Magneto-Optical Trap for Continuous Reloading of a Neutral Atom Quantum Computer

Ria Rosenauer

M. Sc. Quantum Science & Technology

Technische Universität München
Ludwig-Maximilians-Universität München
California Institute of Technology

Pasadena, April 2026

Erstgutachter: Prof. Dr. Johannes Zeiher
Zweitgutachter: Prof. Dr. Manuel Endres

Abstract

Neutral atom arrays have emerged as a highly scalable platform for fault-tolerant quantum computing, offering dynamic reconfigurability and the all-to-all connectivity required for advanced error correction schemes, such as qLDPC codes. However, the platform faces a fundamental limitation: irreversible atom loss induced by background gas collisions, photon scattering, and recoil heating from mid-circuit measurements. This continuous loss strictly limits the achievable circuit depth. Overcoming this bottleneck requires architectures capable of continuous, mid-circuit atom reloading without decohering the qubits. This thesis presents the design, theoretical analysis, and partial experimental implementation of a high-flux continuous-reloading magneto-optical trap (MOT) tailored for a next-generation >10,000-atom ^{133}Cs tweezer array.

This work begins by providing an overview of continuous reloading architectures and detailing the rationale behind our dual-chamber vacuum system, which features a 45-degree optical lattice handshake designed to geometrically isolate the computational array from MOT fluorescence. We then explore the physical constraints of extracting atoms from an actively running MOT using a continuous "ratchet-type" transport scheme. Using multi-level semi-classical rate equations, we analyze critical system dynamics, including position-dependent damping, stray light decoherence, and the necessity of dark-state pumping to prevent spin-changing collisions. Guided by these simulations, this thesis proposes a large-beam MOT configuration with a "pac-man"-shaped repumper profile. The idea of this custom geometry is to create a dark tube along the transport lattice, allowing atoms to be optically pumped into the absolute m_F ground state. This tailored shielding mitigates spin-changing collisions and facilitates the efficient extraction of atoms directly from the active MOT. Finally, the thesis details the specialized optical setup engineered to realize this continuous, high-flux atom source.

Contents

Chapter 1

Introduction

Neutral atoms have emerged in recent years as a highly promising platform for both quantum computing and quantum simulation [1–3]. A key driver of this progress is the scalability of the platform, with a recent experiment successfully demonstrating the operation of a 6100-atom tweezer array [4]. Beyond raw qubit numbers, neutral atom systems offer a high degree of spatial flexibility. Unlike architectures with fixed physical topologies, such as superconducting qubits [5–7], neutral atoms in tweezers feature dynamic reconfigurability. This flexibility in positioning and moving atoms provides a distinct advantage for building arbitrary geometries tailored for specific quantum information processing tasks and complex many-body quantum simulations. A particularly powerful feature of this reconfigurability is the ability to coherently transport qubits mid-circuit without losing quantum coherence [8]. The flexible, all-to-all connectivity enables a variety of novel fault-tolerant architectures [9] and error correction schemes that are difficult or impossible to implement on other platforms. A prominent recent example is the implementation of non-local quantum low-density parity-check (qLDPC) codes [10, 11]. These have emerged as highly promising candidates for low-overhead, high-distance error correction, significantly lowering the resource estimates for when quantum computers could successfully break classical cryptography [12]. Non-local qLDPC codes inherently require long-range interactions, which can be very difficult for other platforms to achieve. However, the ability to dynamically reconfigure the atoms within a neutral atom platform allows these long-range interactions to be naturally realized by physically shuttling atoms across the array to interact with distant pairs [13]. In this way, the system can dynamically generate the complex, non-local entanglement graphs required by these codes without the routing overhead or fixed wiring constraints typical of other architectures.

Despite these significant advantages, neutral atom systems face several fundamental bottlenecks. Most notably, physical qubits are irreversibly lost when atoms escape their traps. The loss is driven by a combination of background gas collisions, blackbody radiation, and recoil heating that accumulates during gate operations and readout [1, 14]. This continuous atom loss places a hard limit on achievable circuit depth. Another major limit is the ability to read out qubits mid-circuit without disturbing the rest of the computation. In fault-tolerant architectures, mid-circuit readout is strictly required for continuous error syndrome extraction [8, 15]. However, standard state detection relies on fluorescence imaging, a comparatively slow process that scatters thousands of photons. Scattered light introduces severe crosstalk, causing decoherence in neighboring spectator qubits, and frequently imparting enough recoil heating to eject the measured atoms [16–18]. Consequently, the very operations required for active error correction (repeated syndrome measurements over deep circuits) intrinsically exacerbate the atom loss problem.

To overcome this limitation and enable the sustained execution of fault-tolerant quantum algorithms, the system must be able to efficiently replenish lost or overheated atoms during the computation, a process known as continuous reloading. Recently, significant progress has been made toward this capability, with three major demonstrations of continuous reloading already published [19–21]. Complementary approaches to continuous operation have also been realized in various other architectures, including dynamically reloaded optical lattices [22], dual-species continuous-mode arrays [23], and zero dead-time optical clocks [24]. Building upon these advancements, this thesis contributes to the broader effort of developing a continuous-reloading ^{133}Cs tweezer array with a very high flux that does not decohere active qubits, thereby enabling sustained computation. A crucial first step toward this goal is the MOT, as the MOT's loading rate, atom number, and atom temperature fundamentally limit the highest achievable continuous reloading flux. For this reason, this thesis focuses on the design, theoretical analysis, and initial experimental implementation of this specialized continuous-reloading MOT.

Chapter 2

Experimental Context and Theoretical Background

The core objective of this thesis is the design and initial construction of a MOT specifically engineered for continuous reloading. However, the architectural choices, spatial constraints, and performance targets of this MOT cannot be understood in isolation. They are dictated by broader system-level strategies, such as vacuum design, the choice of atomic transport, and transport timing schemes. This chapter provides the necessary background to contextualize these design decisions. First, we outline the fundamental principles and challenges of continuous reloading. Next, we detail the specific vacuum design, transport geometry, and timing schemes chosen for our experiment. Establishing this experimental framework allows us, in section 3, to systematically derive the strict performance requirements that the MOT must satisfy.

2.1 Continuous Reloading Introduction

Neutral atom tweezer arrays represent a highly promising platform for fault-tolerant quantum computing. Their all-to-all connectivity, dynamic reconfigurability, and flexible coherent atom transport enable a variety of novel error correction schemes that are difficult or impossible to implement on other platforms, such as superconducting qubits or trapped ions. However, neutral atom systems face a distinct and fundamental challenge: the physical qubits themselves can be irreversibly lost when atoms escapes their traps. The primary loss mechanisms include collisions with residual background gas, off-resonant photon scattering from nearby laser beams, and motional heating induced by gate operations and fluorescence imaging [2, 4, 25]. This continuous atom loss places a hard limit on the achievable circuit depth. Furthermore, neutral atom platforms are burdened by limited mid-circuit readout capabilities, slow readout speeds, and long cycle times [5, 9]. Significant experimental efforts are currently dedicated to mitigating these loss rates. Advances such as cryogenic vacuum environments to drastically reduce background gas collisions, improved cooling protocols to counteract gate-induced heating, and novel non-destructive readout schemes all aim to extend physical qubit lifetimes [17, 26]. However, as system sizes scale toward our target of $> 10{,}000$ atoms and quantum error correction circuits become increasingly deep, the absolute rate of atom loss across the array becomes non-negligible. For truly fault-tolerant, long-duration algorithms, simply reducing the loss rate is insufficient. The ability to efficiently replace lost or overheated atoms mid-computation is therefore essential, as it addresses many of these scaling challenges simultaneously. Most directly, the continuous replenishment of lost atoms would substantially extend achievable circuit depths. Looking further ahead, continuous reloading could also unlock faster mid-circuit readout schemes: hot atoms that are rapidly imaged could simply be discarded and replaced rather than recovered via slow cooling protocols. Similarly, ancilla qubits could be subjected to fast, destructive imaging with the assurance that any lost ancillas would be promptly replenished. While continuous atom reloading is thus a highly desirable capability, it remains technically challenging. At present, reloading atoms into the computation zone of a conventional neutral atom tweezer apparatus (where the MOT is co-located within the science chamber [4, 27]) is both slow and highly disruptive. Most critically, the standard reloading process destroys the coherence of existing qubits due to the strong magnetic field gradients and near-resonant stray light produced by the MOT [19].

To overcome this limitation and enable fault-tolerant quantum computation, several groups have recently begun working on continuous reloading schemes to load atoms continuously mid-circuit with high flux, all without destroying existing qubits. The requirements these continuous reloading schemes must fulfill are:

1. The reloading should be continuously possible mid-circuit without destroying the coherence of existing qubits. Experiments that do not fulfill this requirement suffer from decoherence during the reloading process, severely limiting their practical utility for mid-circuit operations.
2. The arriving atoms should be as cold as possible so that minimal additional

cooling in the tweezer region is needed.

3. The higher the flux of incoming atoms, the better. As quantum processors scale to larger array sizes and shorter cycle times, both of which are critical parameters to optimize for practical quantum computing, the absolute number of atoms lost per operation increases, necessitating a proportionately higher continuous flux to sustain operation.

These requirements present competing design challenges. First, MOT placement involves a fundamental trade-off: positioning the MOT closer to the science chamber facilitates atom transport with reduced heating and loss, but increases stray light contamination in the computation zone. Second, higher atomic flux often comes at the cost of higher atomic temperatures, requiring careful optimization of the MOT parameters and cooling stages.

To the best of our knowledge, there are currently three published single-species experiments that fulfill these requirements: continuous reloading of a ^{171}Yb 256-tweezer array by the Thompson group [19], a 3000-qubit Rb tweezer array by the Lukin group [20], and a $\approx$ 100-atom ^{171}Yb platform by Atom Computing [21]. While other continuous reloading demonstrations exist (e.g., [22, 28]), they do not satisfy requirement 1. In our experiment, we aim to demonstrate the continuous reloading of a >10,000-atom ^{133}Cs tweezer array, building upon the previous 6,100-atom demonstration [4] in our group. The choice of ^{133}Cs is primarily motivated by its high polarizability ($\alpha_{\mathrm{Cs}} \approx 59.4 \times 10^{-24}\,\mathrm{cm}^3$ at 1064 nm). Compared to other alkali atoms such as Rb ($\alpha_{\mathrm{Rb}} \approx 26.4 \times 10^{-24}\,\mathrm{cm}^3$), this enables deeper optical traps and tighter spatial confinement for a given laser power. Furthermore, the standard tweezer trapping wavelength of 1064 nm benefits from the commercial availability of high-power continuous-wave lasers, which significantly facilitates scaling. This combination of high polarizability and abundant laser power allowed our group to previously demonstrate a 6,100-atom tweezer array [4], the largest reported at the time. Because efficient use of optical power is a strict physical constraint when scaling to our target of >10,000 atoms, ^{133}Cs is well suited for this architecture.

^{133}Cs is the only naturally occurring stable isotope of cesium. The qubit is encoded in the hyperfine ground states of the ^{133}Cs $6S_{1/2}$ ground state, specifically the magnetic field-insensitive clock states $|0\rangle \equiv |F = 3, m_F = 0\rangle$ and $|1\rangle \equiv |F = 4, m_F = 0\rangle$, which are separated by 9.2 GHz [4, 29]. The transition most commonly used for magneto-optical trapping of ^{133}Cs is the D_2 line at 852 nm, corresponding to the $6S_{1/2} \to 6P_{3/2}$

transition. Figure 2.1, taken from [30], shows the hyperfine structure of the ground state and the D_2 line level diagram. Furthermore, figure 2.2 provides a comprehensive overview of the relevant transitions, states, and wavelengths in ^{133}Cs. To achieve our target of continuously reloading >10,000 ^{133}Cs atoms we target higher atom fluxes and lower temperatures than previous implementations. In the following section, we review the continuous reloading techniques employed by other groups and detail the rationale behind our fundamental design choices regarding the vacuum system, atom transport, and timing schemes.

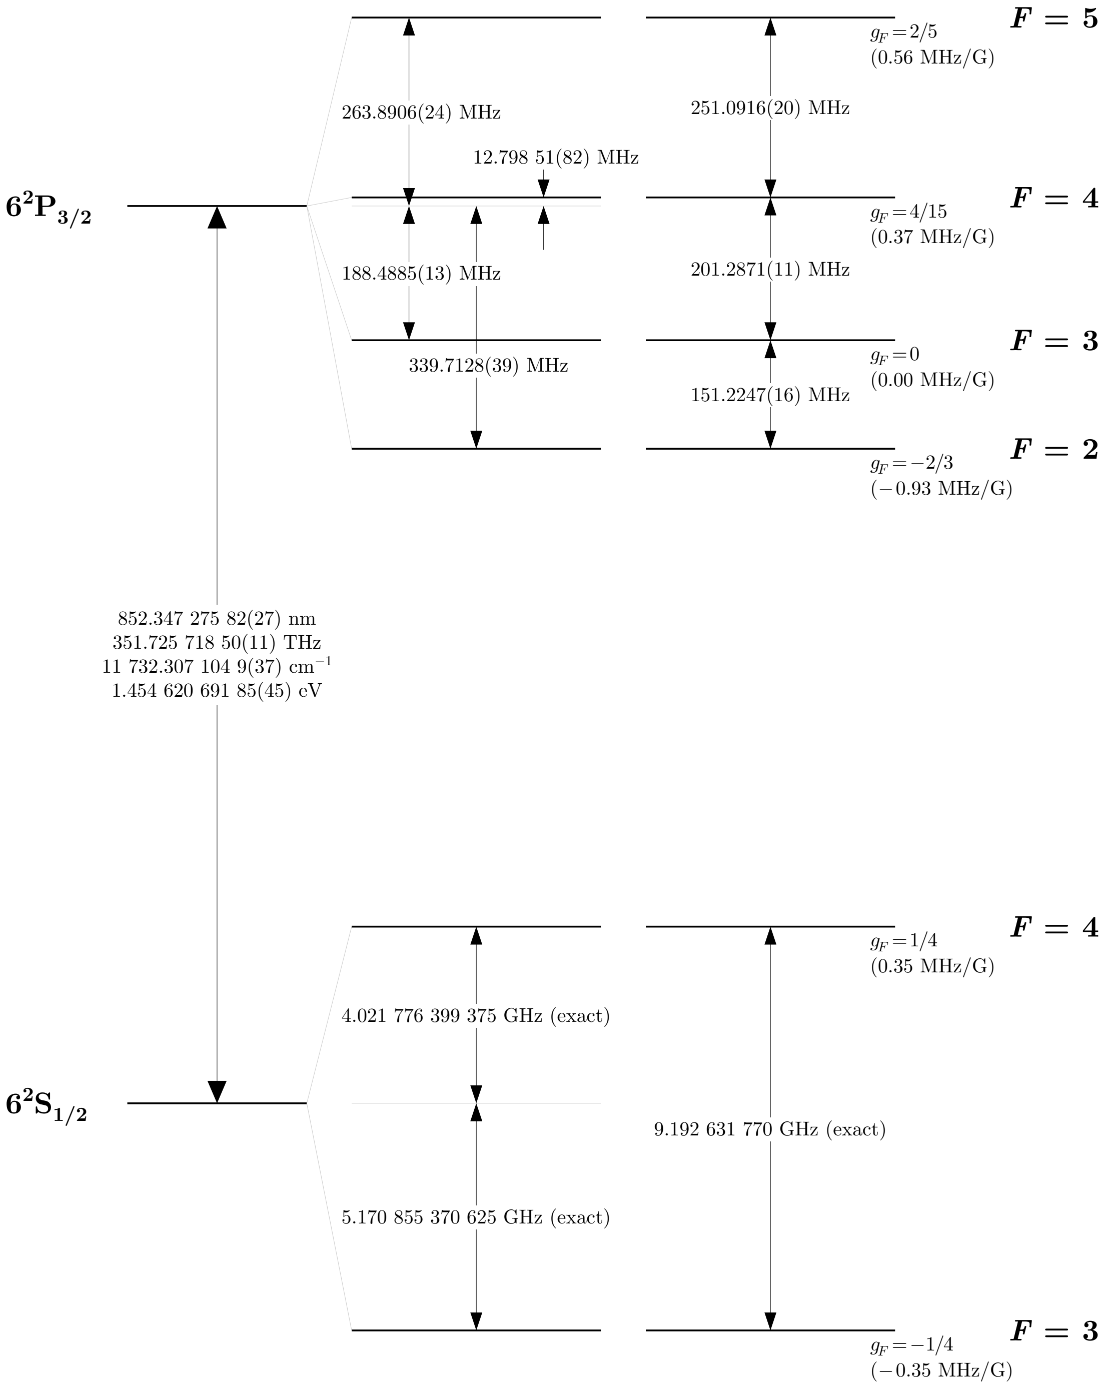

Figure 2.1: ^{133}Cs D_2 line energy diagram, taken from [30]. The qubit is typically encoded in the hyperfine ground states $|F = 3, m_F = 0\rangle$ and $|F = 4, m_F = 0\rangle$. The most important transition for this thesis is the D_2 line at 852 nm, which is used for MOT cooling and state preparation.

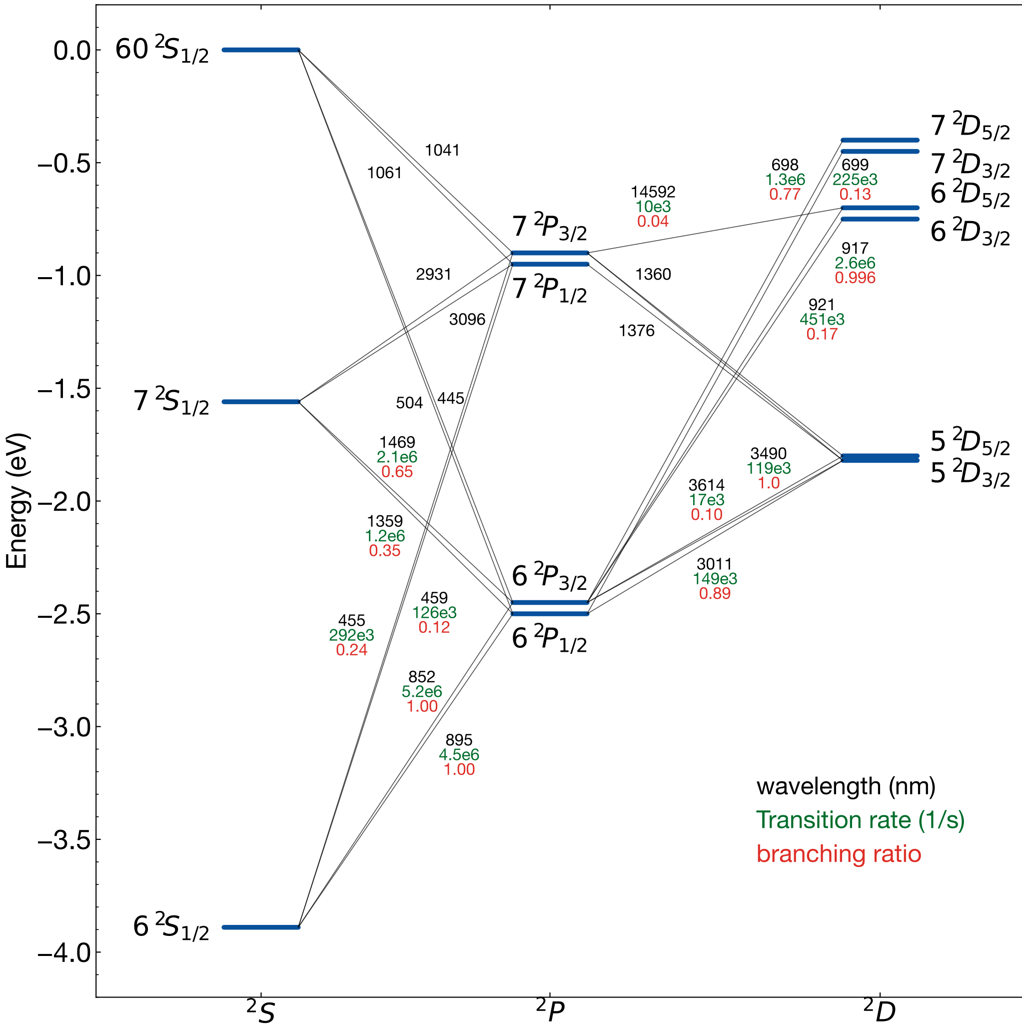

Figure 2.2: Energy level structure of ^{133}Cs, showing the ground state $6S_{1/2}$, the excited P states associated with the D lines, and higher-lying S, P, D, and F states.

2.1.1 Vacuum Design and Geometry of Atom Transport

One of the larger challenges in continuous reloading is protecting the existing qubits from MOT-induced decoherence. Since the qubits reside in the $6S_{1/2}$ ground state, stray 852 nm MOT light causes resonant photon scattering that destroys the encoded quantum information. There are two primary approaches to mitigate this: (1) operating the MOT in the same region and applying local spectral protection to the existing atoms to render them off-resonant with the MOT light, or (2) spatially separating the MOT and tweezer regions to ensure that negligible stray light and magnetic field fluctuations reach the computation zone. If pursuing local spectral protection, the atomic transition frequency can theoretically be shifted using either engineered light shifts (AC Stark shifts) [20] or large magnetic bias fields (Zeeman shifts) [31]. However, relying on magnetic shifts to provide GHz-scale spectral protection is very hard. The required bias fields (exceeding 100 G) would necessitate power supplies with extreme stability to avoid dephasing from current noise. Furthermore, such strong fields would significantly perturb the excited state manifolds (e.g. the $6P_{1/2}$ state in ^{133}Cs), potentially complicating the implementation of high-fidelity Raman gates or Rydberg transitions. Consequently, light-shifting approaches (which can locally shield the optical transition [32, 33]) are generally preferred for single-chamber continuous reloading. While Rydberg gates and Raman operations are theoretically compatible with these shifts, performing high-fidelity state-resolved imaging and state preparation while the atoms are spectrally shifted remains an experimental challenge. For these reasons, current state-of-the-art continuous reloading experiments [19–21] utilize a dual-chamber architecture, isolating the 3D MOT in a dedicated chamber away from the science region. This spatial separation, however, introduces new design challenges. Atoms must be efficiently transported from the MOT chamber to the science chamber, raising three critical questions: how should the vacuum chambers be designed to accommodate this transport, which transport method is optimal, and what timing scheme best synchronizes this delivery?

Starting with the first question raised above of how to transport atoms from the MOT chamber: As during transport atoms have to be accelerated, it inherently introduces both atom loss and heating, necessitating careful selection of the transport method to minimize these effects. Several transport schemes were considered:

- Gravity: One approach allows atoms to fall from the 3D MOT chamber under gravitational acceleration. However, this requires a slowing beam in the science chamber, adding complexity and reintroducing some of the stray light issues that the dual-chamber design was intended to avoid.

- Continuous Dipole Trap: An alternative approach, demonstrated by the Thompson group in [20], uses an off-resonant dipole trap focused near the tweezer region, causing atoms to migrate there continuously. A second 1D molasses beam close to the tweezer region then confines the atoms so they can be loaded into tweezers. This approach relies on the fact that, in ^{171}Yb, the molasses can be applied on the $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ intercombination line while the qubit is safely stored in the metastable $^3\mathrm{P}_0$ state and thus unaffected by the cooling light. This fails for ^{133}Cs, where the qubit is stored in the $6^2\mathrm{S}_{1/2}$ hyperfine states $|F=3, m_F=0\rangle$ and $|F=4, m_F=0\rangle$, which are directly affected by the MOT light.
- Moving Dipole Trap: A related variation uses a moving dipole trap (e.g., scanning with a galvo mirror) to transport discrete bunches of atoms to the science chamber. However, such traps are slow to move and can only transport atoms over short distances (roughly 5 cm to 10 cm), which is insufficient to adequately attenuate the MOT stray light.
- Optical Lattice Transport: This approach utilizes a 1D optical lattice, which is formed by the interference of two counterpropagating laser beams to create a periodic trapping potential. By introducing a controlled frequency difference between these beams, the resulting standing wave pattern is translated, smoothly moving the trapped atoms along the beam axis. The authors in [19] use this method for their continuous reloading. They first load the MOT, turn it off, load the atom bunch into the lattice, and then transport it to the science chamber. Optical lattice transport offers several advantages, including precise velocity control through lattice acceleration and a large capture volume for high flux.

Because of the different advantages and disadvantages explained above, we choose optical lattice transport, the same approach as used in [19]. For lattice transport, one can use either a near-resonant or far-detuned lattice wavelength. While near-resonant lattices require less optical power for the same trap depth, they introduce significant stray light in the science chamber and can cause photon scattering during transport, which may lead to heating. Therefore, similar to [19], we plan to implement a far-detuned lattice at a wavelength of either 1030 nm or 1064 nm using approximately 60 W of optical power per beam. The beam waist radius of the lattice will be on the order of 200 µm (resulting in a trap depth of roughly 1 mK) or 300 µm (yielding a trap depth of 444 µK).

An additional design choice for lattice transport is whether to use a single lattice for direct transport or to employ a two-lattice handshake scheme. Direct transport is simpler and more efficient, as lattice handshakes can introduce additional heating and atom loss. However, direct transport requires an unobstructed optical path between the MOT chamber and the science region, creating a direct line of sight between the two. As shown in our decoherence simulations in Appendix C, even minimal MOT stray light could cause significant qubit decoherence in the computation zone. While light shifts could potentially mitigate this, the more conservative approach is to introduce a second lattice at an angle to the first, breaking the direct line of sight between the MOT and the science region. The handshake angle is then a critical design parameter. Shallow angles reduce the velocity mismatch between the two lattices and allow for greater spatial overlap during transfer, both of which improve transfer efficiency and reduce heating. Larger angles more effectively eliminate the line of sight between the computation zone and the MOT, but at the cost of transfer efficiency. Advanced techniques such as adiabatic ramps or spatially varying lattice phases could potentially recover some of these losses at larger angles.

Initially we considered the shallow-angle geometry to provide a more robust solution with lower technical risk. Additionally, we aimed to allow switching between direct transport (in the event our light-shifting techniques perform well) and shallow-angle transport (as a reliable fallback if they do not). Meeting these requirements through extensive iteration resulted in a first version of the vacuum chamber design shown in figure 2.3. This system consists of an octagonal MOT chamber, where the number of ports is optimized to provide beam access for cooling and imaging, vacuum pump connections, and the 2D MOT. The MOT chamber connects to the science chamber (glass cell from Akatsuki) via a black-coated (Acktar Magic Black) tube. A bellow between the tube and the MOT chamber allows for fine adjustment of the transport angle, enabling us to switch between direct and indirect transport configurations. The lattice handshake is implemented using an intermediate chamber positioned along the tube, where the second lattice beam enters and exits the vacuum system.

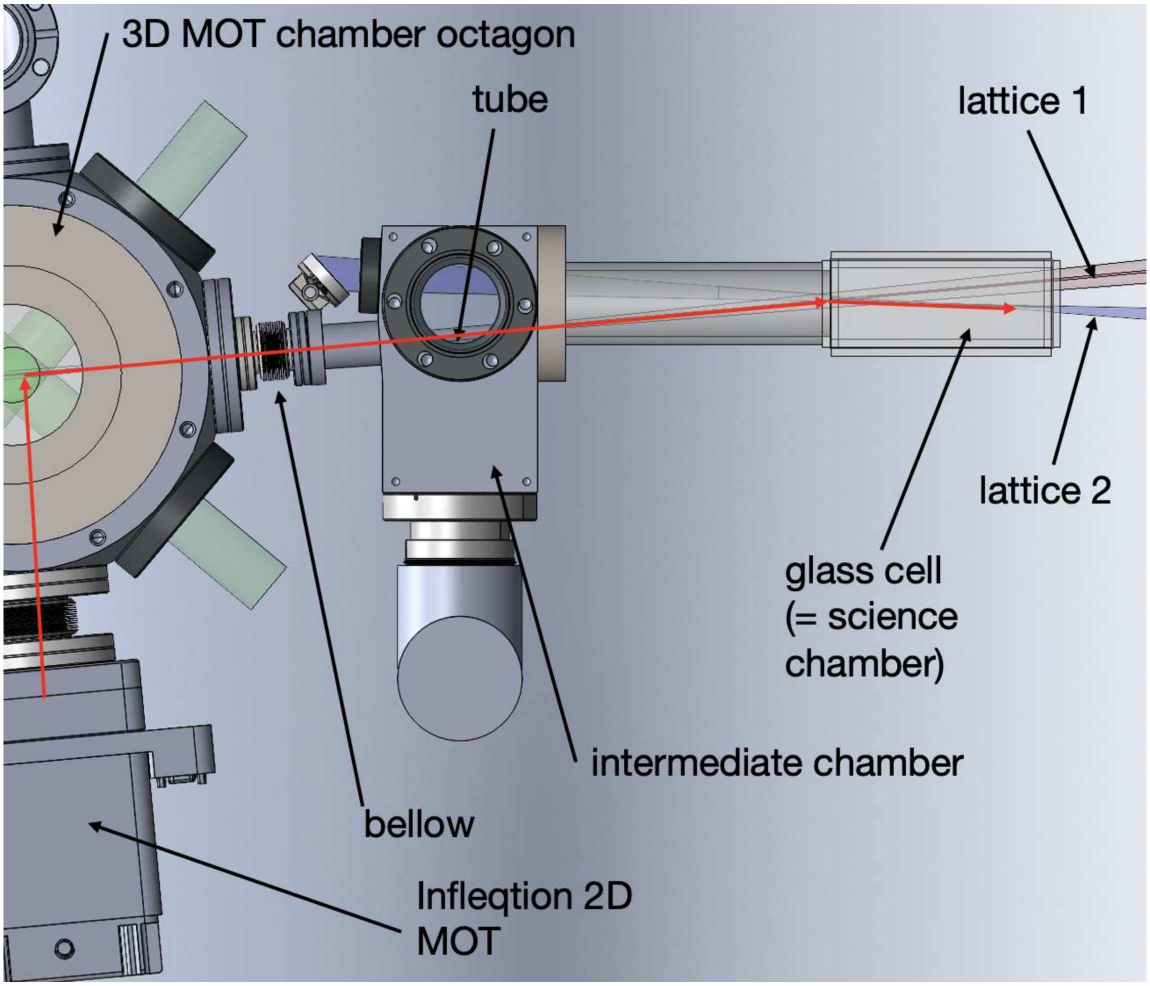

Figure 2.3: Initial vacuum chamber design with coated tube featuring the octagonal MOT chamber, intermediate handoff chamber, and glass cell science chamber. A bellow between the MOT chamber and the science chamber allows to switch between direct and indirect transport. The red arrows indicate the atomic transport path from the 2D MOT, through the 3D MOT chamber, along the transport tube via optical lattice beams, and into the science chamber where atoms are loaded into optical tweezers. To protect the qubits from decoherence, the design relied on a shallow-angle lattice handshake and an Acktar Magic Black-coated tube to place the computation zone in the geometric shadow of the MOT. This approach assumed the coating would inhibit internal reflections.

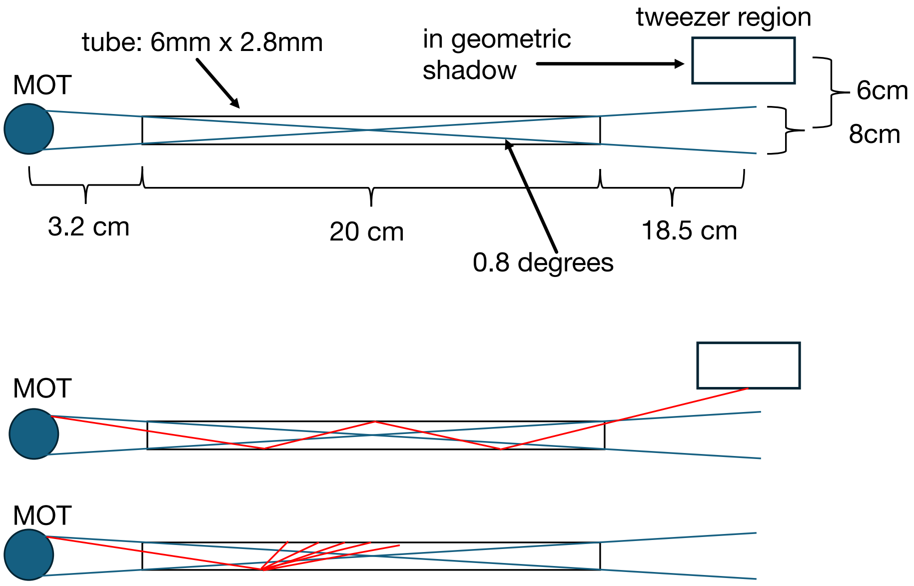

Figure 2.4: Illustration of the geometric shadow concept for indirect transport. The top panel shows the ideal geometric shadow assumption, which hinges on the fact that neither specular internal reflections (middle panel) nor diffusive reflections (bottom panel, indicated by the red example light rays) would reach the computation zone.

A key design idea for the indirect transport involves placing the tweezer array in the geometric shadow of the MOT, as shown in figure 2.4. However, this relies on the assumption that the tube coating would perform good at shallow grazing angles, such that no specular internal reflections (middle scenario in figure 2.4) or diffusive reflections (bottom scenario in figure 2.4) would reach the tweezer array. However, it turns out that the Acktar Magic Black coating exhibits exactly these two reflection features at shallow angles. Detailed experimental analysis of these reflections is given in appendix A. For this reason, we propose a second iteration of the vacuum system, shown in figure 2.5. In this design, we implement a 45-degree lattice handshake within the intermediate chamber, completely eliminating any direct line of sight between the MOT and the science chamber. As an additional precaution, we install an aperture between the MOT chamber and the intermediate chamber. This ensures that any direct MOT stray light geometrically exits through the opposite bucket window of the intermediate chamber. Furthermore, the interior of the MOT chamber is coated with Acktar Magic Black to minimize non-specular, wide-angle reflections into the handshake chamber as much as possible.

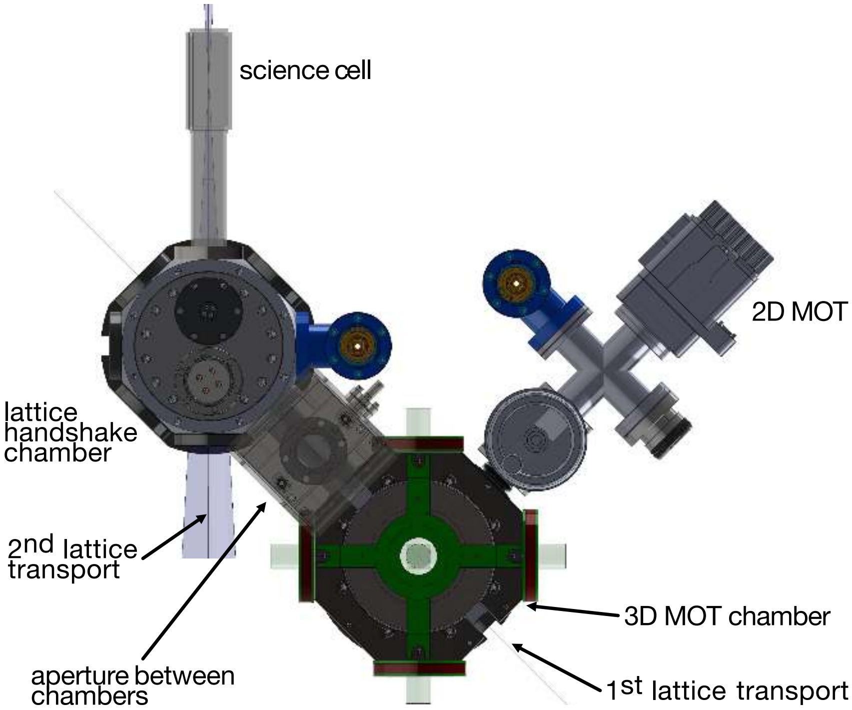

Figure 2.5: Revised vacuum system design featuring a 45-degree lattice handshake. This geometry, combined with Acktar Magic Black coating in the MOT chamber and an aperture between MOT chamber and handshake chamber, ensures that stray MOT light is either absorbed or geometrically directed out of the intermediate chamber, completely eliminating any direct or reflected line of sight to the science chamber.

2.1.2 Transport Timing Schemes

Having addressed how we plan to transport atoms from the MOT to the science chamber, we now turn to the question of how to time that transport. The timing and sequencing of operations significantly impacts both atomic flux and system complexity. Three distinct approaches are possible:

- Continuous Transport: In [19] the transport is implemented using a static, far-detuned dipole trap focused near the tweezer region. Atoms naturally migrate along the intensity gradient toward the science chamber. However, as argued earlier, while this approach works for Yb, it does not work for ^{133}Cs, so we do not pursue this.
- Pulsed Transport with MOT Shutdown: One approach with lattice transport is to load an atom bunch into the MOT, shut off the MOT, load that bunch into the lattice, transport it out of the MOT region, hand it off to the second lattice, and only then turn the MOT back on and begin loading a second bunch. This has the advantage of avoiding the intricacies of transporting atoms out of a running MOT, making it a clean solution. The Harvard Atom Array 2 system [20] uses exactly this approach.
- Ratchet-Type Transport with Continuous MOT Operation: An alternative approach maintains the MOT in continuous operation while the lattice periodically extracts atoms from the MOT center. In this scheme, as soon as a bunch of atoms has been transported approximately 2 mm away from the MOT center, a new loading cycle can begin immediately. The effective atomic flux is then:

  $$\text{Flux} = \frac{N_{\text{atoms}}}{\max(T_{\text{MOT}}^{\text{reload}}, T_{\text{extract}})} \tag{2.1}$$

  where $T_{\text{MOT}}^{\text{reload}}$ is the MOT reloading time and T_{extract} is the time required to move atoms out of the MOT capture region. By eliminating the dead time present in the Atom Array 2 approach (where the MOT remains off during the 72 ms of lattice transport and handover), this scheme could potentially achieve significantly higher flux, provided that downstream processing (such as loading atoms into tweezers or state preparation, as in [20]) is not the limiting bottleneck. Ratchet-type transport therefore represents a promising technique that we plan to implement in our system. However, extracting atoms from an actively running MOT requires careful management of the lattice-MOT interaction, which will be discussed in detail in section 3.1.

2.2 Introduction to MOT Theory

Before discussing the MOT design, it is important to first understand the underlying theoretical foundations and working principles of a MOT. In this chapter we first introduce the theory behind how to describe the interaction of polarized light with atoms in a magnetic field and then introduce the basics of how to model the forces in a MOT first analytically and then computationally.

2.2.1 Driven Hyperfine Structure Calculations

For a lot of the following simulations and theoretical considerations, it is important to understand how polarized light interacts with the hyperfine structure (HFS) of the ^{133}Cs D_2 line in the presence of an applied magnetic field. This section outlines the relevant calculations and underlying theory. While comprehensive derivations can be found in standard textbooks, such as [34, Chapters 5 and 7], a brief review is included here for completeness and to clarify a few subtle aspects that can otherwise cause confusion. We begin by establishing an atomic reference frame defined by the magnetic field quantization axis. In this framework, the coordinate system is consistently oriented such that the external magnetic field ($\vec{B}$) points along the positive z-axis. The propagation direction and polarization vectors of the incident laser light are then expressed relative to this coordinate system, as illustrated in Figure 2.6.

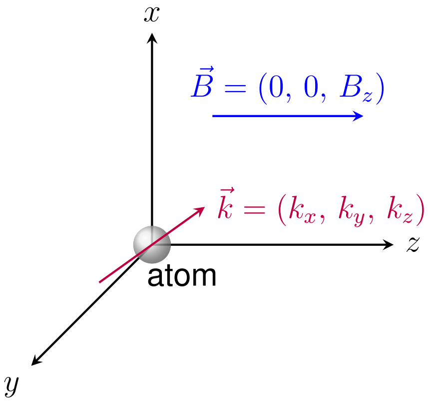

Figure 2.6: The “atomic coordinate system", where the z-axis is chosen such that the magnetic field points along the positive z-axis and the atom is at the origin. The wavevector $\vec{k}$ and polarization vector $\vec{\varepsilon}$ of the light are defined relative to this coordinate system.

The (real) electric field is written as:

$$\vec{\mathcal{E}} = \frac{E_0}{2}\left(\vec{\varepsilon}\, e^{-i\omega t + i\vec{k}\cdot\vec{R}} + \vec{\varepsilon}^{\,*} e^{+i\omega t - i\vec{k}\cdot\vec{R}}\right) \tag{2.2}$$

where the complex polarization vector $\vec{\varepsilon}$ satisfies the constraints:

$$\vec{\varepsilon}\cdot\vec{k} = 0, \qquad \|\vec{\varepsilon}\| = 1. \tag{2.3}$$

In Cartesian components:

$$\vec{\varepsilon} = \varepsilon_x \vec{e}_x + \varepsilon_y \vec{e}_y + \varepsilon_z \vec{e}_z. \tag{2.4}$$

For later calculations, it is important to know that we can expand $\vec{\varepsilon}$ in the spherical (circular and linear) basis:

$$\vec{\varepsilon} = \varepsilon_+ \vec{e}_+ + \varepsilon_0 \vec{e}_0 + \varepsilon_- \vec{e}_-, \tag{2.5}$$

where the basis vectors are defined as:

$$\vec{e}_+ = -\frac{1}{\sqrt{2}}\left(\vec{e}_x + i\vec{e}_y\right), \tag{2.6}$$

$$\vec{e}_- = \frac{1}{\sqrt{2}}\left(\vec{e}_x - i\vec{e}_y\right), \tag{2.7}$$

$$\vec{e}_0 = \vec{e}_z. \tag{2.8}$$

The intensity of the field is given by:

$$I = c\langle \epsilon_0 \vec{\mathcal{E}}^2 \rangle = \frac{1}{2} c\epsilon_0 E_0^2. \tag{2.9}$$

Now consider the Hamiltonian. The full system is described by:

$$\hat{\mathcal{H}} = \hat{H}_0 + \hat{H}_Z + \hat{H}_{\text{int}}. \tag{2.10}$$

Here, $\hat{H}_0 = \sum_m E_m |m\rangle\langle m|$ represents the unperturbed atomic Hamiltonian for however many hyperfine sublevels m we are interested in, fully described by $|m\rangle = |n, l, j, F, m_F\rangle$. In the case of the ^{133}Cs D_2 line, the entire system is described by 48 states. $\hat{H}_Z$ is the Zeeman Hamiltonian, and $\hat{H}_{\text{int}}$ is the laser interaction.

Zeeman and Interaction Hamiltonian

The Zeeman interaction with the external magnetic field $\vec{B}(\vec{r}) = B_z(\vec{r})\vec{e}_z$ is:

$$\hat{H}_Z = g_F \mu_B \hat{\vec{F}} \cdot \vec{B}(\vec{r}) \tag{2.11}$$

$$= g_F \mu_B \hat{F}_z B_z \tag{2.12}$$

where g_F is the Landé g-factor for the relevant state, μ_B is the Bohr magneton, and $\hat{\vec{F}}$ is the total angular momentum operator. For example, consider a ground manifold with $N_g = 3$ states ($m_F \in \{-1, 0, +1\}$) and a uniform magnetic field $\vec{B} = B_z\vec{e}_z$. The Hamiltonian simplifies to $\hat{H}_Z^g = g_F^g \mu_B B_z \hat{F}_z$. The matrix takes the form:

$$\hat{F}_z = \begin{pmatrix} -1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{2.13}$$

The interaction between the atom and the laser field is governed by the electric dipole operator. We adopt the semi-classical approximation, treating the atomic internal degrees of freedom quantum mechanically while treating the electromagnetic field as a classical wave.

The interaction Hamiltonian is:

$$\hat{H}_{\text{int}} = -\hat{\vec{d}} \cdot \vec{\mathcal{E}}(\vec{R}, t) \tag{2.14}$$

where:

- $\hat{\vec{d}} = -e\hat{\vec{r}}$ is the electric dipole operator. Here, $\hat{\vec{r}} = (\hat{x}, \hat{y}, \hat{z})$ represents the position operator of the electron relative to the nucleus.
- $\vec{\mathcal{E}}(\vec{R}, t)$ is the classical electric field vector evaluated at the position of the atom's center of mass $\vec{R}$.

By placing the atom at the origin $\vec{R} = \mathbf{0}$ and assuming the dipole approximation ($\vec{k} \cdot \hat{\vec{r}} \approx 0$), the interaction term simplifies to:

$$\hat{H}_{\text{int}} = \frac{eE_0}{2} \hat{\vec{r}} \cdot \left(\vec{\varepsilon} e^{-i\omega t} + \text{h.c.} \right). \tag{2.15}$$

Rabi Frequency

Now, consider a laser that is nearly resonant with a transition between two hyperfine manifolds $|n, l, j, F\rangle \to |n', l', j', F'\rangle$. To calculate the Rabi frequency for each individual transition $|F, m_F\rangle \rightarrow |F', m_{F'}\rangle$, we first define the projection operator:

$$\hat{P}_F = \sum_{m_F=-F}^{F} |n, l, j, F, m_F\rangle\langle n, l, j, F, m_F|, \tag{2.16}$$

which maps the entire Hilbert space onto the specific hyperfine level F. Projecting the total Hamiltonian $\hat{\mathcal{H}}$ into the subspace spanned by the ground manifold F and excited manifold F', we obtain the effective Hamiltonian $\hat{\mathcal{H}}_{\text{sub}}$:

$$\hat{\mathcal{H}}_{\text{sub}} = \begin{pmatrix} \hat{H}_0^F + \hat{H}_Z^F & \hat{V}_{\text{int}} \\ \hat{V}_{\text{int}}^\dagger & \hat{H}_0^{F'} + \hat{H}_Z^{F'} \end{pmatrix}, \tag{2.17}$$

where the coupling block is defined by the projection of the interaction Hamiltonian:

$$\hat{V}_{\text{int}} = \hat{P}_F \hat{H}_{\text{int}} \hat{P}_{F'}. \tag{2.18}$$

To remove the explicit time dependence of the laser field, we transition into the rotating frame by moving into the interaction picture $\hat{\mathcal{H}}_{\text{rot}} = \hat{U}\hat{\mathcal{H}}\hat{U}^\dagger + \hat{A}$, where we choose the operator $\hat{A}$ as a diagonal matrix representing an arbitrary energy offset. Here we choose the unitary transformation:

$$\hat{U}(t) = \exp\left(i\omega_L t \hat{P}_{F'}\right). \tag{2.19}$$

Additionally, we choose $\hat{A}$ such that we set the energy $E_{n,l,j,F}$ of the ground state manifold to 0. Now, applying the interaction picture plus the Rotating Wave Approximation (RWA), where we neglect terms oscillating at approximately $2\omega_L$, the Hamiltonian in the rotating frame becomes:

$$\hat{\mathcal{H}}_{\text{rot}} = \begin{pmatrix} \hat{H}_Z^F & \frac{\hat{\Omega}_{F,F'}}{2} \\ \frac{\hat{\Omega}_{F,F'}^\dagger}{2} & -\hbar\Delta\hat{I} + \hat{H}_Z^{F'} \end{pmatrix}. \tag{2.20}$$

Here, $\Delta = \omega_L - \omega_0$ is the laser detuning from the $F, m_F = 0 \rightarrow F', m_{F'} = 0$ transition, and

$$\hat{\Omega}_{F,F'} = \sum_{m_F, m_{F'}} \Omega_{n,l,j,F,m_F,n',l',j',F',m_{F'}} |n, l, j, F, m_F\rangle\langle n', l', j', F', m_{F'}| \tag{2.21}$$

where the individual matrix elements are given by:

$$\Omega_{n,l,j,F,m_F,n',l',j',F',m_{F'}} = \frac{eE_0}{\hbar}\langle n,l,j,F,m_F|\hat{\vec{r}}\cdot\vec{\varepsilon}|n',l',j',F',m_{F'}\rangle. \tag{2.22}$$

To evaluate this, we use the spherical decomposition of the scalar product $\vec{\varepsilon}\cdot\hat{\vec{r}} = \sum_q(-1)^q\varepsilon_{-q}\hat{r}_q$, allowing the application of the Wigner-Eckart theorem:

$$\langle n,l,j,F,m_F|\hat{r}_q|n',l',j',F',m_{F'}\rangle = \langle n,l,j,F||\hat{r}||n',l',j',F'\rangle\langle F',m_{F'};1,q|F,m_F\rangle \tag{2.23}$$

where $\langle n,l,j,F||\hat{r}||n',l',j',F'\rangle$ is the m_F-independent reduced matrix element and $\langle F',m_{F'};1,q|F,m_F\rangle$ is a Clebsch-Gordan coefficient, which can be explicitly calculated (see e.g. [34] equation (7.51)) or looked up in tables.

Substituting this back into the expression for the Rabi frequency yields the final form:

$$\Omega_{n,l,j,F,m_F,n',l',j',F',m_{F'}} = \frac{eE_0}{\hbar}\langle n,l,j,F||\hat{r}||n',l',j',F'\rangle\sum_{q=-1}^{1}(-1)^q\varepsilon_{-q}\langle F',m_{F'};1,q|F,m_F\rangle. \tag{2.24}$$

Both the Clebsch-Gordan coefficients ([34], equation (7.51)) and the reduced hyperfine matrix element ([34], equation (7.273)) can be explicitly and analytically computed.

These relations reveal exactly how light of a specific polarization interacts with the atom in the presence of an external magnetic field. By choosing the atomic coordinate system such that the magnetic field defines the quantization axis (as explained at the beginning of this section and illustrated in figure 2.6), and by correctly expressing the polarization vector $\vec{\varepsilon}$ from equation (2.2) in the spherical basis, equation (2.24) dictates precisely which transitions are driven, by which polarization components, and at what coupling strength.

Spontaneous Emission and Decay Rates

Having established how the laser drives transitions between specific states, we now turn to the spontaneous emission process to understand how an atom decays between two hyperfine levels. According to [30], the decay rate from an excited fine-structure level $|e\rangle = |n',l',j',m_j'\rangle$ to a ground level $|g\rangle = |n,l,j,m_j\rangle$ is given by:

$$\Gamma_{n',l',j',m'_j \to n,l,j,m_j} = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} |\langle n,l,j,m_j|\hat{\vec{d}}|n',l',j',m'_j\rangle|^2 \tag{2.25}$$

where $\hat{\vec{d}} = -e\hat{\vec{r}}$ is the electric dipole operator.

We now define the total decay rate out of a specific excited state $|n',l',j',m'_j\rangle$ by summing over all possible final ground states $|m_j\rangle$:

$$\Gamma_{n',l',j',m'_j} = \sum_{m_j} \Gamma_{n',l',j',m'_j \to n,l,j,m_j} \tag{2.26}$$

From there, it can be shown that for any $|m'_{j1}\rangle, |m'_{j2}\rangle$ within the excited manifold, $\Gamma_{m'_{j1}} = \Gamma_{m'_{j2}}$. Thus, the total decay rate from any excited fine-structure level to the lower manifold is uniform. The exact same principle applies to the hyperfine structure. For example, for the ^{133}Cs D_2 line, we define the total natural linewidth as $\Gamma_0 = \Gamma_{6\mathrm{P}_{3/2} \to 6\mathrm{S}_{1/2}} = 2\pi \times 5.234\,\mathrm{MHz}$ [30]. This Γ_0 is identical to the total decay rate from any hyperfine level $|F'\rangle$ of the excited state.

It can also be shown (e.g. [34] equation (11.33)) that the decay rate between two hyperfine manifolds is related to the reduced dipole matrix element by:

$$\begin{aligned}\Gamma_{F' \to F} &= \sum_{m_F} \Gamma_{F',m'_F \to F,m_F} \\ &= \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} \frac{2F+1}{2F'+1} |\langle n,l,j,F||\hat{\vec{d}}||n',l',j',F'\rangle|^2 \\ &= \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3} |\langle n',l',j',F'||\hat{\vec{d}}||n,l,j,F\rangle|^2\end{aligned} \tag{2.27}$$

For the first sum, it does not matter which m'_F we choose; as explained above, the decay rate is the same for all m'_F, but we sum over all ground m_F states. The term $|\langle n,l,j,F||\hat{\vec{d}}||n',l',j',F'\rangle|^2$ is the so-called reduced matrix element because it does not depend on the magnetic quantum numbers m_F and m'_F.

This implies that the decay rate from a specific excited hyperfine state to a specific ground hyperfine state is always less than or equal to the total fine-structure decay rate: $\Gamma_{J',F' \to J,F} \leq \Gamma_{J' \to J}$. Looking again at the ^{133}Cs D_2 line ($6\mathrm{P}_{3/2} \to 6\mathrm{S}_{1/2}$) as an example, dipole selection rules allow an excited state like $F' = 5$ to decay to only one specific lower hyperfine state ($F = 4$). In this case, $\Gamma_{F'=5 \to F=4} = \Gamma_0$. However, for an excited state like $F' = 4$, decays to multiple lower hyperfine states ($F = 4$ and $F = 3$)

are allowed. The individual decay rates to these specific lower states are strictly less than the total natural linewidth (i.e. $\Gamma_{F'=4\rightarrow F=4} < \Gamma_0$ and $\Gamma_{F'=4\rightarrow F=3} < \Gamma_0$). Nevertheless, their sum must equal the total decay rate: $\Gamma_{F'=4\rightarrow F=4} + \Gamma_{F'=4\rightarrow F=3} = \Gamma_0$. The relative decay probabilities are known as branching ratios and are determined by the square of the reduced matrix elements $|\langle n, l, j, F||\hat{\vec{d}}||n', l', j', F'\rangle|^2$. A figure showing the calculated branching ratios for the ^{133}Cs D_2 line can be found in figure 2.7.

The last step in equation (2.27) was done using the relation:

$$\langle F||\hat{\vec{d}}||F'\rangle = \langle F'||\hat{\vec{d}}||F\rangle^*(-1)^{F-F'}\sqrt{\frac{2F'+1}{2F+1}} \tag{2.28}$$

as shown in [34] equation (7.250). Exactly as we saw with the Rabi frequency (equation (2.24)), this reduced dipole matrix element can be analytically computed ([34], equation (7.273)).

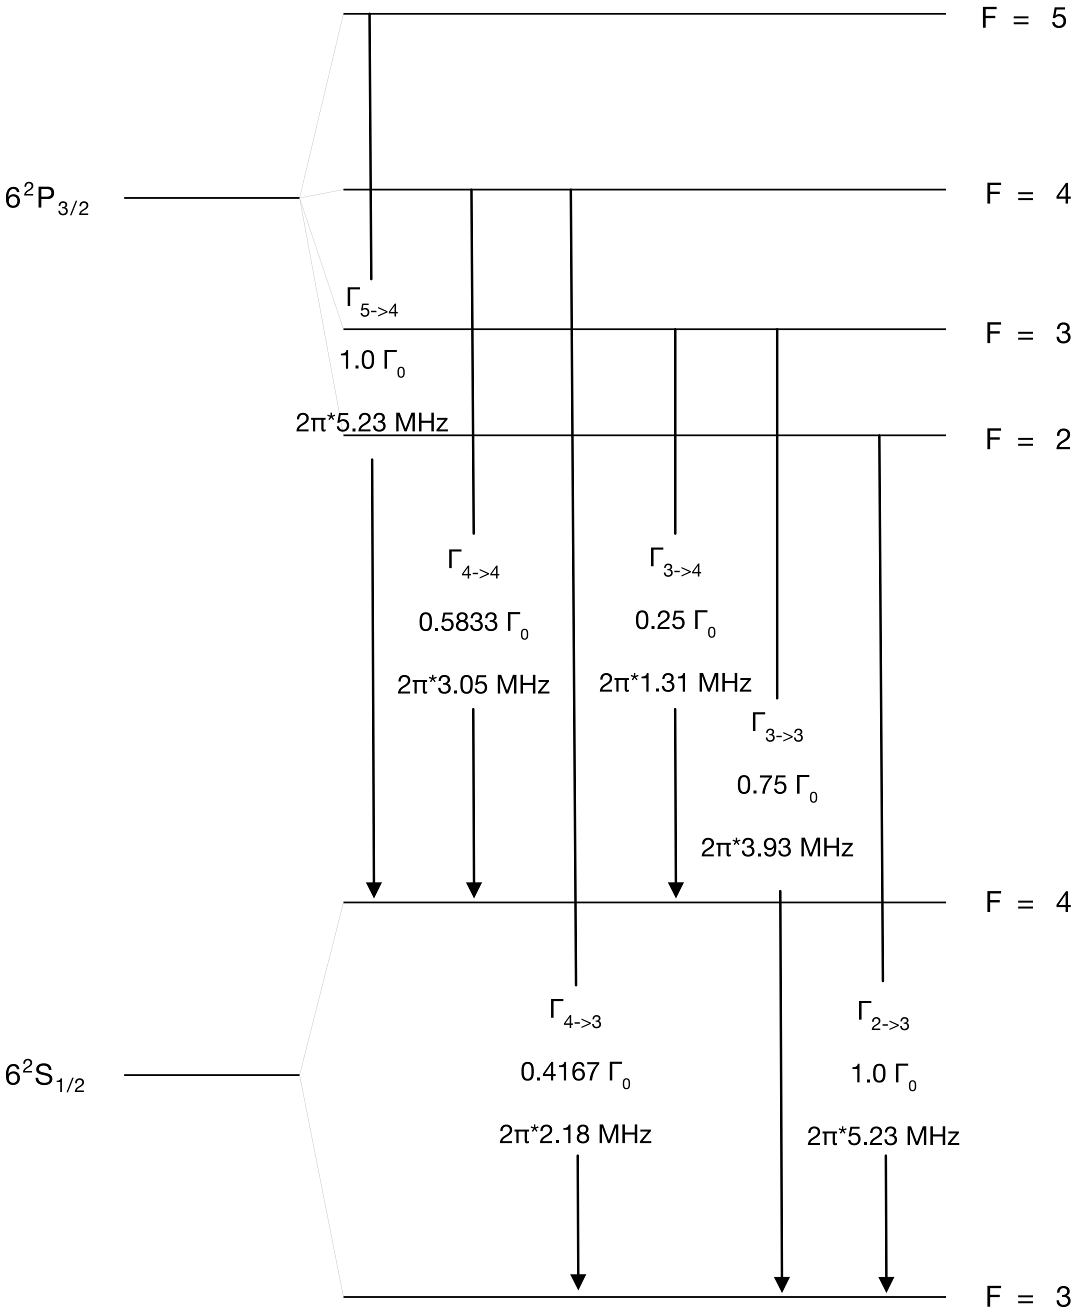

Figure 2.7: Hyperfine structure and decay branching ratios of the ^{133}Cs D_2 line ($6^2\mathrm{P}_{3/2} \rightarrow 6^2\mathrm{S}_{1/2}$). The diagram illustrates the allowed spontaneous emission channels and their corresponding decay rates as fractions of the total natural linewidth ($\Gamma_0 = 2\pi \times 5.23\,\mathrm{MHz}$). As described in the text, the cooling transition $F' = 5 \rightarrow F = 4$ is closed with a branching ratio of 1.0, whereas the $F' = 4$ state decays into both the $F = 4$ and $F = 3$ ground states with branching ratios of 0.5833 and 0.4167, respectively. Figure adapted from [30].

2.2.2 MOT Working Principle and Basic Formulas

A MOT fulfills two primary purposes: cooling and trapping atoms. The cooling aspect is achieved using three pairs of counter-propagating, red-detuned optical molasses beams. These beams cool the atoms via the Doppler cooling mechanism: due to the Doppler shift, moving atoms preferentially absorb photons from the laser beam opposing their motion. The atoms then spontaneously emit photons in random directions, resulting in a net momentum transfer that dampens the atomic velocity. While the optical molasses cools the atoms in all three dimensions, it lacks a restoring force. Without a position-dependent confinement mechanism, atoms would slowly diffuse out of the cooling region. To create this spatial confinement, a quadrupole magnetic field is applied, typically generated by a pair of anti-Helmholtz coils as illustrated in the 3D schematic in figure 2.8. This field introduces a spatially dependent Zeeman shift in the atomic energy levels, as shown in the 1D simplification in figure 2.9. Furthermore, the six intersecting counter-propagating laser beams are circularly polarized with opposite handedness (σ^+ and σ^-, for more details on the polarization scheme, see Appendix B). Due to this combination of the Zeeman shift and circular polarization, atoms displaced from the center preferentially absorb photons from the specific circularly polarized beam that pushes them back toward the trap center. This position-dependent absorption creates a robust restoring force, effectively trapping the cooled atoms. The mechanism is illustrated in figure 2.9

Fundamentally, the force exerted on an atom in a MOT is the result of momentum conservation during photon scattering. In the semi-classical regime, the total force is simply the vector sum of the radiation pressure forces from all interacting laser beams.

The net force is given by:

$$\vec{F}_{total} = \sum_j \hbar \vec{k}_j \times R_{scatt,j}(\vec{r}, \vec{v}) \tag{2.29}$$

where:

- $\vec{k}_j$ is the wavevector of the j-th laser beam.
- $R_{scatt,j}$ is the photon scattering rate for that specific beam.

While this formulation appears straightforward, the complexity lies entirely within the scattering rate R_{scatt}. In a realistic experimental scenario involving alkali atoms, the scattering rate depends on the internal multilevel hyperfine structure, the local polarization of the light field, and the coherence between atomic states.

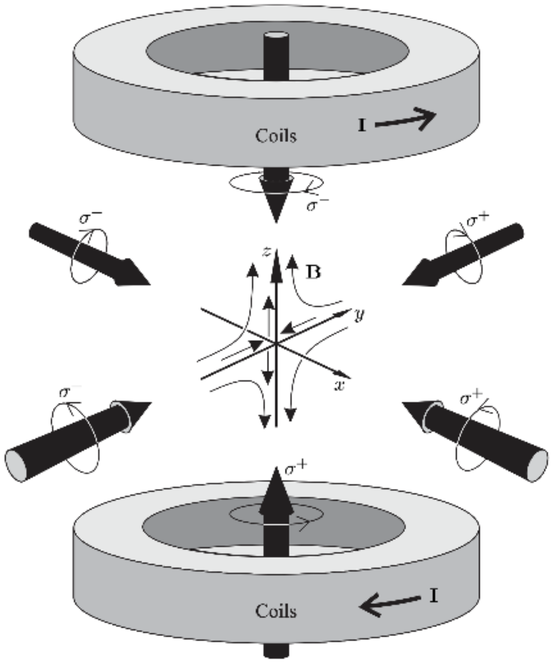

Figure 2.8: 3D configuration of coils and beams for a MOT. Two coils in an anti-Helmholtz configuration (with currents I flowing in opposite directions) create a quadrupole magnetic field with a zero-point at the origin. Three pairs of counter-propagating, red-detuned laser beams intersect at the trap center to provide three-dimensional cooling and spatial confinement. The circular polarizations (σ^+ and σ^-) are chosen to drive the appropriate Zeeman transitions relative to the local magnetic field axes, pushing displaced atoms back toward the trap center. Figure taken from [35], figure 9.9.

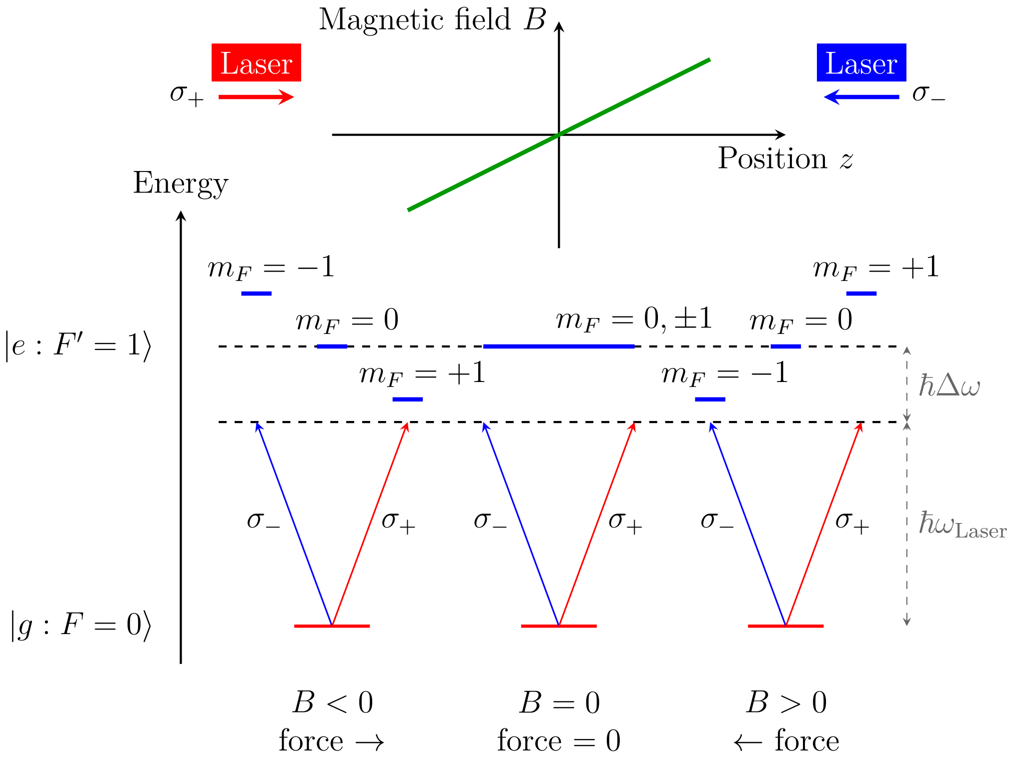

Figure 2.9: Simplified illustration of the operating principle of a MOT. For clarity, a transition between $F = 0$ and $F' = 1$ is assumed; in a realistic ^{133}Cs MOT, the $F = 4 \rightarrow F' = 5$ transition is used, involving significantly more m_F states where atoms are eventually pumped into the stretched states. The diagram illustrates the 1D spatial confinement achieved through the position-dependent Zeeman shift. This shift, combined with the σ^+/σ^- circular polarization of the counter-propagating beams, ensures that displaced atoms preferentially absorb photons from the beam that pushes them back toward the trap center. This figure primarily focuses on spatial trapping and does not explicitly show the velocity-dependent Doppler shift: for an atom moving toward a laser beam, the light is blue-shifted in the atom's rest frame, bringing the red-detuned laser closer to resonance. This results in a higher scattering rate from the beam opposing the motion, leading to an additional velocity-dependent damping force. Figure inspired from [36].

Derivation of the MOT force for a TLS

To derive an instructive analytic expression for the MOT force, we simplify the problem to a 1D model involving a hypothetical two-level system (TLS).

We apply the following approximations:

1. 1D Geometry: We consider motion along a single axis (denoted $\hat{z}$) defined by a pair of counter-propagating beams.
2. TLS: The atom is modeled as having a single ground state $|g\rangle$ and excited state $|e\rangle$.
3. Low Saturation: We assume the laser intensity is low ($s \ll 1$), allowing us to treat the scattering from the two beams independently.

For a two-level atom, the scattering rate follows a Lorentzian profile:

$$R_{scatt} = \frac{\Gamma}{2} \frac{s}{1 + s + (2\delta_{eff}/\Gamma)^2} \tag{2.30}$$

where Γ is the natural linewidth, $s = I/I_{sat}$ is the saturation parameter, and δ_{eff} is the effective detuning.

The magnetic field gradient creates a position-dependent Zeeman shift, while the atomic velocity creates a Doppler shift. For a standard quadrupole field approximated near the center ($B \approx b'z$), the effective detunings seen by the atom are:

$$\delta_{right} = \delta_0 - kv - \beta z \tag{2.31}$$

$$\delta_{left} = \delta_0 + kv + \beta z \tag{2.32}$$

Here, δ_0 is the laser detuning, and $\beta = \mu_B(g_e m_e - g_g m_g)b'/\hbar$ characterizes the Zeeman shift magnitude, with μ_B the Bohr magneton, g_e and g_g the Landé g-factors for the excited and ground states, and m_e and m_g the magnetic quantum numbers of the specific transition (typically the stretched states $m_g = -F_g$ and $m_e = -F_e$ for the cooling transition).

Summing the forces from the right-traveling ($+\hbar k$) and left-traveling ($-\hbar k$) beams:

$$F_{total} = \hbar k(R_{right} - R_{left}) \tag{2.33}$$

For small velocities ($kv \ll \Gamma$) and small displacements, we can Taylor expand the scattering rate around δ_0. This linearization yields the damped harmonic oscillator equation:

$$\boxed{F_{MOT} \approx -\alpha v - \kappa z} \tag{2.34}$$

The damping coefficient α is given by:

$$\alpha = -4\hbar k^2 s \frac{2\delta_0/\Gamma}{[1+s+(2\delta_0/\Gamma)^2]^2} \tag{2.35}$$

This result highlights the dual nature of the MOT: the velocity-dependent term ($-\alpha v$) provides the viscous damping (cooling), while the position-dependent term ($-\kappa z$) provides the restoring force (trapping).

It is important to note at this point, that the derivation above relies on the "semi-classical" approximation, where the atomic internal degrees of freedom are treated quantum mechanically, but the center-of-mass motion is treated classically. This is valid because the internal equilibration time ($\sim 1/\Gamma$) is significantly faster than the timescale of atomic motion. However, the TLS model derived here predicts a minimum temperature known as the Doppler Limit:

$$T_D = \frac{\hbar\Gamma}{2k_B} \tag{2.36}$$

For ^{78}Rb, this predicts $T_D \approx 146\,\mu\mathrm{K}$ and for ^{133}Cs $T_D \approx 125\,\mu\mathrm{K}$. In practice, experimental MOT temperatures are often found to be lower. Even when considering multi-level structures, this big difference cannot be explained. The reason are so-called sub-doppler cooling mechanisms, such as sisyphus cooling (polarization gradient cooling). In these mechanisms, atoms moving through the polarization gradient undergo optical pumping cycles that dissipate energy more efficiently than the doppler mechanism alone.

The Role of the Repumper

When modeling a ^{133}Cs MOT, the 852 nm D_2 transition from the $6S_{1/2}, F = 4$ ground state to the $6P_{3/2}, F' = 5$ excited state is typically treated as an ideal two-level system. Due to the $\Delta F = 0, \pm 1$ selection rule, the $6P_{3/2}, F' = 5$ state can only decay back to the $6S_{1/2}, F = 4$ state, forming a closed cycling transition. This makes it highly preferable for MOT operations, and it serves as the primary cooling transition in most ^{133}Cs experiments. However, the TLS approximation does not hold true in reality. The MOT cooling laser, while near-resonant with the $F = 4 \rightarrow F' = 5$ transition, can off-resonantly excite atoms from the $6S_{1/2}, F = 4$ ground state to the $6P_{3/2}, F' = 4$ excited state. Once in the $F' = 4$ state, the atom can spontaneously decay into either the $6S_{1/2}, F = 4$ or the $6S_{1/2}, F = 3$ ground state, according to their respective branching ratios. Consequently, the ideal two-level system description derived in the previous section breaks down. Because the $F = 3$ state is dark to the MOT cooling light, the atomic population will eventually accumulate there (as shown in figure 2.10), rendering the atoms invisible to the trapping forces. To counteract this optical pumping into the dark state, a repumper laser is required. In principle, any laser that couples the $6S_{1/2}, F = 3$ level to an excited state capable of decaying back into the $F = 4$ ground state can be used. A standard choice is to drive the $6S_{1/2}, F = 3 \rightarrow 6P_{3/2}, F' = 4$ transition [4, 30, 37]. To quantify this leakage into the dark state, we perform a quantum Monte Carlo (QMC) wave function simulation of the multilevel dynamics (as shown in figure 2.10). This simulation models a four-level system comprising the ground states $|0\rangle \equiv 6S_{1/2}, F = 3$ and $|1\rangle \equiv 6S_{1/2}, F = 4$, along with the excited states $|2\rangle \equiv 6P_{3/2}, F' = 4$ and $|3\rangle \equiv 6P_{3/2}, F' = 5$. The system is driven by a red-detuned MOT laser near-resonant with the $|1\rangle \leftrightarrow |3\rangle$ transition. The QMC wave function method treats the atom-light interaction as a stochastic process: the system evolves deterministically under an effective non-Hermitian Hamiltonian, interrupted by random quantum jumps that represent spontaneous emission events. The coherent evolution is governed by the effective non-hermitian Hamiltonian:

$$H_{\text{eff}} = H_0 - \frac{i\hbar}{2} \sum_k \hat{L}_k^\dagger \hat{L}_k \tag{2.37}$$

where H_0 is the coherent part describing laser coupling in the rotating wave approximation:

$$H_0 = \hbar \begin{pmatrix} 0 & 0 & -\frac{\Omega_{02}}{2} & -\frac{\Omega_{03}}{2} \\ 0 & \omega_{01} & -\frac{\Omega_{12}}{2} & -\frac{\Omega_{13}}{2} \\ -\frac{\Omega_{02}^*}{2} & -\frac{\Omega_{12}^*}{2} & \delta_2 & 0 \\ -\frac{\Omega_{03}^*}{2} & -\frac{\Omega_{13}^*}{2} & 0 & \delta_3 \end{pmatrix} . \tag{2.38}$$

Here, $\delta_2 = \omega_{02} - \omega_L$ and $\delta_3 = \omega_{03} - \omega_L$ are the detunings from the excited states, where ω_L is the laser frequency, and Ω_{ij} are the complex Rabi frequencies for the respective transitions. The jump operators $\hat{L}_k$ represent the spontaneous decay channels from the excited states:

$$\hat{L}_{02} = \sqrt{\Gamma_{02}}\, |0\rangle\langle 2| , \qquad \hat{L}_{03} = \sqrt{\Gamma_{03}}\, |0\rangle\langle 3| \tag{2.39}$$

$$\hat{L}_{12} = \sqrt{\Gamma_{12}}\, |1\rangle\langle 2| , \qquad \hat{L}_{13} = \sqrt{\Gamma_{13}}\, |1\rangle\langle 3| \tag{2.40}$$

The wave function evolves deterministically between quantum jumps according to:

$$|\psi(t + \mathrm{d}t)\rangle = \frac{\exp(-iH_{\text{eff}}\mathrm{d}t/\hbar)\, |\psi(t)\rangle}{\| \exp(-iH_{\text{eff}}\mathrm{d}t/\hbar)\, |\psi(t)\rangle \|} . \tag{2.41}$$

At each time step, quantum jumps occur stochastically with probabilities $p_k = \Gamma_k \langle\psi|\hat{L}_k^\dagger \hat{L}_k|\psi\rangle \, \mathrm{d}t$. When a jump occurs on channel k, the wave function undergoes an instantaneous transition $|\psi\rangle \to \hat{L}_k |\psi\rangle / \| \hat{L}_k |\psi\rangle \|$. Averaging over many such trajectories reproduces the master equation dynamics while preserving the quantum character of individual measurement outcomes.

Figure 2.10 shows the population dynamics for an atom initially prepared in the $F = 4$ ground state and exposed to the MOT laser without a repumper beam, using typical experimental parameters. The laser is detuned by $\delta_3 = -2\Gamma = -10.44\,\mathrm{MHz}$ (which corresponds to $-65.63 \times 10^6\,\mathrm{rad/s}$) from the $|1\rangle \to |3\rangle$ ($F = 4 \to F' = 5$) transition. Assuming a saturation parameter $s_0 \approx 1$, this leads to a Rabi frequency of:

$$\Omega_{13} = \sqrt{\frac{s_0 \Gamma^2}{2}} = \sqrt{\frac{1 \cdot (32.8150 \times 10^6\,\mathrm{rad/s})^2}{2}} \approx 23.2037 \times 10^6\,\mathrm{rad/s} \tag{2.42}$$

The results clearly demonstrate optical pumping into the dark $F = 3$ state: the population in $F = 4$ (where atoms experience the cooling force) steadily decreases, while population accumulates in $F = 3$ (which is decoupled from the MOT light). In Figure 2.10, one can observe a brief initial period of coherent Rabi oscillations, after which the system stabilizes into a quasi-steady state. It remains in this state until it undergoes spontaneous decay, at which point the atom either falls into the dark state or the coherent Rabi oscillations restart.

To get a rough estimate for the decay time, one can analytically model this process by assuming an effective pumping rate $\Gamma_{\mathrm{pump}} = \rho_{ee}\Gamma_{\mathrm{leak}}$. Here, the leak rate into the dark state is determined by the total natural linewidth and the branching ratio: $\Gamma_{\mathrm{leak}} = \Gamma(|2\rangle \to |0\rangle) = 2\pi \times 5.2227\,\mathrm{MHz} \times 0.4167 \approx 2\pi \times 2.1763\,\mathrm{MHz}$. The excited state population ρ_{ee} in the off-resonantly driven $|2\rangle$ state ($F' = 4$) can be approximated using the steady-state solution of the two-level Optical Bloch Equations:

$$\rho_{ee} = \frac{\Omega_{12}^2}{\Gamma^2 + 2\Omega_{12}^2 + 4\Delta_{12}^2} \approx 0.00341 \tag{2.43}$$

This assumes that the $|1\rangle \leftrightarrow |2\rangle$ transition reaches a quasi-steady state much faster than the population leaks into the $|0\rangle$ state. Despite treating a multi-level structure as an effective TLS with a slow leak, this approximation turns out to be quite accurate. Inserting the simulation parameters yields an analytical lifetime of $\tau = 1/\Gamma_{\mathrm{pump}} \approx 21.46\,\mathrm{\mu s}$, which fits the simulated ensemble average very well, as indicated by the black dashed line in Figure 2.10.

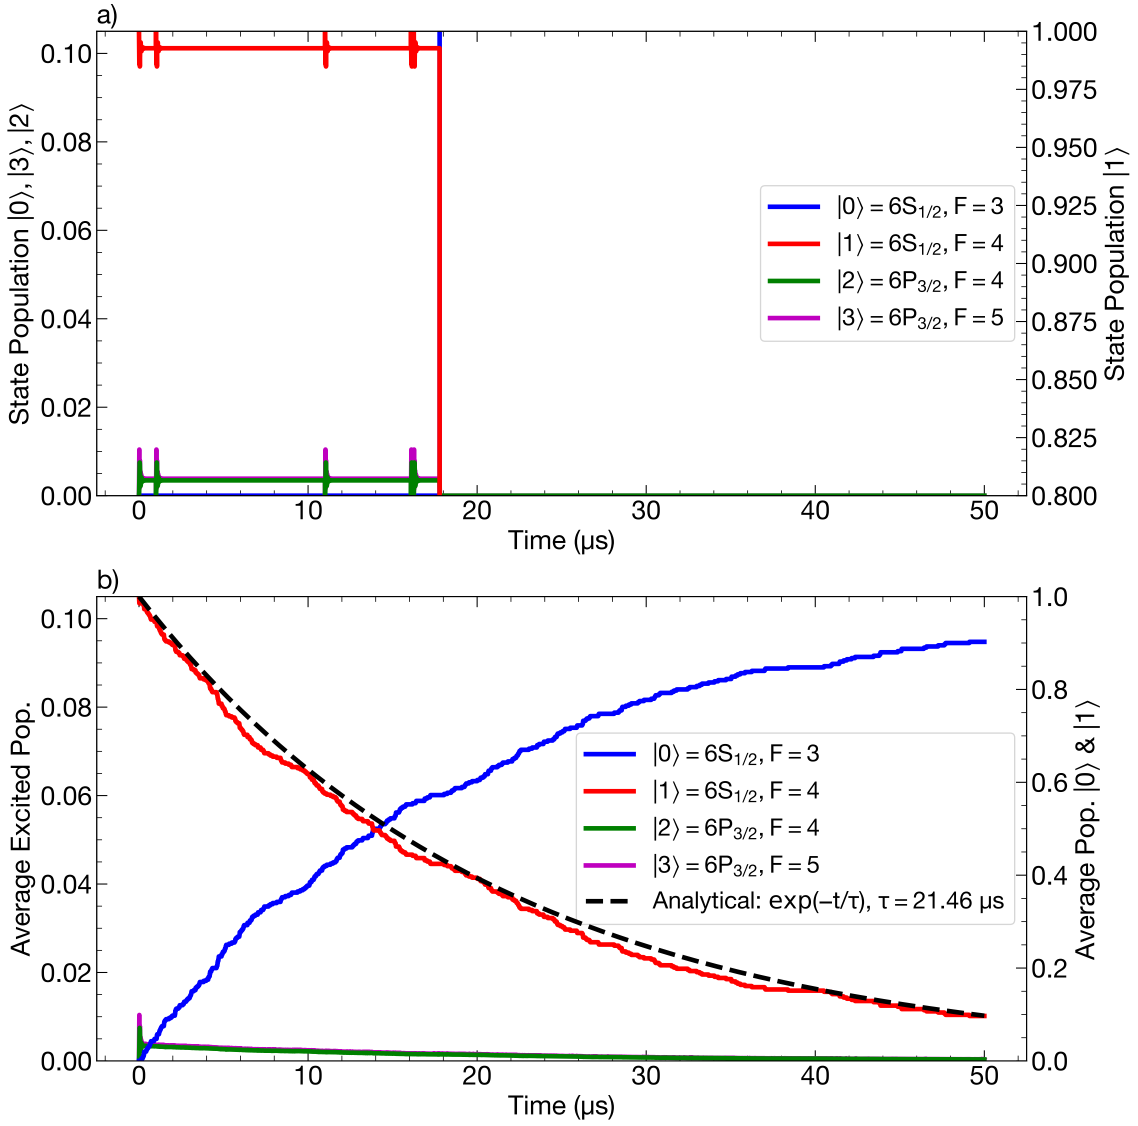

Figure 2.10: Quantum Monte Carlo simulation showing population dynamics in a four-level ^{133}Cs system driven by the MOT laser resonant to the $|1\rangle \to |3\rangle$ ($F = 4 \to F' = 5$) transition with $\delta_3 = -2\Gamma = -10.44\,\mathrm{MHz}$ and $s_0 \approx 1$ without a repumper. Top: Single trajectory showing a stochastic quantum jump between states. Bottom: Ensemble average over 200 trajectories. The atom starts in $|1\rangle \equiv F = 4$ (red) and gradually accumulates in the dark $|0\rangle \equiv F = 3$ state (blue) through off-resonant excitation to $|2\rangle \equiv F' = 4$ (green) followed by decay. The excited state populations (green and purple) remain small throughout. Without a repumper beam to return atoms from $F = 3$ to $F = 4$, the MOT progressively loses atoms from the cooling cycle.

Rate Equation Model

To rigorously describe the dynamics of open quantum systems (systems that are not isolated but instead interact with an external environment, leading to dissipation and decoherence) one typically employs Markovian master equations, such as the Lindblad equation [38–40]. This dictates the time evolution of the system's density matrix ρ, ensuring that the map is trace-preserving and completely positive so that ρ remains a valid physical state at all times. However, for a Hilbert space of dimension N, the density matrix contains N^2 elements. This quadratic scaling leads to a system of coupled differential equations that quickly becomes computationally very hard for large quantum systems. For certain effects discussed later in this thesis, it becomes necessary to simulate the individual Zeeman sublevels. The primary cooling transition alone (considering only the $6^2S_{1/2}, F = 4$ ground state and the $6^2P_{3/2}, F' = 5$ excited state) comprises 20 distinct Zeeman states. To overcome the computational bottleneck of the full master equation, we introduce an approach that scales much more favorably: the rate equation model. This method reduces computational complexity by neglecting the off-diagonal terms of the density matrix (coherences) and tracking only the diagonal terms (populations). While rigorous derivations exist (e.g. via the secular approximation), rate equations are frequently constructed heuristically. To ensure accuracy within this simplified model, the rate coefficients are often calibrated such that the steady-state solution of the rate equations perfectly matches that of the full Lindblad equation.

Example 1: Open, driven TLS rate equation

To illustrate the connection between the Lindblad equation and the rate equation, we will try to derive the rate equation for a TLS from the Lindblad equation. This motivates that a very similar approach can be taken for larger system size. Consider a ground state $|g\rangle$ and an excited state $|e\rangle$. The system is driven by a laser with Rabi frequency Ω and detuning Δ, and decays via spontaneous emission at a rate Γ.

The time evolution is governed by the Lindblad master equation:

$$\dot{\rho} = -i[H, \rho] + \mathcal{L}[\rho], \tag{2.44}$$

where the Hamiltonian in the rotating frame is $H = -\frac{\Delta}{2}\sigma_z + \frac{\Omega}{2}\sigma_x$. Solving the resulting Optical Bloch Equations for the steady state ($\dot{\rho} = 0$) yields the excited state population $P_e = \rho_{ee}$:

$$P_e^{\mathrm{ss,Lindblad}} = \frac{\Omega^2/4}{\Delta^2 + \Gamma^2/4 + \Omega^2/2}. \tag{2.45}$$

Defining the on-resonance saturation parameter as $s = \frac{2\Omega^2}{\Gamma^2}$, the dimensionless steady-state population is:

$$P_e^{\text{ss,Lindblad}} = \frac{s/2}{1 + s + 4\Delta^2/\Gamma^2}. \tag{2.46}$$

The rate equation for this same TLS focuses only on populations, assuming a balance between flows into and out of the excited state (P_e):

$$\dot{P}_e = 0 \quad \Longrightarrow \quad R_{\text{spont}} + R_{\text{stim,em}} = R_{\text{stim,abs}}. \tag{2.47}$$

Using the phenomenological stimulated rate α, we have $\Gamma P_e + \alpha P_e = \alpha P_g$. This leads to the steady-state solution:

$$P_e^{\text{ss,rate}} = \frac{\alpha}{\Gamma + 2\alpha}. \tag{2.48}$$

By demanding that $P_e^{\text{ss,rate}} = P_e^{\text{ss,Lindblad}}$, we calibrate the heuristic rate α to match the quantum result:

$$\alpha_{\text{heuristic}} = \frac{\Gamma s}{2(1 + 4\Delta^2/\Gamma^2)} = \frac{\Omega^2}{\Gamma}\frac{1}{1 + 4\Delta^2/\Gamma^2}. \tag{2.49}$$

The rate equation can be directly derived from the Lindblad equations in the limit where the coherences decay much faster than the populations, known as the secular approximation or adiabatic elimination. The equations of motion for the population P_e and the coherence ρ_{ge} are:

$$\begin{aligned}
\dot{P}_e &= -\Gamma P_e + i\frac{\Omega}{2}(\rho_{ge} - \rho_{eg}) \\
\dot{\rho}_{ge} &= -\left(i\Delta + \frac{\Gamma}{2}\right)\rho_{ge} + i\frac{\Omega}{2}(P_e - P_g)
\end{aligned}$$

In the fast-decoherence limit, we set $\dot{\rho}_{ge} \approx 0$ to adiabatically eliminate the coherence:

$$\rho_{ge} \approx \frac{i\Omega/2}{i\Delta + \Gamma/2}(P_e - P_g)$$

Substituting this expression back into the population equation for $\dot{P}_e$:

$$\begin{aligned}
\dot{P}_e &= -\Gamma P_e + i\frac{\Omega}{2}\left[\frac{i\Omega/2}{i\Delta+\Gamma/2}(P_e - P_g) - \text{c.c.}\right] \\
&= -\Gamma P_e - \frac{\Omega^2}{4}(P_e - P_g)\left[\frac{1}{i\Delta+\Gamma/2} + \frac{1}{-i\Delta+\Gamma/2}\right] \\
&= -\Gamma P_e - \frac{\Omega^2}{4}(P_e - P_g)\left[\frac{\Gamma}{\Delta^2+\Gamma^2/4}\right] \\
&= -\Gamma P_e + \underbrace{\left[\frac{\Omega^2\Gamma}{\Gamma^2+4\Delta^2}\right]}_{\alpha_{\text{rigorous}}}(P_g - P_e)
\end{aligned}$$

Factoring the Γ term out of the denominator yields:

$$\alpha_{\text{rigorous}} = \frac{\Omega^2}{\Gamma}\frac{1}{1+4\Delta^2/\Gamma^2}. \tag{2.50}$$

Comparing equation (2.49) and (2.50) shows that the rate coefficient derived from the adiabatic elimination of coherences is identical to the rate coefficient obtained by matching the steady-state solutions. The rate equation is therefore the accurate population dynamics model in the fast-decoherence limit. A relevant remaining question is, when can we actually ignore coherences? This approximation generally holds true in two specific regimes. First, in the overdamped limit. If the decoherence rate (transverse relaxation rate, Γ_2) is significantly faster than the driving rate (Ω), the off-diagonal elements of the density matrix decay almost instantly. In this limit, $\Omega \ll \Gamma_2$, the system never completes a coherent Rabi oscillation, and the dynamics are overdamped, allowing the population-only (rate) description to be accurate. The second case where this approximation is true is for broadband (so incoherent) driving light. If an atom is driven by a highly monochromatic laser, the phase relationship between the atom and the light is preserved. This leads to Rabi oscillations. A rate equation cannot capture this oscillation, it only captures the exponential approach to equilibrium. On the other hand, if the incoming light has a broad spectrum (a short coherence time compared to the atomic dynamics), the phase of the light fluctuates rapidly. These fluctuations average out the atomic coherences ($\rho_{ge} \to 0$), making the population description valid. In our case for the MOT, we are not strictly in either of these extreme limits. For the ^{133}Cs D_2 line, the natural linewidth is $\Gamma = 2\pi \times 5.22\,\text{MHz}$, resulting in a transverse relaxation rate of $\Gamma_2 = \Gamma/2 \approx 2\pi \times 2.61\,\text{MHz}$. Standard MOTs operate near the saturation intensity, yielding a Rabi frequency of

$\Omega \approx \Gamma/\sqrt{2} \approx 2\pi \times 3.7\,\mathrm{MHz}$. Because $\Omega \sim \Gamma_2$, the system is not strongly overdamped, and transient Rabi oscillations do occur. The internal transient dynamics (the Rabi oscillations and decaying coherences) damp out on the timescale of the excited state lifetime ($1/\Gamma \approx 30$ ns). In contrast, the actual mechanical cooling process, where the atom accumulates enough photon momenta to slow down, occurs over microseconds to milliseconds. As the atom moves through the trap, its internal state effectively reaches equilibrium instantaneously compared to its physical motion. Therefore, by calibrating our rate equations to match the exact steady-state of the Lindblad model, we can safely ignore the nanosecond-scale transients and accurately simulate the long-term scattering forces that govern the MOT dynamics.

Example 2: Open, driven 20-Level ^{133}Cs System rate equation

Having seen the example of how to derive and think about the rate equation for a TLS, we can apply a very similar approach for ^{133}Cs. The system we aim to simulate is the ^{133}Cs D$_2$ transition, specifically the $6^2S_{1/2}, F = 4 \rightarrow 6^2P_{3/2}, F' = 5$ cycling transition. This constitutes a 20-level system ($N = N_g + N_e = 9 + 11 = 20$). To simulate the motion of atoms in this multilevel MOT, the standard method is the semi-classical approximation. This approach treats the atomic internal degrees of freedom quantum mechanically while treating the external degrees of freedom (position $\vec{r}$ and momentum $\vec{p}$) classically. This approximation is justified under two conditions. First, the atom must be sufficiently localized such that its thermal de Broglie wavelength is negligibly small compared to the relevant spatial variations of the environment, allowing us to approximate the position operator $\hat{r}$ with a classical vector $\vec{r}(t)$. Assuming the ^{133}Cs atoms are cooled to the Doppler temperature $T_D \approx 125\,\mu\mathrm{K}$ and their velocities follow a Maxwell-Boltzmann distribution, they have a most probable speed of $v = \sqrt{\frac{2k_BT_D}{m}} \approx 0.125\,\mathrm{m/s}$. This corresponds to a thermal de Broglie wavelength of roughly $\lambda = \frac{h}{mv} \approx 24\,\mathrm{nm}$. This wavelength is vastly smaller than the characteristic length scale of the MOT trapping potential, which is typically on the order of millimeters, and comfortably smaller than the 852 nm wavelength of the trapping light itself. And second, if the driving field can be assumed to be a classical field. The laser field is a coherent state with a large average photon number, allowing us to treat the electric field classically (negligible quantum fluctuations of the field). Note that this assumption fails inside high-finesse cavities, where the field must be quantized.

To create the semi-classical Hamiltonian of this system, one simply has to follow the steps detailed in 2.2.1. However, a full quantum treatment requires solving for the density matrix ρ, which has dimensions $N \times N = 400$ complex elements. This makes a full quantum simulation computationally expensive, scaling as N^4 or worse depending on the solver. While simpler models (e.g. an effective $F = 1 \to F' = 2$ system) are often used as a first approximation, we aim to model the full 20-level structure to capture the correct Zeeman sub-structure and polarization gradients. As an alternative, we can employ the already mentioned rate equation model. Tarbutt [41] and Eckel [42] present an explicit multistate rate equation model for simulating a MOT. While they do not derive it, this model can in principle be derived in a very similar way with adiabatic elimination as we saw for the two level system in Example 1. In the following we will go through their rate equation shortly, as this will become relevant again in section 3.1.1.

For the rate equation, one first needs the scattering rate $R_{l,u,p}$ between a lower level l and an upper level u driven by a laser component p. With this scattering rate, one can then derive the rate equation for the excited states and the ground states separately. So starting with the scattering rate, Tarbutt (equation (2) in [41]) defines the excitation rate $R_{l,u,p}$ between a lower level l and an upper level u driven by a laser component p as:

$$R_{l,u,p} = \frac{\Gamma}{2} \frac{f_{l,u,p} s_p}{1 + 4\left(\frac{\delta_{l,u,p}}{\Gamma}\right)^2} \tag{2.51}$$

Here, s_p is the saturation parameter, $f_{l,u,p}$ is the fractional transition strength and Γ is the natural linewidth Γ_0 of the excited state. As noted in the text of [41] (below equation (3)), these parameters are directly related to the Rabi frequency $\Omega_{l,u,p}$:

$$f_{l,u,p} s_p = \frac{2\Omega_{l,u,p}^2}{\Gamma^2} \tag{2.52}$$

Substituting this identity into Eq. 2.51, we obtain the more familiar Lorentzian form used in quantum optics:

$$\begin{aligned} R_{l,u,p} &= \frac{\Gamma}{2} \frac{\left(\frac{2\Omega_{l,u,p}^2}{\Gamma^2}\right)}{1 + \frac{4\delta_{l,u,p}^2}{\Gamma^2}} \\ &= \frac{\Omega_{l,u,p}^2}{\Gamma^2 + 4\delta_{l,u,p}^2} \Gamma \end{aligned}$$

Thus, adapting this to our notation (ground state $l_{Tarbutt} \to g_j$, excited state $u_{Tarbutt} \to$

e_n, laser component $p_{Tarbutt} \to l$), the scattering rate is:

$$R^l_{n,j} = \frac{(\Omega^l_{n,j})^2}{\Gamma^2 + 4(\delta^l_{n,j})^2}\Gamma \tag{2.53}$$

where $\delta_{n,j,l} = \omega_l - \omega_{n,j} - \vec{k}_l \cdot \vec{v} - \Delta\omega^{\text{Zeeman}}_{n,j}$ is the total detuning including Doppler and Zeeman shifts.

With this scattering rate, the population evolution of an excited state $|e_n\rangle$ can be derived. It is modeled as the sum of three distinct physical processes. We must sum over all possible ground states j and all laser components l (Eq. 2 in [41]):

1. Stimulated Absorption (+): The laser beams drive atoms from any ground state $|g_j\rangle$ up to the excited state $|e_n\rangle$.

$$\text{Term:} \quad + \sum_{j,l} R^l_{n,j} P_{g_j}$$

2. Stimulated Emission (-): The laser beams drive atoms from the excited state $|e_n\rangle$ back down to the ground states.

$$\text{Term:} \quad - \sum_{j,l} R^l_{n,j} P_{e_n}$$

3. Spontaneous Emission (-): The atom naturally decays from the excited state to the ground manifold at rate Γ. This process is independent of the laser field l.

$$\text{Term:} \quad - \Gamma P_{e_n}$$

Combining these yields the differential equation:

$$\frac{d}{dt} P_{e_n} = \sum_{j,l} R^l_{n,j} (P_{g_j} - P_{e_n}) - \Gamma P_{e_n} \tag{2.54}$$

Similarly, the population of a ground state $|g_j\rangle$ evolves according to the inverse processes (equation (1d) in [41]):

1. Spontaneous Decay Gain (+): Population falls into state $|g_j\rangle$ from the various excited states. This depends on the specific branching ratio β_{nj} for the transition $n \to j$.

$$\text{Term:} \quad + \sum_{n} \Gamma \beta_{nj} P_{e_n}$$

2. Net Stimulated Transfer (-): The laser beams drive population out of $|g_j\rangle$ (absorption) and back into it (stimulated emission).

$$\text{Term:} \quad -\sum_{n,l} R^l_{n,j}(P_{g_j} - P_{e_n})$$

Combining these yields the ground state equation:

$$\frac{d}{dt}P_{g_j} = \sum_n \Gamma\beta_{nj}P_{e_n} + \sum_{n,l} R^l_{n,j}(P_{e_n} - P_{g_j}) \tag{2.55}$$

With equations (2.54) and (2.55), we have successfully established the coupled differential equations for the ground state and excited state populations. In section 3.1, we will use these equations to calculate the scattering MOT force on the atoms, as P_g and P_e allow us to calculate the scattering rates at each individual level and, consequently, the total MOT force (see equation (3.15)).

Chapter 3

Deriving MOT Requirements

Having established the architectural basics of our continuous reloading experiment and the theoretical foundations/working principle of a MOT, we can now outline the critical design considerations for the MOT:

1. High atomic density and total atom number to ensure efficient loading into the optical lattice.
2. Low atomic temperature to minimize thermal losses during lattice transport.
3. Rapid reloading times to maintain the continuous operation of the experiment.
4. Mitigation of loss mechanisms arising from interactions between the atoms and the MOT magnetic or optical fields during extraction. This applies to both pulsed and ratchet-type lattice transport schemes.

In the following section, we first examine the primary loss mechanisms that must be addressed. Subsequently, we dive into MOT physics to explore how the aforementioned parameters can be systematically optimized.

3.1 Loss and Heating Mechanisms

Successfully extracting atoms from the MOT into a lattice without introducing excessive loss or heating requires careful consideration. Two primary loss mechanisms must be addressed. First, during ratchet-type transport, atoms must be pulled out of an actively running MOT while the magnetic field gradient and cooling beams remain on. This necessitates holding the atoms statically off-center within the MOT, which can lead to substantial heating, as detailed below. Second, particularly for cesium, spin-changing collisions during lattice transport can cause significant atom loss and heating if the atomic ensemble is not properly prepared in the absolute m_F ground state. The following subsections examine these two effects in greater detail.

3.1.1 Static Off-Center Holding in the MOT

During semi-classical simulations of lattice transport, previous work in our group by Zunqi Li revealed that statically holding atoms in regions where the local magnetic field exceeds a critical threshold, B_{critical}, leads to significant atom loss. This finding is particularly relevant for our planned ratchet-type transport scheme, which requires holding atoms off the MOT center for extended periods. This atom loss cannot be attributed solely to the velocity-independent, position-dependent force $F(v = 0, r)$. Because the optical lattice potential should be sufficiently deep to confine atoms to a stable equilibrium position without significant spatial displacement, Zunqi hypothesized that the loss is instead driven by an anomalous, position-dependent damping effect. Specifically, he proposed that the total force includes a velocity-dependent damping term $\alpha_d(r)$ that varies strongly with the atom's position in the magnetic field:

$$F(v, r) = F(v = 0, r) - \alpha_d(r)v \tag{3.1}$$

Analytical Two-Level Model

To understand the origin of the position-dependent damping coefficient $\alpha_d(r)$, we perform a Taylor expansion of the MOT force for a TLS. Starting from the total force given by equation (2.29) and using the scattering rate from equation (2.30), we write:

$$F = F_+ - F_- = F_+(v) - F_-(v) \tag{3.2}$$

where $F_\pm$ denotes the force from the right-propagating (+) and left-propagating (−) beams. Using the effective detunings from equations. (2.31) and (2.32):

$$\delta_+ = \delta_0 - kv + \beta z \tag{3.3}$$

$$\delta_- = \delta_0 + kv - \beta z\,, \tag{3.4}$$

where $\beta = \mu' b'/\hbar$ characterizes the Zeeman shift due to the magnetic field gradient $b' = dB/dz$.

The force from each beam is:

$$F_\pm = \pm\hbar k \frac{\Gamma}{2}\frac{s_0}{1 + s_0 + (2\delta_\pm/\Gamma)^2} \tag{3.5}$$

The net force becomes:

$$F = \hbar k \frac{\Gamma}{2} s_0 \left[\frac{1}{1 + s_0 + (2\delta_+/\Gamma)^2} - \frac{1}{1 + s_0 + (2\delta_-/\Gamma)^2}\right] \tag{3.6}$$

$$= \hbar k \frac{\Gamma}{2} s_0 \left[\frac{-8\delta_+/\Gamma^2}{(1 + s_0 + (2\delta_+/\Gamma)^2)^2}(-kv) - \frac{-8\delta_-/\Gamma^2}{(1 + s_0 + (2\delta_-/\Gamma)^2)^2}(+kv)\right] . \tag{3.7}$$

Expanding to first order in velocity v around position z (or equivalently, around the effective detunings $\delta_\pm(z, v = 0)$):

$$F(v, z) \approx F(v = 0, z) + \hbar k^2 \frac{\Gamma}{2} s_0 v \left[\frac{-8\delta_+/\Gamma^2}{(1 + s_0 + (2\delta_+/\Gamma)^2)^2} + \frac{8\delta_-/\Gamma^2}{(1 + s_0 + (2\delta_-/\Gamma)^2)^2}\right] \tag{3.8}$$

This yields the form:

$$F(v, z) = F(v = 0, z) - \alpha_d(z)v \tag{3.9}$$

where the damping coefficient is:

$$\boxed{\alpha_d(z) = -\hbar k^2 \frac{\Gamma}{2} s_0 \left[\frac{8\delta_+/\Gamma^2}{(1 + s_0 + (2\delta_+/\Gamma)^2)^2} + \frac{8\delta_-/\Gamma^2}{(1 + s_0 + (2\delta_-/\Gamma)^2)^2}\right]} \tag{3.10}$$

Since $\delta_+ = \delta_0 + \beta z$ and $\delta_- = \delta_0 - \beta z$ (at $v = 0$), the damping coefficient explicitly depends on position through the Zeeman-shifted detunings.

For the two-level case, plotting the damping coefficient $\alpha_d(z)$ as a function of magnetic field (or equivalently, as a function of position x in the quadrupole field) reveals that α_d becomes negative at a certain critical field strength. When $\alpha_d < 0$, the damping force reverses sign, becoming a heating force that amplifies atomic motion rather than suppressing it.

Our working hypothesis is that this critical magnetic field occurs approximately where the Zeeman shift becomes large enough that the absolute value of the effective laser detuning begins to increase rather than decrease as atoms move toward the laser source (see figure 3.1). Specifically, when the Zeeman shift βz becomes comparable to or larger than the laser detuning δ_0, an atom moving toward one beam may experience an effective detuning that moves further from resonance rather than closer to it. For faster atoms, this results in weaker deceleration, and when $\alpha_d < 0$, net acceleration occurs, producing a heating effect.

$$m\frac{d^2\vec{r}}{dt^2} = \vec{F}_{\text{gravity}} + \vec{F}_{\text{scatt}} + \vec{F}_{\text{dipole}} \tag{3.11}$$

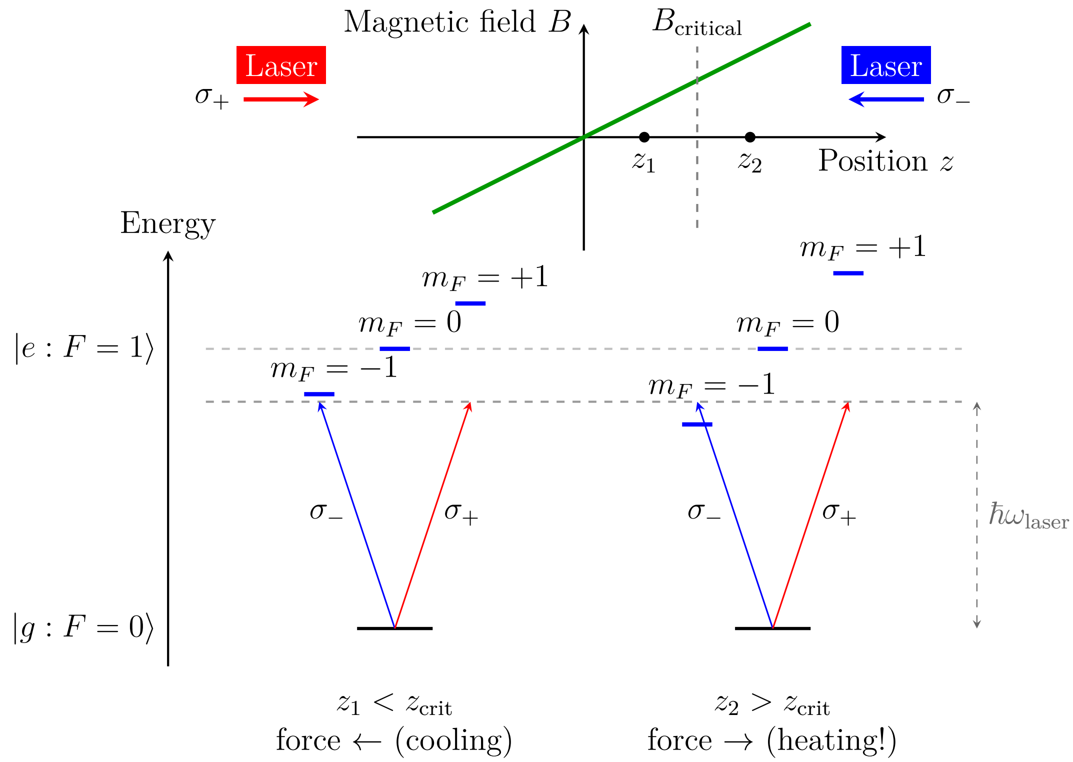

Figure 3.1: Illustration of the heating mechanism in a MOT for atoms held statically off-center. At position z_1, the σ_- laser is nearly resonant but remains red-detuned from the $m_F = -1$ excited state, providing a restoring force toward the center. Beyond the critical magnetic field (B_{critical}) at position z_2, the Zeeman shift is so large that the atomic transition becomes blue-detuned relative to the laser frequency. In this regime, the scattering rate decreases as the atom moves further from the center, resulting in a net force in the outward direction (heating) and subsequent atom loss.

This hypothesis can be investigated analytically using the two-level model. Figure 3.2 shows the damping coefficient as a function of magnetic field for typical MOT parameters ($s_0 = 1.30$, $\delta_0 = -15.7\,\text{MHz} = -3.02\,\Gamma$).

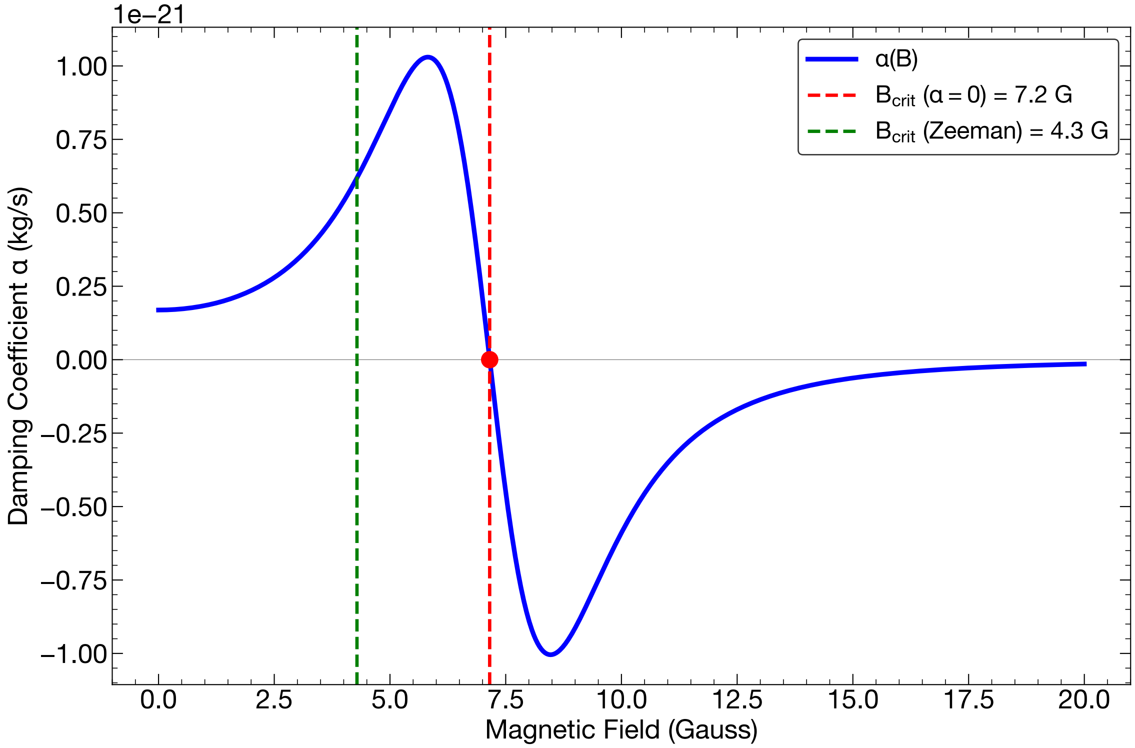

Figure 3.2: Damping coefficient α_d vs. magnetic field for the two-level approximation of the $F = 4 \rightarrow F' = 5$ transition (stretch states $m_g = -4 \rightarrow m_e = -5$). Used parameters: $s_0 = 1.3$, $\delta_0 = -15.7\,\text{MHz}$, $\lambda = 852.3\,\text{nm}$. The red dashed line indicates B_{crit} where $\alpha_d(B) = 0$, marking the transition from cooling ($\alpha_d > 0$) to heating ($\alpha_d < 0$). The green dashed line shows $B_{\text{crit,Zeeman}}$ where the Zeeman shift equals the laser detuning magnitude.

The blue curve shows $\alpha_d(B)$ calculated from equation (3.10). At $B_{\text{crit}} \approx 11.2\,\text{G}$ (red dashed line), the damping coefficient crosses zero, indicating the transition from cooling to heating.

To check if our hypothesis is plausible, we compare this to the field strength at which the Zeeman shift equals the laser detuning magnitude. For our ^{133}Cs MOT operating on the $F = 4 \rightarrow F' = 5$ transition with stretch states ($m_g = -4$, $m_e = -5$), the

Zeeman shift at magnetic field B is:

$$\Delta E_{\text{Zeeman}} = \mu_B(g_e m_e - g_g m_g)B = \mu_B\left(\frac{5}{6}\cdot(-5) - \frac{1}{4}\cdot(-4)\right)B = -\frac{5\mu_B}{3}B\,, \quad (3.12)$$

where we used $g_F = g_J\frac{F(F+1)+J(J+1)-I(I+1)}{2F(F+1)}$ with the appropriate values for the $6^2S_{1/2}$ ground state ($g_g = 1/4$) and $6^2P_{3/2}$ excited state ($g_e = 5/6$) of ^{133}Cs.

Setting the Zeeman shift equal in magnitude to the laser detuning, $|\Delta E_{\text{Zeeman}}| = \hbar|\delta_0|$, gives:

$$B_{\text{crit,Zeeman}} = \frac{3\hbar|\delta_0|}{5\mu_B} \approx 6.5\,\text{G}\,. \quad (3.13)$$

This is indicated by the green dashed line in Figure 3.2. As one can see, this magnetic field value is on the same order of magnitude as the analytical B_{critical} derived from $\alpha_d(B)$, although they are not strictly identical.

Multilevel simulation of the critical B-field

Because determining the precise critical magnetic field is highly relevant for our transport scheme, we performed a comprehensive numerical simulation of the full 20-level system ($F = 4$ with 9 ground Zeeman sublevels and $F' = 5$ with 11 excited Zeeman sublevels). This simulation utilizes the optical Bloch equation rate model introduced in section 2.2.2. The key finding is that the heating mechanism observed in the two-level toy model indeed persists in the full multilevel system. Furthermore, the critical magnetic field occurs at similar values, particularly for the moderate laser detunings ($|\delta_0| \sim 2$–$4\,\Gamma$) and saturation parameters ($s_0 \lesssim 2$) typical of our MOT operation.

Solving the rate equations detailed in section 2.2.2 yields the scattering rates and the steady-state populations of the atoms in the MOT. A remaining question, however, in order to do this multilevel simulation is: what is the total force exerted on an atom if it is held stationary at a specific off-center position z_0? To simulate the transport dynamics, we again treat the atomic motion semi-classically. The atom is considered a point particle moving under the influence of forces calculated from its internal state populations. The equation of motion is:

Here, $\vec{F}_{\text{gravity}} = mg\hat{e}_y$. We distinguish between the dissipative scattering force (radiation pressure) and the conservative dipole force (lattice potential).

We treat the far-detuned optical lattice as a conservative scalar potential $U_{\text{lat}}(\vec{r})$ acting on the atomic center of mass. Neglecting the state-dependent tensor light shifts (which are small for far-detuned lattices) [43], the lattice force is simply the gradient of this potential:

$$\vec{F}_{\text{lattice}} = -\nabla U_{\text{lat}}(\vec{r}) \tag{3.14}$$

The scattering force arises from the momentum transfer during photon absorption and stimulated emission. Spontaneous emission is isotropic on average and does not contribute to the mean force (though it contributes to heating/diffusion).

Using the rates $R^l_{n,j}$ defined in subsection 2.2.2, the force from a specific laser beam l is the photon momentum $\hbar\vec{k}_l$ multiplied by the net rate of photon redistribution into that mode. The total scattering force is the sum over all beams:

$$\vec{F}_{\text{scatt}} = \sum_l \hbar\vec{k}_l \underbrace{\left(\sum_{n,j} R^l_{n,j}(P_{g_j} - P_{e_n})\right)}_{\text{Net Interaction Rate}} \tag{3.15}$$

Here, the term $\sum R^l_{n,j} P_{g_j}$ represents the total absorption rate from beam l (momentum gain $+\hbar\vec{k}_l$), and $\sum R^l_{n,j} P_{e_n}$ represents the total stimulated emission rate into beam l (momentum loss $-\hbar\vec{k}_l$).

To calculate the populations P_i required for equation (3.15) and (3.14), we rely on the separation of timescales. The internal atomic state evolves at the rate of the natural linewidth, $\Gamma \approx 2\pi \times 5.2\,\text{MHz}$ for cesium (D_2). In contrast, the external motion in the lattice is characterized by the trap frequencies: the longitudinal frequency is typically $\omega_{z,\text{trap}} \approx 2\pi \times 100\,\text{kHz}$, while the radial frequency is $\omega_{r,\text{trap}} \approx 2\pi \times 100\,\text{Hz}$. Since $\Gamma \gg \omega_{z,\text{trap}}$ and $\Gamma \gg \omega_{r,\text{trap}}$, the internal populations reach equilibrium virtually instantaneously compared to the atomic motion. Therefore, at every time step of the simulation, we assume $\dot{P} = 0$. We construct the rate matrix M from the transition rates derived in equation (2.2.2) and solve the linear system:

$$M(\vec{r}, \vec{v}) \cdot \vec{P} = 0 \tag{3.16}$$

subject to the normalization condition $\sum P_i = 1$. This allows us to calculate the instantaneous force $\vec{F}(\vec{r}, \vec{v})$ and integrate the trajectory using a standard ODE solver (e.g. Runge-Kutta). The steady-state equation $M\vec{P} = 0$ involves a 20×20 rate matrix composed of four blocks. These dimensions arise from the 9 Zeeman sublevels of the $6\mathrm{S}_{1/2}, F = 4$ ground state and the 11 Zeeman sublevels of the $6\mathrm{P}_{3/2}, F' = 5$ excited state:

$$M = \begin{pmatrix} M_{gg} & M_{ge} \\ M_{eg} & M_{ee} \end{pmatrix} \tag{3.17}$$

where:

- M_{gg} (9×9) describes the depletion of the ground states.
- M_{ge} (9×11) describes the population transfer from the excited states to the ground states (spontaneous emission and stimulated emission).
- M_{eg} (11×9) describes the population transfer from the ground states to the excited states (stimulated absorption).
- M_{ee} (11×11) describes the depletion of the excited states.

Using the rate equations derived in section 2.2.2, the specific elements are populated as follows. Note that the sum over l accounts for all laser beams, where selection rules are enforced (i.e., $R^l_{n,j} = 0$ if the transition is forbidden by the polarization of beam l).

- Diagonal Elements (Depletion): The diagonal elements represent the total rate of leaving a state. For a ground state j (indices $1 - 5$):

  $$M_{jj} = -\sum_{n,l} R^l_{n,j} \tag{3.18}$$

  For an excited state n (indices $6 - 12$):

  $$M_{nn} = -\Gamma - \sum_{j,l} R^l_{n,j} \tag{3.19}$$

- Off-Diagonal Elements (Transfer): These elements represent the rate of arrival into row r from column c.

  For transfer from Ground j to Excited n (Block M_{eg}):

  $$M_{n,j} = \sum_{l} R^l_{n,j} \quad \text{(Stimulated Absorption)} \tag{3.20}$$

  For transfer from Excited n to Ground j (Block M_{ge}):

  $$M_{j,n} = \Gamma\beta_{nj} + \sum_{l} R^l_{n,j} \quad \text{(Spontaneous + Stimulated Emission)} \tag{3.21}$$

Here, β_{nj} is the branching ratio determined by the squared Clebsch-Gordan coefficients. For this specific closed transition, $\sum_j \beta_{nj} = 1$.

The system $M\vec{P} = 0$ is singular because the equations are linearly dependent (conservation of population). To find the unique physical solution, we must enforce the normalization constraint:

$$\sum_{k=1}^{12} P_k = 1 \tag{3.22}$$

In the numerical implementation, this is achieved by replacing the last row of M (row 12) with a row of ones, and setting the corresponding element in the solution vector to 1:

$$\tilde{M}\vec{P} = \vec{b} \implies \begin{pmatrix} & M_{1..11,:} & & \\ & & & \\ 1 & 1 & \dots & 1 \end{pmatrix} \begin{pmatrix} P_1 \\ \vdots \\ P_{12} \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ 1 \end{pmatrix} \tag{3.23}$$

This linear system $\tilde{M}\vec{P} = \vec{b}$ can be solved using standard linear algebra routines.

Simulation results

Figure 3.3 compares the twenty-level rate equation simulation with the two-level analytical prediction. Both models predict nearly identical critical magnetic fields where $\alpha = 0$, but the twenty-level simulation yields systematically lower damping amplitudes. This discrepancy increases with saturation because optical pumping redistributes population across multiple Zeeman sublevels, reducing the fraction of atoms in the cycling transition that produces the damping force. Figure 3.4 illustrates this effect: at low saturation, most population remains in the cycling states, but at high saturation, power broadening and off-resonant scattering spread population to other m_F states. Since damping is proportional to the cycling state population, this redistribution directly reduces α. The two-level model cannot capture this effect and therefore overestimates damping.

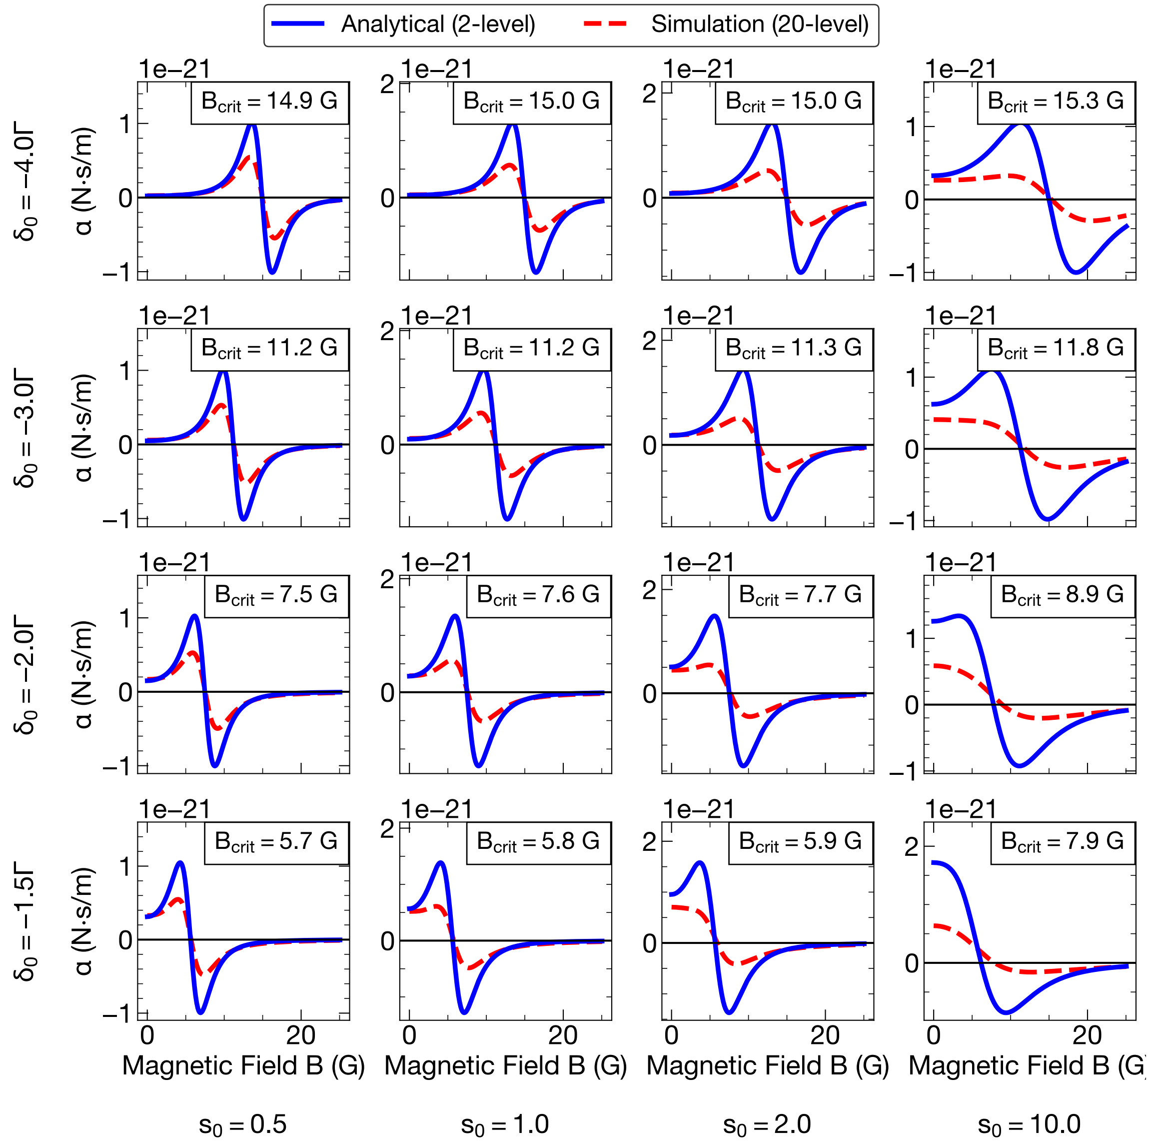

Figure 3.3: Damping coefficient $\alpha(B)$ as a function of magnetic field for different detunings δ_0 (rows) and saturation parameters s_0 (columns). The blue solid line shows the two-level analytical prediction, while the red dashed line shows the twenty-level rate equation simulation. The indicated B_{crit} is the critical B-field given by the 20-level simulation. Both models predict similar critical fields, but the twenty-level simulation yields lower damping amplitudes, with the discrepancy increasing at higher saturation.

This result has critical implications for the ratchet-type transport scheme. In this protocol, the atomic cloud must be held at several locations in the MOT, for example 1 cm off the MOT center while a new ensemble is loaded. Depending on the applied magnetic field gradient, the local magnetic field strength at this holding position can realistically reach $B \approx 10\,\mathrm{G}$ to $15\,\mathrm{G}$, thereby exceeding B_{critical}. Consequently, atoms in this region would experience heating rather than cooling from the MOT beams, leading to a substantial temperature increase and potential ejection from the optical lattice. Ultimately, this heating mechanism would severely limit overall transport efficiency and compromise the stringent temperature requirements necessary for subsequent tweezer loading in the science chamber. To mitigate this effect, we propose implementing a “dark tube” MOT (very similar to a dark SPOT MOT) along the lattice axis by spatially excluding the repump beam from the region where atoms are transported out of the MOT, as detailed in section 3.3 and illustrated in figure 3.7.

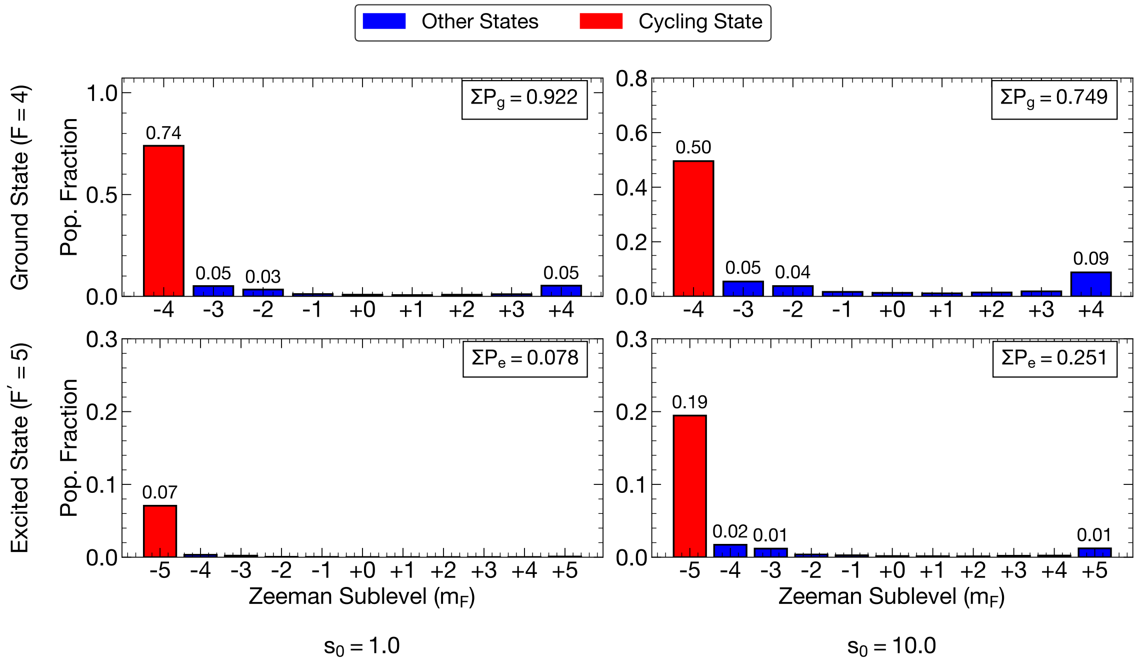

Figure 3.4: Steady-state population distribution across Zeeman sublevels for $s_0 = 1.0$ (left) and $s_0 = 10$ (right) at position $z = 0.5\,\mathrm{cm}$ with zero velocity. Simulation parameters: position $z = 0.5\,\mathrm{cm}$, velocity $v_z = 0.0\,\mathrm{m/s}$, B-field $B = 7.5\,\mathrm{G}$, detuning $\delta_0 = -20.9\,\mathrm{MHz}$. Red bars highlight the cycling transition states ($m_F = -4$ in ground state, $m_F' = -5$ in excited state) that the two-level model assumes contain all the population. At low saturation ($s_0 = 1.0$), most population concentrates in the cycling ground state ($m_F = -4$), validating the two-level approximation. At high saturation ($s_0 = 10$), significant population spreads across other Zeeman sublevels through optical pumping, reducing the fraction in the cycling transition.

3.1.2 Spin-Changing Collisions in ^{133}Cs

A second critical loss mechanism during transport arises from inelastic spin-changing collisions among ^{133}Cs atoms. When atoms are trapped in any Zeeman sublevel other than the absolute ground state ($|F = 3, m_F = +3\rangle$), they become susceptible to spin-relaxation. During such an event, the internal Zeeman energy released as the atoms decay to lower-energy m_F states is converted into translational energy. To conserve total energy, this is imparted to the colliding pair as a momentum kick, significantly heating the atoms (as illustrated in figure 3.5). For a transition of $\Delta m_F = 1$, the energy released per collision is $g_F \mu_B B$, where g_F is the Landé g-factor, μ_B is the Bohr magneton, and B is the applied magnetic field. For ^{133}Cs operating in a moderate background magnetic field of 10 G, this internal energy difference corresponds to a thermal energy equivalent of roughly:

$$\frac{g_F \mu_B B}{k_B} \approx 168\,\mu\text{K} \tag{3.24}$$

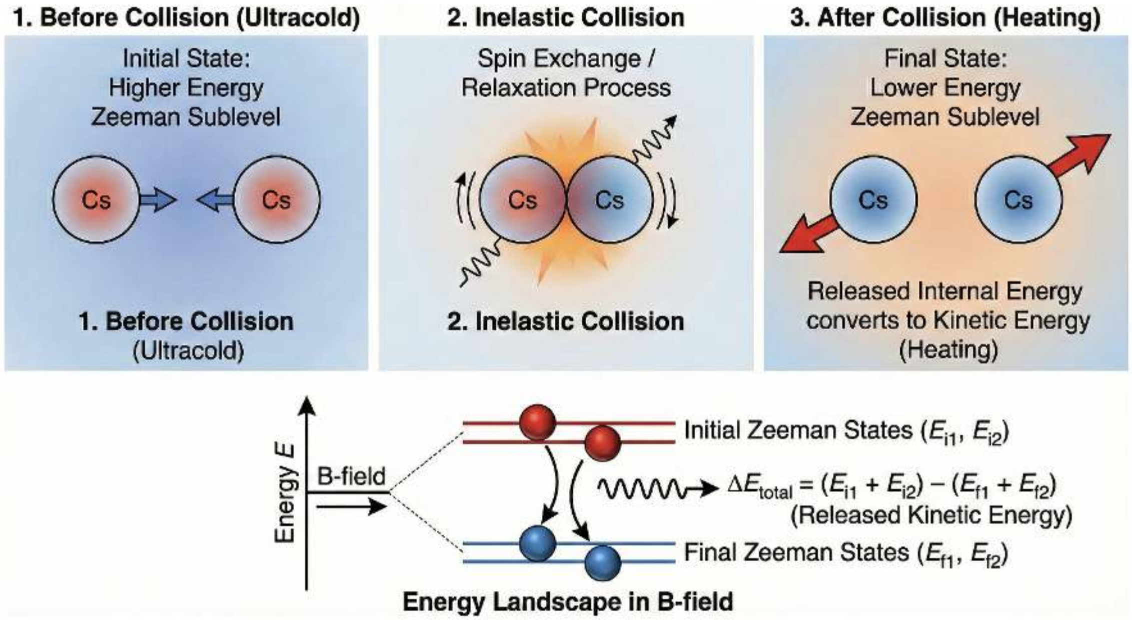

Figure 3.5: Mechanism of an inelastic spin-changing collision in ^{133}Cs. Two ultracold atoms in a higher-energy Zeeman sublevel collide and decay to a lower-energy state. The difference in internal Zeeman energy is released as kinetic energy, heating the atoms and typically ejecting them from the trap. Figure generated using Gemini.

Our optical lattice trap depths (U_{trap}) used during transport are on the order of 400 µK to 1 mK. While this energy is strictly lower than the lattice trap depth ($E_{\text{kick}} < U_{\text{trap}}$) and will not directly eject the atoms, it represents a massive injection of heat relative to the initial microkelvin atomic temperatures. Consequently, inelastic spin-changing collision leads to heating which increases the probability of subsequent trap escape.

This heating mechanism is particularly problematic for ^{133}Cs compared to other alkali metals like ^{87}Rb. The complex scattering properties of ^{133}Cs are characterized by a dense spectrum of broad Feshbach resonances that extend into the low magnetic field region [44, 45]. These Feshbach resonances result in large inelastic scattering cross-sections. If the atoms are not prepared in the absolute ground state, these resonances can drive above mentioned spin-changing collisions. Standard ^{133}Cs experiments mitigate this by optically pumping the ensemble into the $6S_{1/2}\ |F = 3, m_f = 3\rangle$ state before transport or evaporative cooling [46–48]. In this absolute ground state, inelastic spin-relaxation is energetically forbidden, as no lower-energy states exist for the atoms to decay into, thereby suppressing this heating channel.

3.2 Cesium MOT in literature

Having identified the transport-specific challenges in section 3.1, we now examine how MOT performance can be optimized to achieve high atom number, high density, and low temperature. Extensive theoretical efforts have been dedicated to understanding MOT dynamics and predicting fundamental parameters such as the resulting temperature, spatial shape, and density distribution [49]. However, the underlying physics is highly complex, and no universally accurate analytical model exists for many aspects of the trap. An example of this difficulty is the attempt to describe the atomic density in the multiple-scattering regime. Early work by [50] proposed an analytical model for this density limit, however, [51] subsequently showed experimentally that this model was incomplete. Instead of a universal analytical framework, they proposed a largely empirical model whose parameters strongly depend on the specific experimental apparatus. Despite the lack of a comprehensive and universal theoretical model, it remains entirely possible to derive robust empirical strategies and identify specific operating regimes to optimize desired trap parameters. In the following sections, we present the standard techniques and operation sequences developed by the cold atom community, providing the necessary context for our specific design choices. Most MOT-based experiments follow a similar operational sequence, regardless of whether the goal is BEC production, lattice loading, or tweezer array assembly. The sequence typically consists of distinct phases, each optimized for different physical objectives that require mutually incompatible parameter regimes.

3.2.1 Standard MOT Operation Phases

Loading Phase

The MOT initially operates in “capture mode” with parameters optimized for maximum atom number rather than density or temperature. As explained by [52], the total atom number is predominantly determined by the capture process, which means the slowing of individual atoms by Doppler cooling forces in the intersection region of the laser beams. Ultimately, the steady-state atom number is dictated by the balance between the trap’s loading rate and its loss mechanisms, a standard theoretical model that was experimentally confirmed many times, e.g. by [53]. Optimizing this capture process typically requires relatively small laser detuning ($\delta \approx -2\Gamma$ to -3Γ), moderate to high intensity ($s_0 \approx 1-3$), and modest magnetic field gradients ($5\,\mathrm{G/cm}$ to $15\,\mathrm{G/cm}$) [52, 54]. Loading times vary from $0.1\,\mathrm{s}$ to several seconds depending on background vapor pressure and beam geometry, with final atom numbers ranging from 10^6 to 10^{10} and temperatures of $100\,\mu\mathrm{K}$ to $300\,\mu\mathrm{K}$ [4, 52, 54, 55]. To push the atom number even higher, [56] argue that because the capture velocity depends on the beam diameter, using larger beams strongly increases the loading rate and final atom number. By applying this principle with large beam diameters ($d \geq 4\,\mathrm{cm}$), a detuning of $\delta = -20\,\mathrm{MHz}$ ($\approx -3.8\Gamma$ for Cs), peak intensities of $I = 25\,\mathrm{mW/cm^2}$, and a gradient of $7\,\mathrm{G/cm}$, [56] achieved 3.6×10^{10} Cs atoms in $0.2\,\mathrm{s}$. This scaling behavior was also explicitly demonstrated by [57], who showed that larger MOT beam diameters directly yield significantly more trapped atoms.

In terms of temperature after this initial cooling phase, for ^{133}Cs cooling on the D_2 line, the theoretical Doppler temperature limit is given by $T_D = \hbar\Gamma/(2k_B) \approx 125\,\mu\mathrm{K}$. The fundamental recoil limit is $T_R = (\hbar k)^2/(mk_B) \approx 198\,\mathrm{nK}$. MOTs typically achieve temperatures around the Doppler temperature limit. Sub-Doppler cooling within a MOT is possible by increasing the detuning and carefully reducing the magnetic field gradient, but only to a certain point before the trap becomes unstable.

Compression Phase (Optional)

After loading, many experiments implement a compression stage to increase atomic density. [52] emphasize that density and temperature are determined by the polarization gradient (sub-Doppler) forces at the trap center, which require different conditions than the capture phase. The compressed MOT (CMOT) technique (e.g. explained by [55]) separates collection and compression by first loading at a small magnetic field gradient and detuning, and then ramping to large detuning (typically $\delta \approx -6\Gamma$ to -10Γ, e.g., -44 MHz $\approx -8.5\Gamma$ for Cs) while increasing the magnetic field gradient (from loading values of $\sim$10 G/cm up to 30–60 G/cm). The reason for switching to large detunings and field gradients is that large detuning reduces repelling reradiation forces much more strongly than it reduces the trapping force. The simultaneous increase in the magnetic field gradient increases the trap's spatial confinement, which pushes the atoms toward the trap center. [55] showed that this allows for densities up to 5×10^{11} cm^{-3} with Rb for fields below 60 G/cm. This compression phase typically lasts 10–50 ms [55, 58].

This is a good point to discuss typical ^{133}Cs MOT densities, shapes, and sizes. As mentioned in the introduction of this section, accurately modeling these macroscopic cloud parameters is theoretically complex. According to [51], a MOT operates across several distinct density regimes depending on the atom number and trap parameters. At low atom numbers, the trap is temperature-limited and behaves essentially like a harmonic oscillator, yielding a Gaussian spatial profile where density scales linearly with the number of trapped atoms. However, as atom numbers increase, the trap enters the multiple-scattering regime. Here, the reabsorption of scattered cooling photons creates an outward radiation pressure that caps the maximum density, forcing the cloud to expand into a flat-topped, uniform spatial distribution. Most experiments aiming for high atom numbers operate within this multiple-scattering regime. For example, [56] observed a spherical, nearly uniform ^{133}Cs cloud in this regime with a full width at half maximum (FWHM) of 1.1 cm and a density of $3.6 \times 10^{10}\,\mathrm{cm}^{-3}$. In contrast, [4] report a much smaller, tighter trap when loading approximately 10^7 ^{133}Cs atoms from a 2D MOT in 100 ms. Using 2.5 cm diameter beams with a detuning of $\Delta = -3.17\Gamma$, a total intensity of $10 I_0$, and a steeper magnetic field gradient of 20 G/cm, they produced a MOT cloud with a $1/e^2$ diameter of roughly 1.6 mm. This already shows that there is quite some range in the physical size and spatial shape of the MOT depending heavily on the chosen operating parameters.

Cooling Phase

Temperature reduction below the Doppler limit is achieved through various sub-Doppler cooling techniques, each with distinct advantages and experimental requirements. The choice of cooling method depends on the experimental application and available infrastructure. For BEC experiments, resolved Raman sideband cooling (RSC) is frequently employed due to its ability to reach ultra-low temperatures. This technique requires confining atoms in a 3D optical lattice with sufficient trap frequency that individual vibrational sidebands are spectrally resolved [59]. [58] demonstrated degenerate Raman sideband cooling (dRSC) for ^{133}Cs, achieving $< 1\,\mu\mathrm{K}$ after 6 ms of cooling in a resonant optical lattice. While RSC achieves temperatures approaching the recoil limit, it requires additional experimental complexity (lattice infrastructure and precise frequency control) and is typically implemented only when sub-µK temperatures are essential, such as for efficient evaporative cooling toward BEC.

For tweezer-based quantum computing experiments and many lattice loading applications, polarization gradient cooling (PGC) is the most commonly used technique. PGC exploits sub-Doppler Sisyphus cooling mechanisms in the polarization gradients created by counter-propagating beams. The magnetic field gradient must be turned off, the laser detuning is increased, and the intensity is reduced, allowing atoms to cool through preferential optical pumping in the spatially varying light field. For ^{133}Cs, [60] report temperatures of 10 µK, and [61] report temperatures of 2.5 µK.

Another relatively recent approach, proposed by [62] and demonstrated for continuous reloading by [20], is Λ-enhanced gray molasses (LGM). Because it has a lower velocity capture range than PGC, it typically requires some precooling of the atoms, as demonstrated in [20]. LGM can achieve lower temperatures than PGC by exploiting dark state shelving in a three-level Λ configuration to reduce photon scattering for already-cold atoms. [63] reached 1.7 µK with LGM and ^{133}Cs. To achieve the lowest possible temperatures, LGM requires phase coherence between the cooling and repumper beams. Hence, most groups modulate the repumper into the cooling beam using an electro-optic modulator (EOM). Like standard molasses, LGM requires operation at zero magnetic field (or a small, constant bias field, but strictly no gradients), making it incompatible with simultaneous MOT loading and thus unsuitable for ratchet-type continuous reloading schemes where the MOT magnetic field must remain on during atom extraction.

3.2.2 Advanced MOT Architectures

As established, a fundamental limitation to atomic density in the large atom number regime is the re-radiation and rescattering of photons between trapped atoms, which creates an effective outward radiation pressure [55].

One method to mitigate this rescattering is the dark SPOT (Spontaneous Force Optical Trap) MOT [64]. The principle relies on pumping atoms in the dense center of the trap into a hyperfine state that is "dark" (off-resonant) to the main MOT cooling light. A common implementation simply places a physical block in the center of the repump beam profile, creating a dark spot. Atoms that migrate to the trap center fall into the $F = 3$ ground state due to the missing repumper and cease scattering cooling photons, dramatically reducing radiation pressure while remaining magnetically trapped. Proposed by [64], [52] later implemented this for Cs and achieved a density of $10^{12}\,\mathrm{cm}^{-3}$.

A related approach is the core-shell MOT, a method very popular for alkaline earth atoms [65]. Here, cooling light operating on a broad transition is used in the outer spatial region for a large capture velocity, while cooling light on a much narrower transition is used in the inner region. The narrow transition naturally scatters far fewer photons, reducing radiation pressure in the dense core and leading to a much lower Doppler temperature limit. For example, [66] combined standard 852 nm cooling light with the exceedingly narrow $6S_{1/2} \to 5D_{5/2}$ electric quadrupole transition at 685 nm. Calculating the Doppler temperature of this transition with its effective linewidth of $\Gamma \approx 2\pi \times 170\,\mathrm{kHz}$ leads to a theoretical limit of $T_D \approx 4.1\,\mu\mathrm{K}$ for ^{133}Cs. [66] achieved 5.29 µK directly in the MOT, though they do not report the corresponding atom density or final number.

Table 3.1 provides an overview of these different methods for ^{133}Cs atoms, detailing the reported benchmarks across literature.

Reference	**Method / Technique**	**Max. Atoms (N)**	**Min. Temperature**	**Max. Density (cm^{-3})**	**Loading Time**
Gibble *et al.* [56]	Large beams for high atom number	3.6×10^{10}	–	3.6×10^{10}	0.2 s
Petrich *et al.* [55]	Compression phase	–	–	5.0×10^{11}	–
Klostermann *et al.* [58]	cMOT with molasses and dRSC	3.0×10^{7}	30 µK to 40 µK (after cMOT) 10 µK (after Molasses) < 1 µK (after dRSC)	5.0×10^{10}	3.0 s
Hsiao *et al.* [63]	Gray molasses cooling + LGM	3.2×10^{8}	120 µK (after MOT) ~ 8 µK (after GMC) 1.7(2) µK (after LGM)	–	~ 100 ms
Ketterle *et al.* [64]	dark SPOT MOT	–	105 µK	1.0×10^{12}	–
Saffman *et al.* [66]	Narrow line $6S_{1/2} \rightarrow 5D_{5/2}$ core-shell (685 nm) MOT	–	5.29 µK	–	–

Table 3.1: Overview of advanced trapping and cooling techniques for ^{133}Cs highlighting peak performance metrics reported in literature.

3.3 Derived MOT Design Choices

Based on the preceding information, we can now evaluate the requirements for the proposed MOT and the physical reasoning behind them. The conclusions drawn in this section are summarized in table 4.1 .

3.3.1 MOT Cooling Beams

Starting with the MOT cooling beams, [56] states that achieving the highest possible atom number requires maximizing the capture velocity. As described in their work, the primary method to achieve this is by using very large MOT beams. In our setup, the side windows of the MOT chamber are mounted on standard 2.75 in ConFlat (CF) flanges, which provide a clear viewport diameter of approximately 1.4 in to 1.5 in. Typically, mirrors and apertures should have a diameter at least 1.5 times larger than the beam diameter ($1/e^2$) to avoid undesirable diffraction effects. An illustration of this can be found in the appendix in figure E.1. Considering that standard 2 in mirrors (such as the BB2-E03) only guarantee a clear aperture of $> 85\%$ of their diameter, the MOT beam diameter should be roughly 1.1 in. This still provides a substantial capture volume. Another interesting consideration is the ideal beam shape. In theory, the same MOT beam size can achieve a higher capture velocity if the beam is more flat-top-like (as opposed to Gaussian), because the atoms at the edges of the beam are cooled more due to increased scattering. However, creating a flat-top beam is difficult, as it requires phase or intensity modulation rather than simple refractive optics. These methods are often power-inefficient and necessitate a more complex optical setup. One idea we are considering is to first expand the Gaussian beam to about a 0.6 in diameter, cut it by a certain percentage, and then expand it again. This approach would theoretically yield a flatter Gaussian profile. However, if we do not want to image the cutting aperture onto the atoms (as this increases setup complexity), a question arises as to whether the diffraction pattern of an aperture is acceptable for the MOT.

Figure 3.6 shows a simulation of the diffraction pattern created by cutting a Gaussian beam with a waist radius of $w_r = 8.609\,\mathrm{mm}$ by 30%, 40%, 50%, and 60%, and then expanding that beam again with a 4f telescope of $f_1 = -50\,\mathrm{mm}$ and $f_2 = 20\,\mathrm{cm}$. The reason for choosing this expansion telescope is that if you cut the beam by 50%, the resulting beam radius is 1.3 in, which is roughly what our 2 in optics allow us to accommodate. The diffraction pattern was calculated using the Fresnel propagation method. Figure 3.6 shows that there is some intensity ringing at the edges of the diffraction pattern, and for the 60% case, one starts to see a sharp middle peak. Both the ringing at the edges and especially the middle peak could potentially be non-ideal for a MOT, but it is hard to quantify. For that reason, we decided that we will initially try to implement this cutting technique for our MOT using a simple size-changeable aperture. Should the diffraction-induced intensity structure prove detrimental to MOT performance, the aperture can be removed.

For all subsequent purposes, we assume that the MOT beam is first expanded to $w_r = 8.609$ mm, cut by 50%, and then expanded again to 1.3 in. As can also be seen in figure 3.6, cutting by 50% means we lose about 60% of the power. With that, we can estimate how much power for the MOT cooling laser is needed. To start, I assume that the maximum saturation parameter we could possibly need is $s_0 = I/I_s = 10$, so $I = 10 \times 1.1\ \mathrm{mW/cm^2} = 11\ \mathrm{mW/cm^2}$. That is already far too much to be optimal. Usually MOTs are operated at $s_0 \approx 1$ as too high powers broaden the atomic transition and increase re-radiation pressure, which limits the minimum temperature and maximum density of the cloud. However [56] also used significantly more power than usual for their high-atom-number MOT. In order for the peak of the 50% cutting case in figure 3.6 to be at $s_0 = I/I_s = 10$, the total power required in that Gaussian beam at 1 m is about 75 mW. Therefore, we want about 75 mW of power for each of the six beams.

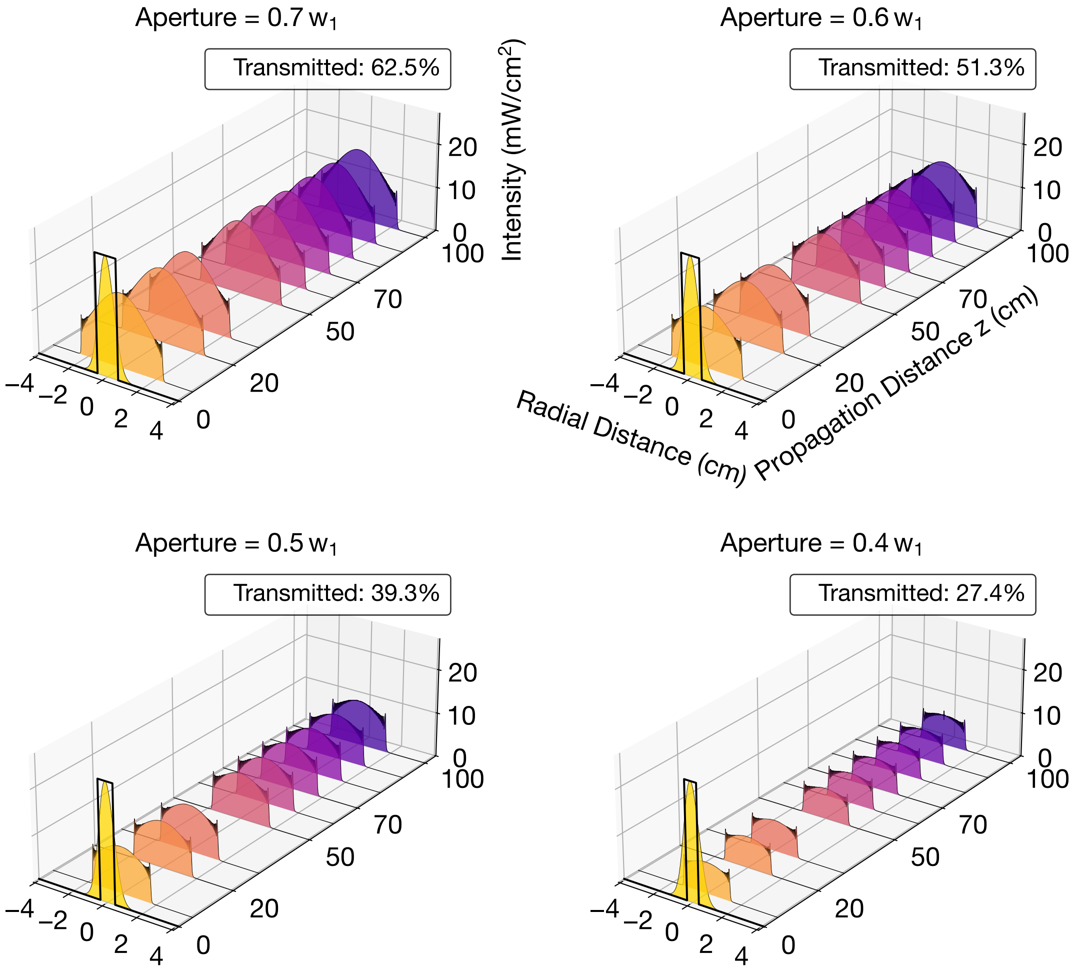

Figure 3.6: Simulation of the diffraction patterns created by aperture clipping at various radii. The first yellow Gaussian beam is the incoming $w_r = 8.609\,\text{mm}$ beam, shown just for reference. The two black vertical lines indicate the aperture size relative to the incoming Gaussian. The simulation then models the beam passing through a 4f telescope created by two lenses with $f_1 = -50\,\text{mm}$ and $f_2 = 20\,\text{cm}$ (these lenses are not pictured). The diffraction patterns shown along the propagation distance z represent the beam after the second lens (f_2) of the 4f telescope. In reality, there is a significant physical distance between the initial reference Gaussian beam and the subsequent diffraction patterns.

The last open question is the required detuning range for the MOT beams and the requirements for the LGM cooling mentioned in section 3.2. Summarizing all the different regimes discussed in that section, the detuning range should be approximately from -2Γ to -10Γ. As described in the same section, we would like the option to perform LGM cooling with the MOT beams. LGM requires two beams: one cooling beam driving the $|6S_{1/2}, F = 4\rangle \rightarrow |6P_{3/2}, F = 4\rangle$ transition, and one repumper beam driving the $|6S_{1/2}, F = 3\rangle \rightarrow |6P_{3/2}, F = 4\rangle$ transition [63]. The two beams must be coherent with each other. A reliable way to achieve this is by using an EOM to modulate the repumper onto the cooling beam (the repumper intensity should be about 10% of the cooling intensity). This means that we have the additional requirement for our cooling beams to be able to switch the cooling beam resonance from $|F = 4\rangle \rightarrow |F' = 5\rangle$ to $|F = 4\rangle \rightarrow |F' = 4\rangle$ and have an EOM in the beam path.

3.3.2 Repumper

The repumper requirements depend on the transport timing scheme we intend to use (see section 2.1.2). If we want to do pulsed transport with MOT shutdown, no special repumper is required, as we can perform the m_F-ground state preparation (required to avoid spin-changing collisions during ^{133}Cs lattice transport) after shutting down the MOT and before transporting the atoms into the lattice. In this case, we could simply use one global repumper covering the entire MOT. If, on the other hand, we want to perform ratchet-type transport, two issues arise: first, we statically hold the atoms off-center in the MOT, which introduces the critical B-field heating issue described in section 3.1.1. Second, we must perform the m_F-ground state preparation while the MOT is actively running.

One way to fulfill these two requirements without disrupting the MOT is illustrated in figure 3.7. If we spatially shape the repumper so that it is absent along the lattice direction (where the atoms are extracted), we avoid the critical B-field heating problem because the atoms become dark to the cooling light in the absence of the repumper. Furthermore, we can perform state preparation within this dark tube, as the state preparation and depumper beams will only efficiently manipulate the atomic states in the region devoid of repumper light.

The idea, then, is to apply this “pac-man” shape from two orthogonal directions in the MOT, leaving a continuous tube along the lattice completely free of repumper light. Another interesting aspect is that this setup effectively creates a dark SPOT MOT, as described in section 3.2, potentially leading to a higher-density and lower temperature

(a) MOT beam geometry cross-section (b) ^{133}Cs D_2 level scheme with transitions

Figure 3.7: Overview of the involved beams and MOT geometry. **(a)** Schematic cross-section of the beams in the MOT chamber (not to scale). The DMD-shaped repumper (green) is formed into a "pac-man" profile, leaving a dark tube along the lattice axis (yellow, 400 µm width). The MOT cooling beams span an area of roughly 1.2 in. The σ^+ state preparation and π depumper beams overlap with the lattice axis at an angle of approximately 3°, operating effectively only within the dark tube where the repumper light is absent. **(b)** Relevant cesium D_2 energy level scheme ($6S_{1/2} \rightarrow 6P_{3/2}$) indicating the specific transitions used for the experiment. The MOT cooling beam drives the $F = 4 \rightarrow F' = 5$ transition, while the repumper drives $F = 3 \rightarrow F' = 4$. State preparation and depumping are performed on the $F = 4 \rightarrow F' = 4$ transition.

in the MOT. If we can control the repumper profile, we could try to optimize its shape to increase atom flux into the tweezer region. A good way to accomplish this is by using a digital micromirror device (DMD), which allows individual micromirrors in the image plane to be toggled on and off. Generally, the repumper should ideally cover the entire MOT region (the full 1.3 in) and be formed into a pac-man shape using a DMD. The pac-man inset should extend to half the MOT region's length (about 0.75 in), and the radius of the inner cylinder should approximately match the lattice radius (100 µm to 200 µm), but the exact size will be subject to optimization later in the experiment. The edges should be as sharp as possible, ideally on the order of tens of micrometers. Related to edge sharpness is the question of how much leaked repumper light is permissible within the dark tube, and how much repumper intensity is required in the surrounding regions. This is estimated by taking the atomic Hamiltonian and decay parameters described in section 2.2.2, adding another driving term for the repumper, and then directly solving for the steady-state of the four-level density matrix. This allows us to evaluate the fractional population in the

$|6S_{1/2}, F = 3\rangle$ dark state as a function of the local repump intensity. The results are shown in figure 3.8. It can be seen that at an intensity of around $5\,\mu\mathrm{W/cm}^2$, the $F = 3$ population is already depleted by over 90%. Consequently, the DMD profile must have a quite high extinction ratio within the dark tube to ensure the atoms remain completely unaddressed by the repumper for successful state preparation, which is possible and evaluated in more detail later. Conversely, the bulk MOT region requires only a modest intensity of at least $50\,\mu\mathrm{W/cm}^2$ to effectively clear out the dark state and keep the atoms actively cycling. Assuming an intensity of $100\,\mu\mathrm{W/cm}^2$ and a repumper beam diameter of 1.3 in, this requires approximately 0.86 mW per repumper arm.

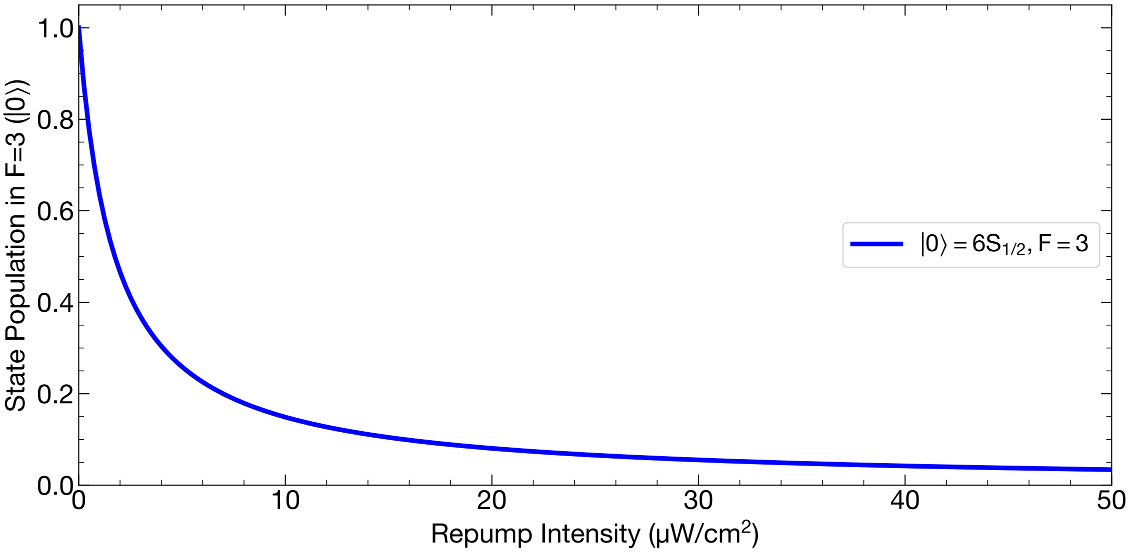

Figure 3.8: Fractional population in the ^{133}Cs $|6S_{1/2}, F = 3\rangle$ ground state as a function of local repump intensity, obtained from a full four-level density matrix steady-state solution. The rapid decline indicates that even small amounts of leaked repumper light ($\sim 5\,\mu\mathrm{W/cm}^2$) significantly deplete the dark state population, highlighting the need for a high-extinction DMD profile to maintain the dark tube for state preparation.

3.3.3 Depumper and State Preparation

As these two beams are physically integrated into the lattice optical setup, a full discussion of their simulated dynamics is omitted here in favor of summarizing the key results. Their primary function is to optically pump the atomic ensemble into the $|6^2S_{1/2}, F = 3, m_F = +3\rangle$ absolute ground state to avoid inelastic spin-exchange collisions during transport. To achieve this, we employ a standard two-beam optical pumping scheme widely used in ^{133}Cs BEC experiments [47]. The state-preparation beam is resonant with the $F = 3 \rightarrow F' = 3$ transition at 852 nm. By utilizing σ^+ polarization, this beam drives transitions that increase the magnetic quantum number ($m_F \rightarrow m_F + 1$), progressively pumping atoms toward the stretched state $|F = 3, m_F = +3\rangle$, which acts as a dark state for further σ^+ absorption. However, the state-preparation beam alone is insufficient because the $F' = 3$ excited state can decay to both the $F = 3$ and $F = 4$ ground state manifolds. To recover atoms that decay into $F = 4$, a depumper beam resonant with the $F = 4 \rightarrow F' = 4$ transition is applied. This drives the atoms back into the $F = 3$ manifold, where the state-preparation beam can continue accumulating population toward $m_F = +3$. This closed pumping cycle efficiently transfers the atomic population into the absolute ground state. Implementing this scheme requires specific geometric considerations. While both beams should propagate along the local magnetic field (which in our setup happens to be collinear with the lattice) we do not want them present along the entire transport direction. Furthermore, they must not be perfectly parallel to the local magnetic field, as this can lead to the emergence of unwanted dark states. Consequently, they are introduced at a slight angle of approximately 3° relative to the lattice axis. In this atomic frame geometry, the state-preparation beam is σ^+-polarized, whereas the depumper beam is π-polarized. To ensure complete state preparation, these beams are spatially matched to cover the entire radial extent of the lattice (with a radius of approximately 200 µm). The optical power requirements for these beams are comparable to those of the MOT repumper.

3.3.4 Absorption and Fluorescence Imaging

The final aspect to consider is how to image the MOT. We want to image it for two main reasons: first, to roughly estimate the total atom number and the loading rate and second, and more importantly, to monitor the loading from the MOT into the lattice to evaluate the transfer efficiency and identify any potential issues. There are two main imaging methods available: fluorescence imaging and absorption imaging. For fluorescence imaging, we can simply use the standard cooling beams. All we need are two cameras positioned along two orthogonal axes that are free of other beam paths. For absorption imaging, on the other hand, we require separate probe beams that pass through the atomic cloud and are subsequently imaged onto a camera. Directly comparing fluorescence and absorption imaging is nuanced, as their optimal use cases depend heavily on the optical depth of the atomic sample [67]. The optical depth τ is calculated as $\tau = n_0\sigma_0 L$, where $\sigma_0 \approx 1.1 \times 10^{-9}\,\mathrm{cm}^2$ for the ^{133}Cs D$_2$ line [30]. Assuming a ^{133}Cs MOT with 10^{12} atoms and a density of 10^{11} atoms/cm^3, the optical depth is $\tau \approx 10$ (which is generally too high for standard absorption imaging). Conversely, for the atoms loaded into the lattice (assuming roughly 10^6 to 10^7 atoms within a 500 µm waist diameter) the density is approximately 10^9 atoms/cm^3. This yields an optical depth of $\tau \approx 0.1$ to 1, which is ideal for absorption imaging. Because absorption imaging excels in the moderate optical depth regime ($0.1 < \tau < 5$), we intend to utilize both methods: fluorescence imaging for monitoring the bulk MOT, and absorption imaging to image the atoms in the lattice and align the lattice with the MOT position.

For the lattice alignment, it is particularly important to image along two orthogonal axes, one vertical (top/bottom) and one horizontal (side), allowing for 3D spatial tracking. Determining the optimal probe beam size involves a trade-off. In principle, a larger beam is preferable to cover the entire atomic cloud and to track the atoms as they are transported out of the MOT. However, especially for absorption imaging, a smaller beam generally yields better imaging quality due to higher spatial uniformity, which is critical for extracting accurate absorption profiles. As a compromise, we decided to aim for absorption imaging beams with a diameter of approximately 0.6 in. Their power requirement is quite low. To avoid non-linear absorption effects from saturating the atomic transition, imaging is typically performed well below the saturation intensity ($I \ll I_{\mathrm{sat}}$). For the ^{133}Cs D_2 transition ($I_{\mathrm{sat}} \approx 1.1\,\mathrm{mW/cm^2}$), probing at $0.1 I_{\mathrm{sat}}$ with a 0.6 in diameter beam requires only about 0.2 mW of optical power, though having up to 2 mW available is beneficial to allow for high-intensity imaging techniques if necessary.

In summary, the derived MOT design establishes several stringent requirements for the optical system. The six cooling beams must be expanded to a 1.3 in diameter (utilizing a 50% clipping technique for a flatter profile), requiring approximately 75 mW of power each, a detuning range of -2Γ to -10Γ, and EOM-driven LGM capabilities. The repumper beams will use a DMD to project a 1.3 in "pac-man" profile to create a high-extinction dark tube along the transport axis. Within this dark tube, 3° off-axis σ^+ state-preparation and π depumper beams matched to the 200 µm lattice radius will continuously drive the atoms into the absolute ground state. Finally, the system must accommodate dual-axis fluorescence imaging for the bulk MOT alongside 0.6 in, low-power ($\approx 0.2\,\mathrm{mW}$) absorption imaging beams to monitor lattice loading.

Chapter 4

Experimental Setup

In section 3.3 we already discussed a lot of the requirements we would like the MOT beams to fulfill. In order to achieve these beams at the atoms, there are three required stages:

- Setting up the required lasers and frequency stabilization
- “Distribution breadboards” preparing the correct wavelengths and enabling the independent detuning of these wavelengths and shutters to control the different beams
- Experimental table optics, bringing the beams in the correct size and correct polarization to the atoms in the MOT chamber and, if needed, meeting the required imaging condition

In the following we will go through each part. As the first two are heavily influenced by the experimental table optics and the requirements for the optics on that side, we will start with that part.

4.1 Experimental Table Optics

As a preliminary note, the optical layout presented in this section was originally developed for the first iteration of our vacuum system, which featured the black-coated tube shown in figure 2.3. The initial design included custom optical breadboards tailored specifically to that version of the chamber. Because the pivot to the new vacuum design occurred late in the timeline of this thesis, the SolidWorks models have not been fully adapted to reflect this update. However, the core optical design (including beam sizes, lens selections, and general specifications) is directly transferable to the new chamber. Only minor adjustments to the overall layout and port routing will be necessary.

4.1.1 General Layout and Design Choices

There are several fundamental design considerations when laying out the main optical table. A primary question is whether we load the 3D MOT directly from a dispenser or if we instead load from a 2D MOT. While direct dispenser loading provides a high initial atomic flux, the high thermal velocities of the emitted atoms severely limit the fraction that falls within the capture velocity of the 3D MOT. Furthermore, the resulting background gas load significantly degrades the vacuum lifetime. In contrast, a 2D MOT pre-cools the atoms, yielding a much slower atomic beam ($\approx 10\,\mathrm{m/s}$), which dramatically increases capture efficiency while preserving the ultra-high vacuum necessary to prevent collision-induced losses during subsequent lattice transport. The vacuum lifetime in the main MOT chamber is especially critical for the ratchet-type transport scheme, where the atomic ensembles will experience extended hold times in this region. For these reasons, we elect to implement a 2D MOT. Because a neighboring experiment in our laboratory has already demonstrated good reliability with the commercial ^{133}Cs 2D MOT cell from Infleqtion, we choose to purchase our atomic source from the same manufacturer.

Another critical design choice concerns the optical delivery of the MOT beams, specifically, whether to use independent counterpropagating beams or a retroreflected geometry. While retroreflection significantly reduces the required optical footprint and needed power, a notable advantage given the large diameter of our MOT beams and the spatial constraints of routing them to individual chamber ports, it presents a challenge for our spatially shaped repump light. Because the repumper propagates along two of the MOT axes and features a precisely shaped “pac-man” spatial profile, retroreflecting the cooling beams would necessitate retroreflecting the shaped repumper as well. To maintain the correct spatial profile and sharp edges at the atomic cloud upon reflection, the retroreflected path would require an additional $4f$ relay imaging system. Given the difficulty of overlapping the counterpropagating spatial dark tubes, we opt against retroreflection and instead design independent beam paths for all six MOT axes.

Finally, the integration of the absorption imaging system requires careful planning to minimize the number of dedicated vacuum viewports. Because the exact polarization state of the absorption imaging beam is an experimental degree of freedom, we can multiplex the imaging beam with the MOT cooling beam on the input side by utilizing orthogonal polarizations. This configuration allows us to easily separate the beams using a waveplate and polarizing beam splitter. However, a potential complication arises regarding how to cleanly isolate the weak transmitted imaging signal from the strong counterpropagating MOT cooling beam that shares the exact same optical axis. While this optical separation is physically realizable, it requires a rigorous accounting of the MOT cooling polarization definitions, specifically the distinction between the laboratory frame and the atomic frame. A detailed theoretical analysis demonstrating how this beam separation is effectively achieved is provided in Appendix B.

Determining the optimal port assignment for each beam relies on experimental considerations introduced in the previous chapter. Specifically, both absorption imaging and the shaped repumper beam must be applied along two axes: one transverse (side) axis and the vertical (top-to-bottom) axis. Both optical systems impose stringent imaging constraints: the structured repumper profile must be sharply imaged onto the atomic cloud, and the absorption image of the atomic shadow must be relay-imaged onto a camera sensor. By arranging the repumper and the imaging light such that they counterpropagate along identical axes, we effectively reduce the number of vacuum viewports that require imaging optics to just two.

The final major design decision concerns the spatial shaping of the repumper beam. While static shaping with a physical amplitude mask is the simplest approach, we favored dynamic shaping. Real-time control allows us to empirically optimize the MOT loading efficiency by systematically varying the length and diameter of the "pac-man" dark tube. To achieve this, we selected a DMD rather than a liquid-crystal spatial light modulator (SLM). A DMD offers several advantages for this application: it provides an exceptionally high extinction ratio in the "off" regions, is more cost-effective, and operates via direct amplitude modulation in the image plane. This provides highly intuitive, pixel-by-pixel control over the beam shape, bypassing the need to continuously compute complex Fourier-plane phase holograms (e.g. via the Gerchberg-Saxton algorithm) as would be required for an SLM. These considerations, together with all the requirements mentioned in the previous section 3 are summarized in table 4.1. The port enumeration can be seen in figure 4.1.

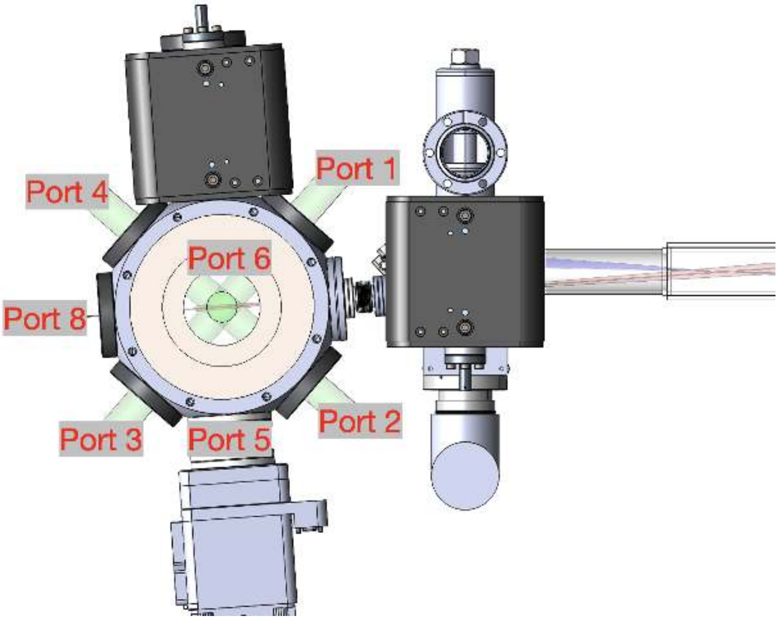

(a) Top view

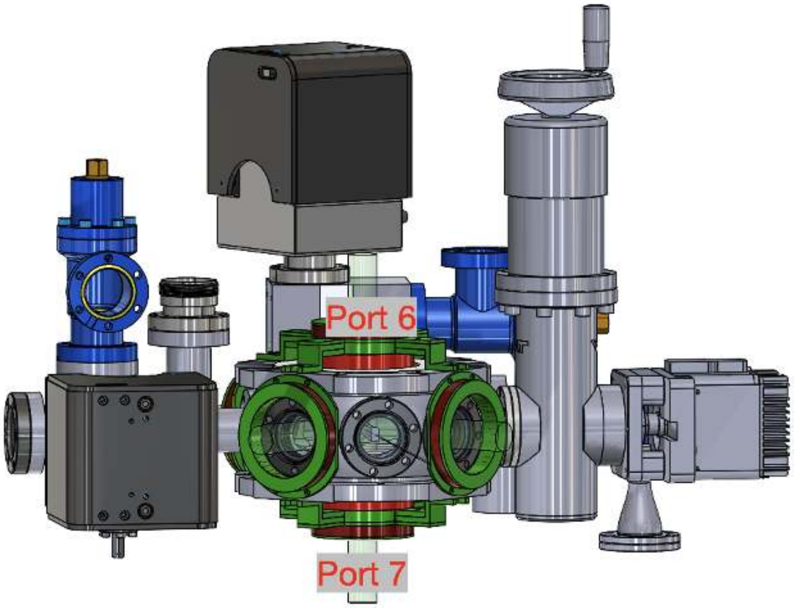

(b) Side view

Figure 4.1: Port enumeration showing (a) top view and (b) side view of the MOT chamber

Port	**Purpose**	**Wavelength**	**Approx. Detuning Range**	**Pol.**	**Power**	**Beam Size (diameter)**	**Notes**
p1, p2, p3, p4, p6, p7	MOT cooling	852 nm, $4 \to 5$	$-2\Gamma =$ 10.4 MHz to $-10\Gamma =$ 52 MHz	$\sigma+/\sigma-$	75 mW per beam	30.48 mm = 1.2 in	
	LGM cooling	852 nm, blue detuned $4 \to 4$	$+4\ \Gamma =$ 20.8 MHz to $+12\ \Gamma =$ 62.4 MHz	$\sigma+/\sigma-$ or linear			MOT cooling beam is switched to LGM cooling beam by switching the vescent laser offset-lock point from $4- > 5$ to $4- > 4$; LGM cooling does not work under magnetic field gradient, only under constant magnetic field, so is only compatible with atom array 2 type transport;
	LGM repump	852 nm, $3 \to 4$	-0.01Γ	does not matter	$\approx 10\,\%$ of LGM cooling		As the LGM repump should be coherent to the LGM cooling beam, this one is modulated onto LGM cooling beam using an EOM on the laser table
p2 (axis 2)	absorption imaging camera						
	MOT repumper	852 nm, $3 \to 4$	On resonance	does not matter	≈ 1 mW per arm	30.48 mm = 1.2 in	Shapable to “pac-man” shape with DMD
p4 (axis 2)	absorption imaging	852 nm, $4 \to 5$	$\pm 2\Gamma =$ ± 10.4 MHz	does not matter	$\frac{1}{10} I_{sat} \Rightarrow$ 0.2 mW	15.24 mm = 0.6 in	Because the polarization does not matter, we can give this beam the opposite circular polarization of the parallel cooling beam and then separate it on the other side
p5	fluoresence imaging camera						
	2D MOT beam with atoms						

Port	**Purpose**	**Wavelength**	**Approx. Detuning Range**	**Pol.**	**Power**	**Beam Size (diameter)**	**Notes**
p6 (top)	MOT repumper	852 nm, 3→4	On resonance	does not matter	≈ 1 mW per arm	30.48 mm = 1.2 in	Shapable to “pac-man” shape with DMD
	absorption imaging camera						
p7 (bottom)	absorption imaging	852 nm, 4→5	$\pm 2\Gamma =$ $\pm 10.4\,\mathrm{MHz}$	does not matter	$\frac{1}{10} I_{sat} \Rightarrow$ 0.2 mW	15.24 mm = 0.6 in	Because the polarization does not matter, we can give this beam the opposite circular polarization of the parallel cooling beam and then separate it on the other side
p8 (lattice axis)	fluoresence imaging camera						Use a dichroic to seperate lattice wavelength and 852 nm for fluoresence imaging
	lattice	1063 nm / 1030 nm		linear	60 W per beam	400 µm to 500 µm	1 mK with 200 µm radius or 444 µK with 300 µm
	state prep	852 nm, 3→3	On resonance	$\sigma+$	few mW	≈ 400 µm	Slightly angled (by ≈ 3° such that it exits the MOT chamber bucket window and it’s not propagating along the lattice
	depumper	852 nm, 4→4	On resonance	π/ linear	few mW	≈ 400 µm	Slightly angled (by ≈ 3° such that it exits the MOT chamber bucket window and it’s not propagating along the lattice

Table 4.1: Summary of beam requirements and port assignments for the main MOT and imaging systems.

The physical design of the main experimental table imposes several constraints on the optical layout. The table consists of three layers of breadboards, as illustrated in figure 4.2, with the glass cell positioned in the center of the breadboards. The entire left side of all three breadboard layers must remain unobstructed to accommodate tweezer paths, zone beams, and subsequent beamlines. Additionally, the entire right side of the middle layer is dedicated to the lattice optics. Consequently, the MOT optics are restricted to either the top-right or the bottom-right sections of the setup. Because the bottom section features a highly limited vertical clearance of only 13.3 cm, optical alignment in this region would be highly impractical. Therefore, the majority of the MOT optical assembly is designed to sit on the upper layer, with only the minimal essential routing components situated on the bottom layer (see figure 4.3). Based on these spatial considerations, figure 4.4 illustrates the final optical layout developed after multiple iterative design phases. Due to extended manufacturing lead times for the custom non-magnetic stainless steel honeycomb breadboards (Newport), combined with the late pivot in the vacuum system design, the physical assembly of this subsystem was not completed within the timeframe of this thesis. However, the comprehensive SolidWorks model ensures that the designed setup is directly implementable on the main table. The focal lengths of the lenses shown in figure 4.4 can be found in table 4.2. In the following, each beam path is explained in detail.

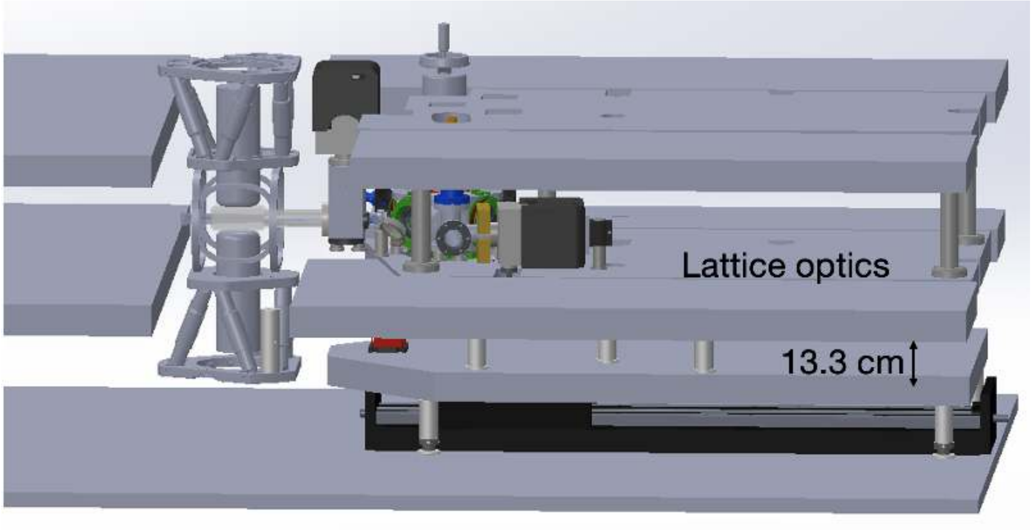

Figure 4.2: Main table breadboards without optics

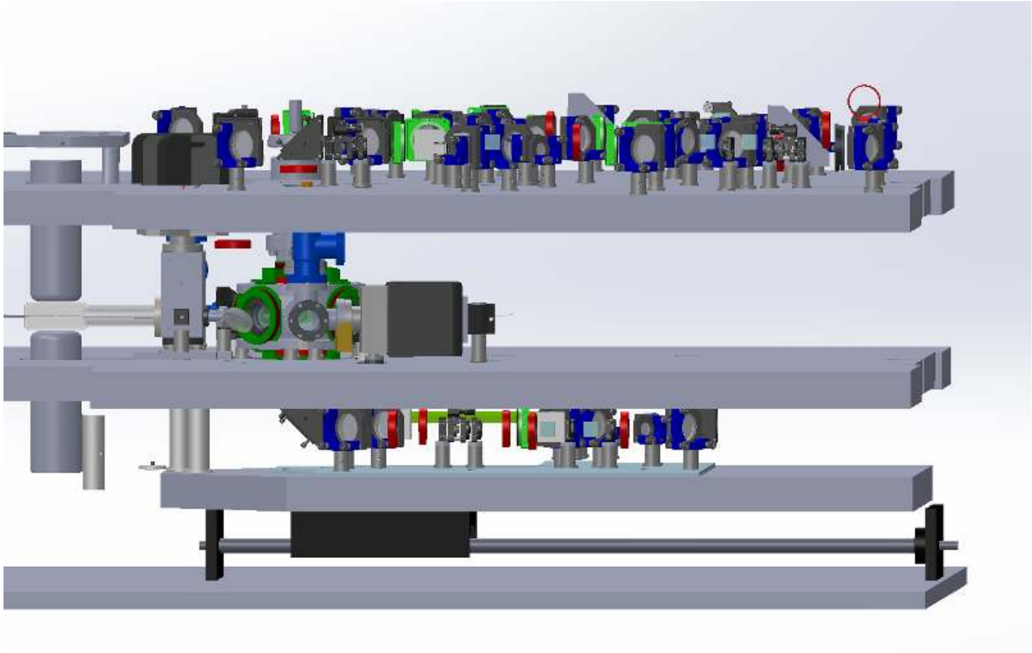

Figure 4.3: Main table breadboards with optics

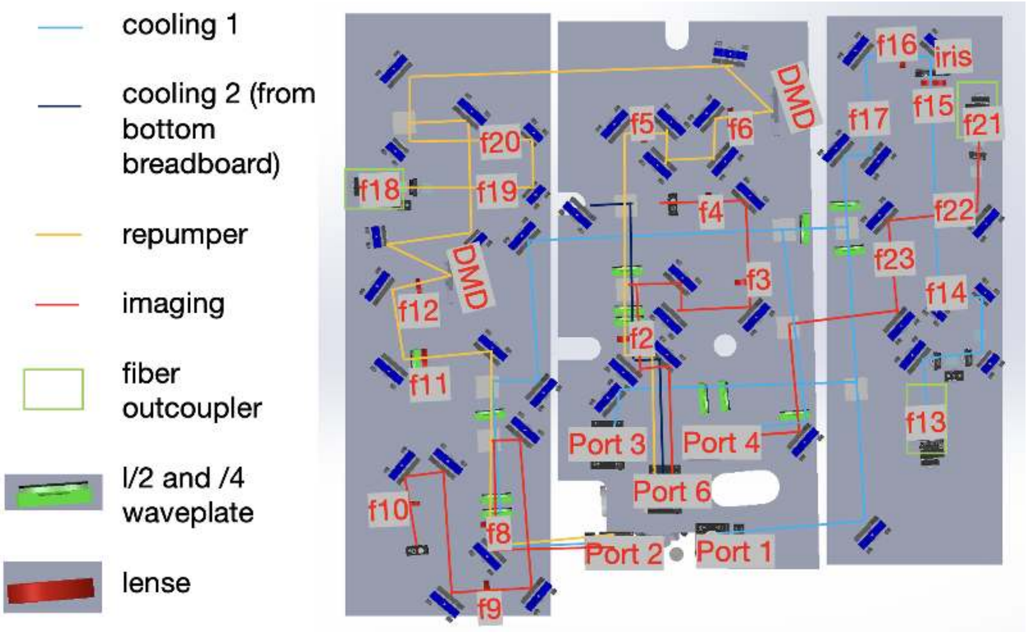

(a) MOT optics - top breadboard

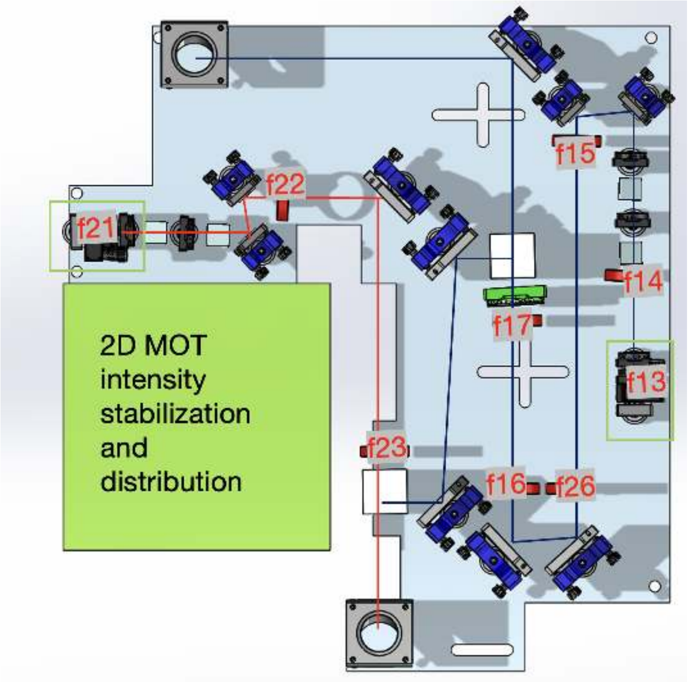

(b) MOT optics - bottom breadboard

Figure 4.4: MOT optics showing (a) top breadboard optics and (b) bottom breadboard optics.

Table 4.2: Lens specifications and part numbers

Lens Number	**Effective Focal Length (mm)**	**Thorlabs Part No.**
f1	250	AC508-250-B-ML
f2	250	AC508-250-B-ML
f3	200	ACT508-200-B-ML
f4	50	AC254-050-B
f5	250	AC508-250-B
f6	65	AC254-065-B
f7	250	AC508-250-B
f8	250	AC508-250-B
f9	400	AC254-400-B
f10	100	AC254-100-B
f11	250	AC508-250-B
f12	50	AC254-050-B
f13	25	AC127-025-B-ML
f14	-50	ACN127-050-B
f15	400	ACT508-400-B-ML
f16	-50	ACN254-050-B
f17	200	ACT508-200-B-ML
f18	18.4	C280TMD-B
f21	9.6	C060TMD-B
f22	-50	ACN127-050-B
f23	200	ACT508-200-B-ML

4.1.2 Repumper Path

To plan the repumper path, the first crucial decision is the choice of DMD, as it dictates the beam size and the so called blazing angle condition (for more details on this, see the DMD appendix section D). The options for a DMD at 852 nm are limited. We selected the highest resolution NIR-compatible model from Texas Instruments, the DLP650LNIR. We ordered two DLPLCR65NEVM units (the pre-packaged DMD) and two corresponding DLPLCRC410EVM driver boards, which interface with a computer via USB-mini. According to the DLP650LNIR datasheet, this DMD has a 10.8 µm micromirror pitch and a 1280 × 800 pixel array. The active area is approximately 13.8 mm × 8.6 mm. Hence to shape the entire beam profile, the beam incident on the DMD must have a maximum $1/e^2$ waist (radius) of approximately 4.3 mm. To achieve this, the output from a standard Thorlabs fiber (e.g. P3-780PM-FC-10) with a mode field diameter (MFD) of 5.3 µm is collimated using an $f = 18.4\,\mathrm{mm}$ lens (f_{18}), resulting in a collimated beam with a 1.883 mm radius. This beam is then expanded by 2.25× using a $4f$ telescope to a 4.237 mm radius. A sketch of this expansion is shown in figure 4.6. For all following optics schematics the legend is given in figure 4.5.

The subsequent challenge is relay imaging the DMD onto the atoms. As shown in table 4.1, we target a repumper size of approximately 1.2 inch at the atoms. After optimization, including the absorption imaging constraints, the solutions for the top and side paths are shown in figures 4.7 and 4.8. For the top path, as the vertical distance to the atoms is slightly shorter than for the side path, we choose a relay $4f$ with $f_1 = f_2 = 25\,\mathrm{cm}$ leading to the atoms and then an expanding $4f$ telescope with $f_5 = 25\,\mathrm{cm}$ and $f_6 = 6.5\,\mathrm{cm}$. This means we have a 3.86× beam expansion, leading 4.23 mm → 0.64 in radius at the atoms, which is a beam of 1.28 in diameter from the top. For the side path (Port 2), the setup is similar. Here we also have an $f_7 = f_8 = 25\,\mathrm{cm}$ relaying $4f$ telescope from the atoms. The DMD is then imaged onto this $4f$ with a magnifying $4f$ of $f_{12} = 5\,\mathrm{cm}$ and $f_{11} = 25\,\mathrm{cm}$, meaning a 5× magnification from 4.23 mm radius → 0.83 in radius at the atoms. So this beam is slightly larger than the other repumper, but that should be solvable by, e.g. turning off more pixels around the rim than on the other one.

Now that we have established an imaging path for both the top and the side, we can discuss the resolution of the DMD and the alignment sensitivity. The on-off contrast (extinction ratio) is discussed in Appendix D. The achievable extinction ratio is expected to be better than 10^{-6}, meaning it will definitely not be the limiting factor. The primary concern is whether, under realistic alignment and aberration conditions, the projected edge of the “pac-man” shape will remain sufficiently sharp. The larger of the two DMDs is imaged with a magnification of $M = 5$. Given a micromirror pitch of $d = 10\,\mu\mathrm{m}$, the projected pixel size at the plane of the atoms is $50\,\mu\mathrm{m}$. Consequently, the $\approx 400\,\mu\mathrm{m}$ straight dark tube in the “pac-man” shape is defined by essentially 8 deactivated rows on the DMD, ensuring the intrinsic geometric resolution of the DMD is more than sufficient. The next consideration is whether the optical imaging system imposes a diffraction limit. The longest focal length in the beam path is $f = 25\,\mathrm{cm}$. The diffraction-limited spot radius, given by the Rayleigh criterion, is $\Delta x = 1.22\frac{\lambda f}{D}$, where D is the limiting aperture. Assuming an effective aperture of approximately 1.2 in (30.48 mm) for the 852 nm repump beam at this lens, the diffraction limit evaluates to:

$$\Delta x = 1.22\frac{852 \times 10^{-9}\,\mathrm{m} \cdot 0.25\,\mathrm{m}}{0.030\,48\,\mathrm{m}} \approx 8.5\,\mu\mathrm{m}$$

This means that diffraction is also not the limiting factor. The final remaining question is the z-axis alignment sensitivity (depth of focus) of the “pac-man” projection. Because it can be experimentally challenging to perfectly conjugate the DMD imaging plane to the atom/lattice plane, the actual focal plane might deviate by 1 cm to 2 cm. To analyze how this defocusing affects the pattern’s edge sharpness, we perform a Fresnel propagation simulation of the optical field. The simulation models the DMD as a discrete binary amplitude mask with $10\,\mu\mathrm{m}$ pixels and numerically propagates an 852 nm plane wave through a 4f imaging system ($f_1 = 5\,\mathrm{cm}$, $f_2 = 25\,\mathrm{cm}$). The results of this simulation are presented in figure 4.9, which shows a 1D cross-section of the normalized intensity across the “pac-man” edge at various observation planes. At the ideal focal plane ($z = 25\,\mathrm{cm}$, red curve), the intensity exhibits a sharp, steep transition that closely traces the ideal geometric edge, confirming that the 4f system cleanly images the pixelated boundary. When the observation plane is shifted out of focus, the optical field blurs symmetrically along the propagation axis. At a focal shift of $\pm 2\,\mathrm{cm}$, the edge of the repumper beam bleeds into the nominally dark region by approximately 0.1 mm. This level of diffraction is entirely acceptable. Because the total designed width of the dark tube is 0.4 mm, a 0.1 mm intrusion on either side leaves a functional, albeit narrower, trapping region. If necessary, this blurring could be effortlessly compensated for by deactivating an additional 1 to 2 micromirror rows on the DMD to artificially widen the initial dark channel. However, at a more extreme defocus of $\pm 5\,\mathrm{cm}$, the light bleeding extends nearly 0.4 mm into the dark region, which would start to very strongly wash out the trapping channel and compromise the experiment. Fortunately, standard experimental alignment can be performed with higher precision than ±5 cm. Ultimately, this analysis confirms that there are no theoretical limitations to this optical layout, and aligning the DMD repumper to achieve a sufficiently sharp geometric projection is highly feasible in practice.

Figure 4.5: Legend of optical elements used in all subsequent setup schematics.

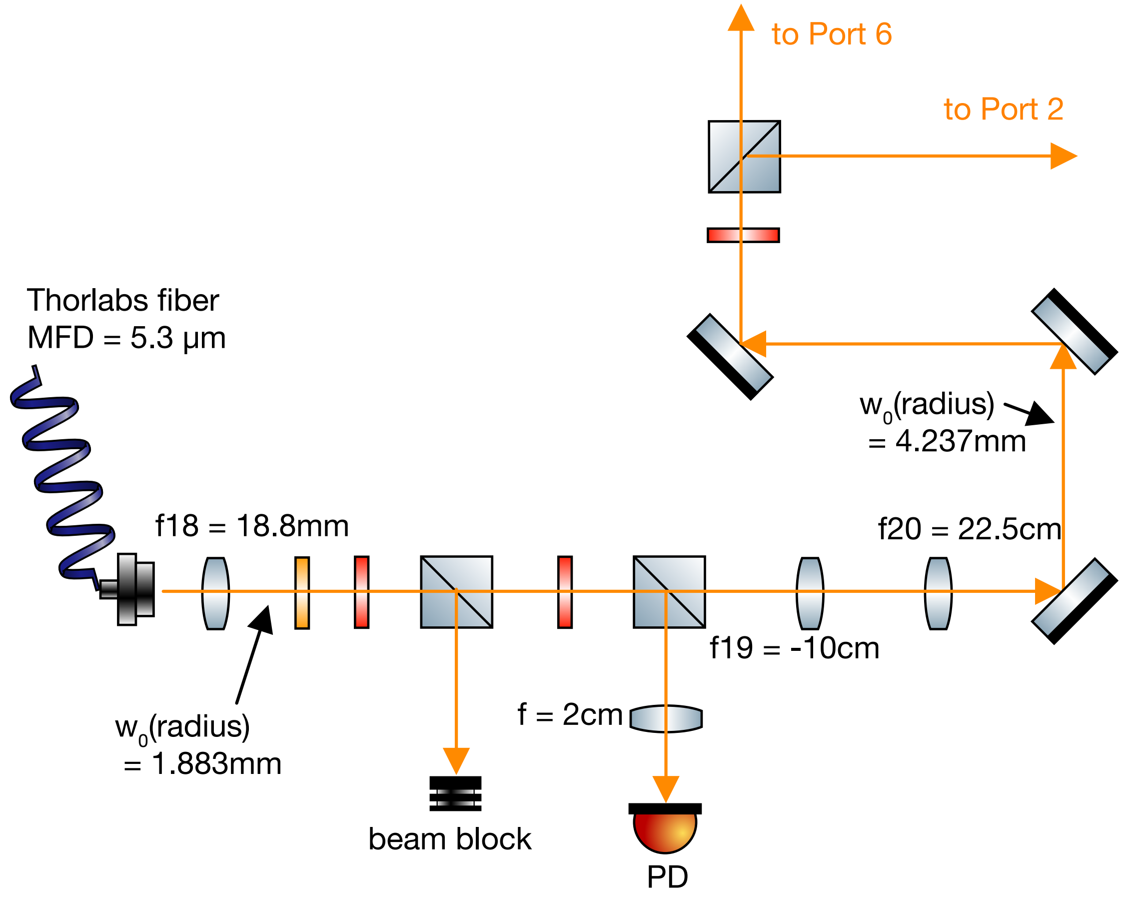

Figure 4.6: Schematic of the initial beam collimation and 4.5× expansion telescope of the repumper beam prior to the DMD. Beam block is for polarization cleaning, photo diode (PD) is for intensity stabilization. Legend can be found in figure 4.5.

Figure 4.7: Annotated top path schematic including all optics (waveplates, etc.) for the repumper and absorption imaging paths. Legend can be found in figure 4.5.

Figure 4.8: Annotated side path schematic including all optics (waveplates, etc.) for the repumper and absorption imaging paths. Legend can be found in figure 4.5.

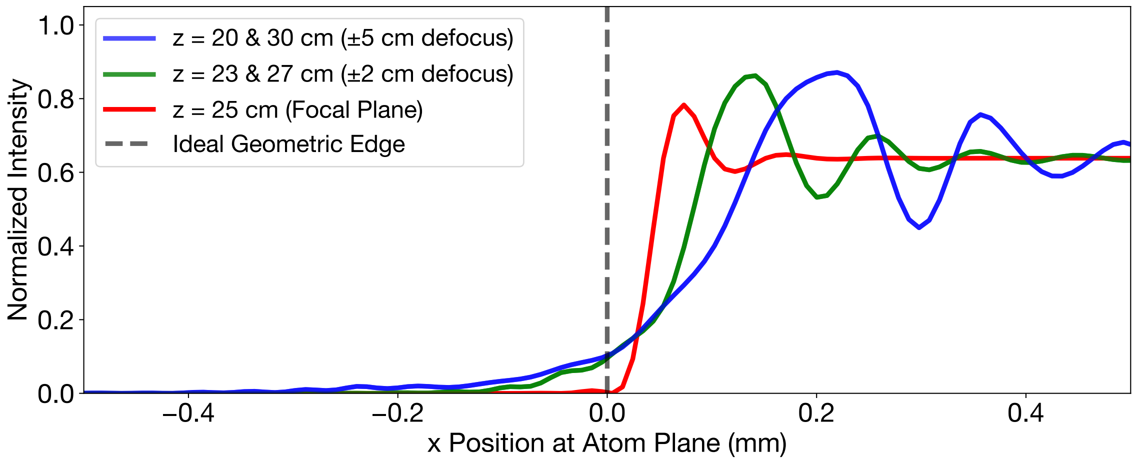

Figure 4.9: Simulation of the 1D cross-section of the normalized intensity profile at the edge of the "pac-man" dark channel. At the ideal focal plane ($z = 25\,\text{cm}$), there is a sharp transition near the geometric edge ($x = 0$). Shifting the focus by $\pm 2\,\text{cm}$ and $\pm 5\,\text{cm}$ causes blurring and diffraction ringing at the edge. The intensity ringing in the bright region (fluctuations of about 20% to 30%) is not a problem. The profile does not need to be perfectly flat, as long as the overall repumper intensity stays above the threshold needed to keep the atoms out of the dark state.

4.1.3 Imaging Beam Path

There are two absorption imaging outcouplers and expansion beam paths, one on the top and one on the bottom breadboard. However, both of them use the same lenses, shown in figure 4.10. With $f_{21} = 18.4\,\text{mm}$ the beam is expanded to a radius of $0.94\,\text{mm}$ for a fiber MFD of $5.3\,\mu\text{m}$ coming out of a Thorlabs fiber. Subsequently, the beam is expanded with a $4f$ telescope with magnification

$$M = \left|\frac{f_{23}}{f_{22}}\right| = \frac{40\,\text{cm}}{5\,\text{cm}} = 8,$$

using $f_{22} = -5\,\text{cm}$ and $f_{23} = 40\,\text{cm}$, leading to a large beam with a radius of

$$w = 0.94\,\text{mm} \times 8 = 7.52\,\text{mm} \approx 0.30\,\text{in}.$$

On the other side, both axes (port 2 and port 6) go through the same relay $4f$ system already described in the repumper path. After that, the imaging beam is de-expanded with $4f$ de-expansion optics. As shown in figure 4.7, for the path on port

6 the imaging beam is reduced by lenses $f_3 = 20\,\mathrm{cm}$ and $f_4 = 5\,\mathrm{cm}$, corresponding to a magnification of

$$M = \frac{f_4}{f_3} = \frac{1}{4}.$$

This leads to a beam radius of

$$w = \frac{7.52\,\mathrm{mm}}{4} = 1.88\,\mathrm{mm},$$

which comfortably fits on the camera sensor.

For the beam through port 2, the beam is de-expanded by $f_9 = 40\,\mathrm{cm}$ and $f_{10} = 10\,\mathrm{cm}$, again corresponding to a 4× reduction. This results in the same final beam radius of $w = 1.88\,\mathrm{mm}$.

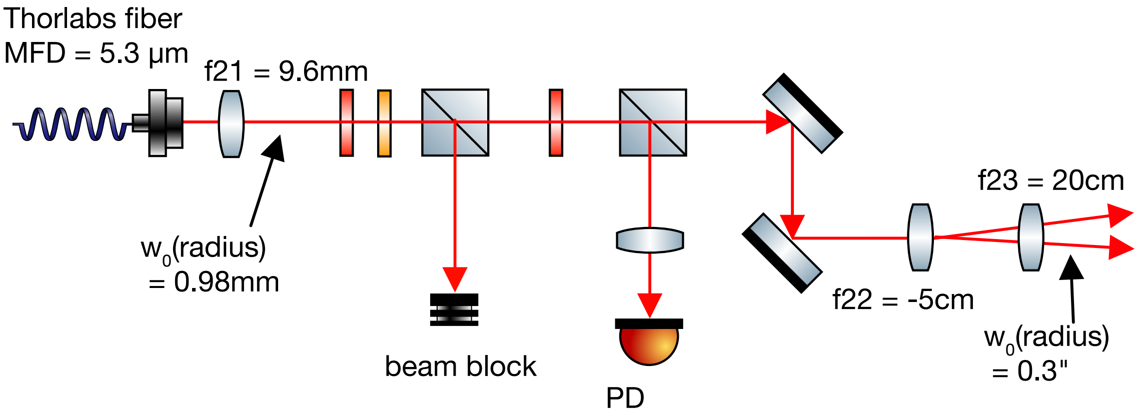

Figure 4.10: Schematic of the absorption imaging beam expansion and collimation optics. The beam block is for polarization cleaning and the photodiode for intensity stabilization. Legend can be found in figure 4.5.

4.1.4 Cooling Beam Path

The cooling optics require a significant amount of space due to their size and the fact that the beams are not retroreflected. There are two options: either one can deliver the light trough one fiber or with two/ three individual fibers for each axis. The latter approach allows for more individual power control and independent intensity stabilization per axis. We ultimately decided to route the top-bottom axis through a separate fiber from the side axes. An important experimental consideration is that two counter-propagating beams must always originate from the same fiber. Even with active intensity stabilization, relative intensity fluctuations between separate fibers can induce small force imbalances. In large MOTs, these imbalances can cause "spatiotemporal instabilities" that make the MOT center fluctuate and drift [68]. For this reason, the top-bottom beam is routed free-space vertically between the top and bottom optical layers. The expansion optics on the bottom breadboard are identical to those on the top breadboard. The schematic of the beam expansion and preparation is shown in figure 4.11. As mentioned in section 3.3, we aim to generate a "flat-top" Gaussian cooling beam to maximize the MOT loading rate. This is achieved by implementing an initial expansion stage, followed by spatial truncation with an iris, and then a second expansion stage, exactly as simulated in section 3.3. The beam emerges from an Alphanov fiber (LMA-PM-15), which has a core diameter of approximately 12.6(15) µm and a 7.6° FC/APC connection on the experiment side (SMA connection on the distribution table side). This divergent output is initially collimated to a beam radius of $w_0 = 1.076\,\mathrm{mm}$ using an $f_{13} = 25\,\mathrm{mm}$ lens. Following collimation, the beam passes through a polarization cleaning stage consisting of a $\lambda/2$ waveplate and a polarizing beam splitter (PBS), with the rejected polarization dumped into a beam block. A second $\lambda/2$ waveplate and PBS pick off a small fraction of the light, which is focused by an $f = 5\,\mathrm{cm}$ lens onto a photodiode (PD) for active intensity stabilization.

After the preparation stage, the initial expansion is performed using a magnifying telescope with a magnification of $M = |f_{15}/f_{14}| = 40\,\mathrm{cm}/5\,\mathrm{cm} = 8$, yielding a beam radius of 8.61 mm. To achieve the desired "flat-top" Gaussian profile, this expanded beam is truncated by an iris that cuts off approximately 50% of the beam, resulting in a transmitted beam radius of $w_0 \approx 4.3\,\mathrm{mm}$. Finally, a second telescope expands the beam one last time with a magnification of $M = |f_{17}/f_{16}| = 20\,\mathrm{cm}/5\,\mathrm{cm} = 4$. This sequence results in a final beam diameter of approximately 1.3 in with a flattened intensity profile, which is then distributed to the respective vacuum chamber ports.

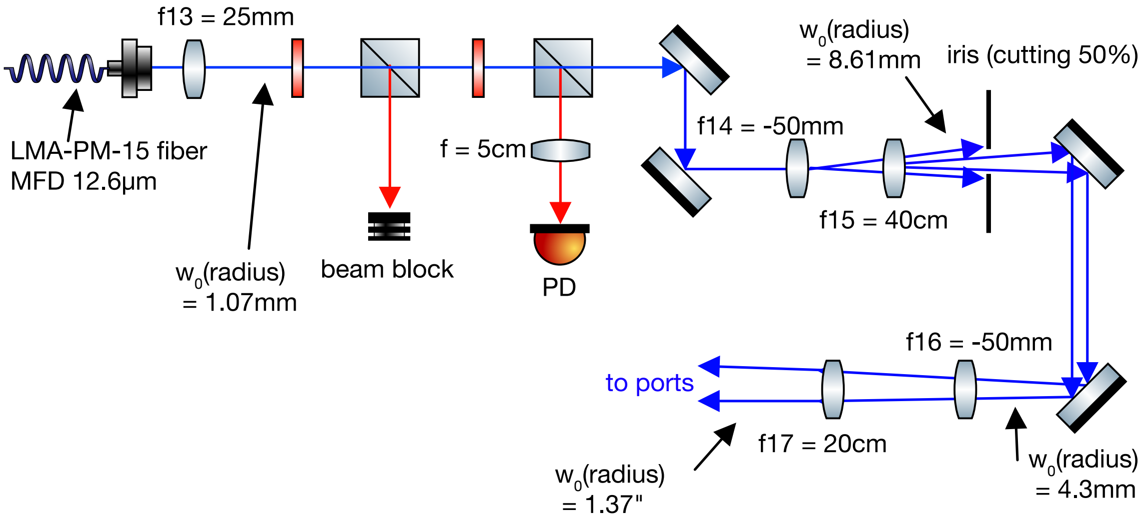

Figure 4.11: Schematic of the cooling beam preparation and expansion optical layout. The beam is collimated from the fiber, passes through polarization cleaning and active intensity stabilization stages, and is subsequently expanded by two successive telescopes. An intermediate iris truncates the beam to flatten the final intensity profile before it is routed to the experimental ports. Legend can be found in figure 4.5.

4.1.5 2D MOT Distribution Breadboard

According to the Infleqtion 2D MOT datasheet, the module features three input fibers. Two of these are designated as the main 2D MOT fibers, each requiring a dual-wavelength input: $> 20\,\text{mW}$ of cooling power and $> 2\,\text{mW}$ of repump power per arm. The third is the push beam fiber, which requires a minimum of $2\,\text{mW}$ for cooling and $0.2\,\text{mW}$ for repumping. However, experience from the E2 experiment indicates that achieving optimal operating conditions often requires two to three times the optical power specified by Infleqtion. Furthermore, the datasheet recommends a cooling beam detuning of -3Γ, which corresponds to $-3 \times 2\pi \times 5.23\,\text{MHz}$ for Cesium. Additionally, previous work on E2 demonstrated that the 2D MOT operates most stably when the cooling and repump beams are actively intensity-stabilized and spatially combined on the main experimental table. By utilizing very short optical fibers to then deliver this combined light to the Infleqtion module, polarization and intensity fluctuations within the fibers are minimized, leading to a highly stable 2D MOT. An additional optimization derived from the E2 setup is the use of 99:1 fiber-optic splitters (Thorlabs PN850R1A1). By continuously monitoring the 1% output port with a fiber-coupled photodiode, it is possible to track the coupling stability and the exact power delivered to the 2D MOT in real time, without ever needing to disconnect the fibers from the vacuum cell. The complete optical layout for this recombination and distribution setup is detailed in figure 4.12.

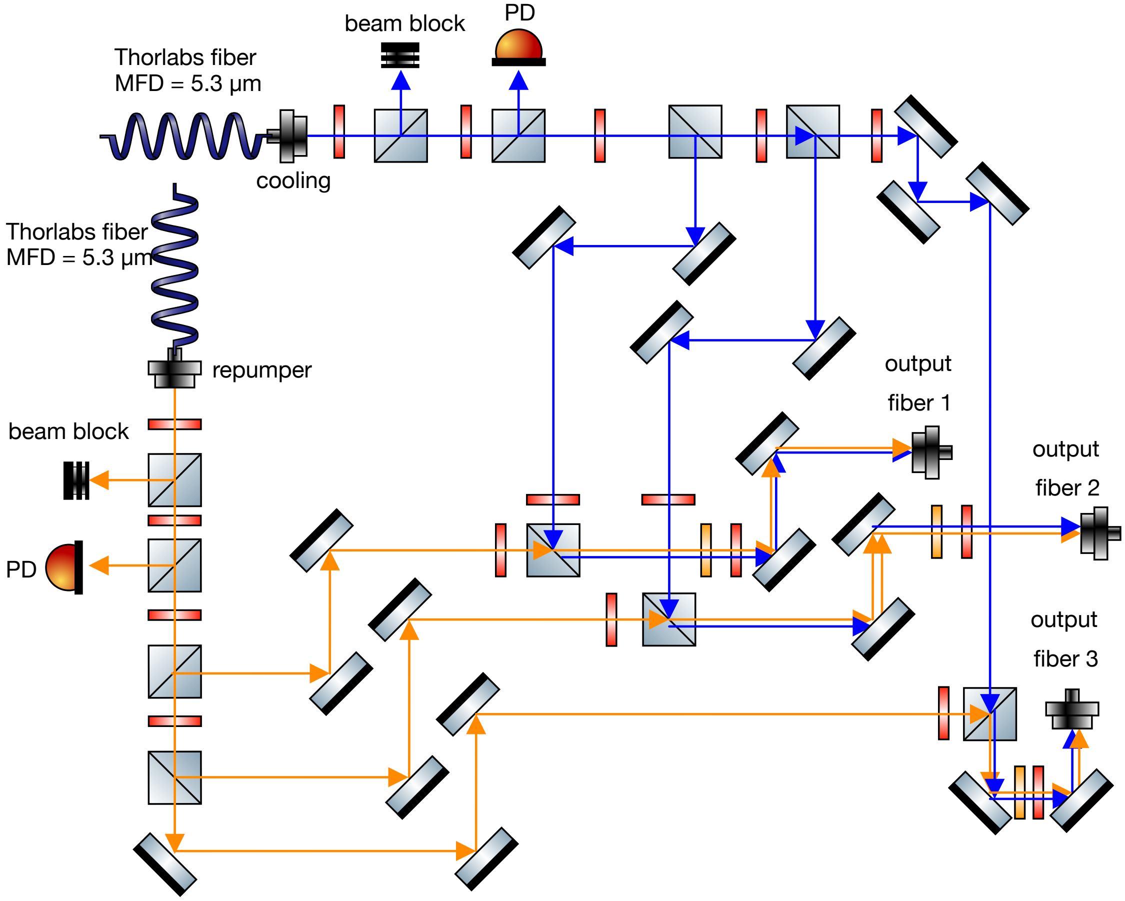

Figure 4.12: Optical layout of the 2D MOT distribution breadboard on the main experimental table. The cooling and repump beams are intensity-stabilized and combined before being coupled into short fibers. A 99:1 fiber splitter allows for continuous power monitoring without disconnecting the main output fibers from the 2D MOT module. Legend can be found in figure 4.5.

4.2 Laser Locking and Laser Distribution Setup

4.2.1 Master-Slave Locking Scheme and Spectroscopy

In section 4.1 we learned in detail which laser wavelengths we need, at what power, and with what detuning capabilities. In this section we look at how the laser setup is designed to meet these requirements. The first question is which lasers to use and how to lock them. Our group already has a ^{133}Cs experiment in which very good experience has been made using the Vescent D2 DBR laser module in combination with Vescent locking electronics, as it runs reliably and we already have the expertise in our group. We therefore decided to use a similar system. This means we lock a "master" laser on a Vescent D2 spectroscopy module and then offset-lock two "slave" lasers to it. Vescent offers locking electronics and controllers for all of this. The main question is which laser to lock to which transition. From table 4.1 we know that we need $F = 3 \rightarrow 4$, $F = 4 \rightarrow 4$, and $F = 4 \rightarrow 5$ light. As we need a certain tuning range for all of them, we want to reach these precise wavelengths using single- and double-pass acousto-optic modulators (AOMs). AOMs typically have a center frequency of around 80 MHz to 200 MHz and a tuning range of ±15 MHz to 25 MHz. Generally, we prefer double-pass AOMs for two main reasons: first, they allow a larger detuning range, and second, if a cat's-eye configuration is used (explained in section 4.2.2), sweeping the AOM frequency does not change the output beam angle. Maintaining a stable beam pointing is important as angular deviations degrade the efficiency of subsequent fiber coupling. The main disadvantage of a double-pass configuration is the increased optical power loss incurred by diffracting the light through the AOM crystal twice. However, given our available laser power, this loss is an acceptable trade-off that we can afford. The largest tuning range is required for the cooling beams, so the cooling beam should definitely be on a double pass, and ideally the lock point of the cooling laser can also be adjusted to extend the accessible range further. For this reason it is advantageous to use the cooling laser as the slave laser. The slave laser is offset-locked electronically to the master laser, so its frequency can be adjusted relatively easily. As the Vescent D2-260 beat-note detector has a bandwidth of 250 MHz–9 GHz, the frequency difference between the master and slave lasers must fall within that range. Taking all these constraints into account, a good locking scheme (also recommended to us by E2) is illustrated in figure 4.13. The master laser is locked to the $F = 2 \rightarrow F' = 2/3$ crossover of a Vescent D2-210 Cs absorption spectroscopy cell, indicated by the red arrow in figure 4.13. The two slave lasers are then each offset-locked to the master laser such that they are +70 MHz detuned from the $F = 2 \rightarrow F' = 4$ transition, shown by the blue arrows. This allows a combination of 70 MHz, 80 MHz, and 140 MHz single- and double-pass AOMs, as illustrated in figure 4.13, to achieve all the required wavelengths and tuning ranges. As described in section 4.1, we also want the capability to perform LGM cooling in the MOT chamber. To achieve this, we need an $F = 4 \rightarrow F' = 4$ cooling beam and an $F = 3 \rightarrow F' = 4$ LGM repumper. Since we want to use the main MOT cooling beams for LGM cooling, both of these wavelengths should be accessible via the same double-pass AOM setup used for the cooling beams. This requires shifting the slave laser lock point slightly, such that the double-pass AOM now addresses the $F = 4 \rightarrow F' = 4$ transition instead. This alternative operating point is indicated by the pink arrow in figure 4.13: by slightly decreasing the slave laser offset frequency, the required LGM wavelengths are reached without any changes to the optical setup. The $F = 3 \rightarrow F' = 4$ repummper is achieved using an EOM, the setup for that can be seed as breadboard 8 in figure 4.14.

Because we had to move between two labs in the middle of this master's thesis, and because we may also want to switch lasers at some point, we built the whole locking setup to be as modular as possible. To achieve this, each Vescent laser is placed on its own breadboard with two fiber-coupled outputs: one goes directly to a wavemeter, and the other is routed to one of the distribution breadboards (where the double-pass AOMs are), from which a small portion of the light is then sent to the beat breadboard containing the ^{133}Cs cell. The individual optical layouts for the Vescent laser system (breadboards 1–4) are detailed in figures 4.15, 4.16, 4.17, and 4.18, respectively. A photo of their experimental realization is shown in figure F.3 in the appendix.

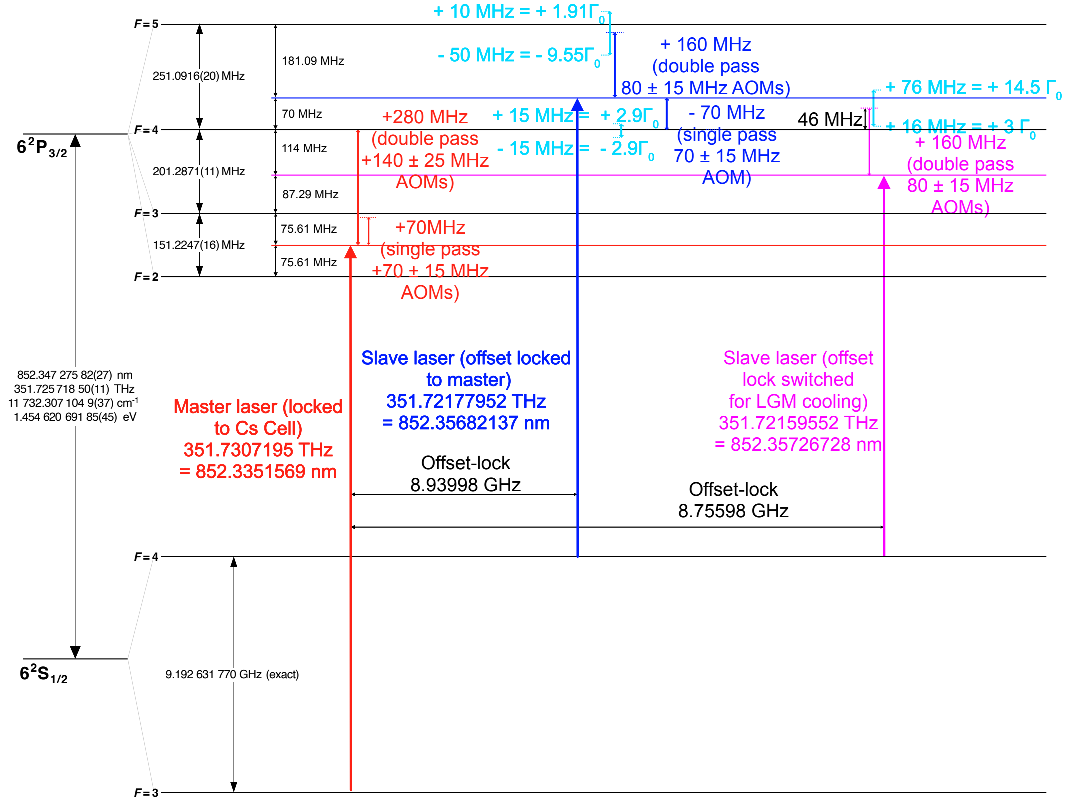

Figure 4.13: Vescent laser locking scheme. The red arrow represents the master laser, which is locked to the $F = 3 \to F' = 2/3$ crossover using the Vescent D2-210 Cs cell. This master laser is used for the $3 \to 4$ repumper and the $3 \to 3$ state preparation beams. The $3 \to 4$ transition is addressed with a double-pass AOM, whereas $3 \to 3$ is addressed with a single-pass AOM, as indicated by the small red arrows. For MOT cooling, the slave laser is offset-locked by $\approx 8.94\,\mathrm{GHz}$ to be blue-detuned by $+70\,\mathrm{MHz}$ from the $4 \to 4$ transition. This provides a large tuning range on the $4 \to 5$ MOT cooling transition ($\approx -9\,\Gamma_0$ to $+2\,\Gamma_0$) via a double-pass AOM. The $4 \to 4$ depumper is addressed with a single-pass AOM. If LGM cooling in the main chamber is desired, the offset lock point must be switched to $8.76\,\mathrm{GHz}$, as indicated by the pink arrow. At this offset, the cooling double-pass AOMs become resonant with the $4 \to 4$ line. An additional EOM placed in that double-pass path can then modulate the LGM repump power onto the beam. Figure adapted from [30]. Legend can be found in figure 4.5.

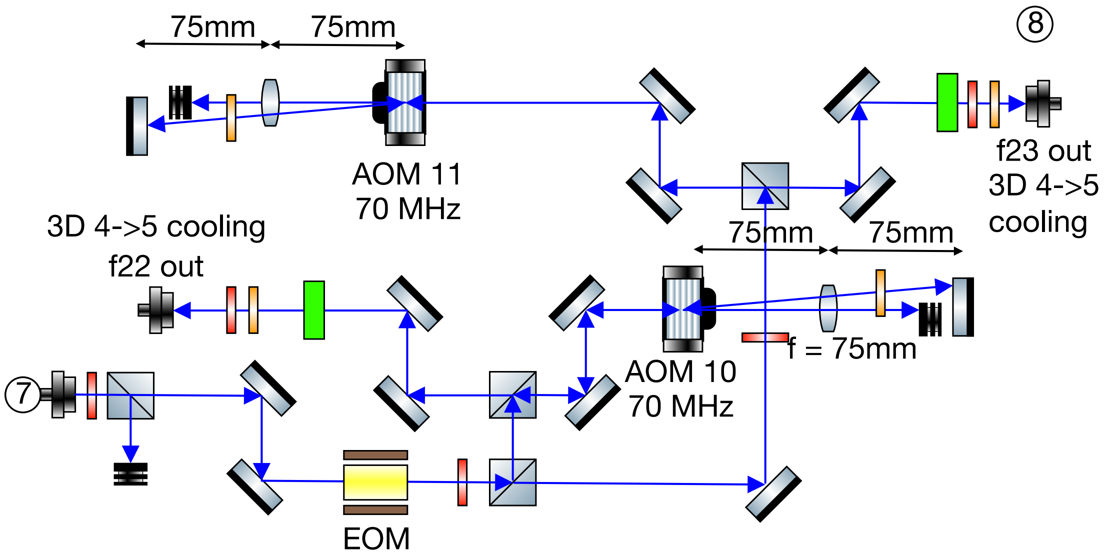

Figure 4.14: 3D cooling distribution breadboard (breadboard 8). Receives amplified light from slave laser 2 via the Toptica TA (breadboard 3) and distributes it to the 3D MOT cooling beams. Legend can be found in figure 4.5.

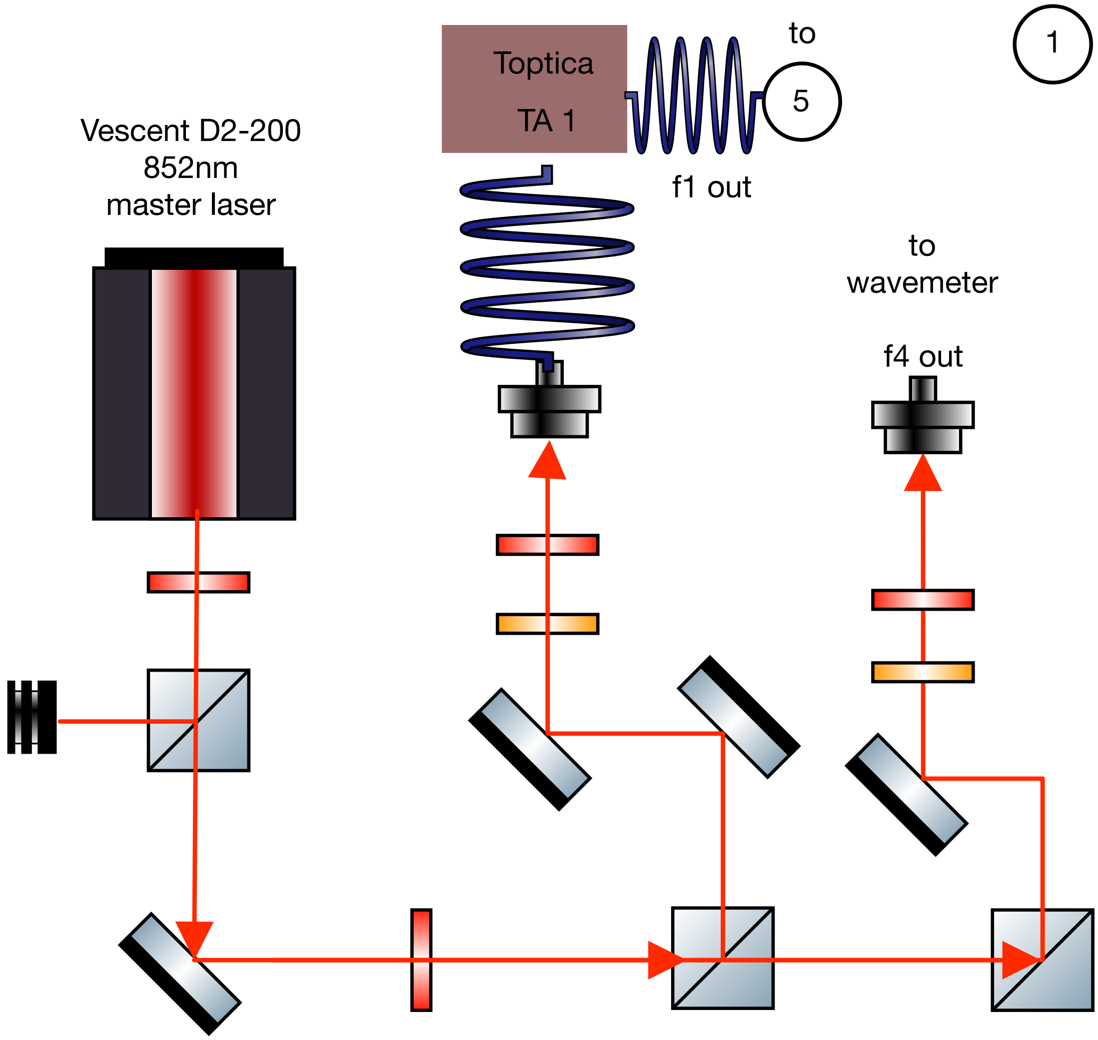

Figure 4.15: master laser (breadboard 1). Output split to wavemeter and AOM distribution breadboard (breadboard 5); part of the light is routed to the beat breadboard (breadboard 4). Legend can be found in figure 4.5.

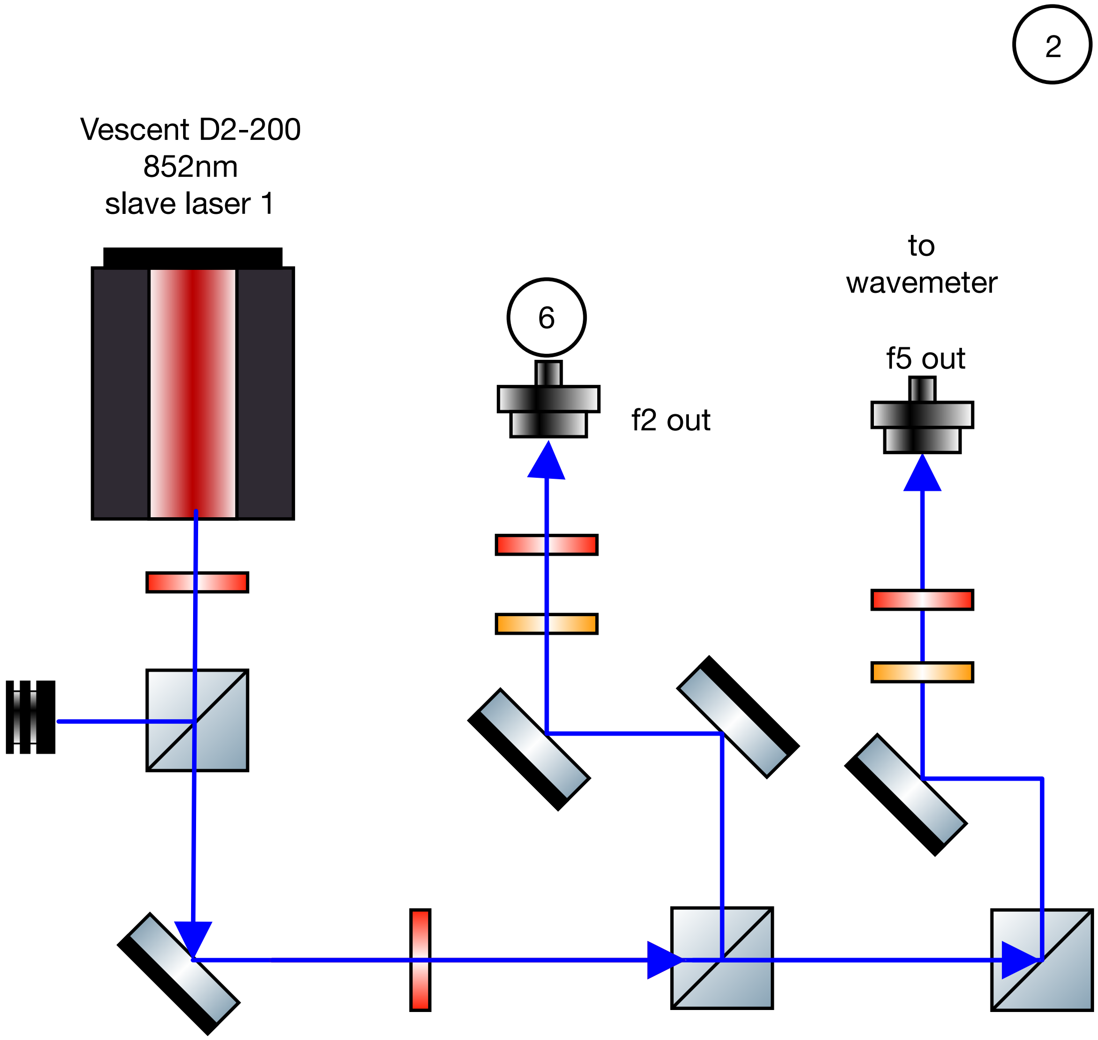

Figure 4.16: slave laser 1 (breadboard 2). Output split to wavemeter and AOM distribution breadboard, feeding the 2D MOT distribution breadboard (breadboard 6). Legend can be found in figure 4.5.

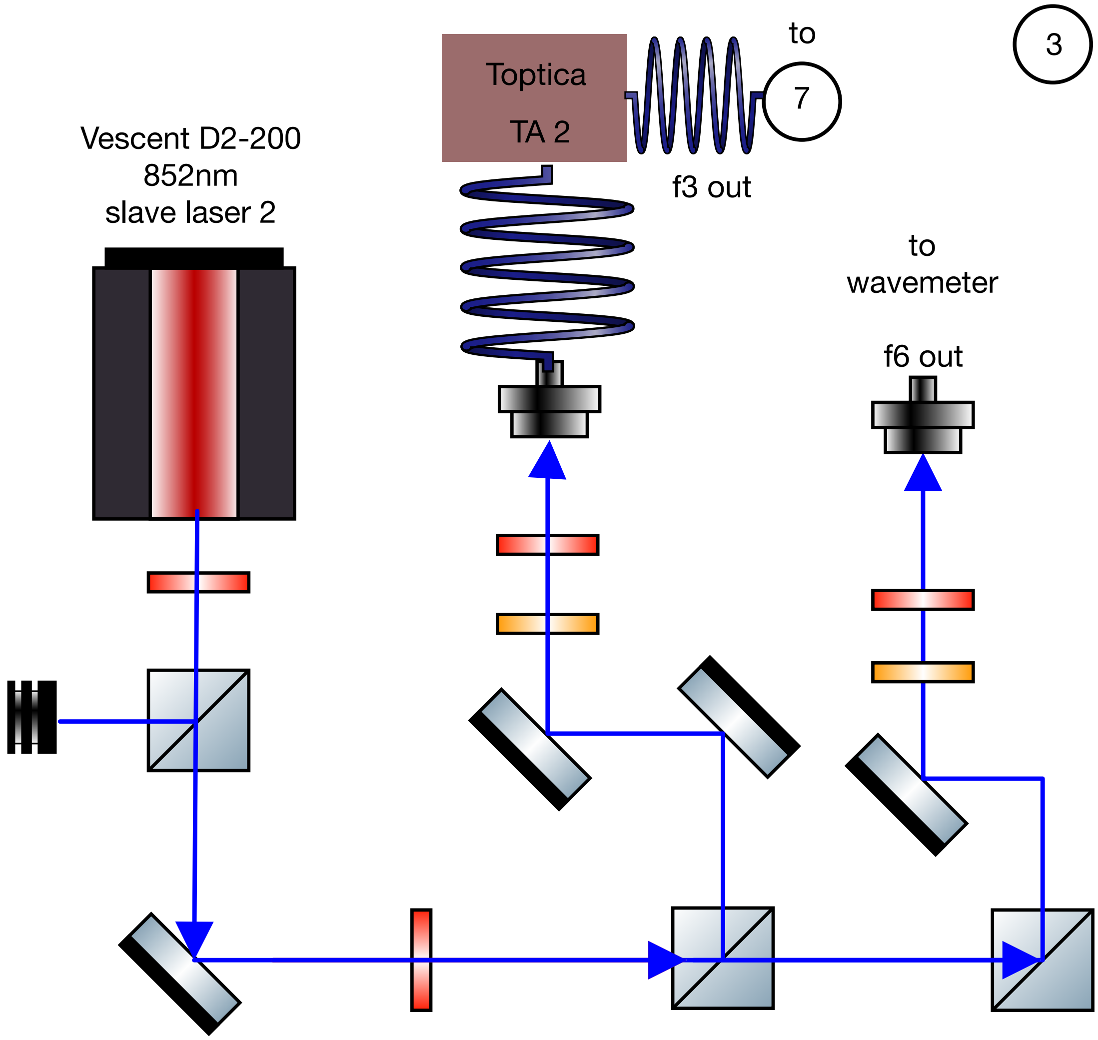

Figure 4.17: slave laser 2 (breadboard 3). Output amplified by a Toptica TA and distributed via breadboard 7 to the 3D cooling breadboard (breadboard 8). Legend can be found in figure 4.5.

Figure 4.18: Beat breadboard (breadboard 4). The master laser is locked to the ^{133}Cs spectroscopy cell; slave lasers 1 and 2 are offset-locked to the Master via beat-note detectors (see figure 4.13). Legend can be found in figure 4.5.

The Vescent D2-210 spectroscopy module is connected to the error input of a Vescent D2-125-PL laser servo. The servo implements a PI^{2}D controller with the peak-locking option, both described in detail below. The RF output of the D2-125-PL is connected to the RF modulation input of the Vescent SLICE-DLC laser controller, which in turn drives the Vescent laser itself. Additionally, the servo output of the D2-125-PL is connected to the servo input of the SLICE-DLC. To observe the error signal, one connects the DC Error output and Ramp TTL output of the D2-125-PL directly to an oscilloscope using standard BNC cables, setting the oscilloscope to trigger on the Ramp TTL. In ramp mode, the laser servo sends a voltage ramp to the laser controller, which sweeps the laser current and therefore its frequency. If the spectroscopy cell is well-aligned and sufficient laser power enters the cell, the resulting error signal on the oscilloscope displays the absorption spectrum. The ramp amplitude is set with the Ramp Amp knob, and the ramp center can be adjusted over a $\pm 5\,\mathrm{V}$ range with the ramp offset knob, both located on the front panel. Any coarser frequency tuning must be performed directly on the SLICE-DLC by adjusting the laser current or temperature.

PID Basics

A standard PID controller produces a control signal $u(t)$ from the error signal $e(t)$ according to

$$u(t) = K_\mathrm{p}\, e(t) + K_\mathrm{i} \int_0^t e(t')\,\mathrm{d}t' + K_\mathrm{d}\,\frac{\mathrm{d}e}{\mathrm{d}t}, \tag{4.1}$$

where K_p, K_i, and K_d are the proportional, integral, and derivative gains, respectively. The integral term ensures that any constant (DC) frequency offset is driven to zero over time. However, for a linearly drifting error, such as that caused by slow thermal or mechanical drift, a single integrator can only keep up with the drift but cannot eliminate the residual error. The D2-125-PL therefore implements a PI^{2}D controller, which adds a second integrator:

$$u(t) = K_\mathrm{p}\, e(t) + K_{\mathrm{i}_1} \int_0^t e(t')\,\mathrm{d}t' + K_{\mathrm{i}_2} \int_0^t \int_0^{t'} e(t'')\,\mathrm{d}t''\,\mathrm{d}t' + K_\mathrm{d}\,\frac{\mathrm{d}e}{\mathrm{d}t}. \tag{4.2}$$

The second integrator builds up a correction that itself grows in time, allowing the servo to track and eliminate ramp disturbances and thereby achieve tighter long-term lock stability.

A PID servo requires an error signal that passes through zero with a sign change at the desired lock point, so that the servo knows in which direction to push: a negative error signal should increase the laser frequency, and a positive error signal should decrease it. A raw spectroscopy absorption peak does not provide such a zero crossing, it has the same sign on both sides of the maximum, so without further processing the servo cannot lock directly to the peak and would instead have to lock to the side of a spectroscopy feature. Side-of-fringe locking works, but it has significant drawbacks: the lock point is not at the exact transition frequency, and it is sensitive to intensity noise and electronic drift, both of which shift the apparent zero crossing and hence the laser frequency. The peak-locking option addresses this by electronically computing the derivative of the spectroscopy signal, turning each peak into a linear, sign-changing zero crossing. The principle is identical to that of a lock-in amplifier: the D2-125-PL modulates an additional RF signal at $\omega_{\mathrm{RF}} = 2\pi \times 4\,\mathrm{MHz}$ onto the laser current via the RF output, producing a small frequency modulation. The spectroscopy signal U then becomes

$$U\Big(\nu_0 + A\sin(\omega_{\mathrm{RF}}t)\Big) \approx U(\nu_0) + A\sin(\omega_{\mathrm{RF}}t)\,\frac{\mathrm{d}U}{\mathrm{d}\nu}\bigg|_{\nu_0} + \mathcal{O}(A^2), \tag{4.3}$$

where A is the modulation amplitude. This signal is then demodulated by multiplying with a reference $\sin(\omega_{\mathrm{RF}}t + \varphi)$ and low-pass filtering. Using the identity $\sin(\omega_{\mathrm{RF}}t)\sin(\omega_{\mathrm{RF}}t + \varphi) = \frac{1}{2}\Big[\cos\varphi - \cos(2\omega_{\mathrm{RF}}t + \varphi)\Big]$, the low-pass filter removes the $2\omega_{\mathrm{RF}}$ term, leaving

$$\text{error signal} \propto \frac{A}{2}\cos\varphi\,\frac{\mathrm{d}U}{\mathrm{d}\nu}\bigg|_{\nu_0}. \tag{4.4}$$

The error signal is thus proportional to the derivative of the spectroscopy signal. To maximise it, the phase difference φ between the modulation and demodulation references should be set to 0 or π, which is adjusted via the phase knob on the front panel of the D2-125-PL.

To generate the reference signal, we use a Vescent D2-210 spectroscopy module to perform saturated absorption spectroscopy. This technique uses a strong pump beam to deplete the ground state of atoms in a specific velocity class, saturating the transition [69]. A weak, counter-propagating probe beam derived from the same laser scans across the same frequencies. When the laser frequency matches the atomic resonance for atoms with zero longitudinal velocity ($v = 0$), both beams interact with the same atoms, and the probe beam experiences reduced absorption, resulting in a narrow transmission peak. A typical spectrum in absorption spectroscopy is shown in figure 4.19(a), displaying true transmission peaks (T) alongside crossover resonances (C). Crossovers occur when an atom moving with a non-zero velocity v is Doppler-shifted such that the pump beam is resonant with one transition (ν_1) and the counter-propagating probe beam is resonant with another transition (ν_2) sharing the same ground state. The resonance conditions in the laboratory frame are $\nu \approx \nu_1(1+v/c)$ for the pump and $\nu \approx \nu_2(1-v/c)$ for the probe. Solving this system of equations for the laser frequency ν yields the crossover frequency:

$$\nu_c \approx \frac{\nu_1+\nu_2}{2}. \tag{4.5}$$

This demonstrates that the crossover resonance appears exactly halfway between the two true transitions. By substituting ν_c back into the resonance conditions, we can solve for the specific velocities of the atoms contributing to this signal. We find that the crossover resonance is generated by two distinct, symmetric velocity classes:

$$v_1 = c\,\frac{\nu_2-\nu_1}{\nu_1+\nu_2} \quad \text{and} \quad v_2 = -c\,\frac{\nu_2-\nu_1}{\nu_1+\nu_2}. \tag{4.6}$$

Because the crossover signal draws from these two distinct velocity groups (v_1 and v_2) in the Maxwell-Boltzmann thermal distribution, rather than the single zero-velocity class ($v = 0$) probed by the main transitions, these crossovers are often stronger than the actual transmission peaks.

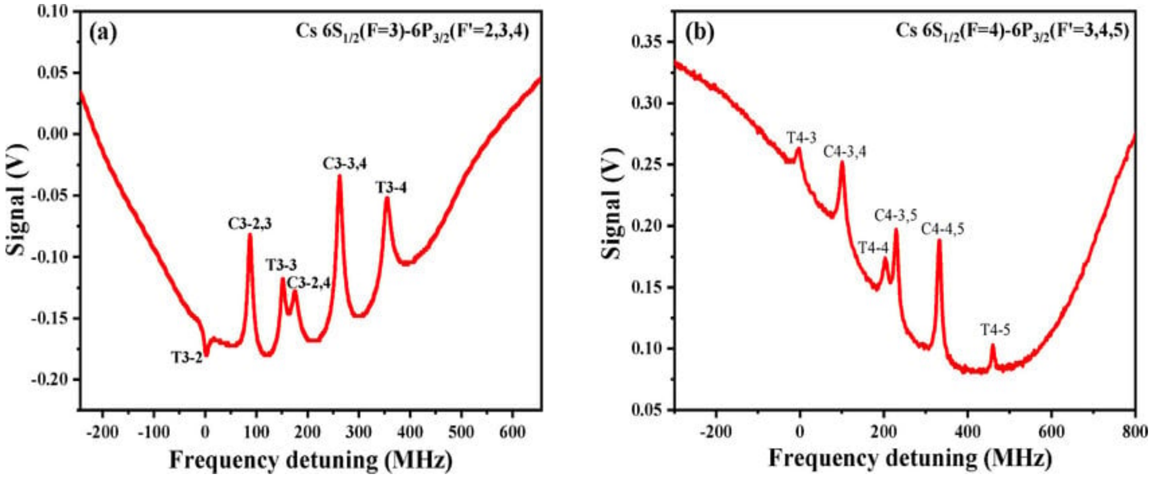

Figure 4.19: Saturated absorption spectra of the ^{133}Cs D_2 line. Panel (a) shows the $6S_{1/2}(F = 3) \rightarrow 6P_{3/2}(F' = 2, 3, 4)$ transitions and corresponding crossover resonances, and panel (b) shows the $6S_{1/2}(F = 4) \rightarrow 6P_{3/2}(F' = 3, 4, 5)$ transitions alongside their crossover resonances. Figure taken from [70].

During the setup, we put the servo into ramp mode to sweep the laser frequency and monitor the resulting signals on an oscilloscope. As seen in figure 4.20, the green trace represents the ramp voltage driving the laser controller. As it moves from left to right, it increases the injection current to the laser diode, slightly heating it and changing its refractive index. This makes the effective optical cavity longer, meaning the laser emits a longer wavelength (λ). Because frequency and wavelength are inversely proportional ($f = c/\lambda$), moving from left to right on the screen corresponds to sweeping from higher to lower frequencies. The blue trace indicates the ramp TTL signal, and we aim to lock during its "low" state, which corresponds to the positive slope of the green voltage ramp.

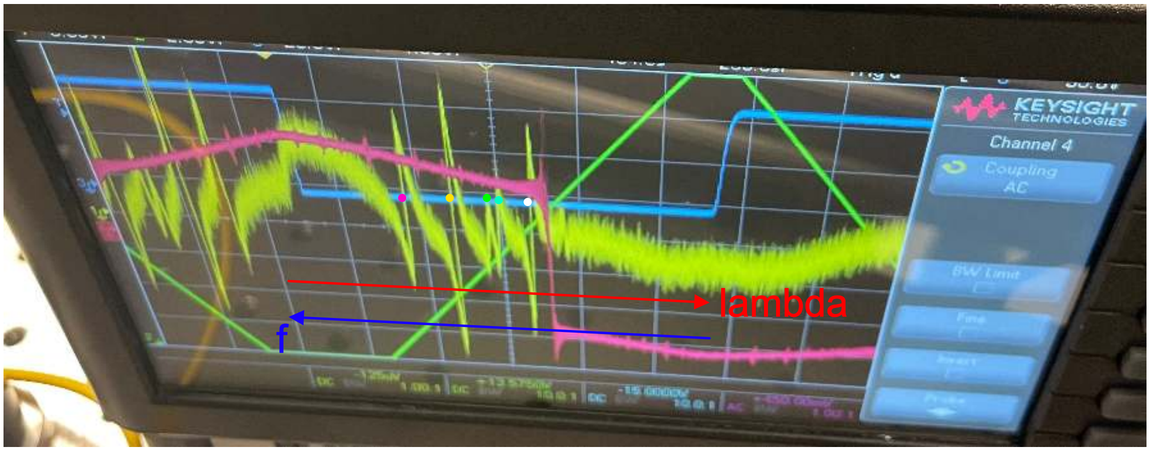

Figure 4.20: Oscilloscope trace of the laser lock setup. The green trace is the voltage ramp driving the laser diode, the blue trace is the ramp TTL, and the orange trace is the DC-coupled error signal from the D2-125-PL. The gain sign switch on the servo is set to positive (pointing up), requiring a physical positive slope on the frequency spectrum to lock. Colored markers indicate the zero-crossings corresponding to specific atomic transitions and crossovers: pink ($T_{3\to4}$), orange ($C_{3\to3,4}$), green ($C_{3\to2,4}$), blue ($T_{3\to3}$), and white ($C_{3\to2,3}$).

The orange trace in figure 4.20 is the DC-coupled error signal from the D2-125-PL. Assuming the gain sign switch is set to positive (pointing up), the servo requires an error signal with a positive slope to lock. Since the physical frequency decreases from left to right on the scope, locking to a physical positive slope on the frequency spectrum means we must lock to the region where the error signal displays a negative slope on the oscilloscope's time axis. A wavemeter reading of approximately 852.3354 nm confirms we are operating on the Cesium D2 line, specifically the $6S_{1/2}(F = 3) \to 6P_{3/2}(F' = 2, 3, 4)$ hyperfine transitions. Comparing the peak-lock error signal to the reference spectroscopy signal in figure 4.19(a), we can identify the specific zero-crossings marked on the oscilloscope trace. The pink dot corresponds to the $T_{3\to4}$ transition, orange to the $C_{3\to3,4}$ crossover, green to the $C_{3\to2,4}$ crossover, blue to the $T_{3\to3}$ transition, and white to the $C_{3\to2,3}$ crossover. Finally, to achieve a robust lock, the D2-125-PL PID parameters were carefully tuned. The typical tuning procedure begins by turning off all integral and derivative gains. The proportional gain K_p, which provides the immediate restoring force, is increased until the lock loop begins to oscillate, and is then slightly reduced to a stable margin.

To characterize the effectiveness of the locking scheme, we record the output wavelength of both lasers using a High finesse WS7-60 wavemeter (absolute accuracy of 60 MHz and wavelength deviation sensitivity of 2 MHz) under four conditions: neither laser locked, master laser locked only, slave laser locked only, and both lasers locked. Figure 4.21 shows the frequency deviation from the mean as a function of measurement index for each condition. When both lasers are free-running, both exhibit large, slow drifts on the order of several MHz, reflecting the thermal and mechanical instabilities of the free-running lasers. Locking the master laser to the saturated absorption spectroscopy signal immediately suppresses these fluctuations, confining its frequency to a narrow band around the cesium resonance. When only the slave laser is locked, that is, the master is free-running while the slave is offset-locked, the slave inherits the master's drift and shows only modest improvement over the free-running case.

The statistical distributions of these frequency deviations are shown in figure 4.22, where each panel displays a normalized histogram together with a Gaussian fit and the fitted standard deviation σ. In the free-running case, both lasers exhibit a standard deviation of approximately 2.9 MHz. Locking the master laser alone reduces its standard deviation to $\sigma = 0.29\,\mathrm{MHz}$, an improvement of roughly an order of magnitude, while the free-running slave laser remains at $\sigma = 1.51\,\mathrm{MHz}$, partially benefiting from reduced master drift through the injection path. When both lasers are locked, the master reaches $\sigma = 0.32\,\mathrm{MHz}$ and the slave narrows to $\sigma = 0.41\,\mathrm{MHz}$. It should be noted that these measured standard deviations for the fully locked configuration are artificially limited by the deviation sensitivity and internal noise floor of the wavemeter itself, meaning the true laser frequency stability is likely much better. Regardless, even these measured upper bounds are well below the natural linewidth of the cesium D_2 transition ($\Gamma/2\pi \approx 5.2\,\mathrm{MHz}$), demonstrating that the locking scheme is sufficient for the MOT.

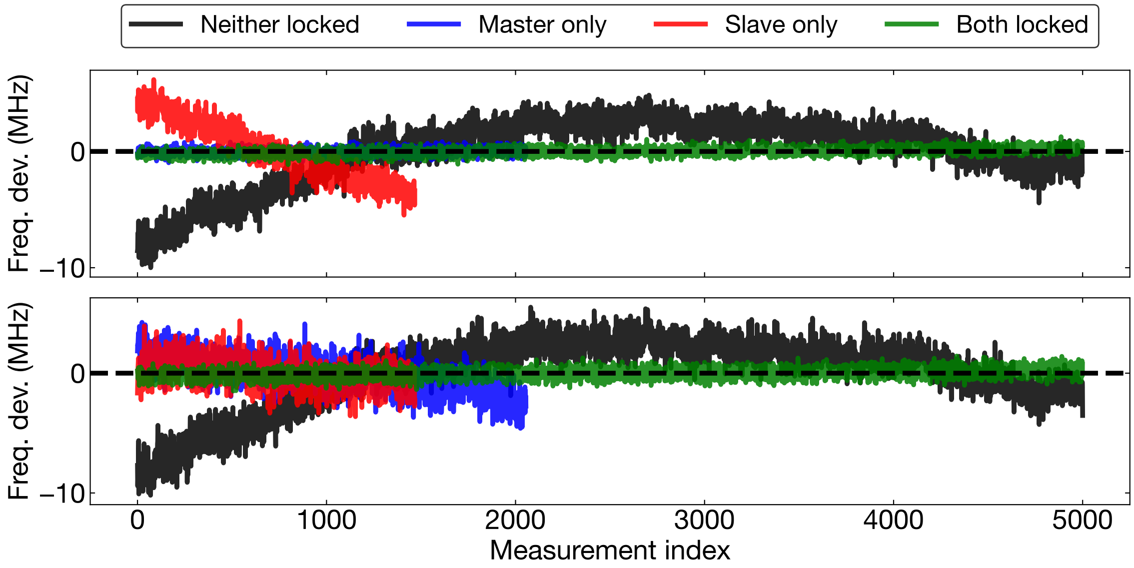

Figure 4.21: Frequency deviation from the mean as a function of measurement index for the master laser (top) and slave laser (bottom), recorded under four locking conditions. Each trace represents a continuous wavemeter acquisition. The free-running lasers (grey) exhibit slow drifts of order ±5 MHz, while the fully locked configuration (green) suppresses fluctuations to below ±1 MHz for both lasers.

In terms of laser power, the Vescent lasers each provide roughly 200 mW. For all beams except the MOT cooling beam this is in principle sufficient. However, significant power is lost at each stage: double-pass AOMs achieve at most $\sim 64\%$ efficiency (the AOMs themselves have approximately 80% single-pass efficiency), and fiber coupling stages typically add another $\sim 50\%$ loss. To compensate, we have two fiber-coupled Toptica BoosTA amplifiers (TAs), each providing up to 1.5 W of fiber-coupled output power. Since the Vescent lasers are themselves fiber-coupled, the TAs can be placed after any of them depending on where the power is needed most. For now, we plan to place one of the TAs after the Master laser on breadboard 1 and the second TA after the slave laser 2 on breadboard 3.

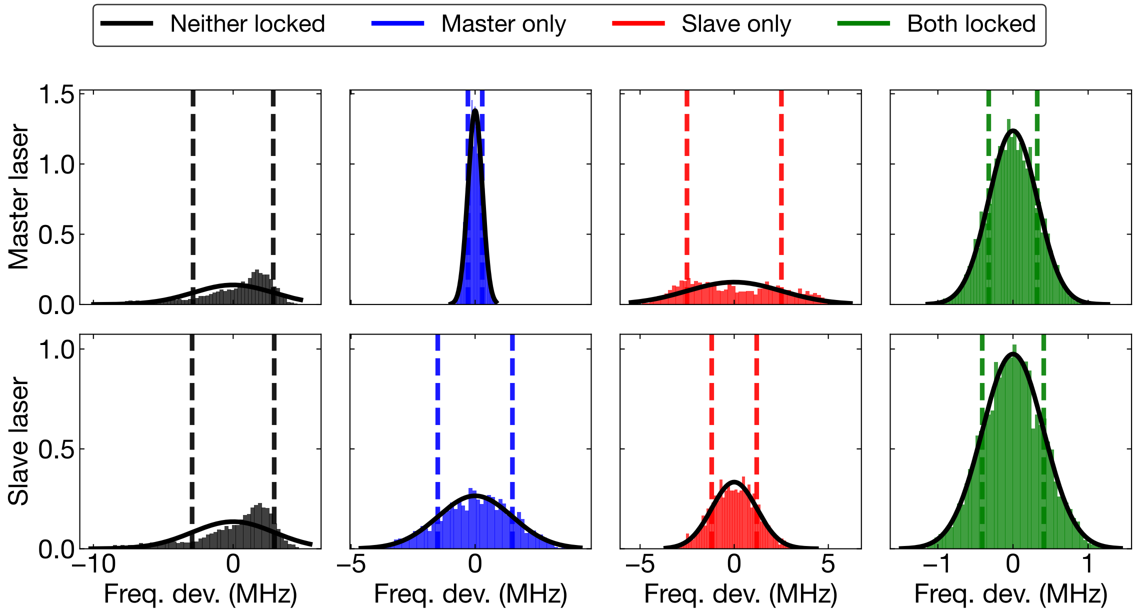

Figure 4.22: Frequency deviation distributions for the master laser (top row) and slave laser (bottom row) under four locking conditions. Histograms are normalized to unit area; black curves are Gaussian fits; dashed lines indicate $\pm\sigma$. The fitted standard deviations σ are quoted in each panel. Locking the master laser to the spectroscopy signal reduces its standard deviation from 2.87 MHz to 0.32 MHz, and the fully locked slave narrows from 2.94 MHz to 0.41 MHz.

4.2.2 AOM Distribution Breadboards

After locking the Vescent lasers and amplifying them, the wavelengths explained and required in table 4.1 for the main experimental table optics are prepared using so-called "distribution breadboards" shown schematically in figure 4.23, 4.24, 4.25 and 4.14. Photographs of these fully implemented distribution breadboards can be seen in figure F.4 in the appendix. As previously explained for the laser breadboards, the need to move between two labs in the middle of the thesis motivated building everything on modular breadboards that can easily be transported. Additionally, because the setup is entirely fiber-coupled, it is straightforward to upgrade or replace components. For example, one could replace the fiber coupling from breadboard 5 to breadboard 7 with a new Vescent laser to eventually achieve higher powers, or place the TA amplifier after breadboard 2 instead of breadboard 3. breadboard 5 prepares all the $F = 3 \rightarrow F'$ wavelengths and is connected to the master laser.

Slave laser 1 is connected to breadboard 6, which exclusively supplies the 2D MOT. The rationale for this separation is to maintain independent control over the lasers locking the 2D and 3D MOTs. slave laser 2 then supplies breadboard 7, which itself requires very little power. It relays most of its light to the MOT cooling breadboard (breadboard 8) and directs a small fraction to the beat breadboard (breadboard 4).

The double-pass AOMs are all arranged in a "cat's eye" configuration. In this setup, the AOM is placed at the front focal plane of a lens, and a retro-reflecting mirror is positioned at the back focal plane [71, 72]. When the RF driving frequency of the AOM is swept, the diffraction angle of the first-order beam changes. Because the AOM is located at the focal point, the lens converts these changing angles into parallel spatial translations. Since the mirror is placed exactly at the opposite focal plane, it retro-reflects the beam perfectly back through the lens, converting the parallel translation back into the exact same angle required to re-enter the AOM. This ensures that the outgoing, double-passed beam remains completely stationary (i.e. it does not walk spatially) even as its frequency is actively swept. This stationary beam is critical for maintaining a constant coupling efficiency into single-mode fibers downstream. Additionally, the lens focuses the beam, making it significantly easier to spatially filter and block the undiffracted zeroth-order beam.

In summary, this chapter presents the complete optical and laser infrastructure required to operate the cesium MOT. Starting from the requirements derived in section 3, the main table optics are designed to deliver six independent, flat-top cooling beams of approximately 1.3 in in diameter, two spatially shaped repumper beams via DMD with verified diffraction-limited edge sharpness, and a multiplexed absorption imaging system. The laser locking scheme, based on Vescent DBR lasers locked via saturated absorption spectroscopy and offset beat-note locking is fully implemented and characterized. The measured frequency stability of the fully locked system ($\sigma = 0.32\,\mathrm{MHz}$ for the master laser and $\sigma = 0.41\,\mathrm{MHz}$ for the slave laser) is well below the cesium D_2 natural linewidth of $\Gamma/2\pi \approx 5.2\,\mathrm{MHz}$, demonstrating that the locking scheme is sufficient for stable MOT operation. The modular, fiber-coupled distribution breadboards are likewise build and operational.

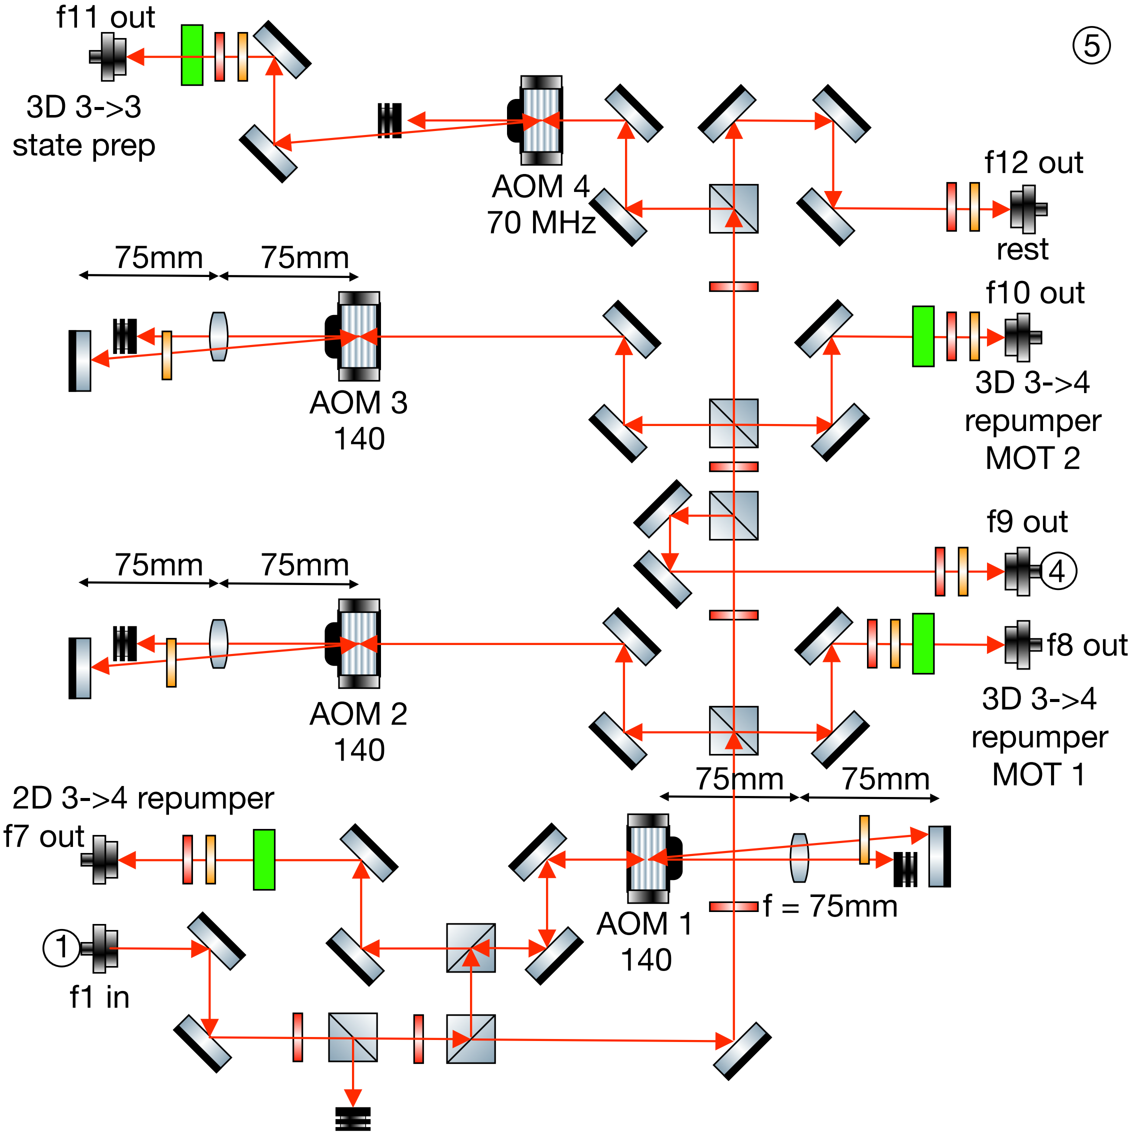

Figure 4.23: Master laser distribution breadboard (breadboard 5). Receives input from breadboard 1 and distributes light to the beat breadboard (breadboard 4) and further downstream. Legend can be found in figure 4.5.

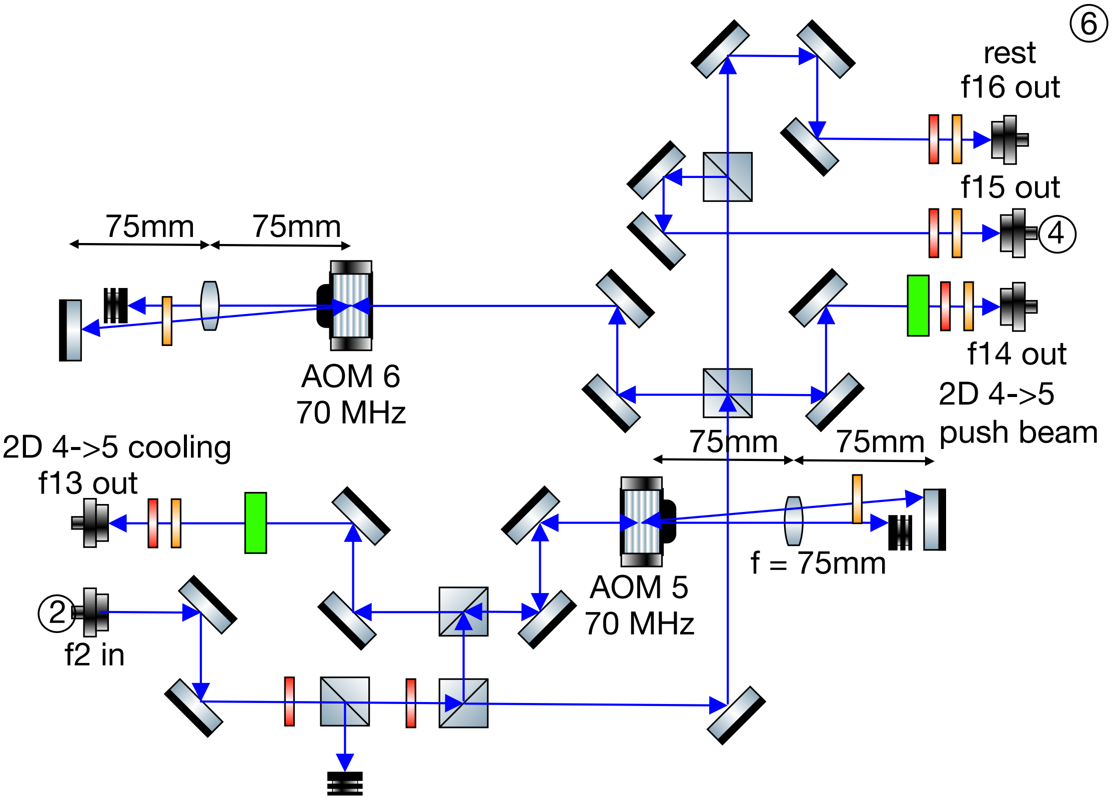

Figure 4.24: 2D MOT distribution breadboard (breadboard 6). Receives input from slave laser 1 (breadboard 2) and distributes light to the 2D MOT setup. Legend can be found in figure 4.5.

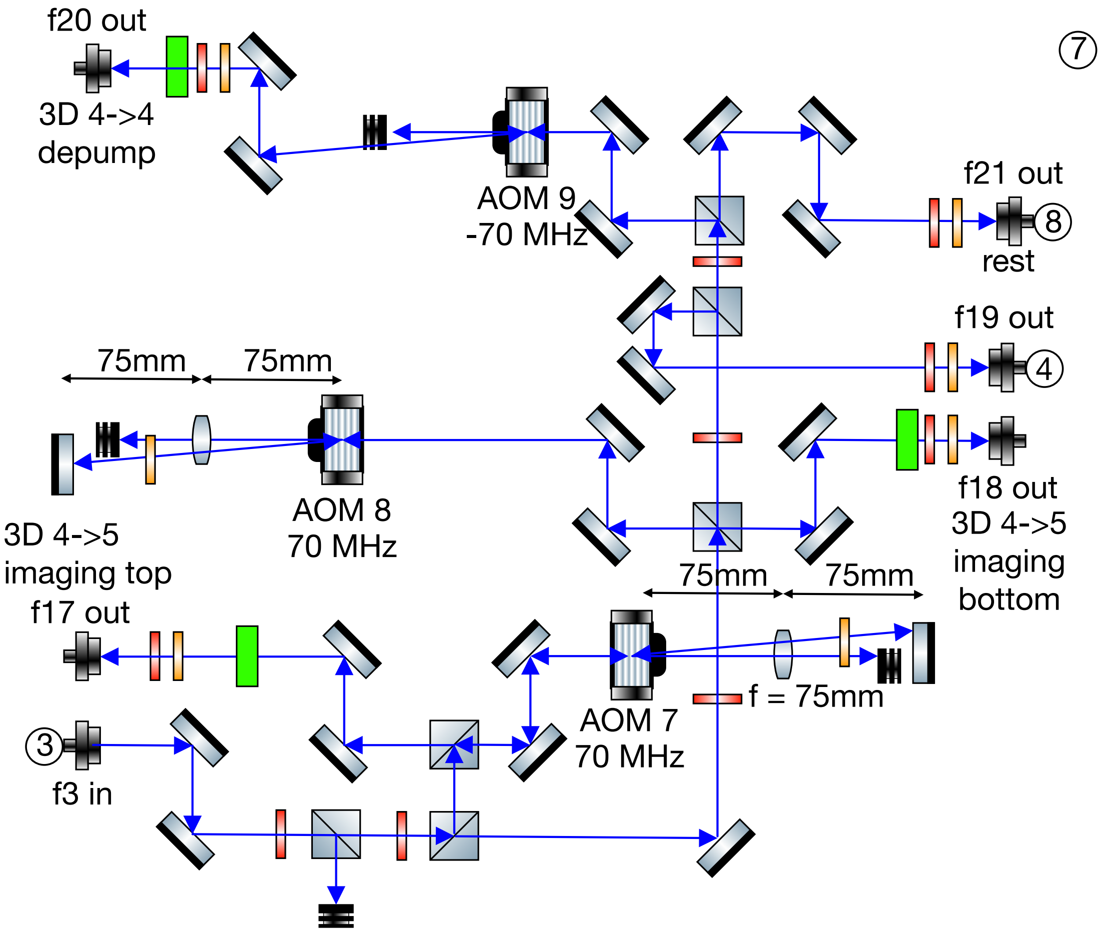

Figure 4.25: Imaging and depump distribution breadboard (breadboard 7). Distributes light for absorption imaging and the depump beam paths. Legend can be found in figure 4.5.

Chapter 5

Conclusion and Outlook

This thesis presents the theoretical analysis, design, and initial experimental implementation of a MOT optimized for the continuous reloading of a large-scale neutral atom quantum computer. In this work, we establish the fundamental requirements for continuous reloading in our experimental geometry and develop novel solutions to address the associated challenges. Section 3 derives the key MOT design considerations from first principles, analyzing loss and heating mechanisms that are particularly critical for ratchet-type continuous transport. Most notably, we identify the position-dependent damping reversal mechanism that occurs when atoms are held off-center in a MOT during transport. To address this, we propose a DMD-based shaped repumper as a dual-purpose solution to both mitigate this heating and enable in situ optical pumping to the absolute ground state. We further analyze spin-changing collision losses in ^{133}Cs and propose an optical pumping scheme from ^{133}Cs BEC literature to prepare atoms in the collision-safe $|F = 3, m_F = +3\rangle$ state within the repumper dark tube.

On the experimental front, the laser locking system (breadboards 1-4) has been fully assembled and characterized, employing three Vescent D2 DBR lasers: one master laser stabilized via saturated absorption spectroscopy and two offset-locked slave lasers for flexible frequency control. Frequency stability measurements confirmed performance well below the ^{133}Cs D_2 natural linewidth, with measured standard deviations of $\sigma = 0.32\,\text{MHz}$ (master) and $0.41\,\text{MHz}$ (slave lasers). The distribution breadboards (breadboards 5-8) have likewise been fully implemented, providing independently tunable frequencies for all required atomic transitions via combinations of double-pass and single-pass acousto-optic modulators in cat's-eye configurations.

Furthermore, a comprehensive SolidWorks-based 3D optical design of the main table optics is presented, including all MOT cooling beams, the DMD-shaped repumper relay imaging paths, state preparation and depumper beams, and absorption imaging systems. This design includes detailed specifications for beam sizes, lens selections, polarization control, and spatial routing across a three-layer breadboard architecture, ensuring compatibility with the spatial constraints imposed by tweezer and lattice optics.

Due to unforeseen delays in laboratory construction and iterative revisions to the vacuum chamber design, several critical experimental milestones could not be completed within the timeframe of this thesis. As a result, the physical assembly and alignment of the main table optics could not be performed. Additionally, the vacuum system has not yet undergone its bakeout procedure. Consequently, while the theoretical framework and optical design are complete, and the experimentally validated subsystems (laser locking and distribution) are operational, no atomic cloud has been loaded into the 3D MOT as of the conclusion of this thesis. With the arrival and installation of the main table breadboards, the next immediate steps will be the assembly of the MOT cooling and repumper optical paths, followed by the alignment of the shaped repumper relay imaging system and verification of the DMD spatial profiles at the location of the atoms. Once the vacuum chamber is baked, we anticipate successfully loading our first 3D MOT.

From that point forward, several exciting questions will become experimentally accessible. First, we will investigate whether, and to what extent, the digitally shapable repumper can be used to optimize the MOT loading rate and temperature by empirically varying the dark tube geometry and repumper intensity profile. Second, we will characterize the performance of ratchet-type continuous MOT extraction by measuring the atom flux into the optical lattice, the transport efficiency through the 45-degree handshake geometry, and the extent to which continuous MOT operation can sustain high-fidelity atom delivery to the tweezer loading region without introducing decoherence to pre-existing qubit arrays. Successful demonstration of these capabilities will represent a significant step toward realizing fault-tolerant quantum computation on neutral atom platforms, enabling the long-circuit-depth algorithms required for practical quantum advantage.

Appendices

Appendix A

Acktar Magic Black Shallow Angle Data

As discussed in section 2.1, our initial vacuum design incorporated a black-coated tube with the intention of eliminating internal reflections such that only the geometric shadow of the light before the tube remains as illustrated in figure 2.4. However, experimental observations revealed that this assumption did not hold true. The first indication of reflection at shallow grazing angles is shown in figure A.1, where illuminating the tube entrance with a flashlight produces clear internal reflections visible at the output. This prompted us to investigate the efficacy of the tube coating at shallow angles, especially since the manufacturer (Acktar) only provides reflectivity data for angles of incidence above 10°. To address this, we constructed a custom reflectivity test setup using 852 nm laser light, depicted in figure A.2. An 852 nm outcoupler was securely clamped to the optical table. A sample of the coating was mounted on a rotation stage and aligned perfectly parallel to the outcoupled beam. In this configuration, a fraction of the beam grazes the coating while the remainder is recorded by a camera positioned behind the sample. By rotating the stage, we systematically increased the angle of incidence between the coating and the beam. A limitation of this setup was the absence of a translation mount for the camera due to time constraints, which instead had to be displaced manually. Consequently, we could not extract precise, spatially-resolved angular data. Nevertheless, we successfully recorded camera images for various angles to perform both a qualitative analysis of the beam profile (determining how specular or diffuse the reflection is) and a quantitative estimation of the total reflected power by integrating the intensity on the camera sensor. Figure A.3 presents the raw camera images acquired at different angles. One can already qualitative see that while at low angles (surprisingly strong) specular reflection dominates, at higher angles one can observe very clear diffuse reflection.

For a more quantitative analysis, we applied a background subtraction and thresholding routine to isolate the reflected peaks, as shown in figure A.4. By integrating these pixel intensities, we calculated the relative reflection I_{n°/I_{0°. To validate the accuracy of using the camera as a power sensor, we compared these integrated intensities against independent power meter measurements. As shown in figure A.5, the power meter data aligns relatively well with the integrated camera intensities. Most measurements were performed twice, once with low exposure and once with high exposure, which accounts for the dual camera data points at each angle. While this confirms the validity of our camera-based integration method, the critical parameter of interest is the power contained strictly within the specular reflection peak. By isolating and integrating only the specular peak (excluding the primary transmitted beam), we determined the fraction of power reflected at each angle. The results, plotted in figure A.6, indicate significant specular reflection (up to 25%) at small grazing angles. As the angle increases, the reflection drops considerably, reaching approximately 10% at 10°. It is worth noting that our 10° data point exhibits characteristics of an outlier as at this angle, the exposure time had to be increased drastically (up to 35.65 ms), yet the measured intensity remained low, resulting in a low signal-to-noise ratio for the integrated data.

While the aforementioned measurements conclude that, especially at small angles of incidence, the Acktar Magic Black coating exhibits significant specular (as well as diffuse) reflection, they do not fully address the question of how much intensity will actually arrive at the atoms in the tweezer. Even with precise reflectivity data, accurately simulating the behavior of the MOT light within the tube proves to be highly challenging. Consequently, we sought to determine if this could be measured directly. A primary difficulty arises from the ambiguity of how exactly the MOT scatters light towards the tube, the number of scattering atoms, and how to appropriately model this experimentally. Ultimately, we chose to model the MOT using a divergent fiber tip. By placing a collimation lens on a rotatable mount, we were able to initially align the collimated beam to the entrance of the tube (a significantly easier task than aligning a divergent beam). After alignment, the lens could be removed to switch to a divergent beam, allowing us to observe the resulting intensity profile on a camera placed behind the tube at a distance corresponding to the planned tweezer region in the original vacuum design.

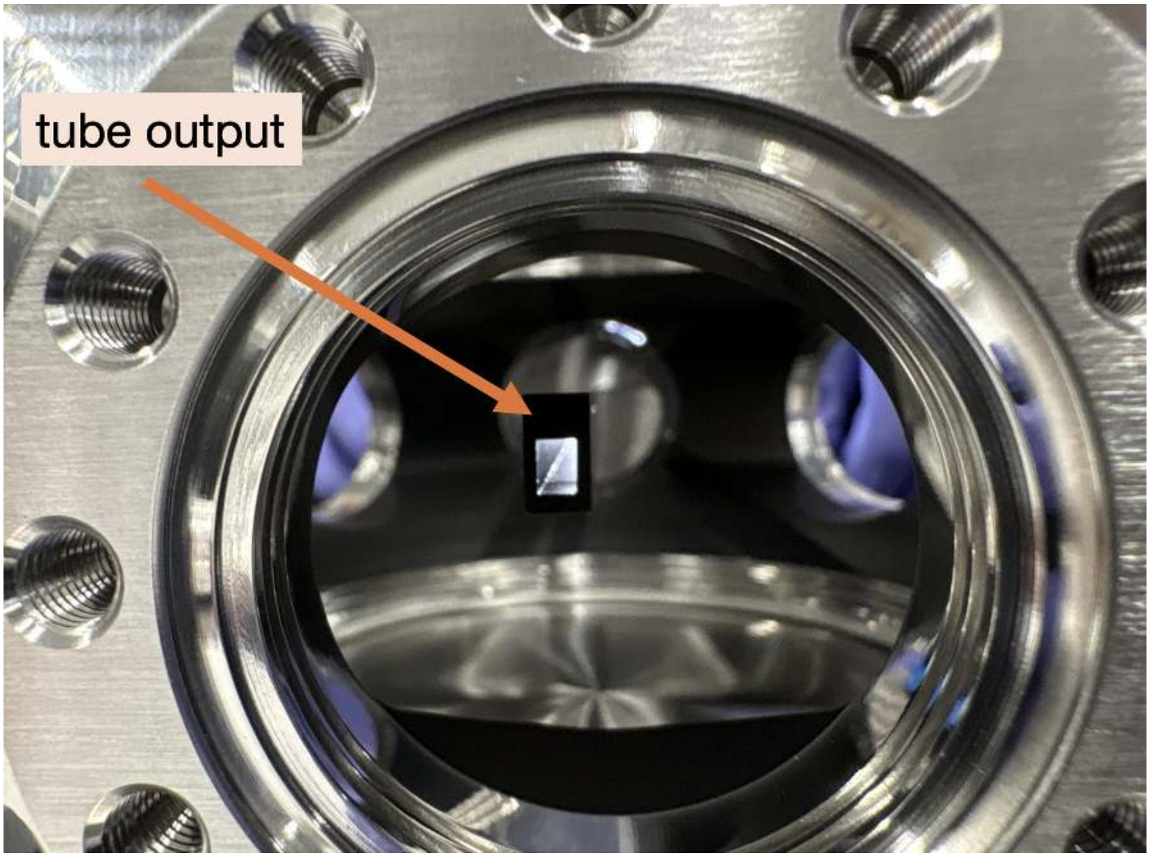

Figure A.1: Visual observation of internal light reflections when illuminating the entrance of the black-coated tube with a flashlight.

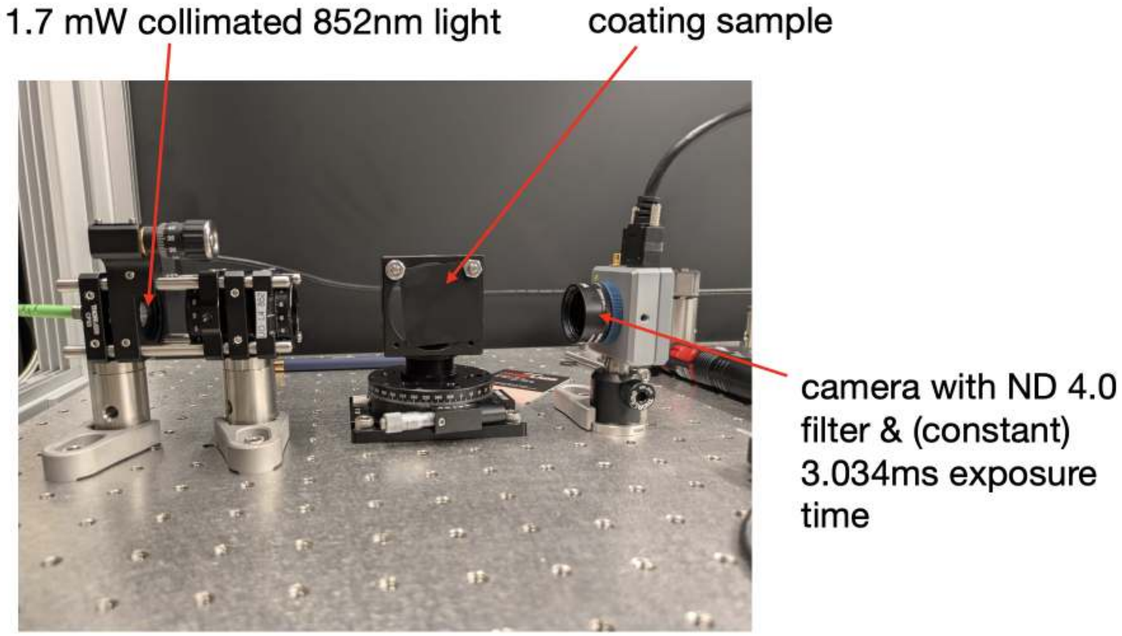

(a) Setup overview

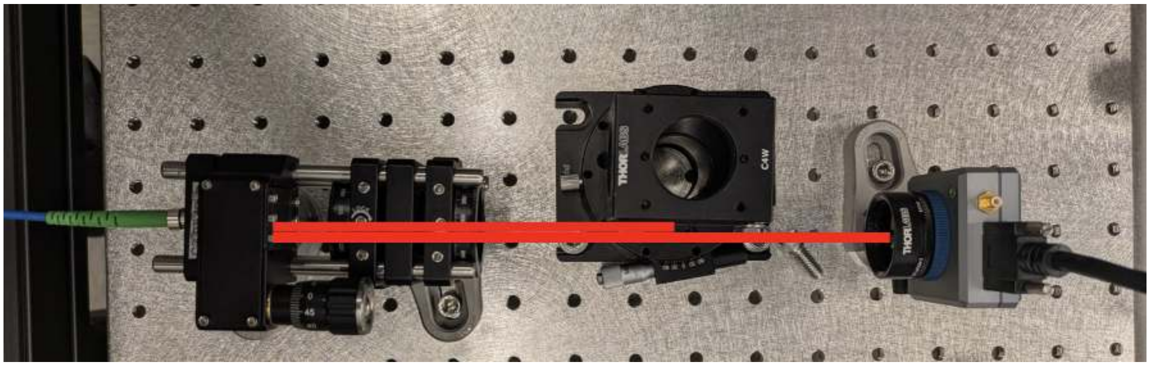

(b) Close-up view

Figure A.2: Photographs of the custom test setup used to measure the reflectivity of the tube coating at shallow angles using an 852 nm laser source and a camera.

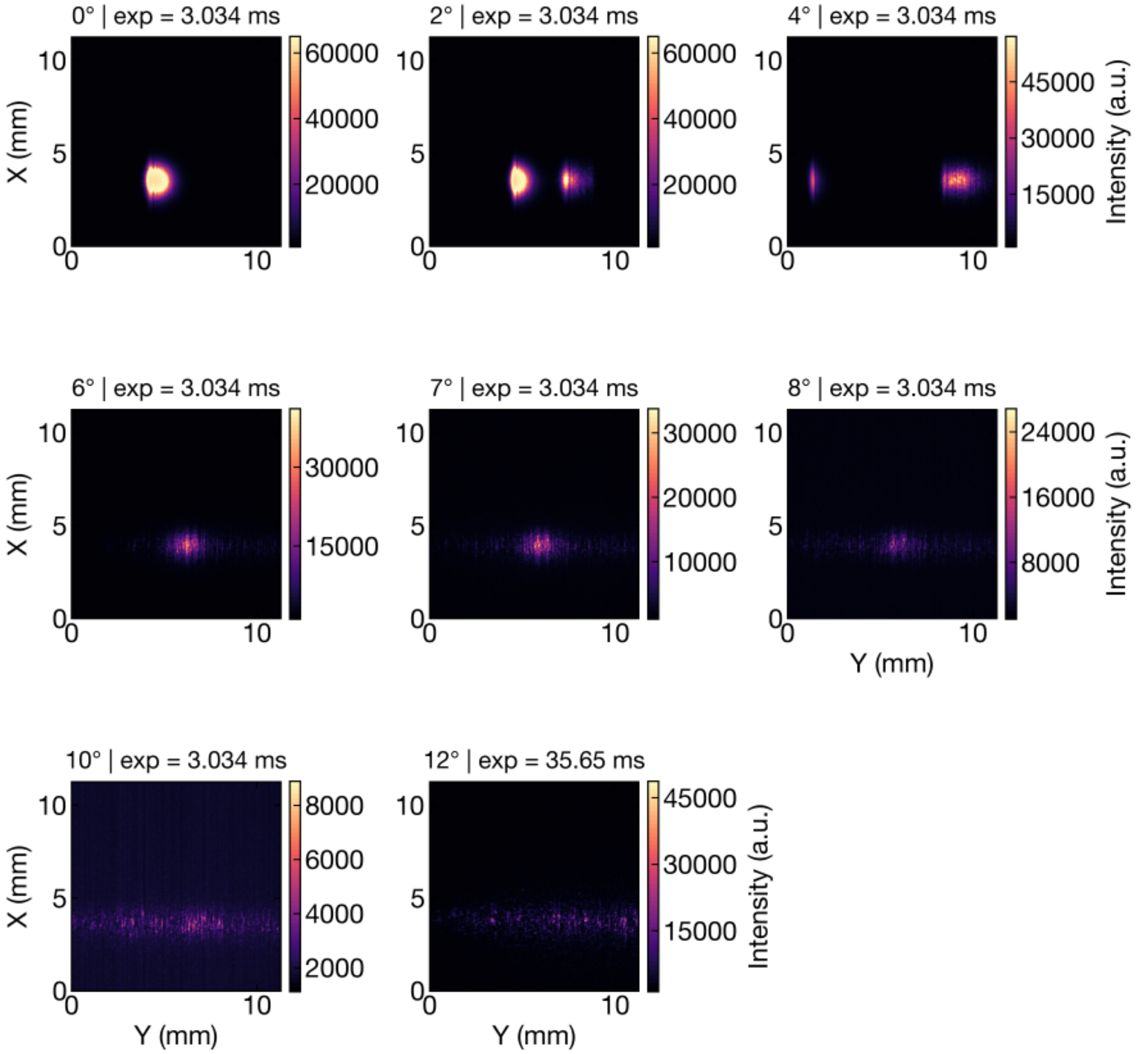

Figure A.3: Grid of raw camera images capturing the reflected and transmitted 852 nm light at various shallow grazing angles.

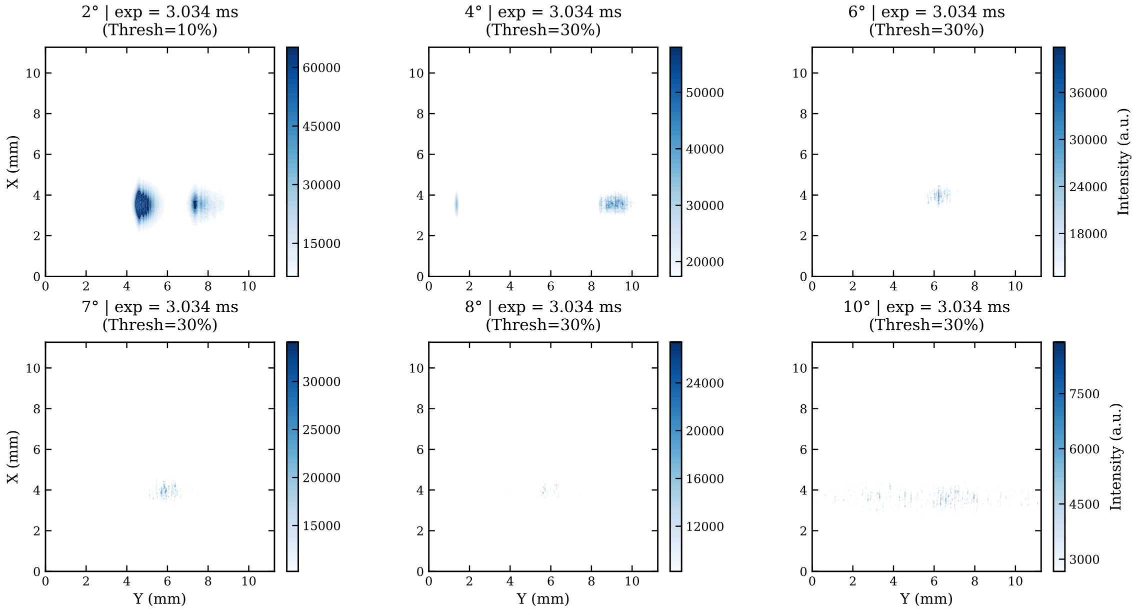

Figure A.4: Processed camera data displaying isolated peaks after applying background subtraction and thresholding algorithms.

A central question of interest was whether the stray light problem could be mitigated by simply displacing the MOT off-center relative to the tube axis, and if so, what magnitude of displacement would be required. Figure A.7 illustrates the experimental setup. The fiber was mounted on a translation stage to allow for controlled lateral displacement. A technical complication with this setup is that as we translate the fiber off-center, the tip must be rotated to maintain its pointing towards the tube. Otherwise, the transmitted signal is lost. Because the fiber tip does not coincide perfectly with the axis of rotation, viewing the fiber tip as the MOT implies that the effective off-center displacement is smaller than the translation stage reading suggests. This geometric discrepancy is illustrated in figure A.8.

Furthermore, moving the MOT off-center spans several orders of magnitude in intensity behind the tube, making it impossible to capture the full range with a single set of camera settings. Therefore, we recorded the neutral density (ND) filter value, camera exposure time in milliseconds (t_{exposure}), and the laser power exiting the fiber (P_{fiber}) for each off-center distance. The resulting camera data is shown in figure A.9.

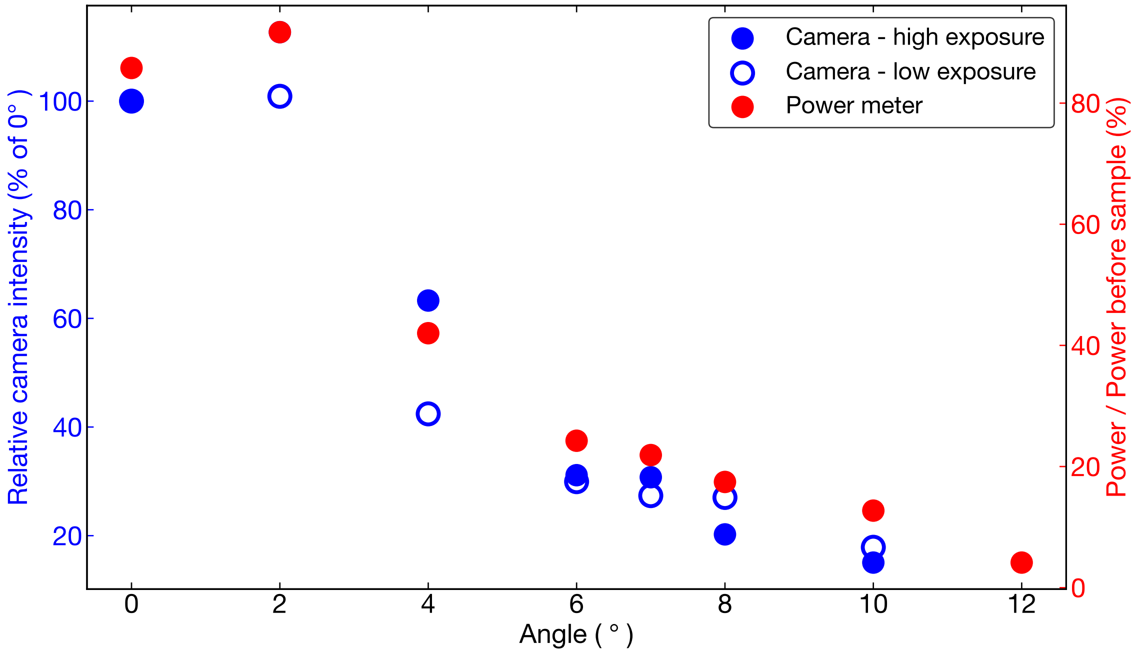

Figure A.5: Comparison of the integrated relative intensity obtained via the camera with independent power meter measurements across varying angles, demonstrating the reliability of the camera as an intensity sensor.

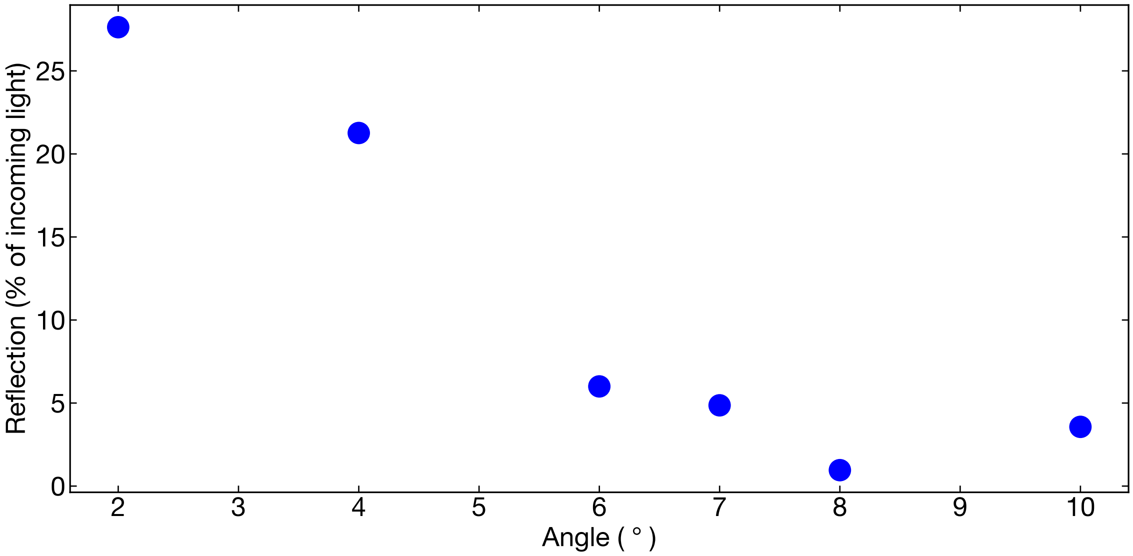

Figure A.6: Measured power reflected specifically within the specular peak as a function of the angle of incidence. Note the relatively high reflection at grazing angles that gradually drops as the angle increases.

By applying the scaling formula:

$$I_{\text{scaled}} = \left(\frac{1.5\,\text{mW}}{P_{\text{fiber}}}\right) \cdot I_{\text{measured}} \cdot \frac{10^{\text{ND}}}{t_{\text{exposure}}} \tag{A.1}$$

and subtracting the independently measured dark background, we obtain the normalized data presented in figure A.10.

Several conclusions can be drawn from these measurements. First, at zero off-center displacement, the resulting diffraction pattern in figure A.9 clearly shows intensity distributions extending far beyond the geometric shadow. The intensity outside this shadow is substantial, indicating that the original shielding concept is insufficient. Second, even though off-centering the MOT suppresses the scattered light, the residual intensity remains significant. Relying on this tube geometry in the final setup poses too high a risk of stray light contamination. Additionally, achieving a meaningful reduction in scattered light requires a substantial off-center displacement. This introduces severe practical drawbacks, such as clipping the large MOT beams and necessitating fundamental changes to the MOT design proposed in this thesis. Finally, practical challenges arose during the assembly of the vacuum components. The tube is held in place by a groove-grabber installed in one of the bucket windows of

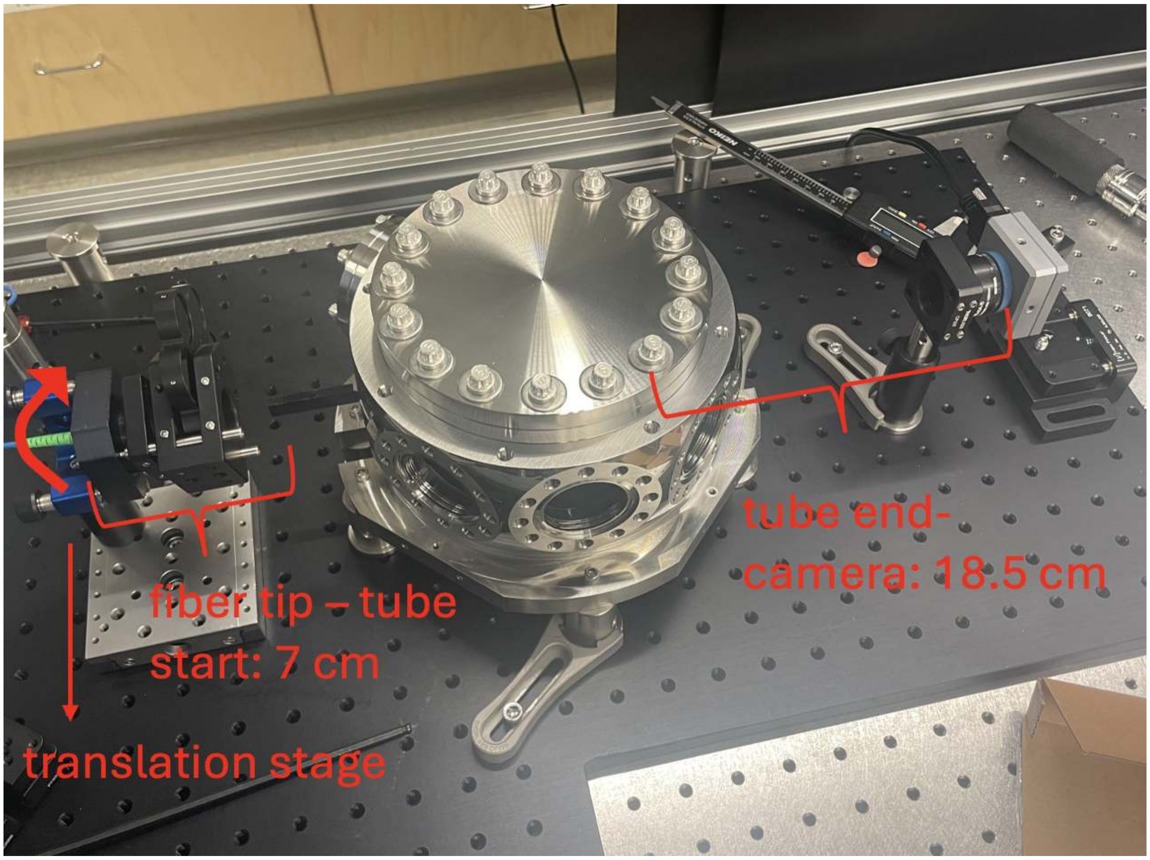

Figure A.7: Experimental setup used to model the MOT scattering into the tube using a divergent fiber tip.

the MOT chamber. In practice, this mounting strategy proved to be highly finicky. Not only was inserting the tube into the groove-grabber cumbersome, but the tube also ended up with a significant angular deviation once secured. Precisely aligning the tube's angle and position using this mechanism proved to be exceptionally difficult, revealing substantial mechanical flaws in this vacuum design over and above the stray light issues.

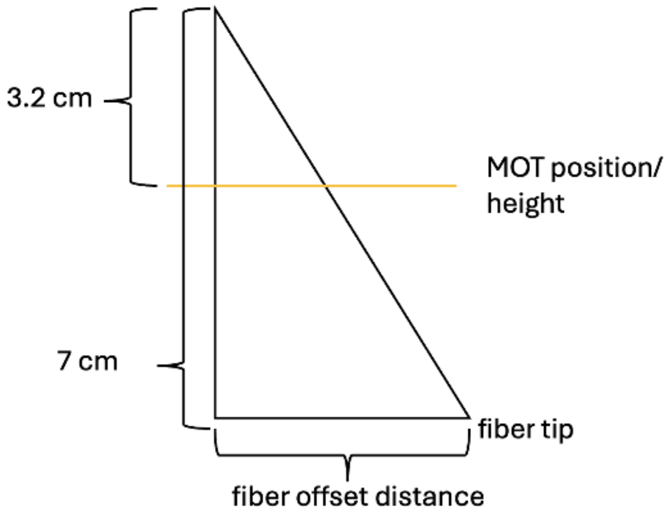

Figure A.8: Geometric illustration of the fiber tip displacement. The effective off-center distance is smaller than the translation stage indicates due to the offset between the fiber tip and the rotation axis.

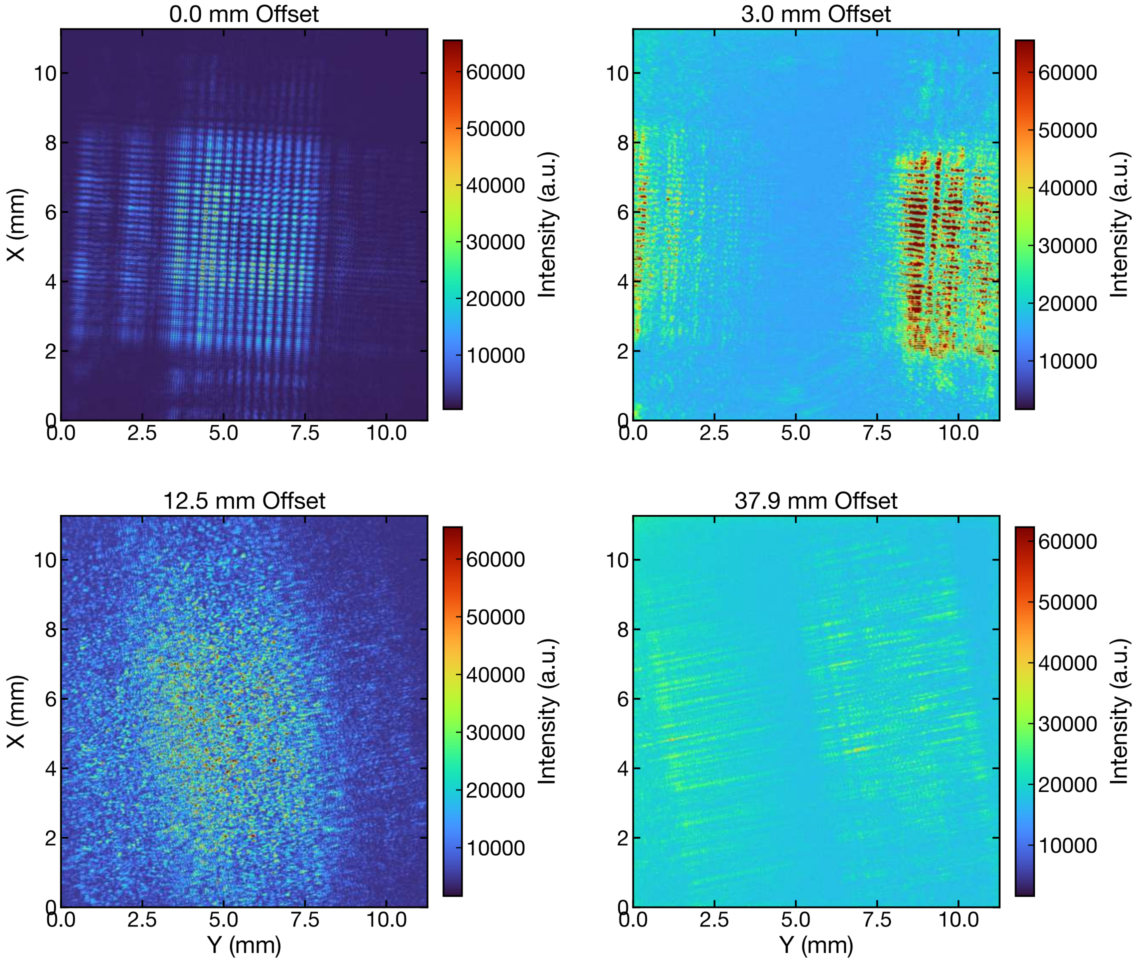

Figure A.9: Camera data recorded at various off-center distances, illustrating the intensity profile of the light behind the tube.

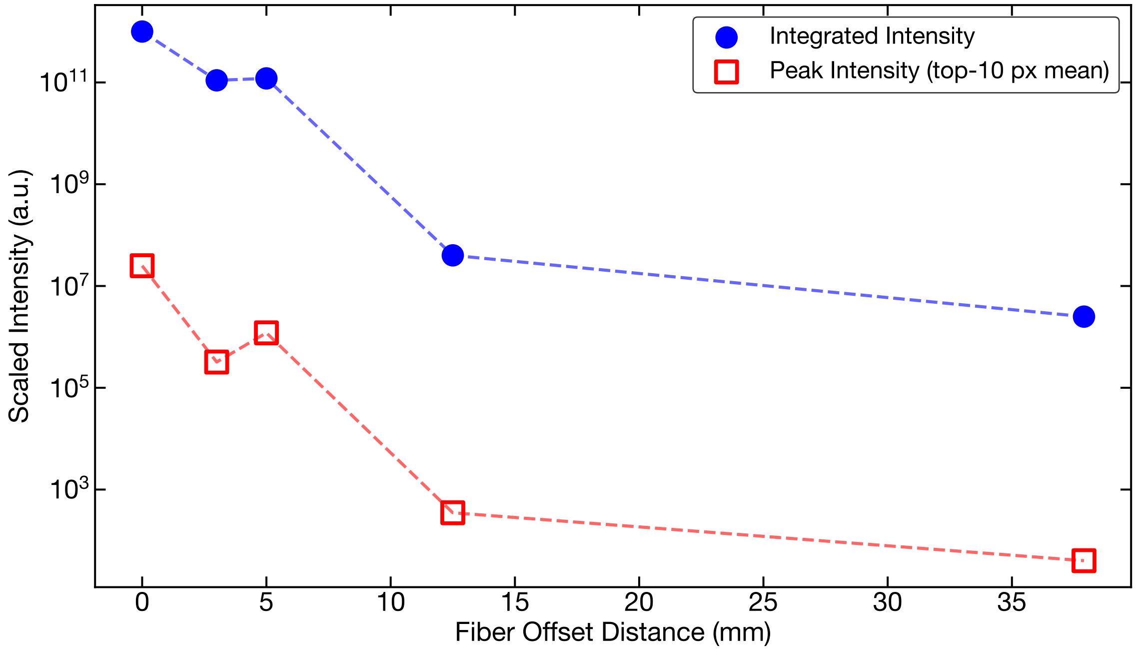

Figure A.10: Normalized and background-subtracted integrated intensity as a function of the off-center distance.

Appendix B

Separating MOT and Imaging Light

The question at hand is if and how we can separate the imaging beam from the other beams, to only have the imaging beam on a camera after it passed through the atom chamber. This in the end boils down to the question, if the two counterpropagating cooling beams are separable or not (see figure B.5 a). The remainder of this section addresses this question in detail. In order to answer that question we have to carefully analyze the polarization requirements at the atoms of all involved beams in the beampath and then see how you can separate them. As shown in figure 4.4 and table 4.1, we do absorption imaging from along two axes. On both axes the setup is the same: from one port (p7 and p4) the imaging light is copropagating with the cooling light, going through the chamber and then exiting on the other side's port (p6 and p2). From the opposite site through p6 and p2, the repumper is copropagating with the cooling beam. Both beam paths involve exactly the same beams, illustrated in figure B.1.

So the precise question is "can we separate the imaging beam from the counterpropagating cooling beams and the repump beam?". To answer this question, we have to understand the polarization requirements of the beams at the atoms. However, before doing that, it is useful to introduce the notation of polarization used in the following chapter, because this can be a bit confusing. Polarization describes the orientation of the electric field vector $\vec{E}$ of an electromagnetic wave. For a plane wave propagating in the z-direction, the electric field oscillates in the xy-plane:

$$\vec{E}(z,t) = \frac{1}{2}\left[\left(E_x e^{i\phi_x}\,\vec{e}_x + E_y e^{i\phi_y}\,\vec{e}_y\right) e^{i(\omega t - kz)} + \text{c.c.}\right] \tag{B.1}$$

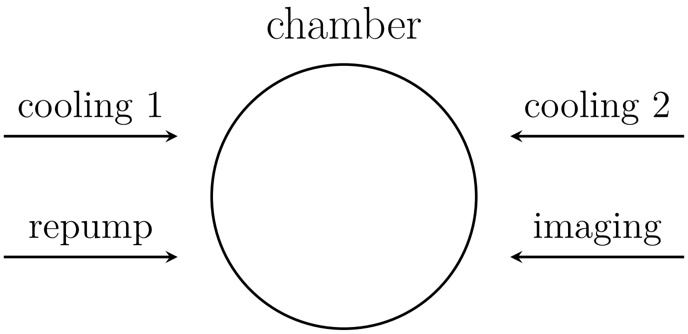

Figure B.1: The beam configuration along the two imaging axis are the same and can be represented like this

where E_x and E_y are real amplitudes and ϕ_x, ϕ_y are phases. The polarization state depends on the relative amplitudes and phase difference $\Delta\phi = \phi_x - \phi_y$:

- **Linear**: $\Delta\phi = 0$ or π — field oscillates along a fixed direction
- **Circular**: $E_x = E_y$ and $\Delta\phi = \pm\pi/2$ — field vector rotates in a circle
- **Elliptical**: General case — field vector traces an ellipse

The description of the rotation direction depends critically on the coordinate system you are using to observe the light. When talking about polarization of light in respect to atoms there are two important coordinate systems, as illustrated in figure B.2:

B.1 Light Coordinate System (Optical Convention):

This is the right-handed coordinate system defined by the light propagation direction, where the z-axis points along the direction of light propagation (i.e., the direction of the wavevector $\vec{k}$). As shown in figure B.2, in this convention:

- **Right-handed circular polarization (RCP or RHP)**: The electric field vector rotates counterclockwise when viewed looking *into* the direction of propagation (i.e., looking in the $-z$ direction with light coming toward you). Using the right-hand rule: thumb points in the propagation direction ($+z$), fingers curl in the direction of $\vec{E}$ field rotation.
- **Left-handed circular polarization (LCP or LHP)**: The electric field vector rotates clockwise when viewed looking into the direction of propagation.

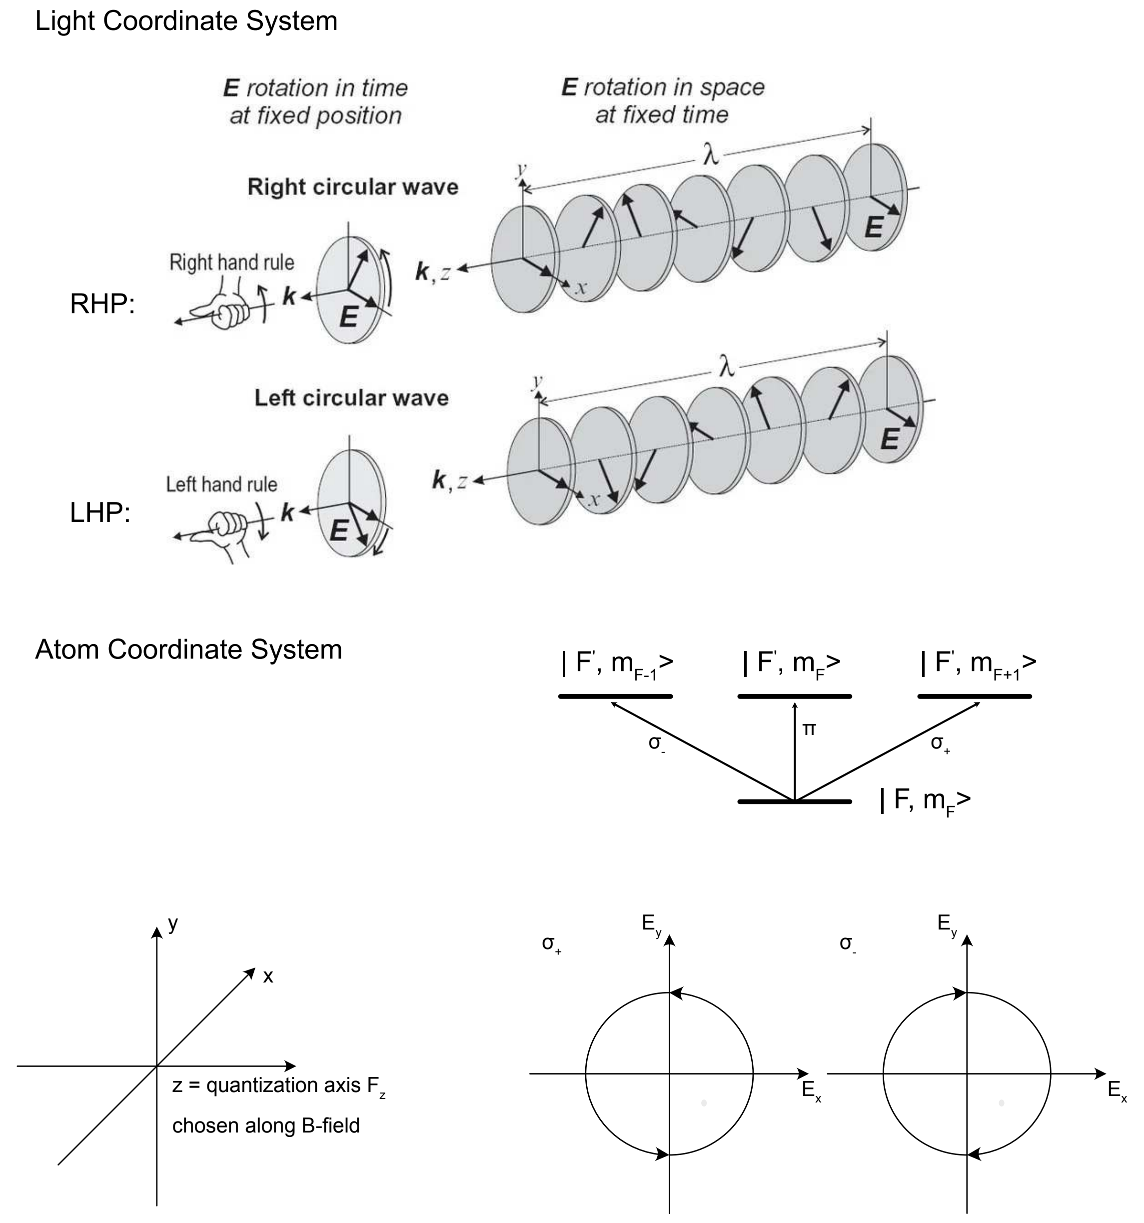

Figure B.2: Comparison of the two coordinate systems used for describing polarization. ight coordinate system: Shows right-handed polarization (RHP) and left-handed polarization (LHP) defined by pointing your thumb in the direction of the light propagation. Atom coordinate system (quantization axis coordinate system): As explained and derived in the text, σ^- polarization is defined as the polarization of the light that drives the $\Delta m_F = -1$ transition, and it can be shown that in the right-handed coordinate system defined by the quantization axis, this is equal to the electric field rotating clockwise in the x-y plane. Equally, σ^+ polarization is defined as the polarization that drives the $\Delta m_F = +1$ transition, corresponding to the electric field rotating counter-clockwise in the x-y plane. The upper figure is taken about the light coordinate system is taken from [73].

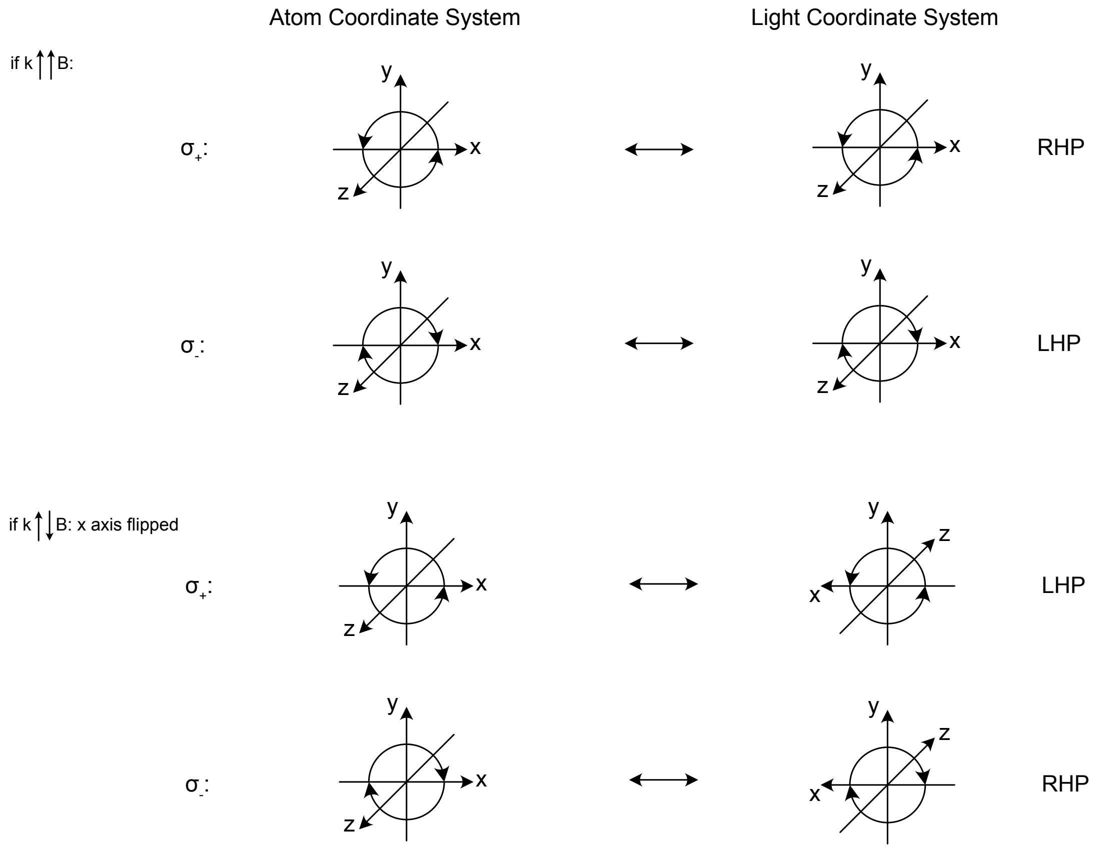

Figure B.3: Transformation between coordinate systems: Demonstrates how σ^+ and σ^- polarizations in the atom frame correspond to RHP and LHP in the light frame when light propagates along $+z$, but this mapping reverses when the light-coordinate system is flipped.

Mathematically, in the light coordinate system:

$$\text{RCP:} \quad \vec{E}_{\text{RCP}} = \frac{1}{2}\left[E_0(\vec{e}_x + i\vec{e}_y)e^{i(kz-\omega t)} + \text{c.c.}\right] \tag{B.2}$$

$$\text{LCP:} \quad \vec{E}_{\text{LCP}} = \frac{1}{2}\left[E_0(\vec{e}_x - i\vec{e}_y)e^{i(kz-\omega t)} + \text{c.c.}\right] \tag{B.3}$$

To verify the rotation direction, take the real part at fixed position $z = 0$ from equation (B.2):

$$\vec{E}_{\text{RCP}}(t) = \frac{1}{2}\left[E_0(\vec{e}_x + i\vec{e}_y)e^{-i\omega t} + \text{c.c.}\right] = E_0(\vec{e}_x \cos(\omega t) + \vec{e}_y \sin(\omega t)) \tag{B.4}$$

As time increases, the field vector traces $+x \to +y \to -x \to -y$, which is counterclockwise (CCW) when looking into the propagation direction; similarly, LCP with the $-i$ term produces clockwise (CW) rotation.

B.2 Atom Coordinate System (Quantization Axis Convention)

This coordinate system is defined by the atomic quantization axis, which is typically chosen to align with the magnetic field direction. This is exactly also the coordinate system introduced in 2.2.1. The quantization axis defines the basis for the Zeeman sublevels (m_F quantum numbers) of the atomic states, as illustrated in figure B.2. In this frame, the z-axis is chosen along the direction of the magnetic field $\vec{B}$ (or more precisely, the direction of the chosen angular momentum quantization, typically $\hat{z} \parallel \vec{B}$).

In the atomic coordinate system, we define:

- σ^+ polarization: Drives transitions with $\Delta m_F = +1$ (increases the magnetic quantum number)
- σ^- polarization: Drives transitions with $\Delta m_F = -1$ (decreases the magnetic quantum number)
- π polarization: Drives transitions with $\Delta m_F = 0$ (no change in magnetic quantum number)

Now how do we know which polarization direction of the light (in the coordinate system defined by the right handed coordinate system of the quantization axis) drives which transition? This can be derived using the interaction Hamiltonian:

The interaction Hamiltonian for an atom in an electric field is:

$$\hat{H}_{\text{int}} = -\hat{\vec{d}} \cdot \vec{E} \tag{B.5}$$

where $\hat{\vec{d}}$ is the electric dipole operator and $\vec{E}$ is the electric field. Defining:

$$\begin{aligned}
\hat{d}_0 &= \hat{d}_z && \text{(B.6)}\\
\hat{d}_{+1} &= -\frac{1}{\sqrt{2}}(\hat{d}_x + i\hat{d}_y) && \text{(B.7)}\\
\hat{d}_{-1} &= +\frac{1}{\sqrt{2}}(\hat{d}_x - i\hat{d}_y) && \text{(B.8)}
\end{aligned}$$

It can be shown that $\hat{d}_{\pm 1}$ are essentially ladder operators for angular momentum:

$$\begin{aligned}
\hat{d}_0|F, m_F\rangle &\propto |F', m_F\rangle \quad (\Delta m_F = 0) && \text{(B.9)}\\
\hat{d}_{+1}|F, m_F\rangle &\propto |F', m_F + 1\rangle \quad (\Delta m_F = +1) && \text{(B.10)}\\
\hat{d}_{-1}|F, m_F\rangle &\propto |F', m_F - 1\rangle \quad (\Delta m_F = -1) && \text{(B.11)}
\end{aligned}$$

Now, the electric field can also be decomposed in the same spherical basis:

$$\begin{aligned}
\vec{E}(t) &= \frac{1}{2}\left[(E_x\vec{e}_x + E_y\vec{e}_y + E_z\vec{e}_z)\, e^{i(kz-\omega t)} + \text{c.c.}\right]\\
&= \frac{1}{2}\left[E_z\vec{e}_z + \frac{1}{2}(E_x + iE_y)(\vec{e}_x - i\vec{e}_y) + \frac{1}{2}(E_x - iE_y)(\vec{e}_x + i\vec{e}_y) + \text{c.c.}\right] && \text{(B.12)}\\
&= \frac{1}{2}\left[E_0\vec{e}_0 - E_{+1}\vec{e}_{-1} - E_{-1}\vec{e}_{+1} + \text{c.c.}\right],
\end{aligned}$$

with

$$\vec{e}_0 = \vec{e}_z, \qquad \vec{e}_{\pm 1} = \mp\frac{1}{\sqrt{2}}(\vec{e}_x \pm i\vec{e}_y), \tag{B.13}$$

and

$$E_0 = E_z, \qquad E_{-1} = \frac{1}{\sqrt{2}}(E_x - iE_y), \qquad E_{+1} = -\frac{1}{\sqrt{2}}(E_x + iE_y). \tag{B.14}$$

The explicit time dependence $e^{-i\omega t}$ of the electric field has been factored out and will be omitted in the following. This is justified because the polarization of the light is fully determined by the relative amplitudes and relative phases of the field components multiplying the basis vectors. The common oscillatory factor $e^{-i\omega t}$ is the same for all polarization components and therefore does not affect which spherical components (E_0, $E_{\pm 1}$) are present. Since the atomic selection rules and the classification into π, σ^+, and σ^- polarizations depend only on these relative amplitudes and phases, the time dependence is not required for the discussion of polarization and will be suppressed for clarity.

Using this the interaction Hamiltonian can be decomposed into:

$$\hat{H}_{\text{int}} = -\hat{\vec{d}} \cdot \vec{E} = -(\hat{d}_x E_x + \hat{d}_y E_y + \hat{d}_z E_z) = -\hat{d}_0 E_0 - \hat{d}_{+1} E_{-1} - \hat{d}_{-1} E_{+1} \tag{B.15}$$

From equation (B.15), we see that:

- $E_{+1} \neq 0$ (with $E_0 = E_{-1} = 0$) couples through $\hat{d}_{-1}$, driving $\Delta m_F = -1$ transitions
- $E_{-1} \neq 0$ (with $E_0 = E_{+1} = 0$) couples through $\hat{d}_{+1}$, driving $\Delta m_F = +1$ transitions
- $E_0 \neq 0$ (with $E_{\pm 1} = 0$) couples through $\hat{d}_0$, driving $\Delta m_F = 0$ transitions

Therefore, we define the polarization states by which atomic transitions they drive:

$$\sigma^+ \text{ polarization} \equiv \vec{e}_{\sigma^+} = \vec{e}_{+1} \text{ component present} \rightarrow \Delta m_F = +1 \tag{B.16}$$

$$\sigma^- \text{ polarization} \equiv \vec{e}_{\sigma^-} = \vec{e}_{-1} \text{ component present} \rightarrow \Delta m_F = -1 \tag{B.17}$$

$$\pi \text{ polarization} \equiv E_0 = E_z \text{ component present} \rightarrow \Delta m_F = 0 \tag{B.18}$$

B.3 Transforming Between the Two Coordinate Systems

Now that we described these two coordinate systems, we have to understand how they can be translated into each other. As illustrated in figure B.3, the mapping between RHP/LHP and σ^+/σ^- depends on the relative orientation of the two coordinate systems.

Case 1: $\vec{k}$ parallel to $\vec{B}$

When light propagates along $+z$, both coordinate systems are aligned with the same z-axis. In this case:

- σ^+ polarization (atom frame) corresponds to **RHP** (light frame)
- σ^- polarization (atom frame) corresponds to **LHP** (light frame)

This can be verified mathematically: The σ^+ polarization is given by $\vec{E} \propto \vec{e}_{+1} = -\frac{1}{\sqrt{2}}(\hat{x}+i\hat{y})$, which corresponds to the definition of RCP in equation (B.2). Similarly, the σ^- polarization is given by $\vec{E} \propto \vec{e}_{-1} = \frac{1}{\sqrt{2}}(\hat{x} - i\hat{y})$, which corresponds to the definition of LCP in equation (B.3).

Case 2: $\vec{k}$ antiparallel to $\vec{B}$

When the coordinate system is flipped (either by reversing the propagation direction to $-z$ or by flipping the x-axis), the correspondence reverses. In this case, the z-axis of the coordinate system of the light is opposite to that of the atom, meaning that the x-axis is also flipped. Because of this:

- σ^+ polarization (atom frame) corresponds to **LHP** (light frame)
- σ^- polarization (atom frame) corresponds to **RHP** (light frame)

This is explicitly shown in figure B.3, where the diagrams demonstrate the coordinate transformation for both cases.

With these writing conventions (RHP and LHP are in the light coordinate system; σ^+ and σ^- are in the atom coordinate system, for which it is always true that σ^+ drives $\Delta m_F = +1$ and σ^- drives $\Delta m_F = -1$) and equipped with the transformation between the two coordinate systems, we can derive the polarization requirements for all the involved beams.

B.4 Polarization Requirements for the Cooling Beams

To understand the polarization requirements, let's consider a concrete example: an atom at position $z > 0$ in a linear magnetic field gradient with $B_z > 0$ along the z-axis. For ^{133}Cs, both $g_{F,g}$ and $g_{F,e}$ are positive. Figure B.4 illustrates the relevant $m_F = -1, 0,$ and $+1$ states, showing how the MOT operates with one laser driving the $\Delta m_F = -1$ transition and the other driving the $\Delta m_F = +1$ transition. At the position $z > 0$, we want Laser 2 (cooling beam 2, coming from the $+z$ side with $\vec{k}$ propagating along $-z$) to be more resonant, so it should drive the $\Delta m_F = -1$ transition. In the atomic frame, this requires σ^- polarization. Since $\vec{k}$ is antiparallel to $\vec{B}$ in this case, the transformation rules from figure B.3(3) tell us that σ^- corresponds to RHP in the light frame. Conversely, Laser 1 (coming from the $-z$ side with $\vec{k}$ parallel to $\vec{B}$) must drive the $\Delta m_F = +1$ transition, requiring σ^+ polarization in the

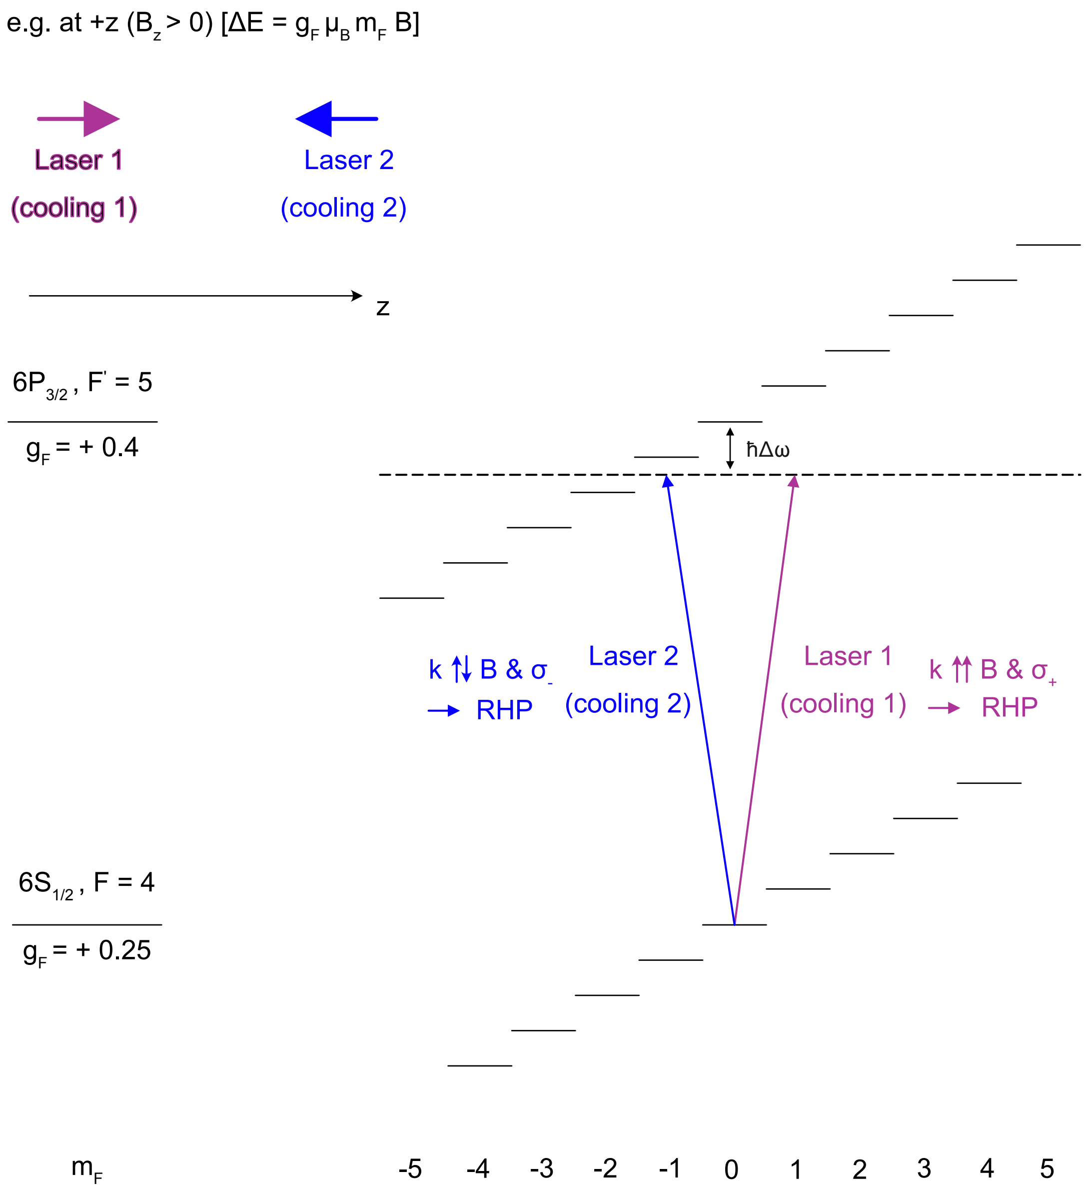

Figure B.4: Zeeman structure and laser configuration for MOT cooling. The figure shows the relevant magnetic sublevels for an atom at position $z > 0$ in a linear magnetic field gradient ($B_z > 0$ along the z-axis). For ^{133}Cs, both $g_{F,g}$ and $g_{F,e}$ are positive. At this position, the more resonant laser should be Laser 2, so we want laser 2 to drive the $\Delta m_F = -1$ transition (σ^-). The right side shows the required polarization in both coordinate systems: Laser 2 requires σ^- (atom frame) = RHP (light frame) with $\vec{k} \parallel \vec{B}$, while Laser 1 requires σ^+ (atom frame) = RHP (light frame) with $\vec{k}$ antiparallel to $\vec{B}$

atomic frame. Applying the transformation for $\vec{k} \parallel \vec{B}$, we find that cooling beam 1 also requires RHP in its light frame. Therefore, both cooling beams require RHP in their respective light frames (for a linearly increasing B-field along the z-axis) to properly drive the MOT transitions. Note that if the direction of the magnetic field gradient were reversed, the required polarizations would switch from RHP to LHP to maintain the proper MOT operation.

B.5 Polarization Requirements for Repumper and Imaging

The polarization states of the repumper and imaging beams are not strictly constrained by the atomic physics of the system, which offers an opportunity to simplify the optical setup. The repumper's sole function is to continuously pump atoms out of the dark ground state back into the cooling cycle; since any polarization will sufficiently drive this process, we can set the repumper to share the same polarization as cooling beam 1, eliminating the need to separate an additional polarization state. Similarly, standard imaging does not require a specific polarization, as atoms distributed across multiple magnetic sublevels will absorb nearly any incident light. By deliberately choosing the imaging beam polarization to be orthogonal to cooling beam 2, the two beams can be easily combined and separated along the same physical axis using a polarizing beam splitter (PBS).

B.6 Can Imaging and Cooling Be Separated?

Given this information, one can finally answer the initial question: Can we separate the imaging beam from cooling beam 1 using a PBS, see figure B.5(a)? This question reduces to asking whether cooling beams 1 and 2 themselves can be separated on a PBS. The answer is no—cooling beams 1 and 2 exit through the same port of a PBS, as illustrated in figure B.5(c). This is excellent news, because it means that we can indeed separate the imaging beam from cooling beams 1 and 2, as well as from the repumper.

How do we arrive at this conclusion? There are two ways to understand this:

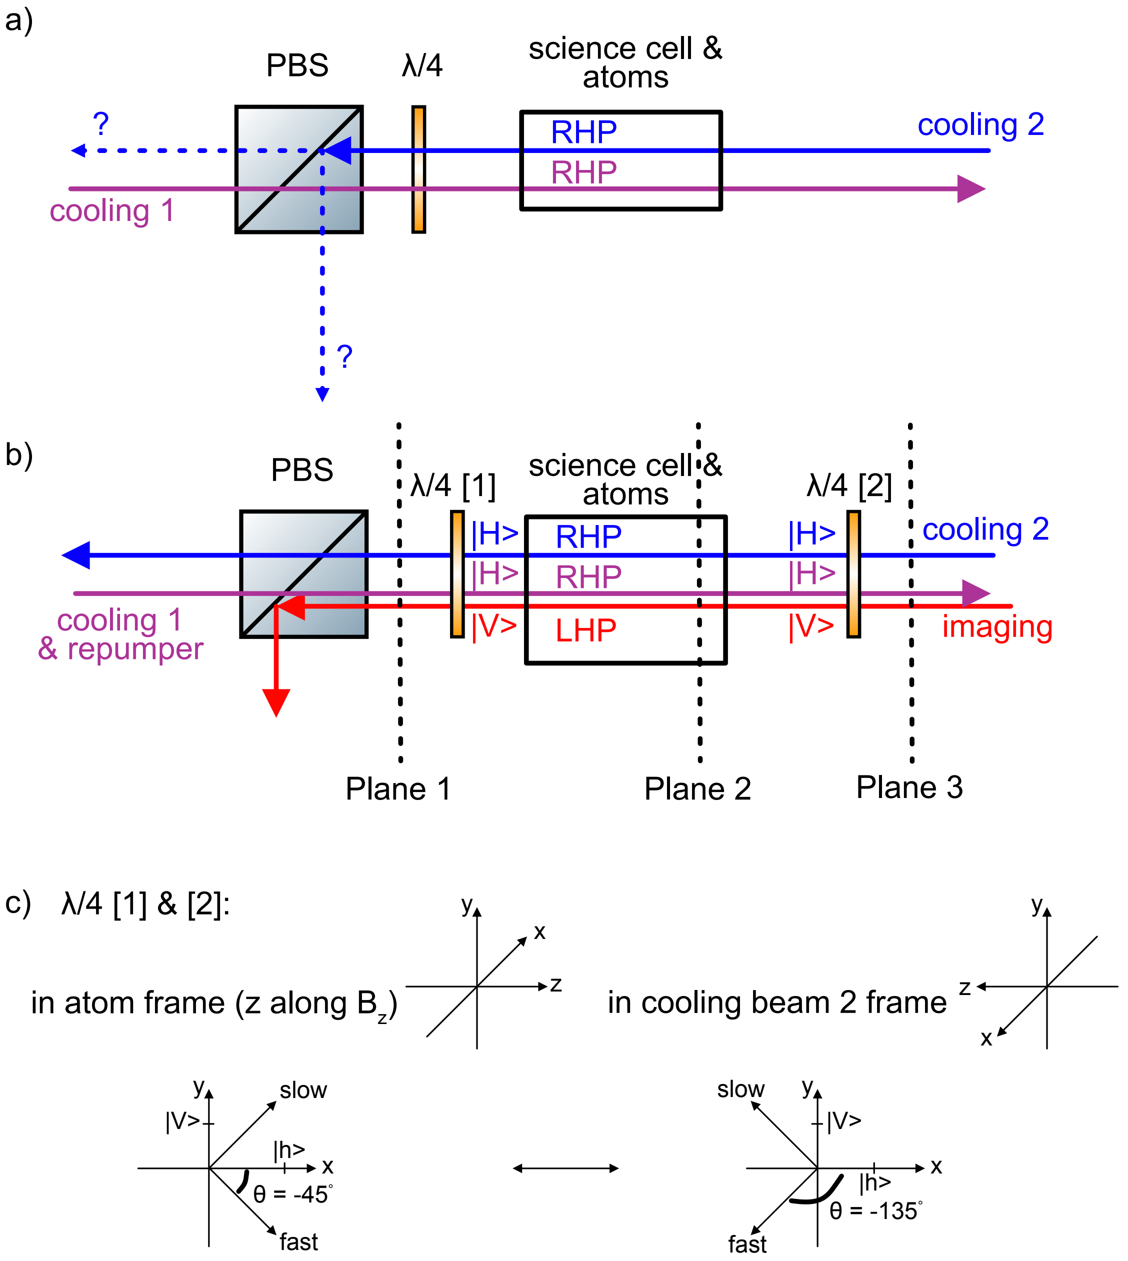

Figure B.5: (a) Illustration of the main question: can cooling beams 1 and 2 be separated on a PBS or not? (b) Complete beam configuration showing all beams (cooling 1, cooling 2, repumper, and imaging) passing through the system, which consists of a PBS and two $\lambda/4$ plates. As explained above, the two cooling beams must have RHP at the atoms, which then determines the polarization in planes 1 and 3. (c) Coordinate system transformation showing how the $\lambda/4$ waveplate axes relate to both the atomic frame (with z along $\vec{B}_z$) and the cooling beam 2 frame. $\lambda/4$ plates [1] and [2] have the same axis configuration, since both cooling beams must be RHP in the atomic frame. Assuming they are initially in $|H\rangle$ polarization, the given angles θ lead to the correct polarization at the atoms. More details are provided in the section on Jones calculus.

B.6.1 The Time-Reversal Argument

Intuitively, two beams that are time-reversals of each other cannot be separated by a PBS. Are cooling beams 1 and 2 time-reversals of each other? Yes, they are. Cooling beam 1 propagates in the $+z$ direction with RHP. Its time-reversal is a beam propagating in the $-z$ direction, and the sense in which the electric-field vector rotates is also reversed. As a result, in the light frame the time-reversed beam is again RHP—which is precisely cooling beam 2. Therefore, both beams must exit through the same port of the PBS.

B.6.2 Jones Calculus

Slightly less elegant, but a bit more rigorous and a good exercise for revisiting basic optics courses, one can explicitly verify the result using Jones calculus.

Before applying Jones calculus to the concrete beam geometry, it is useful to briefly recall how polarization states are represented in this formalism and how this connects to the electric-field description used above.

We start from the electric field of a monochromatic plane wave propagating along the z-direction,

$$\vec{E}(z,t) = \frac{1}{2}\left[\left(E_x e^{i\phi_x}\,\vec{e}_x + E_y e^{i\phi_y}\,\vec{e}_y\right) e^{i(kz-\omega t)} + \text{c.c.}\right], \tag{B.19}$$

where E_x and E_y are real amplitudes and ϕ_x, ϕ_y are the corresponding phases. Since the common space- and time-dependent factor $e^{i(kz-\omega t)}$ does not affect the polarization state, all polarization information is contained in the complex amplitudes multiplying $\vec{e}_x$ and $\vec{e}_y$.

Jones calculus makes this explicit by collecting these complex amplitudes into a two-component vector, the *Jones vector*,

$$\vec{E}_{\text{Jones}} = \begin{pmatrix} E_x e^{i\phi_x} \\ E_y e^{i\phi_y} \end{pmatrix}. \tag{B.20}$$

This vector fully characterizes the polarization state of the light in the transverse plane. Assuming $E_x = E_y$ and dropping the overall complex phase factor, the horizontal and vertical polarization states are represented as

$$|H\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad |V\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{B.21}$$

More general polarization states correspond to arbitrary superpositions of these basis states. For example, right- and left-handed circular polarization in the light coordinate system are given by

$$|\mathrm{RHP}\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}, \qquad |\mathrm{LHP}\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix}, \tag{B.22}$$

which reproduces the electric-field expressions discussed earlier in equations (B.2) and (B.3).

A quarter-wave plate with its fast axis aligned to the horizontal axis is described by

$$Q = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \tag{B.23}$$

corresponding to a relative phase shift of $\pi/2$ between the fast and slow axes. A quarter-wave plate whose fast axis is rotated by an angle α with respect to the horizontal axis is then given by

$$Q(\alpha) = R(-\alpha)\, Q\, R(\alpha). \tag{B.24}$$

where $R(\alpha)$ is the rotation matrix that transforms between linear polarization bases rotated by an angle α,

$$R(\alpha) = \begin{pmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{pmatrix}. \tag{B.25}$$

Before explicitly applying the Jones matrices, it is useful to understand why the fast and slow axes of both quarter-wave plates must be oriented at an angle of $\theta = -45°$

with respect to the horizontal axis. The polarization requirements of the cooling beams are fixed by two independent constraints. First, as derived above, both cooling beams must arrive at the atoms with right-handed circular polarization (RHP) in their respective light coordinate systems in order to correctly drive the MOT transitions. This fixes the required polarization state in plane 2 of figure B.5(b). Second, due to the presence of the polarizing beam splitter (PBS), both cooling beams must be linearly polarized along the horizontal axis in planes 1 and 3. Any vertical polarization component would be reflected by the PBS and thus removed from the cooling beam path. Therefore, the role of each quarter-wave plate is uniquely defined: it must transform a horizontally polarized input state $|H\rangle$ into a right-handed circularly polarized state at the position of the atoms. In Jones calculus, this transformation is only possible if the horizontal polarization has equal projections onto the fast and slow axes of the waveplate, such that the two components acquire a relative phase shift of $\pi/2$ while maintaining equal amplitudes. This condition is fulfilled precisely when the fast (and slow) axes of the quarter-wave plate are oriented at ±45° with respect to the horizontal axis. Choosing $\theta = -45°$ ensures that an incoming $|H\rangle$ state is decomposed into equal-amplitude components along the fast and slow axes, with the slow axis acquiring the additional phase shift of $\pi/2$. The resulting superposition corresponds to right-handed circular polarization in the light frame. Since both cooling beams enter the setup with the same horizontal linear polarization and must be converted into RHP at the atoms, both quarter-wave plates must share the same physical axis orientation in the laboratory frame. The apparent difference in the effective waveplate angles for the two beams arises solely from their opposite propagation directions and the associated change of the photon coordinate system, not from any difference in the actual alignment of the optical elements.

We now consider the action of the waveplates on the two counterpropagating cooling beams. Both beams are assumed to enter the setup in the same linear polarization state $|H\rangle$, as shown in figure B.5(c). The only difference between the two beams is their propagation direction, which determines how the waveplate axes are seen in the respective light coordinate systems. For cooling beam 1, propagating in the $+z$ direction, the laboratory frame coincides with the light frame. Choosing the waveplate angle $\alpha = -45°$, we obtain

$$Q(-45°)\,|H\rangle = R(45°)\begin{pmatrix}1 & 0\\ 0 & i\end{pmatrix}R(-45°)\begin{pmatrix}1\\ 0\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}1\\ i\end{pmatrix}, \tag{B.26}$$

which corresponds to right-handed circular polarization (RHP) in the light frame as defined in equation (B.22).

For cooling beam 2, propagating in the $-z$ direction, the photon coordinate system is flipped with respect to the laboratory frame. In particular, the transverse x-axis changes sign, which effectively changes the rotation angle of the waveplate as seen by the beam. As sketched in figure B.5(c), this leads to an effective waveplate angle $\alpha = -135°$. Applying the same Jones formalism, we find

$$Q(-135°)\,|H\rangle = R(135°)\begin{pmatrix}1 & 0\\ 0 & i\end{pmatrix}R(-135°)\begin{pmatrix}1\\ 0\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}1\\ i\end{pmatrix}, \tag{B.27}$$

which again corresponds to right-handed circular polarization in the respective light frame. Thus, although the two cooling beams propagate in opposite directions and experience different effective waveplate angles, both emerge with identical Jones vectors corresponding to RHP. Since a PBS separates light purely based on its linear polarization components in a fixed laboratory basis, two beams with identical polarization states cannot be separated and must exit through the same port. This explicitly confirms the conclusion drawn from the time-reversal argument.

Appendix C

MOT Stray Light Decoherence

To evaluate how MOT stray light impacts qubit coherence, we must distinguish between the distinct physical mechanisms of quantum information loss. Our qubit is encoded in a superposition of two ground states, $|\downarrow\rangle$ and $|\uparrow\rangle$. When the atom interacts with the stray light field from the MOT, the light acts as an uncontrolled environment, leading to decoherence. We model this open quantum system dynamics using the Bloch-Redfield master equation, which naturally separates this environmental noise into two primary effects:

- Population Decay (T_1): Also known as longitudinal relaxation, this is driven by the intensity of the stray light. Stray MOT photons can be directly absorbed by the atom, exciting it to the $6P_{3/2}$ manifold. When it spontaneously emits and decays back down, it may land in a different state, effectively pumping the population out of our defined qubit subspace. The rate of this process is defined as $\Gamma_1 = 1/T_1$.
- Pure Dephasing (T_ϕ): Even if a photon doesn't pump the atom to a different state, the act of scattering a photon where the atom ends in the same state it started still provides the environment with information about the atom's state. This effectively "measures" the atom, collapsing the superposition and leading to dephasing without any net change in population. The rate of this pure dephasing is $\Gamma_\phi = 1/T_\phi$.

The total decoherence time of the system (T_2), often visualized as the decay of the transverse components on the Bloch sphere, is a combination of both the population loss and the pure phase scrambling. It is defined as:

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi} \tag{C.1}$$

In the Bloch-Redfield model, the time evolution of the density matrix for a two-level system is essentially fully described by T_1, T_2, and T_ϕ in the following way:

$$\dot{\rho} = \begin{pmatrix} \dot{\rho}_{00} \\ \dot{\rho}_{01} \\ \dot{\rho}_{10} \\ \dot{\rho}_{11} \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & \Gamma_1 \\ 0 & -\left(\frac{\Gamma_1}{2} + \Gamma_\phi\right) & 0 & 0 \\ 0 & 0 & -\left(\frac{\Gamma_1}{2} + \Gamma_\phi\right) & 0 \\ 0 & 0 & 0 & -\Gamma_1 \end{pmatrix} \begin{pmatrix} \rho_{00} \\ \rho_{01} \\ \rho_{10} \\ \rho_{11} \end{pmatrix} \tag{C.2}$$

To evaluate how much the MOT stray light influences Γ_ϕ and Γ_1 in the qubit subspace, we construct a full open-quantum-system model. This approach can be broken down into three main steps: We first create the Hamiltonian describing the coherent energy landscape of the atom within the optical tweezer. This is calculated for the full 48-level system (comprising the $6S_{1/2}$ ground and $6P_{3/2}$ excited manifolds):

$$H = H_0 + H_{\text{Zeeman}} + H_{\text{AC}} \tag{C.3}$$

Here, H_0 contains the bare atomic energies and hyperfine structure, H_{Zeeman} accounts for the linear Zeeman shifts from the external quantization magnetic field, and H_{AC} represents the AC Stark shifts induced by the 1064 nm tweezer. Calculating H_{AC} is relevant because the deep tweezer potential shifts the atomic resonance, which increases the effective detuning of the MOT stray light and renders it less resonant. Next, we construct the Lindblad jump operators to model the environmental noise. Two families of incoherent processes are modeled:

a) Spontaneous Emission: We define one jump operator per allowed $|e\rangle \to |g\rangle$ decay channel:

$$L_{eg} = \sqrt{\Gamma_{\text{tot}} \cdot \text{BR}_{eg}}\; |g\rangle\langle e| \tag{C.4}$$

where $\Gamma_{\text{tot}} = 1/\tau(6P_{3/2}) \approx 32.8$ MHz is the total natural linewidth, and BR_{eg} is the branching ratio for the specific transition.

b) MOT Rescattered Light Absorption: Four laser channels (cooling, repumper, depumper, state-prep) drive incoherent $|g\rangle \to |e\rangle$ transitions. Each channel has its own intensity I_k and detuning Δ_k. The steady-state scattering rate per sublevel pair is calculated using a two-level saturation formula:

$$R_{ge} = \frac{\Gamma_{\text{tot}}}{2} \frac{s \cdot \text{BR}_{ge}}{1 + s + 4\left(\frac{\delta_{\text{eff}}}{\Gamma_{\text{tot}}}\right)^2} \tag{C.5}$$

where $s = I_k/I_{\text{sat}}$ is the saturation parameter, BR_{ge} is the branching ratio, and δ_{eff} is the effective detuning that accounts for the AC-Stark-shifted transition frequency. The corresponding jump operator is:

$$L_{ge} = \sqrt{R_{ge}}\ |e\rangle\langle g| \tag{C.6}$$

From this full master equation, we can extract the effective decoherence parameters T_1, T_2, and T_ϕ. We assume that the dynamics of our encoded qubit subspace (defined by the states $|{\downarrow}\rangle = |6S_{1/2}, F = 3, m_F = 0\rangle$ and $|{\uparrow}\rangle = |6S_{1/2}, F = 4, m_F = 0\rangle$) can be accurately mapped onto the two-level Bloch-Redfield model described previously. Therefore, we simply need to equate the elements of the calculated $\dot{\rho}$ corresponding to the qubit subspace with the rates Γ_1, Γ_2, and Γ_ϕ. By vectorizing the density matrix, the master equation is recast into Liouville space, where $\mathcal{L}$ becomes an $N^2 \times N^2$ matrix (with $N = 48$ being the total number of simulated states). For a two-level qubit embedded within this larger N-level system, the relevant decay rates appear directly on the diagonal of $\mathcal{L}$ in its superoperator representation. Consequently, no computationally expensive eigensolver is required; we can extract the rates by reading two diagonal entries directly from the sparse matrix:

$$\Gamma_1 = -\text{Re}\left[\mathcal{L}_{q_1 N + q_1,\, q_1 N + q_1}\right] \tag{C.7}$$

$$\Gamma_2 = -\text{Re}\left[\mathcal{L}_{q_1 N + q_0,\, q_1 N + q_0}\right] \tag{C.8}$$

Here, q_0 and q_1 are the respective basis indices of the $|6S_{1/2}, F = 3, m_F = +3\rangle$ and $|6S_{1/2}, F = 4, m_F = +4\rangle$ states. Γ_1 yields the total depopulation rate of the $|6S_{1/2}, F = 4, m_F = +4\rangle$ state, while Γ_2 yields the total decoherence rate of the off-diagonal coherence term between these two states (i.e., $\rho_{q_0 q_1}$). Once Γ_1 and Γ_2 are extracted, the pure dephasing rate is simply determined by $\Gamma_\phi = \Gamma_2 - \Gamma_1/2$. Figure C.1 shows the results from this master equation simulation, plotting the extracted decoherence rates against different intensities of the scattered MOT light. For this simulation, we assumed a quantization magnetic field of 3.0 G and applied all four

MOT laser channels (cooling, repumper, depumper, and state-prep) at zero detuning relative to the bare atomic transitions. We evaluated the dynamics under three different 1064 nm tweezer trap depths: 0.0 mK (free space), 0.5 mK, and 1.0 mK. To contextualize these results, we compare the full open-quantum-system model against a simplified "Absolute Worst Case" model (labeled as the Zunqi calculation). This baseline model treats the ^{133}Cs atom as a simple two-level system that is perfectly resonant with the stray 852 nm D_2 light. It assumes the light interacts with the atom using the maximum theoretical scattering cross-section, $\sigma_{\text{std}} = 3\lambda^2/(2\pi)$. Under this approximation, the total scattering rate is simply given by $\Gamma = I\sigma_{\text{std}}/(\hbar\omega)$. Because this worst-case model ignores the complex multi-level structure of the atom, Clebsch-Gordan coefficients, and external energy shifts, it serves as a strict theoretical upper bound for the decoherence rate. If our experimental goal is to maintain a coherence time of $T_2 = 1$ s (the teal horizontal line at $\Gamma = 1\ \text{s}^{-1}$), the worst-case model suggests the stray MOT light must be kept strictly below $\sim 10^{-7}$ mW/cm^2. However, according to the more accurate master equation model, atoms trapped in a deep 1.0 mK tweezer, we can tolerate nearly 10^{-6} mW/cm^2 of stray light, a full order of magnitude more, before breaking that coherence threshold.

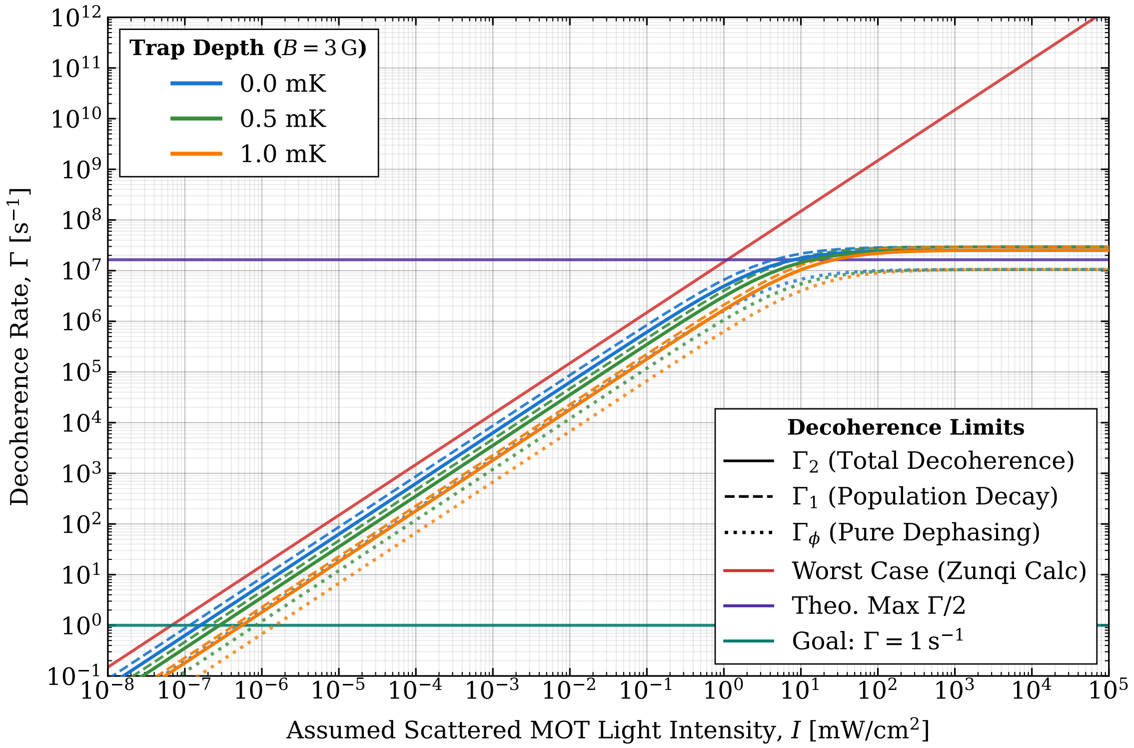

Figure C.1: Extracted decoherence rates versus scattered MOT light intensity based on a master equation simulation. The simulation assumes a quantization magnetic field of 3.0 G with all four MOT laser channels (cooling, repumper, depumper, and state-prep) at zero detuning relative to the bare atomic transitions. Dynamics are evaluated under three 1064 nm tweezer trap depths: 0.0 mK (free space), 0.5 mK, and 1.0 mK. The full open-quantum-system model is compared against a simplified “Absolute Worst Case” baseline model (Zunqi calculation), which assumes resonant interaction with stray 852 nm D_2 light using the maximum theoretical scattering cross-section. The teal horizontal line denotes a target coherence time of $T_2 = 1$ s ($\Gamma = 1$ s^{-1}).

Appendix D

DMD

This section details the operational principles of DMDs and their practical alignment. A DMD essentially functions as a blazed grating because its individual micromirrors are tilted with respect to the device plane by an angle $\gamma = 12^\circ$ (see figure D.2). Conceptually, the DMD can be modeled as the combination of two optical elements: a standard diffraction grating with an effective line spacing of $d = d' \cos(\gamma)$, and a reflective mirror tilted by γ relative to the grating plane. Each of these conceptual elements governs a distinct aspect of the resulting angular intensity distribution. For the pure grating component, plotting the intensity $I_g(\beta)$ against the diffraction angle β yields the periodic interference pattern governed by the grating equation, where the discrete maxima occur at $\sin(\beta_m) = m\frac{\lambda}{d}$. Simultaneously, the finite size of the tilted mirror produces a broader diffraction envelope, $I_m(\theta_r)$, as a function of the angle of reflection θ_r. To achieve maximum diffraction efficiency, a peak of the grating's interference pattern must coincide with the peak of the mirror's diffraction envelope. By relating the local mirror angles (θ_i, θ_r) to the global grating angles (α, β), we find that $\theta_i = \alpha - \gamma$ and $\theta_r = \beta + \gamma$. Equating these yields the geometric requirement for maximum intensity: $\alpha - \beta = 2\gamma$. When this geometric condition is satisfied simultaneously with the grating equation $\sin(\beta_m) = m\frac{\lambda}{d}$ for an integer diffraction order m, the "blazing condition" is met. For the specific DMD used in this setup (the DLP650LNIR), the micromirrors rest at -12° in the "off" state and tilt to $+12^\circ$ in the "on" state (see figure D.1). To achieve the blazing condition in practice, the DMD is first mounted on a rotation stage and roughly positioned so that the incident beam arrives at $\alpha = 12^\circ$ relative to the surface normal. In the "on" state, the light will primarily exit near $\beta = 36^\circ$. However, plugging these initial angles along with the operating wavelength λ and the effective mirror pitch into the grating equation yields a non-integer value for the diffraction order (e.g., $m = 7.288$). Because m must be an integer, this indicates the blazing condition is not yet perfectly satisfied. To finalize the alignment, the DMD is slightly rotated to adjust the angle of incidence while monitoring the output beam. The resulting diffraction pattern will appear as a grid of discrete dots. By fine-tuning the rotation until one of these diffraction spots reaches maximum brightness and clarity, an integer solution for m is found. This confirms the blazing condition is met, and this optimized diffracted beam can then be used for the subsequent optical setup.

This explains how to align the DMD and how to find the blazing condition. However, it does not address how strong the extinction ratio can be or how to theoretically predict it. A valuable resource for this is a recent master's thesis [74], which builds upon the work of Zhang et al. [75]. In the supplementary materials of [75] (specifically Equation S11), an explicit formula is provided for the intensity $I(m)$ as a function of the diffraction order m. Using this formula, it becomes clear that the theoretically achievable on/off contrast is on the order of 10^{-6}, a value that the authors also report achieving experimentally. For our intended application of the DMD—creating a repumper dark tube with an "on" state intensity of approximately $100\,\mu\mathrm{W/cm}^2$ and an "off" state intensity below $5\,\mu\mathrm{W/cm}^2$—an extinction ratio of this magnitude is more than sufficient.

Although the main experimental optical tables were not yet fully assembled, we constructed an independent DMD test setup. This allowed us to verify our optical understanding of the device and to establish reliable programmatic control. During this process, we encountered significant compatibility issues regarding the control software and firmware versions. These challenges are documented here to aid future troubleshooting. Different versions of the DMD driver boards are compatible with specific Graphical User Interface (GUI) software releases, and identifying the correct pairing can be difficult. Texas Instruments provides three primary GUI versions across two different download pages: two versions (from 2016 and 2018) available within the dlpc104a ZIP file and a separate 'Version 2" GUI. Through systematic testing, we determined that only the 2018 release (Software Version 2.0, OCX Version 1.2) from the dlpc104a package successfully communicated with all of our DMD driver boards. The other versions failed to establish a connection. Furthermore, the installation of this specific 2018 GUI initially failed on our systems; it required the prior installation of the Microsoft Visual C++ 2010 Redistributable Package. Finally, if multiple versions of the GUI are installed on the same computer, care must be taken to ensure the explicitly correct 2018 version is launched to avoid communication errors. Once the correct software environment was established, we proceeded with hardware testing. Notably, before connecting the driver board to the computer, the DMD independently cycled through pre-programmed test patterns. Depressing the "SW3" switch on the board caused the micromirrors to park permanently, a state that persisted even after the switch was released. Subsequent programmatic control of the DMD was achieved via a Python interface, utilizing the codebase developed by Nadine Meister. By strictly following the repository's documentation, programatic python operation was successfully established. Using this software infrastructure alongside the test setup shown in figure D.3, we were able to individually control the micromirrors, as demonstrated by the sample camera images of generated patterns shown in figure D.4.

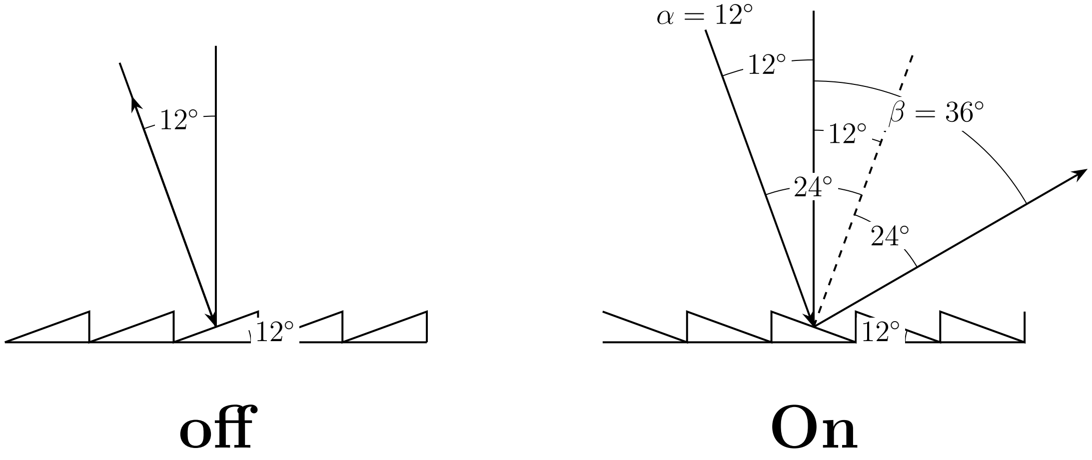

Figure D.1: Practical angular configuration for the DLP650LNIR DMD in the off and on states. For the On state, the general blazed condition implies $\sin(\beta) = \sin(\alpha + 2\gamma) = \sin(12° + 2 \cdot 12°) = \sin(36°) = m\frac{\lambda}{d}$. Substituting the effective pitch yields an actual device alignment of $m = \frac{\sin(36°) \cdot d}{\lambda} = \frac{\sin(36°) \cdot 10.8\,\mu\mathrm{m} \cdot \cos(12°)}{\lambda} = 7.288$.

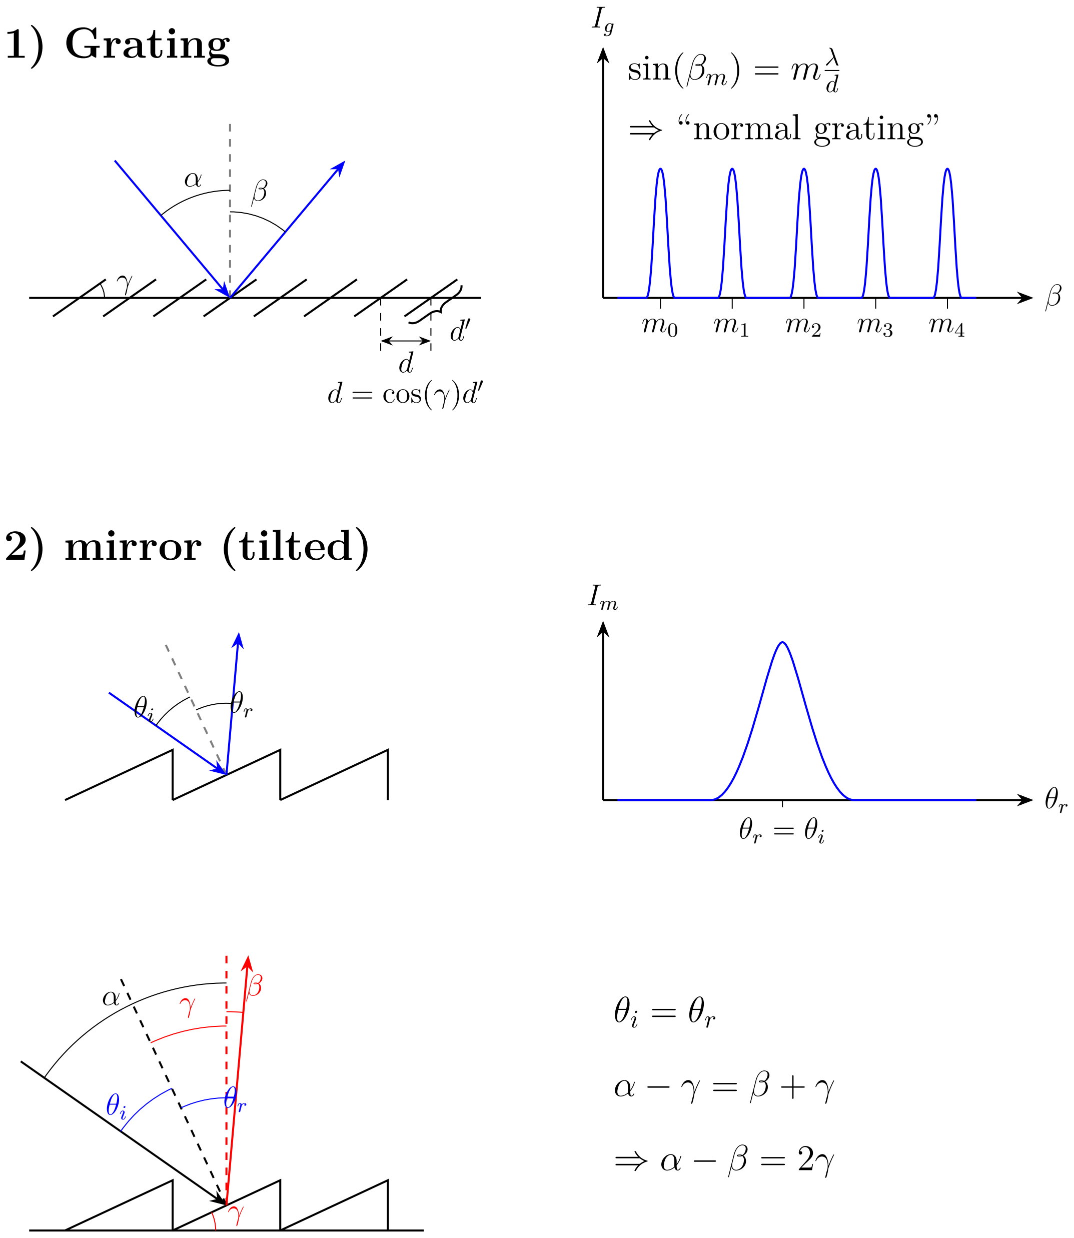

Figure D.2: Theoretical model of a DMD acting as a blazed grating, showing the combination of grating diffraction and mirror reflection. The best diffraction efficiency is achieved if the peak in I_m coincides with a peak in I_g. This occurs when the blazing condition is met: $\sin(\beta_m) = m\frac{\lambda}{d}$ and $\alpha - \beta = 2\gamma$.

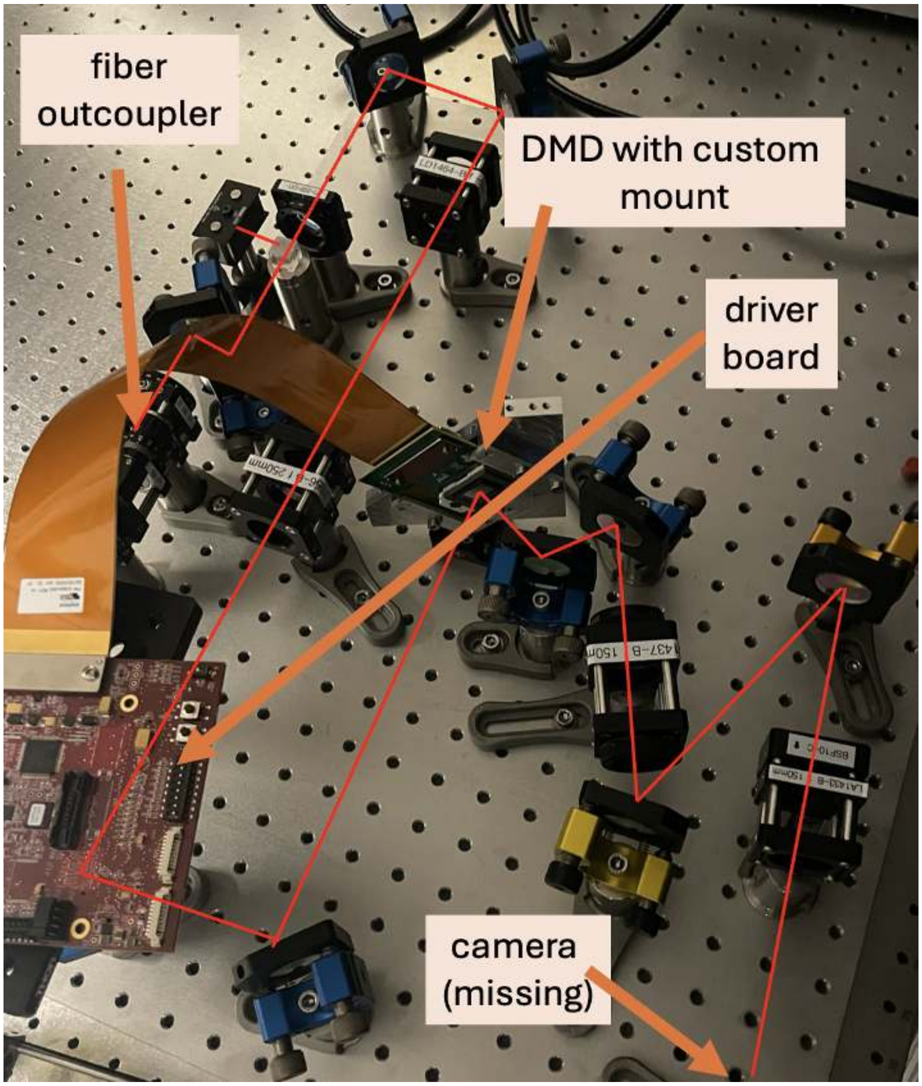

Figure D.3: The independent DMD test setup built to verify optical alignment and programmatic control prior to integration into the main experiment.

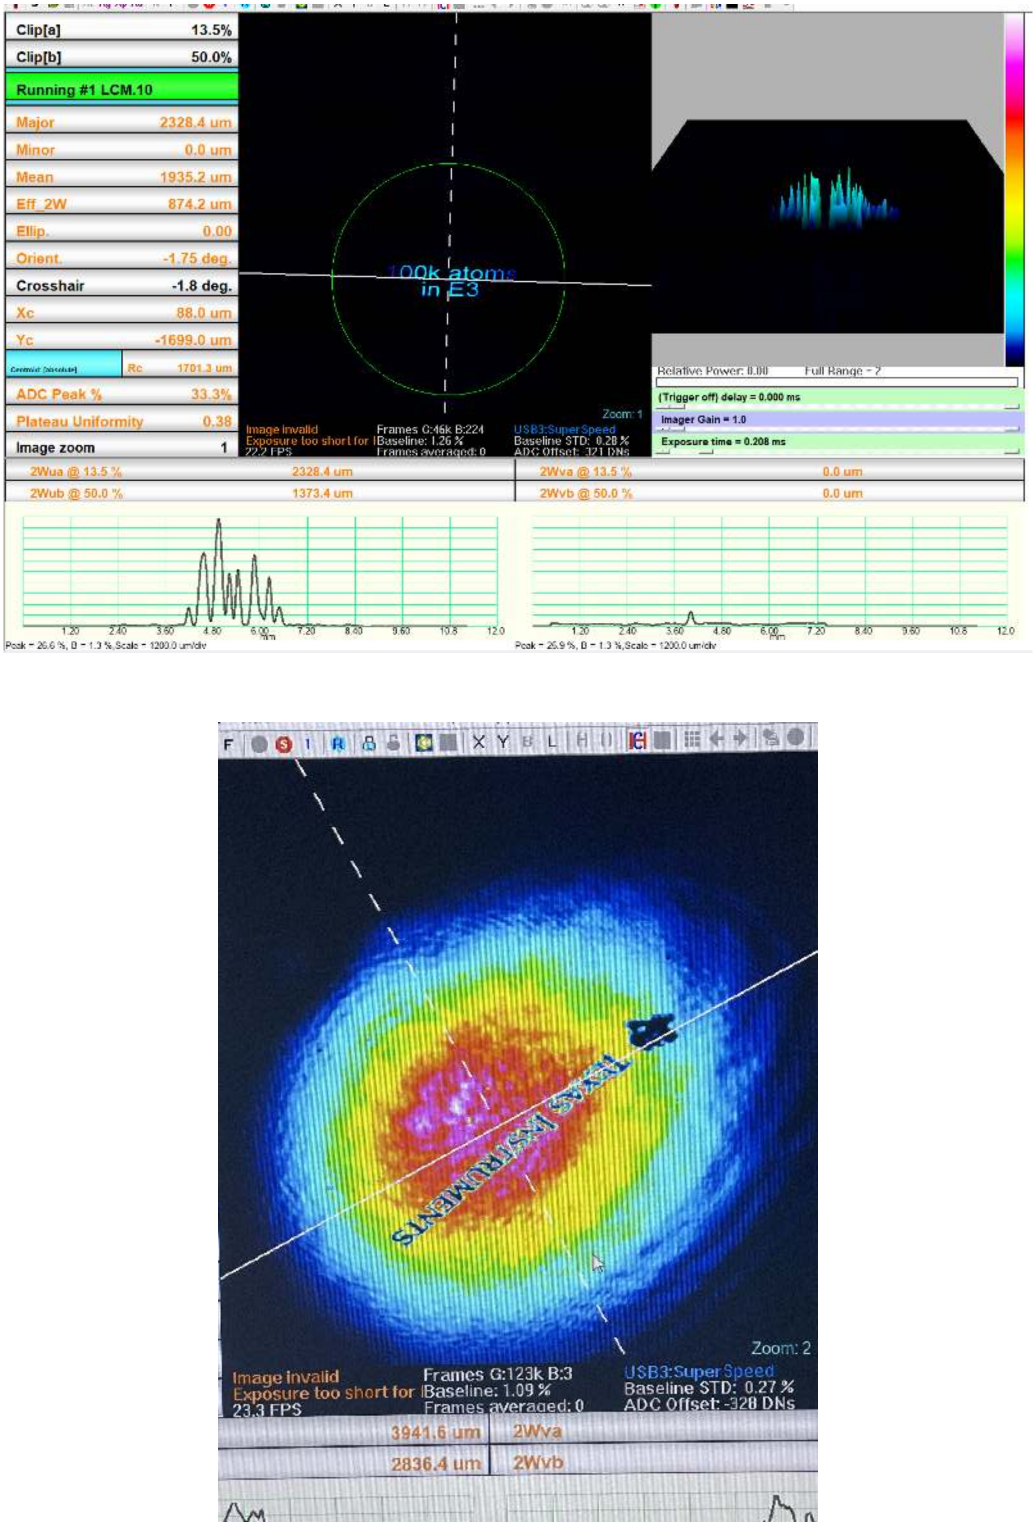

Clip[a] 13.5%
Clip[b] 50.0%
Running #1 LCM.10
Major 2328.4 um
Minor 0.0 um
Mean 1935.2 um
Eff_2W 874.2 um
Ellip. 0.00
Orient. -1.75 deg.
Crosshair -1.8 deg.
Xc 88.0 um
Yc -1699.0 um
Centroid [absolute] Rc 1701.3 um
ADC Peak % 33.3%
Plateau Uniformity 0.38
Image zoom 1

Image invalid
Exposure too short for l
22.2 FPS
Frames G:46k B:224
Baseline: 1.26 %
Frames averaged: 0
USB3:SuperSpeed
Baseline STD: 0.28 %
ADC Offset: 321 DNs

Relative Power: 0.00 Full Range = 2
(Trigger off) delay = 0.000 ms
Imager Gain = 1.0
Exposure time = 0.208 ms

2Wua @ 13.5 % 2328.4 um
2Wub @ 50.0 % 1373.4 um
2Wva @ 13.5 % 0.0 um
2Wvb @ 50.0 % 0.0 um

Peak = 26.6 %, B = 1.3 %,Scale = 1200.0 um/div
Peak = 26.9 %, B = 1.3 %,Scale = 1200.0 um/div

Image invalid
Exposure too short for
23.3 FPS
Frames G:123k B:3
Baseline: 1.09 %
Frames averaged: 0
USB3:SuperSpeed
Baseline STD: 0.27 %
ADC Offset: -328 DNs

3941.6 um 2Wva
2836.4 um 2Wvb

Figure D.4: Example images captured on the camera, demonstrating successful programmatic control over individual micromirrors to generate specific spatial patterns.

Appendix E

Effect of Aperture on Gaussian Beam

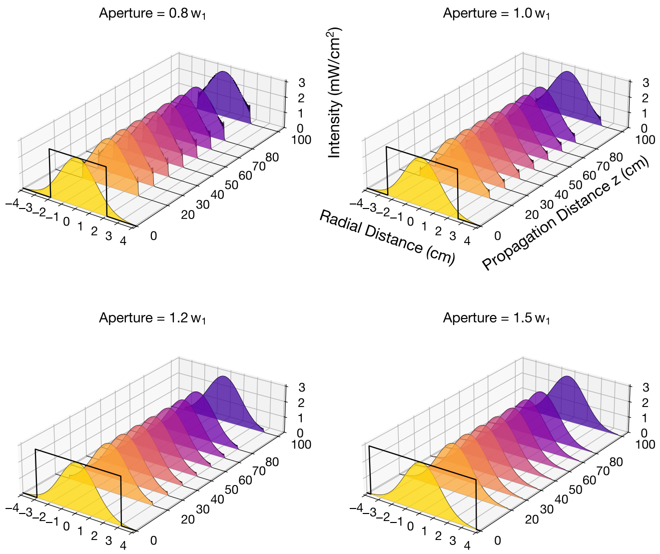

Figure E.1: Simulation demonstrating the undesirable diffraction effects (clipping) that occur when a Gaussian beam is passed through an aperture or reflected off a mirror that is not sufficiently larger than the beam's cross-section.

Appendix F

Additional Images

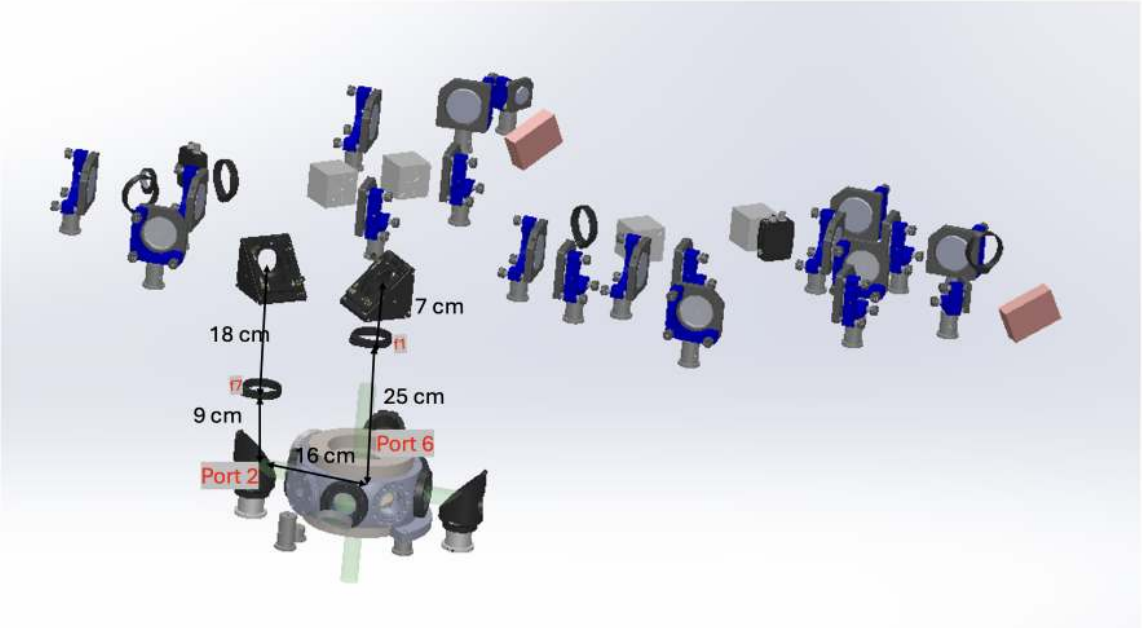

Figure F.1: Side view of the SolidWorks model of the imaging paths for both the repumper and imaging systems (polarization optics omitted for clarity).

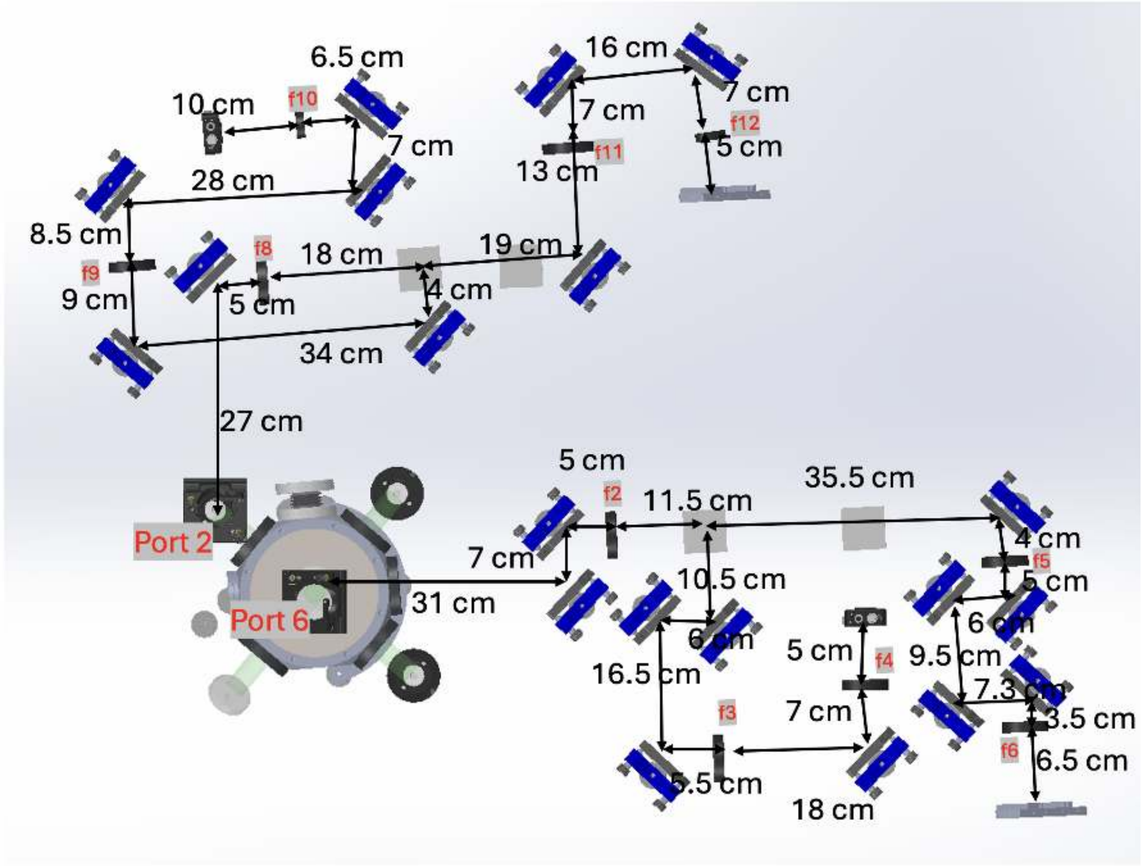

Figure F.2: Top view of the SolidWorks model of the imaging paths for both the repumper and imaging systems (polarization optics omitted for clarity).

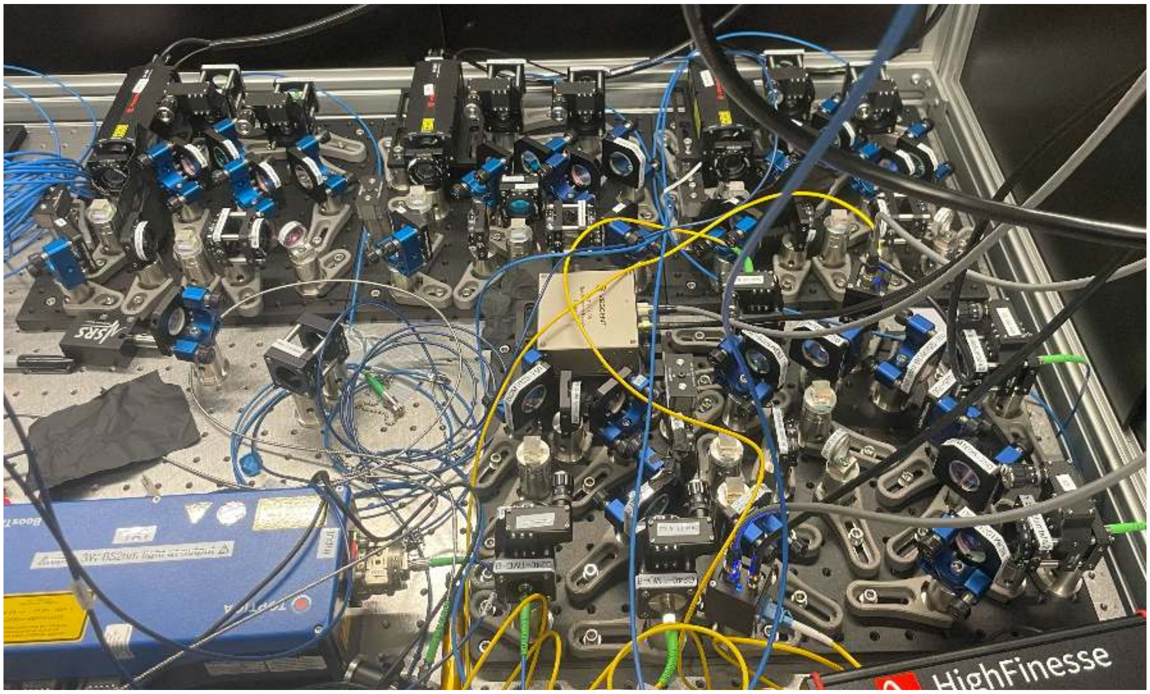

Figure F.3: Experimental realization of breadboards 1–4. The three upper breadboards correspond to breadboards 1–3; the lower-right breadboard is the beat breadboard (breadboard 4).

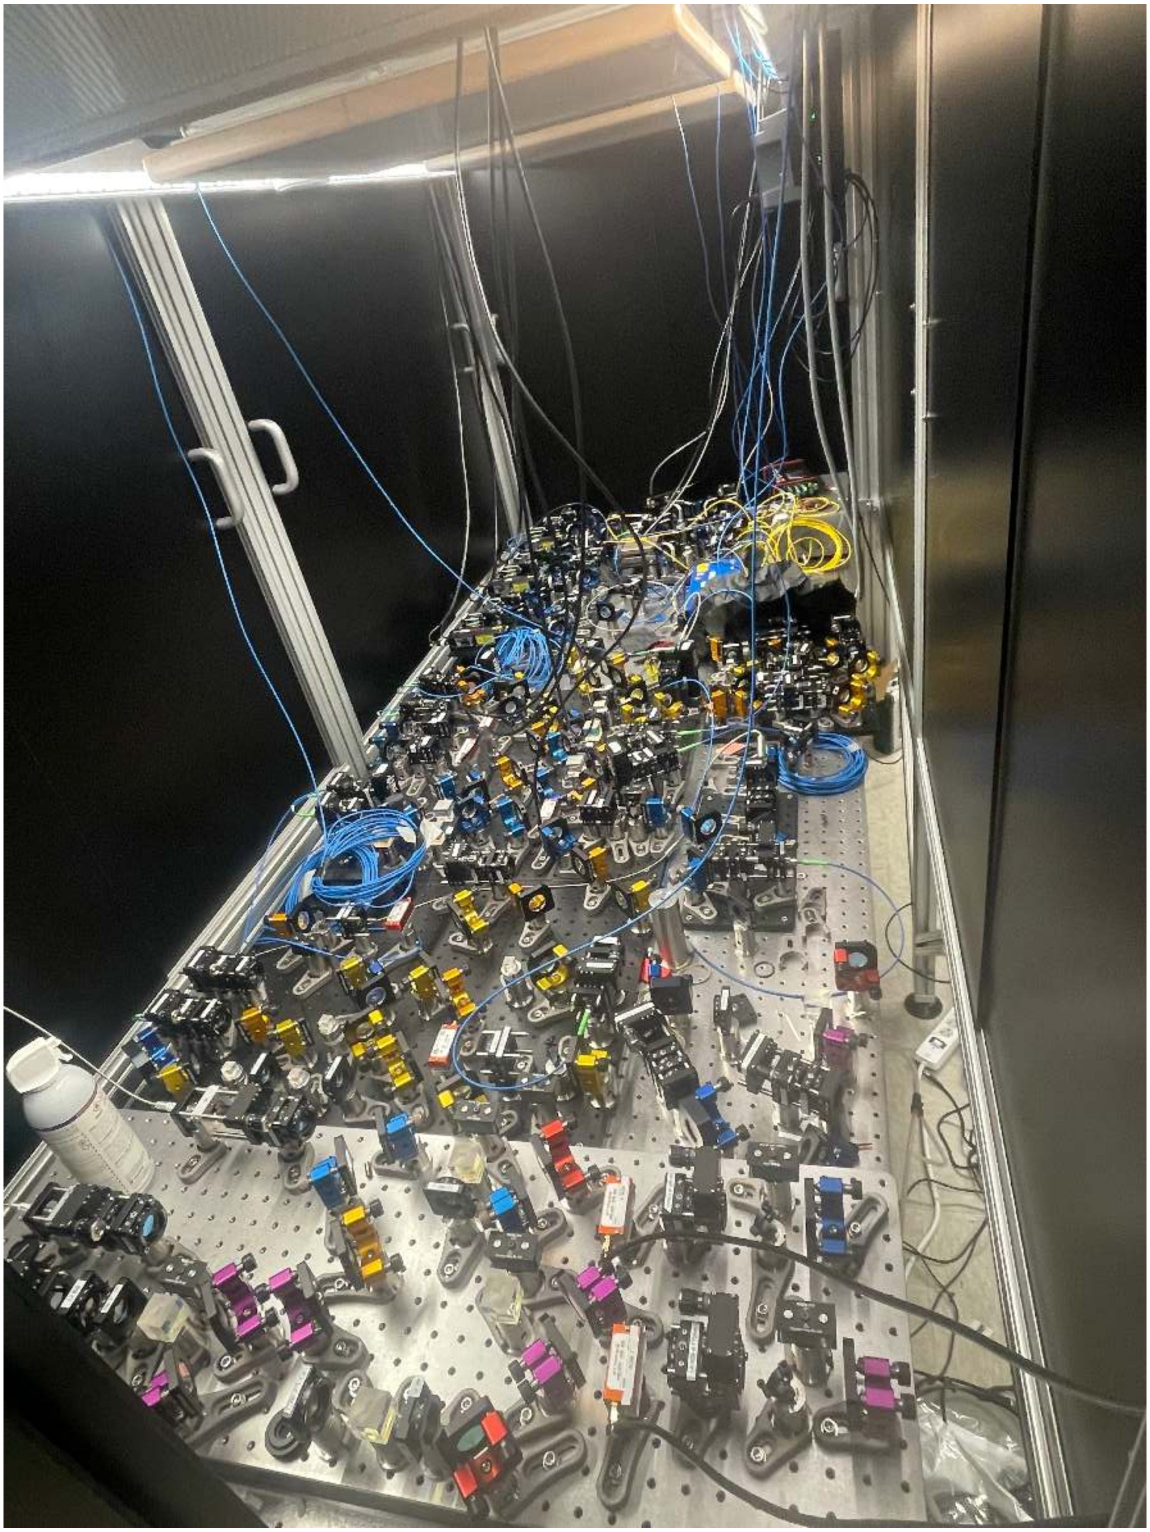

Figure F.4: Photograph of the fully implemented distribution breadboards.

List of Figures

Bibliography

[1] Henriet, L. *et al.* Quantum computing with neutral atoms. en. *Quantum* **4,** 327 (Sept. 2020).

[2] Browaeys, A. & Lahaye, T. Many-body physics with individually controlled Rydberg atoms. en. *Nat. Phys.* **16,** 132–142 (Feb. 2020).

[3] Saffman, M. Quantum computing with atomic qubits and Rydberg interactions: progress and challenges. *J. Phys. B At. Mol. Opt. Phys.* **49,** 202001 (Oct. 2016).

[4] Manetsch, H. J. *et al.* A tweezer array with 6,100 highly coherent atomic qubits. en. *Nature* **647,** 60–67 (Nov. 2025).

[5] Wintersperger, K. *et al.* Neutral atom quantum computing hardware: performance and end-user perspective. en. *EPJ Quantum Technol.* **10** (Dec. 2023).

[6] Arute, F. *et al.* Quantum supremacy using a programmable superconducting processor. en. *Nature* **574,** 505–510 (Oct. 2019).

[7] Krantz, P. *et al.* A quantum engineer's guide to superconducting qubits. en. *Appl. Phys. Rev.* **6,** 021318 (June 2019).

[8] Bluvstein, D. *et al.* A quantum processor based on coherent transport of entangled atom arrays. en. *Nature* **604,** 451–456 (Apr. 2022).

[9] Bluvstein, D. *et al.* Architectural mechanisms of a universal fault-tolerant quantum computer. arXiv: `2506.20661 [quant-ph]` (June 2025).

[10] Panteleev, P. & Kalachev, G. Asymptotically good quantum and locally testable classical LDPC codes. arXiv: `2111.03654 [cs.IT]` (Nov. 2021).

[11] Bravyi, S. *et al.* High-threshold and low-overhead fault-tolerant quantum memory. en. *Nature* **627,** 778–782 (Mar. 2024).

[12] Cain, M. *et al.* Shor's algorithm is possible with as few as 10,000 reconfigurable atomic qubits. arXiv: `2603.28627 [quant-ph]` (Mar. 2026).

[13] Bluvstein, D. *et al.* Logical quantum processor based on reconfigurable atom arrays. en. *Nature* **626,** 58–65 (Feb. 2024).

[14] Morgado, M. & Whitlock, S. Quantum simulation and computing with Rydberg-interacting qubits. en. *AVS Quantum Sci.* **3,** 023501 (June 2021).

[15] Fowler, A. G. *et al.* Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A* **86** (Sept. 2012).

[16] Singh, K. *et al.* Mid-circuit correction of correlated phase errors using an array of spectator qubits. en. *Science* **380,** 1265–1269 (June 2023).

[17] Lis, J. W. *et al.* Mid-circuit operations using the omg-architecture in neutral atom arrays. arXiv: `2305.19266 [quant-ph]` (May 2023).

[18] Evered, S. J. *et al.* High-fidelity parallel entangling gates on a neutral-atom quantum computer. en. *Nature* **622,** 268–272 (Oct. 2023).

[19] Li, Y. *et al.* Fast, continuous and coherent atom replacement in a neutral atom qubit array. arXiv: `2506.15633 [quant-ph]` (June 2025).

[20] Chiu, N.-C. *et al.* Continuous operation of a coherent 3,000-qubit system. en. *Nature* **646,** 1075–1080 (Oct. 2025).

[21] Muniz, J. A. *et al.* Repeated Ancilla reuse for logical computation on a neutral atom quantum computer. en. *Phys. Rev. X.* **15** (Dec. 2025).

[22] Gyger, F. *et al.* Continuous operation of large-scale atom arrays in optical lattices. *Physical Review Research* **6,** 033104 (July 25, 2024).

[23] Singh, K. *et al.* Dual-element, two-dimensional atom array with continuous-mode operation. en. *Phys. Rev. X.* **12** (Mar. 2022).

[24] Schioppo, M. *et al.* Ultrastable optical clock with two cold-atom ensembles. en. *Nat. Photonics* **11,** 48–52 (Jan. 2017).

[25] Saffman, M., Walker, T. G. & Mølmer, K. Quantum information with Rydberg atoms. en. *Rev. Mod. Phys.* **82,** 2313–2363 (Aug. 2010).

[26] Deist, E. *et al.* Mid-circuit cavity measurement in a neutral atom array. en. *Phys. Rev. Lett.* **129,** 203602 (Nov. 2022).

[27] Cooper, A. *et al.* Alkaline-earth atoms in optical tweezers. en. *Phys. Rev. X.* **8** (Dec. 2018).

[28] Norcia, M. A. *et al.* Iterative assembly of 171 Yb atom arrays with cavity-enhanced optical lattices. en. *PRX quantum* **5** (July 2024).

[29] Bureau International des Poids et Mesures. *Résolution 1 (2018): On the revision of the International System of Units (SI)* Resolution (BIPM, Nov. 16, 2018).

[30] Steck, D. A. *Cesium D Line Data* version 2.2.1. University of Oregon. https://steck.us/alkalidata/cesiumnumbers.pdf (2026).

[31] Metcalf, H. J. & van der Straten, P. Laser cooling and trapping of atoms. en. *J. Opt. Soc. Am. B* **20,** 887 (May 2003).

[32] Grimm, R., Weidemüller, M. & Ovchinnikov, Y. B. *Optical dipole traps for neutral atoms* Feb. 24, 1999. arXiv: physics/9902072.

[33] Xia, T. *et al.* Randomized Benchmarking of Single-Qubit Gates in a 2D Array of Neutral-Atom Qubits. *Physical Review Letters* **114,** 100503 (Mar. 12, 2015).

[34] Steck, D. A. *Quantum and Atom Optics* https://atomoptics.uoregon.edu/~tbrown/files/relevant_papers/quantum-optics-notes.pdf (2026).

[35] Foot, C. J. *Atomic Physics* en (Oxford University Press, London, England, Aug. 2003).

[36] Wikipedia contributors. *Magneto-optische Falle* Page Version ID: 243054583. https://de.wikipedia.org/w/index.php?title=Magneto-optische_Falle&oldid=243054583 (2026).

[37] Monroe, C. *et al.* Very cold trapped atoms in a vapor cell. *Physical Review Letters* **65,** 1571–1574 (Sept. 24, 1990).

[38] Lindblad, G. On the generators of quantum dynamical semigroups. en. *Commun. Math. Phys.* **48,** 119–130 (June 1976).

[39] Gorini, V., Kossakowski, A. & Sudarshan, E. C. G. Completely positive dynamical semigroups of N-level systems. en. *J. Math. Phys.* **17,** 821–825 (May 1976).

[40] Breuer, H.-P. & Petruccione, F. *The theory of open quantum systems* (Oxford University Press, London, England, Jan. 2007).

[41] Tarbutt, M. R. Magneto-optical trapping forces for atoms and molecules with complex level structures. *New J. Phys.* **17,** 015007 (Jan. 2015).

[42] Eckel, S. *et al.* *PyLCP: A python package for computing laser cooling physics* Nov. 16, 2020. arXiv: 2011.07979[physics].

[43] Heinz, A. *Ultracold Strontium in State-Dependent Optical Lattices* PhD thesis.

[44] Chin, C. *et al.* Feshbach resonances in ultracold gases. *Reviews of Modern Physics* **82,** 1225–1286 (Apr. 29, 2010).

[45] Frye, M. D., Yang, B. C. & Hutson, J. M. Ultracold collisions of Cs atoms in excited Zeeman and hyperfine states. *Physical Review A* **100,** 022702 (Aug. 1, 2019).

[46] Weber, T. *et al.* Bose-Einstein Condensation of Cesium. *Science* **299,** 232–235 (Jan. 10, 2003).

[47] Kraemer, T. *et al.* Optimized production of a cesium Bose–Einstein condensate. en. *Appl. Phys. B* **79,** 1013–1019 (Dec. 2004).

[48] Klostermann, T. *et al.* Fast long-distance transport of cold cesium atoms. *Physical Review A* **105,** 043319 (Apr. 26, 2022).

[49] Arnold, A. S. & Manson, P. J. Atomic density and temperature distributions in magneto-optical traps. en. *J. Opt. Soc. Am. B* **17,** 497 (Apr. 2000).

[50] Walker, T., Sesko, D. & Wieman, C. Collective behavior of optically trapped neutral atoms. en. *Phys. Rev. Lett.* **64,** 408–411 (Jan. 1990).

[51] Townsend, C. G. *et al.* Phase-space density in the magneto-optical trap. en. *Phys. Rev. A* **52,** 1423–1440 (Aug. 1995).

[52] Townsend, C. G. *et al.* High-density trapping of cesium atoms in a dark magneto-optical trap. en. *Phys. Rev. A* **53,** 1702–1714 (Mar. 1996).

[53] Haw, M. *et al.* Magneto-optical trap loading rate dependence on trap depth and vapor density. en. *J. Opt. Soc. Am. B* **29,** 475 (Mar. 2012).

[54] Yan-Xu, H. *et al.* Realization of high optical density Rb magneto-optical trap. *Chin. Physics Lett.* **26,** 023201 (Feb. 2009).

[55] Petrich, W. *et al.* Behavior of atoms in a compressed magneto-optical trap. en. *J. Opt. Soc. Am. B* **11,** 1332 (Aug. 1994).

[56] Gibble, K. E., Kasapi, S. & Chu, S. Improved magneto-optic trapping in a vapor cell. en. *Opt. Lett.* **17,** 526–528 (Apr. 1992).

[57] Camara, A., Kaiser, R. & Labeyrie, G. Scaling behavior of a very large magneto-optical trap. *Physical Review A* **90,** 063404 (Dec. 1, 2014).

[58] Klostermann, T. M. *Construction of a Caesium Quantum Gas Microscope* PhD thesis.

[59] Hamann, S. E. *et al.* Resolved-Sideband Raman Cooling to the Ground State of an Optical Lattice. *Physical Review Letters* **80,** 4149–4152 (May 11, 1998).

[60] Wang, Y. *Coherent operations, entanglement, and progress toward quantum search in a large 2D array of neutral atom qubits* PhD thesis.

[61] Salomon, C. *et al.* Laser cooling of cesium atoms below 3 μK. *Europhys. Lett.* **12,** 683–688 (Aug. 1990).

[62] Rosi, S. *et al.* Λ-enhanced grey molasses on the D2 transition of Rubidium-87 atoms. en. *Sci. Rep.* **8,** 1301 (Jan. 2018).

[63] Hsiao, Y.-F., Lin, Y.-J. & Chen, Y.-C. Λ -enhanced gray-molasses cooling of cesium atoms on the D2 line. en. *Phys. Rev. A* **98** (Sept. 2018).

[64] Ketterle, W. *et al.* High densities of cold atoms in a dark spontaneous-force optical trap. en. *Phys. Rev. Lett.* **70,** 2253–2256 (Apr. 1993).

[65] Lee, J. *et al.* Core-shell magneto-optical trap for alkaline-earth-metal-like atoms. *Physical Review A* **91,** 053405 (May 8, 2015).

[66] Scott, J. *et al.* Laser cooling and qubit measurements on a forbidden transition in neutral Cs atoms. en. *Phys. Rev. Lett.* **135,** 223403 (Nov. 2025).

[67] Ketterle, W., Durfee, D. S. & Stamper-Kurn, D. M. Making, probing and understanding Bose-Einstein condensates. arXiv: `cond-mat/9904034 [cond-mat]` (Apr. 1999).

[68] Gaudesius, M. *et al.* Three-dimensional simulations of spatiotemporal instabilities in a magneto-optical trap. *Physical Review A* **105,** 013112 (Jan. 21, 2022).

[69] MacAdam, K. B., Steinbach, A. & Wieman, C. A narrow-band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb. en. *Am. J. Phys.* **60,** 1098–1111 (Dec. 1992).

[70] Chang, R., He, J. & Wang, J. Interference filter based external-cavity diode laser with combined dual interference filters and largely adjustable feedback range. *Physica Scripta* **98,** 125111 (Nov. 2023).

[71] Richmond, G. *Double-pass acousto-optic modulator system* Technical Note (Stony Brook University, Laser Teaching Center, Mar. 5, 2018).

[72] Donley, E. A. *et al.* Double-pass acousto-optic modulator system. en. *Rev. Sci. Instrum.* **76,** 063112 (June 2005).

[73] Villarraga Gomes, H. *Magneto-optical studies of cobalt-doped nickel oxide thin films* PhD thesis.

[74] Zischka, N. *Single-Atom Addressing in a Strontium Quantum Simulator* Master's thesis.

[75] Zhang, B. *et al.* Scaled local gate controller for optically addressed qubits. en. *Optica* **11,** 227 (Feb. 2024).

Acknowledgements

The experience here at Caltech is something my childhood self would never believe if I told them. This past year has been incredibly full of learning, both scientifically and personally. None of it would have been possible without the remarkable people who supported me throughout this journey. I am grateful beyond words for every one of them, and it is humbling how many extraordinary, generous people I have encountered along the way. The list below captures only a small portion of them.

- **Professor Johannes Zeiher**, my thesis advisor and mentor from Germany, without whom I would likewise not be here. He was incredibly supportive throughout the entire year, calm and present even when things did not go as planned. I truly do not know how he managed to always respond within such a short time, to meet with me on Zoom almost immediately whenever I had a question, and even to proofread this thesis, all while juggling several labs in Germany.
- **Professor Manuel Endres**, without whom this year would also not have been possible. Thank you for giving me the opportunity to write my thesis in this incredible group. The weekly meetings, filled with invaluable scientific input, were deeply formative. I am also deeply grateful for the opportunities beyond the lab, including allowing me to attend DAMOP and CalBay, experiences that are rarely extended to a master's student and through which I grew enormously as a scientist.
- **Professor Andreas Reiserer**, without whom I would not be pursuing physics anymore. His continuous belief in me as a physicist is one of the primary reasons I was able to write this thesis at Caltech.
- **Kevin Mours** for teaching me a lot of the experimental skills that carried me through this thesis and for proofreading.
- **Franz Pfanner**, for being the best possible mentor in this group, and without

whom this thesis would absolutely not have been possible. From choosing this project for me, to always ensuring I understood what I was doing, to answering every question, experimental or theoretical, to managing the entire project with care, making everyone feel welcome, and providing a great deal of off-hours emergency support. Also thank you for proofreading.

- **Zunqi Li**, for invaluable discussions on theory. Without him, none of the simulations in this thesis would exist, and I would know far less about MOT physics. Thank you also for all your thoughtful questions about double-pass AOMs and other optics; I learned a great deal through every one of them. And for proofreading.
- **Aadithya Shankar**, for asking great questions, for taking on the MOT coil design (which I would not have managed alone), and for thoughtful input on papers and physics discussions and for proofreading.
- **Alice Jeffrey**, for the wonderful hours shared in the lab in the early mornings and evenings, for teaching me her lab skills, for double-checking all my MOT optics and asking sharp questions about them, and for working through all the tube testing with me.
- **Dr. Kon Leung**, for his tremendous help in the lab. From teaching me how the group builds fiber couplers in my very first week, to helping me lock Vescent lasers, tune PID parameters, and much more. Without him, none of the optics I built would exist and the Vescent lasers would probably not be locked now.
- **Dr. Lewis Picard**, for his generous help with double-pass AOMs and other general lab questions, and for ensuring my MOT planning made sense.
- **Nadine Meister**, for her discussions and help with the DMD (and, of course, for surviving the bed bug incident in Portland with me).
- **Elie Bataille, Gyohei Nomura, Xiangkai Sun and Richard Tsai**, for answering every question about the 2D/3D MOT and distribution breadboards multiple times over and always doing so with incredible patience and openness.
- **Hannah Manetsch**, for being a wonderful mentor here in Pasadena, for saying "good job" even when I did not believe it myself, for pushing me through the hardest stretches of this thesis, for being an inspiring role model, and for being a great thesis-writing companion.
- **Andrew Winnicki**, for rekindling my love of science through his infectious enthusiasm, and for supporting me through the final weeks of writing with

late-night and early-morning sessions.

- **Felix Schwarzfischer**, for always being there at any hour of the day or night, carrying me through this year and the decisions I had to make, for proofreading, for answering countless physics questions with an almost magical ability to do so from outside AMO, and for the help with the figures.
- **Niklas Zischka**, for literally carrying me through the thesis writing process by being the world's best proofreader for this thesis, helping with figures, helping me figure out how to hand in my thesis remotely, and providing very essential moral support. Also for wonderful AMO discussions, for somehow always having the answer to every technical question imaginable (DMDs, high-power fibers, and everything else) and for showing me the most precise clock in the world in Boulder.
- **Dr. Andrea Rosenauer and Dr. Bertram Layer** for being the world's best parents, supporting me with every stupid dream, always believing in me, and without whom I would definitely not have ever been able to be who and where I am. And especially Andrea also for amazingly detailed proofreading of this thesis.
- **Chiara Rosenauer** for being the world's best sister, being the most considerate and socially capable person I know, always being there for me and knowing the right thing to say to make me feel good, and for loving me unconditionally even when I am too focused on physics to be there for her enough.

To everyone else who walked alongside me this year — fellow students, friends, and colleagues — thank you.