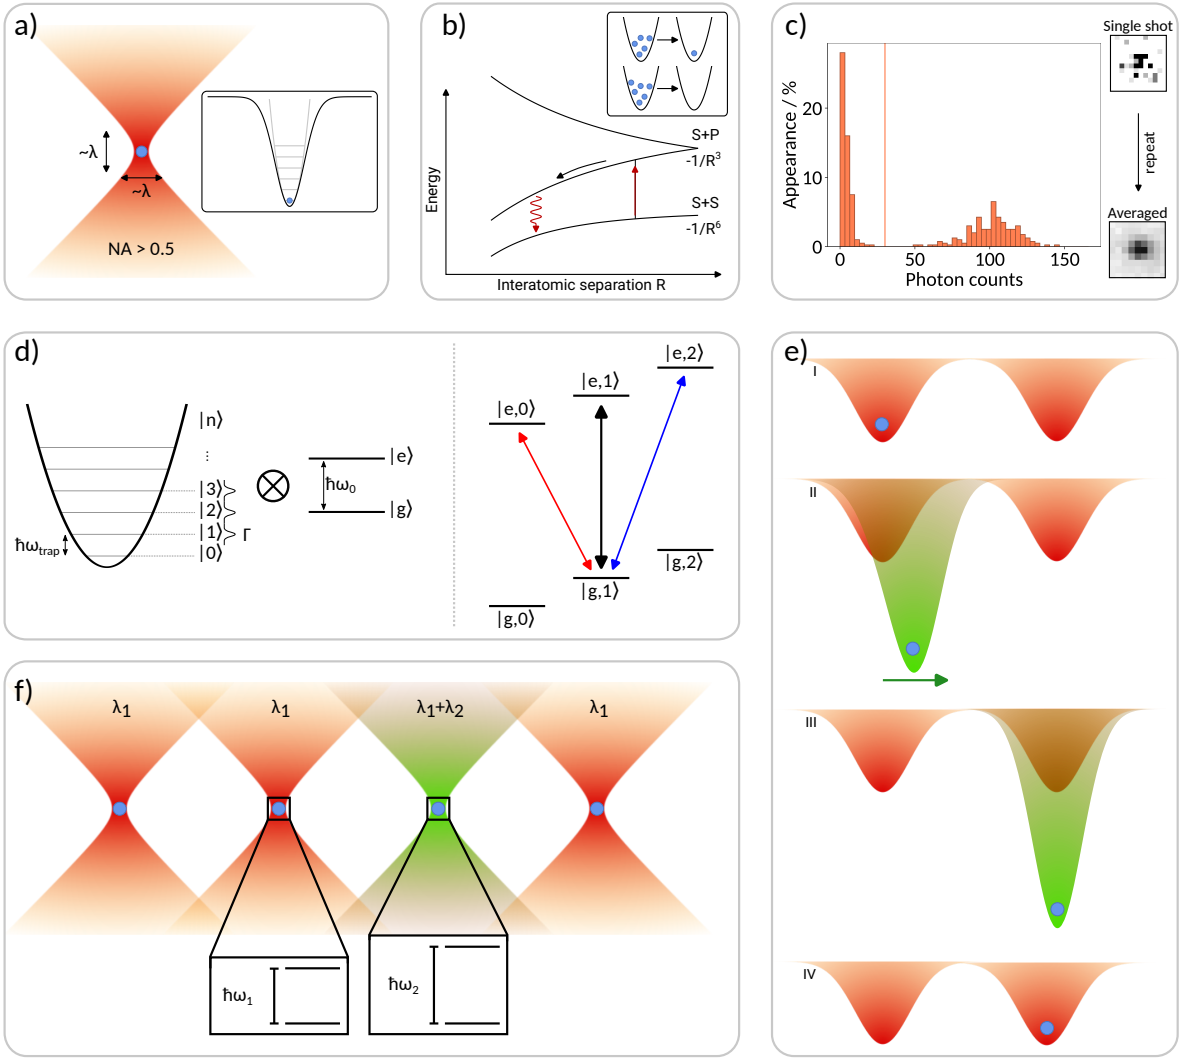

Figure 2.1 – **(a)** A single atom trapped in an optical tweezer, created by a tightly focussed laser beam. The three-dimensional confinement is on the order of the wavelength of the trapping light. Around the focal position, the potential can be approximated as harmonic (inset). **(b)** Light-assisted creation of bound molecular states leads to pairwise loss of atoms, leaving a single atom per trap at most. **(c)** Images of a single atom captured with an electron-multiplying CCD (EMCCD) camera for a single image and averaged over many repetitions of the experiment. The accumulated signal in a region of interest (ROI) on the camera around the trap follows a bimodal distribution, shown in the histogram. The signal can be separated into atom-events, where the scattered number of photons follows a Poissonian distribution and a background peak centered at zero, separated by the vertical line. **(d)** For optical transitions with linewidths γ smaller than the trap frequency ω_{trap} the quantized motional states can be addressed individually. This fact can be exploited to cool atoms to their motional ground state. **(e)** A spatially tunable trap is used to transport a single atom between two traps, which allows for the assembly of low-entropy defect-free arrays. **(f)** An additional local tweezer is used to shift the energy levels at a specific site, shown by the different energy spacing of the idealized two-level system. This shift can be used for site selective manipulation or readout.

CHAPTER 3

Experimental setup and cold atom preparation

We will now take a look at the experimental setup and the cooling stages that are needed to create microkelvin-cold atomic samples. In the first part, we describe the vacuum- and laser- and optical setup, the magnetic field coils and the accompanying control system. In the second part, we discuss the initial cooling stages that are needed to bring evaporated strontium to microkelvin temperatures in several magneto-optical traps (MOTs) to be further trapped in optical tweezers.

The singlet- and triplet level structures of alkaline-earth elements allow for fast and efficiency multi-stage laser cooling protocols on different optical transitions. Specifically, we use an initial broad MOT phase on the $^1S_0 \rightarrow {}^1P_1$-transition where the 31 MHz broad transition allows for rapid cooling of more than 10^6 atoms from room temperature to a few millikelvin in about a second. The resulting temperature prevents loading into a narrow-line MOT on the $^1S_0 \rightarrow {}^3P_1$-transition directly. We therefore insert a stage with an artificially broadened laser that increases the transfer probability to the narrow-line MOT, leading to atom-numbers of around 10^5 atoms at 10 µK.

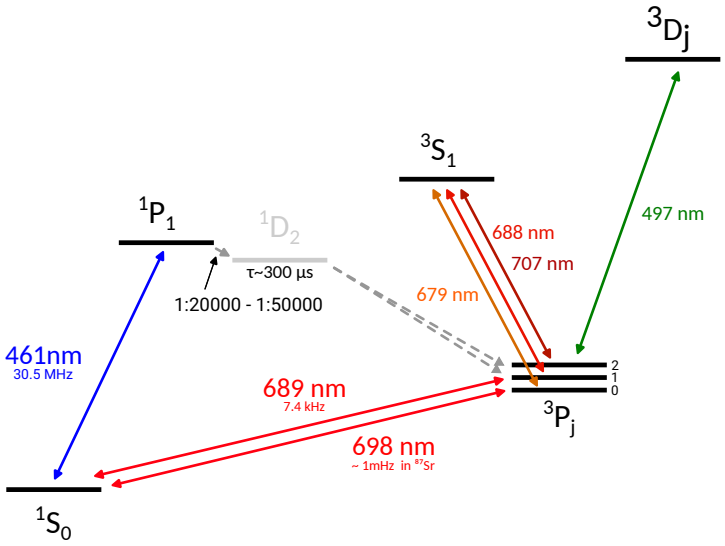

Figure 3.1 – Strontium level scheme [67]

3.1 Vacuum system

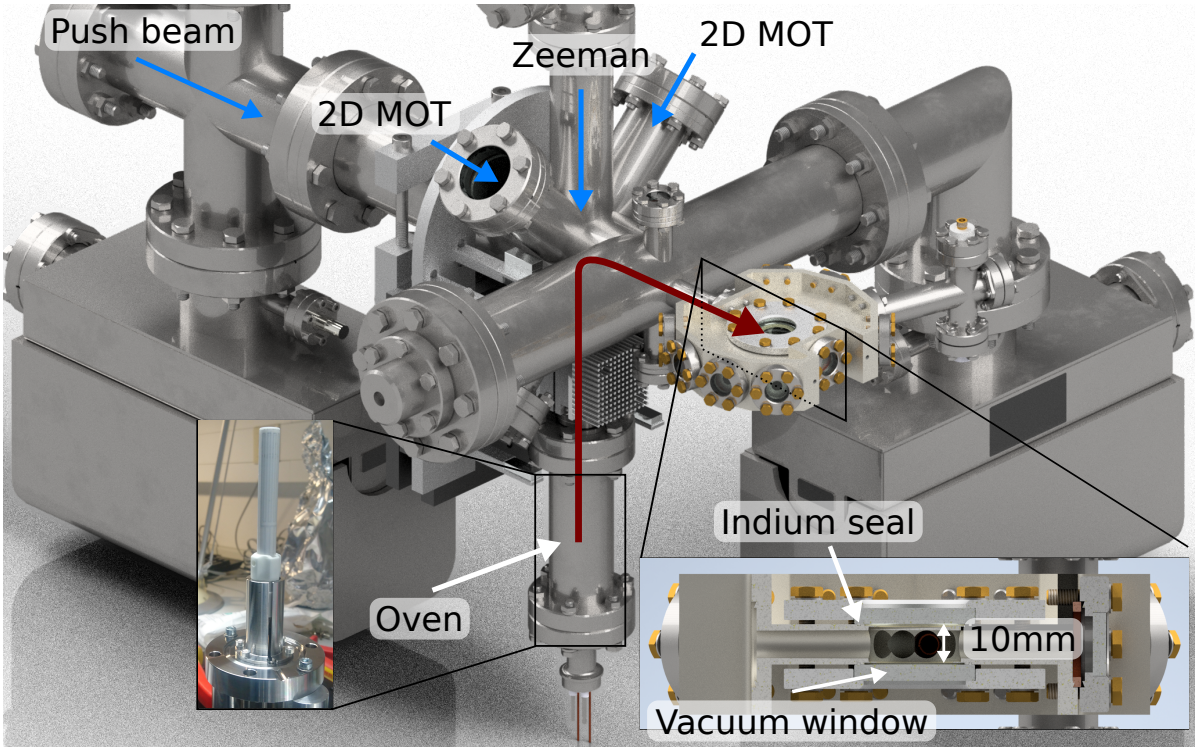

Figure 3.2 – Rendered CAD image of the vacuum system. Beams, optics and other surrounding infrastructure have been omitted.

The vacuum system is separated into two regions: one high pressure region where the oven is located and one low pressure region where the main chamber is attached. The design is inspired (among others) by [68, 69]. The two regions are connected by a differential pumping system (DPS) which has a length of 38.5 mm and increases from a diameter of 1.5 mm initially to 4.2 mm at the end. The initial size is given by the maximum possible divergence angle of the atomic beam from the 2D MOT into the main chamber, where it allows for a full-angle of around 70 mrad. This value is larger than the expected divergence corresponding to the radial Doppler temperature of the 2D MOT and an axial velocity of 40 m/s and does therefore not limit the atomic beam flux.

The DPS, together with a 75 L/s ion-pump[1] on either side, creates a pressure difference of more than two orders of magnitude between the oven and the main chamber side. To compensate outgassing of the vacuum components we additionally make use of titanium sublimation-pumps on each side that we activate every few months. We observe an increase in pressure from 1.5 mBar (8.8×10^{-11} mBar) to 1×10^{-10} mBar (6.4×10^{-10} mBar) with the oven turned off (on) over the course of two months. The pressure on the main chamber site ranges below the detection limit for the pump's pressure sensor of 1.3×10^{-11} mBar. As the pressure on this side does not increase for most of the time, we are confident that the actual pressure is even lower. We use the sublimation pump once the pressure on the main side increases above the detection limit.

[1] *Gamma Vacuum* 75S TiTan Ion Pump

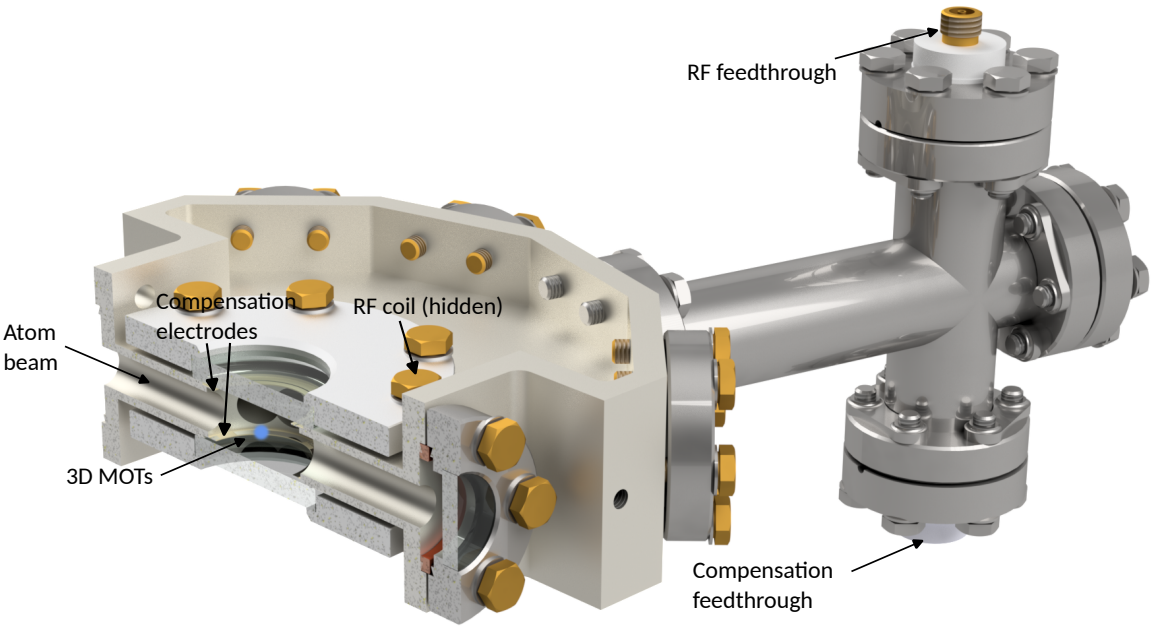

Figure 3.3 – Rendered CAD image of the main chamber. The vertical viewports are sealed using an indium wire. Below are two electrodes to compensate possibly accumulating charges on the windows. The connections are fed through the CF16 cross attached to one port. It also features a small RF-coil to create microwave pulses if hyperfine-state addressing is desired in the future.

Lifetime measurements of atoms in the tweezers Section 5.8 suggest that the pressure at the position of the atoms in the main chamber is up to two orders of magnitude higher than the reading of the pump. We attribute this pressure difference to the complicated geometric structure of the main chamber which can trap particles for a long time before they reach the pump. Additionally, we suspect that there might be one or more virtual leaks in an imperfect weld inside the CF16-cross attached to the main chamber. The weld opened up multiple times over the course of the last four and a half years and had to be sealed with *VacSeal*. A simulation using *Molflow*[2] supports this claim and requires the main chamber to be replaced depending on the future goals of the experiment.

The main chamber is an in-house design built out of titanium shown in Figure 3.3. We chose titanium for its non-magnetic properties compared to stainless steel at a similarly low outgassing rate [70]. The chamber has outer dimensions of 120 mm × 140 mm × 35 mm and features ten CF16 viewports in the $xy-$plane (later also referred to as the *radial*-plane) Behind each window there is a 7 mm wide channel leading to the center of the chamber. The size of this aperture limits the divergence of the atomic beam from the 2D MOT to around 40 mrad. The distance between the 2D MOT and the center of the main chamber was therefore deliberately kept as short as possible at 18 cm. The top and bottom windows are 30 mm wide, 4 mm thick and coated with a novel nano-structure[3] technique. The latter allows for low reflectivities of below 0.05 % over the whole visible spectrum even at large incident angles. This is

[2] https://molflow.web.cern.ch/
[3] *TELAZTEC LLC*

important because two high microscope objectives with numerical apertures (NAs) of 0.7 are used to project the tweezer through the windows and to collect the atomic fluorescence. Standard broadband dielectric anti-reflective coatings have strong dependence on the angle of incidence, which ranges from roughly −45 to 45° for an NA of 0.7. We therefore decided to use the aforementioned nano-structured windows, which have performed well so far. We have yet to explore how well these windows can be cleaned with e.g. Acetone or Methanol.

The windows are sealed using indium wire which is pressed into a bezel engraved into the chamber. Indium chemically bonds to the titanium and the glass which allows for ultra-high vacuum operation [71]. After breaking the surface oxide layer of the indium by immersing it in 5 to 10 % hydrochloric acid for a few minutes, it has to be evenly squeezed between the window and the chamber to chemically bond with both. This process is even more delicate in our case because of the constraints on the parallelity of the windows and the chamber. We aim for an angle of less than 0.05° which we confirm by reflecting a 532 nm laser pointer off of both windows during the sealing process. We overlap the reflections of the windows with the incident beam after traveling a distance of 2 m and observe the number of fringes on the interference pattern while tightening the screws on both flanges. We ensure that the chamber is aligned perpendicular to the alignment beam by glueing a small mirror to a spare flange which we in turn mount flush to the chamber without a window. Placing the chamber on a translation stage allows us to move either the center of the windows or the mirror on the flange into the beam. After moving the stage several times we could not see a deviation in the position of the reflected beam and decided that the chamber is sufficiently well aligned. Although the alignment of the windows to the chamber is not as crucial as the alignment of the windows with respect to each other, it is important for the sealing process. A misalignment can lead to an uneven pressure distribution on the indium wire which can later lead to (virtual) leaks. In the worst case the windows touches the chamber leading to physical damage of the window as they crack due to the additional stress. It is instructive to practice the procedure a few times with low-cost dummy windows, to get a feeling for the behavior and creeping characteristics of the indium wire during the sealing process.

To keep the options for future measurements as wide as possible, we additionally install coils to create static electric fields and radiofrequency (RF) pulses. On the inner side of the vertical windows we have silver-electrodes printed onto Polyimide that can help with compensating stray charges which build up on the windows under the exposure of UV-light. This could be important if we decide to go in to the direction of Rydberg-physics, as these states are sensitive to stray electric fields while also relying on UV-lasers (see Chapter 7).

The fermionic isotope of ^{87}Sr has a rich 11-fold degenerate ground-state state can be used for hyperfine-qubits or $SU(N)$-physics (see e.g. [72, 73]). To manipulate these states, we installed a small RF-coil close to the center of the chamber in the channel of the CF-16 port. We feed the wires through a connector on the bottom port of the CF16 cross, that way the port can still be used for optical access to the center.

3.2 Atom source

A home-built oven serves as the source for the atoms, shown in Figure 3.4(a). The ceramic pellet heating element has been converted into a reservoir with a titanium plug cemented into the bottom of the ceramic base which holds the electrical connections as well. It can hold about 3 to 4 g of 99.99 %-pure strontium

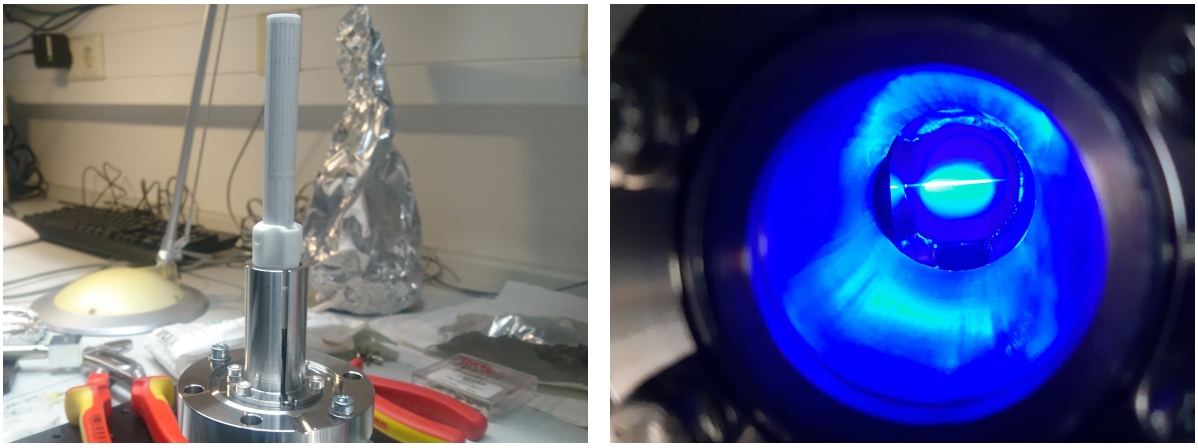

(a) **(b)**

Figure 3.4 – Oven assembly **(a)** and fluorescence **(b)** from the push-beam passing through the atomic gas when viewed through the Zeeman-beam viewport.

chunks[4] and was filled under an inert gas atmosphere to prevent oxidation. We estimate this amount to last for several years with continuous operation under the current experimental conditions.

A thermocouple element embedded into the bottom of the ceramic tube is used to monitor the temperature of the oven. We measured the temperature distribution of the heated system in air and suspect that there is a difference of about 130 °C between the sensor and the hottest part of the filament. A third temperature, calculated from the linearly rising resistance of the heating element suggests a temperature around 30 °C higher than the sensor reading. The actual temperature of the strontium inside the oven probably lies somewhere in between. We estimate the flux after [74] to be:

$$\Phi = \frac{1}{4} n \sqrt{\frac{4}{\pi}} v_p A \tag{3.1}$$

with the atom density of $n = \frac{p}{k_B T}$ with the pressure $p(T)$ determined from the vapor pressure curve [75]. $v_p = \frac{2k_B T}{m}$ is the most probable velocity and $A = 0.5\,\mathrm{cm}^2$ is the aperture of oven. The resulting flux is around $\Phi = 1.2 \times 10^{14}$ atoms/s.

We usually operate the oven at 270 °C sensor temperature which requires around 9 W. This power level increases the temperature of the heat sink attached to the vertical connection tube between the oven and the 2D MOT to around 45 °C.

3.3 Magnetic field control

Around the main chamber there are four pairs of coils to control the magnetic fields (see Figure 3.5). As the name suggests the MOT coils are used to create the magnetic field gradients for both magneto-optical traps (MOTs). They are built from 6 mm × 1 mm wire with 16 turns each, resulting in an inductivity of about 24 µH per coil. The wire is glued to a monolithic slitted aluminium mount that is cooled with two

[4] *Sigma Aldrich* dendritic pieces

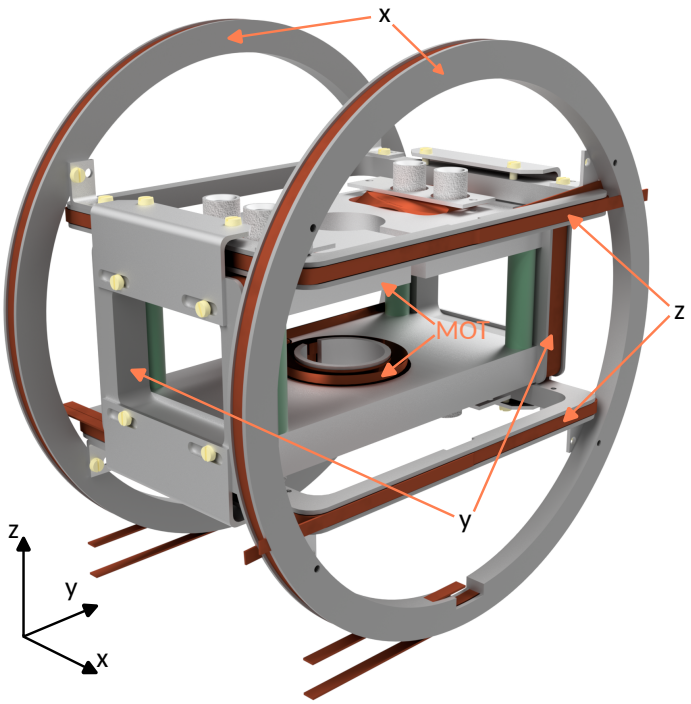

Figure 3.5 – Rendered CAD image of the coils.

commercial CPU cooling blocks per side. For the blue MOT, gradients of around 50 G/cm are required for which around 40 A are needed. During continuous operation with cooling, the coils heat up to about 34 °C which, although not a serious problem, already causes misalignment in the upper objective. We find that a gradient of 24 G/cm leads to a steady-state temperature of 24 °C which is roughly equal to the value that is reached during the tweezer sequences, preventing misalignment.

For the red MOT we have to decrease the gradient from 55 G/cm to about 4 G/cm, ideally on a timescale of a millisecond. To achieve this, we use two power supplies[5]: one to supply the required current for the red MOT gradient of around 4 G/cm, to which we add the much larger current required for the blue MOT. We then switch off the larger current using a power MOSFET and dump the inductive voltage spike into a varistor. With this technique we measure a field decay time of around 10 µs when switched from 50 A which we cross-check by a circuit simulation software[6]. This decay rate is much faster than required for tweezer operation.

To add constant offset fields, we have three pairs of coils in a Helmholtz-configuration. We use the same wire as for the MOT-coils at 8, 20 and 6 turns for the x-, y- and z-coils. In hindsight, this was a suboptimal choice. Although lower inductivity coils can be switched faster, the higher currents also lead to more heat dissipation. The coils were initially designed for fields around 5 G which is much smaller than the large magnetic offset fields around 20 G required for the tweezer operation at 515 nm, where currents around 35 A are necessary. For a resistance of around 15 mΩ these currents result in about 10 to 20 W of dissipated power leading to significant temperature increases. Attempts to

[5] *Delta Elektronika* SM 60-100 and *Delta Elektronika* SM 30-5
[6] LTSpice by *Linear Technology*

cool the coils were not successful and they reach more than 40 °C (depending on the duty cycle) for measurements with the 515 nm tweezer sequence. This temperature itself is again not problematic, but it causes misalignment in the upper objective which must be avoided. In the future, a replacement of the coils with an updated design using more turns of thinner wire could be combined with the replacement of the main vacuum chamber. Using an TTL-controlled H-bridge we can change the direction of the current flow and therefore the orientation of the magnetic field.

3.4 Experiment control system

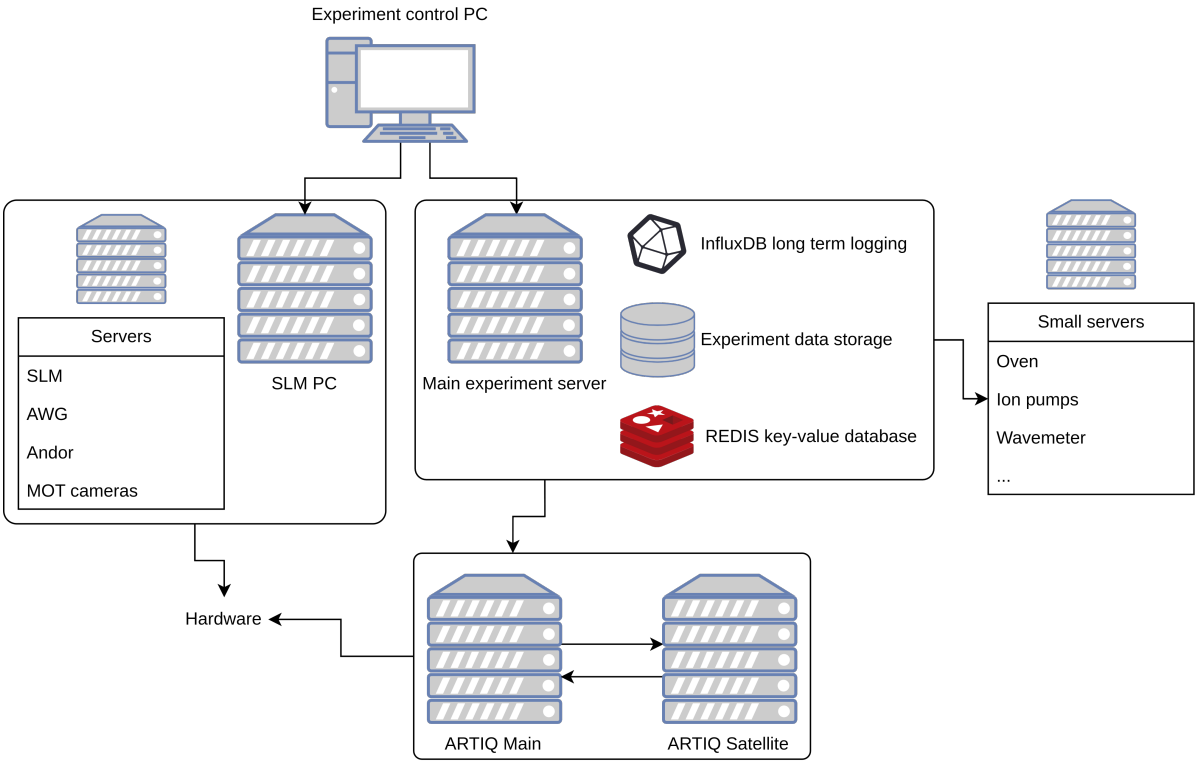

Figure 3.6 – Schematic of the experiment control system. The real-time I/O is handled by the commercial ARTIQ system, which is a combination of a processor and an FPGA.

The experiment is controlled by a combination of a real-time[7] I/O FPGA/CPU system, several non-real-time servers and computers. The former is a commercial system by ARTIQ (Advanced Real-Time Infrastructure for Quantum physics) systems, featuring an Artix-7 FPGA[8]. The system is divided into two 19 "-rack mountable crates linked by a SFP-fiber connection, and the available channels are shown in Table 3.1.

Communication with the ARTIQ-system is done a by server running on the main experiment computer which can be connected to from a client via network. For a measurement sequence (real-time-sequence),

[7] We adapt the common definition of real-time as everything that is guaranteed to take a certain time independent of system or network load.

[8] *Xilinx* XC7A100T

Channel type	Channels	Specifications	Chip
Digital outputs	56	5 V TTL, 150 MHz max toggle rate	SN74BCT25245
Digital inputs	8	5 V TTL, 3 ns min pulse width	SN74BCT25245
Analog outputs	64	±10 V, 16 bit, 20 µs settling time	AD5372
Analog inputs	16	±10 V, 16 bit, 1.5 MHz max sampling rate	LTC2320-16
DDS	12	0 to 400 MHz, −32 to 10 dBm, switch in < 100 ns	AD9910

Table 3.1 – ARTIQ system hardware specifications.

a sequence of commands (sequence-script) together with a list of parameters is send to the ARTIQ-server, which then gets compiled and uploaded to the FPGA. Hardware that is also needed during a real-time-sequence such as the spatial-light-modulator, arbitrary-waveform-generator (see Section 4.3 and Section 4.4) or the EMCCD camera is controlled using a second PC. The PC is only used to initialize the hardware where any changes (like taking an image) during the sequence are triggered by TTL-pulses. The atom rearrangement described in Chapter 6 is an exception to this, as the algorithm that determines how atoms are rearranged, the computation and upload of the RF-waveform is currently run on the CPU. In the future, this could be outsourced to a dedicated real-time-system to be again independent of system load.

3.5 Laser system

3.5.1 461 nm

Beam	Detuning/MHz	Power on experiment table/mW
Spectroscopy	0	30
Zeeman	-220	45
2D MOT	-30	30
3D MOT	-30	10
Push/Excitation	±25	0.1
MOT imaging	±25	10
Tweezer imaging	±25	0.1

Table 3.2 – Laser setup 461 AOM parameters

The 461 nm light used to drive the main cooling and imaging is generated by a home-built frequency-doubling cavity. The latter is pumped by a titanium-sapphire (Ti:sapph) laser[9] giving about 4 W at 922 nm. It is pumped by a 25 W laser at 532 nm[10] which we usually operate at 21 W. We follow the standard practice and place the laser setup on a separate optical table to reduce the mechanical and thermal coupling and to reduce stray light. This also means that all beams are passed through optical fibers to the main experiment table.

[9] *Sirah* Matisse CS

[10] *Spectra Physics* Millenia eV 25

We use a periodically-poled KTP crystal for frequency doubling to reduce the complexity of the cavity design compared to a Brewster-cut crystal. We see that the current design becomes inefficient for higher powers because of *thermal dephasing* of the phase-matching condition caused by the buildup of 461 nm light inside the crystal [76]. The cavity generates around 550 mW of second-harmonic light at a pump power of 3 W which could be increased in the future by changing the radii of the curved intracavity mirrors. We see a drift in the divergence of the infrared beam leaving the Ti:sapph cavity on the order of months which requires us to reposition the incoupling lens from time to time. After such an optimization we reach up to 700 mW of second-harmonic light which then decreases over the course of a week back to the 550 mW level. We further split the light after the pick-off after the spectroscopy according to Table 3.2. We noticed that the order of separation matters as we see that back-reflections from the higher power beams can be coupled into the lower power beams even through multiple beamsplitters. The last three beams are therefore the imaging beams where for tweezer operation microwatt-accuracy and -precision is required. We also shutter the axial beam of the 3D MOT as it otherwise increases stray light on the EMCCD camera which works in direct line of sight.

Laser stabilization - hollow cathode lamp lock The Ti:sapph laser has an internal reference cavity which can be used to stabilize the frequency to a linewidth of 50 kHz. This cavity has a sufficient short-term stability for the 31 MHz wide main transition but it tends to drift on the time-scale of minutes. To compensate this drift we use about 30 mW to stabilize the frequency with a PDH locking scheme [77] to a hollow cathode lamp[11]. For that we employ a pump-probe saturation spectroscopy scheme [63] which generates a Doppler-free signal of the otherwise > 1 GHz broadened transition. Because of the small volume of around 1 mm $\times$ 1 mm $\times$ 20 mm and low density of the atomic vapour the signal strength is low, with a signal-to-noise ratio of about 5. To get rid of the Doppler-broadened background, we chop the pump beam using a 117 kHz square-wave and demodulate the Doppler-free signal with the Doppler-broadened signal (see [78]). We feed the resulting error signal back to the reference cell of the Ti:Sapph cavity. In the future, the hollow cathode lamp will be replaced by a home-built spectroscopy cell that can also be used for spectroscopy of the $^1S_0 \rightarrow {}^3P_1$ transition. We will comment on the design of the cell in Section 3.7.3 where we explain the methods used to locate the $^1S_0 \rightarrow {}^3P_1$ transition frequency. We also did some brief testing by stabilizing the frequency using the reading from our wavemeter. This locking scheme is good enough to see MOTs of ^{84}Sr, ^{86}Sr and ^{87}Sr but is not feasible in practice for tweezer operation.

3.5.2 689 nm

Beam	Detuning/MHz	Power/mW
MOT radial 1	−10 to 1	4
MOT radial 2	−10 to 1	4
MOT axial	−10 to 1	10
Probe beam	−20 to 10	1

Table 3.3 – Parameters of the different beam paths at 689 nm.

[11] *Heraeus* 3BAXSR

To drive the $^1S_0 \rightarrow {}^3P_1$-transition requires it to have a linewidth on the order of a kilohertz at output powers of around 60 mW which makes the 689 nm laser system the most delicate laser used in the experiment. It is based on a Toptica DLPro diode laser which gives around 30 mW of power at a linewidth of less than 5 kHz. To reduce the laser linewidth from the free-running value of 100 kHz to below 5 kHz, it is locked to an ultra high-finess cavity[12] using an offset-locking technique. We use about 50 µW of light that is passed through a fiber-coupled electro-optic modulator (EOM)[13]. The latter generates sidebands at ≈ 370 MHz of which one is locked to the cavity using a standard PDH-lock [77]. Using this technique, we can shift the laser frequency by adjusting the frequency of the EOM which is controlled by a signal generator[14]. After an initial drift of around 14 kHz per day for the first ≈ 20 months, it gradually slowed down to 8 kHz per day for the last 12 months.

To increase the output power to 60 mW which are required to for experiment, we use two additional laser diodes[15] which are injected by the seed [79]. We could not find an influence of the seeding process on the linewidth of the emitted light [59]. The injection status is monitored using a wavemeter[16]. We found the system to be quite robust, not requiring realignment on the order of weeks. When a diode loses its injection status, it is usually enough to slightly tune the injected diode current which can be controlled remotely and automated. The system has been in operation for three years and we did not see any damage or decline in power, something that has been reported to be an issue with tapered amplifiers at this wavelength [80, 81].

We further split the light into four different beams, two for each diode. The first two beams drive the radial MOT beams. To adjust the frequency and power we again use AOMs in a double pass configuration with an efficiency of 70 %. Together with the fiber coupling efficiency, each beam has a power of 6 mW, enough to supply the 4 mW required. The second two beams are the axial MOT beam and an additional *probe*-beam. Due to the objectives, the axial MOT beam cannot be retroreflected and is split after the optical fiber before entering the main chamber. This in turn means that the required power is about a factor of two larger than the radial beams. The probe beam is used for spectroscopy and cooling of atoms in the tweezers with a detuning between ±5 MHz and a power of up to 2 mW. We use this beam for the sisyphus cooling processes described in Section 5.4.2. For the sideband cooling described in Section 5.4.1 three (nearly) orthogonal beams are required. In this case we have to use the MOT beams as there is not enough optical access to add a second set of beams along all spatial directions.

3.5.3 Tweezer lasers 515 nm, 532 nm & 813 nm

For the 515 nm laser, we use a 10 W system by *Azur Light Systems*, for 532 nm a Verdi V-18 by *Coherent* and for 813 nm light a second titanium-sapphire laser by *Sirah*[17]. In addition to the 922 nm model, the latter features an additional intracavity EOM to reduce the linewidth below 30 kHz. Due to the off-resonant nature of the dipole trap, the frequency stability requirements for the tweezer lasers are comparably low and we do not use any additional locking techniques. However, to prevent parametric heating in the traps, the power stability of the tweezer lasers is of great importance. We use intensity-servoing capabilities of the ARTIQ-system to stabilize the lasers to 0.2 % in front of the vacuum chamber.

[12] *Stable Laser Systems*, Finesse > 250000
[13] *Jenoptic* PM705
[14] *Rohde & Schwarz* SMA100A
[15] *Thorlabs* HL6750MG
[16] *HighFinesse* WS7
[17] Matisse CX, linewidth < 30 kHz

We also measure the stability in frequency space using a spectrum analyzer and do not find any peak above −95 dBm between 0 to 300 kHz.

3.5.4 Repumping lasers 679 nm & 707 nm

Figure 3.1 shows that there is a chance of about 1 : 20000 to 1 : 50000. The $^2D_1$-state decays to the $^3P_{\{1,2\}}$ triplet states, and atoms in the $^3P_0$ (which is populated after repumping by the $^3S_1$ decay) and $^3P_2$ are effectively lost due to their long lifetimes. The effect can be mitigated by repumping both states via the $^3S_1$ state back into the $^3P_1$-state, which then decays to the ground state $^1S_0$. We currently do not have a repumping laser for $^3P_1 \rightarrow {}^3S_1$ which could be useful to add in the future at it is required for certain spectroscopy protocols in the tweezers. The other two repumping lasers itself are home-built external-cavity diode lasers.

3.6 Optical setup

3.6.1 Microscope objectives

To create the tweezers and to collect the scattered fluorescence, we use two microscope objectives by *Special Optics*, each with an NA of 0.7. Each objective has an effective focal length (EFL) of 13.477 mm, a clear aperture of 18.88 mm and an f-number of $f/\# = 0.714$. The f-number and the numerical aperture NA are related by $f/\# = \frac{1}{2\mathrm{NA}}$ (see e.g. [82]). The lower objective is used the image the atoms and project tweezers at 515 nm while the upper objective project the tweezer arrays at 813 nm. To ensure that our objective performs like expected, we measure the PSF of a point-like source.

3.6.1.1 PSF and wavefront measurements

To directly measure the point spread function of the objective, a point source has to be imaged. Of the four available wavelengths (461 nm, 515 nm, 689 nm, 813 nm) we decided to perform the measurements with 461 nm, as shorter wavelengths are generally more sensitive to misalignments.

We considered two approaches: a back illuminated pinhole or a scanning-near-field-microscopy fiber (SNOM-fiber) [18]. The latter is a stretched and then etched single mode optical fiber with an aperture around 70 nm. We measure the point-spread-function by capturing a few ten microwatts of light from the SNOM-fiber which we image onto a CMOS camera[19] using a lens with a focal length of 1 m. For the effective focal length of 13.5 mm for our objective this results in a magnification of 75. This setup was helpful for getting used to the alignment procedure, but the overall stability was insufficient. The vertical mounting shown in Fig. 3.7(a) showed to be sensitive to vibrations which made it difficult to capture a good image. An example image is shown in Fig. 3.8(a) and Figure 3.8(b) which confirmed that the objective performed as expected. We explain the residual aberrations by the deviation of the window thickness from the ideal case. We also repeated the procedure by focussing a collimated beam using a different microscope objective[20] with an NA of 0.65 and imaged the focus using the 0.7NA objective. It requires some experience to couple a laser beam orthogonally aligned on the optical axis and additionally the window that simulates the vacuum window (see Appendix B). Due to this complexity and because it

[18] *Spectrum Instruments Ltd.* MF002

[19] Thorlabs CS165MU1

[20] 40X Olympus Plan Achromat Objective, 0.65 NA, 0.6 mm WD

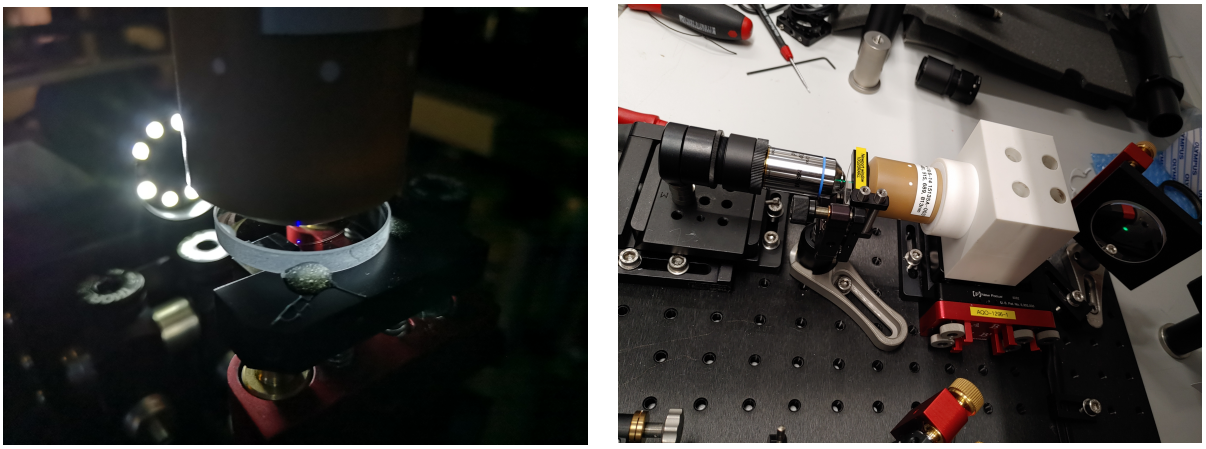

Figure 3.7 – Shining SNOM fiber in our test setup (blue dot) **(a)**. The distance to the window is monitored with a cheap USB microscope. The vacuum window dummy and high NA objective can also be seen. Because of insufficient stability we repeat the procedure using a second microscope objective with an NA of **(b)**.

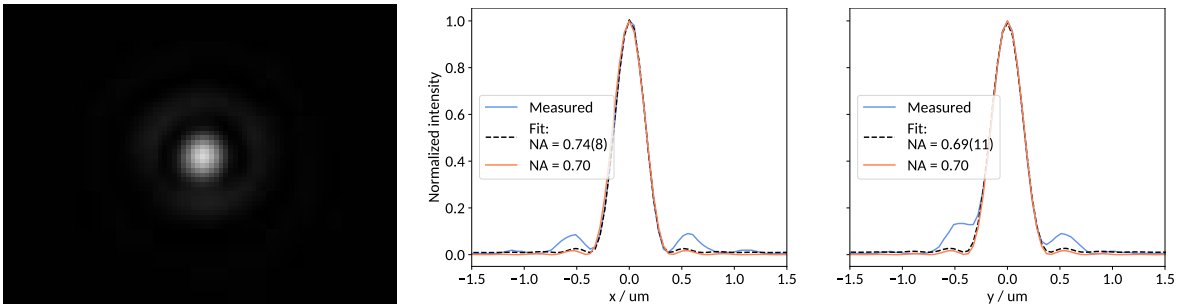

Figure 3.8 – **(a)** Image of the point-spread-function (PSF) measured by imaging a SNOM fiber. **(b)** Fitting the PSF we extract an NA matching the design value of 0.7. The larger side-lobes can be explained by a residual spherical aberration caused by a deviation of the window thickness from the optimal value.

is strictly speaking not a point-source at an NA of 0.7, it was not considered as an initial approach, but it helped later on to investigate the influence of artificially introduced misalignment on the quality of the imaging. The optical setup is shown in Fig. 3.7(b).

3.6.1.2 Focal shifts

The objectives show a focal shift of around 50 µm between 461 nm and 515 nm and 200 µm between 461 nm and 813 nm. This means we cannot use a single objective to image atoms in 813 nm traps and the fluorescence light at 461 nm is aberrated for imaging atoms in 515 nm traps. Additionally, we see a drift of the upper objective in both radial and axial direction caused by the micrometer-stage[21] which further complicates the alignment-process and severely limits the stability of the system. This drift has multiple contributions: first, the screws of the stage show a hysteresis which relaxes over the course of

[21] *Newport* 8081

$\sim$ 24 h. Secondly, drifts caused by thermal coupling to the mot and compensation coils, even though not in direct contact, also leads partly reversible misalignment. In the future, this could be solved by a single new microscope objective that has a focal shift in the region of a few micrometers.

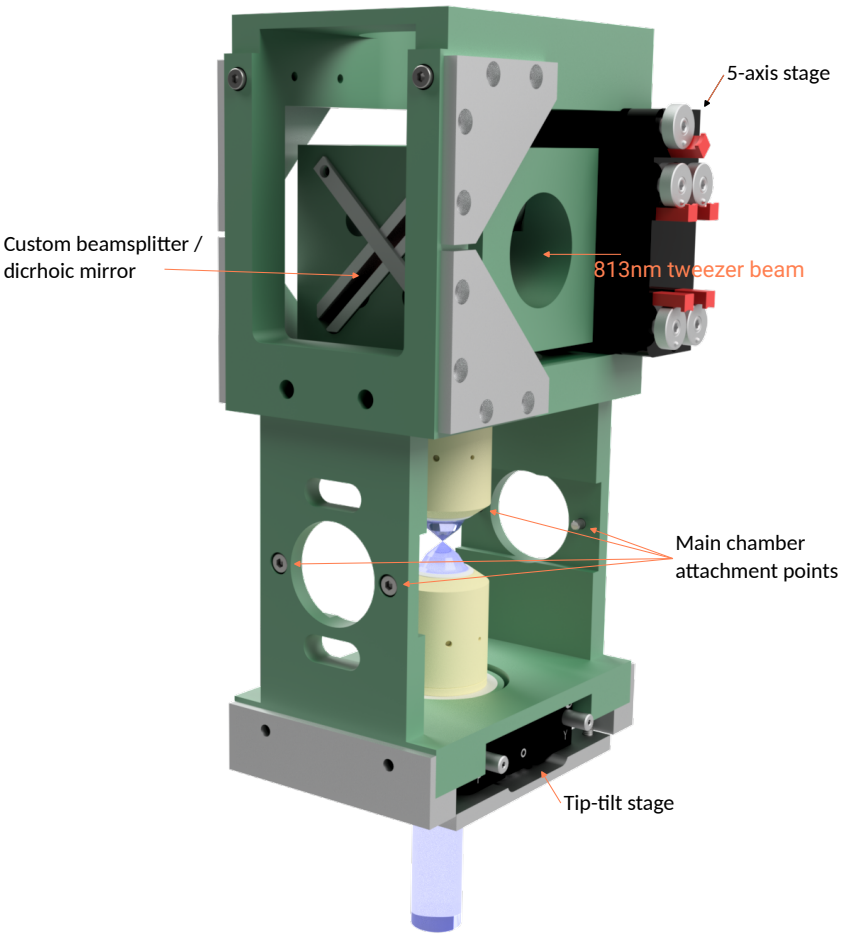

Figure 3.9 – CAD render mounted objectives

3.6.2 Beam path in and around the vacuum chamber

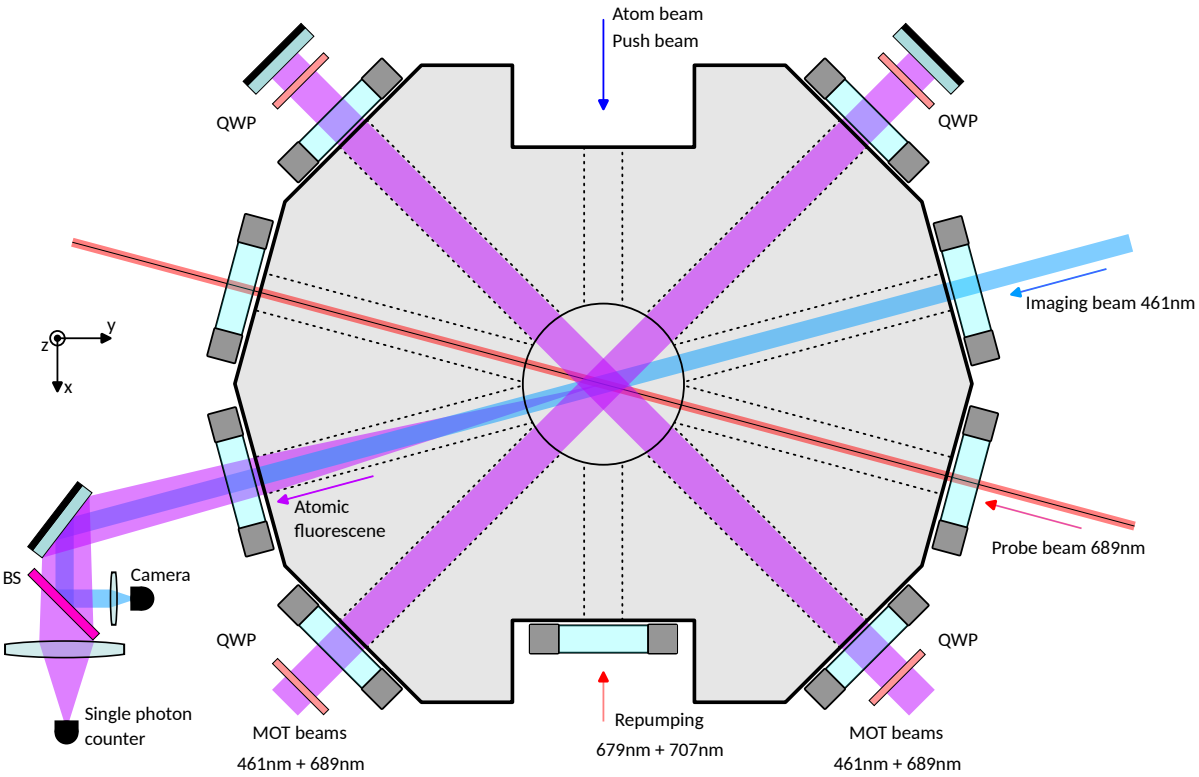

Figure 3.10 – Beams around the chamber passing through the objectives in the radial xy-plane.

To conclude description of the experimental setup, we will now take at the look of the laser beams in and around the main chamber. Figure 3.10 shows the optical setup in the radial plane, whereas Figure 3.11 shows the yz-plane.

Radial plane

The push beam transfers the atoms from the two-dimensional MOT in the main chamber. The blue and red MOT beams use the same ports, where only the radial beams are retroreflected. A dedicated MOT-imaging beam at 461 nm is used to take absorption images of the MOTs, whereas the push-beam is used to excite the atoms in the traps for fluorescence imaging. We use two imaging beams as the power levels and beam sizes are vastly different. Especially the red MOT takes up to several milliwatts for imaging due to its high density, compared to around 10 µW for the tweezer imaging. A probe beam at 689 nm is used for probing and cooling in the traps at 532 nm and 813 nm. It has a power of up to 2 mW for detunings between −5 to 5 MHz at a beam size of 2.7 mm at the position of the atoms. For resolved sideband cooling in 515 nm traps we have to use the MOT beams because of spatial constraints. This is not ideal as standing-wave effects from the retroreflection might lead to unwanted local intensity deviations.

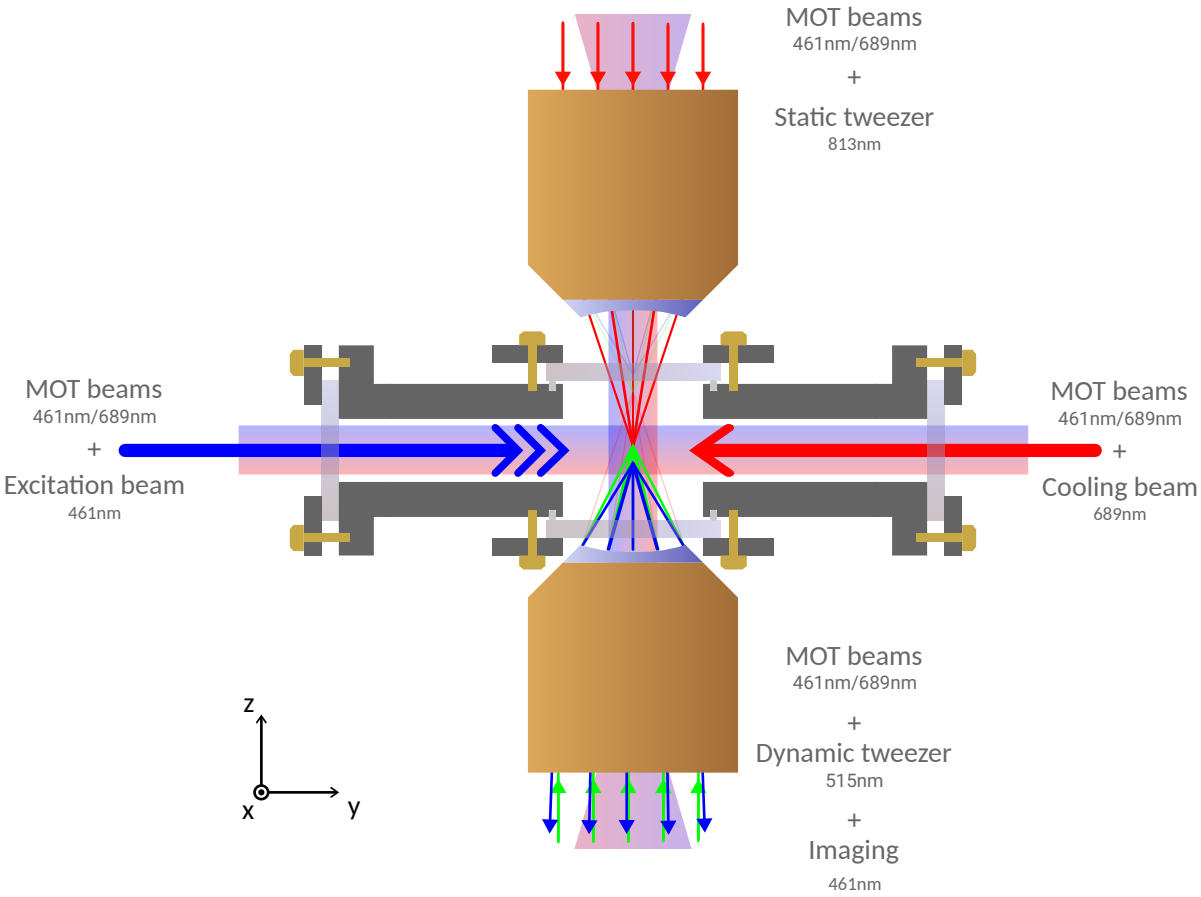

Figure 3.11 – Beams around the chamber passing through the objectives in the yz-plane.

Axial direction

Along the axial direction things get a little crowded. The MOT beams, which have to be collimated at the position of the atoms are focussed into the back focal plane of the objectives. Because the tweezer and fluorescence light have to also pass the objective we have a custom-built beamsplitter in front of each objective. These beamsplitters transmit light at 689 nm and also 10 % at 461 nm independent of the polarization. They reflect 90 % of the light at 461 , around 93 % at 515 and 532 nm and close to 100 % at 813 nm. This ensures, that the tweezer and fluorescence light are reflected instead of transmitted whenever possible to reduce optical aberrations. In fact, the fluorescence light is only reflected all the way to the camera to reduce aberrations as it is convergent due to the difference in focal lengths. The green tweezer light is passed through a dichroic mirror once, separating the fluorescence light from the counterpropagating tweezer light. The tweezer light, compared to the imaging light, can be ensured to be collimated. We find that the coating slightly rotates the polarization axis at 515 nm which is a problem for the *magic*-traps (see Section 5.4.1.4). In the future, this could be solved by designing a coating with a zero phase delay for both polarizations at an angle of 45° at 515 nm.

3.7 Initial cooling stages

We now describe several different magneto-optical trap (MOT) stages that are required to load the tweezers. The MOT has become a standard ingredient in modern cold atom experiments and a detailed description can be found in e.g. [60, 63].

	$^1S_0 \rightarrow {}^1P_1$	$^1S_0 \rightarrow {}^3P_1$
Linewidth $\gamma/2\pi$	30.5 MHz	7.4 kHz
Lifetime $\tau = 1/\gamma$	5 ns	21 µs
Saturation intensity $I_{\text{sat}} = \frac{\pi}{3}\frac{hc\gamma}{\lambda^3}$	42.7 mW/cm^2	2.95 µW/cm^2
Maximum acceleration $a_{\text{max}} = \frac{\hbar k\gamma}{2m}$	9.5×10^5 m/s^2	150 m/s^2
Doppler temperature $T_{\text{D}} = \frac{\hbar\gamma}{2k_{\text{B}}}$	0.7 mK	180 nK
Recoil temperature $T_{\text{R}} = \frac{\hbar^2k^2}{k_{\text{B}}m}$	1 µK	0.5 µK

Table 3.4 – Relevant parameters for laser-cooling and trapping strontium atoms on the $^1S_0 \rightarrow {}^1P_1$ and $^1S_0 \rightarrow {}^3P_1$ transitions [67].

3.7.1 2D MOT

The design of the 2D MOT is heavily influenced by [68] where further details can be found. For the magnetic field of the 2D MOT we use stacks of permanent magnets[22]. These stacks have a surface field strength of up to 1.5 T and their arrangement leads to a gradient of 24.5 G/cm at the position of the 2D MOT. The residual magnetic field and field gradient in the center of the main chamber are around 0.2 G and 0.05 G/cm which is comparable to the earth's magnetic field of 0.5 G.

The 2D MOT uses a single cycling beam with a diameter of 24 mm at a power of 40 mW, resulting in a saturation of $I/I_s \sim 0.4$ at a detuning of -30 MHz. Assuming a capture velocity v_c of around 60 m/s for the 2D MOT, we capture less than around 3×10^{-4} % of the atoms leaving the oven. We arrive at this conclusion by estimating the captured fraction of the atoms following [83] with:

$$\Phi_c = \Phi \times \Omega \times \frac{1}{2}\left(\frac{v_c}{v_p}\right)^4. \tag{3.2}$$

Here Φ is the total flux from the oven, v_c is the capture velocity of the MOT, determined from a simulation and $\Omega = 0.48$ % is the fractional solid angle of the MOT capture region from the oven. $\frac{1}{2}\left(\frac{v_c}{v_p}\right)^4 \approx 0.055$ % is the fraction of atoms with a velocity v lower then the capture velocity v_c for a one-dimensional Maxwell-Boltzmann distribution with $T = 300\,^\circ$C $\Rightarrow v_p \approx 330$ m/s. This leads to a capture rate of 3.2×10^8 atoms/s. We investigate the 2D MOT performance by measuring the flux inside the main chamber with a single photon counter (Fig. 3.10). For these measurements, we use a push beam with 10 µW and a detuning of 20 MHz. We observe a saturation of the flux for 2D MOT beam diameters of 15 to 25 mm and saturations above 0.4 and conclude that the atomic beam from the oven is at least somewhat directed.

[22] *Eclipse* block magnet, 10 mm × 5 mm × 25 mm, neodymium N35, holding force 4.9 kg

Zeeman slower

To increase the fraction of captured atoms we additionally use a Zeeman-like slower beam, focussed into the aperture of the oven. Ideally a square-root shaped magnetic field should be used to keep the slowing atoms on resonance [60]. This is not true in our case and the actual efficiency depends on the magnetic field present in the experiment and the (to some degree unknown) velocity of the atoms exiting the oven. We see an atom number increase of a factor of 5 to 10 in the blue MOT when using this Zeeman beam at a power of 55 mW at an intensity of $\sim 7I_s$.

An annotated CAD image of the yz 2D MOT plane, including the oven and Zeeman slower viewport is shown in Figure 3.12.

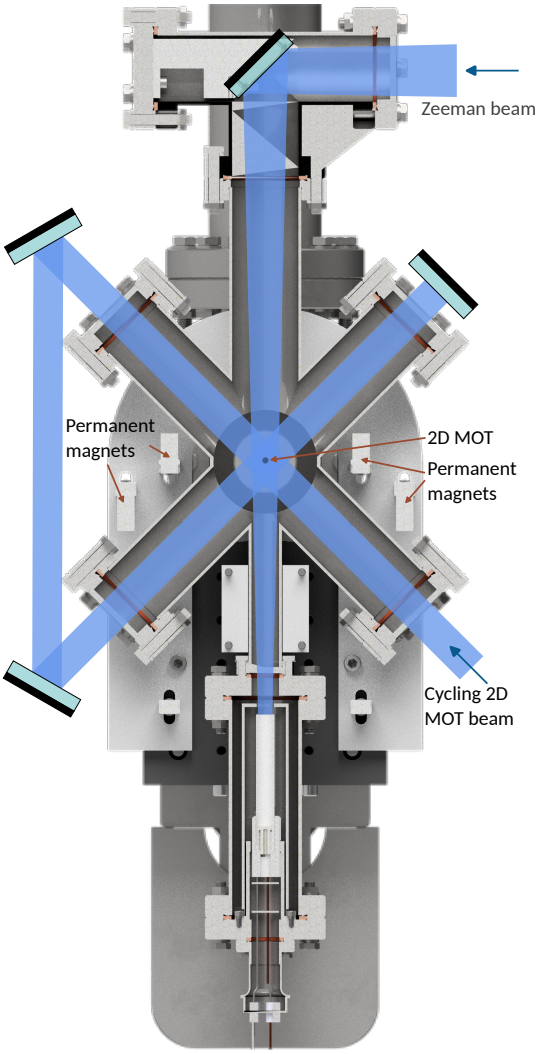

Figure 3.12 – Rendered CAD image of the oven and the 2D MOT. The U-shaped aluminium block is attached to a stage to align the permanent magnets to the chamber. The ceramic oven tube is encased in a enclosure to limit the divergence of the atomic beam and to help with heat dissipation. The Zeeman beam enters through the top viewport, and an intra-vacuum mirror guides it directly into the oven aperture.

3.7.2 Blue MOT

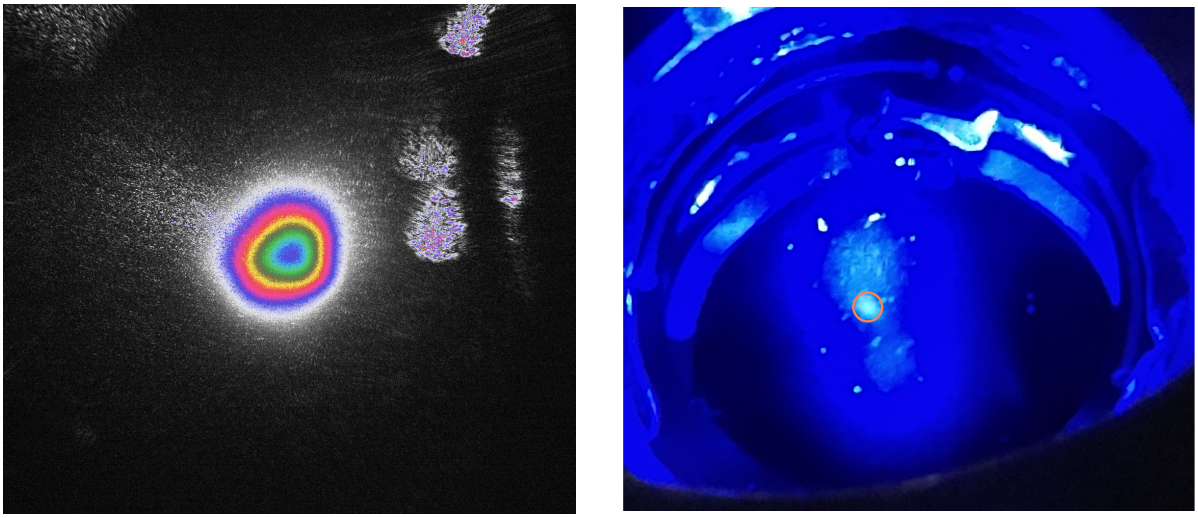

Figure 3.13 – **(a)** Fluorescence image of the blue MOT. The curved structure on the right is the shadow of the small microwave coil with an outer diameter of 6 mm. **(b)** A photo taken with a common smartphone camera. The red circle shows the blue MOT while the diffuse clouds above and below are the MOT beams scattering off of the vacuum windows.

The first three-dimensional cooling stage is the blue MOT formed on the $^1S_0 \rightarrow {}^1P_1$ transition in the main chamber, where it can be continuously operated. The atoms are transported from the 2D MOT into the main chamber using a 20 MHz blue-detuned beam with a power of 10 µW at an intensity of $I/I_s \sim 8$, focussed at the position of the 2D MOT. The resulting velocities are on the order of 30 m/s roughly equal to the capture velocity $v_c \approx 28$ m/s of the blue MOT, and we conclude to capture most of the atoms. We use radial beams with diameters of 6 mm and intensities of $I/I_s = 0.4$ and a magnetic field gradient of 55 G/cm. For the axial beams, the shape is difficult to estimate, as the process of focussing the beam through the beamsplitter into the back focal plane of the objective introduces severe aberrations (see Section 3.6.2). We estimate a beam size of ~ 2 to 3 mm at the position of the atoms at a power of 2 mW per beam.

With the aforementioned parameters, we measure a flux of 3×10^7 atoms/s in the main chamber. The resulting atom number is much higher than needed for tweezer operation and we will measure the tweezer occupation probability depending on the loading time in Section 5.3. A larger atom number is still desirable to reduce the necessary loading time and increase the repetition rate of the experiment. The MOT has a maximum optical density of ~ 1, resulting in a atomic density of 3×10^{10} cm^{-3}. We determine the temperature by performing a time of flight measurement, resulting in temperatures of 2 mK and 3 mK for the radial and axial direction(s) respectively. The final temperature of 2 mK results in a Doppler broadened linewidth of 1.5 MHz for the $^1S_0 \rightarrow {}^3P_1$ transition.

3.7.3 $^1S_0 \rightarrow {}^3P_1$ spectroscopy

To determine the frequency $^1S_0 \rightarrow {}^1P_1$ transition, our wavemeter with an accuracy of 60 MHz can be used, and we can detect a signal from the hollow cathode lamp after that. This, however, does not work

for the $^1S_0 \rightarrow {}^3P_1$ transition frequency, as the signal is not strong enough (a possible solution in the form of a proper spectroscopy cell will be discussed at the end of this section). We are thus required to come up with a different technique, where we have to use the atoms in the blue MOT to measure the $^1S_0 \rightarrow {}^3P_1$ transition frequency. To locate the transition frequency from the 100 MHz-level of the wavemeter down 100 kHz-level required for the MOT, we use a two step process. We first perform spectroscopy on the $^1S_0 \rightarrow {}^3P_1$ transition in the blue MOT first and then on the cold atomic cloud immediately after switching off the magnetic field gradient of the blue MOT.

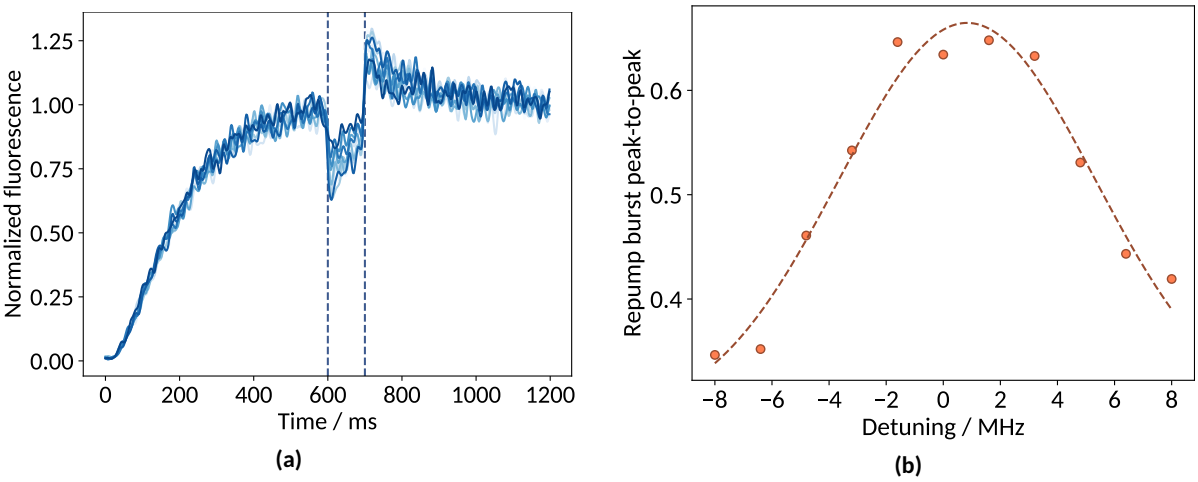

Figure 3.14 – **(a)** Fluorescence measurement using a single photon counter while loading the blue MOT. After 600 ms, a 100 ms long pulse of 689 nm light is applied which shelves atoms to the $^3P_1$ state, reducing the detected fluorescence. After the pulse ends, the atoms return to the ground state leading to a spike in fluorescence. **(b)** Peak-to-peak variation of the fluorescence spike between 550 ms and 750 ms for different probe detunings.

For the first step, we use a photon counter[23] at 461 nm to collect the fluorescence of the blue MOT. We shine in a 100 ms long pulse of 689 nm light while measuring the blue fluorescence. Because the pulse shelves atoms to the $^3P_1$ state, the detected fluorescence signal reduces. After the pulse stops, the atoms can be excited to the $^1P_1$ state again, leading to a spike in fluorescence. An example of the measured fluorescence against time is shown in Figure 3.14. Because atoms in the $^3P_1$ state cannot be lost via the $^1P_1 \rightarrow {}^1D_2$ channel, this method has recently been described as an *atom-number enhancement* in the MOT in [84]. We change the frequency of the shelving pulse and calculate the peak-to-peak variation in the signal around the pulse normalized to the steady state value at 1.2 s. With this measurement we can determine the actual frequency on the 1 MHz level as shown in Figure 3.14. At this level, the maximum Zeeman shift of ~ 3.1 MHz over the size of the blue MOT limits the accuracy of the measurement. We therefore move on to perform spectroscopy of the atomic cloud without magnetic fields which does not show a Zeeman shift but only the Doppler-broadened linewidth. At the stated temperature, the atoms are around for about 2 ms which gives us enough time to collect the fluorescence at 689 nm using the photon counter shown in Figure 3.10. The extracted frequency is accurate enough to be used for the red broadband and single frequency MOTs. After the red MOTs are implemented, we can repeat the procedure using a depletion technique to determine the frequency on the 10 kHz, described at the end of Section 3.7.4.

[23] *Excelitas* SPCM AQRH 15

In the future, a custom-built spectroscopy cell which can to be heated to around 500 °C could help to simplify the process. Standard CF vacuum components, especially common copper gaskets can be heated to 450 °C [70], so care has to be taken when exceeding this limit. Following [85], we designed a cell that features an inner custom-built stainless steel chamber which sits inside an outer chamber built from standard CF-components. The inner cell sits on two *Macor*-rods which thermally decouple it from the outer rest of the system. It uses uncoated sapphire windows which are pressed against the inner chamber using two heated copper flanges to prevent the strontium vapour to escape to the outer chamber. These flanges can be heated to several hundred degrees Celsius and the expansion of the material seals the inner chamber while also heating the metallic strontium inside it. Keeping the windows hot is important to prevent deposition of strontium on the windows which would essentially result in an absorptive coating.

3.7.4 Red MOT

The Doppler-broadened linewidth of 1.5 MHz cannot be compensated by power-broadening alone, as several watts of laser power would be required. We therefore have to insert an additional broadband phase to precool the atoms before we can trap them inside a single frequency operated red MOT, reaching temperatures around 10 µK.

Another aspect to consider is the large mismatch in magnetic field gradients between the blue and red MOTs of 55 G/cm to 4 G/cm. The strong magnetic field gradient in the blue MOT phase leads to a Zeeman shift of more than 10^3 linewidths which, together with most probably velocity of $v_p = 1\,\mathrm{m/s}$, requires the magnetic field gradient to be reduced in around 1 ms. We achieve this by using two power supplies in parallel where we switch the large current using a power MOSFET and dump the inductive voltage spike into a varistor (see Section 3.3). Compared to the blue MOT, the red MOT is also sensitive to offset fields and 0.5 G from the earth's magnetic field can lead to shifts larger than the MOT size. It is therefore crucial to compensate this field appropriately.

3.7.4.1 Broadband MOT

In recent years, a technique called sawtooth-wave adiabatic passage (SWAP)-cooling has been demonstrated in optical molasses [86, 87]. SWAP cooling makes use of an adiabatic transfer to an excited state and cools more efficiently compared to traditional modulation techniques, like triangular ramps or frequency combs. The adiabatic transfer prevents spontaneous emission by forcing stimulated emission from the counter-propagating beam. In result, a larger directed momentum transfer is achieved which leads to faster cooling compared to a process relying on spontaneous emission alone.

Inside a MOT, the magnetic field gradient in combination with the polarization selection rules prevent exploiting stimulated emission. Compared to free-space SWAP cooling, it still relies on spontaneous emission but it suppresses stimulated emission from the same beam that transferred the atom to the excited state [88, 89]. The technique therefore provides colder and denser samples of atoms faster while being less sensitive to frequency and intensity fluctuations.

We start the SWAP-phase 5 ms − 10 ms before the blue MOT is switched off. We use a power of 4 mW in each of the radial beams and 3.5 mW in the axial beam. The latter cannot be retro-reflected and is therefore split into two beams. After capturing the atoms, we reduce the power to 1 mW (1 mW) radially (axially) while keeping the frequency broadening.

3.7.4.2 Single frequency MOT

We stop the SWAP ramp after 50 ms and enter a single frequency phase. We reduce the power instantly to 360 µW (300 µW) radially (axially) and ramp the detuning linearly over 40 ms from −600 kHz to −60 kHz. During that time, the power is also ramped to a final value of 7 µW (5.3 µW). We find the final temperature to be roughly constant for final detunings between −250 to −60 kHz with a difference of around 10 % in the atom number. Towards the resonance, the temperature quickly increases while the atom number decreases. Together with the broadband phase, this results in around 1.5×10^5 atoms at a final temperature of 1.3 µK (1.0 µK) in radial (axial) direction which can be achieved in ~200 ms. The MOT has a size of 110 µm (75 µm) in radial (axial) direction and a central density of $n_0 = 6.9 \times 10^{12}$ cm^{-3}.

3.7.4.3 The red MOT as a frequency reference

The cold sample of atoms serves as an excellent system to perform further spectroscopy on. At 1 µK, the atoms have a velocity of ~ 1 cm/s and can therefore be observed for durations of more than 10 ms[24]. Similar to the spectroscopy on the blue MOT described above, we switch off the magnetic field gradient to avoid any Zeeman shifts. To measure the $^1S_0 \rightarrow {}^1P_1$-transition frequency, we scan the frequency of the imaging pulse while taking an absorption image. We measure the strongest absorption and therefore the *largest* perceived atom number on resonance. For the $^1S_0 \rightarrow {}^3P_1$-transition, we make use of an depletion scheme (see Figure 5.8(a)). We use a 100 µs long pulse at 689 nm to excite atoms to the $^3P_1$ state and take an absorption image at 461 nm. On resonance we therefore detect the *lowest* perceived number of atoms. With this technique we can determine the frequency of the $^1S_0 \rightarrow {}^1P_1$-transition with an accuracy of ~3 MHz and ~5 kHz for the $^1S_0 \rightarrow {}^3P_1$-transition. An example of such a measurement is shown in Fig. 3.15.

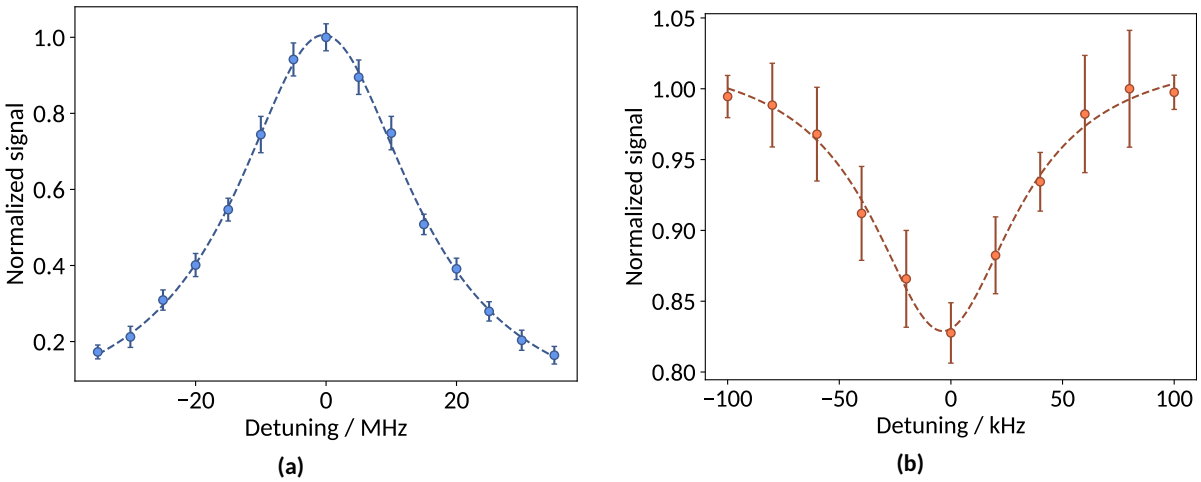

Figure 3.15 – Spectroscopy of **(a)** the $^1S_0 \rightarrow {}^1P_1$- and **(b)** the $^1S_0 \rightarrow {}^3P_1$-transition frequency on the microkelvin cold atom cloud after switching off the magnetic field. This measurement allows us to determine both frequencies on the 3 MHz and ~5 kHz level respectively.

[24] At this point the expansion of the cloud is not the limiting factor but rather the fact that the MOT leaves the field of view due to gravity.

CHAPTER 4

Optical potentials for single atom trapping

The following chapter focuses on the details of creating optical tweezer potentials suitable for single particle trapping and manipulation. In Chapter 2 we introduced the light shift as the origin of the trapping force at the example of the driven two-level system. Chapter 3 covered the experimental setup, specifically the vacuum system and microscope objectives that are used to create the traps and also collect fluorescence light from the atoms. We now delve into the details which are required to project a controllable set of optical traps.

For real atoms, the two-level approximation breaks down and we introduce the *polarizability* to quantify the coupling of between an induced dipole and an electric field. We then motivate the choice of trapping wavelengths used in the experiment based on the polarizability and its consequences. Extending the Gaussian beam model of the trapping beam presented earlier, we show the influence of clipping at finite apertures in real optical systems on the trapping confinement.

In Section 2.4 we motivated the desire to rearrange atoms in an underlying tweezer grid or *array*. It is therefore not enough to use a single unmodulated laser beam to create a trap, but we need the ability to create more than one trap, and additionally, we would like to tune the position of another trap dynamically. Both of these functionalities require a dedicated device and the second half of this chapter is concerned with these beam modulating devices and their description and precise application.. More precisely, we realize static tweezer arrays with a liquid-crystal-based spatial light modulator (LCoS SLM) and a dynamic trap with a pair of acousto-optic deflectors (AODs). Both devices modulate the wavefront of the incident light, which determines the spatial evolution of the beam's amplitude distribution. To quantify to wavefront modulation, this chapter, particularly the section on the liquid-crystal-based spatial light modulator, heavily relies on the Fourier theory of optics. We will keep the discussion concise and reduce the mathematics to a minimum here. A summary of key theoretical insights is given in Appendix A which is itself based on [90]. We close the discussion by combining the knowledge gained throughout this chapter for the design of the tweezer setup used for this work.

4.1 Polarizability and choice of tweezer wavelengths

The two-level system from the description in Section 2.1 is not a good approximation for alkaline-earth(-like) elements, and if more than two levels are involved, all couplings between the levels have to be considered. This effect is described by the polarizability tensor $\alpha\left(\omega, \vec{\epsilon}\right)$ which is related to the induced

dipole moment $\vec{\mu}_e$ by [91]:

$$\vec{\mu}_e = \alpha\left(\omega, \vec{\epsilon}\right) \vec{E} \tag{4.1}$$

with the electric field $\vec{E}$. α depends on the frequency of the light field ω and also on the polarization $\vec{\epsilon}$. For simplicity, we will only look at the scalar part of the polarizability for now by omitting the dependence on $\vec{\epsilon}$ and discuss corrections later on. The potential energy resulting from the coupling is related to α by:

$$V_{\text{dipole}}\left(\omega\right) = -\vec{\mu}_e \cdot \vec{E} = -\text{Re}\left[\alpha\left(\omega\right)\right] |\vec{E}|^2 \tag{4.2}$$

$$= -\text{Re}\left[\alpha\left(\omega\right)\right] \frac{I}{2\epsilon_0 c}, \tag{4.3}$$

where we expressed the absolute square of the electric field by the intensity I with the vacuum polarizability ϵ_0 and the speed of light c. The potential for a state i can be written as a sum over all other states $k \neq i$ by (see e.g. [91]):

$$V_{i,\text{dipole}}\left(\omega\right) = -\sum_k \frac{\omega_{ik} |\left\langle i|\hat{\mu}_e|k\right\rangle|^2 \vec{E}^2}{2\hbar\left(\omega^2 - \omega_{ik}^2\right)} \tag{4.4}$$

which gives

$$\text{Re}\left[\alpha\left(\omega\right)\right] = \sum_k \frac{2\omega_{ik} |\left\langle i|\hat{\mu}_e|k\right\rangle|^2}{\hbar\left(\omega^2 - \omega_{ik}^2\right)}. \tag{4.5}$$

Compared to the two-level system, it is now possible for two states related by an optical transition to be trapped. The polarizability is often expressed in the unit-system of *atomic units* [a.u.] which can be converted to SI-units with the factor $1/4\pi\epsilon_0 a_0^3$ with the Bohr-radius a_0. However, this is again only *almost* the final truth. For states featuring a non-zero total angular momentum J, the polarization of the trapping light and additional static magnetic fields also play a role. This results in different components of the polarizability, namely the scalar α_s, vector α_v and tensor polarizabilities α_t. The polarizability in strontium is of great importance for the optical clock community and we refer to [91, 92] for further details on the calculation and [93–95] for results for strontium.

For this work we are interested in the scalar polarizability α_g of the ground state $^1\text{S}_0$ which is necessary to estimate trap depths from the laser power per trap. Using data from [95], we calculate α_g which we show in Figure 4.1(a). Note that only dipole-allowed transitions contribute to the value, which is why α_g only shows a divergence at 461 nm from the $^1\text{S}_0 \rightarrow {}^1\text{P}_1$-transition but not at 689 nm from the $^1\text{S}_0 \rightarrow {}^3\text{P}_1$-transition. The values for important wavelengths for this work, together with values obtained from the literature, are shown in Table 4.1.

Another important quantity is the difference in polarizabilities $\Delta\alpha$ ("differential light shift") for two given states, if driving a transition between said states is of interest. Depending on the value (and especially the sign) of $\Delta\alpha$, the transition frequency changes compared to the free-space resonance, which in turn can affect the cooling and imaging inside the trap drastically. We calculate the difference

between scalar polarizabilities α_g and α_e for the excited state $^3P_1$ with $|m| = 1$[1], to motivate the choice of our trapping wavelengths. Combined with the 10 kHz-wide transition from $^1S_0$ to $^3P_1$, the quantized motional state of an atom inside the trapping potential can be optically resolved.

Figure 4.1(b) shows the resulting differential polarizability $\Delta\alpha$. It can be seen, that there are two crossings where the difference vanishes: around 515 nm and around 915 nm. These wavelengths are also referred to as *magic wavelengths* and a trapped particle experiences the same potential in both states. This situation can be used to cool an atom and also directly measure the motional state in the trap, a feature that we will investigate in Section 5.4.1[2]. An important aspect is the sign of the difference of the ground and excited state polarizabilities, as it changes the trapping and cooling characteristics depending on the chosen tweezer wavelength. In this work we will make use of three different trapping wavelengths which we now briefly introduce.

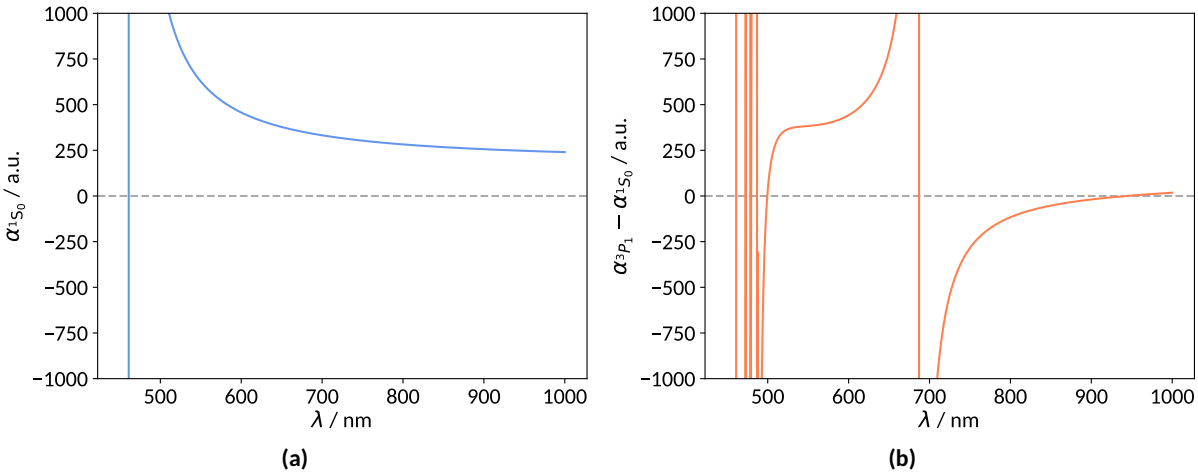

Figure 4.1 – **(a)** Polarizability of the $^1S_0$ ground state. The divergence at 461 nm is caused by the $^1S_0 \rightarrow {}^1P_1$-transition. As the model only includes dipole-allowed transitions, there is no divergence at 689 nm from the $^1S_0 \rightarrow {}^3P_1$-transition. **(b)** Scalar differential light shift $\Delta\alpha$, calculated for the $^1S_0$ and $^3P_1$ states. The divergence at 461 nm from the $^1S_0 \rightarrow {}^1P_1$-transition is visible, and in addition a divergence at 688 nm from the $^3P_1 \rightarrow {}^3S_1$-transition. If vector- and tensor-components are included, the zero-crossings lie at 515 nm and 915 nm respectively.

Wavelength / nm	Our result / a.u.	Literature / a.u.
515	930	910 to 950 [46], 900 [47]
532	745	750 [96]
813	280	286 [56, 93], 280 [96]

Table 4.1 – $^1S_0$ ground state polarizability together with values from the literature.

[1] As discussed in the introduction, the narrow intercombination line in alkaline-earth(-like) elements offer a great advantage for cooling, spectroscopy, and even imaging compared to standard 10 MHz-wide $s \rightarrow p$ transitions in alkali elements.

[2] Due to vector and tensor components the tweezer polarization and external magnetic fields have to be fine-tuned to realize $\Delta\alpha = 0$

515 nm Around 515 nm, the differential light shift $\Delta\alpha = \alpha\left({}^3\mathrm{P}_1\right) - \alpha\left({}^1\mathrm{S}_0\right)$ vanishes ("magic wavelength"). Additionally, the large ground state polarizability allows for deep traps due to the small waist and large polarizability reducing the power requirement per trap. It is also beneficial for the atom rearrangement described in Chapter 6, where a deep tweezer is used to transport an atom between two shallow traps at 813 nm to achieve an atom pick-up, without turning of the static traps. A disadvantage at this wavelength are multiple known loss channels, as the trapping wavelength is not red-detuned to all transitions. First, the $5s4d\,{}^1\mathrm{D}_2$ state, which is populated by a decay from ${}^1\mathrm{P}_1$, is antitrapped at this wavelength which leads to atom loss due to the long lifetime of 300 µs. As imaging is performed on the ${}^1\mathrm{S}_0 \rightarrow {}^1\mathrm{P}_1$-transition, this leads to an increased loss during the imaging process. Recently, a possible repumping scheme inside a MOT for this state has been identified [97], but it has not been investigated in tweezers yet. Furthermore, off-resonant excitation from ${}^3\mathrm{P}_1$ to $5s5d\,{}^3\mathrm{D}_2$ by trap photons can also lead to atom loss during cooling or imaging [47]. The currently established way is therefore to combine 515 nm with a second (near-infrared) wavelength with fewer loss channels at the expense of shallower traps.

813 nm 813 nm is the most well investigated trapping wavelength for strontium, as the differential light shift for the ${}^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$-transition in ${}^{87}\mathrm{Sr}$ vanishes here. In result, it has been investigated for almost two decades for the creation of optical lattices and lattice clocks with strontium [98, 99] and in recent years also for single strontium atoms in optical tweezers [48] and tweezer clocks [53, 55]. Compared to 515 nm, 813 nm is red-detuned to all lower-lying transitions which closes the loss-channels mentioned above and reduces the loss during the imaging process. However, the lower polarizability and larger wavelength also leads to a high power requirement per trap (or shallower traps at similar powers). Although we currently do not plan to work with ${}^{87}\mathrm{Sr}$ and the ${}^3\mathrm{P}_0$ clock state, working at 813 nm keeps this as an option for the future.

532 nm 532 nm can be seen as an alternative to 515 nm, where similarities include the short wavelength and large polarizability but also the loss-channels. The largest limitation compared to 515 nm is the fact the condition of a vanishing differential light shift cannot be realized which blocks access to the experimental tools described above. The benefit that could compensate for this drawback is the fact that 532 nm is a common wavelength used in other cold-atom experiments to create dipole traps, and high-power single-mode lasers and optical elements are readily available. Up to now, strontium has only been investigated in low NA tweezers [100] and the cooling measurements in this work are therefore the first time where strontium is investigated in high NA tweezers at this wavelength. This makes 532 nm a more "experimental" choice but progress at this wavelength for (strontium) tweezer experiments could be of great benefit for the community in the future.

Mentionable future options In addition to the wavelengths mentioned above, two additional candidates that (to our current knowledge) have not yet been explored, are 915 nm and 1 064 nm. First, 1 064 nm could be a viable candidate for replacing 813 nm, especially if the clock state in ${}^{87}\mathrm{Sr}$ is not of interest. 1 064 nm is a common wavelength in cold-atom research and thus shows similar availability for optical components and high-power single-mode lasers as 532 nm. Compared to 813 nm, 1 064 nm has a positive $\Delta\alpha$ for ${}^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ which requires a cooling scheme similar to 532 nm. A lower polarizability than 813 nm requires even more power per trap, but on the other hand lasers with several tens watts of output power are available. The benefit would therefore be similar to 532 nm: Progress at an established wavelength which could lead to reduced complexity and costs. Secondly, the second zero crossing for $\Delta\alpha$

at 915 nm could be used to perform resolved sideband spectroscopy and cooling at this wavelength. This makes 915 nm the most interesting candidate as an alternative to 813 nm from a scientific perspective. The largest drawback is the wavelength itself where high power ($\gtrsim$ 5 W) lasers are not commercially available. However, 915 nm also lies inside the range of the anti-reflective optical coatings of the vacuum windows and the objective and the latter shows nearly diffraction limited performance at this wavelength. It is therefore possible to add the wavelength later on, to investigate it alongside other options.

4.2 Trap potential simulation - Fourier optics calculations

Just as the polarizability extends the rather idealized model of a two-level system to a real atom, this section is dedicated to exploring the experimental realization of an optical tweezer where clipping at finite-sized apertures or finite laser powers has to be considered. Starting from Equation (2.10) with $V \propto I$, the potential close to the center can be approximated as harmonic by:

$$V(r,z) = V_0 + \frac{1}{2}\frac{\partial^2 V}{\partial r^2}\bigg|_{r_0,z_0} r^2 + \frac{1}{2}\frac{\partial^2 V}{\partial z^2}\bigg|_{r_0,z_0} z^2 + \ldots, \tag{4.6}$$

which is shown in Figure 4.2. Comparing the coefficients for the radial and axial direction with the

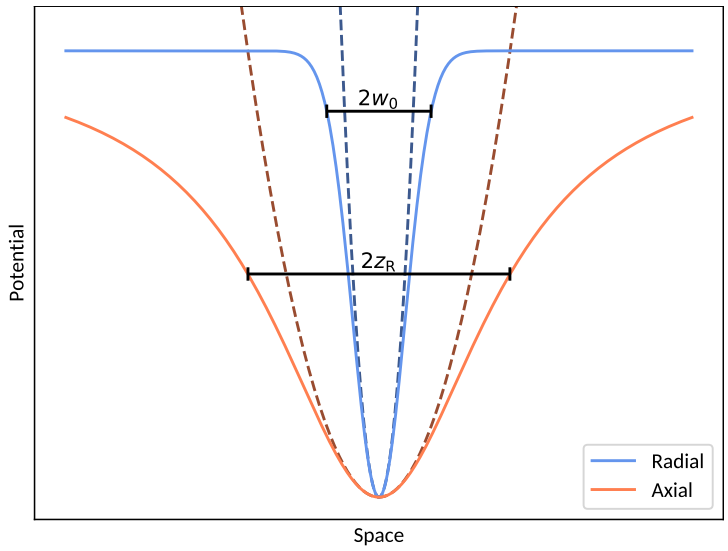

Figure 4.2 – Harmonic approximation of a Gaussian potential in radial and axial direction for $w_0 = 0.8\lambda$.

standard harmonic oscillator potential $V(x) = \frac{1}{2}m\omega^2 x^2$ allows us to deduce the trap frequencies ω_r and ω_a as:

$$\omega_r = \sqrt{\frac{4|V_0|}{m w_0^2}} \quad \text{and} \quad \omega_a = \sqrt{\frac{2|V_0|}{m z_\mathrm{R}^2}}. \tag{4.7}$$

By measuring both trap frequencies we can then fully characterize the trap by calculating the waist

$$w_0 = \frac{\omega_r}{\omega_a} \frac{\lambda}{\sqrt{2}\pi} \tag{4.8}$$

and in result the trap depth V_0 with either frequency from Equation (4.7). Different methods to measure the trap frequencies exist, and we will make use of resolved sideband spectroscopy in Section 5.4.1 and parametric excitation in Section 5.7.

We would now like to investigate which values for experimental parameters, for example the potential depth V_0 or the waist w_0, are preferable. In order to do this, we define the efficiencies ϵ_r and ϵ_a as the second derivative of the potential at the position of the focus. This quantity combines the trap depth, trap frequency, the spot size and the power loss and can therefore be used to determine the parameters which maximize performance. In order to determine ϵ_i, however, we first have to understand how to calculate the trapping potential from an initial electric field distribution E. We already know from Chapter 2, that a lens is required to achieve a small focal spot and we will quantify this statement now.

4.2.1 Numerical calculations - the Fourier lens

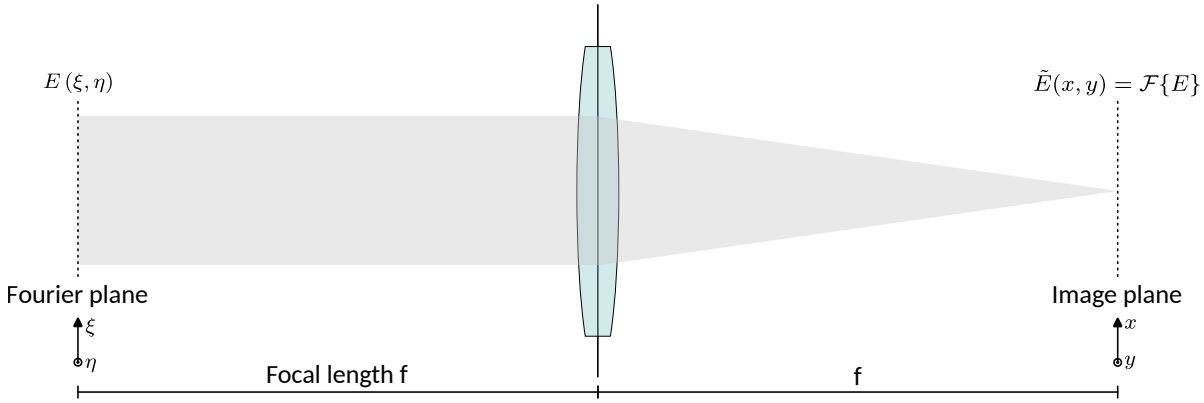

Figure 4.3 – An idealized lens performs a Fourier transformation $\mathcal{F}$ of the electric field distribution in the front focal plane (Fourier plane) to the back focal plane (image plane).

Starting from Kirchhoff's diffraction law, it can be shown that (with certain assumptions) an object that imprints a parabolic phase-change on an incident wavefront performs the mathematical operations of a Fourier transformation. The parabolic phase-changing object is much more commonly known as a *lens*. Appendix A summarizes the main results of the theory of Fourier optics where we also briefly comment on the numerical methods. In summary, the electric field distribution (and therefore the intensity distribution) in the focal- or image-plane is related to the electric field distribution in the lens- or Fourier-plane by a Fourier transformation. We can therefore numerically compute the electric field of a focused beam from the electric field distribution of the incident beam and investigate the trapping parameters derived from the spot size.

4.2.2 Truncation effects

The calculation from the beginning of this section holds true, if the intensity profile can be assumed to be Gaussian without significant truncation effects[3]. In reality, optical elements have finite clear apertures, always resulting in truncation to some degree. We can describe this truncation by a truncation factor ζ, which we define as the ratio of beam waist w ($1/e^2$-radius) to aperture radius r_{ap}, thus $\zeta = w/r_{\mathrm{ap}}$. The limiting aperture can be located anywhere in the beam and essentially determines the NA of the final lens. We can therefore assume, that the aperture is given by the final lens as this determines the achievable spot size. In this picture, the ideal Gaussian beam is described by $\zeta = 0$ and a flat-top beam is described by $\zeta = \infty$. The former produces a perfectly Gaussian spot size, while the latter produces a sinc^2-shaped intensity profile in one dimension and an Airy disk in two dimensions. The Gaussian beam case has perfect power efficiency as no power is clipped at the aperture, but it does not produce the smallest spot size, leading to lower trap frequencies. On the other hand, the flat-top beam achieves the smallest possible potential curvature in the center, at the expense infinite power loss. Furthermore, side-fringes of the Airy disk might form unwanted additional potentials which can lead to interference between traps. These two cases can therefore be seen as limiting cases of the efficiency scale defined above, and in reality on would therefore like to work in an intermediate regime between the two aforementioned edge cases, balancing power loss, spot size and side-fringe depth.

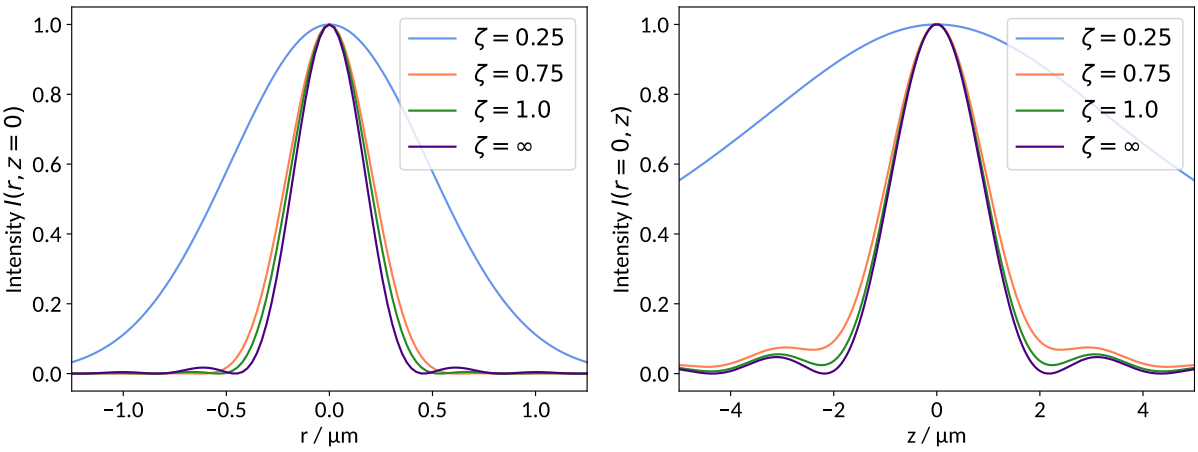

Figure 4.4 – Radial **(a)** and axial **(b)** intensity distribution of a focussed Gaussian beam with initial waist w truncated at an aperture with radius r_{ap} for different truncations ratios $\zeta = w/r_{\mathrm{ap}}$.

To determine the value of the truncation ratio that we would like to work with, we calculate the radial- and axial intensity profiles in the focal plane numerically for a beam with size w, clipped at an aperture with radius r_{ap}, focussed down by a lens with focal length f. From the intensity profile we can extract the trap depth U, trap frequency ω and efficiency ϵ, all in arbitrary units. The power loss $\mathcal{L}$ is given by the transmitted power of a Gaussian beam through a circular disk. The efficiency ϵ combines the trap depth U, given by the power loss $\mathcal{L}$ and the trap frequency ω, given by the curvature of the potential.

[3] We only consider beams with an initial intensity profile following the Gaussian fundamental mode TEM_{00}.

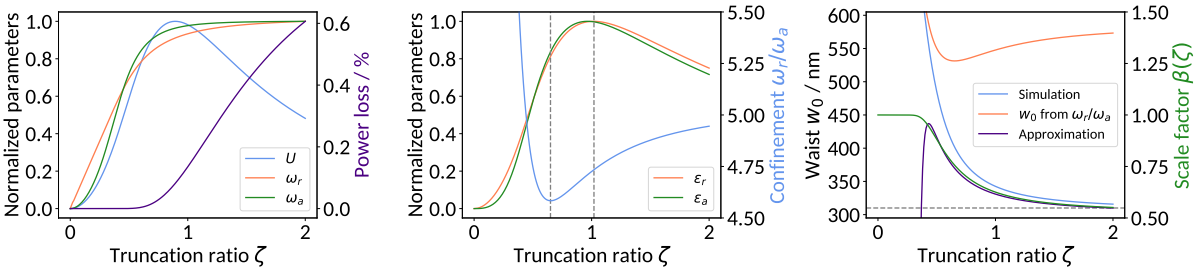

Figure 4.5 – Trap performance considering diffraction at a two-dimensional aperture for different truncation ratios ζ. Fixed parameters are an NA of 0.7 and a wavelength of 515 nm. **Left:** Trap depth, trap frequencies together with the power loss $\mathcal{L}$. For low ζ, no power is lost but the achieved trap frequencies are also low. For large ζ, the trap frequencies increase but trap depth decreases due to increasing power loss. **Center:** Deriving the efficiency ϵ and confinement ω_r/ω_a from the trap depth and the trap frequencies shows that $\zeta \in [0.6, 1.0]$ achieves the best confinement for close to maximum efficiency. **Right:** Comparison of the $1/e^2$-waist extracted from a simulation and the waist calculated from $w_0 = \omega_r/\omega_a \lambda/\sqrt{2}\pi$ for a wavelength of $\lambda = 515$ nm. The green curve shows the ratio β of the simulation and the waist calculated from the aspect ratio. The violet curve shows a common approximation [101] and the grey dashed line shows the limit for $\zeta \to \infty$ of 310 nm.

We calculate U, ω, $\mathcal{L}$ and ϵ from the intensity distribution in the focal plane $I(r, z, \zeta)$ by:

$$U(\zeta) = I(0, 0, \zeta) \tag{4.9}$$

$$\omega_i(\zeta) = \sqrt{-\frac{1}{U(\zeta)} \frac{\partial^2 I(r, z, \zeta)}{\partial i^2}}\Bigg|_{r=0, z=0} \quad \text{for } i \in \{r, z\} \tag{4.10}$$

$$\mathcal{L}(\zeta) = \exp\left[-2/\zeta^2\right] \tag{4.11}$$

$$\epsilon_i(\zeta) = \sqrt{-\frac{\partial^2 I(r, z, \zeta)}{\partial i^2}}\Bigg|_{r=0, z=0} \quad \text{for } i \in \{r, z\} \tag{4.12}$$

where we assume $\omega_z = \omega_a$. The electric field on the lens is given by

$$E(\rho) = \frac{1}{\sqrt{\sqrt{\pi/2} w}} \exp\left[-\frac{\rho^2}{w^2}\right] \times P(\rho) \text{ with } P(\rho) = \begin{cases} 1 & |\rho| < r_{\text{ap}} \\ 0 & \text{else} \end{cases} \tag{4.13}$$

for a beam of waist w, where ρ is the radial coordinate in the Fourier plane. The resulting intensity distribution in the focal plane is then given by:

$$I(r, z) = |\mathcal{F}\{E\}|^2(r, z)\,. \tag{4.14}$$

We have dropped the w-dependence since we start with a normalized electric field in Equation (4.13) *before* applying the truncation and therefore only the truncation ratio ζ matters. The result is shown in Figure 4.5. It can be seen, that for $\zeta \lesssim 0.4$ virtually no power is lost, but the trap frequency also only increases to about 60 % of its maximum value. Maximum efficiency is reached for $\zeta = 1.0$ ($\zeta = 0.95$) for the radial (axial) direction where the combination of trap frequency to trap depth reaches its maximum. After that point, the trap frequency does still increase but at the expense of power losses larger than ~40 %. The strongest confinement, expressed by the ratio of trap frequencies ω_r/ω_a, is reached at

$\zeta = 0.65$, far below the Airy limit of $\zeta \to \infty$.

To still be able to estimate the trap size and depth from the trap frequencies ω_r and ω_a via Equation (4.8), we introduce a scale factor $\beta(\zeta)$ that corrects for the deviation from the Gaussian case. We determine $\beta(\zeta)$ by extracting the ratio ω_r/ω_a and the $1/e^2$-waist $w_{0,s}$ by simulating the intensity distribution for different ζ. The scale factor $\beta(\zeta)$ is then given by the ratio $w_{0,f}/w_{0,g}$ where $w_{0,g}$ is calculated from Equation (4.8). The rightmost plot in Figure 4.5 shows the result for a wavelength of 515 nm. In the limit of $\zeta \to \infty$, a possible approximation for the waist is $w_{0,\text{approx}} \approx 0.84\lambda N$ with the f-number N and the wavelength λ [102]. N is given by f/D with the focal length f and diameter of the lens D and can therefore be calculated from the NA as $N = \frac{1}{2\text{NA}}$. Using this approximation we calculate a waist of $w_{0,\text{approx}} \approx 310$ nm which is shown as the black dashed line in Figure 4.5 and compares well to the simulation.

The simulation represents the best-case scenario, as every real optical system will contain additional aberrations due to misalignments or imperfect optical elements. In this case the simulation can serve as a benchmark to identify misalignments in the optical setup.

4.3 Liquid-crystal based spatial light modulation

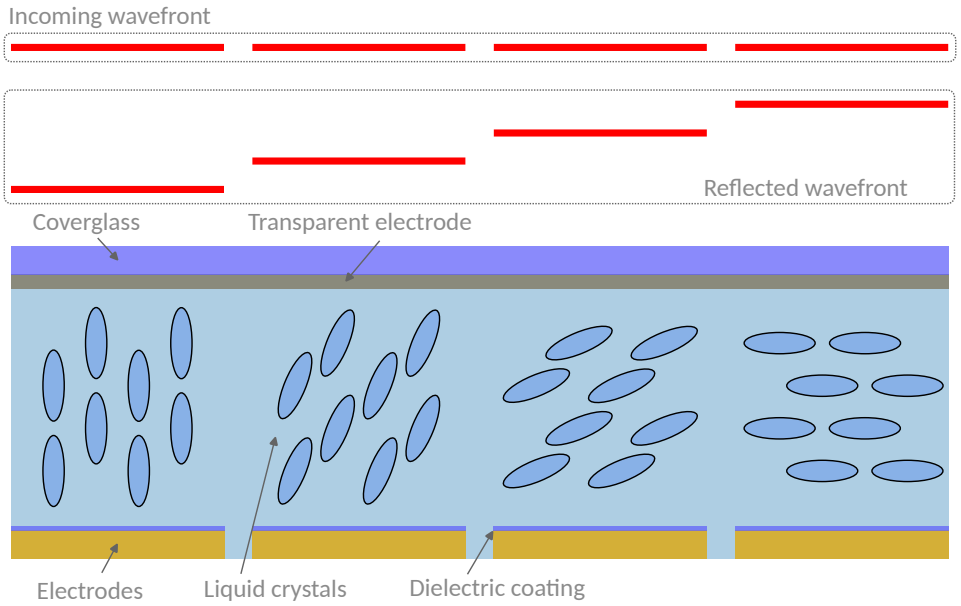

Figure 4.6 – Liquid-crystal based spatial light modulator. Applying a voltage to the electrodes rotates the liquid-crystals, leading to a change in refractive index and therefore optical path lengths. The back electrode is divided into pixels, similar to a liquid-crystal display (LCD), so a spatially varying phase pattern ("phasemask") can be imprinted on the incident wavefront.

The LCoS SLM is a type of liquid crystal display and the general working principle is shown in Figure 4.6. By changing an externally applied voltage, the orientation of the crystal can be rotated, which rotates the polarization axis and in turn changes the refractive index due to the crystal's anisotropy. To achieve a locally varying refractive index, the electrode on the backside is separated into pixels, with common pitches around 10 to 20 µm, resulting in filling fraction > 95 %. When light passes through the crystals, the phase-delay of the wavefront depends on the local crystal orientation, which can be used

to change the wavefront globally. As the liquid crystal layer itself is not pixelated, the effect washes out at the border between adjacent pixels where the electric field is the sum of values at the center of the pixels. This so called *pixel-crosstalk* limits the achievable efficiency of the device. To sustain the dynamic operation the pixels are toggled with frequencies around 1 kHz, otherwise a stable crystal orientation cannot be maintained. The device is used similar to a standard LCD monitor, where an image or *phasemask* is written to the device, which changes the voltage at every pixel electrode. In contrast to standard RGB(A) images, the phasemasks are monochromatic ("grayscale") with an 8 bit modulation depth. The direct mapping of the digital value to applied voltage, which in turn results in a nonlinear change of the refractive index, requires a linearization step, see Section 4.3.1.

For this work we use two SLMs, a full-HD 1 920 px × 1 152 px SLM by Meadowlark[4] for 785 nm[5] and an older model by Hamamatsu[6] with a resolution of 800 px × 600 px with a dielectric coating for (800 ± 50) nm. The following calibration was carried out for the SLM by Meadowlark at 515 nm and 813 nm. During the measurements in Chapter 5, we noticed a strange behavior with the SLM by Meadowlark where the pattern was not stable in time anymore. This resulted in fluctuations in the tweezer arrays and finally atom loss. Meadowlark could not determine the root cause of the issues, and we assume that the SLM is damaged in a way that results in increased phase noise, rendering it inadequate for tweezer projection. We therefore continued with the SLM from Hamamatsu.

Fourier plane operation

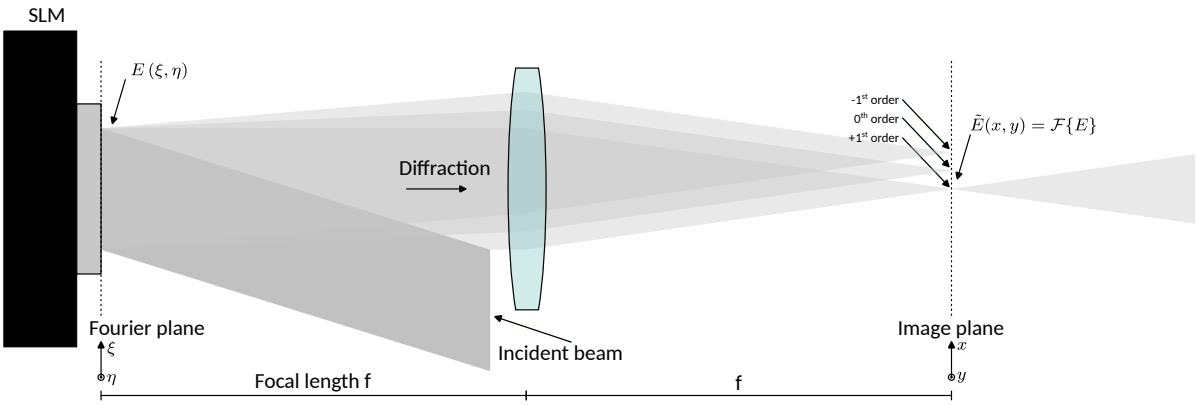

Figure 4.7 – Fourier plane operation of the SLM. A diffractive phase grating on the SLM is used to create different diffraction orders, where the +1st order is aligned to the optical axis of the Fourier lens. With the depth of the grating and its slope the position and amplitude of a diffracted spot can be controlled. Unwanted orders can be blocked by an aperture in the focal plane.

Two methods exist to operate the SLM, one directly in reflection and one in the Fourier plane. The difference is, that in the first case the SLM has to operate in the far-field, which is impractical for our application (see Appendix A). In Fourier plane operation the SLM is placed in the front focal plane of a lens and illuminated with a collimated beam. The incident angle should be as small as possible, otherwise

[4] HSP1920-600-1300-HSP8

[5] The SLM has a broadband coating and is usable between 500 nm and 1 000 nm. The backplane is made from aluminium which reduces the reflectivity to around 60 % around 800 nm.

[6] X10468-02

the efficiency will be reduced as a photons passes through multiple pixels before being reflected by the backplane. As described in Section 4.2.1, we calculate the electric field distribution in the focal plane by Fourier transforming the electric field distribution on the SLM. In this case, the beam amplitude distribution is given by the electric field modulus of the laser beam. The phase-pattern, however, is given by the phasemask displayed on the SLM. A key result that helps to understand the operation of the SLM, is the relation between the deflection angle α and the displacement Δx it causes in the image plane (see Appendix A.1.4). By changing the slope of the wavefront written to the SLM we can therefore move the focus in radial direction in the focal plane.

It is important to ensure that the mapping from the grayscale levels of the digital phasemask to the analog pixel voltage (and in turn the phase delay of the pixels) correctly matches the expectation. Furthermore, the phase pattern we apply later on to create tweezer arrays only yields the desired result if the wavefront of the diffracted laser beam does not have other contributions. The latter could stem from optical aberrations in the system and can severely alter the change of the electric field distribution in the focal plane.

The first step, before creating any tweezers or tweezer arrays, is therefore to calibrate the SLM for the wavelengths used in the experiment to ensure it performs as intended.

4.3.1 Calibration

As mentioned before, the first calibration requires us to determine and correct the linear mapping from $[0, 256) \rightarrow [0, 2\pi)$. This calibration is only necessary for the Meadowlark SLM, as the Hamamatsu SLM is already internally calibrated for the correct wavelength.

In a second calibration step, we need to measure the wavefront interferometrically, and we use the result as a correction that we apply to the SLM to *flatten* the wavefront. Compared to the first calibration step, this one applies to both SLMs.

4.3.1.1 Phase response

To measure the mapping from $[0, 256)$ to $[0, 2\pi)$, we make use of a binary grating and the associated change in intensity when the depth of the grating is changed. Imagine a phase grating like the one shown in Fig. 4.8. By varying the depth of the grating $\Delta\phi = \phi_{\max} - \phi_{\min}$, the intensity in different orders will change. Figure 4.9(a) shows the dependence of the 1^{st}- order intensity on $\Delta\phi$. The functional form of this variation is that of a squared sine and can be intuitively understood by the fact that after $\Delta\phi = \pi$ is reached, the effective grating depth decreases again, until a flat wavefront is reached at $\Delta\phi = 2\pi$ (see also end of Appendix A.3). It can be also seen, that the measured intensity dependence starkly differs from the expectation, where the first minimum, corresponding to a grating depth of 2π, is already reached at ~ 160 compared to 256. Additionally, the peak which corresponds to a grating depth of π is not reached at the expected value of 128 but also not at $160/2$, signalling the nonlinearity of the mapping.

In order to correct for this nonlinearity, a lookup table (LUT) is generated which maps the 8 bit input value to the correct output-value that linearizes the pixel stroke. For our Meadowlark-SLM , the 8 bit input value is internally converted to a 11 bit value. This is necessary to preserve the full 8 bit resolution once the linearization LUT is applied. If we were to correct the mapping in software, we would reduce our resolution by about one third ($160/255$). By expanding to 11 bit in hardware, we still have more than four times of the required phase levels left ($2/3 \cdot 2^{11} \approx 1365$).

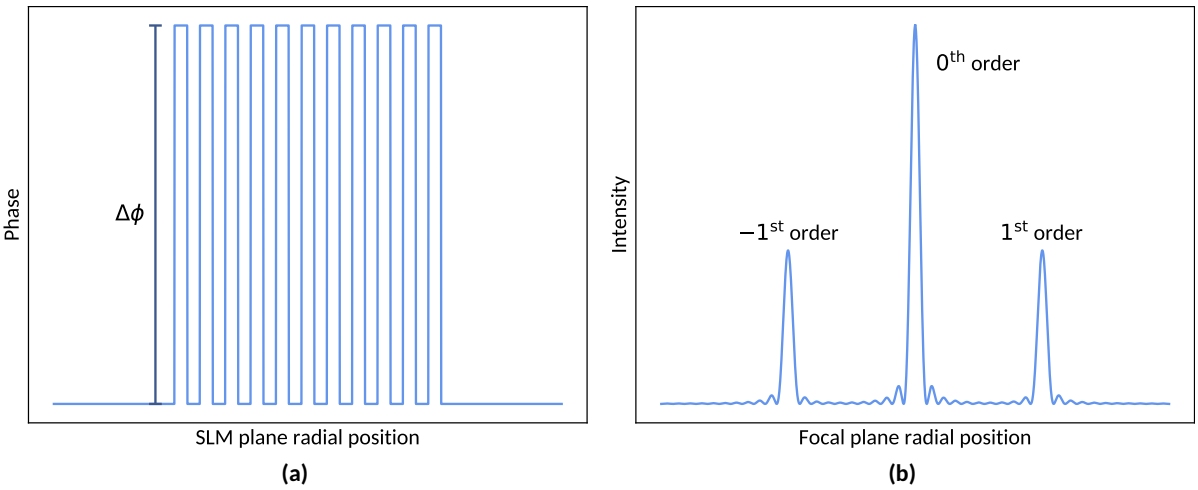

Figure 4.8 – Binary phase grating in the SLM plane **(a)** and the image plane **(b)**.

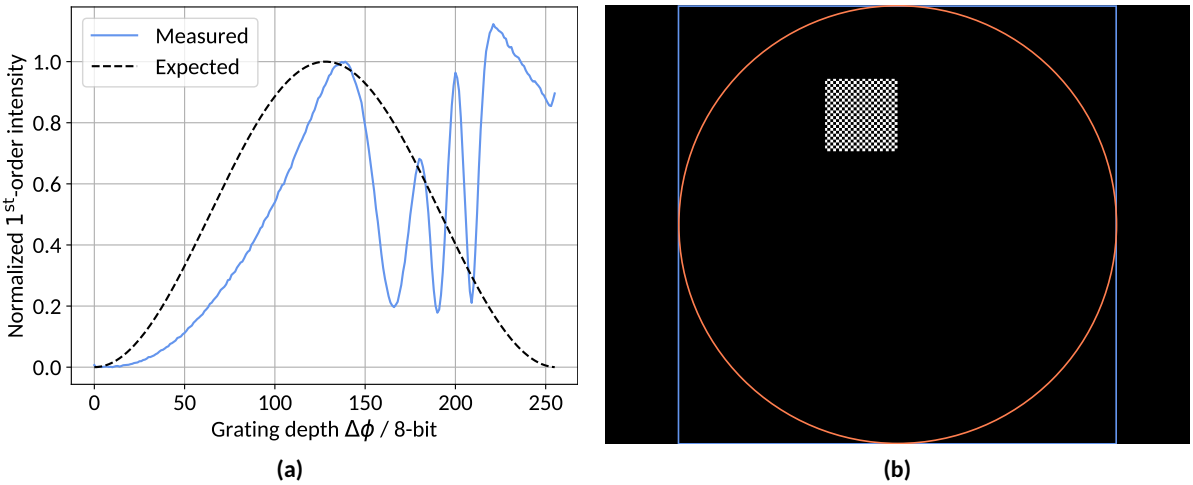

Figure 4.9 – **(a)** Shows the expected and measured first order intensity dependence on the binary grating depth. The expected function form follows a sinusoid. The deviation in the measured intensity can be explained by the (spatially varying) non-linear phase response of the liquid crystals. **(b)** To determine the spatial dependence, the relevant area of the sensor (blue square) sensor is divided into smaller areas (checkerboard pattern) which are then measured sequentially. The checkerboard pattern is the result of a superimposing two binary grating along two orthogonal directions. The red circle shows the aperture that is used for tweezer operation.

The linearization step will procede in a two-step process. The center region of the sensor is divided into 9×9 smaller sections[7] as shown in Figure 4.9(b). In a first step, we will compensate for the nonlinearity on a global level by averaging the measurements from all subregions, but as the nonlinearity varies from pixel to pixel, we have to repeat the correction locally.

[7] For the center square area of the SLM of 1 152 px × 1 152 px this results in patches of 128 px × 128 px.

Global correction To extract the underlying (averaged) phase, the inverse of the analytically known form is applied:

$$\phi = 2 \arcsin(\sqrt{I/I_{\max}}), \tag{4.15}$$

where $I/I_{\max}$ is the intensity normalized to the maximum value. Due to slightly different nonlinearities, the sum of many pixels will show a dephasing effect, which in turn will produce a curve between 0 and $\phi_{\max} < 2\pi$. To correct for this, the phases are rescaled such that $\phi_{\max} = 2\pi$.

Figure 4.10 shows the computed correction for different wavelengths and Figure 4.11(a) shows the dependence with and without the global lookup table applied. It can be seen that although the endpoints of the intervals $[0, 256) \rightarrow [0, 2\pi)$ are now well aligned, the inner part varies significantly for different regions on the sensor. A second fix-point for the calibration is the peak position which should lie at $\pi \hat{=} 128$. The deviation from the π-position with and without the global correction is shown in Figures 4.11(b) and 4.11(d). It can be seen, that the correction lowered the difference from the π-position from $(68 \pm 5)\,\%$ to $(13 \pm 8)\,\%$. We will therefore now apply a second correction where we again divide the sensor in smaller segments and determine the nonlinearity locally *without* the averaging of the first step.

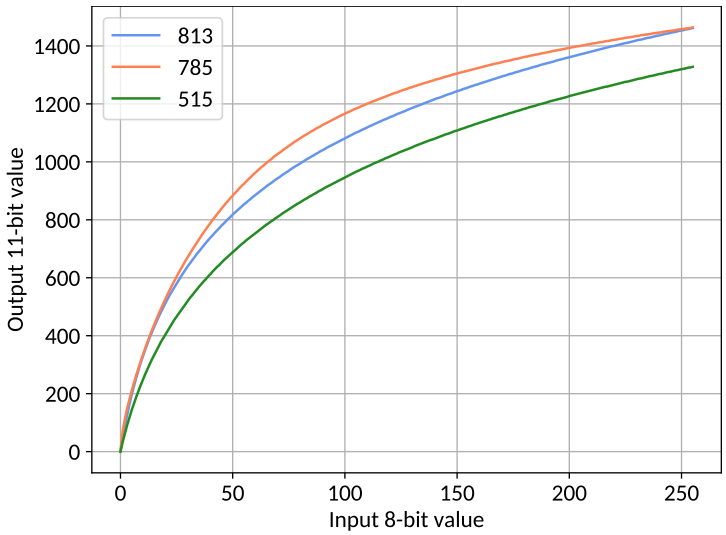

Figure 4.10 – Lookup tables compensating the nonlinear pixel response for different wavelengths. The data for 785 nm is supplied by Meadowlark.

Local correction To further linearize the mapping, we apply a second, spatially resolved LUT which maps input values in the interval $[0, 256)$ to output values in the interval $[0, 256)$. This LUT is only applied in software, i.e. the grayscale values in a phasemask are replaced before it is written to the SLM where the hardware LUT is applied. This slightly reduces the resolution as some of the grayscale values are mapped to the same value. We measure the loss of resolution to be below 5 %, by counting the average number of unique output grayscale values in the calibrated region. With the global and local correction we now get a deviation of $(5 \pm 5)\,\%$ from the real π-position .

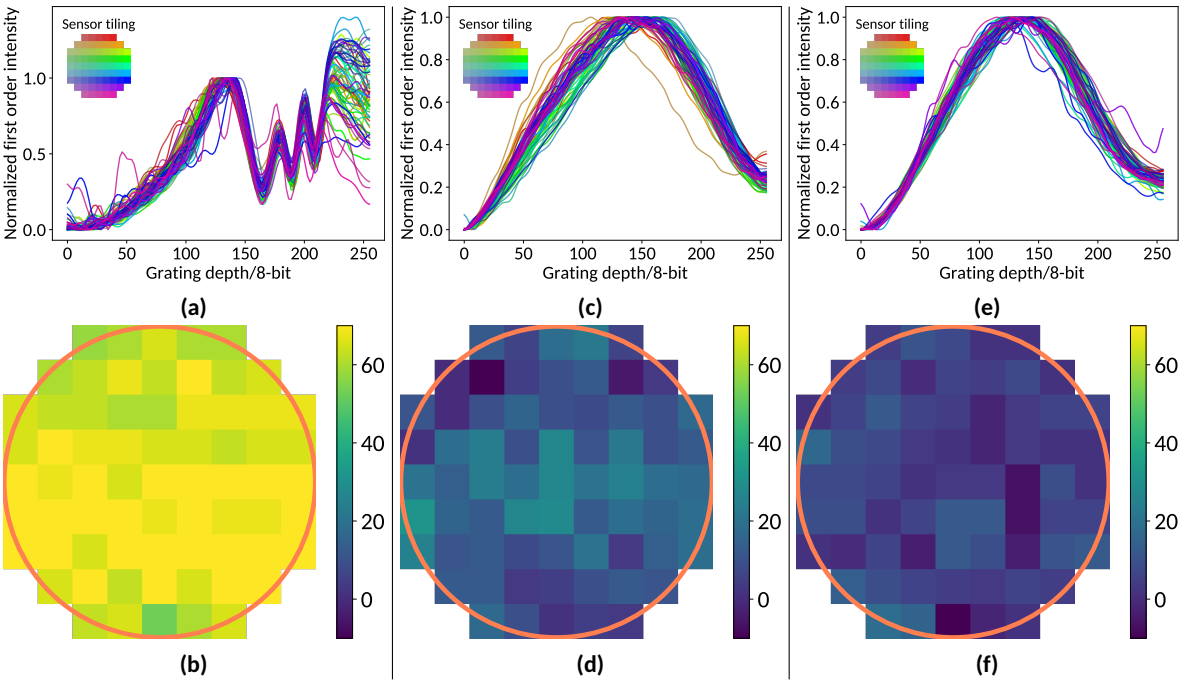

Figure 4.11 – Measurement of the first order intensity dependence on the applied grayscale level without correction **(a)**/**(b)**, with a global correction **(c)**/**(d)** and with a global and a local correction **(e)**/**(f)**. The lower row shows the spatially resolved difference between the measured and the real π-position at a grayscale value of 128 in percent. The red circle shows the center aperture that is used for tweezer creation. From initially (68 ± 5) %, the deviation is reduced to (13 ± 8) % by the global correction and to (5 ± 5) % with the local correction.

4.3.1.2 Wavefront flattening

Aberrations present in every real optical system can lead to severe image distortions which in our case manifest in low intensity efficiencies and imperfect point spread functions for the tweezer array. In the following section, we will present a method to measure and correct these aberrations. It is based on the idea described in [103], where interference of light refracted by different regions of the sensor is used to directly measure the wavefront. Subtracting this measured wavefront from the pattern on the SLM can reduce the peak-to-peak wavefront deviation from $\sim \lambda$ down to $\lambda/10$[8].

Different regions of the sensor are compared to a fixed reference which results in a map of relative phases (a constant global offset can be ignored). To isolate the contribution of different sensor regions, we write a blazed phase grating to the regions of interest ("patch") and observe the interference pattern in first order. Blazed gratings play a key role for the operation of the SLM as they allow to separate different diffraction orders with higher efficiencies than a binary grating. More details for the different gratings are given in Appendix A.3. To measure the wavefront, we now have to extract the phase of the interference pattern for every two-patch-image.

Figure 4.12(a) shows the two patches on the sensor and 4.12(b) shows the corresponding interference pattern in the image plane. We now want to extract the phase of the interference pattern reliably and

[8] This value depends strongly on the number of optical elements and the length of the beam path. $\lambda/10$ can be reached in the first Fourier plane after the SLM, however for the full experimental setup we measure a value $\lambda/6$ right before the vacuum chamber.

(a) **(b)**

Figure 4.12 – **(a)**Two gratings in SLM plane. **(b)** Interference pattern in image plane. We find no difference in performance between circular- and square-shaped gratings.

efficiently. For this we model the pattern using the following functional form:

$$f(x,y) = A \exp\left(-\frac{\left(x-\mu_\mathrm{x}\right)^2}{2\sigma_\mathrm{x}^2} - \frac{\left(y-\mu_\mathrm{y}\right)^2}{2\sigma_\mathrm{y}^2}\right) \tag{4.16}$$

$$\times \left[1 + C\cos\left(2\pi\left(k_\mathrm{x}\left(x-\mu_\mathrm{x}\right) + k_\mathrm{y}\left(y-\mu_\mathrm{y}\right)\right) + \phi\right)\right], \tag{4.17}$$

where A is the amplitude and C is the contrast of the pattern. The above equation is only an approximation, as the real envelope is either an Airy-disk for circular patches, or a two-dimensional sinc^2 for square/rectangular patches instead of a two-dimensional Gaussian function. This leads to deviations in the wings but as it turns out we can neglect that. We can now try to fit the image using a two-dimensional least-square algorithm. In order for the fit to work, we have to estimate initial values for the fit parameters, especially the frequency components k_x and k_y, as precisely as possible. The latter can easily be determined from a two-dimensional numerical Fourier transform (FFT).

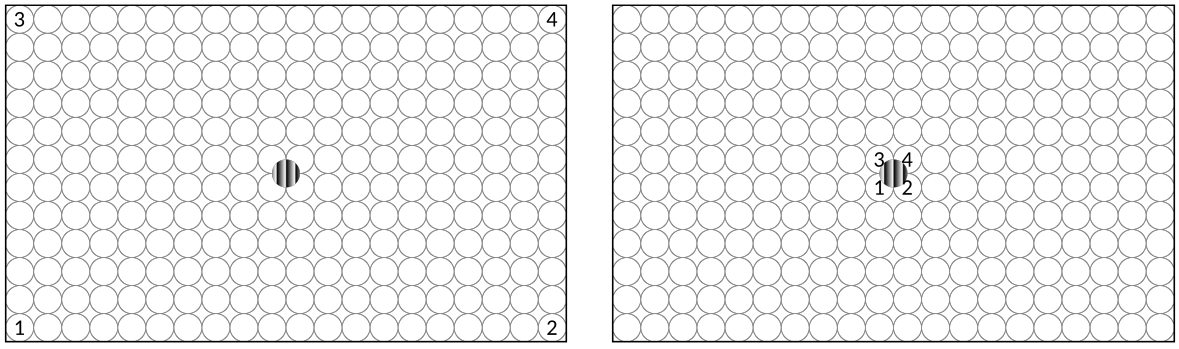

(a) **(b)**

Figure 4.13 – **(**a) Moving patch positions which maximize the k-vector components. **(b)** Center positions

To avoid scanning the whole Fourier space to extract the frequencies we estimate the frequency from

the images alone without any dependence on external parameters. The interference pattern will have the largest $|\vec{k}|$ if the distance between the patches on the sensor is maximal. We can therefore look at the four *corner* patch positions (see Fig. 4.13(a)) to estimate the maximum k-vector components in units of the FFT pixel spacing. From the position of the moving patch compared to the reference patch in the center, we can then estimate the magnitude and sign of the k-vector components to restrict our search to a small area in the output of our FFT. In principle, it would now be possible to also extract the phase at the detected position. However, we find that even for perfectly accurate data without any kind of noise[9] determining the phases from an FFT is unreliable. We therefore still have to fit the original image to extract the phase. Given the fact that we already have the frequencies, we can fit the interference pattern but with the frequencies as fixed parameters this time. It turns out that this indeed works well both in terms of runtime speed and reliability of the fit.

There are several optimizations we can apply to Eq. (4.17) to reduce correlations between the fit parameters to accelerate the computation. The simplest one is to combine ϕ_x and ϕ_y into a single phase ϕ. However, there is an important second simplification. The cosine term depends on three free parameters $\left(\mu_x, \mu_y, \phi\right)$ which are therefore strongly correlated. We can reduce the corrlations by making the substitution $\phi' := -2\pi\left(\mu_x k_x + \mu_y k_y\right) + \phi$. That way, we only have a single free parameter in the cosine as k_x and k_y are fixed and Eq. (4.17) simplifies to:

$$f(x,y) = A\exp\left(-\frac{\left(x-\mu_x\right)^2}{2\sigma_x^2} - \frac{\left(y-\mu_y\right)^2}{2\sigma_y^2}\right) \times \left[1 + C\cos\left(2\pi\left(k_x x + k_y y\right) + \phi'\right)\right]. \tag{4.18}$$

Of course we have to undo this replacement after the fit, otherwise our phase would not be referenced to the center of the interference pattern. We find, however, that it is not precise enough to reference every phase to the current interference pattern and that we have to determine the global reference point to which we reference all phases. We get an initial value from the four *center*-patches which overlap with the reference patch (see Figure 4.13(b)). These images show only one antinode and by averaging the positions of the maxima we get a good starting point $\vec{\mu}_{\text{est}}$ for the real reference position. We can now rereference to this position using the following transformation:

$$\phi_{\text{new}} = \phi + 2\pi\vec{k}\cdot\left(\vec{\mu}_{\text{real}} - \vec{\mu}_{\text{fit}}\right) \mod 2\pi \tag{4.19}$$

where $\vec{\mu}_{\text{real}}$ is the real center position and $\vec{\mu}_{\text{fit}} = \begin{pmatrix}\mu_x\\ \mu_y\end{pmatrix}$ is the center position extracted from the respective fit. We vary $\vec{\mu}_{\text{real}}$ around $\vec{\mu}_{\text{est}}$ while calculating the global standard deviation of the pattern. The real reference position will be the position that minimizes the global standard deviation of the wavefront. This procedure is justified by the following argument: If we assume a well aligned optical system, the aberrations (see Appendix A.4 and Figure A.7) are rotationally symmetric about the optical axis. By referencing all phases to a point radially displaced from the optical axis, we introduce an additional phase gradient along the displacement vector. This gradient is not radially symmetric and can therefore not cancel the "real" aberration in the system, thus enhancing it.

Finally, we repeat the measurement for different positions of the static "center" patch and determine the wavefront for two reference positions. To determine the actual wavefront, the raw measurement has to

[9] We generated the interference pattern artificially, by performing a numerical FFT on the SLM plane electric field distribution.

be phase-unwrapped, expanded to a single pixel level and then smoothed. We unwrap the phases using a *Python* package called `unwrap` and apply a Gaussian filter with a standard deviation of half the patch size. This correction can now be subtracted from the phasemask written to the SLM. The measured wavefront with and without correction is shown in Figure 4.14. Inside the aperture, the peak-to-peak wavefront deviation is reduced from λ to $\lambda/6$ and the standard deviation reduces from $\lambda/7$ to $\lambda/40$. Figure 4.15(a) shows the peak-to-peak deviation depending on the radius of the center disk. It can be seen, that an error of no more than $\lambda/10$ is maintained until 75 % of the maximum radius, where it rises to $\lambda/6$ at the end. Further reduction is possible by repeating the measurement and combining the resulting corrections. We find, however, that noise in the measurement data increasingly becomes an issue, especially along the edges of the center aperture. At this point, the measurement and evaluation take a significant amount of time, but peak-to-peak deviations of $\lambda/10$ can be reached. The correction is stable for weeks but only on the $\lambda/6$ level. For this work, we perform two iterations to reach the $\lambda/6$ level which, compared to the error the alignment of the objective to the vacuum window causes, is more than enough.

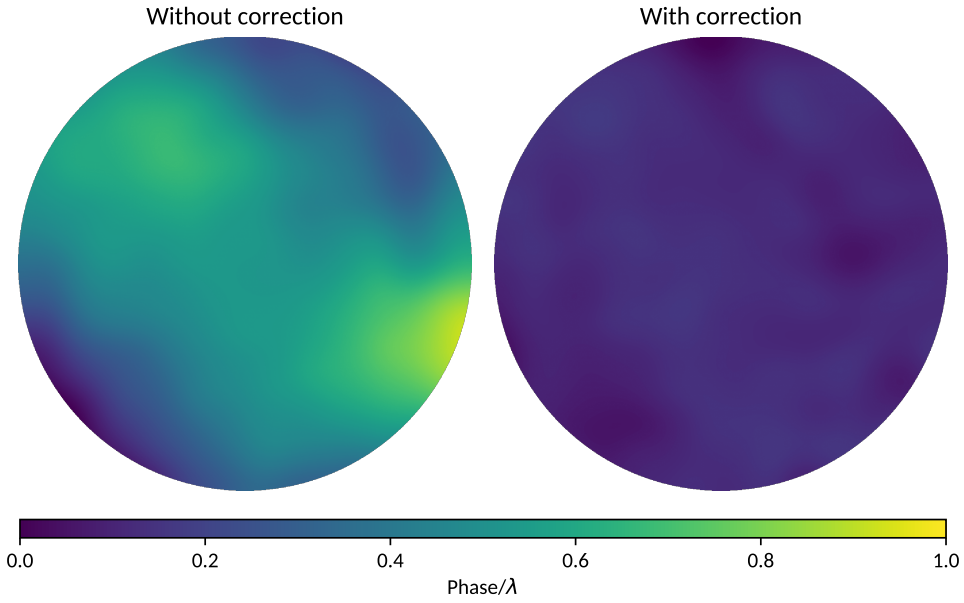

Figure 4.14 – **(Left)** Wavefront measurement without correction and with correction **(right)**. Inside the aperture the peak-to-peak deviation is reduced from λ to $\lambda/6$ and the standard deviation reduces from $\lambda/6$ to $\lambda/40$.

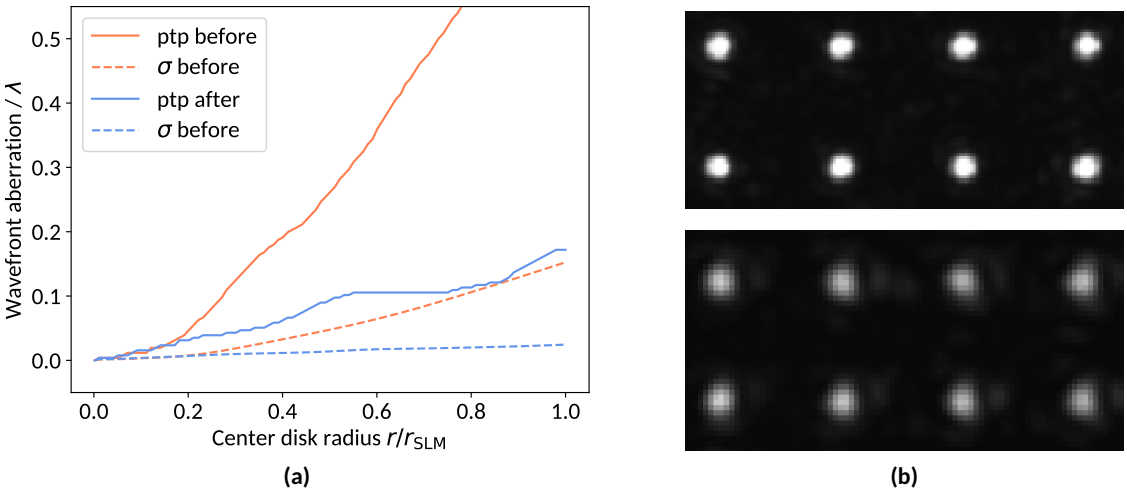

Figure 4.15 – **(a)** Wavefront deviation (peak-to-peak and standard deviation σ) depending on aperture size. **(b)** Point spread function with (**top**) and without (**bottom**) wavefront correction. Different magnifications have been used, leading to different resolutions.

4.3.2 Tweezer array generation

Displacing the beam in the image plane can be achieved by a blazed grating in the fourier plane. Hence, the superposition of a set of blazed gratings with different slopes and angles creates an array of foci in the image plane. The problem with this approach is, that it leads to a poor homogeneity of the intensity distribution in the array. The reason for this is, that the pattern is sensitive to small deviations of the phases and that the computation does not account for interference effects. A variety of different approaches have been developed and a good overview of the results is given in [104]. For this work, we will use an algorithm based on an iterative feedback, known as the Gerchberg-Saxton algorithm (GS-algorithm) [105]. Since its initial proposal in 1972, several modifications to increase convergence rate and quality of the final image have been developed. We will first discuss the basic algorithm and two of these extensions afterwards.

4.3.2.1 Standard Gerchberg-Saxton algorithm

Figure 4.16 shows the general layout of the phase-retrieval algorithm. We start in the upper left corner in the SLM plane. Here, the electric field modulus is given by the radial intensity distribution of the incident laser beam. The phase distribution is initially assumed to be random and we find that a flat phase distribution can lead to significantly longer convergence times. The electric field distribution is propagated to the image plane with a two-dimensional numerical Fourier transformation. We compare the amplitudes of the field at the desired trap positions, and if a certain difference threshold is crossed, the computation is stopped. Otherwise, we replace the amplitudes by the target amplitudes while keeping the phases and transform back to the SLM plane. Here, we again replace the amplitudes by the intensity distribution of the laser beam and again keep the phases. The phasemask is the free parameter of the optimization process.

The SLMs, that we use, have a center aperture diameter of around 1 000 px, which we further pad

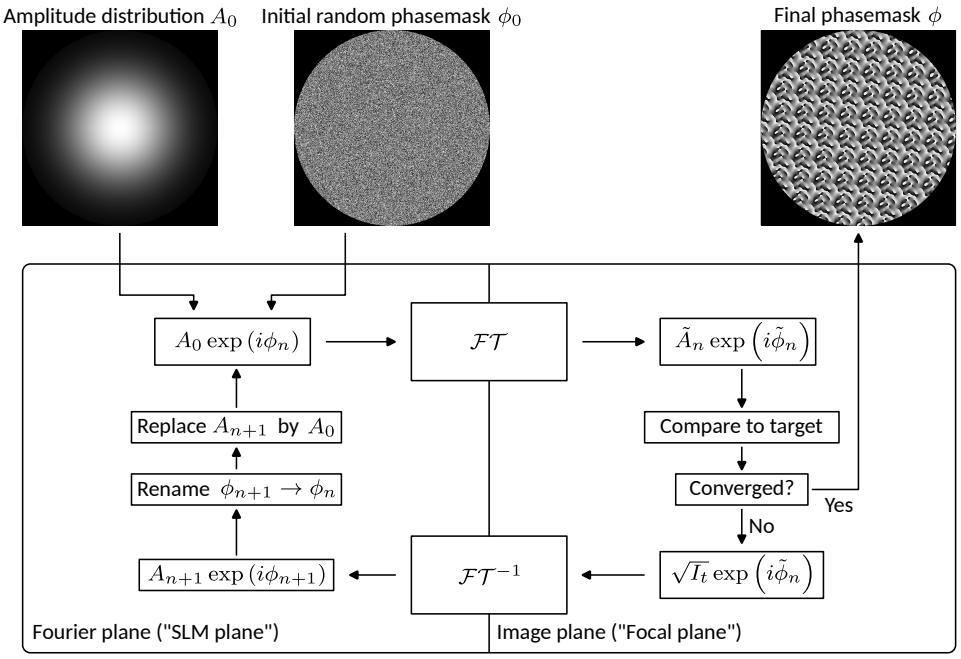

Figure 4.16 – Iterative Gerchberg-Saxton phase-retrieval algorithm [105]. The algorithm starts with a random phase distribution ϕ_0 and an amplitude distribution A_0 given by the beam. As long as the desired target intensity distribution is not sufficiently well approximated, it projects between the SLM plane and the focal plane connected by a Fourier transformation. In the focal plane the calculated amplitude distribution $\tilde{A}$ is discarded and replaced by the target intensity distribution $\sqrt{I_t}$. In the SLM plane the calculated amplitude distribution A is discarded and replaced by the electric field modulus of the laser beam A_0. The phase distribution ϕ is the parameter to be optimized.

to 8 192 px[10] to increase the resolution in the image plane ("zero padding", see Appendix A.2.2). With a complex `float` or `double` value for every pixel, the SLM- and image-plane both use around $16\,\mathrm{B} \cdot (8192)^2 \approx 1\,\mathrm{GB}$ of memory. To compute the Fourier transformations efficiently, we implement the algorithm in C++ and compute the Fourier transformations on a graphics processing unit (GPU) with the *CUDA toolkit*[11]. This decreases to computation time by a factor of 50 to 100 to around 100 ms per iteration (compared to a CPU implementation), depending in the GPU.

We quantify the homogeneity of the array with the *relative standard deviation* $c_v = \sigma/\langle I \rangle$, where σ is the standard deviation of the trap intensities and $\langle I \rangle$ is the mean peak intensity. For this work, we assume that all traps have the same size and therefore power and intensity are interchangeable. We will later see, that this is not true for small spacings where interference between adjacent sites becomes an issue. Although the standard GS-algorithm is able to produce c_vs below 40 %, this is still far from what we can achieve at best. We will therefore first discuss the modifications that have been developed to improve uniformity and convergence rate, before we present the results.

[10] Powers of two are computationally desirable for the fast fourier transformation algorithm.
[11] https://developer.nvidia.com/cuda-toolkit

4.3.2.2 Modified Gerchberg-Saxton algorithm

Adaptive target amplitude weighting The first modification is to let the algorithm adapt to the current result, that is, it actively updates the target intensities based on the current intensities. This is also known as the "adaptive Gerchberg-Saxton algorithm" and has been presented in [106, 107][12]. To increase the resulting homogeneity, we replace the amplitudes in the image plane by a scaled version of the target intensities, see Figure 4.17.

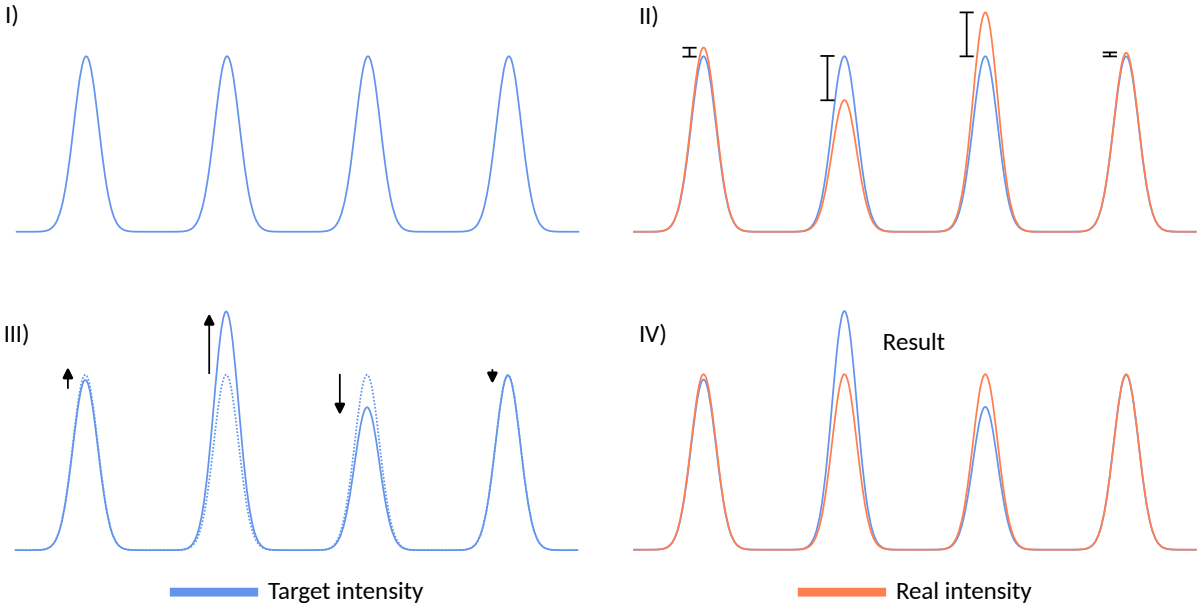

Figure 4.17 – Visualization of the weighting process. **I)** Initially, all target intensities are equal. **II)** After an iteration of the algorithm, we compare the current intensities to the target intensities. **III)** Compared to the standard GS-algorithm, the target intensities are updated according to the current intensities. The updated target intensities are then used during the update step in the focal plane in Figure 4.16. **IV)** Weighting the target intensities improves the resulting homogeneity.

We calculate the weighting factor for each trap i according to:

$$w_i = \frac{1}{1 - g_0 \left(1 - \left(\frac{I_i'}{I_{t,i}}\right)^{g_1}\right)} \tag{4.20}$$

where I_i' is the current intensity of the $i-$th tweezer, normalized to the mean of the current intensities. Similarly, $I_{t,i}$ is the target intensity of tweezer number i normalized to the mean of the initial target intensities. The gain factor g_0 determines how much of the current result is mixed into the new target intensity. In cases where all target intensities are equal, $I_i'/I_{t,i}$ simplifies to $I_i'/\langle I \rangle$. In contrast to [106], we include a second gain g_1 which allows us to change the strength of the feedback[13]. We find that g_1 becomes useful to increase convergence rates, if data from the trapped atoms is used for further intensity homogeneization (see Section 5.6).

[12] These are also the references that we will compare our results to.

[13] We also considered to use the field amplitudes instead of intensities but found it to perform worse.

Phase fixing The second modification, referred to as *phase-fixing*, increases the convergence rate at the cost of a lower modulation efficiency[14] and has been described in [107].

The weighting step only rescales the amplitude depending on the current amplitude or intensity, but the computation is also sensitive to changes of the phases. Changes in the latter are thus not accounted for but can have large influences on the uniformity and convergence rate, as they appear in the exponent of the electric field distribution. A possible solution to this problem is to keep the phases ϕ_i for each trap constant while rescaling the amplitudes. This leads to a better homogeneity which is also achieved with less iterations, but it reduces the modulation efficiency by a few percent. A common approach is therefore to fix the phases after a certain number of iterations, or when a certain threshold is crossed. This greatly improves the final homogeneity, whereas the reduced modulation efficiency can be neglected compared to other losses in the system. Figure 4.18(a) shows the convergence for a rectangular 30×50 array with a spacing of $4w_0$ with and without the phase fixing. The phase is fixed at a non-uniformity of 3 % which limits the modulation efficiency to about 91 %. It can be seen, that the phase-fixing achieves about three times better uniformity in one-third of the iterations. Additionally, it can be seen that the result is deterministic for the same initial random phase distribution. We also compare the convergence for different weighting factors where we fix the phase at a non-uniformity of 1.5 %. We achieve the best result with an intensity feedback with gains of $g_0 = 1, g_1 = 1$, compared $g_0 = 0.7, g_1 = 1$ in [106] and $g_0 = 1, g_1 = 1$ with amplitudes instead of intensities in [107]. Despite initial overshoots, we reach the same non-uniformity faster, as shown in Figure 4.18(b).

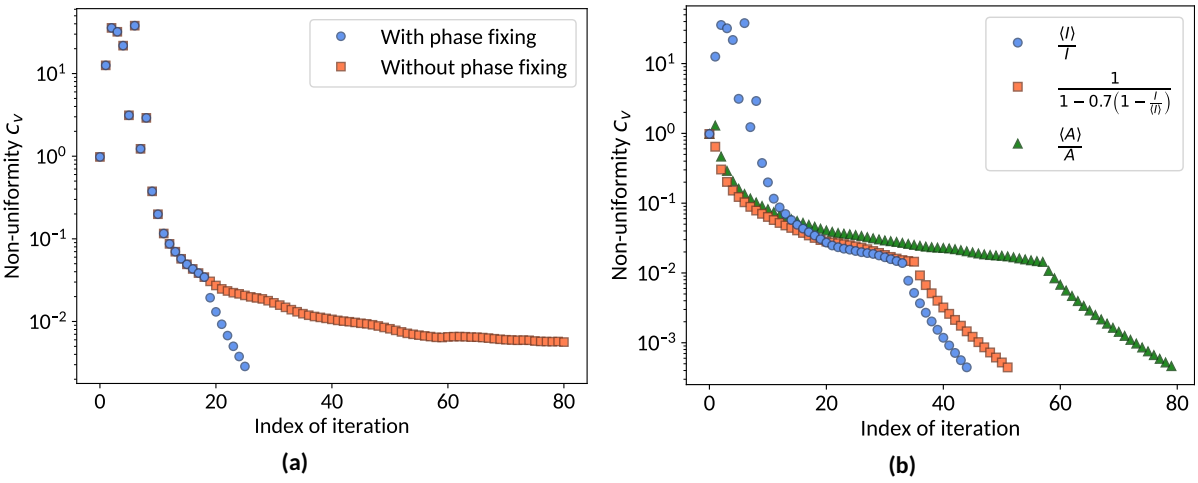

Figure 4.18 – **(a)** Convergence with and without phase-fixing for a 30×50 array with a spacing of $4w_0$. The phase is fixed when $c_v \leq 3$ %. **(b)** Convergence for the same array as **(a)** but for different weighting functions. Blue dots show our best result, red squares show a computation with parameters from [106] and green triangles show a computation with parameters from [107].

4.3.2.3 Camera feedback

Although theoretical non-uniformities of less than 0.5 % can be achieved, in reality the nonuniformity will be higher due to effects that are not accounted for in the computation, including aberrations in the optical path and pixel cross talk in the SLM. We therefore perform an additional feedback step where we

[14] The modulation efficiency is defined as the fraction of the diffracted light that is modulated into the traps $\epsilon = \sum_n I_n / \int I(\mathbf{x}) \mathrm{d}^2 x$.

use intensities measured by a camera in the optical setup in the algorithm, once a certain threshold is reached. We apply the phasemask to the SLM, together with an underlying blazed grating to separate it from the 0^{th} order. Afterwards, we capture the image with a camera placed in the image plane.

Locating the array To extract the individual intensities I_i, we first have to determine the correct orientation of the camera image before we have to locate the array in the camera image. This can be achieved by creating an array without mirror-symmetries for example a character or a word. We compare the camera image to the orientation of the theoretical intensity distribution and apply the correct transformation.

To locate the array in the camera image, we make use of a convolution algorithm explained below. The advantage of this method compared to a peak detection algorithm is that we get the correct assignment "for free", whereas the peak detection might return the positions in an arbitrary order. We know the spacing in the image plane in SI-units and from this we compute the spacing in units of the camera pixel pitch. Next, we define a binary kernel that is non-zero at the tweezer positions and zero everywhere else. To reduce the influence of noise, we apply a Gaussian-filter to the image which can be thought of as replacing every pixel by a local average of the intensities around it. We convolve the kernel with the filtered camera image and determine the position of the maximum value which is equal to the offset that we are looking for. This offset is the vector that has to be added to the kernel, to maximize the overlap. The position for each trap is then computed by the sum of the offset vector and the trap position inside the kernel. To account for small inaccuracies, for example in the spacing, we create small square *lookup-windows* around the approximate trap positions and perform a peak detection in these lookup-windows. We extract the intensities and in turn the amplitudes. The latter are combined with the phases from the theoretical iteration to be used in the feedback loop. Figure 4.19(a) shows the convergence of the camera feedback for different weights and array spacings. It can be seen, that the final non-uniformity is limited by the array spacing and not the feedback function. We also find no difference in convergence or resulting non-uniformity for arrays of sizes 10×10 to 30×30. Smaller arrays converge faster but with the same resulting non-uniformity of around $\approx 1.5\,\%$. The latter number has to be taken with a grain of salt, as it depends on the method that is used to extract the intensity from an image.

To compare our result with values from the literature, we evaluate the arrays from the embedded images in [106] and [107] where the authors state a c_v of 1.4 % and 1.4 % respectively. With our method, we determine a c_v of 1.9 % for the 10×10 array and 2.5 % for the 30×50 array. For comparison, we compute a c_v of 1.1 % (3.0 %) for a 10×10 (30×50) array generated with our algorithm. The distribution of trap depths with and without the camera feedback for the 10×10 array is shown in Figure 4.19(b). It can be seen that the camera feedback significantly reduces the spread of the trap intensities and the result is comparable to [106].

We conclude that we can achieve the same non-uniformities for different array sizes as observed in the literature. Care has to be taken to evaluate images with the same algorithm to get comparable results. Examples of tweezer arrays used in this work are shown in Figure 5.22.

Spacing limitations We find that for spacings below $4w_0$, the interference between adjacent traps dominates the intensity distribution and traps start to coalesce, see Figure 4.20. This can be problematic for e.g. Rydberg- or Hubbard-experiments where small spacings for the blockage or tunneling are required. These smaller spacings can still be created with an acousto-optic device where each trap has a slightly different frequency and interference is therefore suppressed. This type of tweezer array

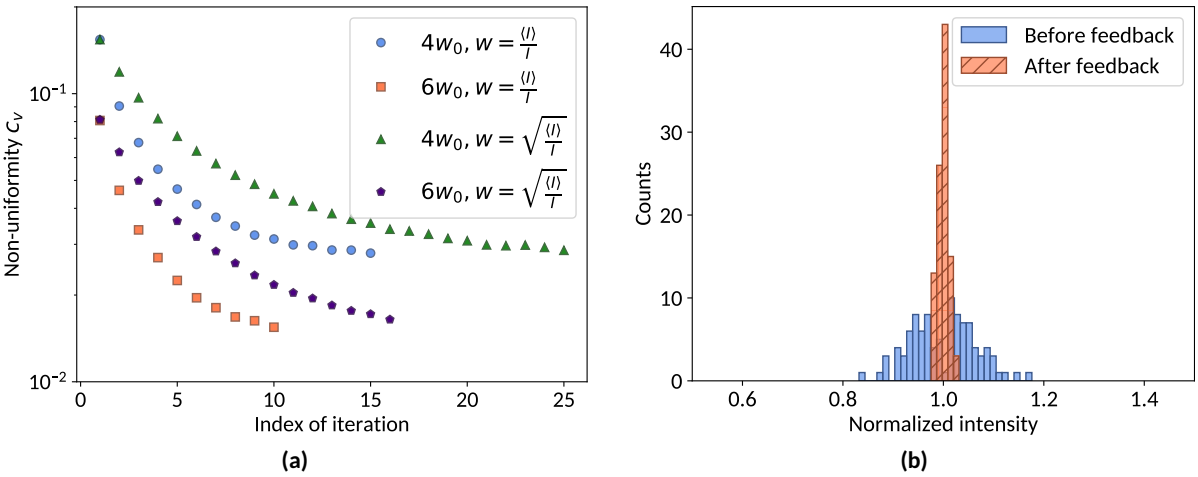

Figure 4.19 – **(a)** Camera feedback convergence for a 30×50 array for different weighting functions and array spacings. **(b)** Distribution of trap intensities with and without camera feedback for a 10×10 array. The nonuniformity c_v (relative peak-to-peak deviation) was reduced from 6.5 % (34.4 %) to 1.1 % (5.5 %).

generation is covered in e.g. [108].

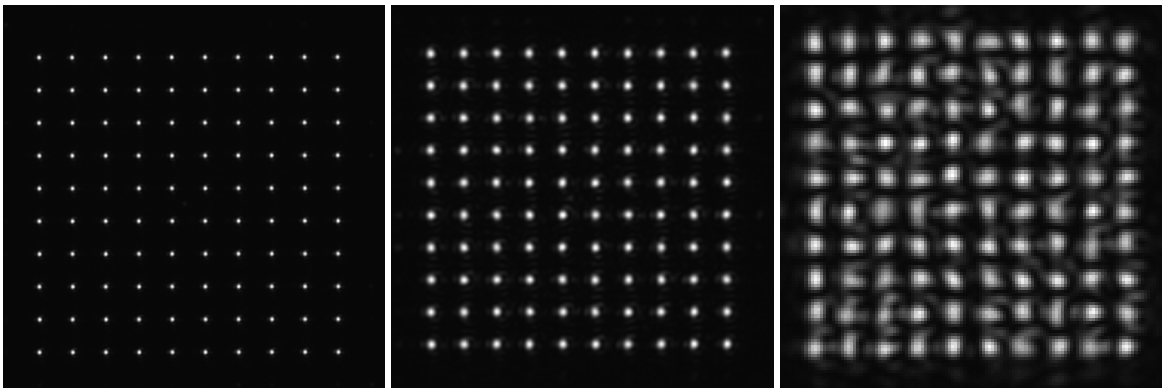

Figure 4.20 – Influence of interference between adjacent spots for different lattice spacings. Spacings are **(a)** $13.5w_0$, **(b)** $6.5w_0$ and **(c)** $3.5w_0$. For small separations the traps start to merge due to interference and continuous structures begin to form.

4.3.3 Combined phasemask and diffraction efficiency

The phasemask that is finally written to the SLM is a superposition of several constituents. We combine the wavefront correction that we calculated in Section 4.3.1.2, a blazed grating to separate different diffraction orders and the tweezer phasemask. The blazed grating has a period of 7 px leading to a separation of $\lesssim 1$ mm in the first focal plane (see Appendix A.1.4). Here, a tradeoff between the size of the image plane and in turn the maximum array size and the diffraction efficiency has to be made. Details for the grating efficiency can be found in Appendix A.3. A blazed grating with seven phase steps has a maximum diffraction efficiency of $\epsilon = 94\,\%$ which is squared as it is applied along both axes, i.e.

$\epsilon^2 = 87\,\%$ (see Appendix A.3 and Figure A.5). Together with the zeroth-order diffraction efficiency (given by the reflectivity of the backplane) of around 97 % and a modulation efficiency of around 90 % this leads to a fraction of the power of about 75 % that is modulated into the traps. With additional losses through optics of 80 %, the whole setup has an efficiency of 60 % (The clipping loss at the aperture of the SLM is around 1 % which is negligible, see Section 4.2).

To fine tune the radial- and axial position of the arrays, we use additional linear phase gradients and Fresnel lenses, see Appendix A. More generally, a set of orthogonal polynomials, called *Zernike* polynomials, can be used to add and compensate certain optical aberrations. An overview of these is given in Appendix A.4. The superposition of the different contributions and the final result are shown in Figure 4.21.

Figure 4.21 – Combined phasemask. The grating is used to separate different diffraction orders. In addition, a small phase gradient is used to fine tune radial position of the array. Similarly, the Fresnel lens is used to fine tune the axial position of the array.

4.4 Acousto-optic spatial light modulation

The second dynamic device that we will use for this work to create spatially varying intensity patterns is a crossed pair of acousto-optic deflectors (AODs)[15]. Acousto-optic deflectors or modulators are a common tool in modern cold atom experiments where they are often used as frequency shifters and intensity stabilizers or switches. In this work, we will create dynamic optical tweezers with a pair of crossed AODs. We will see that by changing the frequency of the radio-frequency drive we can displace the trap spatially in the atom plane, which allows us to pick up and to rearrange atoms at will.

4.4.1 Working principle

AODs are based on Bragg scattering of photons off acoustic phonons inside a crystal [82], which leads to momentum transfer causing a deflection of the photon. The optimal deflection angle Θ_B can be calculated from the Bragg condition:

$$\Theta_B = \frac{\lambda}{2\Lambda}, \tag{4.21}$$

where the wavelength of the acoustic wave Λ is given by

$$\Lambda = v_a / f_{RF}. \tag{4.22}$$

[15] DTSXY-400-515 by AA OPTO-ELECTRONIC

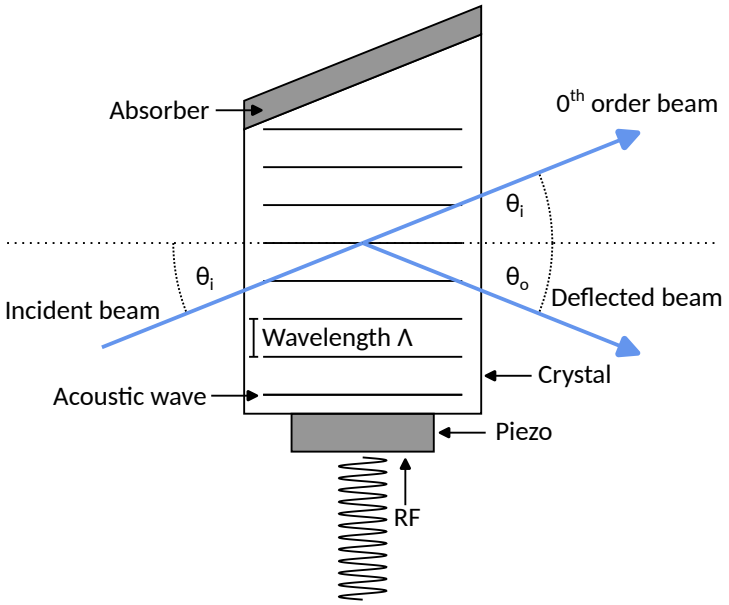

Figure 4.22 – AOD working principle. A radio-frequency (RF) wave is translated to an acoustic wave using a piezo-crystal. Photons of an incident laser beam can scatter on the acoustic wave with wavelength Λ, similar to Bragg-scattering inside a crystal. The deflection-angle and the intensity of the deflected part can be controlled by the frequency and the power of the RF-wave. For optimal efficiency, Θ_i should be tuned to the Bragg angle $\Theta_B = \lambda/2\Lambda$ for an incident beam with wavelength λ.

With the speed of sound v_a in the crystal we arrive at:

$$\Theta_B = \frac{\lambda}{2v_a} \cdot f_{RF}. \tag{4.23}$$

A schematic of the process is shown in Figure 4.22. We already saw in Section 4.2.1, that a change in angle in the front focal plane of a lens is translated to a change in position in the backfocal plane (see also Appendix A.1.4). It is therefore possible to tune the *output angle* Θ_o of a beam by changing the *frequency* of the RF-signal. The amount of light that is modulated into the deflected order depends on the RF-amplitude but also on the frequency itself and the bandwidth of the device[16]. Compared to the SLM, which operates on timescales of 100 ms, an RF-frequency can be changed on a microsecond timescale and we can therefore use the AOD to create a dynamically moveable trap. The functional form of the change in frequency is of great importance if an atom inside a trap is to be moved, and we will discuss it in detail in Section 6.1.

We can calculate the displacement by combining Equation (4.23) and Equation (A.22)[17]. The angle is changed by the magnification M of the optical system which is $M = 750/150 \cdot 200/100 = 10$ in our case. With $\lambda = 515$ nm and $f = 13.48$ mm we get a displacement of around 1.25 µm/MHz.

We will now briefly describe the principle behind the RF-wave generation using an arbitrary waveform generator. Afterwards, we discuss a method of intensity stabilization, which is needed to compensate for changes in the deflected beam power, if the RF-frequency is changed.

[16] A common value is a bandwidth of 10 MHz for RF-frequencies around 100 MHz.

[17] For $\lambda = 515$ nm, $v_a = 650$ m/s for TeO_2 and $f_{RF} = 82$ MHz we compute $\Theta_B = 1.86°$.

4.4.2 Arbitrary waveform generation

The pair of crossed AODs is driven by an arbitrary waveform generator[18] (AWG). The generation of different waveforms and the implementation of the atom rearrangement scheme in Chapter 6 is a result of the intrinsics of the AWG card. To understand the calculation of the different waveforms used throughout this work, we will first therefore discuss some of the relevant parts of the AWG card and the software implementation.

The output converter stage of the AWG has a resolution of 16-bit with a maximum voltage output of ±2.0 V into 50 Ω. We further amplify the signal with an amplifier[19] to achieve a maximum power of 1 W that we apply to the AODs. For 515 to 532 nm the AODs have their maximum diffraction efficiency at around 82 MHz. The maximum sampling rate of $f_S = 1.25\,\text{GHz}$ therefore allows to sample one sine period with roughly 15 points. For this work, we will only consider the case of driving both outputs, in which case the data is written in an interleaved format like shown in Table 4.2. Every change in amplitude and frequency has to be computed explicitly for every timestep $1/f_S \approx 1\,\text{ns}$ which requires a substantial amount of samples to compute even for a few milliseconds. We again implement the computation in C++ making use of CPU-multithreading without GPU-acceleration this time. We find the latter to not be appropriate for the amounts of data and the more-or-less *real-time* nature of the streaming process in Chapter 6. From March 2024 on, a new DDS-mode is available for the AWG-card which simplifies the implementation greatly as it only requires the computation of values for the amplitude and frequency without sampling the full waveform. It still has to be seen if this mode can achieve the same stability in amplitude and frequency as the explicit computation.

sample 0		sample 1		···	sample N	
16-bit data channel 0	16-bit data channel 1	16-bit data channel 0	16-bit data channel 1	··· ···	16-bit data channel 0	16-bit data channel 1

Table 4.2 – AWG data layout. For every timestep, both channels are sampled and written in an interleaved data format to the card.

Sequency replay mode

The AWG has several different operating modes that can be used to stream data to the outputs. The mode we will use is called *sequence replay mode* where the available memory is divided into N different segments, each with length L_i, where N is an integer power of two, that we are free to choose. The length L_i of a segment can also be chosen freely but the hardware dictates a minimum size of 384 and divisibility by 32 (with the maximum size given by the total memory divided by N.). The sequence replay mode then leaves us with the freedom to connect these segments however we want as Fig. 4.23 shows.

Phase matching

To use different segmented signals as building-blocks for more complex sequences we would to like to arrange them arbitrarily. This requires the phases at the transition from one segment to the next one to be

[18] M4i.6631-x8 by Spectrum Instrumentation

[19] *AA OptoElectronics* AMPB-B-34-10.500, 34 dBm, 10 to 500 MHz

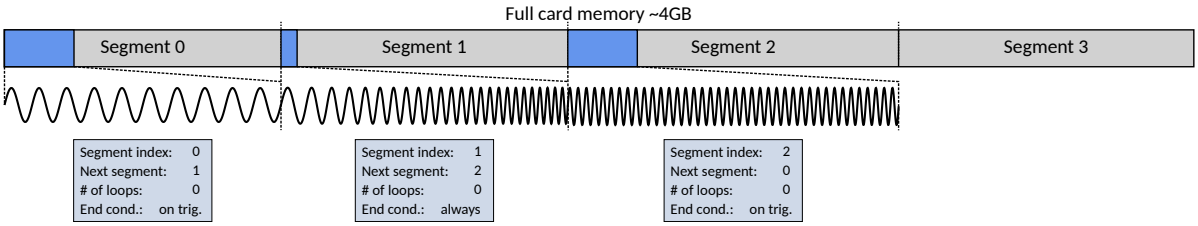

Figure 4.23 – Example of the sequence replay mode. The total number of segments has to be a power of two but not all segments have to be used. The full memory is divided into the number of segments, but the memory of each segment can also be used partially and each segment that is currently not active can be updated while the card is running. Segments can loop a specific number of times or indefinitely if the next segment number is set to the current segment number. In this example segment will loop until a trigger is detected which will change the frequency in segment 1 before entering an infinite loop in segment 2.

matched (difference $\lesssim \pi/12$) otherwise the RF-intensity will drop inside the modulator. Depending on the frequency and segment lengths, the phase difference can be anything between 0 and 2π, and we have to explicitly incorporate this into the computation. One possible solution is to find a segment size L_S that, together with a certain sampling rate f_S, allows to sample a certain set of frequencies with a phase mismatch of 0. Specifically we would like to sample frequencies $f = f_b + n\Delta f$ for a base frequency f_b and a "resolution" of Δf. We determine how many full wave sine samples with frequency f_b we need to fill a memory segment with a size of an integer multiple of the minimum memory size M_{min} by:

$$n = \left\lceil \frac{M_{\text{min}} \times f_b}{f_{S,\text{max}}} \right\rceil . \tag{4.24}$$

For example, for a base frequency of $f_b = 80\,\text{MHz}$ and the maximum sampling rate $1\,250\,\text{MHz}$ we compute $n = \lceil 24.6 \rceil$. To get an integer value we scale the sampling rate which in this case is $f_S = 1\,228.8\,\text{MHz}$. The segment length now depends on the desired frequency spacing Δf. We compute the number of samples per channel (twice the segment size) as $N = 4\frac{f_b}{n\Delta f} M_{\text{min}} = 49152$ for a frequency resolution of $100\,\text{kHz}$. This setup works well for single frequencies or superpositions but not if a frequency is changed continuously in time. The latter becomes important for the rearrangement of atoms and we will discuss a solution for that case in Section 6.1.2.

4.4.3 Deflection efficiency and passive intensity stabilization

When using the AODs to create foci at different positions in the image plane, it is important to ensure a constant trap intensity. The bandwidth of the AODs, together with the output stage of the AWG card, yields a frequency dependent maximum amplitude that we have to compensate for, as seen in Figure 4.24(a). An active intensity stabilization cannot be used in this case, as we would like to perform amplitude ramps in combination with changing the frequency, so we decided for a passive stabilization.

The first approach we tried was to fit the deflected intensity $I(f_x, f_y)$ with a two dimensional function. Converting the measured intensity back to a correction however did not work reliably, probably due to the nonlinearity of the output stage and the amplifiers. We therefore decided to determine the correction with a direct measurement, where we determine the output amplitude for each frequency that results in a certain reference intensity.

To stabilize the amplitude in a frequency interval $\left[f_0, f_1\right]$, we measure the deflected intensity for an

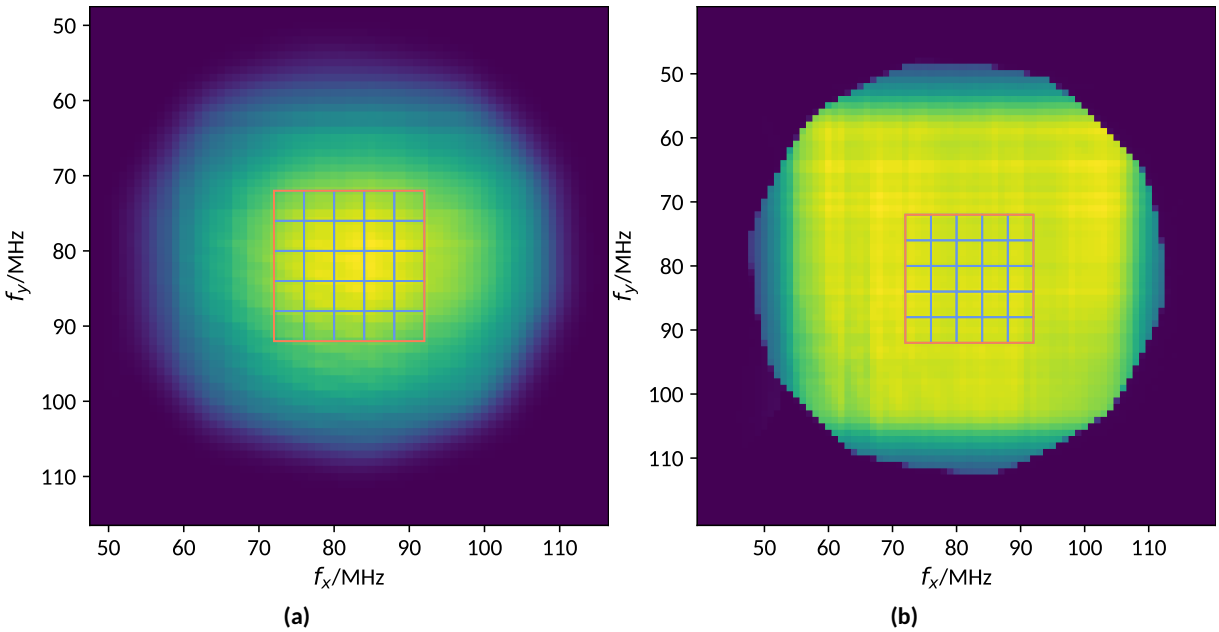

Figure 4.24 – Normalized deflected intensity without **(a)** and with **(b)** passive intensity stabilization. Inside the red square of 20 MHz × 20 MHz the intensity shows a peak-to-peak deviation of 29 % and 10 % respectively. In 4 MHz × 4 MHz (blue) subdivisions the maximum peak-to-peak deviation is 18 % (5 %) without (with) stabilization. The horizontal and vertical patterns are a result of the measurement process, where the frequencies for each channel are only scanned once for the center frequency of the other.

amplitude ramp from $[0, A_{\max}]$ with $N_{\text{steps}} = 20$ steps for frequencies spaced by 1 MHz. We do not raster the whole frequency space but only vary the frequency of one axis at a time while keeping the other at the center frequency. We then determine the frequency $f_{\min}$ that has the lowest $A_{\max}$, which is the amplitude that all other amplitudes have to be scaled to. For every frequency f inside the interval, we interpolate the amplitude ramps and determine the amplitude A for which $A(f) = A_{\max}(f_{\min})$. We create a lookup-table for all frequencies for both channels by storing the determined amplitude scaling factors in an array, where the frequency in megahertz is used as the index (we pad this array towards 0 MHz). During the waveform computation, we then only have to look up the appropriate amplitude scaling factor in the lookup-table without the need for a, possible expensive, computation.

Figure 4.24 shows the resulting intensity on a photodiode in dependence of the RF-frequency with and without the passive stabilization enabled. For this work, we are interested in the stabilization in a square region of 20 MHz × 20 MHz, as this is the size of the static tweezer array in which we would like be able to address every site. Inside the square centered around $f_x = f_y = 82$ MHz, we find a peak-to-peak deviation of around 10 % with a standard deviation of 1 %. However, for the rearrangement process in Chapter 6, the frequency never changes by 20 MHz in a single frequency chirp and in more than 80 % of the cases we stay inside a 4 MHz × 4 MHz-square. Tiling the center square in 4 MHz × 4 MHz-squares, we find a maximum peak-to-peak variation of 5 % at a mean of 3 %. This is a good improvement to the values of a maximum peak-to-peak variation of 29 % at a mean of 10 % without stabilization. In the future, the deviation can be further reduced by performing a full two-dimensional measurement to determine a scale factor for every (f_y, f_x)-combination. We will measure the intensity stability and estimate its influence on the movement of atoms in Section 6.3.3.

4.5 Tweezer beam path

We conclude the description of the SLM and AODs with an overview of the optical setup around the main chamber. Both setups receive the light through an optical fiber from the laser table. Behind the fiber, a half-waveplate is be used to fine tune the polarization before the light passes through a polarizing beamsplitter (PBS). The PBS converts polarization drifts into intensity drifts which, together with other intensity drifts, can be reduced by an active stabilization.

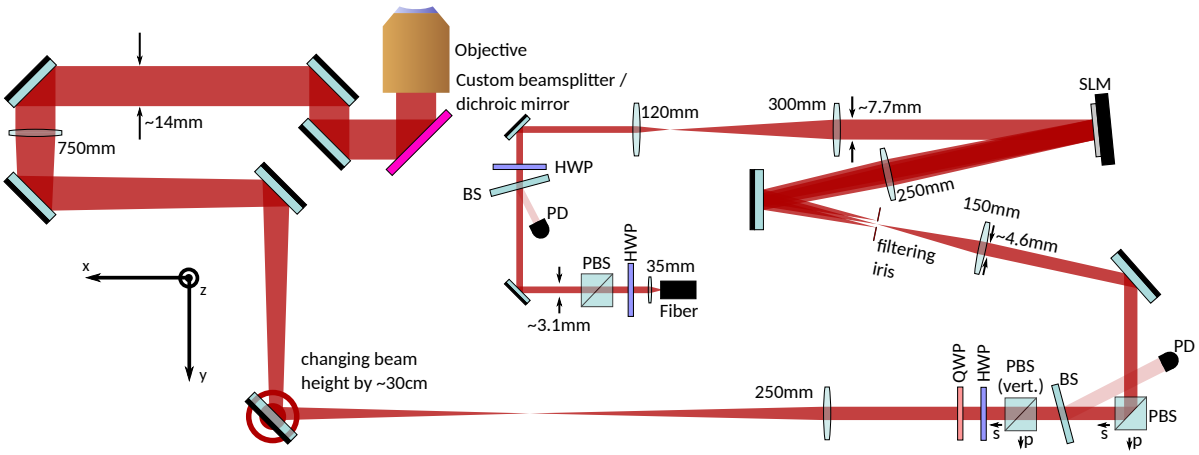

Figure 4.25 – Optical setup for the SLM

SLM setup - Figure 4.25 The beam exiting the fiber is collimated to by a 1 " achromatic lens with a focal length of 35 mm to a waist of 3.1 mm. The beam is expanded to a waist of 3.9 mm on the SLM, which has a aperture radius of 6 mm. The setup is build such that the sensor is imaged to the back-focal plane of the objective through a series of telescopes. This ensures that angles are mapped to angles between the two planes while minimizing changes in position. We choose the magnification such that the SLM aperture matches the clear aperture of the back-focal plane of the objective, thus $\zeta \approx 0.65$. Different diffraction orders are filtered with a pinhole in the first image-plane after the SLM. The beam is collimated again to a waist of 4.6 mm to fit through 10 mm PBSs and the center of 20 mm waveplates to reduce aberrations. From the collimated beam, we also pick-off a small percentage of light for the active stabilization which can therefore also compensate the flickering noise of the SLM. After the pick-off, a half- and quarter-waveplate can be used to tune the polarization at the position of the atoms, which we measure with a polarization analyzer shortly before the objective. For the wavefront flattening and the camera feedback for the tweezer array generation, we place an additional mirror in front of the objective and guide the light through a 500 mm lens to a camera.

AOD setup - Figure 4.26 The AOD setup follows the same design principles as the SLM setup, however, the power has to be stabilized before the AODs to be able to drive amplitude ramps correctly. Notice that the relay to the back-focal plane is less precise, as the deflection for both axis happen at two different locations. We did not see an influence of this mismatch yet but in the future this can be optimized by placing the two AODs in a $4f$-configuration. Due to spatial constraints and the requirement to relay the center positions of the AODs to the back focal plane of the objective, the magnification ratios

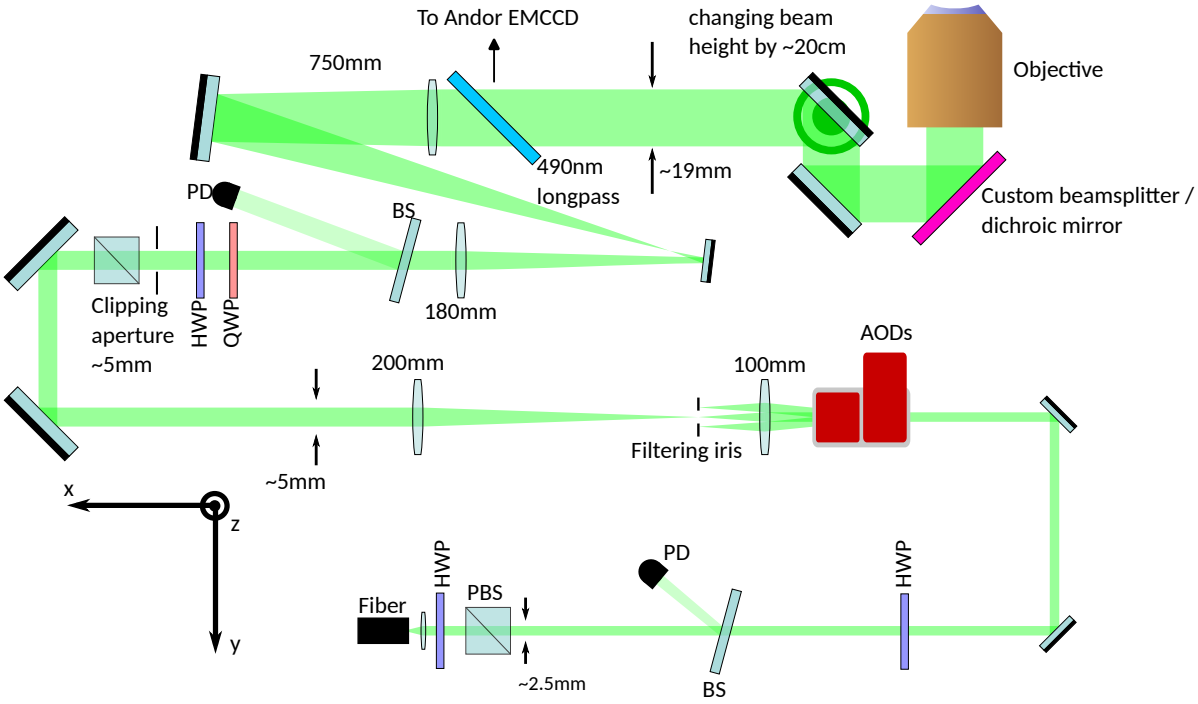

Figure 4.26 – Optical setup for the AODs

could not be chosen freely and in result the truncation ratio of $\zeta = 1.5$ deviates from the optimal value around $\zeta \approx 0.65$.

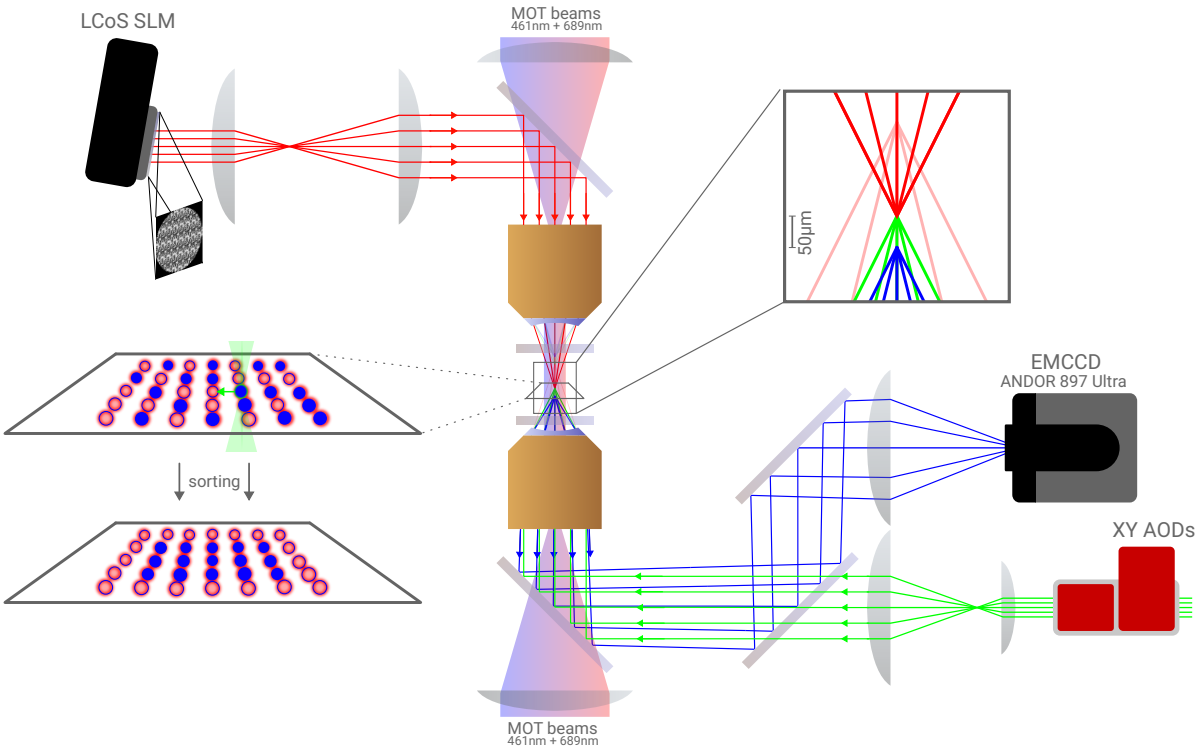

Figure 4.27 – Optical setup for the atom rearrangement

4.6 Summary and conclusion

In this chapter, we discussed the details for the creation of optical potentials for single atom trapping. We presented and characterized two devices to create static and dynamic optical tweezers. First, we discussed a liquid-crystal based SLM that we use to create static optical tweezer arrays at a wavelength of 813 nm. Using an interferometric wavefront measurement, we reduce aberrations in the optical system by up to a factor of 10. We discussed an algorithm for the generation of tweezer arrays and compared the results to the literature. The comparison showed, that we can create tweezer arrays of the same homogeneity of more than 98 % for up to 1 500 tweezers, similar to the literature. We finished the first part with a summary of the experimental parameters and efficiencies.

In the second part, we presented the tweezer generation using acousto-optic deflectors (AODs). We presented the underlying implementation of the arbitrary waveform generation that provides the radio-frequency signal. To realize single site addressing in a tweezer array, we developed a method of passive intensity stabilization, to ensure a constant trap depths in cases where an active stabilization cannot be used.

CHAPTER 5

Single atoms in optical tweezers

The previous chapter discussed the creation of optical tweezer potentials, and we will now populate these potentials with atoms. Building a quantum simulation platform requires the ability to prepare the system in a well-defined initial state, to maximize reproducibility. All deviations from a single initial state can be seen as a form of increased entropy, for example uncertainties associated with different atom numbers or the motional state of an atom inside a trap. The core objective of this chapter is to *decrease* this entropy, for which we need to learn to prepare a single atom in or close to the motional ground state.

The following description is loosely divided into two parts: We begin by investigating the loading process of the tweezers from the red narrow-line MOT and then move on to the preparation, cooling and detection of single cold atoms. In order to do this, we first discuss how light-assisted collisions (LACs) can be used to reduce the atom number to a single atom per trap at most. After preparing a single atom, we then cool it towards the motional ground state of the confining potential. Finally, we discuss how the detection of these single, ultracold atoms works. In general, we make no distinction between atoms in a tweezer array or in a single trap, except where mentioned explicitly.

As motivated in Section 4.1, we use tweezer wavelengths of 515 nm, 532 nm and 813 nm for this work. The focus for the goal of this work lies on tweezers at 813 nm, as this is the wavelength in which atoms will be rearranged and detected in Chapter 6. Tweezers at 515 nm or 532 nm are used to move atoms between sites, and during the transfer-phase no cooling light is applied. Even though it is not directly relevant for the main goal of this work, it is still instructive to also investigate the light-assisted-collisions and sideband cooling in traps at 515 nm and separately sisyphus cooling at 532 nm. At 515 nm, it is possible to realize a condition where the ground state and the excited state experience the same trapping potential, which can be exploited to address the quantized motional states of a trapped atom for cooling or thermometry. For 532 nm, the cooling measurements in this work present the first time where strontium is investigated in high NA tweezers at 532 nm. 532 nm is an attractive choice of wavelength due to commonly available optical elements and lasers, and progress at this wavelength for (strontium) tweezer experiments could be of great benefit for the community in the future.

The second part aims to further characterize and and benchmark the system. We present a fast and robust method to determine the trap depths in a tweezer array which we use in a feedback step to increase the homogeneity of the array. We then investigate the trap quality where we use a parametric excitation method to measure and infer trap parameters. Finally, we determine the vacuum lifetime of an atom in a tweezer, as this quantity inherently limits the deterministic preparation of an atom array and hence the final entropy.

5.1 Overview of the experimental sequence

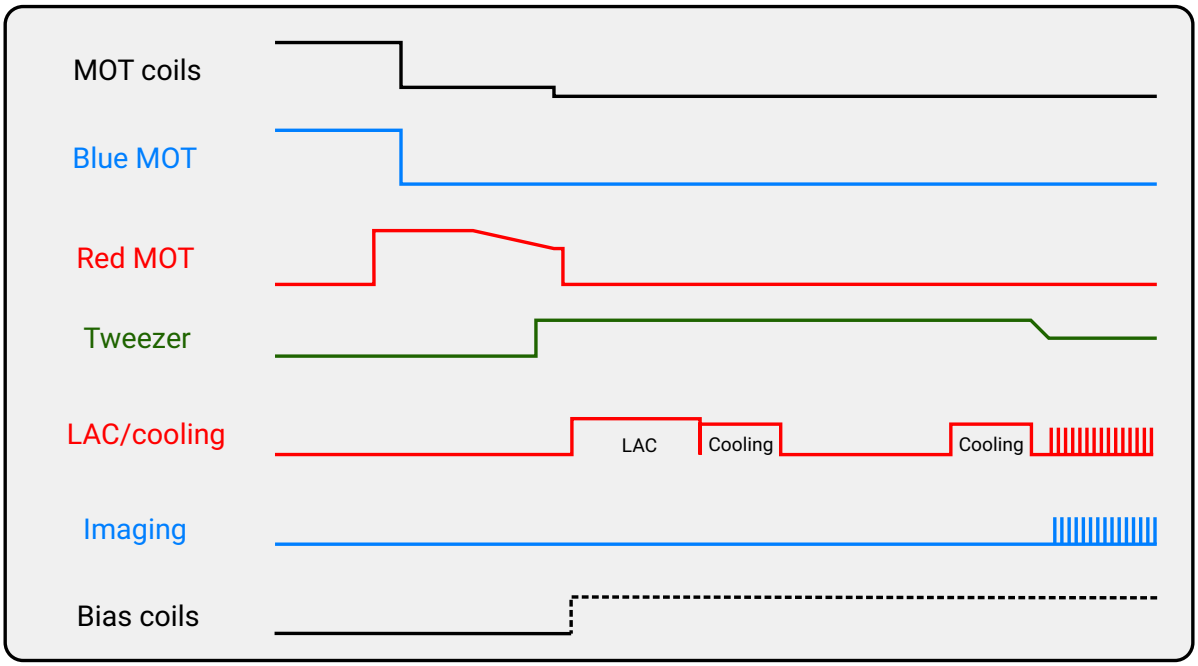

Figure 5.1 – Full schematic of the tweezer sequence. A repetition of the experiment takes around 500 ms.

A visual representation of the full tweezer sequence is given in Figure 5.1. We load the blue MOT for around 100 ms before moving to a 50 ms long SWAP-phase with an overlap of 5 ms. The single frequency red MOT phase has a duration of 50 ms and we enable the tweezers 10 ms before the end. After the red MOT follow several different blocks of operations performed on atoms in the tweezers and each block is described in detail in this chapter. First, a 30 to 50 ms long block of light-assisted collisions (Section 5.3) is used to reduce the atom number to a single atom at most. We then cool the single atom for around 30 to 50 ms where the details of the cooling process depend on the tweezer wavelength (Section 5.4). On the single cold atom(s) we can now perform measurements (Sections 5.6 to 5.8, also Chapter 6) which can take up to several seconds when e.g measuring the trap lifetime. Before we enter the detection phase, we apply a second block of cooling for around 30 to 50 ms, such that atoms always enter the imaging step with the same temperature, irrespective of the measurements performed earlier in the sequence. The experimental sequence finishes with a 100 ms long imaging block, where we collect blue fluorescence light to determine if an atom is present or not (Section 5.5). In total the sequence takes around 500 ms, depending on the performed measurements.

5.2 Atom transfer from MOT to tweezers

The first step for preparing a single atom is to populate the tweezer potential by a transfer from the red MOT operating on the $^1S_1 \rightarrow {}^3P_1$-transition. The narrow-line MOT produces a cold cloud of 1×10^5 atoms with a center density of roughly $7 \times 10^{12}\,\mathrm{cm}^{-3}$. By shining in the tweezer light, several

atoms[1] get trapped in the millikelvin-deep potential which is a visible signature for the EMCCD camera. Two components are required to detect fluorescence signal from the atoms trapped in a tweezer: First, the EMCCD camera and the accompanying beam path has to be aligned such that the atomic fluorescence is imaged onto the camera chip. Secondly, the red MOT and the tweezer potential have to be overlapped by displacing the MOT using magnetic offset fields. The displacements are on the order of the red MOT and therefore do not influence its performance.

To align the imaging system, we align a laser beam to the optical axis of the lower objective, as this is also the axis along which the fluorescence light exits the objective. We then retroreflect this beam by placing a mirror just below the objective and guide the reflection to the camera using a non-polarizing dichroic beamsplitter. The detailed procedure is described in Appendix B.2. The second step is to optimize the spatial overlap of the tweezers and the MOT. In order to do this, we shine in light at 461 nm with a power of 20 µW during the single frequency MOT phase, in addition to the tweezer light. The red MOT has a size around 150 µm which is larger than the field of view (FOV) of the camera of around 100 µm. We displace the red MOT by applying magnetic offset fields on the order of 1 G, until we see a comet-shaped signal, depicted in Figure 5.2(a). The comet is caused by atoms that are first trapped in the tweezer and which are then accelerated out of the trap by the imaging beam. The direction of the tail is thus opposite to the direction of the exciting beam. By increasing the time point of the image step-by-step beyond the end of the MOT phase, the fluorescence signal from the MOT faints and the signal from atoms inside the tweezer becomes more prominent. This allows us to fine tune the magnetic fields and the axial position of the camera such that signals from single atoms can be optimized (see Figure 5.2(b)).

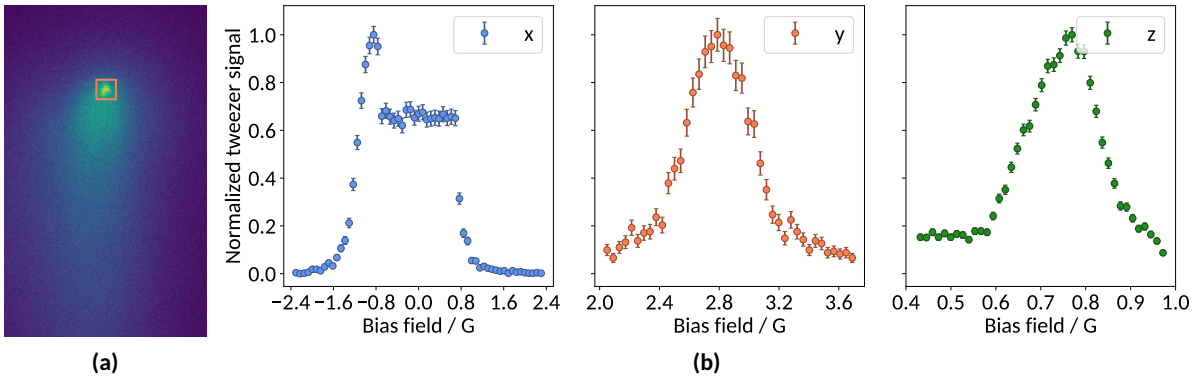

Figure 5.2 – **(a)** Fluorescence signal from a tweezer (red square) in the red MOT. As atoms leave the trap, they still scatter light which shows up as the comet-shaped vertical signal, opposite to the direction of the imaging beam. **(b)** Magnetic offset field scan used to overlap the tweezers and the red MOT. The displacement is on the order of the size of the red MOT of ~ 150 µm.

5.3 Single atom preparation using light-assisted collisions

We load the tweezers by overlapping the tweezer light with the red MOT for 10 to 20 ms before disabling the latter. This results in a distribution of atom numbers inside the trap with a mean of more than one. To reduce the uncertainty associated with the distribution of atom numbers N_{atom}, we employ a phase of light-assisted collisions (LACs) leading to photoassociation. Two atoms in the tightly confining potential can be excited to a bound molecular state [109, 110] which can leave the trap. Repeating the process

[1] probably ten(s) of atoms given the 1 µm^3 volume of the tweezer

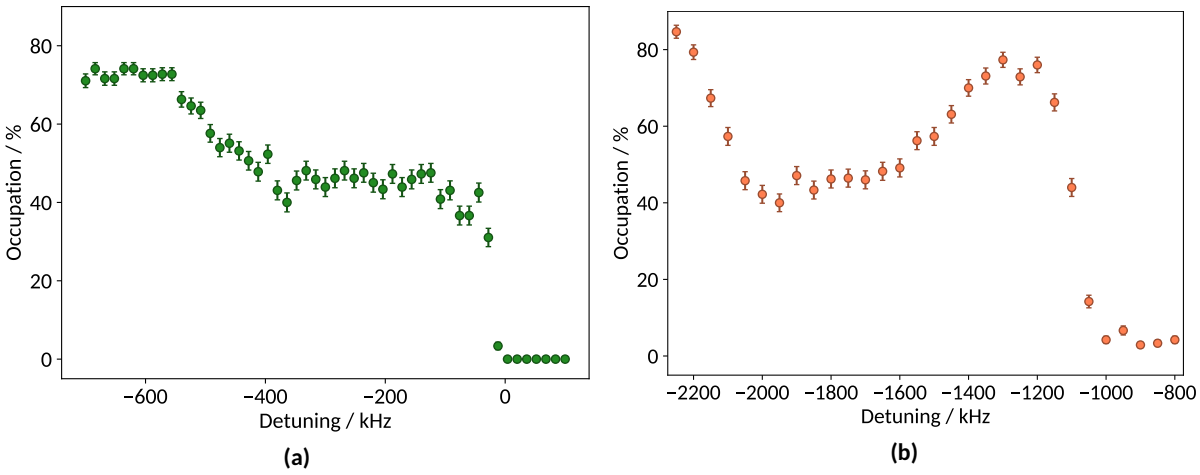

Figure 5.3 – Light-assisted collision frequency dependence for **(a)** 515 nm and **(b)** 813 nm. Detunings are measured against the respective free-space resonance. Further light-assisted collisions during the imaging process prevent the occupation to reach 100 %, even for large detunings.

leads to a pairwise loss and results in either zero or one atom per trap, depending on whether the initial number of atoms was even or odd. It is therefore also referred to as a *parity-projection* (PP) process and is one of the reasons we need the rearrangement step presented in Chapter 6 to produce defect-free arrays. The resulting atom number distribution is highly sub-Poissonian as $P(N_{\text{atom}} > 1) = 0$ [32, 111] .

The probabilities for obtaining zero atoms $P(0)$ and one atom $P(1)$ depend on the details of the LAC process. In the case described above, where a bound molecular state leaves the trap, there is no preference for an initially even or odd atom number, so we can expect a single atom in 50 % of the cases, thus $P(0) = P(1) = 0.5$. In other atom species, for example rubidium, nearly dark states can be used to increase the probability of loading a single atom [112, 113]. The following discussion is restricted to 515 nm and 813 nm, as we have not investigated light-assisted collisions in tweezers at 532 nm yet. We leave a detailed investigation for the future and load the trap at 532 nm by a handover from 813 nm to investigate the cooling process in Section 5.4.2.

Following [46], we use red detuned light at 689 nm to trigger LACs, but [47] suggests that 461 nm can be used as well. We measure the occupation depending on the detuning of the LAC beam for both 515 nm and 813 nm for durations between 30 ms and 50 ms respectively, averaged over several hundred repetitions of the experiment. As discussed in Section 2.3, we place a threshold at a certain number of scattered photons to determine if the trap is empty. A rigorous treatment of this placement will be presented in Section 5.5. Figure 5.3(a) shows the occupation in a trap at 515 nm where it can be seen that for detunings between −400 to −150 kHz with respect to the free-space resonance, we detect an atom in roughly 50 % of the cases. Going more towards and then beyond the resonance at 0 kHz heats the atoms and ultimately leads to atom loss. For detunings below −400 kHz we are too far detuned from the molecular resonance and can therefore not trigger the LAC process. The fact that 100 % occupation is not reached is caused by additional LACs during the detection process. The frequency dependence is in good agreement with the data from [46]. For 813 nm, shown in Figure 5.3(b), the frequency dependence looks similar but is found at a more negative detuning of around −1 750 kHz for a trap depth of 300 µK. Compared to the resonance, which is light-shifted to around −1 500 kHz, the process is therefore again red-detuned by a few hundred kilohertz.

To confirm that we do not accidentally stop at the time where 50 % occupation is reached while atoms are continuously lost, we measure the occupation depending on the duration of the LAC phase, shown in Figure 5.4(a). As it can be seen from the figure, a plateau of 50 % is reached after around 30 ms which stays until at least 100 ms. A similar argument could be constructed for the loading time of the blue MOT, such that only half an atom is trapped on average. To refute this, we also measure the occupation while varying the loading time of the blue MOT (Figure 5.4(b)). In this case, the average occupation again saturates to 50 % after 40 ms and we conclude that a loading time of 50 ms is sufficient.

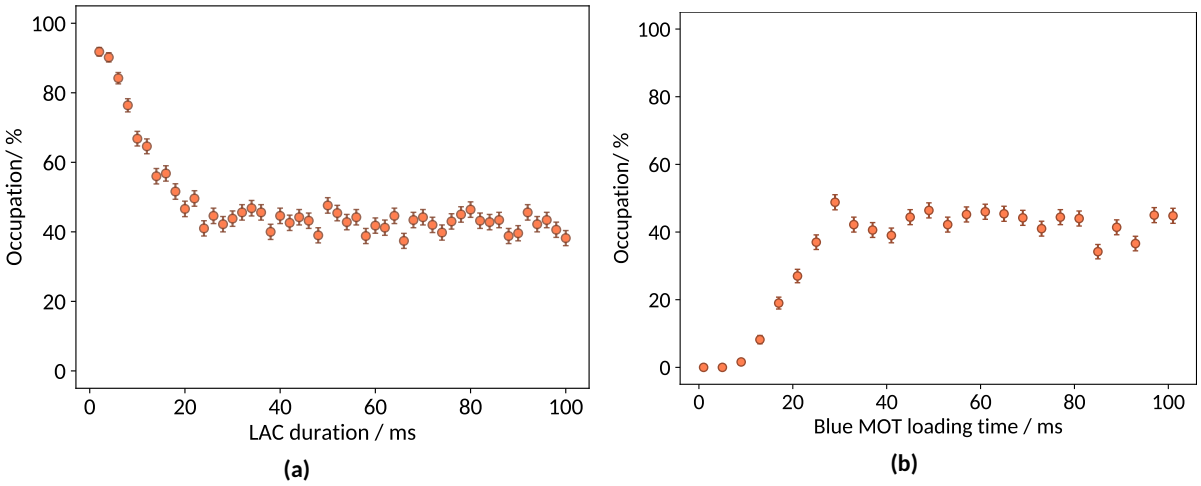

Figure 5.4 – **(a)** Time dependence of the light-assisted collision process. A plateau is reached after 30 ms which shows that the atoms are not continuously lost. **(b)** Tweezer occupation dependency on blue MOT loading time. Again a plateau can be seen as the light-assisted collisions cap the maximum number of atoms per trap. Both measurements are for a single trap at 813 nm.

To verify the binomial nature of the occupation we investigate the histogram of occupied sites in an array of tweezers. The distribution should follow a binomial distribution around the mean $\langle N \rangle = N_{\text{sites}}\, p$ as the occupation for every site in all N_{sites} sites is a Bernoulli-experiment with probabilities p and $1 - p$. An example of such a histogram is shown in Figure 5.5. We fit a binomial distribution to the data and extract a value for p of 48 %. We explain the deviation from the ideal case of 50 % by additional losses due to the finite lifetime and losses during the imaging process. Both of these loss mechanisms are further discussed in Section 5.5 and Section 5.8 respectively.

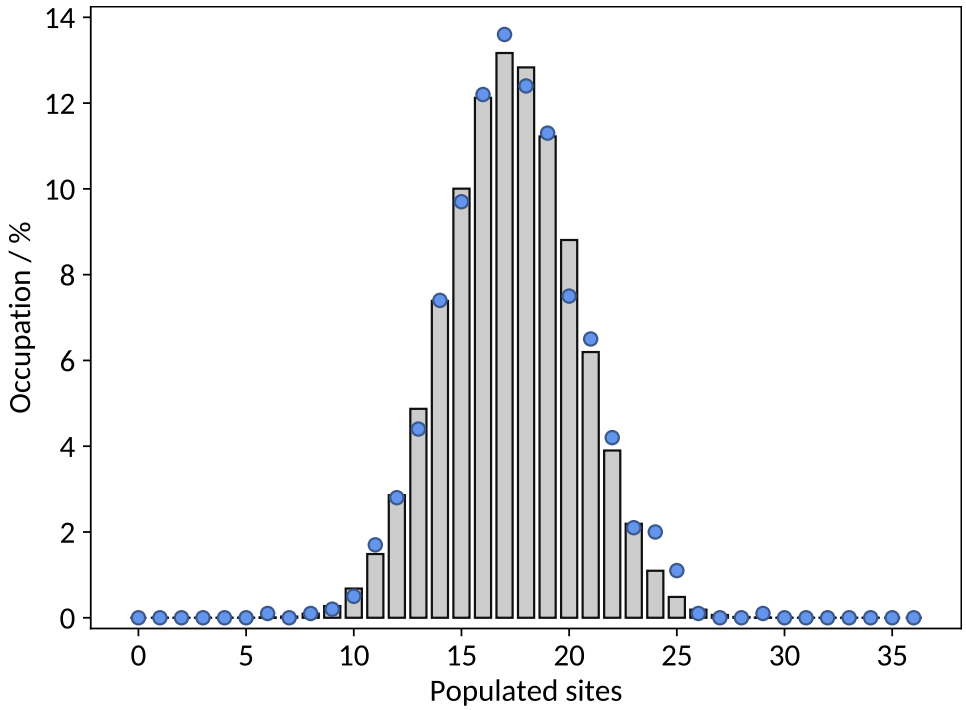

Figure 5.5 – Histogram of occupied sites in a tweezer array with 36 sites for 1 000 realizations of the experiment. The gray bars show a fitted Binomial distribution from which we extract $p \approx 48\,\%$.

5.4 Cooling and thermometry

To experiment with single atoms, we strive to control all degrees of freedom, internal ($|g\rangle$ and $|e\rangle$) and external ($|n\rangle$). In our case this means that we would like to prepare the system in the electronic ground state of the atom in the motional ground state of the harmonic potential. Preparing the electronic ground state is comparably simple, as we can just wait for other states ($^1\mathrm{P}_1$ or $^3\mathrm{P}_1$ specifically) to decay within 21 µs at most. Reaching the motional ground state of the trap is considerably more complex, and we have to actively extract energy from an atom to reduce its temperature. The cooling process is inherently tied to the differential light shift $\Delta\alpha = \alpha_e - \alpha_g$ between the ground state $|g\rangle \equiv {}^1\mathrm{S}_0$ and excited state $|e\rangle \equiv {}^3\mathrm{P}_1$. Three fundamentally different cases exist: a vanishing differential light shift $\alpha_g = \alpha_e$, a stronger confined ground state $\alpha_g < \alpha_e$ and a weaker confined ground state $\alpha_g > \alpha_e$. We can realize these cases for the wavelengths of 515 nm, 532 nm and 813 nm respectively (see Figure 4.1(b)).

5.4.1 Resolved sideband cooling at 515 nm

If the ground state polarizability α_g and the excited state polarizability α_e are equal, the differential light shift vanishes and trap frequencies for both states are equal. In this case, transitions between different motional states are independent of the current state and these *sideband*-transitions $|g, n\rangle \rightarrow |e, n \pm 1\rangle$ can be driven directly. This process is known as resolved sideband cooling and a common cooling technique for ions trapped in radio-frequency traps [114, 115]. We begin with a theoretical description of the process, together with a numerical simulation to estimate parameter limits for the experiment. Afterwards, we present the experimental realization and the result, comparing it to the literature. A

schematic overview is shown in Figure 5.6.

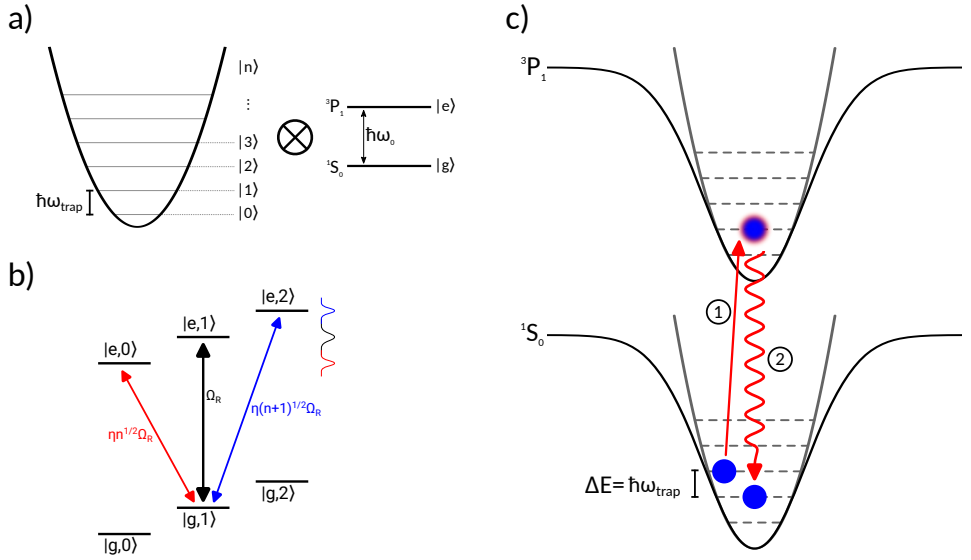

Figure 5.6 – Resolved sideband cooling schematic. **(a)** The state of the trapped atom consists of the external contribution from the confining potential $|n\rangle$ and the internal electronic levels $|g\rangle$ and $|e\rangle$) **(b)** For a dipole transition linewidth γ smaller than the trap frequency ω_{trap} transitions to different motional states $|g,n\rangle \rightarrow |e,n'\rangle$ can be driven directly. The transition strength depends on the Lamb-Dicke parameter $\eta \ll 1$ where transitions with $n' = n \pm 1$ ("sideband-transitions") are weaker by a factor of η compared to $n' = n$ ("carrier-transition"). **(c)** The different transition strengths cause the spontaneous emission to primarily happen on the carrier-transition. Tuning a laser to the red-sideband transition $|g,n\rangle \rightarrow |e,n-1\rangle$ followed by spontaneous emission $|e,n-1\rangle \rightarrow |g,n-1\rangle$ leads to an average reduction in energy per cycle of $\hbar\omega_{\text{trap}}$.

5.4.1.1 Theory

The system can be modeled by the Hamiltonian $\hat{H} = \hat{H}_0 + \hat{H}_{\text{int}}$, where $\hat{H}_0$ describes the trapped two-level atom and $\hat{H}_{\text{int}}$ describes the interaction with an external electric field. $\hat{H}_0$ takes the internal as well as the external degrees of freedom into account [4, 115]:

$$\hat{H}_0 = \hbar\omega_0 \left|e\right\rangle \left\langle e\right| + \hbar\omega_{\text{trap}} \left(\hat{a}^\dagger \hat{a} + \frac{1}{2}\right), \tag{5.1}$$

where ω_0 describes the energy difference between the ground state $|g\rangle$ ($^1\text{S}_0$ in our case) and the excited state $|e\rangle$ ($^3\text{P}_1$ for us). $\alpha_g = \alpha_e$ means that the ground and excited state trapping potentials are equal and we only need a single trap frequency ω_{trap} to describe the harmonic oscillator potential approximation. The interaction Hamiltonian $\hat{H}_{\text{int}}$ describing the coupling of internal and external degrees of freedom by a classical ($\hat{=}$ consisting of many photons) field reads:

$$\hat{H}_{\text{int}} = -\vec{E}\left(\vec{r},t\right) \cdot \vec{\mu}_e \tag{5.2}$$

$$= \frac{\hat{d}E_0}{2}\left(e^{ik\hat{x}} + e^{-ik\hat{x}}\right)\hat{\sigma}_x \tag{5.3}$$

$$\tag{5.4}$$

where we expressed the dipole operator $\vec{\mu}_e$ through $\hat{\sigma}_x = (|g\rangle\langle e| + |e\rangle\langle g|)$ and the electric field in terms of the position operator $\hat{x} = \frac{x_0}{\sqrt{2}}\left(\hat{a} + \hat{a}^\dagger\right)$. $k = 2\pi/\lambda$ is the (absolute value of the) wavevector and $x_0 = \sqrt{\hbar/m\omega}$ is the ground state harmonic oscillator length. Similar to the description in Section 2.1, we remove the trivial time-evolution by transforming into an interaction picture with respect to $\hat{H}_0$:

$$\hat{H}^I_{\text{int}} = -\hbar\Delta\,|e\rangle\langle e| + \frac{1}{2}\hbar\Omega_R\left(e^{i\eta\left(\hat{a}+\hat{a}^\dagger\right)} + e^{-i\eta\left(\hat{a}+\hat{a}^\dagger\right)}\right). \tag{5.5}$$

We defined the Lamb-Dicke parameter $\eta = kx_0$, which measures the extend of the ground state wavefunction x_0 in terms of the wavelength λ. The first requirement for the resolved sideband cooling technique is $\eta \ll 1$ (also called the *Lamb-Dicke*-regime [116]) where the atom is stronger localized than the wavelength of the exciting light. Expressing η as the (square-root of the) ratio of the trap frequency ω_{trap} and the single photon recoil energy $E_r = \frac{\hbar^2k^2}{2m}$ as $\eta = \sqrt{\frac{E_r}{\hbar\omega_{\text{trap}}}}$ shows, that η describes how much the motional state is influenced by the absorption or emission of a photon. The ratio is an important quantity to consider, as it expresses the heating effect due to the photon recoil. The emission follows a dipole pattern oriented along the quantization axis of the system [117] which leads to a coupling between the different axes. Cooling along one axis can therefore increase the motional state of the others, however, with $\eta \ll 1$ this effect is negligible if all axes are cooled simultaneously or alternated with a short cycle duration. Using $\eta \ll 1$, we can further approximate the dipole operator by:

$$e^{i\eta\left(\hat{a}+\hat{a}^\dagger\right)} \overset{\eta\ll1}{\approx} 1 + i\eta\left(\hat{a} + \hat{a}^\dagger\right) + O\left(\eta^2\right). \tag{5.6}$$

This shows, that the dipole operator in the Lamb-Dicke-limit of $\eta \ll 1$ only couples states $|n\rangle$ to states separated by at most one motional quantum $\left|n'\right\rangle \in \{|n\rangle, |n \pm 1\rangle\}$ with higher orders being strongly suppressed.

The second ingredient for the cooling scheme is called the *festina-lente*-regime [118, 119], where the linewidth γ of the driven dipole transition is smaller than the trap frequency ω_{trap}. In this case, it is possible to drive these transitions directly – the *sidebands* can be *resolved*[2]. The key factor for the cooling effect is the mismatch of Rabi frequencies for the three transitions $|g,n\rangle \leftrightarrow |e,n\rangle$ and $|g,n\rangle \leftrightarrow |e,n \pm 1\rangle$. The Rabi frequencies for the sidebands scale with η, so the Rabi frequencies for the blue- and red-sidebands $\Omega_{\text{RSB/BSB}}$ are related to the Rabi-frequency main- or *carrier*-transition Ω_{C} by $\Omega_{\text{RSB/BSB}} \sim \eta\Omega_{\text{C}}$. This means, that an excited state $|e,n\rangle$ is more likely to decay to $|g,n\rangle$ than to $|g,n \pm 1\rangle$, as the latter transition is weaker by a factor of η. On average the atom loses an energy of $\hbar\omega_{\text{trap}}$ when an excitation on the red-sideband $|g,n\rangle \rightarrow |e,n-1\rangle$ is followed by an emission on the carrier-transition $|e,n-1\rangle \rightarrow |g,n-1\rangle$. The cooling stops when $|g,n=0\rangle$ is reached, as the red-sideband transition cannot be driven anymore. Two first order heating processes limit the achievable temperature: Off-resonant excitation on the carrier transition, followed by spontaneous emission on the red sideband $|e,n\rangle \rightarrow |g,n+1\rangle$ or off-resonant excitation on the blue sideband $|g,n\rangle \rightarrow |e,n+1\rangle$ followed by emission on the carrier transition. Combining these processes in a heating rate allows us to

[2] If this is not the case, the cooling loses its efficiency as other (heating) transitions are also driven which limits the achievable minimum quantum number as shown in Equation (5.7). This is the case for the $^1\text{S}_0 \rightarrow {}^3\text{P}_1$ transition in ytterbium with a linewidth of around 180 kHz. In this case Raman sideband cooling can be used [120–122].

express the limit as [123]:

$$\langle n \rangle_{\text{min}} = \frac{5}{16} \frac{\gamma^2}{\omega_{\text{trap}}^2} \tag{5.7}$$

with the transition of the linewidth γ. For $\omega_{\text{trap}} \approx 10\gamma$ we compute $\langle n \rangle_{\text{min}} < 0.01$ which means a motional ground state fraction of around 99 % can be reached.

We can simulate the time-evolution of the system using the Lindblad Master-equation [124] which can model the coupling of a closed system to the environment, leading to non-unitary time-evolution. Using a Master-equation solver from the *Python*-package `Qutip`, we simulate the cooling process numerically with the results shown in Figure 5.7. It can be seen, that the initial coherent Rabi-oscillations quickly decay and that the population accumulates in the ground state. Heating between different axes was omitted as the real sequence alternates between different axes, and in the Lamb-Dicke regime this leads to negligible heating rates, as the energy increase due to photon recoil is small compared to the energy loss per scattered photon. The result compares well to [47] and suggests that a few ten milliseconds of cooling are enough even in the presence of laser frequency and intensity noise.

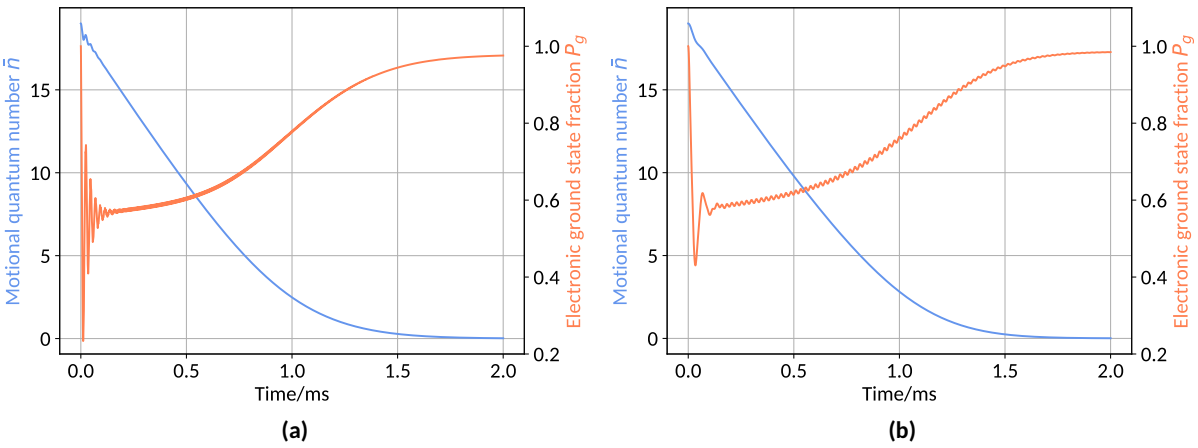

Figure 5.7 – Numerical Master equation simulation of the sideband cooling process in one-dimension for the radial **(a)** and axial direction **(b)** . We choose parameters close to the experimental parameters with $\omega_r = 2\pi \times 217$ kHz and $\Omega_r = 2\pi \times 65$ kHz (radial) and $\omega_a = 2\pi \times 43$ kHz and $\Omega_a = 2\pi \times 10$ kHz (axial) respectively. We truncate the trap at $N = 20$ motional states.

5.4.1.2 Experimental implementation and calibration

The differential polarizability vanishes for the $^1S_0 \rightarrow {}^3P_1$ transition for a wavelength around 515 nm ("magic wavelength"). However, to fine-tune the compensation down to a 10 kHz-level, external magnetic fields and the polarization of the tweezer beam also have to be considered. The correction arises from vector and tensor components α_v and α_t in the polarizability, see Section 4.1.

Two approaches to fine-tune the differential light-shift have been demonstrated in the literature: The first one achieves the *magic* condition by using an elliptical polarization with the ellipticity angle γ ("magic ellipticity angle"). In this case, the elliptic polarization acts as an effective magnetic field $\vec{B}_{\text{eff}} \propto \alpha_v E_0^2 \sin(2\gamma)\, \hat{e}_z$, with the vector polarizability α_v, and is directed in the direction of tweezer propagation z (axial direction). $\vec{B}_{\text{eff}}$ couples to the total angular momentum vector $\vec{J}$ like usual. In

addition, the elliptic tweezer-polarization vector $\vec{\epsilon} = (\cos\gamma, i\sin\gamma, 0)$ also couples to $\vec{J}$ and in turn to the vector and tensor polarizabilities α_v and α_t. By changing γ, the relative contributions of the vector and tensor coupling can be tuned to compensate the difference in polarizabilities for the ground and excited states. Further details for the method can be found in [46]. The second method uses a strong bias field $\vec{B}_{\text{bias}}$ in the tweezer plane (radial plane) at an angle θ_B ("magic angle") to the linear tweezer polarization, described in [47]. Similar to $\vec{B}_{\text{eff}}$, $\vec{B}_{\text{bias}}$ couples to $\vec{J}$, however no coupling to the vector polarizability exists. In contrast to the elliptic angle method, the strong magnetic field induces a Zeeman-shift which dominates the light shift and defines the quantization axis of the system. By tuning the magnetic field angle, the Zeeman-resonances can be shifted to compensate for the light-shift, however, the required magnetic field strength also depends on the tweezer depth. Details for the second method can be found in [47]. Up to now, we only implemented the second method experimentally where we made this choice as we found it more convenient to change the field strength and angle, which we could easily automate compared to fine-tuning the polarization.

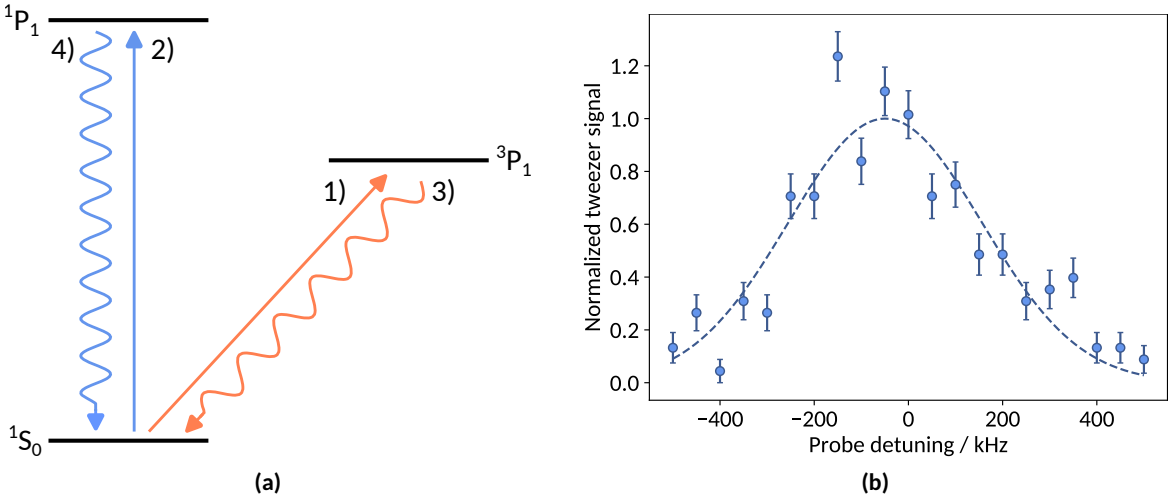

Figure 5.8 – **(a)** Depletion spectroscopy scheme used to determine the $^1S_0 \rightarrow {}^3P_1$ frequency. **1)** Atoms are shelved to $^3P_1$ by a 100 µs long probe pulse. **2)** During the probe pulse, a short (< 10 µs) blow away pulse, resonant to the $^1S_0 \rightarrow {}^1P_1$ transition, is applied. This pulse removes all atoms in the electronic ground state from the trapping region. **3)** 200 µs after the probe pulse atoms have returned to electronic the ground state. **4)** The previously shelved fraction is detected on the standard imaging transiton. **(b)** Spectrum of the $^1S_0 \rightarrow {}^3P_1$ transition, measured with many atoms using the depletion technique. The resonance can be shifted by changing the magnetic field angle with a sensitivity of 100 kHz/°. After the peak is aligned with the free-space resonance at zero detuning, the measurement is repeated with a single atom to reduce the width of the resonance and increase the resolution.

To determine the correct angle magnetic field angle θ_B, we start with a field of 20 G at an angle of about 30° relative to the tweezer polarization axis [47]. As described in Section 5.3, we load the tweezer by overlapping it with the red MOT for 20 ms. To simplify the initial determination of θ_B, we do not employ light-assisted-collisions and determine the light shift by measuring the $^1S_0 \rightarrow {}^3P_1$-transition frequency using a depletion technique, shown in Figure 5.8(a). After employing a 689 nm pulse along any direction for 100 µs, we switch on a strong ($s \gtrsim 20$) *blow-away*-pulse resonant to the $^1S_0 \rightarrow {}^1P_1$ transition for 5 µs. During this pulse, we briefly switch off the trap for 1 µs which removes all atoms in the electronic ground state from the trapping region. We let the atoms return to the ground state by waiting for an additional 200 µs, and image with a lossy one-shot imaging pulse on the 461 nm

transition afterwards. The resulting spectrum is shown in Figure 5.8(b). The transition is broadened by the temperature of the atoms as they enter the dipole trap with a potential energy up to roughly the trap depth[3]. We can tune the carrier transition using the magnetic field angle relative to the tweezer polarization with a sensitivity of about 100 kHz/°. After the peak is aligned with the free-space resonance at zero detuning, the measurement is repeated with a single atom to reduce the width of the resonance and increase the resolution to a level of 10 kHz.

With the carrier-transition aligned to the free-space resonance, we can now proceed with the cooling procedure. To calibrate the cooling power for three orthogonal axes, we drive Rabi-oscillations from which we determine the local intensity at the position of the atoms. An example of a fitted Rabi-oscillation is shown in Figure 5.9(a) where we determine a Rabi-frequency of $2\pi \times (46 \pm 2)$ kHz and a decay constant of $\tau = (15 \pm 3)$ µs. The latter is lower than the lifetime of the $^3P_1$-state of 21 µs which we explain by fluctuating intensity during the probe time, leading to shorter coherence times. Due to limited optical access we have to use the retroreflected red MOT beams, which might not be ideal as standing wave or polarization-gradient effects could influence the measurement. This claim is supported by the fact that we observe longer coherence times when using a non-retroreflected beam. For the cooling process, however, itself it seems to not be a problem.

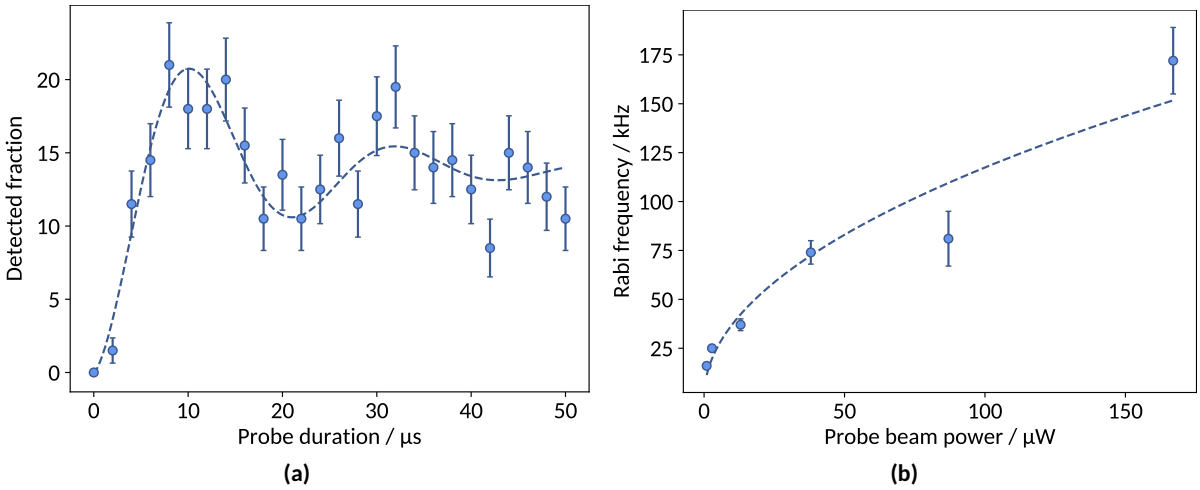

Figure 5.9 – Beam power calibration with Rabi oscillations. **(a)** Oscillation with $\Omega = 2\pi \times (46 \pm 2)$ kHz. **(b)** Measured Rabi frequencies for different beam powers.

5.4.1.3 Resolved sideband spectroscopy

After loading the tweezer and reducing the atom number using light-assisted collisions, we shine in cooling light along three orthogonal directions for 50 ms. Radially, we cool at a detuning of $-\omega_r$ with a Rabi frequency of $\Omega_r = 2\pi \times 65$ kHz. Axially, we start at $-5\omega_a$ at a Rabi frequency of $\Omega_a = 2\pi \times 50$ kHz and switch to $-\omega_a$ and $\Omega_a = 2\pi \times 10$ kHz after 2/3 of the cooling duration. Figure 5.10 shows the measured sideband spectra after employing sideband cooling for 50 ms. Probing is done at the final cooling Rabi frequencies with a 100 µs (200 µs) long pulse radially (axially). From the extracted trap

[3] The fitted width of 200 kHz corresponds to a Doppler temperature of 0.2 mK which is a bit lower than the trap depth of 1 mK. Additional cooling or removal of hot atoms during the imaging process can explain this deviation.

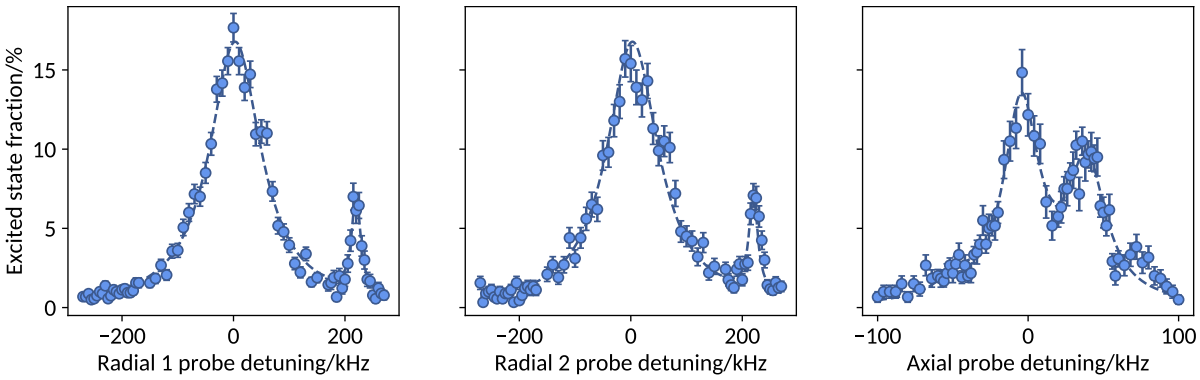

Figure 5.10 – Sideband spectra after employing sideband cooling for 50 ms. In the axial direction, second order sideband-transitions can be seen due to the Lamb-Dicke parameter of $\eta = 0.33$ compared to the radial direction with $\eta = 0.15$. From the fits, we extract trap frequencies of (217 ± 1) kHz, (217 ± 2) kHz and (43 ± 1) kHz for two orthogonal radial directions and the axial direction respectively.

frequencies of $\omega_{r_1} = 2\pi \times 217\,\text{kHz}$, $\omega_{r_2} = 2\pi \times 217\,\text{kHz}$ and $\omega_a = 2\pi \times 43\,\text{kHz}$ we can calculate the aspect ration ω_r/ω_a, the waist of the trap after Equation (4.8) and the trap depth U_0 to be:

$$\frac{\omega_r}{\omega_a} = \frac{2\pi \times (217 \pm 2)\,\text{kHz}}{2\pi \times (43 \pm 1)\,\text{kHz}} = 5.05 \pm 0.13$$
$$w_0 = \frac{\omega_r}{\omega_a}\frac{\lambda}{\sqrt{2}\pi}\beta = (439 \pm 11)\,\text{nm}$$
$$U_0 = \frac{\omega_r^2 m w_0^2}{4k_\text{B}} = 0.95\,\text{mK}$$

where $\beta\,(\xi \approx 0.76) = 0.75$ is the scale factor to account for the truncation of the Gaussian beam (see Section 4.2). The aspect ratio is slightly higher compared to the value obtained from the simulation of 4.9 at this truncation ratio, but this is to be expected due to optical aberrations[4]. It compares well to the values in the literature summarized in Table 5.1. The lower aspect ratio in our case compared to the literature for a similar NA could be explained the fact that we used a single trap at 515 nm without the AODs. As the AODs require more optical elements for beam shaping and polarization rotation we expect larger wavefront aberrations. We support this claim by the fact that we observe a larger aspect ratio after the installation for the measurements in Section 5.7.

5.4.1.4 Thermometry and probe influence

We can directly determine the mean occupation number $\langle n \rangle$ from the height of the sidebands [115]:

$$\langle n \rangle = \frac{A_\text{RSB}}{A_\text{BSB} - A_\text{RSB}} \tag{5.8}$$

[4] In addition to misalignments in the setup causing astigmatisms and comas, there is a residual spherical aberration from the mismatch of the vacuum window thickness from the value the objective is designed for. See Appendix A.4 for an overview of optical aberrations.

where A_{RSB} is the height of the red-sideband and A_{BSB} is the height of the blue sideband. We fit a superposition of Lorentzian-peaks to the spectrum with seven free parameters: one amplitude for each peak, and the carrier- and sideband-positions and -widths. We assume the red- and blue-sideband positions to be equal in magnitude with respect to the carrier frequency, and we also assume an equal width for both sidebands. The resulting $\langle n \rangle$ are $\langle n_{r_1} \rangle = 0.00^{+0.07}_{-0.00}$, $\langle n_{r_2} \rangle = 0.00^{+0.07}_{-0.00}$ and $\langle n_a \rangle = 0.14^{+0.05}_{-0.05}$.

Care has to be taken to not underestimate $\langle n \rangle$: During the probing process the atom is further cooled (heated) when probing the red (blue) sideband which has to be taken into account. To estimate the influence of the probing process, we again use a numerical simulation. We simulate a sideband spectrum with an initial $\langle n \rangle_{\mathrm{initial}}$ that we then fit to extract the fitted $\langle n \rangle_{\mathrm{fitted}}$. The result is shown in Figure 5.11 and compares well to the literature [47]. Note that the simulation fails to predict the behavior for the axial direction at $\langle n_a \rangle \lesssim 0.05$ correctly, as the fitted value is *larger* than the underlying value. It shows that the radial (axial) direction is underestimated by a factor of ∼1/3 (3/4) and the corrected values read:

$$\langle n_{r_1} \rangle = \langle n_{r_2} \rangle = 0.00^{+0.21}_{-0.00}$$
$$\langle n_a \rangle = 0.15^{+0.09}_{-0.07}.$$

Comparing this result to [47] shows that we can achieve a better result in the radial direction. In the axial direction our final $\langle n \rangle$ is higher as [47] but lower than [46]. We followed the latter for the cooling protocol where the cooling detuning and amplitude are changed during the process. This might not be optimal and the constant cooling approach from [47] could result in better results at that level. Nevertheless, we conclude that the result compares well to the literature and that we can prepare an atom in the three-dimensional ground state with 95^{+2}_{-10} %.

Source	NA	Aspect ratio ω_r/ω_a	Ground state fraction / %		
			Radial	Axial	Total
[46]	0.5	6.6	83 ± 14	52 ± 7	72 ± 7
[47]	> 0.65	5.5	95^{+4}_{-16}	100^{+0}_{-10}	97^{+2}_{-11}
Our result	0.7	5.1	100^{+0}_{-15}	85^{+7}_{-9}	95^{+2}_{-10}

Table 5.1 – Comparison of our results to values from the literature.

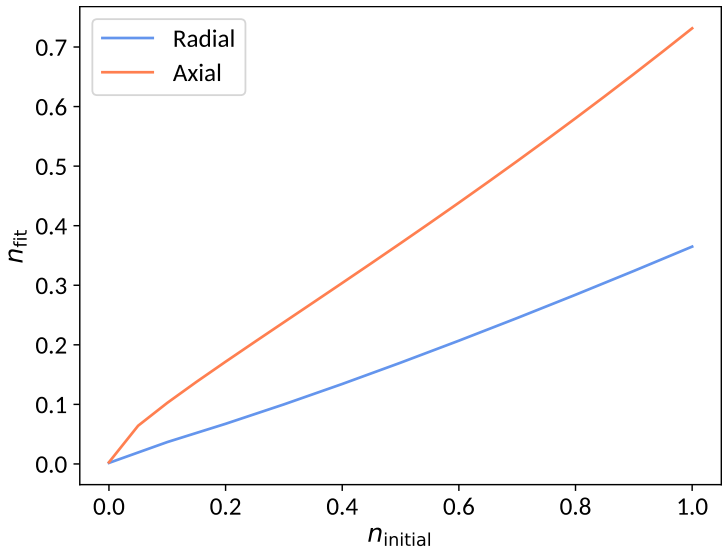

Figure 5.11 – Probe beam influence simulation. To estimate the effect of the probing pulse on the perceived motional state, we simulate the probe pulse on a system with an initial state n_{initial}. We fit the spectrum and compare the fitted n_{fit} to n_{initial} to determine a correction factor for the experimentally determine values.

Carrier shift

Although not strictly needed for this work, we also investigated tweezer arrays at 515 nm created with both the SLM and the AODs. We noted an interesting behaviour of the positions of the carrier peaks when we tried to perform sideband cooling in an array that is worth mentioning. Figure 5.12(a) shows the carrier detuning extracted from the sideband spectrum for a 6×6 rectangular array. A left-to-right gradient, corresponding to the y-axis in the experiment, is clearly visible. The absolute values can be shifted with the magnetic field angle with a slope of 100 kHz/° and the gradient can be inverted if the angle is inverted with respect to the tweezer polarization axis. We explain this phenomenon by a slight rotation of the tweezer polarization for different traps. We exclude a difference in trap depths, as this would also lead to gradient in sideband frequencies, however, no such signature is observed. Every trap corresponds to a beam component with a unique k-vector which results in slightly different propagation angles and therefore directions. The complex beamsplitter coating that is used to combine (or split) the different wavelengths before or after the objective shows a strong dependence of the phaseshift on the incident beam angle (see Figure 5.12(b)). Translating the different k-vectors into angles leads to differences on the order of a few degrees, leading to the observed rotation of the polarization of around ±0.5°. For the rearrangement of atoms described in Chapter 6 this is not an issue, as we do not cool the atom in the dynamic trap at 515 nm but in general it is problematic. By designing a new coating with a zero phase-shift for both s- and p-polarization around 45° at 515 nm the problem could be solved.

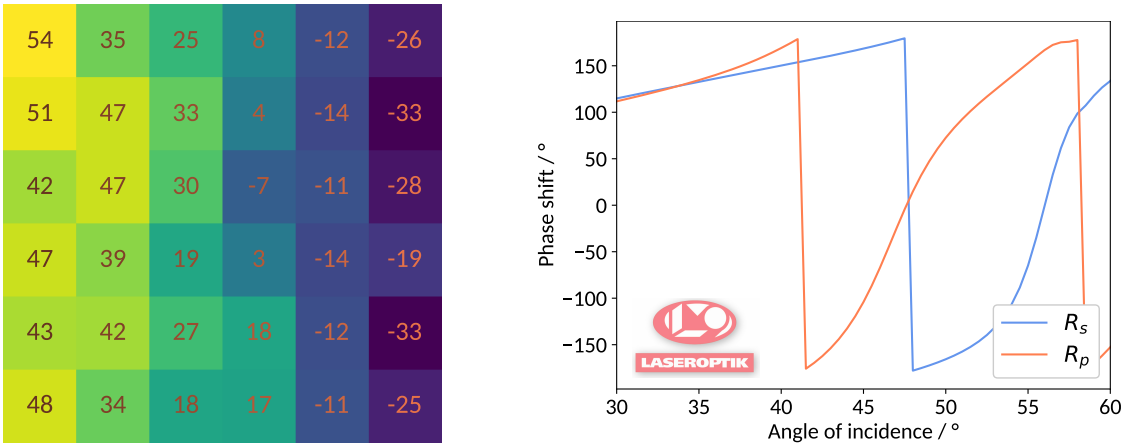

Figure 5.12 – **(a)** Carrier positions in kilohertz, extracted from the sideband spectrum for a 6×6 array. **(b)** Phaseshift of the dielectric coating in front of the objective, for s- and p-polarization, depending on the incident angle. Data courtesy of *Laseroptik GmbH*.

5.4.2 Sisyphus cooling at 532 nm and 813 nm

In general, the polarizabilities for the ground-and the excited state are not equal. This means that the ground and excited state trap-frequencies are different and the sideband transitions cannot be driven as before, as the transition-frequency depends on the current state. Therefore, a different cooling scheme has to be used, for example the sisyphus-like process depicted in Figure 5.13. This cooling protocol has its name from the Greek myth, where Sisyphos has to roll a boulder up a hill time and again only to see it falling down just before reaching the top. The term has been in use since the late 1980's for different sub-Doppler cooling techniques like polarization-gradient cooling [60, 125]. A periodic potential landscape constructed by clever choices for the polarization causes atoms to absorb only in a valley of the potential which results in a net-loss of energy by continuously *climbing the hill*[5]. Later, it was studied theoretically for laser-induced potentials in general [117] and specifically for strontium and ytterbium atoms in optical tweezers [126]. Experimental realizations for strontium in optical tweezers include [46, 48, 56].

5.4.2.1 Theory

$\alpha_g < \alpha_e$ The left part in Figure 5.13 shows the process as a sequence of four steps: A laser tuned to the transition frequency at the center of the trap excites the atom (1). Once excited, the atom moves in the trap, potentially "climbing the hill" to the outer regions of the trap (2). The atom de-excites (relaxes?) and emits a photon with *at least* the energy it absorbed during the excitation (3). Back in the ground state the atom moves around until is excited again while being close to the center of the trap (4). To understand the differences between this cooling protocol and other (sub-Doppler) cooling protocols, we will briefly summarize the main results from [126]. In this letter the authors compute and simulate characteristics of the cooling process for strontium and ytterbium in optical tweezers using a classical model.

[5] The process was actually discovered after experimenters observed temperatures lower than the Doppler-limit.

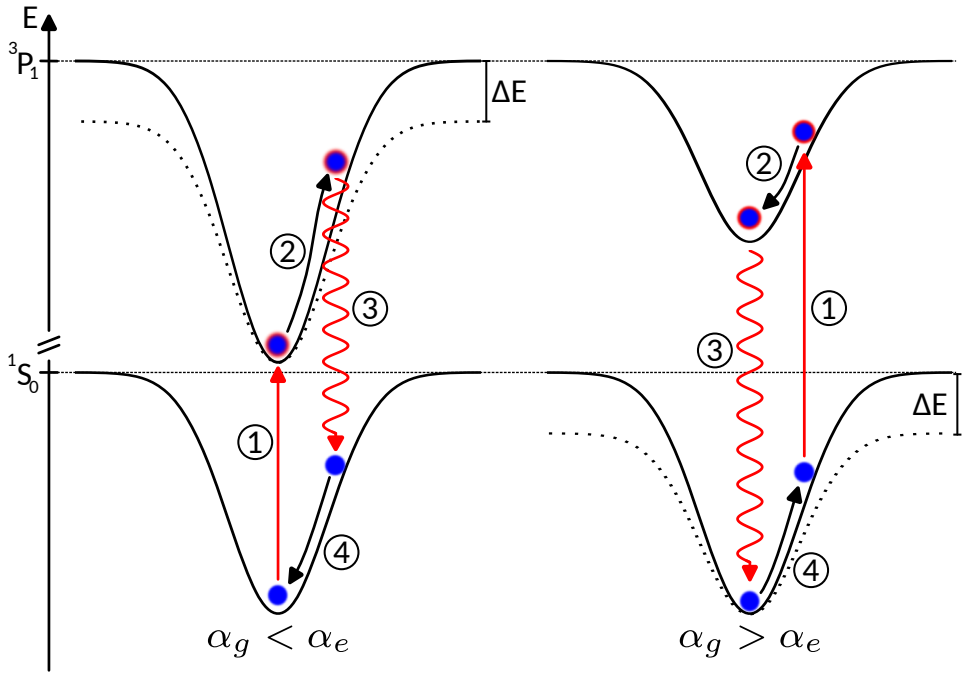

Figure 5.13 – Classical picture of sisyphus cooling processes in an optical tweezer. An energy difference between the excitation (red solid arrow) and the spontaneous emission (red wiggly line) leads to a cooling effect, if the excitation happens at the correct position inside the trap. **Left:** For $\alpha_g < \alpha_e$ a laser tuned to the energy difference at the trap center excites an atom. As it moves inside the trap, it emits at a position with a larger energy difference, leading to a net energy loss. **Right:** For $\alpha_g > \alpha_e$ the excitation has to happen away from the trap center where the energy difference is largest. Compared to $\alpha_g < \alpha_e$, the optimal detuning depends on the current temperature of the atom. The separation between the ground- and the excited states compared to the trap depth is not drawn to scale.

The average energy lost per scattering event is given by:

$$\langle E \rangle = k_B T \left(1 - \frac{\alpha_g}{\alpha_e}\right) \frac{1}{2} \frac{1}{1 + \left(\gamma/2\omega_e\right)^2} \tag{5.9}$$

computed for polarizabilities α_g and α_e and trap frequencies ω_g and ω_e for the ground- and excited states respectively. where $\alpha_g/\alpha_e = \omega_g^2/\omega_e^2 < 1$. For $\omega_e^2 \gg \gamma^2$ (in our case $\omega_e \approx 10\gamma$) we get an average reduction of:

$$\langle E \rangle \approx \frac{1}{2} k_B T \left(1 - \frac{\alpha_g}{\alpha_e}\right). \tag{5.10}$$

This is an interesting result, as it shows that the energy reduction per cycle is proportional to the temperature of the atom, in contrast to e.g. the sideband cooling protocol where it is constant at $\hbar\omega_{\text{trap}}$ or to conventional Doppler-cooling where it is given by photon momentum $\hbar^2 k^2/2m$. Intuitively this is also clear: An atom with higher energy can "climb" higher in the potential of the excited state and thus lose more energy in a single event on average. Emission at higher potential energies are favorable as the probability for the position of emission is inversely proportional to the velocity at that point. The process can therefore reduce the temperature significantly with only a few tens of scattered photons and

is finally limited by either the recoil- or the Doppler-temperature, whichever is higher. While the former is obvious, the latter is due to the uncertainty in the position of excitation due to the finite linewidth of the transition, which ultimately limits the cooling efficiency. For cooling strontium atoms on the $^1S_0 \rightarrow {}^3P_1$ transition, both temperatures of $T_R = 500\,\mathrm{nK}$ and $T_D = 180\,\mathrm{nK}$ (see Table 3.4) are smaller than the ground state temperature of several millikelvin, and to understand the cooling process at this level a quantum mechanical description is required.

$\alpha_g > \alpha_e$ The other option is that the excited state experiences a weaker trapping force than the ground state. In this case a sisyphus-cooling process can again be used but the physics changes significantly. The process has been investigated for strontium in optical tweezers at 515 nm [46] and 540 nm [127] before and also in long working distance tweezers at 532 nm [100]. Compared to the $\alpha_g < \alpha_e$ case, it was shown that tuning the cooling parameters during the cooling process is required to achieve similarly low temperatures.

For a positive differential light shift, exciting an atom at the center of the trap would *heat* the atom. The sequence therefore has to be changed to the following: A laser blue-tuned to the free-space transition frequency excites the atom (1). Once excited, the atom moves in the trap (2). The atom de-excites and emits a photon with *at least* the energy it absorbed during the excitation with a high probability (3). Back in the ground state the atom moves around until is excited again on the outside of the trap (4).

The main difference can intuitively be understood as the optimal detuning of the cooling beam now depends on the current temperature of the atom which suggests that a static cooling process as before is suboptimal. It can also happen that the atom is heated, if the excitation happens to close to the center of the trap or if the emission is on the far outside. A second difference lies in the fact that for maximum energy loss per cycle, the atom has to emit close to the center of the trap which is the position it spends the least time at, thus reducing the cooling efficiency.

5.4.2.2 Release-and-recapture technique

To probe the resulting temperature, we have to resort to a different method than the resolved sideband spectroscopy used before. Although it is possible to use this method at trapping wavelengths with non-zero differential light shifts [56], the lack of a repumping laser for the $^3P_1 \rightarrow {}^3S_1$-transition prevents us from using this spectroscopy scheme. We therefore utilize a release-and-recapture method which has been used to probe single atoms in optical lattices [128] and also in optical tweezers [34]. By rapidly switching off the potential, the energy of the atoms is unaltered and they disperse according to their current energy. After a time t_d ("dark-time"), the traps are turned on again and an atom is recaptured if the kinetic energy is less than the potential depth at the new position. For single atoms this measurement gives a robust, binary signal, as an atom is either lost or recaptured and no information must be extracted from e.g. varying levels of fluorescence. To ensure that every atom that is recaptured is also detected, we insert a second cooling block after the dark-time before imaging.

A common approach to estimate the temperature of the atoms is to compare the measurement results to a classical Monte Carlo simulation [34, 48, 127]. To simulate the system, a thermal Maxwell-Boltzmann-distribution is assumed and position-velocity pairs are randomly drawn. We calculate the evolution in steps of 1 µs and average 10^5 trajectories for the final result. The positions follow normal distributions with $\sigma_{r/a} = \sqrt{k_B T/m\omega_{r/a}^2}$ with the trap frequencies $\omega_{r/a}$, and the velocity components follow normal distributions with standard deviations of $\sigma_v = \sqrt{k_B T/m}$. Free parameters in the simulation are the trap

size, the trap depth (or trap frequency) and the temperature.

Figure 5.14(a) shows a simulation for a trap depth of 300 µK for different temperatures T. The latter can be compared to the temperature associated with the ground state $T_{\mathrm{gs}} = E_{\mathrm{gs}}/k_{\mathrm{B}}$ with Boltzmann's constant k_{B}. The three dimensional ground state energy is $E_{\mathrm{gs}} = 2E_r + E_a$ with $E_i = \hbar\omega_i/2$. With $\omega_r \gtrsim \omega_a/5$ for our optical tweezers, we have $E_{\mathrm{gs}} \approx 2E_r$ from which $T_{\mathrm{gs}} = \hbar\omega_r/k_{\mathrm{B}}$ follows.

We repeat the simulation for different trap depths (see Figure 5.14(b)), aspect ratios ω_r/ω_a and also trap frequencies ω_r, but we find the result to be rather insensitive to all three quantities. We attribute this observation to the fact that the harmonic oscillator length is much smaller than the size of the potential. The mismatch between ω_r and ω_a leads to easier recapture along the axial direction and the recapture is therefore mainly determined by the radial direction.

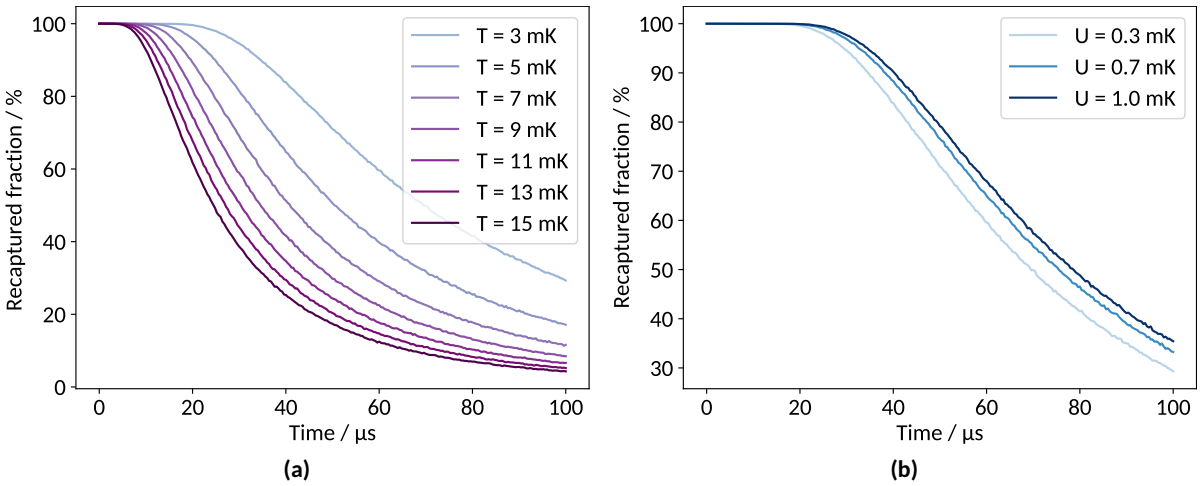

Figure 5.14 – Release-recapture simulation for different temperatures **(a)** and trap depths **(b)**. It can be seen that the recapture behavior is mainly depending on the temperature T and not on the depth of the potential. We further find the simulation to be insensitive to changes in trap frequencies over an order of magnitude, as the harmonic oscillator length is much smaller than the size of the potential. The difference in potential size between the radial and axial direction also makes the behavior to be insensitive to the aspect ratio ω_r/ω_a, as the dynamics is governed by the radial direction.

5.4.2.3 Thermometry for 813 nm

We optimize the cooling parameters at a dark-time of 50 µs. Figure 5.15(a) shows the dependence detuning for a power of 200 µW where a light-shifted resonance lies at −2 500 kHz. Figure 5.15(b) shows the time dependence at a detuning of −2 750 kHz for three difference powers. The recaptured fraction of ∼ 60 % after 60 µs compares well to the literature [48, 56]. Converting the optimal power of 200 µW to an intensity for the 1.35 mm ($1/e^2$-radius) cooling beam gives a value of $2500 I_s$ with $I_s = 3.4\,\mu\mathrm{W/cm^2}$. This result seems unrealistic, as the scattering rate saturates to $\gamma/2$ at such high intensities and the differences in Figure 5.15(b) should be less prominent. We suspect that the cooling beam is not centered on the tweezers and to verify this, the cooling beam can be displaced in the radial plane to map out the correct positioning. Comparison to the literature with values of $200 I_s$ [48] and $90 I_s$ [56] supports this argument. Another option would be to calibrate the cooling intensity by measuring Rabi oscillations, similar to Section 5.4.1.

To estimate the temperature using the release-and-recapture method we calculate the trap size and

depth from the radial trap frequency of $\omega_r = 2\pi \times 60\,\mathrm{kHz}$ measured in Section 5.7. Together with a known optical power of 9 mW per trap, we compute a waist of 950 nm and a depth of 300 µK. This waist is larger than the theoretical limit, where the deviation is probably caused by aberrations due to a slight misalignment of the upper objective (for the alignment procedure see Appendix B). We fit the data with our simulation for different temperatures, but also for different trap depths and aspect ratio. We find the resulting temperature to be is insensitive to the axial trap frequency and trap depth and we observe differences of less than 0.5 µK when varying the aspect ratio between 4 to 10 and the trap depth between 300 to 600 µK. We show a measurement for a single atom, a 5×5 array and a best-fit simulation with a temperature of $T = 2.8\,\mu\mathrm{K}$ for the single atom case in Figure 5.16.

In result, the determined temperature is close to the temperature corresponding to the ground state $T_{\mathrm{gs}} = \hbar\omega_r / k_{\mathrm{B}} \approx 3\,\mu\mathrm{K}$. The extracted temperature is also consistent with the literature [56]. For a more accurate comparison, the trap waist has to be determined more precisely and a fully quantum mechanical model is needed to include the discrete level structure and ground state energy. To avoid this complication it would be helpful to develop a method that does not rely on a comparison with a simulation.

This could be possible by closer inspection of the difference between the single atom versus array behavior, also shown in Figure 5.16. In an array, we measure larger recaptured fractions at longer dark times compared to the single tweezer case. We explain this observation by the fact that atoms can be re-trapped in other tweezers as they mainly expand in the radial direction. The change of trend around 60 µs compared to the single tweezer can be used to roughly estimate the temperature of the sample. Assuming a lattice spacing of 5.2 µm, atoms have to travel at least half that distance to be re-trapped in a different trap. From the distance and time we calculate a velocity of 40 mm/s which equals a temperature of 8 µK. This temperature is not too far off for such a quick *back-of-the-envelope* calculation and by repeating the measurement for different lattice spacings it could be possible to infer the temperature without relying on a simulation with free parameters.

Additionally, the value for the cooling detuning can be used to estimate the trap depth as a cross-check by comparing the result to the literature [48, 56]. However, we have to be careful when comparing the absolute values of the detunings as they depend on the light-shifted sublevel. In addition to the cooling resonance at −2 750 kHz in Figure 5.15(a), we measure a second cooling frequency at −1 600 kHz. The optimal cooling frequency depends linearly on the trap depth and we compute a trap depth of $-1\,600\,\mathrm{kHz} \times \frac{450\,\mu\mathrm{K}}{-2\,600\,\mathrm{kHz}} \approx 280\,\mu\mathrm{K}$ where 450 µK (−2 600 µK) is the trap depth (cooling detuning) from [48]. Comparing to [56], the authors state optimum cooling at −775 kHz at a depth of 135 µK, which translates to a depth of 280 µK in our case. Both estimates are in excellent agreement with our trap depth estimation from the polarizability used for the simulation above. Cooling on the far detuned resonance is favorable, as the energy loss depends on the differential light shift (see Equation (5.10)) and it profits from a larger mismatch. For this work, we therefore always cool around −2 750 kHz for 300 µK deep traps at 813 nm.

5.4.2.4 Thermometry for 532 nm

Similar to the $\alpha_g < \alpha_e$ case, we perform thermometry with a release-and-recapture technique. Figure 5.17(a) shows the survival probability after a dark-time of 6 µs for different detunings of the cooling beam. The measurement follows the expectation: The strongest cooling effect happens between the free-space resonance and the blue-shifted resonance (dashed and dotted vertical lines). Comparing the temperature to the result from 813 nm in Figure 5.17(b) shows that the final temperature is higher. From the Monte-Carlo simulation for a waist of 500 nm and a trap depth of 2 mK we obtain a temperature

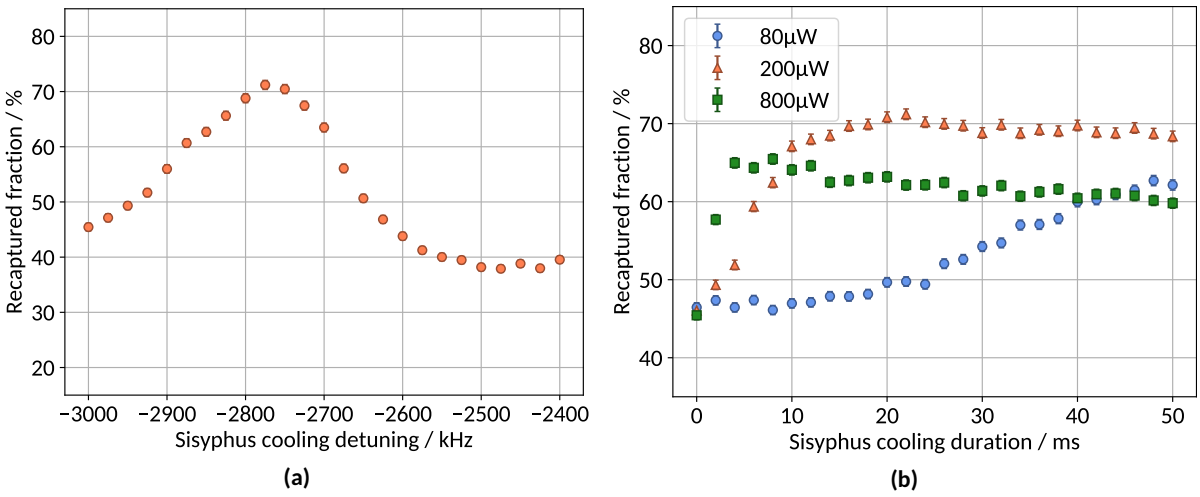

Figure 5.15 – Sisyphus detuning and cooling duration 300 µK deep tweezers at 813 nm, measured after a darktime of 50 µs in a 5×5 array at 813 nm. Data in **(a)** is taken at a cooling power of 200 µW and **(b)** is at a detuning of −2 750 kHz.

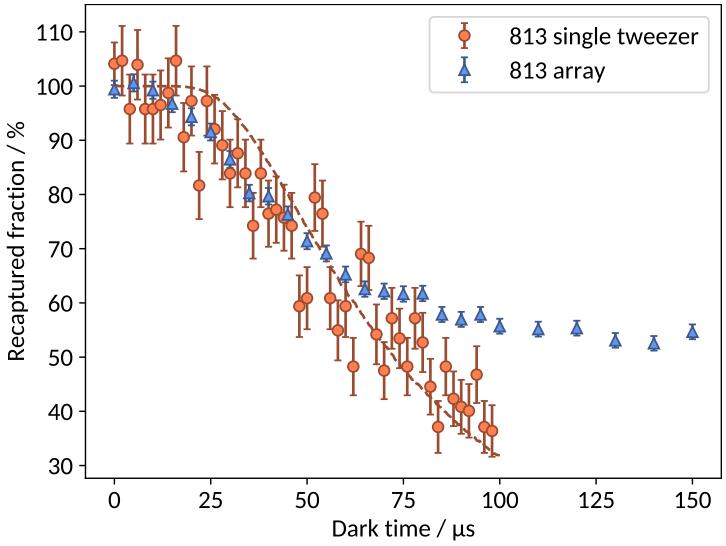

Figure 5.16 – Release and recapture measurements in 300 µK deep tweezers at 813 nm for a variable dark time at 200 µW, −2 750 kHz and 30 ms. Data is taken for a single trap (red) and a 5×5 array and shown with a best-fit Monte-Carlo simulation for 2.8 µK. The change of trend around 60 µs hints towards a recapture effect in an array, which artificially decreases the temperature compared to the single tweezer case.

of 14 µK, determined by fitting the simulation to the data. This is close to the radial ground state at 14 µK for a radial trap frequency of $\omega_r = 2\pi \times 280$ kHz, however, it can be seen that the classical simulation does not model the data accurately. As a reference, we can compare the result to [127], where the cooling is investigated in traps at 540 nm. In this letter, the authors use a chirp-cooling technique, which achieves better results than the static approach. Our result with static cooling lies between their chirp-measurements with final detunings of $\delta_f = -5.5\omega_g$ and $\delta_f = -2.4\omega_g$. The result shows that the static cooling can perform better than dynamic cooling, if suboptimal parameters are chosen for the latter.

With the correct parameters, the dynamic approach can significantly reduce the temperature approaching the limit of the ground state.

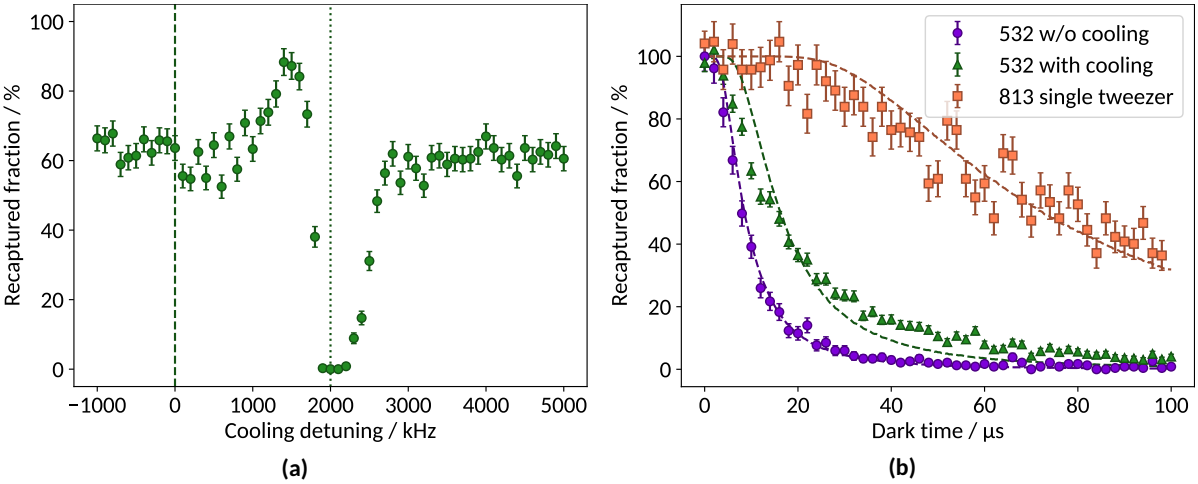

Figure 5.17 – Release and recapture thermometry in a 532 nm tweezer. **(a)** Dependence of the cooling efficiency on the cooling detuning for a duration of 30 ms at a dark-time of 6 µs. Optimal cooling happens at 1 250 kHz between the free-space resonance at 0 kHz (dashed line) and the light-shifted resonance at 2 000 kHz (dotted line). **(b)** Comparison of the recaptured fraction after sisyphus cooling in 532 nm- versus 813 nm-traps. The trap depth of 532 nm is around 2 mK compared to 300 µK at 813 nm. The dashed lines show the simulated dependence for 13 µK (33 µK) with (without) cooling at 532 nm and 2.8 µK at 813 nm.

5.5 Detection

Until now, we have been quite vague about the detection process and we will now proceed to a detailed and precise description. As described in Section 2.3, we detect an atom in a trap by collecting its fluorescence after an excitation pulse on an electron-multiplying CCD (EMCCD) camera[6]. We ensure optimal alignment of the imaging system with the procedure described in Section 5.2 (see also Appendix B) and define a region-of-interest (ROI) around a signal from the trap on the camera. For every repetition of the experiment, we determine the total signal in a ROI and investigate the distribution of the signal strength in a histogram. A ROI and a corresponding histogram are shown in Figure 5.18.

The number of photons scattered from a single atom follows a Poissonian distribution centered around a mean photon number $\langle N_{\text{photon}} \rangle > 0$. Electronically, this photon signal in converted to a certain number of electrons by the EMCCD camera with a certain gain G, again leading to a Poissonian distribution. However, even when no fluorescence light from an atom is observed, an electronic signal is measured. This background signal includes everything that does *not* stem from an atom: stray light, dark counts and also electronic noise from the chip itself. The background in units of electron counts also shows a Poissonian distribution around a certain mean larger zero. Converting the electronic signal back to photon counts is done by the camera itself, where the mean count number from the background peak is subtracted from the signal which is then divided by the gain. (see e.g. [129, 130]). This also explains why histograms of the signal, for example the one shown in Figure 5.18(c), show only half of a

[6] Andor iXon Ultra 897

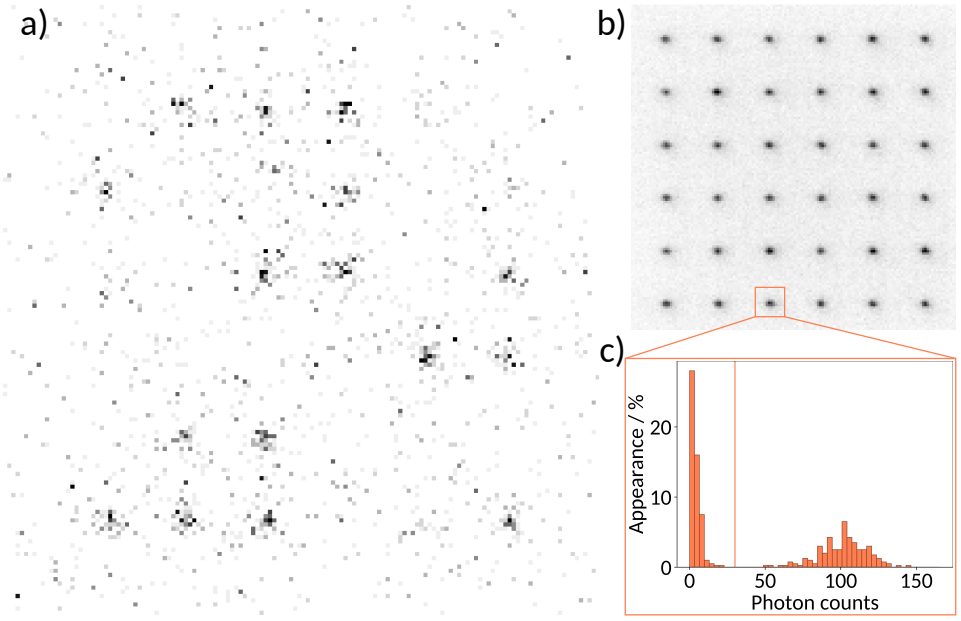

Figure 5.18 – Summary of the imaging principle. **(a)** A single shot of a 6×6 tweezer array. Due to the light assisted collisions half of the sites are occupied on average. **(b)** Superposition of many images. **(c)** The detected photon number in a small region of interest (ROI) around a tweezer follows a bimodal distribution consisting of a background- and an atom-peak.

Poissonian/Gaussian peak. It is known that the internal gain calibration of the camera might become inaccurate over time [131], however, we believe the measured photons counts to be reasonable, as we measure values as other groups [46–48].

An important quantity to describe the quality of the detection is the *fidelity*, which describes how well the atom-signal is separated from the background. The task boils down to the placement of a threshold for the detected number of photons above which we can assume with some certainty that an atom is present. As both distributions have an infinite extend, they overlap and there will be a contribution from both of either side of the threshold. To determine the fidelity, we have to determine all four elements for *binary classification*, namely true positive (TP), true negative (TN), false positive (FP) and false negative (FN) where TP + FN = 1 and FP + TN = 1. In the end, we would like to keep TP, i.e. the fraction of correctly identified atom-events, as close to unity as possible, while keeping FP as low as possible.

Different methods to set the threshold exist, and in this work we use a fit to the data to extract the threshold that separates the peaks optimally. We model the signal as the superposition of two Gaussian curves that we fit to the data[7]. For the example in Figure 5.19, we extract a mean and a width of the atom-peak of 100 photons and 14 photons respectively. The width is larger than the expectation for a Poissonian distribution of $\sqrt{100}$, and we attribute the deviation to loss during imaging which skews the distribution and results in a broader distribution. To estimate the fidelity, we place a threshold (red vertical bar in Figure 5.19) close to the minimum of the fit. The fidelity can be calculated via $\mathcal{F} = A_{\mathrm{atom},r}$, where $A_{\mathrm{atom},r}$ is the area of the fitted atom-signal to the right of the threshold. The false positives can be determined similarly from the area of the background peak to the right side of the threshold $A_{\mathrm{bg},r}$. Our definition is therefore a relative measure and especially suited for optimizing the imaging process,

[7] A Possionian distribution with a mean $\gg 1$ can be approximated by a normal distribution.

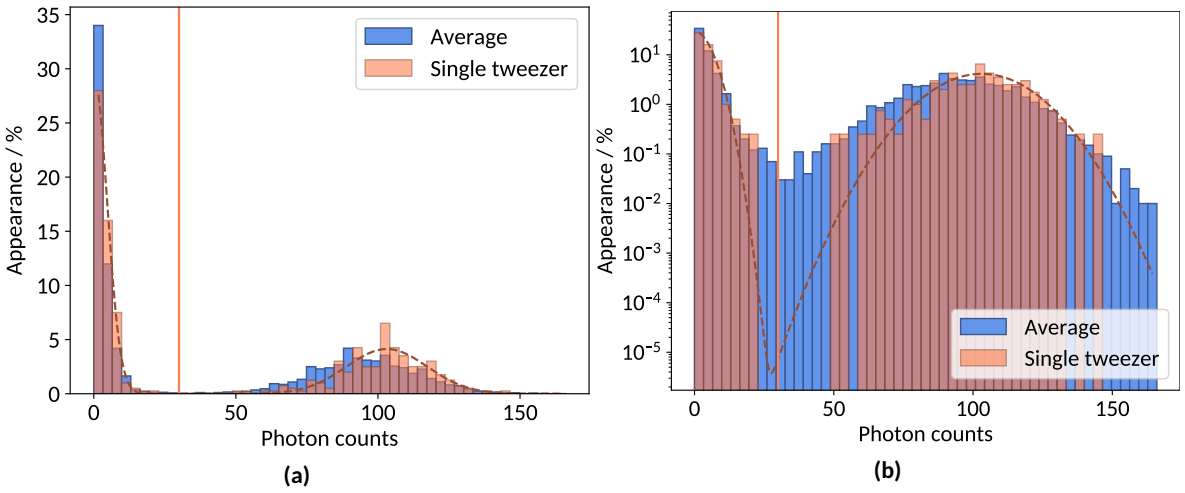

Figure 5.19 – **(a)** Histogram of the scattered photon number averaged over 36 sites (blue) compared to a single tweezer (red), with a fit to the single tweezer distribution (dashed curve). Different sites scatter a different amount of photons and the averaged distribution is therefore wider than the single tweezer distribution. Additionally, loss also leads to a broadening of the Poissonian distribution. **(b)** Logarithmic histogram that shows the deviation from the fit in the region between the peaks if atoms are lost during the imaging process.

as it determines the threshold that separates the two peaks optimally, without the need for a dedicated calibration measurement. A shortcoming of this model is the fact, that it cannot account for loss during the imaging process, which leads two a deviation from the theoretical bimodal distribution. Atoms lost during the imaging process scatter less photons which manifests itself in a *bridge* between the two peaks. In result, our method for quantifying the quality of the imaging might not be optimal for comparison to values from the literature if authors use different definitions for $\mathcal{F}$. In the future, an analysis of our data using a method from the literature will lead to a more comparable result.

5.5.1 Imaging parameters

To compensate the heating due to photon recoil during the imaging, we cool the atoms, where the details depend on the trapping wavelength and in turn the details of the cooling process. Here, we will only comment in detail on the imaging in 813 nm traps, which is relevant wavelength for this work. We briefly investigated imaging in 515 nm and 532 nm, which indicates that high fidelity imaging is possible, albeit with higher losses compared to 813 nm (similar observations have been made in the literature [47]). First, the $^1\mathrm{P}_1$-state decays to the $4\mathrm{d}^1\mathrm{D}_2$-state, which is anti-trapped at these wavelengths and the comparably long lifetime of 300 µs leads to atom loss. A second loss channel is off-resonant excitation from $^3\mathrm{P}_1 \rightarrow 5\mathrm{d}^3\mathrm{D}_2$ which also leads to atom loss.

For imaging, we use a single non-retroreflected, linearly polarized beam at 461 nm with a diameter of around 1 mm. The polarization is chosen to lie in the radial plane to ensure maximum collection efficiency due to the orientation of the emission pattern [117]. We confirm that the signal vanishes for a polarization axis parallel to the axial direction (the tweezer propagation axis) by rotation the polarization axis of the imaging beam. The total collection efficiency is around 6 %, consisting of the collection efficiency of our objective of 11 %, the quantum efficiency of 75 % of the camera and 75 % for the efficiency of the imaging path.

We continuously apply the cooling light while the imaging light is pulsed with a duty cycle of 10 to 15 % at powers of 12 to 25 µW and a cycle time of 1 ms. We find a strong dependence of the imaging loss on the duty-cycle on the per-cent-level and the imaging power on the microwatt-level. Detunings of less than −25 MHz help to fine-tune the scattering rate as they reduce the sensitivity to intensity noise of the imaging beam. To ensure that the atoms survive the imaging process, we take two images with an exposure time of 100 ms each, separated by about 120 ms, while scanning the imaging parameters.

Figure 5.20(a) shows the measurement for the imaging power at a duty cycle of 10 %. We see that the second image is necessary to visualize the loss for higher imaging powers which is not possible with a single image alone. This is to be expected, as atoms that get lost towards the end of the first imaging process still scatter enough photons to be classified as an atom-event. The signal peaks around 20 µW which corresponds to a saturation factor of 0.03 at a scattering rate of $2\pi \times 400$ kHz, assuming the trap sits in the center of the beam. This values are larger than expected from the measured number of photons and also compared to the scattering rates of less than 50 kHz in [48]. We therefore assume that the beam is not centered on the array which reduces the intensity. In fact, by reversing the argument, we can calculate the displacement from the center of a Gaussian beam with a waist of 0.5 mm. For a scattering rate of $R = 2\pi \times 20$ kHz, we compute[8] an off-center position around 0.7 mm, which is well within the error margins of our knowledge of the beam size and of the alignment precision. By displacing the imaging beam in the radial plane the position can be mapped out in the future to get a reliable estimate for the actual imaging beam intensity.

We measure a saturation of the loss at a value around 10 %, see Figure 5.20(b) and Figure 5.20(d), whereas other groups measure losses on the order of 0.1 % (see e.g.[48]). To exclude an imaging independent loss, we wait for the same time without taking the first image. In this case the loss is lower at around 5 % which can be explained by the vacuum lifetime of the traps, see Section 5.8. A likely explanation is insufficient repumping of the $^3P_2$ state during the imaging process as we do not see an influence of the repumping lasers. This could either be caused by insufficient power due to a misalignment or by off-resonant operation caused by light shifts of the $^3S_1/^3P_{2/1}$ states. We can estimate the magnitude of this effect from the branching ratio of the $^1P_1 \rightarrow 4d^1D_2$ decay and the average number of scattered photons. In 100 ms we detect an average number of 100 photons which, together with the detection efficiency of 6 %, means an atoms scatters 1 700 photons on average. Each of these excitations can lead to a decay of the $^1P_1$ state to the $4d^1D_2$ state with a probability between 1 : 20000 to 1 : 50000 (see Figure 3.1). The D-state decays to the $^3P_1$ and $^3P_2$ states with a branching ratio of 2 : 1 and therefore one-third of the the atoms end up in $^3P_2$. The probability of this happening is on the order of a few percent[9], which is a bit lower than the observed loss of 5 to 7 % but the latter is well within the error margin of this estimation.

Another explanation for the large loss due to the imaging could be due to a residual atom stream from the 2D MOT, as we currently use the push beam for imaging. However, at the time of the first image the 2D MOT has been disabled for more than 50 ms which renders the explanation unlikely. It can still be prevented by setting up an additional imaging path that does not pass through the 2D MOT region.

We also measure loss depending on the cooling detuning and amplitude for different imaging powers and cycle times. These parameters all depend on one another but we found the best result for an imaging

[8] R is computed by 120 photons/100 ms at an imaging efficiency of 6 %. From R we calculate the saturation s_0 and it turn the local intensity I_I. Finally, we compute the radius at which a Gaussian beam with a center intensity $I_0 = 2P/\pi w_0^2$ for $P = 20$ µW equals I_I.

[9] $\left(1 - (1 - 1/50000)^{1700}\right) \times 1/3 \approx 1$ % and $\left(1 - (1 - 1/20000)^{1700}\right) \times 1/3 \approx 3$ %

power of 12 µW with a duty cycle of 10 % at a cycle time of 1 ms and a detuning of −25 MHz. We do not see a dependence on the cooling power above a power of 100 µW. A measurement for the cooling detuning during imaging is shown in Figure 5.20(c) which again compares well to the literature [48].

Because the optimal cooling detuning during the imaging process depends on the trap depth in a linear manner for small deviations, it can be used to measure the relative trap depths in a tweezer array. We will exploit this fact in Section 5.6 for an efficient method to quickly homogeneize the trap depths in a tweezer array.

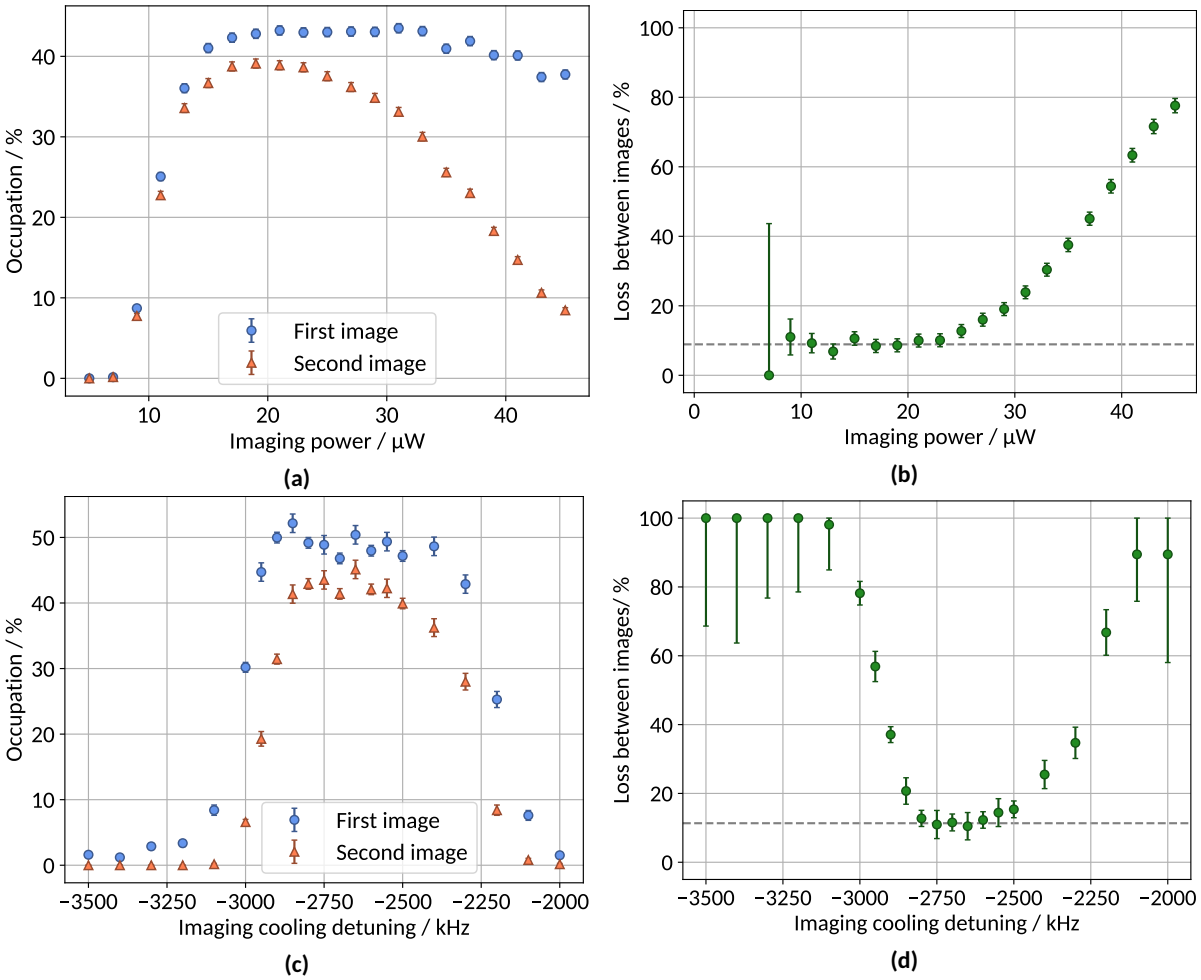

Figure 5.20 – Imaging parameters for atoms in 813 nm tweezers. The parameters are optimized for minimum loss between two images. For all measurements, the imaging light is applied with a duty cycle of 10 % at a cycle duration of 1 ms and at a detuning of −30 MHz. Cooling light is continuously applied with a power of 200 µW but similar results are obtained from 100 µW to 1 500 µW. Results are averaged over 36 tweezers. **(a)** Imaging power averaged over a 6×6-array together with the loss in **(b)**. **(c)** Imaging cooling detuning averaged over a 6×6-array together with the loss in **(d)**. The dashed lines in **(b)** and **(d)** show a minimum loss of 9 % and 11 % respectively. In general, this value depends on other losses (e.g. the vacuum lifetime) as well and slight deviations from the mean value around 10 % occur between measurements.

5.6 Tweezer array trap depth balancing

Although camera images of the tweezer array in front of the objective show non-uniformities on the few percent level, we see that this is not a good indicator for the balancing in the atom plane. The effect of the complex phase-pattern of the SLM is sensitive to misalignments, especially of the high NA objectives with respect to the vacuum windows. An indicator for a potential misalignment could be the visible astigmatism on some traps in the arrays shown in Figure 5.22. The latter does not necessarily mean that the traps projected with the upper objective suffer from the same aberrations as the imaging light is collected with the lower objective but it gives a hint that the alignment process described in Appendix B.1 might be not precise enough. We would therefore like to carry out a final balancing step of the tweezer array with a signal from the atoms that can be used in the Gerchberg-Saxton algorithm (see Section 4.3.2).

In order to do this, we reuse the phases that are calculated initially which are kept constant during the fixed-phase phase (see Section 4.3.2.2). We thus only need to determine the electric-field amplitude of each trap, which we can infer from the trap depth (or trap intensity). Different ways to measure the trap intensity exist and in this work we demonstrate resolved sideband spectroscopy in Section 5.4.1 and modulation spectroscopy by parametric excitation described in Section 5.7. However, we find another robust method that can produce accurate results with less data obtained from less complex measurements.

The non-vanishing differential light shift at 813 nm between the $^1S_0$ and the $^3P_1$ state allows an atom in a deeper trap to be excited at larger negative detunings, compared to a shallower trap (see Figure 5.13). As this shift is proportional to the trap depth, we can infer the trap depth by changing the cooling detuning during the imaging process (see Figure 5.20(c)). Atoms that are not properly cooled due to a suboptimal cooling detuning are lost during the imaging process which leads to a straightforward signature to further work with. We find this method to be more robust for large initial spreads compared to direct trap frequency measurements which rely on proper cooling to gain information about the trap depths.

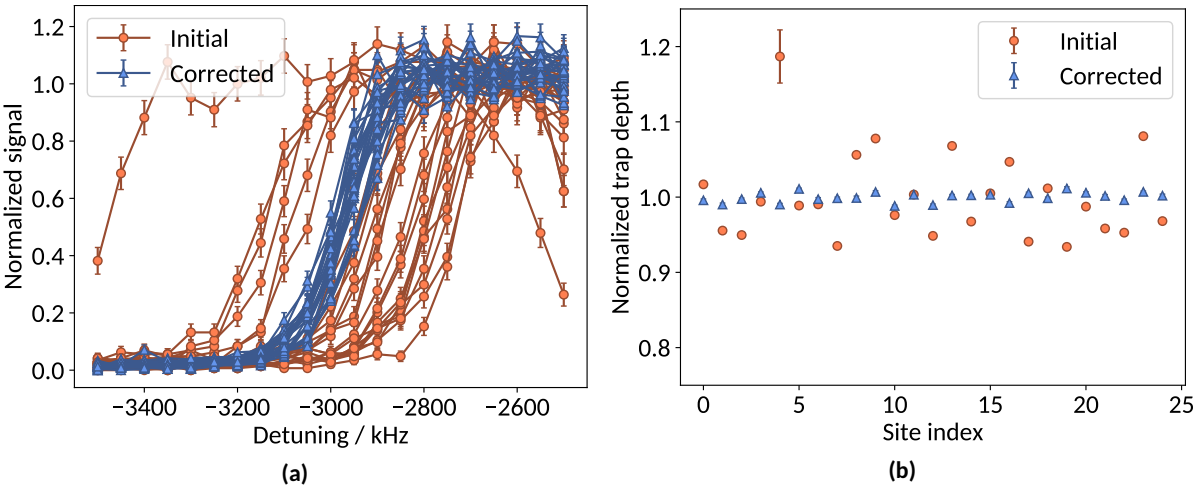

Figure 5.21 – Relative trap depths before and after correction for 25 tweezers at 813 nm. The relative peak-to-peak variation of 25.3 % and standard deviation of 6.0 % are reduced to 2.3 % and 0.7 % respectively by updating the phasemask on the SLM.

Figure 5.21(a) shows the left flanks of the cooling feature for 25 traps before and after the correction step. Fitting a sigmoid curve to the data allows us to extract the inflection point. The latter is normalized

to the mean of all values which is then taken as a relative trap depth deviation. We use the relative trap depths, which are proportional to the relative intensities, in the feedback function Equation (4.20) to optimize the phasemask written to the SLM. For the feedback we use gains of $g_0 = 1.0$ and $g_1 = 2.0$ for the first iteration and after that $g_0 = 1.0$ and $g_1 = 1.2$. Using this technique we can homogenize even large irregular arrays, as shown in Figure 5.22. We find that this leads to the fastest convergence without overshoot for small, rectangular arrays but more complicated geometries (like the ^{88}Sr array in Figure 5.22) can benefit from more iterations at larger gains. We have not investigated the convergence process in detail yet, but for small, rectangular arrays, with no prior optimization, around $4-5$ iterations are enough to reduce the peak-to-peak deviation to the final level of around 3 %. For the example in Figure 5.21, we reduce the initial spread of around 700 kHz to 70 kHz which equals a reduction of the relative peak-to-peak deviation from 25.3 % to 2.3 % and 6 % to 0.7 % for the relative standard deviation.

In addition to a balancing of the relative trap depth, it is also possible to optimize the mean trap depth of all traps at once by an additional wavefront-correcting phasemask on the SLM. Different optical aberrations, like astigmatic or spherical corrections (see Appendix A.4), can be added to the phasemask which can increase the overall trap depths significantly [42]. In the future, this method can be used to reduce the overall power demand of the system.

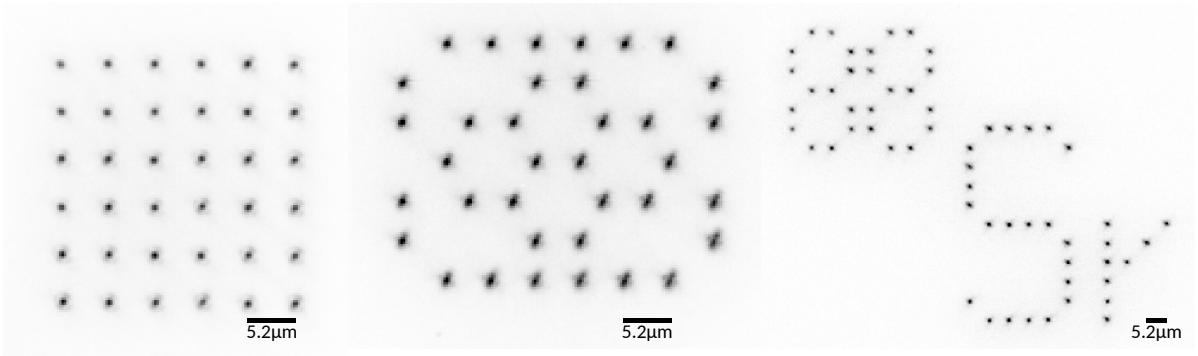

Figure 5.22 – Different tweezer array geometries. The center image shows an astigmatism caused by a slight misalignment of the imaging system at the time.

5.7 Trap characterization - Modulation Spectroscopy

A method to measure the trap frequencies (and in turn the trap size and depth), even for non-zero differential light shifts, is to parametrically heat atoms out of the trap. In contrast to an external driving force, parametric excitation is based on the modulation of the trap parameters itself. Consider the *child on a swing* example: The amplitude of the oscillation can either be increased by periodically pushing the child at a frequency ω which is close to the resonance frequency ω_0, or the child can do it themselves by changing the effective length of the swing by shifting its weight. The latter is an example of parametric excitation as a parameter of the system, in this case the mass distribution, is changed or *modulated*. For sufficiently low damping, the amplitude of the oscillation will increase to infinity, which in our case means that atoms are heated out of the trap.

Theory

In our case we will modulate the depth of the potential V periodically and we can describe the process with the following Hamiltonian:

$$\hat{H} = -\frac{\hbar}{2m}\left(\frac{\partial}{\partial x}\right)^2 + \hat{V}(x,t) \tag{5.11}$$

$$\hat{V}(x,t) = \hat{V}_0\left(1 + \epsilon \sin\left(\omega_{\text{ext}} t\right)\right) \tag{5.12}$$

where $\hat{V}_0 = \frac{1}{2} m\omega_{\text{trap}}^2 \hat{x}^2$ with the trap frequency ω_{trap}. If $\epsilon \neq 0$, the depth of the trap is periodically changed, leading to the parametric excitation. To stay in the perturbative regime, we will assume $\epsilon \ll 1$ and we can therefore model the system using time-dependent perturbation theory (see e.g. [132]). We will not go into a lengthy derivation of the solution here but rather highlight the important aspects from which we derive the expected signal. Separating the unperturbed problem H_0 and the perturbation in Equation (5.11) leads to:

$$\hat{H} = \hat{H}_0 + \epsilon \hat{V}_0 \sin\left(\omega_{\text{ext}} t\right) \tag{5.13}$$

$$= \hat{H}_0 + \frac{\epsilon}{2} m\omega_{\text{trap}}^2 \hat{x}^2 \sin\left(\omega_{\text{ext}} t\right). \tag{5.14}$$

Expressing the Hamiltonian in terms of the ladder operators $\hat{a}^\dagger$ and $\hat{a}$, leads to the well known energy ladder for the unperturbed case of $\hat{H}_0 = \hbar\omega_{\text{trap}}\left(\hat{a}^\dagger \hat{a} + \frac{1}{2}\right)$. The effect of the perturbation is visible when we compute the motional states that it couples:

$$\left\langle m \left| \hat{x}^2 \right| n \right\rangle \propto \left\langle m \left| \left(\hat{a}^\dagger + \hat{a}\right)^2 \right| n \right\rangle = \underbrace{\sqrt{n\,(n-1)}}_{M_-}\,\delta_{n-2,m} + \underbrace{(2n+1)}_{M_0}\,\delta_{n,m} + \underbrace{\sqrt{(n+2)\,(n+1)}}_{M_+}\,\delta_{n+2,m} \tag{5.15}$$

It can be seen, that the perturbation couples levels separated by $\omega_{n,n+2} = 2\omega_{\text{trap}}$ which can be explained by the even parity of the perturbation. From $M_+ > M_-$ we can further see that, for a suitably long excitation at frequency $2\omega_{\text{trap}}$, the average motional quantum number will increase. For the trapped atom in the experiment, this will ultimately lead to atom loss as it is heated out of the trap. The measured width of the resonance depends on the amplitude of the excitation and also on the anharmonic parts of the trapping potential.

Results

Figure 5.23(a) shows a measurement for a single atom in a trap at 515 nm. The first resonance at (82.1 ± 0.3) kHz corresponds to an excitation at twice the axial trap frequency and the resonance at (438.9 ± 0.7) kHz corresponds to twice the radial trap frequency.

With the trap frequencies $\omega_r = 2\pi \times (219.5 \pm 0.4)$ kHz and $\omega_a = 2\pi \times (41.0 \pm 0.2)$ kHz we can

determine the aspect ratio, the trap size and depth:

$$\frac{\omega_r}{\omega_a} = 5.34$$
$$w_0 = \frac{\omega_r}{\omega_a}\frac{\lambda}{\sqrt{2}\pi}\beta = 360\,\text{nm}$$
$$U_0 = \frac{\omega_r^2 m w_0^2}{4k_\text{B}} = 0.7\,\text{mK}$$

where $\beta\,(\xi \approx 1.5) = 0.58$ is the correction factor of the waist due to truncation at the aperture (see Section 4.2). Compared to the theoretical limit for the aspect ratio for a truncation ratio of ≈ 4.9, the measurement shows a larger value. It is also larger than the aspect ratio, determined with resolved sideband spectroscopy in Section 5.4.1. We explain this deviation by the fact, that at the time of the modulation measurement the AODs were already installed. The additional optical elements lead to additional aberrations in the system that reduce the quality of the trap, and we therefore suspect the actual waist to be larger than the one extracted from the aspect ratio. A more realistic value for the waist can be determined from the power per trap of around 3.5 mW, the ground state polarizability of 900 a u and the radial trap frequency of $2\pi \times 220$ kHz. Using these values in Equation (4.3) and Equation (4.7), we compute a waist around 460 nm and a trap depth of 1.1 mK. These values compare well to the values in [47] and show that the 515 nm tweezer beam path, including the objective in relation to the vacuum windows, is well aligned.

We repeat the modulation measurement for a tweezer array at 813 nm and show the result for the radial direction in Figure 5.23(b). We could not resolve the axial sideband yet and can therefore not infer the aspect ratio directly. Although the wavefront is corrected to around $\lambda/6$ peak-to-peak deviation before the objective, we suspect stronger residual aberrations from misalignment with respect to the vacuum window. This can be seen for the single trap in Figure 5.23(b), where we measure a double resonance with frequencies of $\omega_{r_1} = 2\pi \times (47 \pm 1)$ kHz and $\omega_{r_2} = 2\pi \times (60 \pm 1)$ kHz. The underlying aberration is probably a combination of an astigmatism and a coma caused by a tilt of the objective relative to the window, in addition to a residual spherical aberration caused by a deviation from the window thickness from the optimal value. A simulation to estimate the angular misalignment for a correction in the experiment is difficult to carry out due to the large number of free parameters.

The existence of aberrations in the system and no knowledge about the axial trap frequency requires us to estimate the trap parameters from additional quantities. As described in Section 5.4.2 for $\alpha_g < \alpha_e$, we compute the waist from the power per trap of 9 mW, the trap frequency of $2\pi \times 60$ kHz and a ground state polarizability of 286 a u (see Table 4.1). The waist of 950 nm and depth of 300 µK confirm the suspicion that the traps are strongly aberrated. It is not clear, if this can be corrected with the current setup, as our alignment method is apparently not precise enough. The cleanest solution would be to replace both objectives with a single objective which shows a negligible focal shift between the different tweezer wavelengths used in the experiment. This would also prevent the need of frequent realignment caused by a drift in the upper objective stage.

From the data for every site, we calculate a relative standard deviation, normalized to the mean, of 4 %. Note that this result has to be taken with care, as its accuracy is limited by the low resolution of 5 kHz. We expect lower deviations and a better agreement with other methods when repeating the measurement with a frequency resolution < 1 kHz.

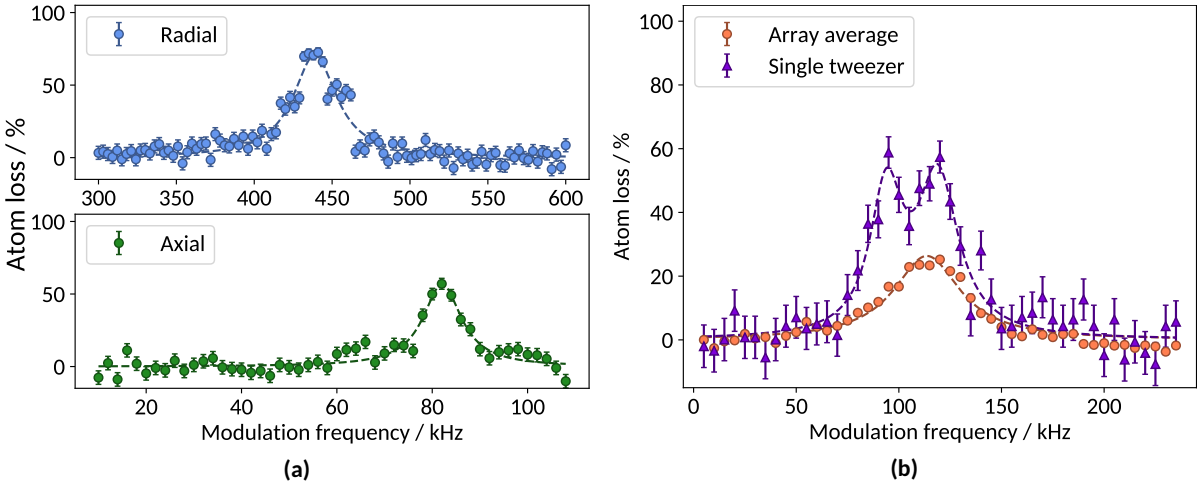

Figure 5.23 – Modulation spectra together with fits to extract the trap frequencies. **(a)** Single tweezer at 515 nm. From the ratio of trap frequencies, the quality of the trap can be inferred. **(b)** Averaged spectrum for 36 tweezers at 813 nm, together with a spectrum for a single representative tweezer in the array.

5.8 Lifetime

Current quantum simulators based on neutral atoms in optical tweezers are assembled in an atom-by-atom manner [39, 41, 42, 133] from an initially disordered state (see Chapter 6). An important factor for the assembly of larger arrays in this *bottom-up*-approach is the lifetime of the traps. If an atom in a trap has a lifetime of τ, the total lifetime of an array of N atoms is given by τ/N. Different mechanism limit the lifetime τ, for example heating due to intensity-noise of the trap but even in the absence of these heating effects the lifetime is ultimately limited by collisions with residual background atoms. At room temperature, pressures of 1×10^{-11} mBar can be reached [70] which can lead to vacuum lifetimes of several minutes. The collision rate can further by reduced if the vacuum system itself is cooled to cryogenic temperatures [134]. For strontium in 813 nm tweezers these minute-long lifetimes at room temperature have been demonstrated [48], which serves as a benchmark for our system. To measure the vacuum lifetime, we vary the delay between the two cooling blocks in Figure 5.1 before continuing with the standard imaging sequence. Figure 5.24 shows the results for atoms in 813 nm traps, but 515 nm and 532 nm traps show similar results. In this measurement, the extracted lifetime is (3.56 ± 0.15) s, but repeating the measurement over several months shows lifetimes between 2 to 10 s. Before discussing the absolute value of the lifetime, we will briefly comment on the differences with and without cooling. We see that cooling, pulsed or continuously applied, does not influence the lifetime in 813 nm traps. It does, however, reduce the lifetime in the 515 nm traps significantly to around 600 ms. We explain this by the fact that 515 nm trap photons can lead to off-resonant excitation to from the $^3P_1$ to the $5d^3D_{1,2}$-state [67]. We have yet to investigate the influence of the cooling light in 532 nm traps, but we expect similar results as for 515 nm, as both wavelengths share the same loss channels.

The absolute value of the lifetime is far from the values reached in the literature of several minutes. First, we exclude parametric heating from the laser by investigating the frequency components with and without intensity stabilization and by repeating the measurement. We measure the frequency spectrum with a spectrum analyzer and find no components above -95 dBm between 0 to 300 kHz even without an active laser stabilization. We also exclude intensity noise (both in amplitude and spatial position)

from the SLM as we observe similar results with the AODs and even without a modulating device at all.

A simulation of the vacuum system using *Molflow*[10] suggests that the pressure at the position of the atoms might be one to two orders of magnitude higher than the pump reading. This is caused by the complex geometry of the chamber with channels that can trap background gas consistents for a long time before they reach the pump. The CF16 connection between the chamber and the rest of system also acts as a differential pumping system which increases the pressure inside the main chamber. We confirm this, by measuring the lifetime after using the titanium-sublimator pump which increased the pressure reading to more than 3×10^{-10} mBar without a change in the lifetime.

Finally, a faulty weld that had to be resealed several times over the last few years could lead to a virtual leak that releases a particle-stream of varying intensity directed at the atoms. A strong indicator for this hypothesis is the fact, that we measure large deviations when repeating the experiment on the order of weeks, consistent with a non-constant particle-stream from a virtual leak [70, 135]. To investigate the influence of a virtual leak, it could help to heat the system locally around the suspected position of the leak by around 10 to 15 °C while measuring the lifetime. The increased temperature should have an influence on the intensity of the particle-stream which should result in a change of the lifetime.

All in all, the problem itself does not seem to be fixable without major changes to the vacuum system, like replacing the vacuum chamber or installing an additional pump to one of the CF16 ports of the main chamber. For the goal of this work, the assembly of small defect-free arrays, the current lifetime has a value that complicates the measurement and evaluation process, but it is long enough to obtain a reliable result.

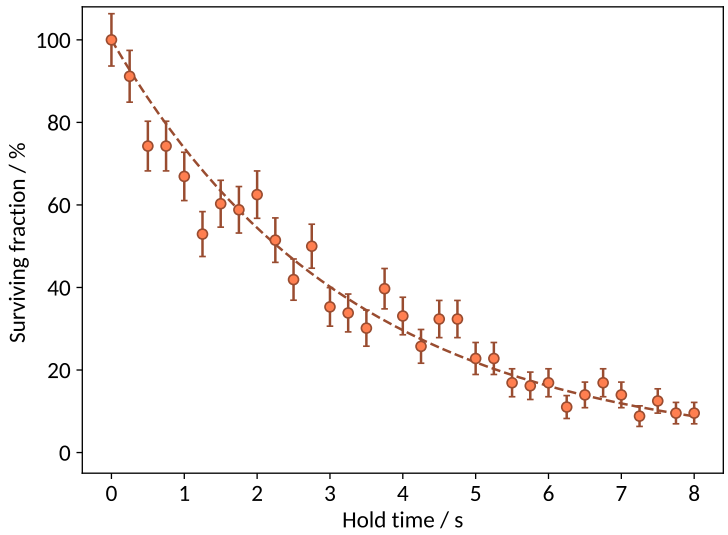

Figure 5.24 – Lifetime measurement of an atom in a tweezer array at 813 nm with an extracted value of (3.56 ± 0.15) s. Similar values are obtained for 515 nm and 532 nm tweezers.

[10] https://molflow.web.cern.ch/

5.9 Summary and open questions

In this chapter, we investigated the preparation, cooling and detection of single ultracold atoms in optical tweezers. We discussed the process of light-assisted collisions that we use to create sub-Poissonian atom number distributions in the tweezers, limiting the atom number to either 0 or 1 atom. We investigated three different cooling protocols for the different wavelengths used in the experiment. For traps at 515 nm, we showed that we can cool atoms to the three-dimensional motional ground state with a probability of 95^{+2}_{-10} %. We also showed that we can cool atoms in tweezers at 813 nm to the level that has been observed in the literature. Finally, we briefly investigated the cooling in traps at 532 nm, where a significant cooling effect can be seen even with a suboptimal protocol.

We showed that we can take high-fidelity images of atoms in 813 nm traps. The high loss of around 7 % during the imaging process is likely caused by insufficient repumping, either due to low power due to misalignment or by off-resonant operation caused by light-shifts. This can be verified by adding acousto-optic modulators to tune the frequency of the repumping lasers and also by displacing the beams in the radial plane to identify a potential misalignment. To compare the absolute values for the imaging fidelity with values from the literature, a different model which accounts for the loss has to be developed. For example multiple images can be used to determine the number of events where an atom is only detected in the second image [47]. This can be done with the current data without additional measurements. In addition, a dedicated imaging beam to prevent a residual atom flow from the 2D MOT region is needed in the future. However, as the error induced by the imaging loss dominates the uncertainty in the detection fidelity, we neglect the latter in the discussion for the rearrangement results in Chapter 6.

A second major issue is the vacuum lifetime. It seems, that this can only be fixed by making major changes to the setup by e.g. adding a small additional pump or by replacing the vacuum chamber all together.

Irrespective of the technical improvements, a more fundamental open question is how well the limit of the sisyphus cooling protocol represents the quantum mechanical motional ground state. In the future, sideband thermometry can be used to investigate this in greater detail. The ability to directly infer the average motional quantum number $\langle n \rangle$ is a robust absolute measure, independent of a possibly inaccurate simulation. The release-and-recapture method can still be used to optimize the cooling parameters in a relative manner as much less data is needed and progress can be made faster. It could therefore be beneficial to cool the atoms in a trap at 532 nm or 813 nm before transferring them to 515 nm where resolved sideband spectroscopy is possible. The details and the influence of this transfer-process can be characterized by a double-handover: First, sideband-cooling at 515 nm ensures a well known initial state. Secondly, a transfer to a different trapping wavelength and back again, followed by sideband thermometry can be used to characterize the transfer process.

Cooling an atom to the motional ground minimizes the uncertainty associated with the initial internal state of the system, and we also reduced the uncertainty associated with the probabilistic distribution by reducing the atom number to a single atom at most. However, this still leaves around 2^n initial configurations for n traps, and the final chapter will now discuss how this uncertainty can actually be *minimized* by rearranging atoms to a defect-free configuration.

CHAPTER 6

Assembling defect-free atom arrays

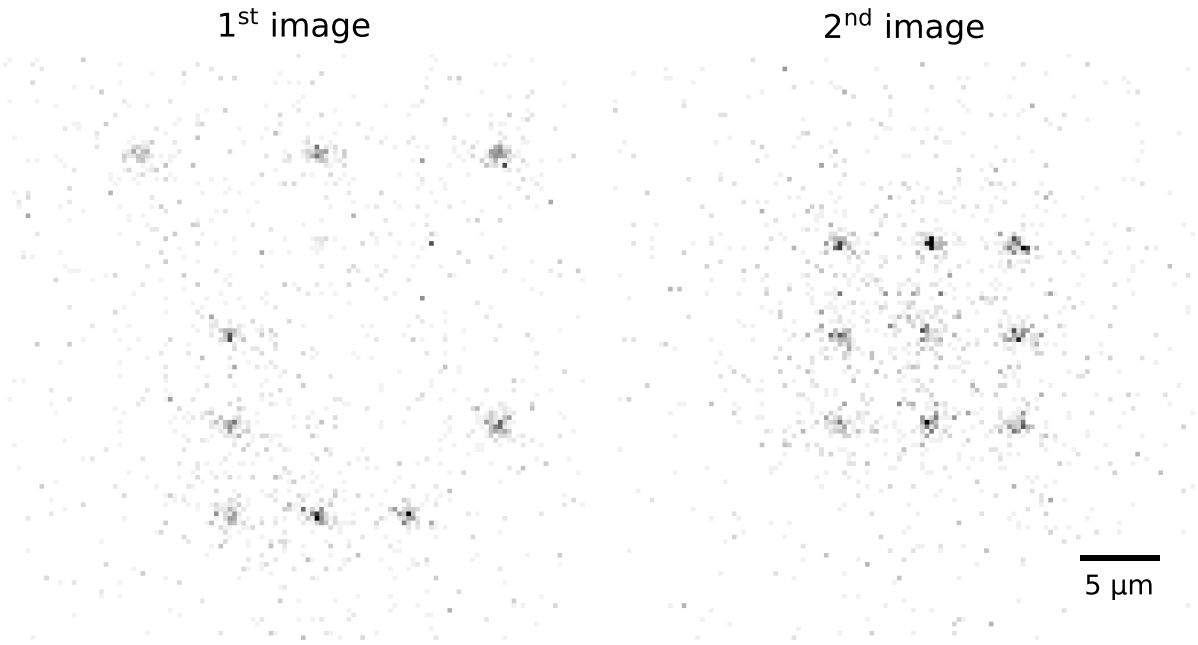

Figure 6.1 – Nine atoms randomly distributed in an underlying 5×5 square array have been rearranged into a defect-free 3×3 square array.

In this last chapter, we will now present the single site addressing capability of the experiment by rearranging partially filled arrays to defect-free configurations. Finite filling fractions are a result of the probabilistic nature of the tweezer loading process which inherently prevents a deterministic preparation of the system. To reduce the uncertainty (or entropy) associated with the binomial distribution of initially populated tweezers (see Section 5.3), the currently established approach is to rearranged atoms between sites to create a defect-free configuration in an atom-by-atom manner to a smaller *target*-region [39, 41, 42, 133]. The assembly is carried out by a dynamic optical tweezer, which moves atoms between different sites of an underlying tweezer array. For this work, we combine two different trapping wavelengths to make use of their respective advantages. The static arrays in which atoms are rearranged are created

at 813 nm, as the wavelength allows for low loss imaging and long lifetimes. For the single dynamic tweezer we use light at 515 nm or 532 nm, as these wavelengths allow us to create deep potentials due to the small waist and large polarizability (see Section 4.1).

We first investigate how to move a single atom with a dynamic tweezer, created with the pair of acousto-optic deflectors (AODs) characterized in Section 4.4. We develop a model based on the kinematic parameters of the motion and construct a control-theory optimal trajectory. Afterwards we discuss the algorithm that is used to rearrange a partially-filled array into a smaller *target*-region of full density, making use of additional *reservoir*-sites. The process has been demonstrated in arbitrary geometries [136], in optical lattices [137] and even in complex large-scale arrangements where target regions are separated for different tasks [133]. We benchmark the performance of the algorithm using a simulation and discuss two modifications that increase the probability to find a solution which minimizes the number of moved atoms. We then present an overview of the combination of all hard- and software components and the experimental parameters that we work at. Finally, we discuss the experimental results where we estimate the single move success probability (also called move *fidelity*) with two different methods to reduce the influence of experimental uncertainties. From the single move fidelity we determine the probability of preparing a defect-free array. We close the chapter by discussing the current bottlenecks and possible improvements.

6.1 Moving a single atom

To create a trap in the focal- or atom-plane, RF-waveforms with frequencies f_x and f_y are applied to the AODs. The position of the trap depends linearly on the angle of the deflected light which itself depends linearly on the RF-frequency (see Figure 4.22 in Section 4.4). We can therefore correlate changes of the position to changes in the frequency as shown in Table 6.1. We start the discussion with the simplest form of motion connecting two positions, namely a linear change in position over time. Afterwards, we identify its flaws and propose a more sophisticated model. For simplicity, we discuss the problem in one dimension first and extend it to the two-dimensional case at the end.

$f(t)$	$\Leftrightarrow$ Position $r(t)$
$f'(t)$	$\Leftrightarrow$ Velocity $v(t)$
$f''(t)$	$\Leftrightarrow$ Acceleration $a(t)$
$f'''(t)$	$\Leftrightarrow$ Jerk $j(t)$
$f''''(t)$	$\Leftrightarrow$ Snap $s(t)$

Table 6.1 – Frequency space correspondence

6.1.1 Constructing an *optimal* trajectory

The aforementioned linear motion corresponds to a linear frequency ramp, or frequency *chirp*, that we apply to the modulator. We can describe the instantaneous frequency $f(t)$ of the chirp by:

$$f(t) = f_0 + (f_1 - f_0)\frac{t}{T} = f_0 + \Delta f \frac{t}{T}, \tag{6.1}$$

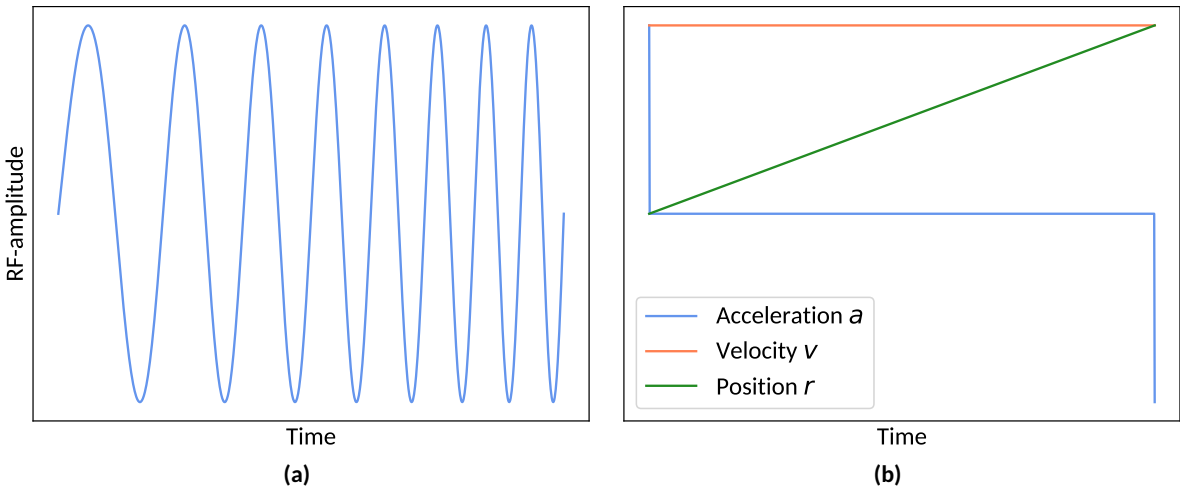

Figure 6.2 – A linear frequency chirp with its waveform **(a)** and the resulting kinematic motion parameters **(b)**. The linear change in position leads to discontinuities in the acceleration which can lead to excess heating or even loss of the atom due to high-frequency components.

where T is the time over which the frequency changes by Δf from f_0 to f_1. We obtain the full waveform $S(t)$ from the phase $\phi(t)$ by integrating out the frequency:

$$\phi(t) = 2\pi \int_0^t f(t')\mathrm{d}t' = 2\pi \left(f_0 t + \frac{1}{2}\frac{\Delta f}{T} t^2 \right) \tag{6.2}$$

$$S(t) = A(t, f(t)) \sin\left(\phi(t)\right) \tag{6.3}$$

where $A(t, f(t))$ is the instantaneous amplitude, which in general depends on the time t and the instantaneous frequency $f(t)$. $S(t)$ is the resulting waveform that will be written to the arbitrary waveform generator (AWG) which drives the AODs (see Figure 6.2(a)). To analyze the effect of the waveform on the atoms, we can investigate the time-dependence of the spatial position and its time-derivatives. As shown in Fig. 6.2(b), the position r varies linearly in time and the velocity v is therefore constant. A problem arises at the start and end of the chirp, where the constant velocity leads to a jump in acceleration a at the very first and last timestep. This impulse carries high-frequency components, that can lead to unwanted excitation of the external (motional) state, resulting in heating or atom loss in the worst case. The precise effect on an atom is difficult to predict, as it depends on the bandwidth of the AWG, the amplifiers and the response-time of the AODs which have to be determined experimentally. To mitigate the effect of instantaneous changes of a, we will therefore design the frequency chirp starting from the jerk j, while respecting maximum values for j, a and v. We already saw that a finite j_{max} is required to avoid high-frequency impulses in the acceleration (and therefore the force) acting on the atom. Finite maximum values for a and v are also necessary, as both can lead to a deformation of the trapping potential. More specifically, an acceleration a leads to a tilt of the tweezer potential which, for a sufficiently large value of a, can lead to atom loss. A large velocity v can also lead to a deformation of the trap, as the wave inside the AOD takes a finite time to cross the clear aperture due to a finite speed of sound which can therefore lead to different deflection angles and in turn an elliptical shape of the trap. Both effects are visualized in Figure 6.3.

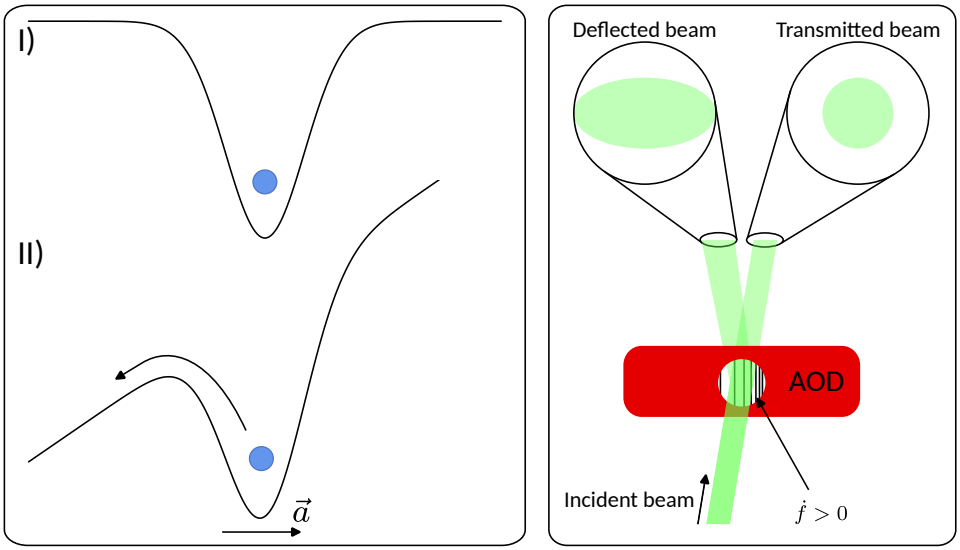

Figure 6.3 – Visualization of trap deformation due to excessive acceleration or velocity. **Left:** Applying an acceleration a leads to an additional linear potential energy slope. If this slope becomes too large, the atom can leave the trap. At the same time, changes in acceleration (jerk j) that are fast compared to the "reaction time" of the wavefunction, given by the inverse trap frequency, can lead to unwanted excitation resulting in heating or atom loss. **Right:** The acoustic wave inside the AOD takes a finite time to travel through the clear aperture. If the change in frequency (and therefore the velocity) is too large, different parts of the beam deflect with different angles, leading to a deformation of the trap. The effect also leads to an axial deformation (not shown) which can lead to out-of-plane trapping or atom loss.

6.1.1.1 Mechanical model - a control-theoretic approach

We can express the problem in the language of control theory in one dimension, where $\vec{x}(t)$ describes the state of the system and $\vec{u}(t)$ describes the control input. The time-evolution is given by [138]:

$$\dot{\vec{x}}(t) = f(\vec{x}(t), \vec{u}(t), t), \tag{6.4}$$

where f is a vector-valued function encoding the dynamics of the system. Control-theory is concerned with the task to find an *optimal* solution to problem where optimal is defined as the optimization of a cost functional[1] J. The task is to find a control input $\vec{u}^*(t)$ that optimizes J which is usually written as:

$$J = h(\vec{x}(t_f), t_f) + \int_{t_0}^{t_f} g(\vec{x}(t), \vec{u}(t), t)\mathrm{d}t \tag{6.5}$$

where h and g are scalar functions, again specific to the system. The form of the cost functional is deeply rooted in the mathematical theory of optimization. Expressed in a language familiar to a physicist, g is also written as L and named the *Lagrangian* of the system, which makes the resemblance of J to the action S encountered in physical problems more obvious.

[1] Sometimes also referred to as *performance measure*.

We define the state $\vec{x}(t)$ and the control input $\vec{u}(t)$ by:

$$\vec{x}(t) = \begin{pmatrix} r(t) \\ v(t) \\ a(t) \\ j(t) \end{pmatrix} \quad \text{and} \quad \vec{u}(t) = \begin{pmatrix} 0 \\ 0 \\ 0 \\ j(t) \end{pmatrix} \tag{6.6}$$

where we truncate the motion at the jerk $j(t)$. We justify this model as a reasonable compromise based on our current knowledge, as we do not know how delicate the process of displacing a quantum-state actually is. The literature [39, 136] suggests that linear chirps, with discontinuities in the acceleration, can work reliably for moving an atom. Other groups use higher-order chirps [139], and recently also other control-theory optimal trajectories have been investigated [140]. One of our future goals is to develop a deeper understanding of the connection between the shape of the trajectory and its influence on a quantum state. We therefore choose an approach that is more refined than the simple linear chirp but at the same time not overly complex before it is proven to be necessary. It also turns out that our approach can be reduced analytically to multiple one-dimensional optimization problems, which simplifies the implementation. As explained above, we impose maximum values for j, a and v:

$$0 \leq v(t) \leq v_{\max} \qquad \text{with } v_{\max} > 0 \tag{6.7}$$

$$-a_{\max} \leq a(t) \leq a_{\max} \qquad \text{with } a_{\max} > 0 \tag{6.8}$$

$$-j_{\max} \leq j(t) \leq j_{\max} \qquad \text{with } j_{\max} > 0, \tag{6.9}$$

which force finite durations for changes of these parameters. $v_{\max}$, $a_{\max}$ and $j_{\max}$ are free parameters in our model that have to be determined experimentally. We will comment on the values used for this work in Section 6.3.3. In addition to the limits, we have to satisfy the boundary conditions:

$$\vec{x}(t_0 = 0) = (r_1, 0, 0, 0) \tag{6.10}$$

$$\vec{x}(t_f = T) = (r_2, 0, 0, 0) \tag{6.11}$$

which state that the atom should be fully at rest at the start and the end of the trajectory.

Solution

Finding an optimal trajectory requires us to determine the duration T and the time dependence of $\vec{u}(t)$, such that an atom is moved by the required distance $d = r(T) - r(0)$. We include higher-order position derivatives by assuming a certain functional dependence for changes of the jerk, which results in continuity to fourth order of the motion, also referred to as *snap* $s(t)$. Although less frequently encountered in the canonical mechanics courses for physicists, the *snap* is a crucial parameter to consider in other disciplines. An example is the field of robotics, where it is used to compute drone trajectories [141]. This could be an interesting field to benefit from by utilizing the developed models and mathematical techniques to design even smoother atom trajectories.

We optimize the control input $\vec{u}^*(t)$ for *minimum time* T. With $h = 0$ and $g = 1$ in Equation (6.5) we have:

$$J = \int_0^T \mathrm{d}t = T. \tag{6.12}$$

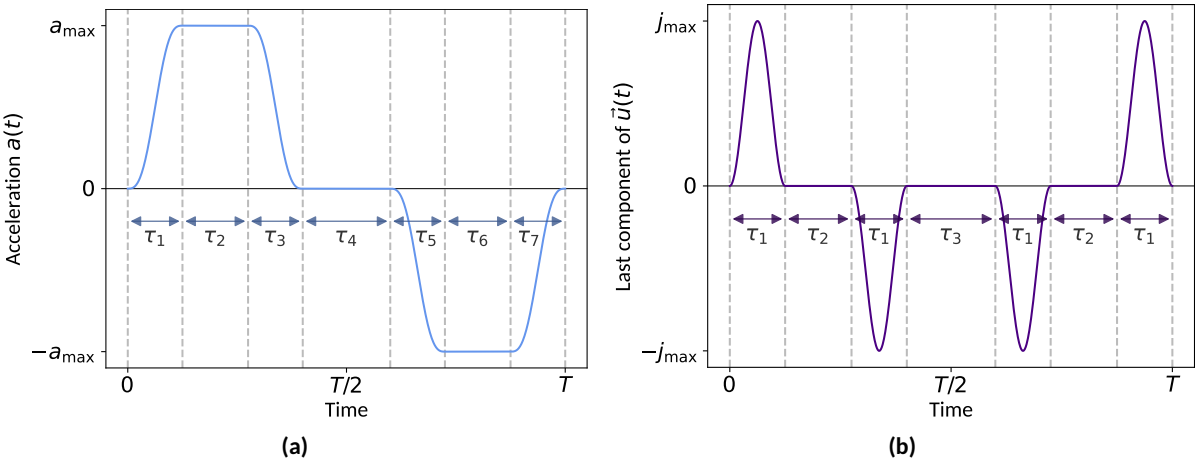

Figure 6.4 – **(a)** Time-series of the acceleration, split into seven segments. **(b)** Control input required to produce the acceleration from **(a)**. By making use of the symmetries of the problem, the complexity can be reduced to three one-dimensional optimization problems to determine the three durations τ_1, τ_2 and τ_3.

It is important to stress that this is merely one of many possible choices. In the future, the measure could be defined as the minimum motional excitation or lowest decoherence based on a fully quantum-mechanical model.

Instead of trying to solve the problem directly, we will simplify it further by making use of the symmetries of the motion. Figure 6.4(a) shows the most general form the acceleration can take, where we assume an equal magnitude for acceleration and deceleration and an equal maximum or minimum change of acceleration (jerk (j)). Of the seven segments, segments 0, 1 and 2 combined describe the process of acceleration to a certain velocity v'_{max}. This velocity is smaller or equal to the maximum velocity v_{max}, which imposes the absolute limit for the velocity. In segment 3 the velocity is constant for τ_3 until we decelerate again in segments 4, 5 and 6. $j(t)$ in the zeroth segment is by:

$$j_0(t, \tau_1, j_{max}) = j_{max} \sin^2\left(\pi \frac{t}{\tau_1}\right) \tag{6.13}$$

where $j_{max} = 2j_{avg}$ and the other components of $\vec{x}(t)$ are obtained by a straightforward analytical integration. $j_0(t, \tau_1, j_{max})$ is chosen such that the motion is continuous up to fourth order in time. τ_1 is the first of three characteristic timescales of the problem. It describes the time it takes to reach maximum acceleration and is given by the ratio a'_{max}/j_{avg}. Depending on the values for j_{max}, a_{max} and v_{max}, it can happen that $v_{max}/2$ is reached before a_{max} is reached. The latter happens after:

$$v_{max}/2 \stackrel{!}{=} v_0(\tau', \tau', j_{max}\tau') \tag{6.14}$$

$$= \frac{1}{2} j_{max} \tau'^2 \tag{6.15}$$

$$\Rightarrow \tau' = \sqrt{\frac{v_{max}}{j_{max}}}, \tag{6.16}$$

where v_0 is obtained by integrating j_0 twice. This statement is making use of one of the symmetries mentioned before: because of equal magnitudes for maximal and minimal changes in acceleration, it

takes the same time to go from $v = 0$ to $v = v_{\max}/2$ as it does from $v = v_{\max}/2$ to $v = v_{\max}$. If we find $\tau' < \tau_1$, the maximum acceleration has to be scaled to $a_{\max} = \sqrt{v_{\max} j_{\max}}$, to have enough time left to decelerate in segment 2, without exceeding $v_{\max}$ at the beginning of segment 3.

Assuming additionally equal magnitudes for the maximum acceleration and deceleration, the duration of segments $\{0, 2, 4, 6\}$ have to be equal and it follows:

$$j_0 = -j_2 = j_4 = -j_6 \tag{6.17}$$

where j_i describes the functional dependence of the jerk in segment i. In segment 1, the constant maximum acceleration is reached, thus $j_1 = 0$. To satisfy the boundary condition from Equation (6.11), the acceleration and velocity have to be symmetric about half of the total duration T[2]. This symmetry argument can be used to show that the durations for segments 1 and 5 are both equal, which we define as τ_2, and thus:

$$j_1 = j_5 = 0. \tag{6.18}$$

The last free parameter is the duration τ_3 of segment 3 where $j_3 = a_3 = 0$ and $v_3 = const$. We now have a piecewise definition of $\vec{u}^*(t)$, shown in Figure 6.4(b), and the next step is to find the durations $\{\tau_1, \tau_2, \tau_3\}$, such that the covered distance d is equal to the required distance d_{target}, and that it is covered in the minimum amount of time. This three-dimensional optimization problem can further be broken down into three one-dimensional problems, which can then be solved independently. The link between v, a and j also connects τ_1, τ_2 and τ_3, leading to three different cases.

Case 1 - Distance limited The first case is characterized by the fact that maximum acceleration or velocity cannot be reached, as the distance is too short. This results in $\tau_2 = \tau_3 = 0$. τ_1 must be varied until the covered distance matches the required distance.

Case 2 - Acceleration limited The second case only exists if $a_{\max} < \sqrt{v_{\max} j_{\text{avg}}}$, otherwise $a_{\max}$ is scaled accordingly. If this holds true, maximum acceleration is reached and there is time to accelerate for $\tau_2/2$ until $v_{\max}$ is reached. If the distance covered is larger than the desired distance, τ_2 is lowered, such that the difference vanishes. If it is smaller, we move to case 3.

Case 3 - Velocity limited The third and last case is reached for long distances where also a constant velocity $v_{\max}$ is reached for a finite time τ_3. After $\tau_3/2$ the motion is reversed.

Case	τ_1	τ_2	τ_3	$\lvert a \rvert$	v	T
Distance-limited	$< \tau_{1,\max}$	0	0	$< a_{\max}$	$< v_{\max}$	$4\tau_1$
Acceleration-limited	$\tau_{1,\max}$	$< \tau_{2,\max}$	0	$a_{\max}$	$< v_{\max}$	$4\tau_1 + 2\tau_2$
Velocity-limited	$\tau_{1,\max}$	$\tau_{2,\max}$	≥ 0	$a_{\max}$	$v_{\max}$	$4\tau_1 + 2\tau_2 + \tau_3$

Table 6.2 – Summary of the trajectory cases

[2] The acceleration is an odd function about the midpoint $T/2$, while the velocity is even.

After the correct case is identified, the respective duration has to be adjusted, such that the correct distance is covered. This can be done analytically or numerically with a one-dimensional optimization algorithm. We chose to solve the problem numerically to ensure greater expandability, as a more complex approach might not be analytically solvable anymore. The parameters for the three cases are summarized in Table 6.2 and the shape of the trajectory, together with the velocity and acceleration are shown in Figure 6.5

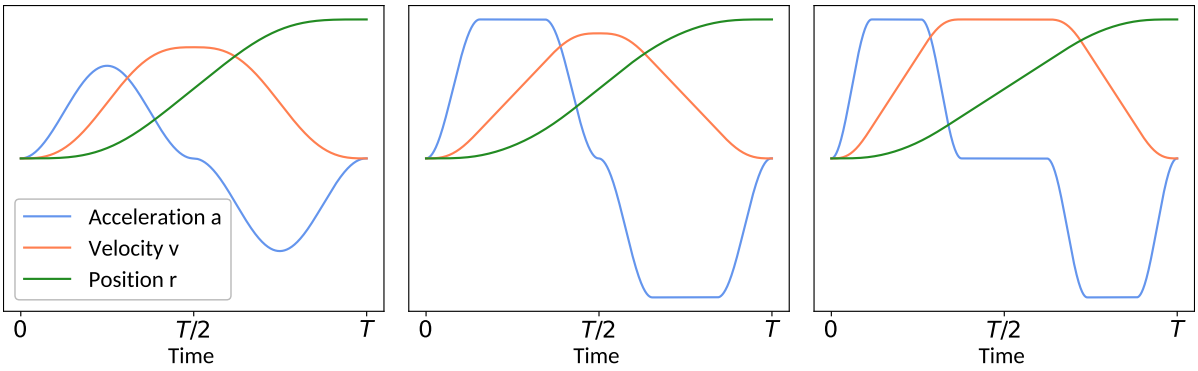

Figure 6.5 – Resulting trajectories for three different cases. **left:** Maximum acceleration and velocity are not reached. **center:** Maximum acceleration is reached but not the maximum velocity. **right:** Maximum acceleration and velocity are reached. Full explanation in main text.

6.1.2 Waveform

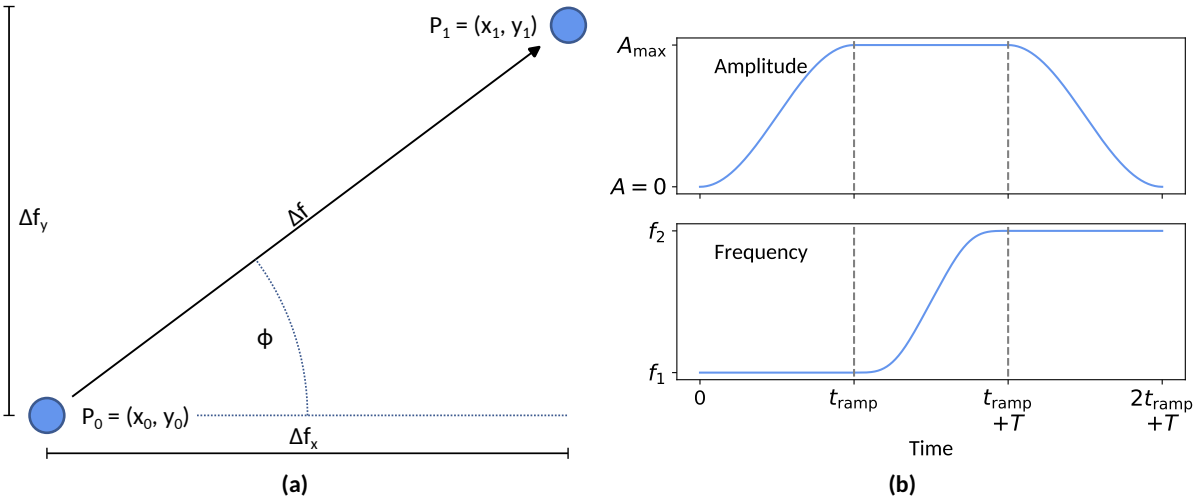

Figure 6.6 – **(a)** The precalculated trajectory is split among the two axes to reach any point in the atom plane. **(b)** Each chirp starts and ends with a $\sin^2$-shaped amplitude waveform to pick up the atom. During the frequency chirp, the amplitude is approximately constant.

With the problem being solved theoretically, the trajectory now has to be converted to a waveform, which can be applied to the AODs. To compute the two-dimensional waveform, we do *not* apply the calculated trajectories to both axes, as this would violate the maximum parameter constraints. Instead,

we use it as a magnitude in polar coordinates and split it among the axes accordingly. Assuming we move from $P_0 = (x_0, y_0)$ to $P_1 = (x_1, y_1)$, we calculate the trajectory for the *distance* between the two points $d := ||P_1 - P_0|| = \sqrt{(x_0 - x_1)^2 + (y_0 - y_1)^2}$ (see Figure 6.6(a)). To map the one-dimensional motion to the distance change *per axis*, we use the angle[3] $\phi = \arctan(y_1 - y_0, x_1 - x_0)$ to compute:

$$d_x(t) = d(t)\cos\phi \tag{6.19}$$
$$d_y(t) = d(t)\sin\phi. \tag{6.20}$$

The final step is to integrate out the trajectory, that we determined earlier, using:

$$f_{x/y}(t) = f_{x/y,0} + \alpha \cdot d_{x/y}(t) \tag{6.21}$$
$$\phi_{x,y}(t) = 2\pi \int_0^t f_{x/y}(t')\mathrm{d}t' \tag{6.22}$$
$$= 2\pi f_{x/y,0} t + 2\pi\alpha \int_0^t d_{x/y}(t')\mathrm{d}t' \tag{6.23}$$
$$= 2\pi f_{x/y,0} t + 2\pi\alpha_{x/y}\chi_{x/y}(t), \tag{6.24}$$

where $\alpha_{x/y}$ is a conversion factor given by the optical system. In our case $\alpha_x = \alpha_y \overset{!}{=} \alpha \approx 1.25\,\mu\mathrm{m/MHz}$ (see Section 4.4). $\chi_{x,y}(t)$ is a "spatial phase", defined by the integral of the position $d_i(t)$. χ and ϕ share the same proportionality constant α as $f(t)$ and $d(t)$. $\chi_{x,y}(t)$ has no immediate interpretation in the mechanical picture and is merely a tool needed to compute the waveform. The discretized AWG signal at index i is calculated by:

$$t = \delta t \cdot i \tag{6.25}$$
$$A(t) = A_{\max}\xi(f(t)) \tag{6.26}$$
$$V(t) = A(t)\sin(\phi(t)), \tag{6.27}$$

where $\delta t = 1/f_S$ with the sampling frequency f_S and $A_{\max}$ is the maximum amplitude. $\xi(f(t))$ is the amplitude scaling factor, which depends on the instantaneous frequency as the AODs and the AWG have finite bandwidths. We round $f(t)$ to the nearest megahertz and use a lookup-table for $\xi(f)$ as discussed in Section 4.4.3.

Frequency chirp computation

Computing the chirp this way leads to two issues, both related to underlying AWG functionality. The first one arises from the length of the chirp, which is arbitrary in the mechanical picture but the AWG requires segment lengths being multiples of 32. We therefore pad the chirp by $[1 - 31]$ single frequency RF-values which prolongs the chirp by $31/f_s$. With a sampling frequency of $f_s = 1\,228.88\,\mathrm{MHz}$ this nanosecond delay is negligible.

The second problem is related to the phase at the end of a chirp, which in general is not equal to zero. We would therefore need to compute the ramp-down amplitude segment depending on the phase for

[3] Care has to be taken to determine the quadrant correctly. For numerical computations many programming languages implement an atan2-function, that ensures that the correct sign is used.

every possible combination of start- and end-sites. To avoid this, we calculate the phase of the signal one step after the final step:

$$\phi_{\text{final}} = 2\pi N_{\text{samples}} \times \alpha \times r_{\text{final}} \mod 2\pi \tag{6.28}$$

and update the phase of the n-th sample according to $\phi_n \rightarrow \phi_n - n \cdot \phi_{\text{final}}/N_{\text{steps}}$. With this substitution the next segment can start with a phase of zero and no phase-jump occurs. The error equates to one oscillation in $T f_b \approx 16000$ oscillations for a chirp duration of $T = 200\,\mu\text{s}$ at a base frequency of $f_b = 80\,\text{MHz}$. As lattice sites are separated by several megahertz, a deviation of a few kilohertz compared to the base frequency is again negligible.

Full move waveform

In addition to the frequency chirp phase, during which the trap depth is constant, two amplitude ramp phases are needed before and after a frequency chirp. During these ramps, the amplitude is smoothly ramped up or down with a $\sin^2$-dependence at the start and end frequency, which again avoids any parametric excitation of the trapped atom (see Section 5.7) due to sudden changes in the trap depth. The duration of these amplitude ramps is a free parameter and we will comment on the value that we use in Section 6.3.3. However, we can still make a statement about the possible values based on the intrinsic timescale of the problem. The latter is given by the inverse trap frequency at around 100 µs, as this is the time it takes the wavefunction to respond to a perturbation. As a rough estimate, a $\sin^2$-ramp with a duration of 100 µs does not carry frequency-components larger −60 dB above 10 kHz which is therefore slow enough to prevent an excitation. A schematic for the full waveform is shown in Figure 6.6(b).

6.2 Rearrangement algorithm

To decide how atoms are rearranged, we make use of a compression algorithm (similar to [39, 136]), which moves atoms from outer regions towards the center. We will introduce the algorithm, by solving an example configuration, before we discuss its shortcomings and two possible improvements.

6.2.1 Compression-algorithm

The compression-algorithm can be thought of as an ordered nearest-neighbor algorithm. Ordered in this case means, that it starts in the center of the array and moves outwards in circles or *shells* which we characterize by their Euclidean distance to the center. This approach solves a flaw of a pure nearest-neighbor approach, where it can happen that some atoms are moved multiple times because they block each other. The compression-algorithm does *not* find the minimum number of moves but it *limits the maximum* number of moves, with the limit being the number of empty target-sites. Details for the performance compared to other approaches can be found in [39]. To create a defect-free array, *target*-sites are defined which will be filled from *reservoir*-sites. On average, the total number of sites should be at least twice the number of target-sites, as the binomial distribution for N sites with $p \approx 0.5$ has a mean of $N/2$. Otherwise, many configurations will not contain enough occupied sites.

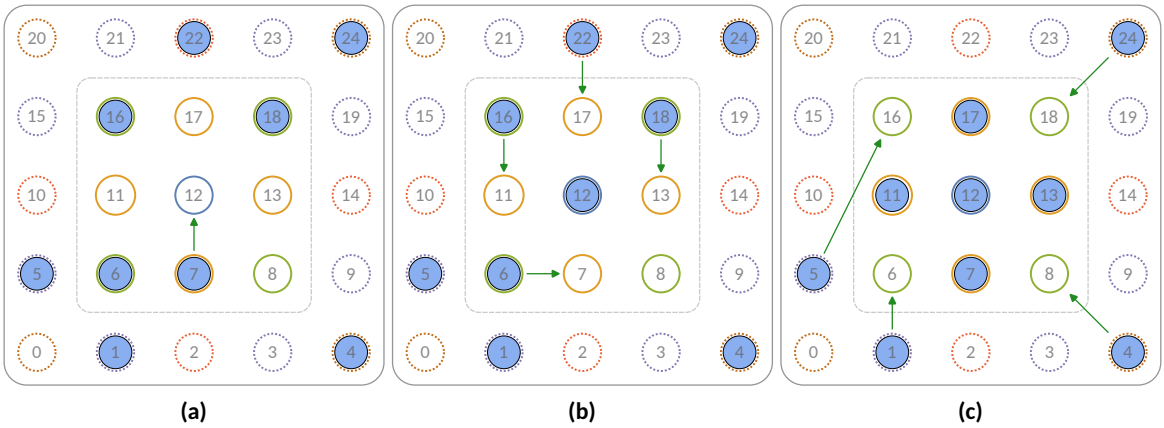

Figure 6.7 – Visualization of the *compression*-algorithm sorting a partially filled array (dashed 3×3 area) to full density using additional *reservoir atoms* (dotted circles). The color of the sites represents the distance to the center of the array. The latter is filled from the inside out to avoid moving an atom twice. The solution filles the innermost (blue) sites first **(a)**, continues with the second set (orange) in **(b)** and finishes with the last (green) sites in **(c)**.

Steps

Figure 6.7 shows the solution for a 5×5 square array, filled with nine sites, that are rearranged to a 3×3 target region. The global shells are represented as solid blue/orange/green circles in the target region and dotted red/violet/brown circles in the reservoir.

The first step is to check if the number of occupied sites is equal or larger than the number of target sites, otherwise the configuration is considered to be *not sortable*. If enough sites are occupied, the algorithm iterates through the sites in the target region to decide how sites are filled.

It starts at the innermost shell, consisting of sites with indices $\{12\}$[4]. Local shells l_n are constructed around the current target site, but only sites that are located *outside* of the current global shell are included. For site 12, this includes sites $l_1 = \{7, 11, 13, 17\}$ for the first, and sites $l_2 = \{6, 8, 16, 18\}$ for the second local shell. We truncate the shells at sites with a distance of 2.85 (in units of lattice spacing) for the adjacent sites. This distance means that for every site a 5×5-block of sites is checked (for site 12 this would be the whole grid)[5]. If the center of the populated site distribution deviates from the center of the array, it is possible that no site to move from is found, and the configuration is considered as *not sortable*. For every site s_i in the first local shell l_1, a local density based on the surrounding occupied sites is computed. For the local density only sites that lie outside the current global shell are considered and again only sites with a relative distance to s_i of less than 2.85. If more than one possible site to move from exists, the local density value is used to pick atoms from higher density regions first. If two sites have the same local density, the site with the smallest index is picked. This choice introduces a bias that we discuss in the next section.

For the case at hand, site 12 is filled from site 7 and the algorithm moves on to the next shell with sites $\{7, 11, 13, 17\}$. The now empty site 7 has to be refilled, and the first local shell is constructed from sites $\{2, 6, 8\}$. Site 12 is not considered, as it is not outside of the global shell of site 7. As site 6 is the only

[4] For an even number of sites along both axes, the center shell would consist of four sites, and two sites if one is even and one is odd.

[5] The distances for sites contained in the local shells l_i are therefore: $d(l_1) = 1$, $d(l_2) = \sqrt{2}$, $d(l_3) = 2$, $d(l_4) = \sqrt{3}$, $d(l_5) = 2\sqrt{2}$. The next possible distance for is $d(l_6) = 3$ and therefore not considered.

occupied site, it is moved to site 7 and we continue. Following the logic, site 11 is filled from site 16, 17 from 22 and 13 from 18, which completes the shell.

The third and last global shell in the target region consists of sites $\{6, 8, 16, 18\}$. In this example, the assignment in Figure 6.7 gives a collision-free[6] solution, but we will now see that this was merely a coincidence.

Optimizations

One problem is the bias in the selection, based on the index of a site, which results in an unwanted ordering. For example, filling the last shell with sites $\{6, 8, 16, 18\}$, the solution would produce a collision, if site 6 is filled from site 5 instead of site 1 as site 16 then has to be filled from site 1. This is not happening purely because for site 6 site 1 and 5 are equal candidates and site 1 is picked because it has a lower index. If the configuration is e.g. rotated by 90° counter-clockwise, the assignment would produce a collision when site 1 is moved through sites 6 and 11 to site 16.

Relying on chance is of course something we would like to avoid and there are modifications to the algorithm which can help in a case like this. Another issue is the fact that target sites that are already filled are moved-from and then have to be refilled. This increases the number of moves, which in general we would like to minimize because every move has a finite success probability.

Identity I	Mirror horizontal H	Mirror vertical V	Rotate 180° R_{180}
$\begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix}$	$\begin{pmatrix} c & b & a \\ f & e & d \\ i & h & g \end{pmatrix}$	$\begin{pmatrix} g & h & i \\ d & e & f \\ a & b & c \end{pmatrix}$	$\begin{pmatrix} i & h & g \\ f & e & d \\ c & b & a \end{pmatrix}$
Transpose T	**Anti-Transpose T_A**	**Rotate 90° R_+**	**Rotate 270° R_-**
$\begin{pmatrix} a & d & g \\ b & e & h \\ c & f & i \end{pmatrix}$	$\begin{pmatrix} i & f & c \\ h & e & b \\ g & d & a \end{pmatrix}$	$\begin{pmatrix} c & f & i \\ b & e & h \\ a & d & g \end{pmatrix}$	$\begin{pmatrix} g & d & a \\ h & e & b \\ i & f & c \end{pmatrix}$

Table 6.3 – Possible array transformations. T, T_A, R_+ and R_- are only applicable to square arrays.

Unique configurations The first change reduces the influence of randomness in the solubility. After checking if enough sites are filled, we transform the array into a different configuration which we then solve. This other configuration is a special configuration out of all configurations that are related by symmetry transformations. Table 6.3 shows the possible transformations at the example of a 3×3 square grid. Note that the result for a transformation is not necessarily unique. Before solving a configuration, we compute all eight transformed configurations and pick the *smallest*[7], which we then solve. After solving the (possibly transformed) configuration, the set of moves has to be transformed back to solve the

[6] We define a collision as a move that passes through an occupied site, or within a certain distance d_{coll}.

[7] We enumerate all configurations and compare them with respect to this ordering. For smaller configurations $\leq 8 \times 8$, we can interpret the flattened array as an unsigned integer and configurations can be directly compared. Otherwise, we can iterate through two configurations and make a decision based on the largest position index where they differ in value. Picking the smallest is just a convention, the largest would also work.

original configuration. This procedure prevents the issue mentioned above, where a rotated configuration results in a different solution. It does however not change the fact, that the solved configuration might be the one that fails because of the index bias. The best option, in terms of success rate, would be to solve *all* possibilities, however this increases the runtime exponentially. Depending on the array size and the frequency of such branching events, it might still be a viable option that is worth to be explored in the future.

Chain optimization The second modification that we include, is a post-selection of the computed moves. If we look at the solution of the grid above, we can detect two chains of more than one move: $\{1 \rightarrow 6 \rightarrow 7 \rightarrow 12\}$ and $\{24 \rightarrow 18 \rightarrow 13\}$. We label these chains by the number of connections, so the first would be a three-move chain and the second one a two-move chain. Let us focus on the two-move chain first. If such a chain is detected, we can check if the intermediate site can be skipped, in this case site 18. For this, we check if there is an occupied site closer to the line connecting the start and end site than a certain distance $d_{\text{collision}}$. If this is the case, we do not make the replacement, otherwise we replace both moves by the move $\{24 \rightarrow 13\}$. We have to be careful though, as it could be that the first move (moving from the intermediate site) and the second move (refilling it) might be at different times in the sequence of steps and it could happen that an atom is moved over the intermediate site in the mean time. Optimizing the moves would therefore lead to a collision, which was not present before. We will investigate the increase in collisions compared to the reduction in moves in the next section.

For three-move (or longer) chains it gets more complex. If we find that connecting the start and end site directly produces a collision, we still have the option to optimize sub-chains. For three-moves this leaves two options and we can check both for possible collisions. If both produce collisions or if only one is collision-free, the choice is obvious but if both are possible, further factors have to be considered. In the future, a more elaborate approach for example by reordering the moves or partially resolving the colliding part probably yields better results, for now chains with more than three sites involved are omitted.

6.2.2 Algorithm performance

To benchmark our modifications, we simulate 10^6 configurations for 6×6 square arrays that we solve[8]. We exclude solutions that contain moves which pass by an occupied site with a distance of $a/\sqrt{2}$, where a is the lattice spacing, or less. Overall, 45.4 % of the configurations are sortable without the chain optimization and 45.8 % with the optimization. We further compare the number of moves in a solution and the number of suboptimal solutions. We define suboptimal solutions as solutions that deviate from the minimum number of moves, where the latter is given by the number of empty sites in the target region. Note that this underestimates the performance of the algorithm, as the existence of an optimal collision-free solution is not guaranteed. Currently, we can only test if a solution is optimal but we cannot prove its existence, let alone compute it[9].

We show the results of the analysis in Figure 6.8. The average length of a solution without optimizing chains is 9.64 compared to 9.03 with the optimization. Compared to the mean minimum number of moves of 8, this is reduction of 40 % in excess number of moves. Additionally, we resolve the deviation from the optimal solution in Figure 6.8(b). Note that we do not normalize the deviation to the optimal

[8] We obtain similar results for other array sizes, e.g. 5×5 or 8×8 arrays.

[9] If we had access to an algorithm that could compute a collision-free optimal solution we would use that instead.

solution, as the absolute number of moves is the quantity to be minimized. As an example, a reduction of an 8 move solution to a 5 move optimal solution is a larger improvement than reducing a 2 move solution to a 1 move optimal solution, even though the relative decrease is less. The improvement is shown in Figure 6.8(c), where a significant improvement for 0 to 2 move solutions can be seen.

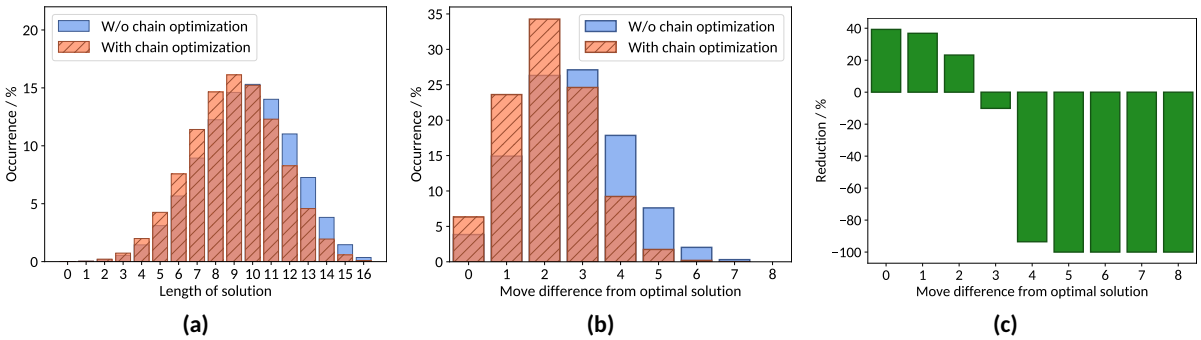

Figure 6.8 – Rearrangement algorithm benchmark for 5×10^5 random configurations of 6×6 square array with 16 target sites. **(a)** Number of moves in a solution. The average number of moves in a solution is reduced from 9.6 to 9, which is a 40 % improvement compared to the mean minimum number of moves of 8. **(b)** Length of a solution compared to the optimal solution. Note that the latter does not necessarily exist and the performance of the algorithm is therefore underestimated. **(c)** Improvement of the chain optimization, where a significant improvement for shorter solutions can be seen. For visualization purposes, values have been clipped to -100 %.

6.3 Experimental implementation

6.3.1 Overview

The combined hard- and software components and their interplay are shown in Fig. 6.9. The sorting process itself can be broken down into six steps:

1. First image
2. Determine site occupation
3. Compute moves
4. Compute and upload waveforms
5. Move atoms
6. Second image

To take high-fidelity images, the process described in Section 5.5 is used, where we optimize the imaging parameters by taking two images with a delay of about 120 ms. The latter is the average time between two images when rearranging a 6×6 array into 4×4 center sites. The frame-transfer together with the binarization takes 26 ms with a jitter of less than 1 ms. The solution for the moves is computed in around 0.5 ms which we determine as the average time from a simulation by solving 10^6 configurations. For a single repetition, we measure around 1 to 2 ms where the difference can be explained by an overhead and jitter. The duration is consistent with values observed in the literature for a CPU-based

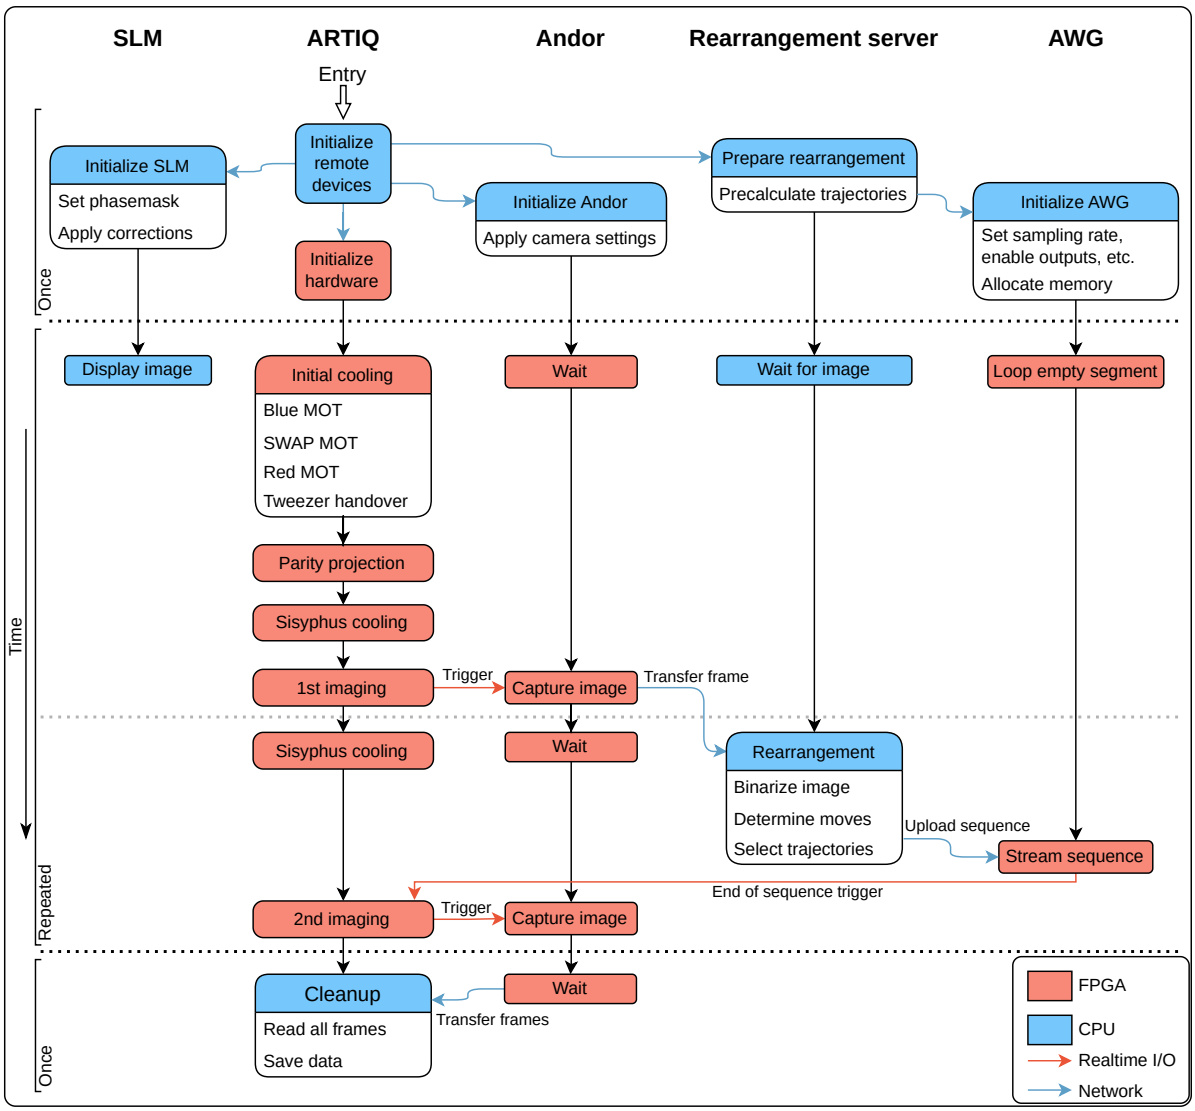

Figure 6.9 – Schematic of the sorting sequence

implementation [136]. The time it takes to compute the waveform and upload it depends linearly on the duration for a single move, which in turn depends on the kinematic parameters. For the parameters in this work (see upcoming Section 6.3.3), a move takes around 800 µs, consisting of two amplitude ramps of 300 µs and a frequency chirp of 200 µs. It takes around 3 ms to compute a single-move waveform and 6 ms to upload it to the card. For an average number of moves of 9 for the 6×6 grid this is a significant overhead[10]. Depending on the array size, it is possible to precompute and even preupload all waveforms. We precompute all moves but do not preupload them, which results in an average duration of 50 ms to upload the sequence. We will discuss possible solutions to this problem at the end of this chapter. The execution of moves scales linearly with the number of moves, so in our case around 7 ms for an average

[10] This value is slightly larger than the theoretical value of 8, as the filling is less than 50 % due to losses.

number of 9 moves. Finally, we wait and cool for an additional 50 ms before taking the second image. The cooling itself is necessary, to ensure that all atoms enter the second imaging phase with the same temperature. However, the duration is longer than required for cooling alone, as the camera is read out using a USB2 connection which limits the time between images to around 100 ms. It also features a *CameraLink* which could be used in the future to speed up the process making the delay obsolete.

We optimize the sequence step by step by using multiple trigger signals from the ARTIQ system, the AWG and the software which we detect with a logic-analyzer[11]. Specifically, we monitor the camera frame trigger, to detect the start of the rearrangement block. To differentiate between the first and the second frame, we use a second trigger which starts at the first image and ends at the start of the second image. We use an option of the AWG to encode the sequence into a two-bit number[12], which we route to two TTL-channels. The output is intrinsically linked to the current segment being replayed and therefore follows the analog output stream in real-time. This allows us to detect the start and end of the stream and the correct order of segments.

To gain insight in the timings of the code running on the CPU, we use an output port from the AWG that can be triggered from software. This non-*real-time* trigger has a higher jitter, as it is not independent on the system load, and it is therefore limited to a resolution of a few milliseconds. The events and durations are shown in Figure 6.10 and summarized in Table 6.4.

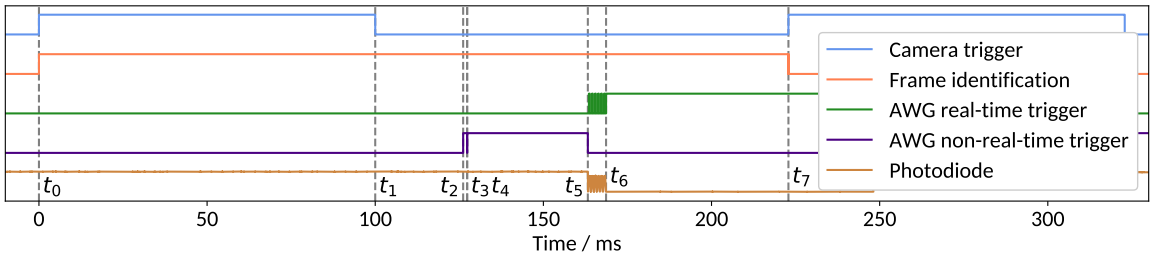

Figure 6.10 – Overview of time events during the sequence, where details for the timings are given in Table 6.4. The physical rearrangement of atoms happens between t_5 and t_6, which is further explained in Figure 6.13. The photodiode-signal is picking up from the main beam after the AODs, see Figure 4.27. We apply an RF-signal to the AODs with a frequency corresponding to a position outside of the array before the moves to avoid thermal switching effects.

6.3.2 Alignment

We find that the dynamic tweezer and the center of the array have to be aligned within 0.25 µm in radial direction, which we ensure by displacing the dynamic trap relative to the static array. We optimize the alignment by sorting 500 – 1000 configurations and determine the success probability given the starting site. Figure 6.11(a) shows an example of a well aligned system, where no spatial dependence is visible compared to Figure 6.11(b) where the frequency spacing in horizontal direction is incorrect. The four center-most sites are never moved-from and therefore do not contribute to the measurement. We repeat the analysis for the end sites of the moves and observe the same behavior. The axial alignment is difficult to optimize, as the dynamic trap cannot be shifted continuously. Moving the static array is also

[11] Saleae Logic Pro 8

[12] The output resolution is reduced by one bit per channel to 15-bit. The value of the free bit can be programmed in the segment and its value is routed to a TTL output.

	Event	Duration
$t_1 - t_0$	First camera trigger	100 ms
$t_2 - t_1$	Camera readout and frame transfer	26 ms
$t_3 - t_2$	Move computation	1 to 2 ms
$t_4 - t_3$	Waveform computation	3 ms per 800 µs move, 25 µs for lookup if precomputed
$t_5 - t_4$	Waveform upload	6 ms per 800 µs move
$t_6 - t_5$	Stream	< 10 ms
$t_7 - t_6$	Additional cooling	50 ms
$t_8 - t_7$	Second camera trigger	100 ms
$t_8 - t_0$	Total duration for 9 moves	340 ms

Table 6.4 – Durations of events in the rearrangement sequence.

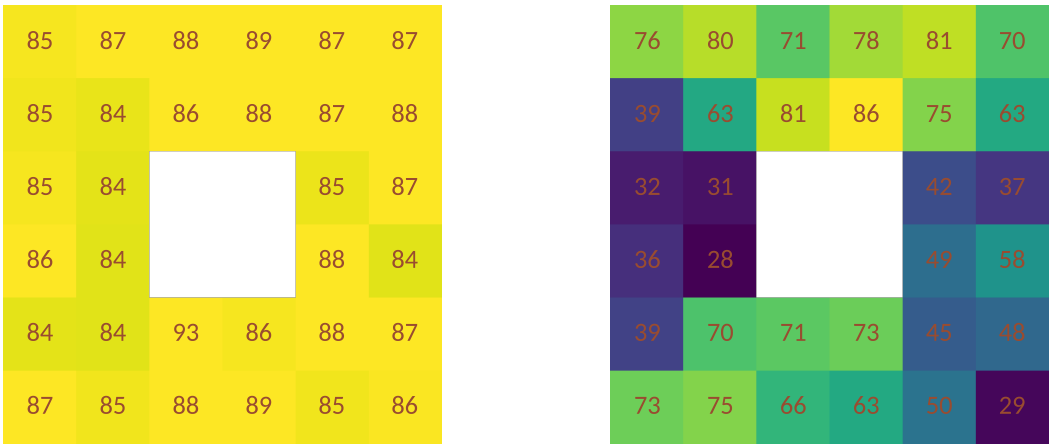

Figure 6.11 – Success probability depending on the starting site in a 6×6 grid. **(a)** For a well aligned system the success probability shows no spatial dependence. **(b)** An incorrect space-frequency conversion factor for the AOD in horizontal direction leads to a spatial dependence. The four center-most sites are never moved-from and do not contribute to the measurement.

complicated to automate, as the intensity balancing of the array changes for axial displacements and additionally the EMCCD camera has to be repositioned[13]. We find that the axial focus positions have to be aligned within a micrometer but for a more precise value the overlap has to be tuned continuously. Overall, we find that the alignment has to be optimized every 6 to 9 h, due to the movement of the stage of the upper objective. Small radial drifts can in principle be compensated automatically, but the unknown axial mismatch is still an issue. Therefore, the proper solution is to use a single microscope objective which is used to project both trapping wavelengths, to avoid any relative movement.

[13] It is also possible to verify this by transferring from a static trap to the dynamic trap and back without moving the dynamic trap.

6.3.3 Dynamic tweezer parameters

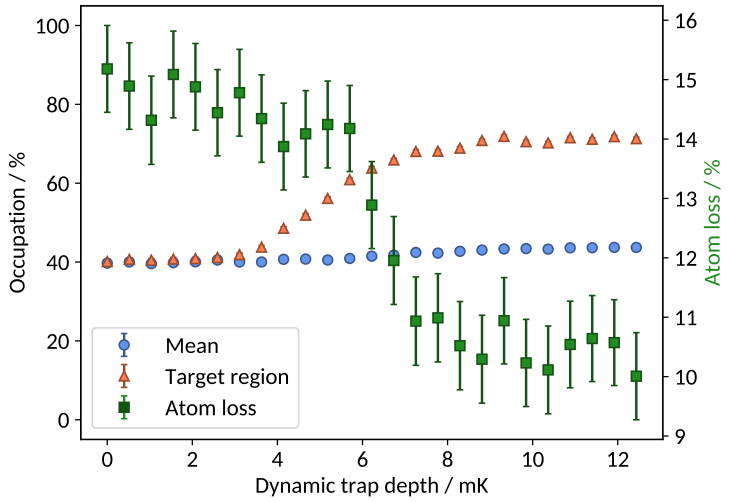

Figure 6.12 – Full array mean (blue) and target-region only (red) occupation in the second image for different dynamic trap depths. For low dynamic trap depths, atoms are not moved but rather heated due to the modulated potential. Data for other measurements in this work is taken at 9.5 mK.

6.3.3.1 Kinematic parameters

The three regimes for the atom trajectories developed in Section 6.1.1.1 describe the intuitive fact that the transfer time is minimized, if we move as fast as we can at every instant, without exceeding the parameter limits. However, this shifts the problem to the determination of these limits. Up to now, we did not map out the whole parameter space, but we found a working point based on physical arguments which also matches values from the literature.

As a conservative measure, we limit $j_{\max}$ to $a_{\max}/100\,\mu s$, although values down to $a_{\max}/10\,\mu s$ might be possible. This is slower than the inverse trap frequency which is the intrinsic timescale of the problem and dictates how quickly the wavefunction can respond to a perturbation. We find an increased atom loss for accelerations above $1\,500\,\mu m/ms^2$ and therefore work at $1\,000\,\mu m/ms^2 \approx 100g$ (and in turn $j_{\max} = 1 \times 10^4\,\mu m/ms^3$). Further investigation is necessary to precisely map out the loss mechanism, as the value for the acceleration is still orders of magnitude below the value obtained for a loss due to a tilted potential alone.

For the velocity, we base our argument on the finite propagation time of the acoustic wave inside the AOD, as shown in Figure 6.3. Assuming a clear aperture for the AOD of 5 mm and a speed of sound inside the crystal of 650 m/s, it takes the around 1 µs to pass the aperture. If the frequency changes significantly during this time, different parts of the incident beam will be deflected at different angles and in turn the trap deforms. This deformation can lead to heating caused by the parametric modulation and to a lower confinement, especially in the axial direction. Changing the frequency by 0.8 MHz in 1 µs leads to an change in RF-frequency around 1 % at a base frequency of 80 MHz and corresponds to a displacement of 1 µm with a velocity of 1 000 µm/ms in the atom plane. As the deflection angle, and therefore the trap deformation, is proportional[14] to the RF-frequency, the trap deforms by the same

[14] To be precise, only the deformation in the radial plane is directly proportional to the angle, the axial direction scales roughly

amount, which we declare to be negligible. To preserve some margin of error, we choose a maximum value of 150 µm/ms, similar to references values of 10 to 100 µm/ms from [136] and [39].

For our lattice spacing of 5.2 µm, moves traveling one lattice spacing along either axis or along the 45°-diagonal lie in the distance limited case. Therefore, $a_{\max}$ and $v_{\max}$ are not reached but lie around 750 µm/m^2 and 60 µm/ms respectively. The precise values that are reached for different distances are shown in Table 6.5.

Parameter	Distance / a			
	1	$\sqrt{2}$	$\sqrt{3}$	2
$a_{\max}$	650 µm/ms^2	730 µm/ms^2	780 µm/ms^2	820 µm/ms^2
$v_{\max}$	42 µm/ms	53 µm/ms	61 µm/ms	67 µm/ms

Table 6.5 – Maximum acceleration and velocity an atom reaches while traveling different distances in a square lattice with spacing $a = 5.2$ µm.

6.3.3.2 Trap depth

We optimize the dynamic tweezer depth by measuring the averaged occupation in the reservoir and target region. If atoms are moved, the fraction of occupied sites in the reservoir-region drops, while it rises in the target-region. Figure 6.12 shows a measurement, where it can be seen that a steady state is reached for trap depths larger than 8 mK. This trap depth is much deeper than the 300 µK deep static array and the values are consistent with the literature [136]. An interesting observation is the loss averaged over all sites depending on the dynamic trap depth. Below 6 mK, we measure an increased loss, before it settles to a value of 10 % above 8 mK. The latter is consistent with values we observe without the rearrangement and gives a hint that the loss due to the latter is close to zero. We explain the increased loss below 6 mK by the fact, that the dynamic and the static tweezer are of comparable depths and the dynamic trap leads to parametric excitation by deformation of the local potential, without picking up an atom. We could not observe a dependence on the ramp duration from 80 to 500 µs and therefore work at 320 µs, similar to the literature [136]. Shorter times might be possible but the changes on the order of 100 µs correspond to frequencies of 10 kHz, which come close to the axial trap frequency. This is not necessarily a problem during the ramp-up, because the atom is trapped in the deeper trap, but for the ramp-down the increased temperature could lead to atom loss in the shallow static tweezers. We therefore continuously apply cooling light on the light-shifted resonance of the static array to mitigate heating after an atom is moved. Due to the strong light-shift in the dynamic trap, the latter is not affected. For this work we work at a trap depth of 9.5 mK.

6.3.3.3 Intensity stability

As the RF-frequency changes by several megahertz during a chirp, the amplitude also changes. We stabilize the intensity with a frequency-dependent lookup-table (LUT) to scale the instantaneous intensity accordingly (see Section 6.1.2). To evaluate the quality of the stabilization, we measure the intensity with a photodiode and a pick-off plate in the beam. Figure 6.13(a) shows an example for six chirps,

with the square of the radial direction.

where the vertical dashed lines mark the start and end of a chirp. The dotted lines mark the transition of one move to the next one. For the chirp, we measure a standard deviation of 0.75 % and a peak-to-peak deviation of 3.6 % when averaged over 640 chirps. Comparing to the loss measurement in Figure 6.12, a fluctuation in depth of this order around the depth of 10 mK is not an issue. To estimate the effect of the intensity noise in terms of heating, we compute the frequency spectrum from a numerical Fourier transformation. The time resolution of 0.64 µs leads to a Nyquist-limit of around 750 kHz, which is large enough to include the radial trap frequency of the deep dynamic trap of ≈ 500 kHz. The spectrum is shown in Figure 6.13(b) and shows no peaks above −70 dB for frequencies larger 35 kHz which is well below the axial trap frequency of the dynamic trap[15].

This is a good result for this type of stabilization, however, it could still be improved. One option would be a more precise LUT, for example using frequency oversampling and filtering. A second option would be to decrease the magnification of the telescope creating the dynamic tweezer, which reduces the frequency-to-space conversion factor, leading to smaller frequency changes. This would also increase the trap size and confinement (see Section 4.2) towards the optimal truncation ratio of $\zeta = 0.65$, which is currently not possible due to spatial constraints on the optical setup. Nevertheless, the passive stabilization is expected to not reach the stability level of 0.1 %, which can be obtained with an active stabilization for our system. However, the active stabilization has a maximum bandwidth around 200 kHz and is therefore probably not fast enough to stabilize chirps with durations of ∼ 200 µs or less.

We therefore conclude that the intensity is sufficiently well stabilized to not limit the rearrangement success probability.

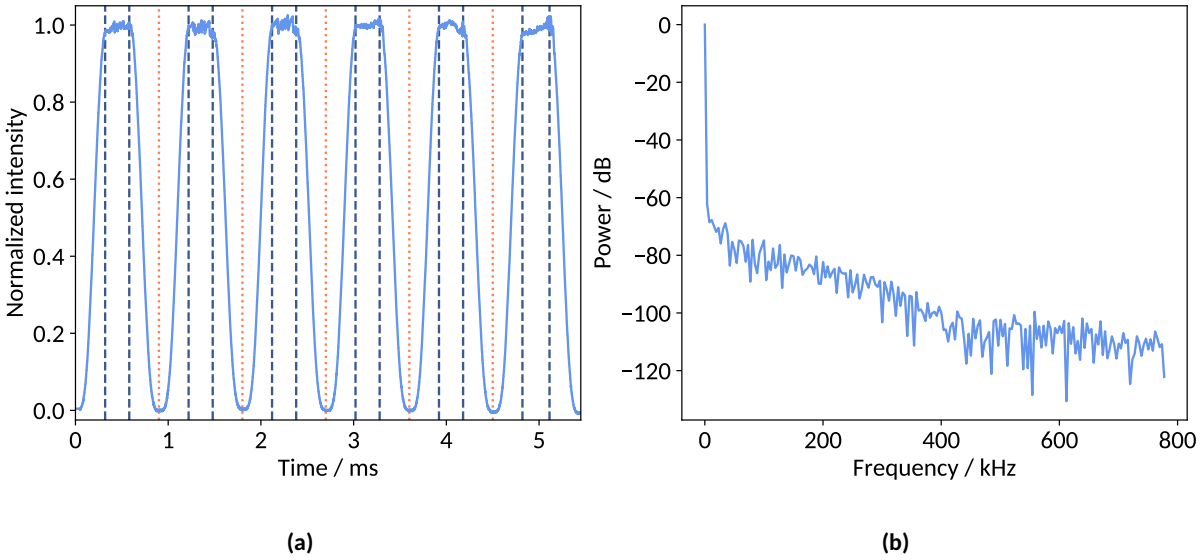

Figure 6.13 – Performance of the passive intensity stabilization using a lookup-table. **(a)** Intensity stability during six moves, separated by red dotted lines. The amplitude is changed with $\sin^2$-ramps between the frequency chirps (dashed lines). We compute a mean relative standard deviation (mean relative peak-to-peak) of 0.75 % (3.6 %) averaged over 650 moves for the intensity during the chirps. **(b)** Power spectrum of the intensity during the chirp, averaged over 650 moves.

[15] In Section 5.7 we measure an aspect ratio of 5.3 for the dynamic tweezer setup. From the calculated radial trap frequency of 500 kHz an axial trap frequency of 90 kHz follows.

6.4 Fidelity estimation

A key question that we would like to answer is how well we can prepare a defect-free array. The answer depends on the single-move fidelity $\mathcal{F}$, defined as the success probability of moving an atom, which has to be determined experimentally. Global losses $\mathcal{L}$, stemming from the finite lifetime and loss during the imaging process, complicate the seemingly straightforward task and require the development of a method to separate $\mathcal{F}$ and $\mathcal{L}$. For the analysis, we do not consider errors associated with a finite imaging fidelity, as the latter is dominated by the uncertainty of other losses $\mathcal{L}$. If $\mathcal{L}$ and its uncertainty are reduced in the future, a finite imaging fidelity has to be accounted for in the analysis. We will present two methods to do so: First, we correct the measured success probability by the losses determined from a separate measurement without the rearrangement step. Secondly, we develop a statistical approach, that determines $\mathcal{F}$ and $\mathcal{L}$ from a single measurement, at the cost of requiring a larger dataset.

For the analysis, we only consider *collision-free* solutions, where we define a collision as a move which passes by an occupied site with a distance of less than $d_{\text{coll}} = a/\sqrt{2}$ with the lattice spacing a. This can be seen as a safety measure, as we have not yet explored what happens if this distance is undercut on a single-move basis. We briefly comment on the observation in the first part of the following discussion, but we exclude colliding solutions for the determination of $\mathcal{F}$ and $\mathcal{L}$.

6.4.1 Distance based analysis

We first investigate the move fidelity depending on the distance of the move. A success is defined as a move, whose target site is not empty afterwards, which combines the move fidelity and the aforementioned imaging and lifetime losses. Figure 6.14 shows the evaluation of a dataset of 700 repetitions. We label moves by their distances, d_1 for a move which travels $(1, 0)$ or $(0, 1)$ sites, $d_{\sqrt{2}}$ for $(1, 1)$ and so on. The left y-axis shows the fraction of moves with a given distance (blue) and the fraction of moves for a given distance that were successful (red striped). Their ratio, the uncorrected success probability, is shown by the green dots referenced to the right y-axis. For d_1 and $d_{\sqrt{2}}$ moves, we determine a success probability of $(87.7 \pm 0.5)\,\%$ and $(87.9 \pm 1.0)\,\%$ respectively. The low occurrence of $d_{\sqrt{3}}$ results in a large uncertainty for an estimate of the success probability, but there are multiple reasons why the latter could actually be lower. First, for the current parameters, the acceleration could be too high, leading to heating and subsequent atom-loss during the handover at the end of a move. A second explanation could be the distance to a static trap for such a move, which would lead to a modulation of the trapping potential, also causing heating. The precise effect on the motional state of the atom strongly depends on how the static potential is passed, where minor spatial deviations can lead to significant differences. A preliminary analysis on a single-move basis, including colliding solutions with d_2 and $d_{2\sqrt{2}}$ moves, shows similar success probabilities around 80 % for both distances. However, with minimum distances between the dynamic and static trap between 0 to $0.44a$, which results in vastly different modulation strengths, this explanation is rather unlikely. In order to locate the issue, repeating the measurement with different values for a_{max}, i.e. limiting it to a maximum value below the value reached for a d_1 move, could be performed.

Coming back to the results for d_1 and $d_{\sqrt{2}}$ moves, we correct the success probability by the known losses of around $\mathcal{L} = (10.5 \pm 1.0)\,\%$ at the time[16]. Correcting the averaged success probability for d_1

[16] This is done by taking two images without the rearrangement step at otherwise equal parameters and timings, followed by a comparison of the populated fraction of *all* sites.

and $d_{\sqrt{2}}$ by this value leads to a fidelity of $\mathcal{F} = (98.1 \pm 1.3)\,\%$. This value is in great agreement with the values of $0.98 - 0.99$ observed in the literature [39, 142]. A drawback of this method is that the calculated value and its uncertainty crucially depend on the precise knowledge of $\mathcal{L}$, which has to be determined separately. However, as these losses are present in the data, we now present a second method, which allows us to isolate the sorting fidelity and global losses without taking additional data.

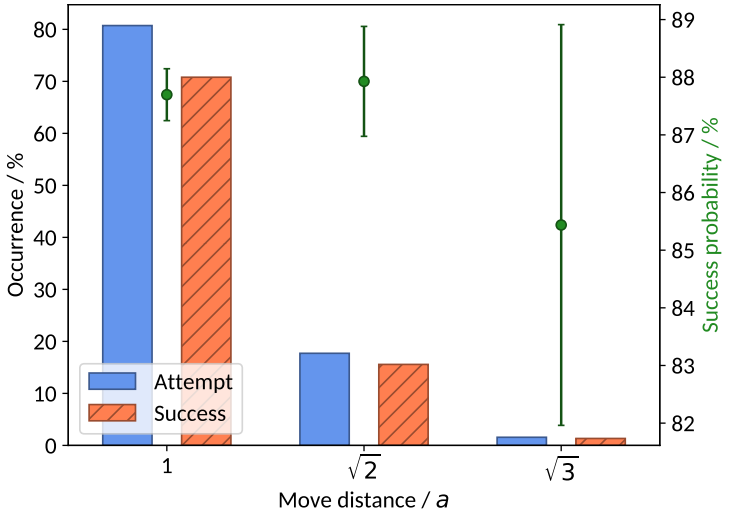

Figure 6.14 – Success probability depending on the move distance including global losses. With knowledge of the latter the single move fidelity can be calculated.

6.4.2 Number-of-moves based analysis

The model

The second approach is based on the distribution of occupied target sites after the rearrangement process. Figure 6.15(a) shows the latter together with the distribution before the rearrangement step. For configurations that have at least N_t populated sites, the mean is increased from 9.2 to 14.2, which corresponds to a probability per site of 0.58 and 0.89 respectively for a binomial distribution. However, the distribution after the rearrangement is not binomial but rather a superposition of binomial distributions $B \propto \sum_n B_n(N_t, p_n)$ with mean-values $\langle N_n \rangle = N_t p_n$, where N_t is the number of target sites[17]. In this case, the p-value of the distribution depends on the global losses but also on the number of moves n. Imagine for a second that there are no losses and that we have a perfect move fidelity, leading to $p_n = 1$ and every target site always being populated. Adding finite global losses reduces p_n by $\mathcal{L}$ and the distribution is given by $B(N_t, 1 - \mathcal{L})$. If we now also include a finite move fidelity $\mathcal{F} < 1$ (independent on the move distance), every move has a probability of $1 - \mathcal{F}$ to fail. The compression algorithm tells us, that every move ends at a target site and if the move fails, this site is empty. Therefore, every move reduces the mean number of occupied target sites $\langle N_n \rangle$ by $1 - \mathcal{F}$ and in turn p_n is reduced by $(1 - \mathcal{F})/N_t$.

[17] The resulting distribution is a *Poisson Binomial Distribution*.

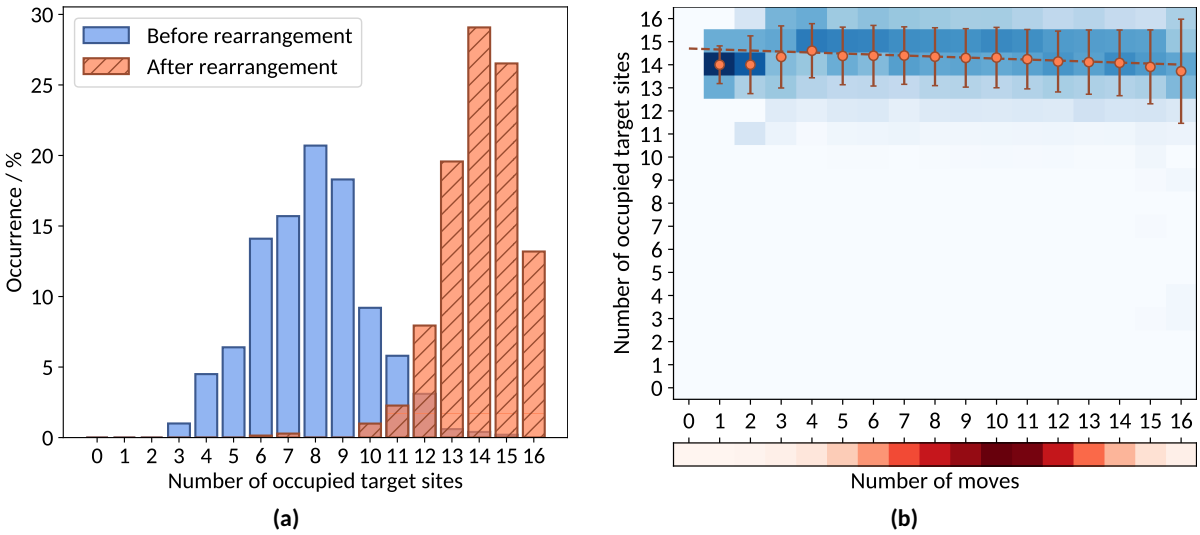

Figure 6.15 – **(a)** Histogram of populated target sites averaged over all solutions. The blue (red) bars show the distribution before (after) the rearrangement process where the mean is increased from 7.9 to 14. Note that the mean of the distribution after the rearrangement step is reduced by imaging losses and finite lifetimes. As the p-parameter depends on the number of moves, the distribution is actually a sum of many binomial distributions, known as a *Poisson binomial distribution*. **(b)** Populated target sites after the rearrangement step, resolved for the number of moves that have been performed. Each column follows a binomial distribution and by a linear fit (dashed line) to the mean value (red dots) of each distribution the fidelity and losses can be calculated. To visualize the result, each column has been normalized to the accumulated number of counts in that column. The occurrence of a solution with a specific number of moves is shown by the red colorbar below. The red errorbars depict the standard deviation of the respective binomial distribution, whereas the fit is performed with weights determined by the accumulated counts from the red colorbar below.

Combining the global losses $\mathcal{L}$ and the loss for n moves leads to:

$$p_n = 1 - \mathcal{L} - \frac{1-\mathcal{F}}{N_t} n \quad n \in \left[0, N_t\right]. \tag{6.29}$$

By counting the number of populated target sites depending on the number of moves n in a solution, we can determine the mean $\langle N_n \rangle$, the probability p_n and in turn $\mathcal{L}$ and $\mathcal{F}$. To satisfy the requirement of a constant $\mathcal{F}$, independent of the move distance, we only consider solutions that contain d_1 and $d_{\sqrt{2}}$-moves (see Figure 6.14). We extract the mean $\langle N_n \rangle$ of the distribution of N_n depending on n. We fit $\langle N_n \rangle$ with the linear model from Equation (6.29), where each mean-value has a weight given by the number of counts in the respective histogram. This includes the distribution of solution-lengths in the model (see also Figure 6.8(a)). An example is shown in Figure 6.15(b).

From two experimental datasets of sizes 700 and 14000 we determine fidelities of (96.3 ± 2.2) % and (95.6 ± 0.5) % and losses of (10.1 ± 1.3) % and (8.1 ± 0.3) % respectively. Before we can further comment on these values and compare them to other results, we have to determine the error introduced by the analysis method.

Systematic error estimation

To determine the systematic error of the method, we simulate configurations using a Monte Carlo method. For a given initial configuration C_i, we calculate the moves $\mathcal{M}$ for the smallest configuration together

with the chain-optimization method (see optimizations in Section 6.2.1) . We apply the post-selection and discard solutions with moves longer than $\sqrt{2}$. We simulate the final configuration C_f by applying the sequence of moves, each with success probability $\mathcal{F}$. In the end, each site has a probability of $\mathcal{L}$ to be emptied.

We generate 2×10^6 configurations with an occupation of at least 16 and simulate the result with a move fidelity of 96 % and losses of 10 %, as these are the approximate values determined above. Out of these configurations, 900000 are sortable and around 720000 have a short-move solution. To estimate the influence of the size of the dataset, we draw subsets of different sizes out of the solutions which we analyze with our method. We keep the total number of configurations per dataset size roughly constant at 250000 by drawing more samples for smaller dataset sizes, i.e. size 2000 is sampled 125 times, whereas size 125000 is only sampled twice (similar to a binning). Averaged over dataset sizes of 500 to 250 000, the fit values reproduce the underlying losses of with a deviation of −0.08 p.p.[18] and the fidelity with a deviation of 0.25 p.p.. To estimate the mean deviation and the standard deviation for our experimental datasets with sizes of 700 and 14000, we average over 360 and 18 repetitions respectively.

In general, the method relies on the correct determination of the mean value for each histogram of n moves. For fidelities approaching unity, these histograms are compressed into the upper few bins, which in turn requires a large number of repetitions for sufficient resolution. The method also benefits from larger array sizes, as more bins are available but for these arrays the number of short-move solutions reduces as well.

Dataset size	Quantity	Fit / %	Systematic error / p.p.	Corrected / %
700	$\mathcal{L}$	(10.1 ± 1.3)	(−0.2 ± 1.5)	(10.3 ± 2.0)
	$\mathcal{F}$	(96.3 ± 2.2)	(0.1 ± 2.6)	(96.2 ± 3.4)
14000	$\mathcal{L}$	(8.1 ± 0.3)	(−0.1 ± 0.4)	(8.2 ± 0.4)
	$\mathcal{F}$	(95.6 ± 0.5)	(0.2 ± 0.4)	(95.4 ± 0.7)

Table 6.6 – Fidelity estimation with statistical method to isolate global losses $\mathcal{L}$ from per-move losses of $1 - \mathcal{F}$.

Evaluation

Table 6.6 shows the fitted values from two datasets and the correction by including the systematic error of the method. The extracted value for the losses is in excellent agreement with the value obtained by other measurements at the time. The fidelity for the first dataset also compares well to the value of 0.98 – 0.99 from [39, 142]. For the second dataset with a measurement duration of ∼ 12 h, a reduced fidelity is to be expected, as the alignment drifts for such long measurement durations.

In general, the second evaluation method works well to extract and isolate losses and fidelities with a single measurement but it requires sufficiently large (> 10000) dataset sizes. This is currently prevented by the stability of the system, where a manual realignment after ∼ 6 to 9 h is required to overlap the static array with the dynamic tweezer. These drifts are on the order of 500 nm both radially and axially and can only partially be automated (Section 6.3.2).

For now, the most reliable result can be obtained by analyzing a single dataset by combining both methods. In this case, we determine the loss with the second method, which we then use

[18] percentage points

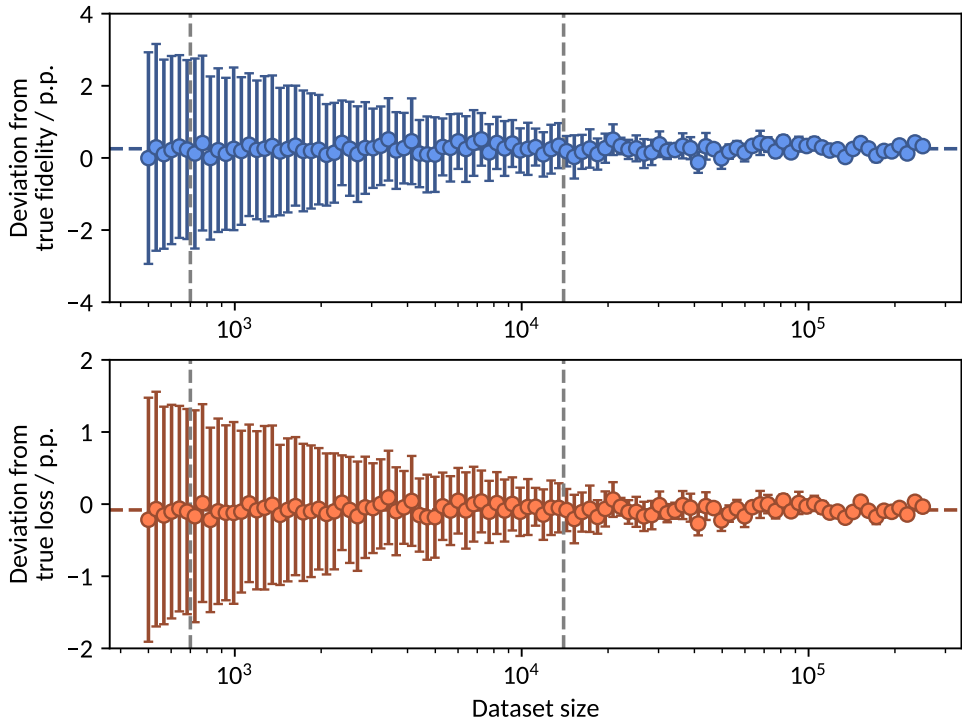

Figure 6.16 – Random initial configurations are created and rearranged by a numerical simulation with a fidelity of 96 % and losses of 10 %. We analyze the result to estimate the systematic error of the analysis method. On average (horizontal dashed lines) the loss is underestimated by 0.08 p.p. whereas the fidelity is overestimated by 0.25 p.p.. The experimental dataset sizes are shown by the dashed vertical lines.

to correct the single move success probability from Figure 6.14. By combining the corrected loss of $\mathcal{L}_{\text{total}} = (10.3 \pm 2.0)\,\%$ with the averaged success probability of $(87.8 \pm 0.6)\,\%$, the final fidelity estimation reads:

$$\mathcal{F} = (97.8 \pm 2.2)\,\%, \tag{6.30}$$

which is consistent with sorting fidelities observed in the literature [39, 142]. Although the uncertainty is larger than the one determined with the first method, we argue that this value is the most reliable, as both $\mathcal{F}$ and $\mathcal{L}$ are determined from a single measurement. From this value it follows that we can create a defect-free array of 16 sites from half filling (nine moves on average, see Figure 6.8) with a probability of $\mathcal{F}^9 = (84 \pm 10)\,\%$ which equals $(15 \pm 6)\,\%$ when including $\mathcal{L}$. However, before possible limitations for the absolute value of $\mathcal{F}$ can be discussed, the large uncertainty has to be reduced by taking more data. In order to do that, it is possible to collect and combine smaller datasets with sizes around ~ 1000 with a re-optimization of the alignment in between.

6.5 Improvements

Before we conclude the chapter, and with that also the technical part of this work, we will now briefly discuss the main bottlenecks and give a technical outlook for future developments.

6.5.1 Sequence duration

As of now, a limitation to the repetition rate of the experiment is the computation and upload time of the RF-waveforms to the AWG card and the camera readout time. The upload and camera readout amount to an additional $\sim 100\,\text{ms}$ which is about 17 % of the full duration. Reducing the sequence by that time would reduce the lifetime losses by a factor of 2.

The camera readout time can be reduced by using the *CameraLink*-interface which increases the bandwidth from 50 MB/s up to 700 MB/s. Additionally, a region of interest can be defined that is applied in hardware and reduces the transferred data payload.

For the computation and upload time, different optimizations are possible. The first is to preupload the full waveforms for all moves. A conservative estimation is that for every target-site 24 (a surrounding 5×5 block) starting sites are possible. For the 6×6 array the total sequence would then consist of 32 ramp-up segments, 16 ramp-down and 16×24 frequency chirps. With the given sampling rate and durations this amounts to data of 1.5 GB which fits into the 4 GB of RAM of the card. During the sequence, the preuploaded segments have to be connected in the correct order which only requires a negligible amount of data to be written.

A second option would be to reduce the amount of samples itself. Picking up the argument from Section 6.3.3, the time resolution of the AODs is around 1 µs. We currently operate the AWG at a sampling rate of 1 228.8 MHz (see Section 4.4) as we are required to sample the 80 MHz wave with sufficient resolution, resulting in a time step of around 10 ns. If the timestep could be increased from 10 ns to 1 µs this would reduce the waveform computation and upload time by a factor of 100 to less than 100 µs. This option requires that the base-frequency is generated by the device itself and only the amplitude- and frequency-changes are supplied, e.g. a DDS-operation. An open question in this case is, if a standalone DDS device can reach the same amplitude and frequency stability as the AWG card.

We recently became aware of a new option for the AWG card available since March 2024, which features a DDS-mode that can be used to program amplitude and frequency ramps without the need to compute the full waveform. If this DDS-mode can be used to realize the complex non-linear frequency chirps, it would not only speed up the sequence but also simplify the software implementation

6.5.2 Algorithm and moves

The algorithm and its implementation are tied to the rearrangement of square arrays investigated in this work. In general, many different geometries and even differently shaped reservoir- and target-regions can be useful [39, 133]. In these cases, different algorithms might be needed to achieve the best sorting fidelity possible. For example, in hexagonal structures longer moves, that do not move across an empty trap could be preferable as the atom or qubit is not heated by the modulated potential.

There are also cases where it could be favorable to move along a non-straight line or multiple atoms at once. For the latter, a superposition of RF-tones can be used to create multiple traps with the AODs [108]. These waveforms have strict requirements for the phases of different tones to prevent interference at reasonable efficiencies. It is therefore difficult to move multiple sites along their connection line as the frequency chirp changes the respective phases and thus the interference. It is however possible to move a chain of traps *perpendicular* to their orientation. This could be useful in the future to sort multiple rows or columns at once, see e.g. [113].

6.5.2.1 Resorting

To increase the probability for obtaining a defect-free array of N sites $P(N)$, it is possible to repeat the rearrangement step one or more times, however, with the current losses this is not feasible. Figure 6.17 shows $P(16)$, given finite losses and fidelities and for different number of defects. Without losses $\lesssim 1\,\%$ and sorting fidelities of $\gtrsim 98\,\%$ obtaining a defect-free array is unlikely, as the global loss dominates.

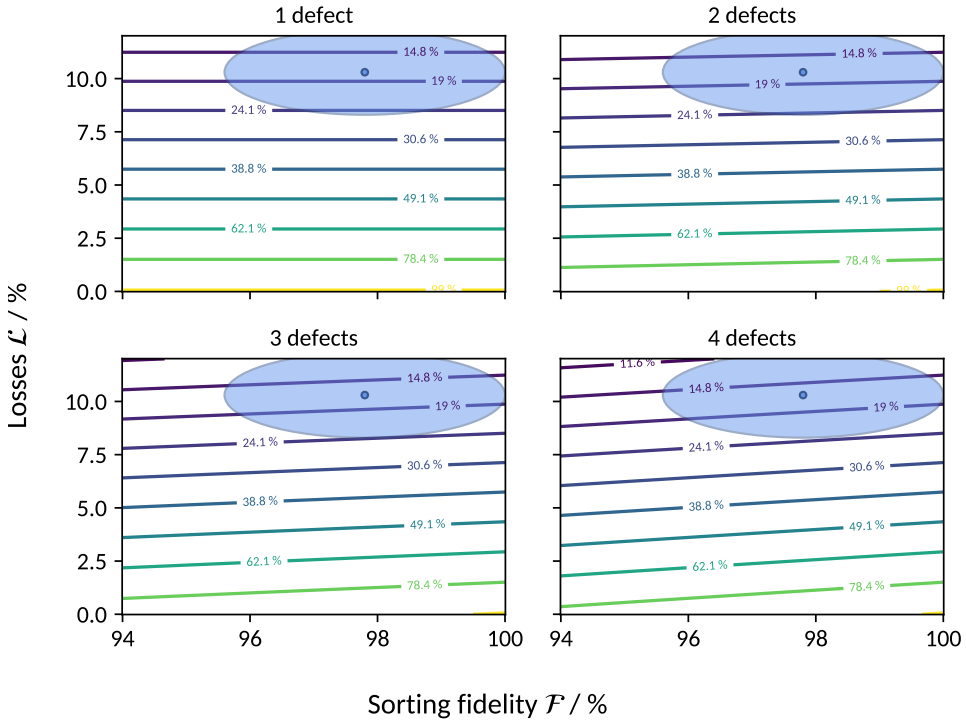

Figure 6.17 – Success probability for obtaining a defect-free arrays of 16 sites after two rearrangement steps, depending on the number of defects given finite losses and sorting fidelity. The dots and ellipses show the current values of $\mathcal{F} = (97.8 \pm 2.2)\,\%$ and $\mathcal{L} = (10.3 \pm 2.0)\,\%$ including the uncertainties.

6.5.2.2 Pattern matching

To increase the number of collision-free solution, it is possible to precompute (parts of) solutions, which can be verified to be collision-free. For arrays with $N \lesssim 30$, it is possible to compute the solutions for *all* configurations. It is not only instructive to do so to evaluate the algorithm with a simulation, but the solutions can also be optimized and used as a lookup-table. One optimization to the compression algorithm that we discussed in this chapter is to always solve the *smallest* configuration of differently transformed arrays (see Table 6.3). This reduces the computational effort by less than the number of transformations, as some transformations produce the same configuration.

We can compute the number of unique configurations using Burnside's lemma:

$$|X/G| = \frac{1}{|G|} \sum_{g \in G} |X^g|, \tag{6.31}$$

where (in colloquial terms) $|X/G|$ is the number of unique configurations of the set X under the action of the group G with $|G|$ elements. The sum runs over all transformations $g \in G = \{I, H, V, R_{180}, T, T_A, R_+, R_-\}$, where $|X^g|$ is the number of elements that are invariant under the transformation g. To determine $|X^g|$ we can look at the following example for a 3×3-array with $g = T$:

$$\begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} \overset{T}{\Rightarrow} \begin{pmatrix} a & d & g \\ b & e & h \\ c & f & i \end{pmatrix} \tag{6.32}$$

For a configuration to be invariant under transposition, $b = d$, $c = g$ and $g = h$ is required, reducing the number of free parameters from 9 to 6. This results in 2^6 configurations that are invariant under the transpose-operation. Extending this analysis to all 8 transformations g results in:

$$|X/G| = \frac{1}{8}\left(2^9 + 2^6 + 2^6 + 2^5 + 2^6 + 2^6 + 2^3 + 2^3\right) = 102 \tag{6.33}$$

unique configurations, which is about a factor of 5 less than the 512 initial number of solutions. We can further exclude configurations which have less than the required number of populated sites and also the ones where all sorting-sites are already populated. Table 6.7 shows the total number of unique configurations for different lattice sizes.

Array size	Initial configurations	Number of unique configurations	Reduction factor
3×3	512	102	~ 5
4×4	65 536	8 552	~ 7.7
5×5	33 554 432	4 211 744	~ 8
6×6	68 719 476 736	8 590 557 312	~ 8

Table 6.7 – Unique configurations example. Not included is the additional reduction due to low occupation or an already fully-filled target region.

Although it is not feasible to compute all solutions for arrays with more than $\gtrsim 30$ sites, the results for smaller arrays can be used to assist with solving larger arrays. Looking at a 6×6 array, we can for example determine the number of configurations for which each quadrant consisting of 9 sites containing at least 4 occupied sites[19]. For the example at hand, the probability of occupying 4 or more sides out of 9, each with independent probability of 50 %, is about 75 % and the probability of every quadrant having 4 or more occupied sites is therefore $0.75^4 = 30\,\%$. If the algorithm detects one of the configurations where this condition is true, we know that it is possible to find a solution without moves between different quadrants. It is therefore possible to create a lookup-table with all possible solutions which can be verified to be collision-free. Whether that is practical depends on different experimental parameters, for example the collision-distance and whether more but shorter moves are preferable.

This train of thought is merely a basic example of a more general approach to the problem making use of *pattern matching*. Enormous progress in the fields of machine learning and neural networks in recent years resulted in tools and frameworks like *tensorflow*[20] which can be used by scientists outside

[19] For arrays with an odd number of sites along one dimension, an asymmetric partition scheme can be chosen.

[20] https://www.tensorflow.org/

of the field of computer science (see [143] as an example). These tools could be used to develop an entirely new approach to the problem, like shown in [144], where the set of moves necessary to create a defect-free array is produced by a trained neural network.

6.6 Conclusion

In this chapter, we presented the single particle control capability of the experiment by moving single ultra-cold atoms between sites in an optical tweezer array. We demonstrated that we can rearrange partially filled arrays into smaller, defect-free arrays with a single move fidelity of $(97.8 \pm 2.2)\,\%$, consistent with the literature [39, 142]. At this level, the large uncertainty makes it difficult to think about and discuss possible limitations for $\mathcal{F}$, however, this can readily be achieved by taking more data. During the course of the experiment, we identified the current bottlenecks and listed possible solutions. With the aforementioned optimizations and a more refined understanding of the atom transfer process, it is possible to increase the repetition rate of the experiment while simultaneously increasing the probability of producing defect-free arrays. Together with a reduction of the imaging loss to less than 1 % as observed in the literature, it will be possible to create large defect-free structures for exciting future research.

CHAPTER 7

Summary and Outlook

This work described the design and characterization of a new optical tweezer setup aimed at preparing, detecting, and manipulating single atoms cooled to their motional ground state.

The first part concentrated on the creation of static and dynamic optical tweezer potentials. We combine a liquid crystal-based spatial light modulator which creates static tweezer arrays, and a pair of crossed acousto-optic deflectors which is used to create a dynamic trap. In combination, the setup is used to prepare and detect atoms in a tweezer array, whereas the dynamic trap is used to address single sites within said array. Our results show, that we can create optical tweezer arrays with homogeneities above 98 % for up to 1 500 tweezers, comparable to existing literature [106, 107]. We further optimized the uniformity with the tweezer depths obtained from spectroscopy with trapped atoms, which we used in a feedback step to ensure the same level of uniformity in the atom plane for up to 60 trapped atoms.

We then discussed the preparation of cold atoms in the tweezer potentials loaded from microkelvin cold atom clouds. We first described the preparation of single atoms in optical tweezers at 515 nm and 813 nm. Using light-assisted collisions, we achieve highly sub-Poissonian atom-number distributions, ensuring each tweezer is occupied by a single atom at most. We then implemented and evaluated different cooling protocols in tweezers at 515 nm, 532 nm, and 813 nm. At these three wavelengths, different relative potential depths for ground and excited states can be created, leading to unique trapping conditions enabling the use and need for different cooling protocols. Using resolved sideband cooling in tweezers at 515 nm, we achieve three-dimensional motional ground-state fractions of 95^{+2}_{-10} % at an aspect ratio of 5.1 ± 0.1, which improves on existing implementations [46, 47] and is close to the theoretical limit. At 813 nm, we cooled atoms using a sisyphus cooling scheme, and we confirm temperatures near the motional ground state by comparison to a classical simulation, similar to results documented in existing research [48]. Additionally, we find that releasing an atom from one trap followed by a capture in another trap can be used to determine temperatures in tweezer arrays without relying on a simulation, offering potential improvements in thermometry accuracy. Our cooling measurements at 532 nm are, to our current knowledge, the first for strontium in high-NA optical tweezers, yielding results comparable to other wavelengths [46, 127]. Our measurements indicate 532 nm as a promising tweezer wavelength for future research, and readily available lasers and optical components at this common wavelength can reduce complexity and cost of the experimental setup.

Finally, we demonstrated single-particle control by assembling defect-free square arrays of 16 sites with separations of 5.2 µm from half-filling using an atom-by-atom approach. To realize the atom rearrangement, we first developed a control-theory optimal trajectory model to reliably move single

atoms between different sites. We implemented an existing algorithm to compute the required atom moves, which we further improved to reduce excess atom moves by 40 % and to increase the probability of finding the solution of minimal length by 40 %. We assessed the success probability of single-atom moves, the *sorting fidelity*, using two methods: a correction of a direct measurement and a novel statistical analysis technique. By combining both methods, we minimize the uncertainty and determine a single-move fidelity of (97.8 ± 2.2) %, on par with current state-of-the-art implementations [39, 142]. This imposes a theoretical limit for obtaining a defect-free array of 16 sites of (84 ± 10) %.

Outlook

In this work, the first step for building a quantum simulation platform — the state preparation — has been realized. Now, interactions between atoms are necessary to create multi-qubit entangled states by building on the basis of isolated single atoms trapped in an optical tweezer array. Collisions between atoms affecting the internal state are largely eliminated, which leads to strong isolation and long coherence times, however, it also means that interactions need to be artificially introduced into the system. Different methods to do so exist and we now sketch two possibilities in the form of Rydberg interactions and interactions mediated by an optical cavity.

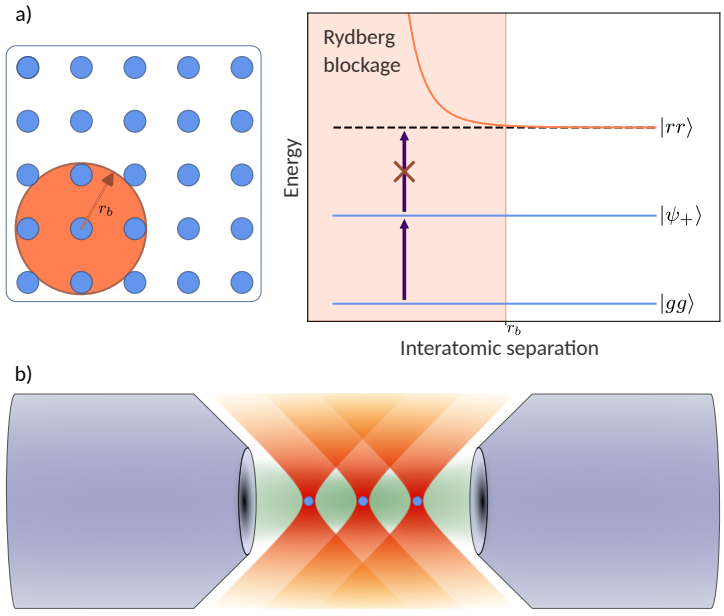

Figure 7.1 – Two examples of effects that can create multi-qubit entangled states. **(a)** The interaction between two Rydberg atoms separated by less than r_b can be used to create an entangled state $|\psi_+\rangle = (|gr\rangle + |rg\rangle)/\sqrt{2}$. For more than two atoms, multi-qubit entangled states can be created. **(b)** The coupling of trapped atoms to the mode of a high-finesse optical cavity can be used to create multi-qubit entangled states.

Rydberg interactions arise from van der Waals interactions between two Rydberg states, where the latter are highly excited atomic states with principal quantum numbers $n \gtrsim 50$ [145, 146]. The interaction distance of Rydberg states can reach several micrometers, thus bridging the gap between adjacent sites in an optical tweezer array. Rydberg states have lifetimes of tens to thousands of microseconds, depending on the Rydberg-level and the decay strength due to radiative processes or blackbody induced

transitions. On the other hand, the energy scale of the interaction is on the order of tens to hundreds of megahertz, depending on the interatomic separation. The use of Rydberg interactions to create two qubit entanglement with neutral atoms has been investigated theoretically [147] and has been applied with great success since [41, 52, 148, 149]. To create an entangled state, the strong dipole-dipole interaction is used to prevent excitation of adjacent ground-state atoms to a collective Rydberg state, by shifting the latter out of resonance (see Figure 7.1(a)).

The excitation of strontium atoms to Rydberg states with $n > 50$ requires wavelengths around 320 nm. Furthermore, Rydberg states are sensitive to perturbation by external electric fields, caused for example by accumulated charges on the vacuum windows. To incorporate such an UV laser into the setup, the coatings of the vacuum cell windows are transmissive at this wavelength, rendering the integration technically possible. To compensate electric stray fields, the setup also includes two electrodes which can be used for compensation. Choosing the path of Rydberg physics requires to add at least one excitation laser but multi-stage excitation schemes have also been used [150]. An open question for this path is the minimum reachable spacing between traps, to increase entanglement fidelity and to (possibly) also create larger entangled states.

Another, less explored option is to use cavity-mediated interactions to realize all-to-all connectivity between multiple qubits. Strong light-matter couplings between a single emitter and the mode of a high-finesse optical cavity have been investigated theoretically in terms of the *Jaynes-Cummings* model [151] which have been realized with ions [152] and neutral atoms [45, 153, 154]. It was shown, that, compared to the Rydberg approach, cavity-mediated interactions allow to create entangled states of the whole system with a single gate operation [155, 156]. To pursue this path, a cavity has to be built and installed in the vacuum chamber. A key requirement in this case is to reach the *strong-coupling*-regime where the cavity-atom coupling g dominates the spontaneous decay γ and the cavity lifetime κ. This requires a small mode volume V with waists on the order of $\lesssim 10\,\mu\text{m}$ and cavity lengths around 1 mm. An experimental challenge lies in the shape of the mirrors itself, to allow for the projection of high NA tweezers into the cavity mode (see Figure 7.1(b)). To ensure a constant coupling along the cavity axis, a near planar cavity configuration is favorable, but mirrors satisfying these geometric requirements are not commercially available and have to be custom-designed and -manufactured. An additional open question is the precision and stability of the tweezer alignment relative to the cavity mode, and it could be necessary to add an optical lattice along the cavity axis to ensure a constant cavity coupling.

Using either Rydberg- or cavity mediated-interactions, it will be possible to create multi-particle entanglement, which enables the new experiment to conduct exciting studies at the frontiers of experimental quantum simulation research.

APPENDIX A

Fourier Optics

This chapter is a compilation of aspects of the Fourier-theory of optics that is required for parts of the tweezer array generation for this work. It recaps the derivation of the *Fourier*-lens and shows the approximations made along the way. The first part is in essence a summary of parts of chapters three, four and six of [90] which are relevant for this thesis. The second part is concerned with methods for performing numerical computations in the field of Fourier optics.

A.1 Huygens-Fresnel principle

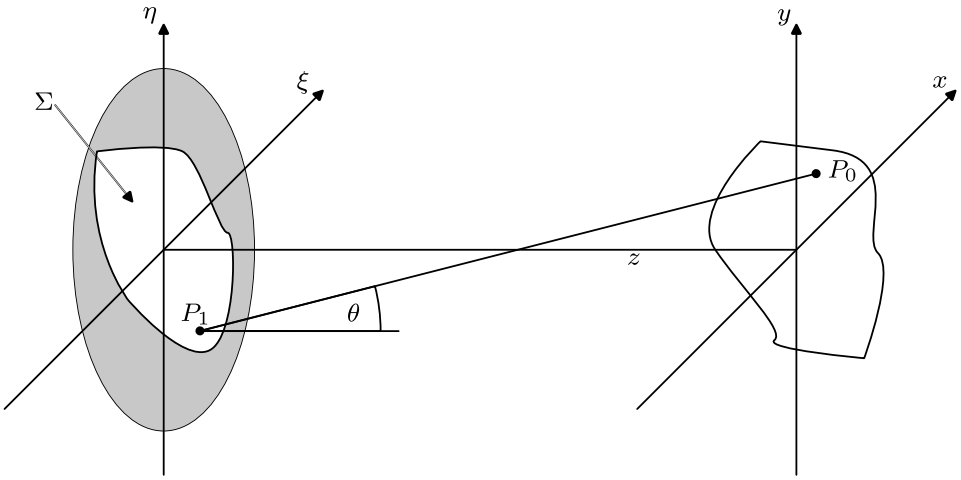

Figure A.1 – A screen in the $\xi\eta$-plane is illuminated from the left. We would like to calculate the electric field distribution in the xy-plane.

Figure A.1 shows an optical setup where a screen ($\xi\eta$-plane) is illuminated from the left and we would like to compute the beam amplitude at a position P_0 in the xy-plane. The Huygens-Fresnel (sometimes just Huygens) principle, states that every point of a wavefront is itself the starting point of a new spherical wave. The field at position P_0 is therefore given by the convolution of the initial field $U(P_1)$ at a point P_1 in the $\xi\eta$-plane, and the propagated field of all elemetary waves progagated from P_1 to P_0, for all

points P_1 inside the aperture Σ. Expressing the explanation mathematically leads to:

$$U_I(P_0) = -\frac{1}{2\pi}\int\int_{\Sigma} U(P_1)\left(ik - \frac{1}{r_{01}}\right)\frac{\exp\left(ikr_{01}\right)}{r_{01}}\cos\left(\vec{n},\vec{r}_{01}\right)\mathrm{d}s \tag{A.1}$$

$$U_I(x,y,z) = h(x,y,z) * U(x,y,0) \tag{A.2}$$

$$h(x,y,z) = \frac{1}{2\pi}\frac{z}{r}\left(\frac{1}{r} - ik\right)\frac{\exp\left(ikr\right)}{r}, \tag{A.3}$$

where $\vec{n}$ is the outward normal vector in negative $z-$direction and $\vec{r}_{01}$ points from P_0 to P_1 and $r = \sqrt{x^2+y^2+z^2}$. $k = 2\pi/\lambda$ is the wavevector of the light with wavelength λ and $\cos\left(\vec{n},\vec{r}_{01}\right) = z/r_{01}$ with $r_{01} = \sqrt{\left(x_1-x_0\right)^2 + \left(y_1-y_0\right)^2 + \left(z_1-z_0\right)^2_{01}}$. Assuming that the separation between the two points in the two planes is large compared to the wavelength, i.e. $r_{01} \gg \lambda$ we can simplify Equation (A.1) further to:

$$U_I(P_0) = \frac{1}{i\lambda}\int\int_{\Sigma} U(P_1)\frac{\exp\left(ikr_{01}\right)}{r_{01}}\cos\left(\theta\right)\mathrm{d}s, \tag{A.4}$$

$$\tag{A.5}$$

where $\cos\left(\theta\right)$ is given by z/r_{01}. We thus arrive at:

$$U(x,y) = \frac{z}{i\lambda}\int\int_{\Sigma} U(\xi,\eta)\frac{\exp\left(ikr_{01}\right)}{r_{01}^2}\mathrm{d}\xi\mathrm{d}\eta, \tag{A.6}$$

where ξ and η are the coordinates describing the position in the aperture plane and x and y describe the position in the image plane.

In addition to the approximations made for the scalar diffraction theory itself and $r_{01} \gg \lambda$ we will now additionally make approximations for $r_{01} = \sqrt{z^2 + (x-\xi)^2 + (y-\eta)^2}$

A.1.1 Fresnel

The Fresnel approximation assumes that the separation between the two planes z is much larger than the changes in radial coordinates. We can therefore approximate r_{01} as:

$$r_{01} = \sqrt{z^2 + (x-\xi)^2 + (y-\eta)^2} \tag{A.7}$$

$$= z\sqrt{1 + \left(\frac{x-\xi}{z}\right)^2 + \left(\frac{y-\eta}{z}\right)^2} \tag{A.8}$$

$$\approx z\left[1 + \frac{1}{2}\left(\frac{x-\xi}{z}\right)^2 + \frac{1}{2}\left(\frac{y-\eta}{z}\right)^2\right]. \tag{A.9}$$

in Equation (A.6) the r_{01}^2 in the denominator can safely be approximated by z^2 but for the exponential have to keep the quadratic terms, leading to:

$$U(x,y,z) = \int_{-\infty}^{\infty}\int U(\xi,\eta,0)\exp\left(j\frac{k}{2z}\left[(x-\xi)^2+(y-\eta)^2\right]\right)\mathrm{d}\xi\mathrm{d}\eta \tag{A.10}$$

which is a convolution of the field $U(\xi,\eta,0)$ with $h(x-\xi,y-\eta)$ with the kernel $h(x,y)$:

$$h(x,y) = \frac{e^{ikz}}{i\lambda z}\exp\left[\frac{ik}{2z}\left(x^2+y^2\right)\right] \tag{A.11}$$

A.1.2 Fraunhofer

The stronger Fraunhofer approximation:

$$z \gg \frac{k\left(\xi^2+\eta^2\right)_{\max}}{2} \tag{A.12}$$

simplifies Eq. (A.10) further by omitting the quadratic terms of ξ and η:

$$U(x,y,z) = \frac{e^{ikz}e^{i\frac{k}{2z}\left(x^2+y^2\right)}}{i\lambda z}\int_{-\infty}^{\infty}U(\xi,\eta,0)\exp\left[-i\frac{2\pi}{\lambda z}\left(x\xi+y\eta\right)\right]\mathrm{d}\xi\mathrm{d}\eta \tag{A.13}$$

which is a Fourier transform with the spatial frequencies $f_X = x/\lambda z$ and $f_Y = y/\lambda z$. However, to fulfill Eq. (A.12) at optical frequencies and apertures around 1 cm requires z to be larger than several hundred meters. The approximation is therefore also referred to as *far-field*-approximation (compared to *near-field* for the Fresnel approximation). To still be able to work in this regime, we will move one step further and add an additional optical element to our setup

A.1.3 The Fourier lens

Let us assume that the $\xi\eta$-plane has the following transparency function:

$$t_l(\xi,\eta) = \exp\left[-i\frac{\pi}{\lambda f}\left(\xi^2+\eta^2\right)\right] \tag{A.14}$$

which we will call a "lens". We will additionally limit its spatial extent by the pupil function

$$P(x,y) = \begin{cases}1 & \text{inside lens aperture}\\ 0 & \text{else}\end{cases}. \tag{A.15}$$

The field inside the aperture is then given by:

$$U_l'(\xi,\eta) = U_l(\xi,\eta)P(\xi,\eta)\exp\left[-i\frac{k}{2f}\left(\xi^2+\eta^2\right)\right], \tag{A.16}$$

where $U_l(\xi,\eta) = A\, t_A(\xi,\eta)$ describes a normally incident, monochromatic plane wave with amplitude A. Inserting Eq. (A.16) into Eq. (A.10) to calculate the field in the back-focal plane of the lens (thus $z = f$) results in:

$$U_f(x,y) = \frac{\exp\left[i\frac{k}{2f}\left(x^2+y^2\right)\right]}{i\lambda f} \times \int_{-\infty}^{\infty} U_l(\xi,\eta)P(\xi,\eta)\exp\left[-i\frac{2\pi}{\lambda f}(x\xi+y\eta)\right]\mathrm{d}\xi\mathrm{d}\eta \tag{A.17}$$

$$\frac{\exp\left[i\frac{k}{2f}\left(x^2+y^2\right)\right]}{i\lambda f} \times \int_{-\infty}^{\infty} U_l(\xi,\eta)P(\xi,\eta)\exp\left[-i\frac{2\pi}{\lambda f}(x\xi+y\eta)\right]\mathrm{d}\xi\mathrm{d}\eta \tag{A.18}$$

which is again the two-dimensional Fourier transformation of our aperture function $U_l(\xi,\eta)P(\xi,\eta)$ (up to a phase factor). We further define the spatial frequencies k_X and k_Y as:

$$\begin{aligned} k_X &= \frac{2\pi}{\lambda f}x \\ k_Y &= \frac{2\pi}{\lambda f}y. \end{aligned} \tag{A.19}$$

A lens thus performs a Fourier transformation so that we get the Fraunhofer diffraction pattern at $z = f$ *without* explicitly fulfilling Eq. (A.12). The condition we have to fulfill is still $r_{01} \gg \lambda$.

Using this result, it is now possible to relate the amplitude distribution in the front-focal plane of a lens to the amplitude field distribution in the back-focal plane, as the two are related by a Fourier transformation.

A.1.4 Angle to position conversion

A known relation of Fourier transformations is:

$$\mathcal{F}\left\{e^{iax}\right\} = \delta\left(k-a\right), \tag{A.20}$$

which relates a phase gradient with slope a to a displacement of a in the plane of the Fourier conjugate variable k[1]. In the case of a wavefront that travels to the right (z-direction) the phase slope in the $\xi\eta$-plane is given by:

$$\exp\left(ia\xi\right) = \exp\left(i\left(2\pi\frac{\tan\alpha}{\lambda}\right)\xi\right), \tag{A.21}$$

where α is the tilt of the wavefront with respect to the ξ-axis in this case, shown in Figure A.2. This leads to a displacement of $-\frac{2\pi\tan\alpha}{\lambda}$ along the k_X direction in the k_Xk_Y-plane. To convert the latter to

[1] For finite extends the input is multiplied by a rectangle function which is Fourier transformed to a sinc shaped peak at position Δx.

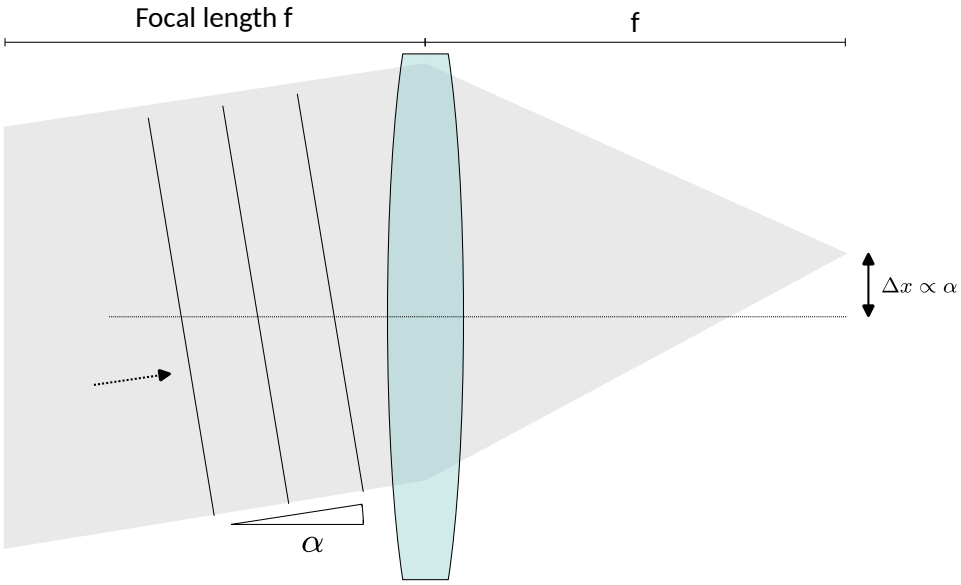

Figure A.2 – Angle to position correspondence of a Fourier lens

coordinates in the xy-plane, we use Equation (A.19) and multiply by $\lambda f/2\pi$ to arrive at

$$|\Delta x| = f \tan\alpha \overset{\alpha \ll 1}{\approx} f\alpha. \tag{A.22}$$

A.1.5 Gauss beam transformation

Using the Fourier-transforming properties of a lens, we can also calculate the waist-size w_0 of a focussed Gaussian beam, by relating it to the beam waist w_1 in the other focal plane of a lens with focal length f. Assuming a Gaussian beam in the front focal plane given by[2]:

$$E(\xi) \propto \exp\left(-\xi^2/w_1^2\right). \tag{A.23}$$

Transforming into k_X-space:

$$E(\xi) = \mathcal{F}\{E(\xi)\} \propto \exp\left(-\frac{1}{4}k_X^2 w_1^2\right). \tag{A.24}$$

[2] We describe the problem in one dimension, as the two dimensional case separates, because of the unique properties of the Gaussian function

Making the substitution $k_X \rightarrow \frac{2\pi}{\lambda f} x$ from Equation (A.19) and comparing to a Gaussian beam $\exp\left(-x^2/w_0^2\right)$ we find:

$$-\frac{1}{4}k_x^2 w_1^2 = -\frac{\pi^2 w_1^2}{f^2\lambda^2}x^2 \stackrel{!}{=} -x^2/w_0^2 \tag{A.25}$$

$$\Rightarrow w_0 = \frac{f\lambda}{\pi w_1}. \tag{A.26}$$

This result shows that the waists of the initial and final field are inversely proportional, i.e. a larger diameter beam can be focused to a tighter spot.

A.2 Numerical computations

Equation (A.18) gives us all we need to know to compute the electric field distribution in the back focal plane of a lens, given the electric field distribution in the front focal plane of a lens. The computation can, however, in general be quite involved and often analytic results cannot be obtained. Luckily, numerical Fourier transformations are among the most well optimized algorithms due to their applications in signal processing with examples of audio or image compression or filtering. One of these algorithms is the fast Fourier transformation (FFT), discovered in 1965 by Cooley and Tukey which efficiently computes the discrete Fourier transformation. FFT implementations are available in many programming languages and for this work we use the implementation from the *CUDA toolkit*[3] to quickly compute large two-dimensional FFTs.

A.2.1 Fourier frequencies

In this work we mainly deal with spatial Fourier transformations where the spatial coordinates in the Fourier plane ξ and η are transformed to spatial frequencies k_X and k_Y. Assuming a one dimensional array of complex values, representing an electric field distribution in the front-focal plane of a lens, we compute the FFT and need to relate the output to the physical world. An FFT of N values, leads to an output N values. The output at index i corresponds to a frequencies f_i, where f is ordered as follows:

$$k_X = [0, 1, \cdots, N/2-1, -N/2, \cdots, -1]\,/(\Delta\xi\ N) \text{ if } N \text{ even} \tag{A.27}$$

$$k_X = [0, 1, \cdots, (N-1)/2, -(N-1)/2, \cdots, -1]\,/(\Delta\xi\ N) \text{ if } N \text{ odd} \tag{A.28}$$

where $\Delta\xi$ is the discrete spacing of the Fourier plane array. To finally compute the spatial resolution in the image plane Δx, we use Equation (A.19), and arrive at:

$$\Delta x = \frac{\lambda f}{2\pi}\frac{1}{\Delta\xi\ N}. \tag{A.29}$$

From this result we can see the resolution in the image plane is inverse proportional to the size of the sampling space $\Delta\xi\ N$ in Fourier space. This is well-known fact for Fourier transformations with fundamental application in physics, for example in form of the Heisenberg uncertainty principle. In the

[3] https://developer.nvidia.com/cuda-toolkit

context at hand, it lays the foundation for a numerical technique to increase the resolution that we will discuss next.

A.2.2 Zero padding

Zero padding is a technique to increase the resolution in the space of the Fourier conjugate by increasing the sample size of the signal. It is a commonly used in digital signal processing and it can be intuitively understood for the time and frequency relation. Inside a certain time interval, we can only make a statement about differences in frequencies, where the frequency resolution gets smaller the longer a signal is monitored. The inverse is also true and the base of the Nyquist theorem: The larger the frequency space is, i.e. if more high-frequencies are included, the higher the time resolution gets. This effect is visible for example at TTL-pulses where finite bandwidths lead to over- or undershoots at the edges. In terms of spatial extend and spatial frequencies the same holds true. It is the reason for the fact that laser beams with larger diameters can be focussed to smaller foci and also why for astronomical telescopes are so large compared to to terrestrial ones, to name two examples.

Imagine a rectangular aperture illuminated by a plane wave (the $\xi = \infty$-limit in Section 4.2), as shown in Figure A.3(a). In the case of a finite aperture, the electric field is zero outside of it, but it still contributes to the resolution in Fourier space. To calculate the shape of the focus when the aperture is imaged with a lens, we have to take the Fourier transformation.[4] Depending on the electric field distribution and the shape of the aperture, this is in general not possible to be done analytically and we have to resort to numerical methods. In this case zero padding helps to replicate the analytical result accurately. This ensure that the image plane resolution is not limited by the computation but by the real physical constraints, for example the size of an aperture. Figure A.3(a) shows the aperture and Figure A.3(b) shows the analytical calculation of the focus along numerical results with different levels of zero padding. The trade-off one has to make is Fourier space resolution compared to the computation time.

[4] For simplicity we will consider the one-dimensional case here but the results apply to the two-dimensional case as well.

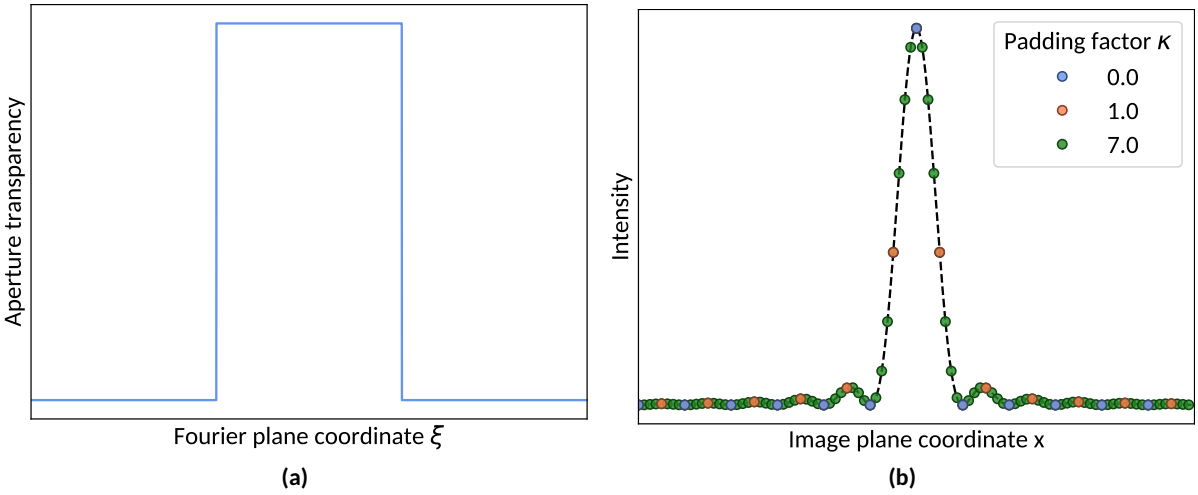

Figure A.3 – One-dimensional aperture **(a)** which gives a sinc^2 power spectrum in the image plane **(b)** when illuminated by a plane wave. The aperture is sampled at N points inside the aperture and at $N\kappa$ points outside, leading to a total of $N\,(\kappa+1)$ samples. The analytical result (dashed line) is successively better approximated by larger values of κ.

A.3 Blazed grating diffraction efficiency

To gain amplitude and phase control with the SLM, we have to move the desired pattern away from the optical axis. This can be achieved by adding a *blazed* phase grating, i.e. a phase slope in Fourier space, as described in Appendix A.1.4. Given the discrete nature of the sensor of our SLM we cannot apply the slope continuously but rather in a discrete form. The deviation from the true phase slope becomes apparent when the full phase range from 0 to 2π has to be sampled by roughly less than 10 px. We will shortly see where this claim is coming from and why we need this large phase slope.

Let us now take a closer look at a single $0-2\pi$ section of the grating. To simplify the calculation and visualizations we will only look at integer numbers of phase steps N. We also assume equality between the pixel width and the pixel pitch, i.e. a fill factor of 100 %. Lower fill factors reduce the overall diffraction efficiency as the pixels itself form a *binary* grating. The phase function of the grating is given by:

$$\phi_m(x) = \sum_{n=0}^{N-1} \frac{2\pi}{N} n \cdot \phi_{s,n}(x) \tag{A.30}$$

$$\phi_s(x) = \mathrm{rect}\left(\frac{x-(w/2+n\cdot w+L\cdot m)}{w}\right), \tag{A.31}$$

where w is the width (or pitch) of a single pixel and $L = N \cdot w$ is the resulting period of the grating. The transparency function of the grating is then given by:

$$g(x) = \exp\left(-i \sum_{m\in\text{aperture}} \phi_m(x)\right). \tag{A.32}$$

Figure A.4 shows a blazed grating sampled with $N = 8$ phase steps and and aperture size of $M = 4$. It

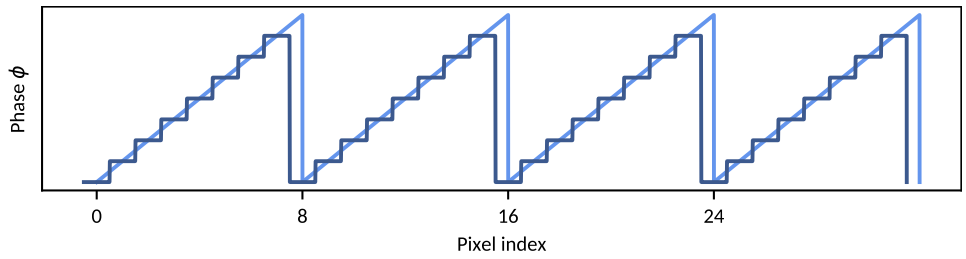

Figure A.4 – Approximation of a continous *blazed* grating with eight phase steps

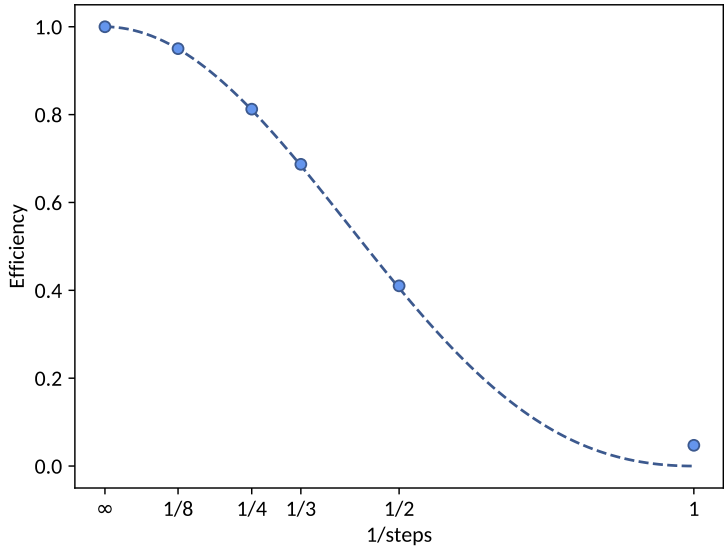

Figure A.5 – Diffraction efficiency of a blazed grating into the first order for different numbers of phase steps N. The dashed line shows the analytical result from Equation (A.33) whereas the data points show the result of a numerical Fourier transformation.

does not matter that we only approximate the phase slope from below as we could easily add a constant phase offset to e.g. center the plateaus on the slope.

The diffraction efficiency can be calculated from the Fourier transformation of the transparency function as [90]:

$$\epsilon\left(k\right) = \mathrm{sinc}^2\left(\frac{k}{N}\right)\left[\frac{\mathrm{sinc}\left(k-1\right)}{\mathrm{sinc}\left(\frac{k-1}{N}\right)}\right]^2 \tag{A.33}$$

A binary grating is a special case of a blazed grating with only two phase steps. In this case the maximum depth is set to π in Equation (A.30). If we substitute the factor $\frac{2\pi}{N}n$ by a variable depth $\Delta\phi$, we can calculate the diffraction efficiency depending on the grating depth. The result is shown in Figure A.6 and it can be seen that increasing the depth $\Delta\phi$ increases the intensity in the first order. At $\Delta\phi = \pi$ the maximum efficiency from Figure A.5 is reached. After that the relative grating depth decreases again which reduces the diffraction efficiency.

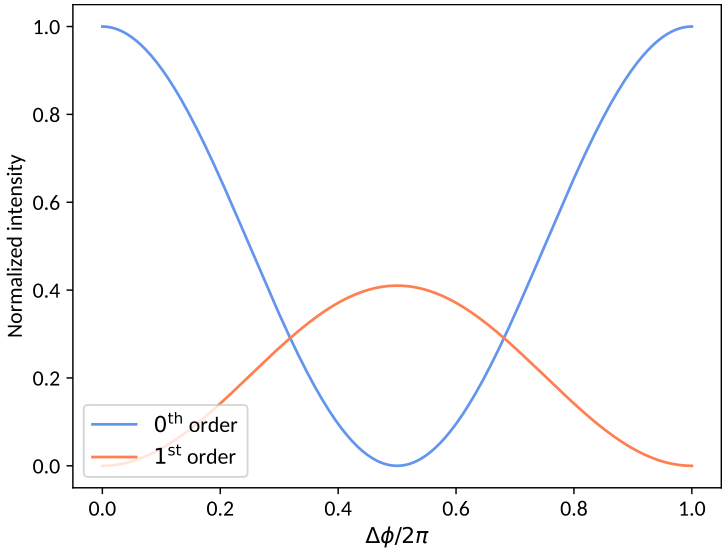

Figure A.6 – Binary grating diffraction efficiency for different depths of the grating $\Delta\phi$.

A.4 Optical aberrations and Zernike polynomials

To fine tune the position of a static array created with the SLM, we use two effects in this work. First, a linear phase gradient which, when taken modulo 2π, has the same effect as a blazed grating. Secondly, a parabolic phase slope which changes the axial position of the projected beam, similar to the Fourier lens.

Both effects can be seen as elements of a set of known wavefront-changes, described by the *Zernike polynomials* (see e.g. [157]). Zernike polynomials $Z_n^m(\rho, \phi)$ are a sequence of orthogonal polymials on the unit disk and they can be used to describe optical aberrations.

$$Z_n^m(\rho, \phi) = R_n^m(\rho, \phi) \cos(m\phi) \tag{A.34}$$

$$Z_n^{-m}(\rho, \phi) = R_n^m(\rho, \phi) \sin(m\phi) \tag{A.35}$$

with the radial polynomial R_n^m:

$$R_n^m(\rho) = \sum_{k=0}^{\frac{n-m}{2}} \frac{(-1)^k (n-k)!}{k! \left(\frac{n+m}{2} - k\right)! \left(\frac{n-m}{2} - k\right)!} \rho^{n-2k} \tag{A.36}$$

if $n - m$ even, else $R_n^m(\rho) = 0$. Also $R_n^m(1) = 1$. The important first order effects that have applications in optics are shown in Table A.1 and visualized in Figure A.7.

n	m	Z_n^m	Name
0	0	1	Piston
1	-1	$\rho \sin \phi$	Y-tilt
1	1	$\rho \cos \phi$	X-tilt
2	-2	$\rho^2 \sin 2\phi$	Oblique astigmatism
2	0	$2\rho^2 - 1$	Defocus
2	2	$\rho^2 \cos 2\phi$	Vertical astigmatism
3	-1	$\left(3\rho^3 - 2\rho\right) \sin \phi$	Vertical coma
3	1	$\left(3\rho^3 - 2\rho\right) \cos \phi$	Horizontal coma
4	0	$6\rho^4 - 6\rho^2 + 1$	Primary spherical

Table A.1 – Common Zeernike polynomials used to describe optical systems. Astigmatism, defocus and coma might be known to those who had to visit an ophthalmologist at some point.

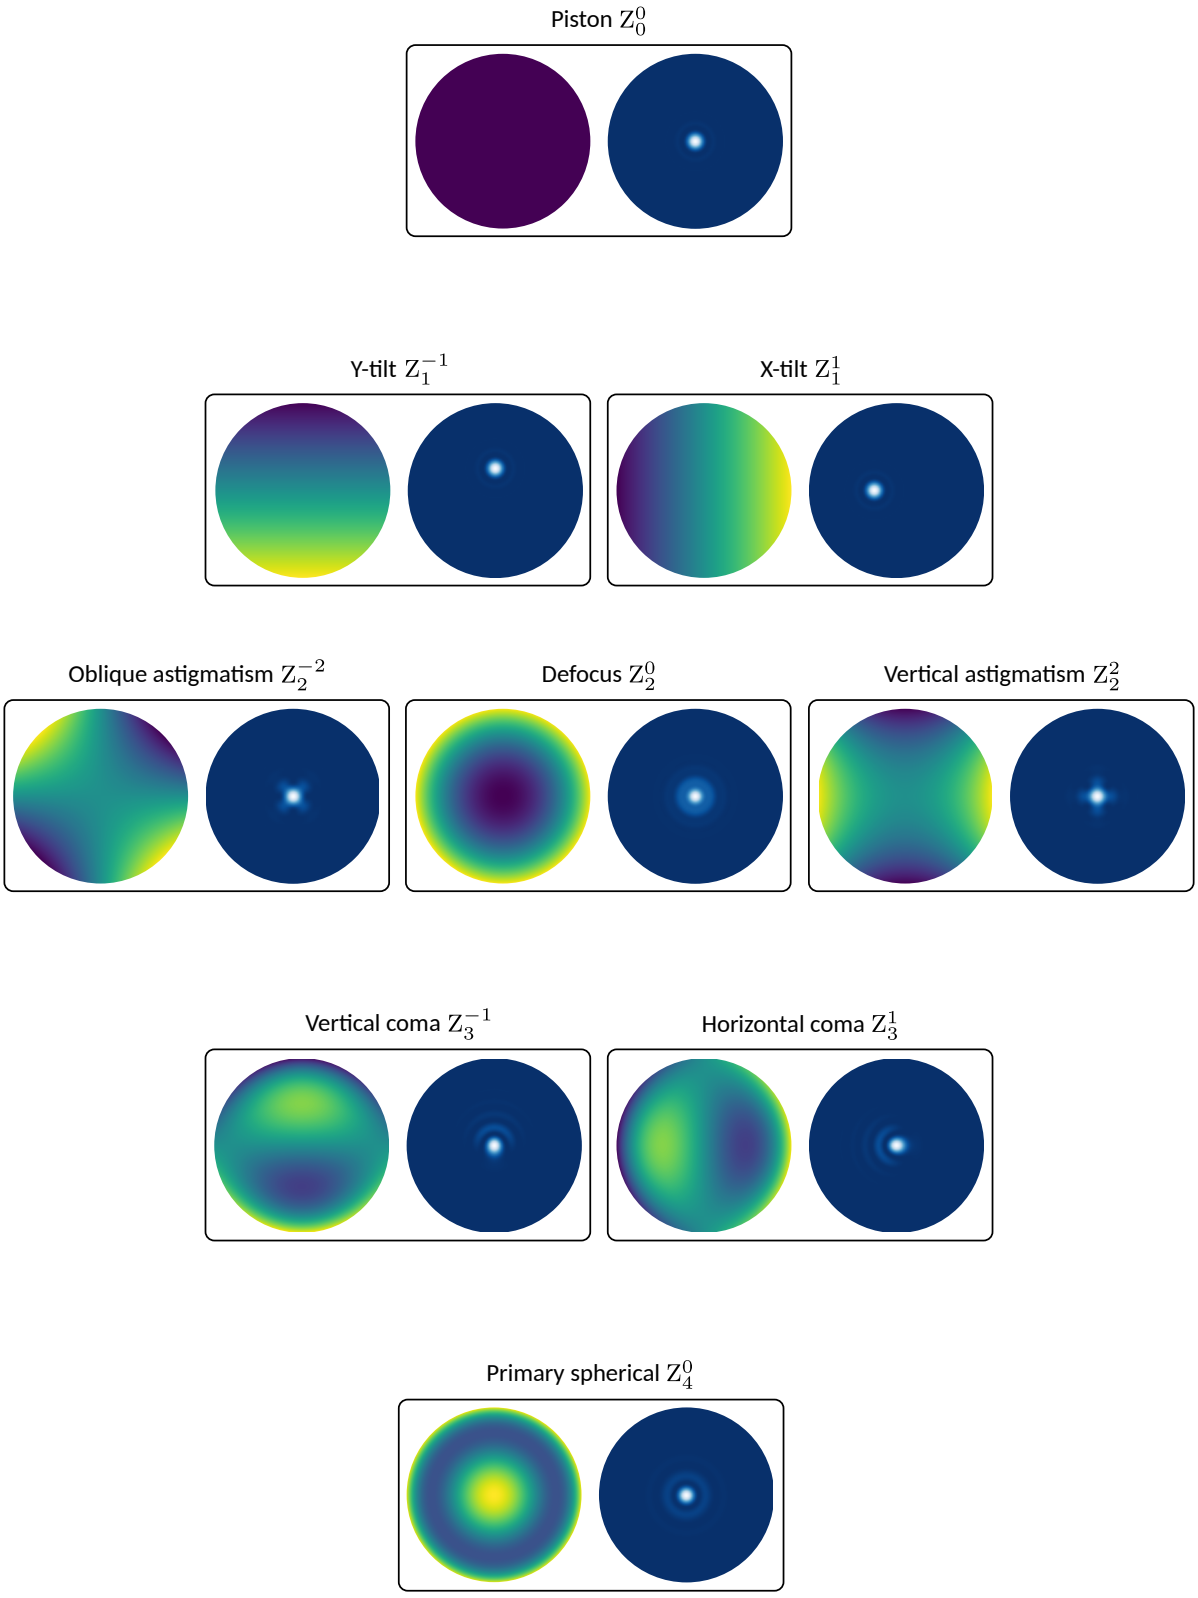

Figure A.7 – Zernike polynomials and their effect on the shape of a point-spread function.

APPENDIX B

Tweezer alignment

B.1 Objective to vacuum chamber

To minimize aberrations due to misalignment, it is important to enter the objective exactly on the optical axis. For this we use a beam at 515 nm and look at the reflections of the different surfaces inside the objective. Because the objective is built out of several lenses, there are multiple circular etalon-like patterns that need to be overlapped. Using light at 515 nm, which is close to the maximum sensitivity of the human eye, makes it possible to look at these rings from the backside of a sheet of paper, where the incident beam is passing through a hole. Using a pair of adjustable mirrors we can then optimize the coupling into the objective. Once the optimum is reached the concentric rings align along a string along the axis of the beam. The visibility crucially depends on the size of the coupling beam where a tradeoff between the size of the rings and locality[1] has to be made. We vary the beam size with an iris aperture to fine tune the contrast during alignment.

We start the alignment process for both objectives with the lower one, without the upper one installed. Compared to the upper one, which is also adjustable in its position, only the angle with respect to the vacuum window can be changed. After an initial visual alignment, we shine in a beam from the top such that we see the reflections of the windows and the interference pattern from the lenses inside the objectives. The incident beam, the reflections and the circular pattern are iteratively adjusted until they fully overlap. This ensures that the objective is aligned perpendicular to the vacuum windows and that the beam is exactly on-axis.

The beam will now serve as a reference to install the upper objective. Again the incident beam and the interference patterns are overlapped but this time by moving the upper objective and not the beam. This fact and the issue that we cannot block the back-reflections from the lower objective make the alignment of the top objective much more complicated and tedious. The measurements in Chapter 5 suggest that further references and alignment techniques are necessary to ensure a high quality alignment.

[1] Locality describes the extent of the objective aperture that is exposed to the coupling beam. A smaller coupling beam shows smaller rings but is more sensitive to the rough alignment of the coupling beam to the objective axis. A larger beam on the other hand produces an interference pattern that is more sensitive to small deviations in alignment but more difficult to read due to its complex structure.

B.2 Imaging path

To initially adjust the imaging beam path of the EMCCD camera such that a signal from the tweezers can be picked up, we use a beam at 515 nm or 532 nm that we align to the optical axis of the lower objective with the procedure above. We then place a retroreflecting mirror in the tweezer path and overlap the reflection with the incident beam. The dichroic longpass mirror[2] reflects about 5 % of the light to the camera which, for 515 nm and 532 nm, is well detectable for the naked eye. We align the reflected beam to the aperture of the camera, while making sure the internal shutter is still closed. Only after lowering the beam intensity as much as possible, we open the shutter and observe the signal. We center the beam on the chip using the mirrors in the imaging path and can therefore assume that the radial alignment is good enough to capture the fluorescence from the atoms in the imaging plane. Due to the focal length at 461 nm being 50 µm shorter than the focal length at 515 nm/532 nm the position of the imaging lens deviates from the optimal position. The magnification of 750 mm/13.5 mm = 55 in radial direction results in an axial magnification of $55^2 = 3000$. The difference of 50 µm thus translates into a difference of 15 cm and the camera has to be placed around 60 cm after the imaging lens. This calculation matches what we see in the experiment and we fine tune the axial position with the micrometer stage mounted to the camera.

[2] Thorlabs DMLP490L

Bibliography

[1] E. Schrödinger, British Journal of Philosophy of Sciences **3** (1952).

[2] T. Hänsch and A. Schawlow, *Cooling of gases by laser radiation*, Optics Communications **13** (1975) 68, ISSN: 0030-4018.

[3] D. Wineland and H. Dehmelt, *Proposed 1014δv/v laser fluorescence spectroscopy on Tl+ mono-ion oscillator*, Bull. Am. Phys. Soc **20** (1975).

[4] S. Haroche and J.-M. Raimond, *Exploring the Quantum: Atoms, Cavities, and Photons*, Oxford University Press, 2006, ISBN: 9780198509141.

[5] R. P. Feynman, *Simulating physics with computers*, International Journal of Theoretical Physics **21** (1982) 467, ISSN: 1572-9575.

[6] J. I. Cirac and P. Zoller, *Goals and opportunities in quantum simulation*, Nature Physics **8** (2012) 264, ISSN: 1745-2481.

[7] I. M. Georgescu, S. Ashhab and F. Nori, *Quantum simulation*, Rev. Mod. Phys. **86** (1 2014) 153.

[8] I. Buluta and F. Nori, *Quantum Simulators*, Science **326** (2009) 108.

[9] A. J. Daley et al., *Practical quantum advantage in quantum simulation*, Nature **607** (2022) 667, ISSN: 1476-4687.

[10] A. Montanaro, *Quantum algorithms: an overview*, npj Quantum Information **2** (2016) 15023, ISSN: 2056-6387.

[11] F. Arute et al., *Quantum supremacy using a programmable superconducting processor*, Nature **574** (2019) 505, ISSN: 1476-4687.

[12] J. Preskill, *Quantum Computing in the NISQ era and beyond*, Quantum **2** (2018) 79, ISSN: 2521-327X.

[13] J. Preskill, *Quantum computing 40 years later*, 2023, arXiv: 2106.10522 [quant-ph].

[14] M. Gall, N. Wurz, J. Samland, C. F. Chan and M. Köhl, *Competing magnetic orders in a bilayer Hubbard model with ultracold atoms*, Nature **589** (2021) 40, ISSN: 1476-4687.

[15] 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** (2002) 39, ISSN: 1476-4687.

[16] 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** (2009) 74, ISSN: 1476-4687.

[17] C. Gross and I. Bloch, *Quantum simulations with ultracold atoms in optical lattices*,
Science **357** (2017) 995.

[18] A. Browaeys and T. Lahaye, *Many-body physics with individually controlled Rydberg atoms*,
Nature Physics **16** (2020) 132, ISSN: 1745-2481.

[19] A. A. Houck, H. E. Türeci and J. Koch,
On-chip quantum simulation with superconducting circuits, Nature Physics **8** (2012) 292,
ISSN: 1745-2481.

[20] M. J. Hartmann, *Quantum simulation with interacting photons*,
Journal of Optics **18** (2016) 104005.

[21] K. Kim et al., *Quantum simulation of frustrated Ising spins with trapped ions*,
Nature **465** (2010) 590, ISSN: 1476-4687.

[22] R. Blatt and C. F. Roos, *Quantum simulations with trapped ions*, Nature Physics **8** (2012) 277,
ISSN: 1745-2481.

[23] C. Monroe et al., *Programmable quantum simulations of spin systems with trapped ions*,
Rev. Mod. Phys. **93** (2 2021) 025001.

[24] A. Aspuru-Guzik and P. Walther, *Photonic quantum simulators*, Nature Physics **8** (2012) 285,
ISSN: 1745-2481.

[25] A. G. White, "Photonic Quantum Simulation", *Research in Optical Sciences*,
Optica Publishing Group, 2014 QTh1A.1.

[26] A. Ashkin, *Acceleration and Trapping of Particles by Radiation Pressure*,
Phys. Rev. Lett. **24** (4 1970) 156.

[27] A. Ashkin and J. M. Dziedzic, *Optical Trapping and Manipulation of Viruses and Bacteria*,
Science **235** (1987) 1517.

[28] Z. Hu and H. J. Kimble, *Observation of a single atom in a magneto-optical trap*,
Opt. Lett. **19** (1994) 1888.

[29] J. D. Miller, R. A. Cline and D. J. Heinzen, *Far-off-resonance optical trapping of atoms*,
Phys. Rev. A **47** (6 1993) R4567.

[30] A. C. Doherty, T. W. Lynn, C. J. Hood and H. J. Kimble,
Trapping of single atoms with single photons in cavity QED, Phys. Rev. A **63** (1 2000) 013401.

[31] D. Frese et al.,
Single Atoms in an Optical Dipole Trap: Towards a Deterministic Source of Cold Atoms,
Physical review letters **85** (2000) 3777.

[32] N. Schlosser, G. Reymond, I. Protsenko and P. Grangier,
Sub-poissonian loading of single atoms in a microscopic dipole trap, Nature **411** (2001) 1024,
ISSN: 1476-4687.

[33] M. Weber, J. Volz, K. Saucke, C. Kurtsiefer and H. Weinfurter,
Analysis of a single-atom dipole trap, Phys. Rev. A **73** (4 2006) 043406.

[34] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais and P. Grangier, *Energy distribution and cooling of a single atom in an optical tweezer*, Phys. Rev. A **78** (3 2008) 033425.

[35] A. M. Kaufman, B. J. Lester and C. A. Regal, *Cooling a Single Atom in an Optical Tweezer to Its Quantum Ground State*, Phys. Rev. X **2** (4 2012) 041014.

[36] J. D. Thompson, T. G. Tiecke, A. S. Zibrov, V. Vuleti ć and M. D. Lukin, *Coherence and Raman Sideband Cooling of a Single Atom in an Optical Tweezer*, Phys. Rev. Lett. **110** (13 2013) 133001.

[37] Y. Miroshnychenko et al., *An atom-sorting machine*, Nature **442** (2006) 151, ISSN: 1476-4687.

[38] M. Endres et al., *Atom-by-atom assembly of defect-free one-dimensional cold atom arrays*, Science **354** (2016) 1024.

[39] K.-N. Schymik et al., *Enhanced atom-by-atom assembly of arbitrary tweezer arrays*, Phys. Rev. A **102** (6 2020) 063107.

[40] D. S. Weiss et al., *Another way to approach zero entropy for a finite system of atoms*, Phys. Rev. A **70** (4 2004) 040302.

[41] P. Scholl et al., *Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms*, Nature **595** (2021) 233, ISSN: 1476-4687.

[42] S. Ebadi et al., *Quantum phases of matter on a 256-atom programmable quantum simulator*, Nature **595** (2021) 227, ISSN: 1476-4687.

[43] M. Saffman, T. G. Walker and K. Mølmer, *Quantum information with Rydberg atoms*, Rev. Mod. Phys. **82** (3 2010) 2313.

[44] T. Wilk et al., *Entanglement of Two Individual Neutral Atoms Using Rydberg Blockade*, Phys. Rev. Lett. **104** (1 2010) 010502.

[45] C. Hamsen, K. N. Tolazzi, T. Wilk and G. Rempe, *Strong coupling between photons of two light fields mediated by one atom*, Nature Physics **14** (2018) 885, ISSN: 1745-2481.

[46] A. Cooper et al., *Alkaline-Earth Atoms in Optical Tweezers*, Phys. Rev. X **8** (4 2018) 041055.

[47] M. A. Norcia, A. W. Young and A. M. Kaufman, *Microscopic Control and Detection of Ultracold Strontium in Optical-Tweezer Arrays*, Phys. Rev. X **8** (4 2018) 041054.

[48] J. P. Covey, I. S. Madjarov, A. Cooper and M. Endres, *2000-Times Repeated Imaging of Strontium Atoms in Clock-Magic Tweezer Arrays*, Phys. Rev. Lett. **122** (17 2019) 173201.

[49] S. Saskin, J. T. Wilson, B. Grinkemeyer and J. D. Thompson, *Narrow-Line Cooling and Imaging of Ytterbium Atoms in an Optical Tweezer Array*, Phys. Rev. Lett. **122** (14 2019) 143002.

[50] A. M. Kaufman and K.-K. Ni, *Quantum science with optical tweezer arrays of ultracold atoms and molecules*, Nature Physics **17** (2021) 1324, ISSN: 1745-2481.

[51] A. W. Young, W. J. Eckner, N. Schine, A. M. Childs and A. M. Kaufman, *Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice*, Science **377** (2022) 885.

[52] P. Scholl et al., *Erasure conversion in a high-fidelity Rydberg quantum simulator*, Nature **622** (2023) 273, ISSN: 1476-4687.

[53] M. A. Norcia et al., *Seconds-scale coherence on an optical clock transition in a tweezer array*, Science **366** (2019) 93.

[54] I. S. Madjarov et al., *An Atomic-Array Optical Clock with Single-Atom Readout*, Phys. Rev. X **9** (4 2019) 041052.

[55] A. W. Young et al., *Half-minute-scale atomic coherence and high relative stability in a tweezer clock*, Nature **588** (2020) 408, ISSN: 1476-4687.

[56] A. Urech, I. H. A. Knottnerus, R. J. C. Spreeuw and F. Schreck, *Narrow-line imaging of single strontium atoms in shallow optical tweezers*, Phys. Rev. Res. **4** (2 2022) 023245.

[57] J. W. Lis et al., *Midcircuit Operations Using the omg Architecture in Neutral Atom Arrays*, Phys. Rev. X **13** (4 2023) 041035.

[58] J. Schmitz, *Spectroscopy of Neutral Sr in Optical Tweezers*, Physikalisches Institut, Universität Bonn, To be submitted.

[59] T. Kree, *A Laser System for Cooling and Trapping of Strontium Atom*, Physikalisches Institut, Universität Bonn, 2021.

[60] H. J. Metcalf and P. van der Straten, *Laser Cooling and Trapping*, Springer-Verlag, 1999.

[61] A. L. Gaunt, T. F. Schmidutz, I. Gotlibovych, R. P. Smith and Z. Hadzibabic, *Bose-Einstein Condensation of Atoms in a Uniform Potential*, Phys. Rev. Lett. **110** (20 2013) 200406.

[62] R. Grimm, M. Weidemüller and Y. B. Ovchinnikov, *Optical dipole traps for neutral atoms*, 1999, arXiv: physics/9902072 [physics.atom-ph].

[63] C. Foot, *Atomic Physics*, Oxford Master Series in Physics, OUP Oxford, 2005, ISBN: 9780198506959.

[64] E. W. Streed, A. Jechow, B. G. Norton and D. Kielpinski, *Absorption imaging of a single atom*, Nature Communications **3** (2012) 933, ISSN: 2041-1723.

[65] D. Leibfried, R. Blatt, C. Monroe and D. Wineland, *Quantum dynamics of single trapped ions*, Rev. Mod. Phys. **75** (1 2003) 281.

[66] D. J. Wineland, *Quantum information processing and quantum control with trapped atomic ions*, Physica Scripta **2009** (2009) 014007.

[67] J. E. Sansonetti and G. Nave, *Wavelengths, Transition Probabilities, and Energy Levels for the Spectrum of Neutral Strontium (SrI)*, Journal of Physical and Chemical Reference Data **39** (2010) 033103, ISSN: 0047-2689.

[68] G. Lamporesi, S. Donadello, S. Serafini and G. Ferrari,
Compact high-flux source of cold sodium atoms,
Review of Scientific Instruments **84** (2013) 063102, ISSN: 0034-6748.

[69] I. Nosske et al., *Two-dimensional magneto-optical trap as a source for cold strontium atoms*,
Phys. Rev. A **96** (5 2017) 053415.

[70] J. H. Moore, C. C. Davis, M. A. Coplan and S. C. Greer, *Building Scientific Apparatus*, 4th ed.,
Cambridge University Press, 2009.

[71] Indium Corporation, *Application Note: Indium For sealing*, [Online; accessed July 2024],
URL: https://www.indium.com/technical-documents/application-notes/download/3379/.

[72] H. Ozawa, S. Taie, Y. Takasu and Y. Takahashi,
Antiferromagnetic Spin Correlation of SU(N) *Fermi Gas in an Optical Superlattice*,
Phys. Rev. Lett. **121** (22 2018) 225303.

[73] M. Hermele, V. Gurarie and A. M. Rey, *Mott Insulators of Ultracold Fermionic Alkaline Earth Atoms: Underconstrained Magnetism and Chiral Spin Liquid*,
Phys. Rev. Lett. **103** (13 2009) 135301.

[74] T. G. Tiecke, S. D. Gensemer, A. Ludewig and J. T. M. Walraven,
High-flux two-dimensional magneto-optical-trap source for cold lithium atoms,
Phys. Rev. A **80** (1 2009) 013409.

[75] S. Stellmer, *Degenerate quantum gases of strontium*, 2013.

[76] Z. M. Liao et al.,
Thermally induced dephasing in periodically poled KTP frequency-doubling crystals,
J. Opt. Soc. Am. B **21** (2004) 2191.

[77] E. Black, *An introduction to Pound-Drever-Hall laser frequency stabilization*,
American Journal of Physics **69** (2001) 79.

[78] R. G. Escudero, *Magnetic Field Control and Laser Frequency Stabilization for Strontium Magneto-Optical Traps*, en, PhD thesis.

[79] G. H. M. van Tartwijk and D. Lenstra, *Semiconductor lasers with optical injection and feedback*,
Quantum and Semiclassical Optics: Journal of the European Optical Society Part B **7** (1995) 87.

[80] A. Kaufman, *Personal conversation*.

[81] S. Blatt, *Personal conversation*.

[82] B. Saleh and M. Teich, *Fundamentals of Photonics, 3rd Edition*, 2019, ISBN: 9781119506874.

[83] K. Lindquist, M. Stephens and C. Wieman,
Experimental and theoretical study of the vapor-cell Zeeman optical trap,
Phys. Rev. A **46** (7 1992) 4082.

[84] J. Höschele, S. Buob, A. Rubio-Abadal, V. Makhalov and L. Tarruell, *Atom-Number Enhancement by Shielding Atoms From Losses in Strontium Magneto-Optical Traps*,
Phys. Rev. Appl. **19** (6 2023) 064011.

[85] Y. Li, T. Ido, T. Eichler and H. Katori, *Narrow-line diode laser system for laser cooling of strontium atoms on the intercombination transition*, Applied Physics B **78** (2004) 315, ISSN: 1432-0649.

[86] M. A. Norcia, J. R. K. Cline, J. P. Bartolotta, M. J. Holland and J. K. Thompson, *Narrow-line laser cooling by adiabatic transfer*, New Journal of Physics **20** (2018) 023021.

[87] J. P. Bartolotta, M. A. Norcia, J. R. K. Cline, J. K. Thompson and M. J. Holland, *Laser cooling by sawtooth-wave adiabatic passage*, Phys. Rev. A **98** (2 2018) 023404.

[88] J. A. Muniz, M. A. Norcia, J. R. K. Cline and J. K. Thompson, *A Robust Narrow-Line Magneto-Optical Trap using Adiabatic Transfer*, 2018, arXiv: 1806.00838 [physics.atom-ph].

[89] S. Snigirev, A. J. Park, A. Heinz, I. Bloch and S. Blatt, *Fast and dense magneto-optical traps for strontium*, Phys. Rev. A **99** (6 2019) 063421.

[90] J. W. Goodman, *Introduction to Fourier optics*, vol. 1, 2017.

[91] D. Steck, *Quantum and Atom Optics*, [Online; accessed 08-04-2024], URL: https://atomoptics.uoregon.edu/~dsteck/teaching/quantum-optics/.

[92] F. Le Kien, P. Schneeweiss and A. Rauschenbeutel, *Dynamical polarizability of atoms in arbitrary light fields: general theory and application to cesium*, The European Physical Journal D **67** (2013) 92, ISSN: 1434-6079.

[93] M. S. Safronova, S. G. Porsev, U. I. Safronova, M. G. Kozlov and C. W. Clark, *Blackbody-radiation shift in the Sr optical atomic clock*, Phys. Rev. A **87** (1 2013) 012509.

[94] S. G. Porsev, M. S. Safronova, U. I. Safronova and M. G. Kozlov, *Multipolar Polarizabilities and Hyperpolarizabilities in the Sr Optical Lattice Clock*, Phys. Rev. Lett. **120** (6 2018) 063204.

[95] M. M. Boyd, *High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time*, 2007.

[96] R. K. Hanley, *Creation of a strontium microtrap: Towards a spin-squeezed atomic clock*, Durham University, 2018.

[97] J. Samland et al., *Optical pumping of* $5s4d\,^1D_2$ *strontium atoms for laser cooling and imaging*, Phys. Rev. Res. **6** (1 2024) 013319.

[98] M. Takamoto, F.-L. Hong, R. Higashi and H. Katori, *An optical lattice clock*, Nature **435** (2005) 321, ISSN: 1476-4687.

[99] R. Le Targat et al., *Experimental realization of an optical second with strontium lattice clocks*, Nature Communications **4** (2013) 2109, ISSN: 2041-1723.

[100] N. C. Jackson et al., *Number-resolved imaging of* ^{88}*Sr atoms in a long working distance optical tweezer*, SciPost Phys. **8** (2020) 038.

[101] CVI Melles Griot, *Gaussian beam optics: Technical Notes*, (2009).

[102] B. Zhang, J. Zerubia and J.-C. Olivo-Marin, *Gaussian approximations of fluorescence microscope point-spread function models*, Appl. Opt. **46** (2007) 1819.

[103] P. Zupancic et al., *Ultra-precise holographic beam shaping for microscopic quantum control*, Opt. Express **24** (2016) 13881.

[104] R. D. Leonardo, F. Ianni and G. Ruocco, *Computer generation of optimal holograms for optical trap arrays*, Opt. Express **15** (2007) 1913.

[105] W. O. S. R. W. Gerchberg, "A practical algorithm for the determination of phase from image and diffraction plane pictures", 1972.

[106] F. Nogrette et al., *Single-Atom Trapping in Holographic 2D Arrays of Microtraps with Arbitrary Geometries*, Phys. Rev. X **4** (2 2014) 021034.

[107] D. Kim et al., *Large-scale uniform optical focus array generation with a phase spatial light modulator*, Opt. Lett. **44** (2019) 3178.

[108] V. Jonas, *Acousto-Optic Sculpturing of Optical Potential Landscapes for Ultracold Fermions*, Physikalisches Institut, Universität Bonn, 2021.

[109] N. Schlosser, G. Reymond and P. Grangier, *Collisional Blockade in Microscopic Optical Dipole Traps*, Phys. Rev. Lett. **89** (2 2002) 023005.

[110] B. J. Reschovsky et al., *Narrow-line photoassociation spectroscopy and mass-scaling of bosonic strontium*, (2018), arXiv: 1808.06507 [physics.atom-ph].

[111] A. Fuhrmanek, R. Bourgain, Y. R. P. Sortais and A. Browaeys, *Light-assisted collisions between a few cold atoms in a microscopic dipole trap*, Phys. Rev. A **85** (6 2012) 062708.

[112] T. Grünzweig, A. Hilliard, M. McGovern and M. F. Andersen, *Near-deterministic preparation of a single atom in an optical microtrap*, en, Nature Physics **6** (2010) 951, ISSN: 1745-2473, 1745-2481, (visited on 17/10/2022).

[113] M. O. Brown, T. Thiele, C. Kiehl, T.-W. Hsu and C. A. Regal, *Gray-Molasses Optical-Tweezer Loading: Controlling Collisions for Scaling Atom-Array Assembly*, Phys. Rev. X **9** (1 2019) 011057.

[114] F. Diedrich, J. C. Bergquist, W. M. Itano and D. J. Wineland, *Laser cooling to the zero-point energy of motion.*, Physical review letters **62 4** (1989) 403.

[115] J. Eschner, G. Morigi, F. Schmidt-Kaler and R. Blatt, *Laser cooling of trapped ions*, J. Opt. Soc. Am. B **20** (2003) 1003.

[116] D. Wineland et al., *Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions*, Journal of Research of the National Institute of Standards and Technology **103** (1997).

[117] Taïeb, Dum, Cirac, Marte and Zoller, *Cooling and localization of atoms in laser-induced potential wells.*, Physical review. A, Atomic, molecular, and optical physics **49 6** (1994) 4876.

[118] J. I. Cirac, M. Lewenstein and P. Zoller, *Collective laser cooling of trapped atoms*, Europhysics Letters **35** (1996) 647.

[119] Y. Castin, J. I. Cirac and M. Lewenstein, *Reabsorption of Light by Trapped Atoms*, Phys. Rev. Lett. **80** (24 1998) 5305.

[120] C. Monroe et al., *Resolved-Sideband Raman Cooling of a Bound Atom to the 3D Zero-Point Energy*, Phys. Rev. Lett. **75** (22 1995) 4011.

[121] S. E. Hamann et al., *Resolved-Sideband Raman Cooling to the Ground State of an Optical Lattice*, Phys. Rev. Lett. **80** (19 1998) 4149.

[122] A. Jenkins, J. W. Lis, A. Senoo, W. F. McGrew and A. M. Kaufman, *Ytterbium Nuclear-Spin Qubits in an Optical Tweezer Array*, Phys. Rev. X **12** (2 2022) 021027.

[123] D. J. Wineland and W. M. Itano, *Laser cooling of atoms*, Phys. Rev. A **20** (4 1979) 1521.

[124] D. Manzano, *A short introduction to the Lindblad master equation*, AIP Advances **10** (2020) 025106, ISSN: 2158-3226.

[125] J. Dalibard and C. Cohen-Tannoudji, *Laser cooling below the Doppler limit by polarization gradients: simple theoretical models*, Journal of The Optical Society of America B-optical Physics **6** (1989) 2023.

[126] V. V. Ivanov and S. Gupta, *Laser-driven Sisyphus cooling in an optical dipole trap*, Phys. Rev. A **84** (6 2011) 063417.

[127] C. Hölzl et al., *Motional ground-state cooling of single atoms in state-dependent optical tweezers*, Phys. Rev. Res. **5** (3 2023) 033093.

[128] W. Alt et al., *Single atoms in a standing-wave dipole trap*, Phys. Rev. A **67** (3 2003) 033403.

[129] M. S. Robbins, "Electron-Multiplying Charge Coupled Devices – EMCCDs", *Single-Photon Imaging*, ed. by P. Seitz and A. J. Theuwissen, Berlin, Heidelberg: Springer Berlin Heidelberg, 2011 103, ISBN: 978-3-642-18443-7.

[130] A. Bergschneider, *Strong correlations in few-fermion systems*, University of Heidelberg, 2017.

[131] Oxford Instruments, *What is EM camera Gain Ageing in EMCCD Sensors?*, [Online; accessed July 2024], URL: https://andor.oxinst.com/learning/view/article/an-overview-of-em-gain-ageing-in-emccd-sensors.

[132] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, 3rd ed., Cambridge University Press, 2020.

[133] D. Bluvstein et al., *Logical quantum processor based on reconfigurable atom arrays*, Nature **626** (2024) 58, ISSN: 1476-4687.

[134] K.-N. Schymik et al., *Single Atoms with 6000-Second Trapping Lifetimes in Optical-Tweezer Arrays at Cryogenic Temperatures*, Phys. Rev. Appl. **16** (3 2021) 034013.

[135] J. O'Hanlon, *A User's Guide to Vacuum Technology*, Wiley, 2005, ISBN: 9780471467151.

[136] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye and A. Browaeys, *An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays*, Science **354** (2016) 1021.

[137] F. Gyger et al., *Continuous operation of large-scale atom arrays in optical lattices*, 2024, arXiv: 2402.04994 [quant-ph].

[138] D. Kirk, *Optimal Control Theory: An Introduction*, Dover Books on Electrical Engineering, Dover Publications, 2012, ISBN: 9780486135076.

[139] D. Bluvstein, *Neutral atom quantum processors and the error correction frontier*, 2023.

[140] A. Pagano, D. Jaschke, W. Weiss and S. Montangero, *Optimal control transport of neutral atoms in optical tweezers at finite temperature*, 2024, arXiv: 2402.17831 [quant-ph].

[141] D. Mellinger and V. Kumar, *Minimum snap trajectory generation and control for quadrotors*, (2011) 2520.

[142] A. W. Young et al., *An atomic boson sampler*, Nature **629** (2024) 311, ISSN: 1476-4687.

[143] M. Link et al., *Machine Learning the Phase Diagram of a Strongly Interacting Fermi Gas*, Phys. Rev. Lett. **130** (20 2023) 203401.

[144] Y. Lee and E. Chae, *Machine learning-enhanced optical tweezers for defect-free rearrangement*, Current Applied Physics **61** (2024) 150, ISSN: 1567-1739.

[145] T. F. Gallagher, *Rydberg Atoms*, Cambridge Monographs on Atomic, Molecular and Chemical Physics, Cambridge University Press, 1994.

[146] N. Šibalić and C. S. Adams, *Rydberg Physics*, 2399-2891, IOP Publishing, 2018, ISBN: 978-0-7503-1635-4.

[147] D. Jaksch et al., *Fast Quantum Gates for Neutral Atoms*, Phys. Rev. Lett. **85** (10 2000) 2208.

[148] A. Pagano et al., *Error budgeting for a controlled-phase gate with strontium-88 Rydberg atoms*, Phys. Rev. Res. **4** (3 2022) 033019.

[149] W. J. Eckner et al., *Realizing spin squeezing with Rydberg interactions in an optical clock*, Nature **621** (2023) 734, ISSN: 1476-4687.

[150] C. Hölzl, A. Götzelmann, E. Pultinevicius, M. Wirth and F. Meinert, *Long-Lived Circular Rydberg Qubits of Alkaline-Earth Atoms in Optical Tweezers*, Phys. Rev. X **14** (2 2024) 021024.

[151] A. Reiserer and G. Rempe, *Cavity-based quantum networks with single atoms and optical photons*, Rev. Mod. Phys. **87** (4 2015) 1379.

[152] M. Steiner, H. M. Meyer, C. Deutsch, J. Reichel and M. Köhl, *Single Ion Coupled to an Optical Fiber Cavity*, Phys. Rev. Lett. **110** (4 2013) 043003.

[153] M. Brune et al., *Quantum Rabi Oscillation: A Direct Test of Field Quantization in a Cavity*, Phys. Rev. Lett. **76** (11 1996) 1800.

[154] A. Boca et al., *Observation of the Vacuum Rabi Spectrum for One Trapped Atom*, Phys. Rev. Lett. **93** (23 2004) 233603.

[155] L.-M. Duan, B. Wang and H. J. Kimble, *Robust quantum gates on neutral atoms with cavity-assisted photon scattering*, Phys. Rev. A **72** (3 2005) 032333.

[156] X.-M. Lin, Z.-W. Zhou, M.-Y. Ye, Y.-F. Xiao and G.-C. Guo, *One-step implementation of a multiqubit controlled-phase-flip gate*, Phys. Rev. A **73** (1 2006) 012323.

[157] M. Born and E. Wolf, *Principles of Optics: 60th Anniversary Edition*, 7th ed., Cambridge University Press, 2019.

Acknowledgements

First and foremost, I would like to thank Michael Köhl for offering me this fascinating project, that perfectly matched my interests. Designing and building a new experiment was a huge challenge, and I thank you for the support, for the trust and for the freedom you gave me during that time. I also thank you for all the physical insights you shared with me and your sometimes astonishingly simple answers to seemingly complicated questions. This not only taught me how to physics but also how to think. Additionally, I thank you for the opportunity to spent parts of my PhD abroad in another group.

Speaking of which, I thank Adam Kaufman for offering me an internship in your group and for an interesting research project to work on. I learned so much from you and during my time in Boulder, and I am incredibly thankful for the experience.

As mentioned in the introduction, a project of this size cannot be conducted by a single person alone, and there are two guys without whom this would not have been possible. I therefore thank the *Fairy*-team of Jonas Schmitz and Tobias Kree for the successes, the failures and the good hours we shared in the lab. It was an amazing time. Amazing times I heard? I thank Thorsten Langerfeld and Leon Reiffert for all the lab time we shared, all the exercise sheets we worked on, but mostly for becoming good friends over the years (and maybe also for having a beer here and there). I thank Andrea Bergschneider for fruitful and heated discussions and for proofreading parts of my work.

I also thank Andreas K., Andreas v. H., Daniel, Janek, Marcell, Martin, Moritz B., Moritz S., Nick, Nicola, Pascal, Santhosh, Steffen and Valentin for interesting group meetings, for weird mensa discussions, for long christmas parties and totally not for any dubious experiments involving a microwave.

During my time in the US, I was welcomed with open arms by Aaron, Alec, Brian, Joanna, Morgan, Nathan, Howie and Will and I thank you for the game nights, the skiing and simply the great time I had in Boulder.

I thank Akos Hoffmann for your help and electronic skills, needed while building the experiment, for your interesting explanations of intricate details, and for your detailed presentations why someone else had no idea what they were doing when I entered your office time and again with broken self-built electronics that you had to fix.

I thank Daniel Spelter and Marvin Lenk for interesting discussions sparking and fueling my research interests during our Bachelor's studies, which eventually led up to this very point.

Special thanks also to you Ludger, for your immediate help when it was most needed.

Finally, but most importantly, I thank my mother, my sister Katharina, my grandma, Julia, and Lasse, but also all my friends, for the support over the years, for enduring my highs and my lows, for all the memories we share, and for the ones yet to be made.